Alma Mater Studiorum – Università di Bologna

DOTTORATO DI RICERCA IN

ECONOMICS

Ciclo 31

Settore Concorsuale: 13/A1

Settore Scientifico Disciplinare: SECS-P/01

Essays on Innovative Activity and the Protection of Innovations

Presentata da: David Tobias Angenendt

Coordinatore Dottorato

Prof.ssa Maria Bigoni

Supervisore

Prof. Vincenzo Denicolò

Esame finale anno 2019

PH.D. DISSERTATION ABSTRACT

DAVID T. ANGENENDT

This dissertation looks at three different widely accepted assumptions about how the patent system works to see whether the incentives are indeed what they are expected to be at the micro level: patent documents disclose inventions; this disclosure happens quickly; and patent owners are able to enforce patents.

The first chapter estimates the effect of stronger trade secret protection on the amount of patented innovations, finding that there is a sizeable reduction in disclosed knowledge per patent. It then shows how this endogeneity of knowledge per patent can affect the measurement of innovation using patent data. The second chapter chapter asks how the introduction of fee shifting in US patent litigation would influence firms' patenting propensity. It finds that manufacturing firms would unambiguously reduce patenting, with small firms affected the most; but that to have the same effect as that attributed to fee shifting in Europe, the US legal system would require a much smaller share of fees to be shifted. Finally, the third chapter contains a theoretical analysis of the influence of delayed disclosure of patent applications by the patent office, and what kind of strategic disclosure incentives such a delay may give when firms compete for more than a single innovations.

Chapter 1. Easy to Keep, but Hard to Find: How Patentable Inventions Are Being Kept Secret. The decision between patenting an invention and keeping it secret has received considerable attention in the theoretical economics, legal, and strategy literature, but empirical contributions suffer from the inherent unobservability of trade secrets. This paper seeks to avoid this problem by exploiting a quasi-natural experiment. It answers two main questions: Is there a trade-off between patents and secrets at all? If yes, how exactly does it affect patenting: Does the number of patents decrease, or is it rather patent content that is affected? If keeping inventions secret becomes a more viable alternative, the patent system is no longer attractive enough to deliver its disclosure function, making research results available to the public. If additionally patent content is the main strategic variable, we will not identify this effect in standard patent count statistics.

The main variable that I use as indicating the inventive content of patents is the number of independent claims, sentences in the patent document that define what exactly it is that is legally protected by the patent. The negative effect, however, also occurs with alternative measures of disclosure. Based on a simple theoretical model, I predict that a) in any situation, patentees will prefer to cover an invention with the smallest possible number of broad patents; b) firms will patent less when legal trade secret protection is increased, but the variable in which the effect shows is determined by firms' pre-treatment patenting behaviour; and c) the effect will be more often in claims rather than in the number of patents.

This paper exploits changes in US state trade secrets law in the context of implementing the Uniform Trade Secrets Act (UTSA). The baseline assumption is that a higher level of trade secret protection will lower a firm's patenting propensity. In contrast to Png's analysis, I find a negative effect on both variables, but it is not affecting both variables in parallel. Computing the text similarity between patent claims as opposed to whole patents, I produce evidence in support of a simple model of allocation of inventive components into patents. The way in which patentees react to a change in the incentive to patent therefore depends on the "design" of their patents, distinguishing many but narrow patents from few but broader ones.

Based on the finding that patentees can influence the number of claims per patent similar to the number of patents, I then show with two examples how using patent claim counts instead of simple

patent counts can affect the measurement of innovation using patent data. I argue that patent claim counts are more often unbiased of changes in patenting propensity compared to patent counts.

Chapter 2. The Role of Firm Size in Choice and Enforcement of Patent Protection: Does Fee-Shifting Help Small Firms? Non-practising entities (NPEs), colloquially known as patent trolls, buy patents not with the intention to produce or license but rather to exact damages or settlement payments from alleged patent infringers. In the US, this activity has become perceived as a problem severe enough to suggest patent policy responses, one of them being the shifting of legal fees to the losing party in patent litigation. The theoretical literature on litigation has shown that fee-shifting can reduce the incidence of unmeritorious lawsuits.

What, however, would be the effect of fee-shifting on meritorious patent lawsuits? There is evidence that small producing firms are generally disadvantaged in enforcing their patents, both in the US as well as in European countries which allow for varying extents of fee-shifting. For the particular influence of fee-shifting, however, it would be misguided to simply compare US patenting behaviour to that of European firms, since the legal systems differ in further dimensions. In the present context, the major difference is the relative importance that legal expenditure plays in determining trial outcomes, which is argued to be much higher in the US. A further difference is in the average quality of patents, due to a more rigorous examination process in Europe.

This paper suggests a model capturing these core parameters of legal systems and then compares predictions with empirical estimates of patenting propensity. In the model, fee-shifting does indeed decrease the profitability of patenting, and small firms are generally the first to forego patenting. The extent of fee-shifting necessary to affect the patenting rate however depends negatively on the importance of spending and on patent quality. To the extent that the United States are characterised by a higher importance of spending in litigation and lower patent quality compared to many European countries, a comparably small amount of fee shifting might be sufficient to reach the effect that requires a much higher amount of fee-shifting in Europe.

Chapter 3. As Soon as Optimal: Delayed Publication of Patent Applications. It is commonly assumed that a patent race ends when the first firm applies for a patent. Even if each application was granted protection with certainty, this does not have to be true since patent offices around the world have almost universally adopted the policy of publishing patent applications with a delay of 18 months. Compared to some estimates of the average length of patent races this can be a sizeable amount of additional time in which each firm other than the winner potentially continues to spend R&D resources that have an effective return of zero. The empirical academic literature has focused on the effect of publication of patent applications at all on the diffusion of knowledge. The theoretical literature on competition in R&D has neglected this feature of the patent system. This paper contributes a first analysis of the effect of a statutory delay in patent publication in races for a single innovation, and then studies cases in which firms may have an incentive to announce their pending patents before they will be published by the patent office. Such incentives may exist when the patentee of a first innovation benefits from her competitor working on a second innovation, which can be the case when, e.g., the first innovation allows generation of licensing revenues (cumulative innovation) or there is some degree of complementarity between the two innovations. Finally, I show that if policy makers were to set the statutory delay to zero there can be an incentive to delay the patenting of some complementary innovations.

Contents

Chapter 1. Easy to Keep, But Hard to Find: How Patentable Inventions Are Being Kept Secret

1 Summary

When Google sued Uber in February 2017 in a dispute over self-driving car technology, the allegations included both misappropriation of trade secrets and infringement of patents.[1] This event, unfortunate as it is for at least one of the two involved companies, serves as a great case study for scholars of intellectual property as it indicates that Google uses both patents and secrecy to cover its inventions in self-driving car technology. There is a growing literature stressing that the two forms of innovation protection cannot be fully understood separately and that their use is in fact closely intertwined. In a recent study, Png (2017b) finds a sizeable reduction of patenting in reaction to stronger trade secret laws, suggesting that with enforcement of trade secrets becoming easier, firms decided against patenting and rather kept a greater share of their inventions secret.

While the full specifications of the trade secrets are naturally hard to determine,[2] the inventions kept secret by Google are related to *printed circuit boards* (PCBs), fundamental computer components used inside of laser systems that enable autonomous cars autonomously navigate through traffic. Google's enormous patent portfolio contains hundreds of patents containing such PCBs, but patents differ in the way in which PCBs appear inside the patent document: in some patents, PCBs are part of the specification, the continuous text describing in detail the invention and its context (such as in US patents 8836922 and 9285464, two of the patents prevalent in the Google v. Uber case); in some other patents, PCBs appear as a minor part of a claim (e.g., in US patent 9123979).

In a third category of patents, the PCB is (one of) the main object(s) of invention, and accordingly a whole claim is dedicated to its description. The claims included in a patent document constitute the legally relevant specification of what it is precisely that is protected by the patent.[3] There are two cases to be distinguished here: those where

[1] A further allegation was that of unfair competition, which could be argued as closely related to trade secret misappropriation.

[2] The trade secrets have not spilled outside of the two companies, therefore the information that Google enclosed in its complaint does not go beyond a general description. The technical evidence in the ensuing case has been mostly filed under seal (Levine, 2017).

[3] Generally, each invention that is included in a patent is represented by an "independent claim", which needs to describe an invention that is patentable in its own right, i.e., its patentability does not depend on the presence of one or more other claims. Each independent claim can then be referred to by a number of "dependent claims" that contain further specifications or special cases of what is covered in the corresponding independent claim. These dependent claims can make the boundaries of patent

the PCB is the sole object of invention (i.e., there is only a single independent claim, or all independent claims are concerned with the PCB; e.g., US patent 9935514), and those where only one of several independent claims deals with a PCB, while the remaining claims cover other components of the same "complex invention" (e.g., US patent 5768103, acquired as part of Motorola's patent portfolio).[4]

This leads to the question of how such differences in patent content come about. An obvious answer, and one that certainly explains a part of this difference, is that the *underlying invention* in one case was focused on the particular PCB, while in the second case the PCB was but one of multiple components of the underlying "complex" invention. The point of this study is to make the argument that this does not explain everything, and that there is instead some discretion on the side of patent applicants regarding what to include in a single patent application. Previous authors have made convincing points in favor of such discretion, and the existence of "optimal patenting strategies", regarding such features of patents as what to patent (e.g., Anton & Yao, 2004; Arundel, 2001), how many redundant claims to file for any part of the invention, and how broad or narrow to phrase any particular claim (e.g., Yiannaka & Fulton, 2006, 2011). The "amount of invention" that is contained, and thereby disclosed to the public, in any patent is, in contrast, commonly regarded as an exogenous parameter of the particular invention or technology.

1.1 Examples of complex inventions and the different ways of patenting them

A very simple example of the kind of innovation that I am thinking about is US patent 216,231, applied for in 1879, and covering "improvements in velocipedes" (see appendix). The patent contains four claims, of which at least the first two are entirely independent of each other: one covering a spring attached to the frame, the other covering a "hammock-seat". Each one could have been filed in a separate patent of its own. A patent is required to disclose all that is necessary to enable a person "skilled in the art" to understand and replicate the invention. (USPTO Manual of Patent Examining Procedure ("Examination Manual"), Section 2164). In this example, this requirement would also be satisfied by a patent only containing the first or the second claim. This example also shows that the "complex" products that we are having in mind do not have to be "high-tech" innovations, even though computers and other modern electronic consumer products are very common examples throughout the literature.

Another example that caught the attention of the legal literature is US patent 7,280,838, assigned to Helferich Patent Licensing, LLC. This patent consists of two sorts of claims: "device claims" and "content claims". Helferich used to enter licensing agreements not including the whole patent, but rather single claims of it.[5] Even though separate licensing

protection easier to understand and to enforce, but they do not represent separate inventions. The distinction between independent and dependent claims is a crucial element of the present study.

[4] A complex product in the present context is a product that consists of multiple independently invented components, but those components have a greater value (or any value at all) if commercialised together. A more detailed definition is given in section 3.

[5] In 2015, the Court of Appeals for the Federal Circuit ruled (in Helferich Patent Licensing LLC v. New York Times Co., 2015 U.S. App. LEXIS 2047, Fed. Cir. Feb. 10, 2015) that the doctrine of "patent

tells use that both types of claims are apparently economically interesting on their own, the patent examiner did not have doubts regarding unity of invention.[6]

This concept of unity of invention implies that all content of a single patent should be an embodiment of the same "inventive concept", i.e., the same basic invention (idea) should be underlying all components (apparatus, devices, "objects") covered by the patent. If a patent application contains two separate inventions, the patent examiner should issue a restriction requirement, and the applicant needs to chose one of the two inventions to be prosecuted further. They are free to file an additional "divisional" patent application for the second invention (Examination Manual, Chapter 0800 and section 1850).

Or, this is what should be happening. Binding resource constraints imply that the process of patent examination is not perfect (Lemley, 2001; Lemley & Shapiro, 2005; Lemley & Sampat, 2012). US patent examiners have a very tight time budget that they can allocate to any patent application that they have to examine (Frakes & Wasserman, 2017). Against this background, the existence of granted patents violating the concept of unity of invention is not hard to imagine. In fact, I show that the number of divisional patent applications decreases as the number of claims per patent decrease.

Beyond that, unity of invention might just be very hard to judge. The concept is not very well defined and leaves ample room for interpretation and discretion by patent examiners (Scott & de Jonge, 2008)[7] But these are "static" considerations of granted patents. The main question of interest to be answered by this study is how an innovative firm files its patents differently if they decide to opt for secrecy for one additional of their new inventions. Would they necessarily forego one or even several of their potential patents?

The literature on intellectual property use agrees that methods are more easily kept secret than physical goods (Arundel & Kabla, 1998). Each of the four patents prevalent in the Google v. Uber case contains a method, but when keeping all those methods secret, Google would still have published a large amount of "device" patents.[8] At the same time, Google would have prevented a sizeable amount of manufacturing knowledge to become available to the public. What about the two patents with more than a single independent device claim: could it be that they had kept one of the devices secret instead? This certainly depends on the similarity of the two claimed inventions and the level of obviousness by which one can be inferred from the other.

exhaustion" does not apply to cases in which a prior license is granted only for a subset of claims, but not for all claims in a patent (Bartz, 2015; Breen & Stockwell, 2015; Ciardullo, 2015).

[6]From the examination records available at the USPTO's Public Pair website, the examiner did have concerns regarding potential "double patenting" of a subset of claims (54–61) that were very similar to those included in another patent application filed by the same applicant in 1997, but granted protection for the remaining claims.

[7]Scott & de Jonge (2008) construct a hypothetical chain of inventions sharing the same inventive concept that leads from a new molecule (chemical formula) via artificial grass made from this molecule to sporting equipment particularly tailored to that artificial grass. They even make the stretch to imagine including a whole stadium, as long as this stadium has features that related to the grass made of the molecule.

[8]According to Ohnsman (2017), Google has 176 patents in total that relate to "LiDAR" systems, the kind of laser systems at the centre of the Google v. Uber lawsuit, and 260 further patents in the broader field of autonomous cars.

To answer the question, it seems not too far-fetched to believe that Google might have kept one or two independent claims secret, rather than dropping patent protection of (parts of) their laser systems completely. Patents are examined and litigated claim-by-claim. Licensing decisions have been made on a per-claim basis as well (Dahlin, 2015; Schmidt, 2015). Making patenting decisions on a per-claim basis therefore only appears natural.

1.2 Empirical Strategy and Results

Building on a recent effort by Png (2017a, 2017b), the enactment of the Uniform Trade Secrets Act across US states is used as a quasi-experiment to study the effect of strengthened protection of trade secrets on firms' patenting behaviour. Empirical results point towards a robust negative effect of increased secrecy protection on the number of claims in complex product industries, whereas estimated effects on patents depend heavily on the regression specification and the selected sample, while several extensions show that the effect on claims depends intuitively on characteristics of invention and assignees.

This means that when we want to learn empirically about the relation between secrets and patents, we have to consider the claims—and patent content more generally—to capture the whole effect. The findings do not only contribute to the literature on the economics of innovation, but will also be of interest to business and legal scholars seeking to understand firms' approach to "designing" patent documents. Furthermore, accepting the idea that the content of patent applications is indeed, to some extent, a strategic variable means that using patent counts alone as proxies for innovative output can severely underestimate (changes to) the amount of knowledge created in complex product industries. The validity of the number of claims as an exogenous explanatory variable in patenting studies is challenged as well.

The next section summarises the literature and clarifies the motivation to focus on patent breadth in complex product industries. Section 3 expands on this by presenting a simple model to derive hypotheses about firms' patenting decisions and how these are affected by the interplay between the number of patents and claims. Section 4 outlines the identification strategy and gives a brief overview of US IP law during the period of study. The specific empirical setup and the data are described in 5, while results are presented in the various parts of section 6. In section 7, I outline with two examples how using claim counts instead of patent counts may change results of studies based on patent data. The final section concludes.

2 State of the Literature...

2.1 ...on interpretation and choice of patent breadth and claims

Compared to the literature on the patenting decision (reviewed in Hall et al., 2014), much less has been written specifically on the choice of patent breadth. It seems advisable to clarify the issue at the heart of this paper by distinguishing between different concepts of patent breadth used in the literature. One strand of literature analyses the incentive effect and optimal extent of *statutory patent breadth* (i. a. Gilbert and Shapiro, 1990;

Klemperer, 1990; Gallini, 1992; Denicolò, 1996, 2000). This kind of breadth represents the technological distance that patent law requires between any single patent and non-infringing imitations and is chosen by policy makers, but not the patentee himself. Even fewer studies exist regarding a patentee's choice about how much information to include in a single patent document, which is the definition of patent breadth used in the present study. On the empirical side, the first such contributions are Lerner (1994), who suggests to use the number of patent technology (IPC) classes that a patent is assigned to by the patent office as a proxy variable for the scope of protection granted by a single patent, and Tong & Frame (1994), who argue that the total number of patent claims is a superior indicator of a country's technological capacity than just the number of patents. A third concept of patent breadth is put forward by Kuhn & Thompson (2017) who show that the scope of a patent *relative to* the—unobservable—true scope of the invention is best approximated by the length of claims rather than their number. By contrast, the present paper is interested in whether the different "pieces" of an invention go into a common patent or are patented separately. The length of the claim text relative to an objectively observable invention indeed changes how narrowly an invention is defined, but there is no obvious reason why this should be affected by secrecy law, assuming that each claim itself is examined regarding full disclosure of the underlying invention.[9]

In most existing empirical studies, the number of claims is commonly regarded as a technologically determined parameter outside the control of the patentee. Potential endogeneity of the number of claims is acknowledged to some extent by Lanjouw & Schankerman (2001) and Sakakibara & Branstetter (2001). The first two authors argue that claims can be regarded as a measure of the amount of knowledge protected by a single patent, and find that additional claims increase the likelihood of the patent being involved in litigation. Sakakibara & Branstetter (2001) analyse the 1988 Japanese patent law reform that allowed patent applicants to include multiple independent claims in a single patent document. The aggregate amount of patent applications to the Japanese Patent Office is found to have levelled off while the number of claims per patent increased considerably, indicating that patentees changed the way of protecting a given invention, combining multiple claims together in a single patent that would have constituted different patents prior to the reform.

More recently, de Rassenfosse & van Pottelsberghe (2007, 2012, 2013) and van Pottelsberghe & François (2009) analyse the effects of patenting costs on applicant behaviour, emphasising the role that patent fees can play as a policy tool. These findings motivate the role of patent fees in the model developed in the following section. Finally, a small number of studies tries to explain the number of claims. Archontopoulos et al. (2007) and van Zeebroeck et al. (2008, 2009) study the determinants of the increase in the "voluminosity" of patents over time, which mainly refers to the observed increase in the number of claims per patent over the past decades. Most importantly, van Zeebroeck et al. (2009) use the number of claims (and pages) contained in patent applications as dependent variable in their empirical specifications. The present study takes their findings as inspiration for control variables.

[9] Out of curiosity and for completeness, I additionally regress claim length, but find no statistically significant effect of the strength of secrecy law.

2.2 ...on the choice between patenting and secrecy

Since trade secrets are per se unobservable, much of the existing evidence stems from business surveys (Risch, 2015). Three major results emerge from virtually all of these surveys that are particularly relevant for this paper: responding firms indicate that 1) a considerable share of patentable inventions is not patented, 2) secrecy is at least as important as patenting, and 3) firms often do not regard the different forms of IP protection as mutually exclusive (Hong et al., 2012; Hall et al., 2014; cf. also Anton et al., 2006; Graham & Somaya, 2006; Levin et al., 1987). These stylised facts are important hints towards the existence of a more complex relationship between R&D output and patents, as well as between patents and secrecy, but the survey results are obtained at the firm level, which prevents inference to be made about the protection decision for single inventions.

Crass et al. (2016) avoid this problem by restricting their sample to responding firms that stated to have introduced a single innovation in the period of study. Empirical studies going beyond survey data have been very sparse. One solution to the problem of unobservability of TS is to examine how firms' patent applications change in response to changes in the law governing the protection of TS (Risch, 2015). Such an restricts identification of the effect of trade secrets law to those inventions that technically lend themselves to both patent protection and secrecy. At the same time, however, besides identifying the additional use of secrecy, this approach allows to generate insights into patenting behaviour itself. Most closely related to the present study, Png (2016; 2017b) uses panel data of patenting firms to study the effect of strengthening TS law across US states. In the 2016 working paper he finds decreased patenting in complex product industries, but no significant effect for discrete product industries; while the 2017 journal article finds a negative effect for both complex and discrete product industries only after instrumenting the change in secrecy law.

Analysing a very similar sample, Dass et al. (2015) find evidence of patenting reacting positively to strengthened patent law and negatively to strengthened TS law. Restricting attention to patent law, Gamba (2016) identifies a positive, though transitory, effect of increased protection on subsequent pharmaceutical patent applications. The validity of the assumption of increased use of a protection measure in reaction to an increased level of legal protection is crucial for similar studies on trade secrets; at the same time it cannot be tested.

The survey by Hall et al. (2014) also contains a number of theoretical contributions that analyse a firm's choice between patenting and secrecy for a single discrete invention. For this kind of inventions, patenting and secrecy represent mutually exclusive alternatives. A small number of contributions—not included in the Hall et al. survey—treats *patents and secrets as complements*[10] and allows firms to combine the two protection measures for a single but complex invention that consists of multiple separately patentable components that are commercialised together[11] (Ottoz & Cugno, 2008, 2011;

[10] Patents and secrets are complements in the way that combining them allows to reach a higher level of total protection of an invention than would be possible to achieve with patent or secrecy protection alone.

[11] Such complex multi-component innovations are also analysed in the literature as "complementary innovations". It seems that this term is used when the focus is put on the incentives to invest in research

Kwon, 2012a, 2012b; Belleflamme & Bloch, 2013; Png, 2017b; Lee, 2017). For each single component, the basic notion remains that patenting and secrecy are mutually exclusive, but it is possible to derive predictions concerning the mixture of patents and secrets at the invention level. This is also what motivates the model presented in the subsequent section.

3 A Simple Model of Patent Design: Allocating Complex Inventions to Patents

3.1 The Need for a Model

The model developed in the following extends the literature summarised in the preceding section by allowing multiple components of a complex invention to be covered by the same patent. Based on a patentee's motive to minimise their patenting cost, it shows that the reaction to a change in the strength of secrecy law can well be in the number of claims per patent rather than in the number of patents themselves. Therefore, the assumption of a 1:1 relation between protected components and the number of patents can no longer be made when interpreting patent data: a change in the number of protected components can leave the number of patents unaffected.

Within this section, an important distinction is made between

k	the number of protected components,
p	the number of patents, and
c	the average number of claims per patent,

where $k = p \cdot c$.[12] Fundamentally, the components of a complex invention can be protected by many narrow patents (high p, low c), by few broad patents (low p, high c), or any combination in between these extremes. The narrower are patents, i.e., the more patents are required to protect a given invention, the higher is the average marginal cost of patenting a single component: applying for an additional patent requires payment of examination, search and issuance fees, maintenance fees later in the life of the patent, as well as some lawyer fees that are computed on a per-patent basis.[13] Adding an independent claim is essentially free for the first claim. If the number of independent claims increases, some "penalty" fees will be charged by the patent office (starting from the second claim until 1981, starting from the fourth claim thereafter), but these are much lower in value than the per-patent fees (see table 1.4 for an overview of the USPTO fee schedule during the sample period). This alone will lead to a lower number of patent components k^* in

and development, while "complex innovations" is used when the focus is on the IP-protection decision.

[12]Since this model is meant to motivate empirical analyses, it neglects the fact that c varies across patents. Most of empirical analysis only considers the mean. As an extension, later in the paper the variance of claims per patent is specifically taken into account.

[13]The total amount of fees does not seem important for patents that are expected to generate millions in revenue, but the average patent value is very low , particularly in the US where patents are less thoroughly examined (thus more "probabilistic" (Lemley & Shapiro, 2013)). There is a literature on the "fee elasticity of patents" (de Rassenfosse & von Pottelsberghe, 2013) finding a significant and non-negligible negative effect of patent fees on patent propensity.

equilibrium. Depending on this number k^*, the impact of a given change in secrecy law differs across firms.

Another element of patenting cost, however, is the effort required to convince the patent examiner of patentability of the filed inventions and to get the patent granted. Apart from novelty and non-obviousness, what is claimed in a single patent needs to satisfy "unity of invention", i.e., it needs to be related by a common inventive concept. A common example of two independent claims featuring unity of invention are a production method and the product resulting from this method. Another example are two different devices that have both been invented by the applicant. These two devices can only be included in the same patent if they are either technically related (e.g., variations of each other), or if they are necessarily combined to achieve a certain result. Now, given a certain complex invention, it seems reasonable that successfully arguing for unity of invention becomes harder the more components are included in a single patent. Similarly, leaving out essential parts by patenting them separately when, e.g., trying to patent results rather than separate mechanical components requires more argumentative effort. From this follows that both the cost of claims and the cost of separate patents are convex functions. The degree of convexity differs across inventions (and, on average, across technologies or industries).

The motivating question for why we need a model is the following: how should we interpret the observation that a firm, in a certain year, is granted three patents with three independent claims each? Does this correspond to three (discrete) inventions, or is it one (complex) invention covered by three patents? Png (2016) finds that the likelihood of a reaction in the number of patents is greater in complex inventions than in discrete inventions.

Furthermore, how many components that belong to the ones included in the three patents are being kept secret? Png (2017b) assumes that the lower is the share of patented components in the original situation, the less strong is the expected reaction to an increase in secrecy protection, all else equal, for strategic patenting motives: a minimum share of patent protection is expected to remain as residual even if secrecy protection is much stronger than patent protection.

Empirically, these two characteristics can be controlled for by the industry-level measures of product complexity and patent and secrecy effectiveness developed by Cohen et al. (2000). R&D spending is only a meaningful control when focussing on the number of patents, but not with respect to patent content, as will be seen. Without controlling for patent and secrecy effectiveness, the following two situations involving a low number of patents are empirically indistinguishable: (1) low number of patented components because the patented technology on average does not lend itself to patenting very well; and (2) low number of patented components because the patented technology requires many separate patents, i.e., few components can be put in a common patent only. This increases the cost of patenting per component.

Patent fees are identical for all applicants,[14] so when controlling for complexity and IP protection effectiveness, observed differences in claim counts can be interpreted as stemming from differences in the degree of convexity of the patent and claim cost func-

[14] They can be reduced by 50% for "small entities", but this discount applies for all relevant fees, so their relative size remains unchanged.

tions. Left as a requirement is a model predicting the effect of increased strength of legal secrecy protection on patents with different patenting cost functions and resulting low as opposed to large numbers of independent claims.

3.2 Model Setup

The model suggested in the following is similar to Png (2017) and Lee (2017), but shifts emphasis to the patenting cost function. As long as the number of patentable components is not very small, the model has an interior solution that maximises aggregate protection of a complex innovation against imitation involving a number $k < K$ of all K components that is patented. Figure 1.1 depicts this general idea.

Neither way of protection is perfect, so a patented component can be invented around, while a secret component can be reverse-engineered, or the secret accidentally leaks. Components differ in the way they benefit from either patent or secrecy protection. Suppose now that the components can be ordered in decreasing probability of inventing around and at the same time increasing probability of reverse engineering/trade secrecy leakage. When the firm patents k of its K components, the total aggregate probability of a competitor inventing around all patents is given by the function $\alpha(k)$, and the aggregate probability that the remaining components that are kept secret are reverse-engineered is given by $\rho(k)$. $\alpha'(k) < 0$, $\alpha''(k) > 0$ and $\rho'(k) > 0$, $\rho''(k) > 0$,[15] implying that the two functions cross each other exactly once.

The smaller is the number of components K, the more likely is a corner solution involving $k^* = K$. This is what characterises purely discrete technologies where $K = 1$. The greater is K, on the other hand, the more likely is it that some components will be kept secret. These latter complex products are in the focus of the present study.

The innovating firm maximises the following profit function:

$$\pi(k) = [1 - \alpha(k)\ \rho(k)]\, M - F_p(p) - pF_c(c). \tag{1.1}$$

The first term is the payoff from successfully protecting the complex innovation from imitation, yielding some profit M. The cost of patent protection consists of two terms: (1) The cost of independent claims per patent, $F_c(c)$, is strictly convex. Its convexity comes from two sources: first, the patent office charges claim-based fees only above some threshold (1 in the first years of the sample period, 3 in the later years). Second, independent claims can only be included in the same patent if they meet the requirement of "unity of invention", which is one of the requirements that the patent examiner needs to assess. Convex costs capture the idea that with an increasing number of claims it becomes ever more difficult to convince a patent examiner of unity of invention. If claim drafting is delegated to legal representatives, this can directly translate to greater monetary costs charged for their service.[16]

[15] ρ could alternatively be defined as $\rho(K - k)$ with $\rho'(K - k) < 0$, $\rho''(K - k) > 0$. This is equivalent to the definition used in the main text when all components of the invention are either patented or kept secret, and there are no "disclosed but unprotected" components as in Lee (2017).

[16] Convexity does not need to capture riskiness of increasing the claim count in a patent application: if a patent examiner feels that unity of invention is violated, they require the applicant to choose between the different inventions identified in the patent application, but are free to include the remaining claims in

(2) The cost of separate patent applications $F_p(p)$ captures the direct monetary and effort costs of drafting and filing patent applications (search, examination, issuing, potentially maintenance, ...) and is assumed to be a convex function. Its convexity can be thought of as being caused by several features. The more inventions are split up into different patents, the less each one of them looks like a whole invention, making it harder to have them patented. Additionally, the more patents are applied for, the more one needs to describe how single components work without referring to the whole underlying innovation from which this might have been obvious, increasing effort needed for each patent application.

The decision about patent protection occurs in two stages: first the decision is on the amount of total patent protection k, then this k needs to be allocated between p and c. Generally, the higher are costs of patenting, the lower will be the number of patented components. Given costs, patenting increases in the value of the function α and decreases in the value of the function ρ, which is how the model captures the strength of trade secrets protection. The focus in the following will be on the cost function.

Sufficiently convex cost functions avoids the trivial results of either a single extremely broad patent or very many single-claim patents. The precise shape of the cost functions is specific to the protected technology, but a profit-maximising strategy can only involve a number of claims c that still yields $F_c'(c) < F_p'(p)$. This means that the shape of F_c affects the marginal cost of patent protection of a component: as long as $F_c'(c) < F_p'(p)$, it is cheaper to protect a component by adding it to an existing patent that would have been applied for anyway as compared to putting it into a new patent of its own. Averaged over all components of the invention, the higher is the highest c that still yields $F_c'(c) < F_p'(p)$, the lower is the marginal cost of patent protection for that specific complex innovation.

The higher will be the marginal cost of patenting of either p or c, the lower will be the optimal value of the respective variable. The corresponding elasticities are

$$\frac{\partial p}{\partial k}\frac{k}{p} = \frac{F_c''(k/p)}{F_c''(k/p) + \frac{p^3}{k^2}F_p''(p)} \quad \text{and} \tag{1.2}$$

$$\frac{\partial c}{\partial k}\frac{k}{c} = \frac{F_p''(k/c)}{F_p''(k/c) + \frac{c^3}{k}F_c''(c)}, \tag{1.3}$$

which are clearly both $\in (0,1)$. The reaction in a variable is stronger the lower is the degree of convexity of its cost function.

An explicit solution is obtained by setting $F_p(p) = f_p p^\delta$ and $F_c(c) = f_c c^\varepsilon$, where $f_p > f_c > 0$ and $\delta, \varepsilon > 1$. k^* is optimally allocated to p and c according to

$$p(k) = \left(\frac{\varepsilon - 1}{\delta}\frac{f_c}{f_p}k^\varepsilon\right)^{1/(\delta+\varepsilon-1)} \tag{1.4}$$

$$c(k) = \left(\frac{\delta}{\varepsilon - 1}\frac{f_p}{f_c}k^{\delta-1}\right)^{1/(\delta+\varepsilon-1)} \tag{1.5}$$

one or multiple separate "divisional" patent applications. As long as this additional administrative step does not significantly increase transaction costs, there is "nothing to lose" by including many independent claims in a patent application.

$p(k)$ increases in ε and decreases in δ, $c(k)$ reacts in the opposite direction. The elasticities are given by the constant terms

$$\frac{\partial p}{\partial k}\frac{k}{p} = \frac{\varepsilon}{\delta + \varepsilon - 1} \tag{1.6}$$

$$\frac{\partial c}{\partial k}\frac{k}{c} = \frac{\delta - 1}{\delta + \varepsilon - 1}. \tag{1.7}$$

Again, elasticities are both $\in (0, 1)$ and determined by the relative degrees of convexity. Increasing the degree of convexity also increases marginal costs, therefore the simple prediction arises that the higher is convexity, (1) the lower will be the optimal value of the variable (p or c), and (2) the weaker will be the reaction in that variable to a change in k.

It shall be noted that the exponential nature of the two cost functions is not strictly required to obtain a differential reaction in patents and claims. If we introduce a probability of patent granting, and if that probability is an inverse-U shaped function of the number of claims included in a patent, then similar predictions as above can be obtained for a wide range of functional parameters.

3.3 Empirical Predictions

This subsection translates the above model into empirically testable hypotheses. The first two hypotheses arise from the assumptions underlying the model:

Hypothesis 1. *The design of patents reacts to changes in the USPTO fee schedule.*

This has been shown by previous studies, so it should happen here as well.

Hypothesis 2. *An increase in secrecy law affects the share of patented components more strongly in complex product industries. If there is a reaction in discrete industries, it will be in claims.*

In the model developed in the previous section, the variable with the less convex cost function reacts more strongly to a change in secrecy law. To the extent that greater convexity also means higher marginal cost, the model yields the following two complementary hypotheses:

Hypothesis 3. *In complex product industries, strengthening secrecy law will reduce the number of patents when the patentee protects their pre-treatment innovations by a high number of separate patents. This effect is particularly strong when the patents have a low number of claims per patent.*

Hypothesis 4. *Equivalently, strengthening secrecy law will reduce the number of claims when the patentee protects their pre-treatment innovations by a low number of separate patents. This effect is particularly strong when the patents have a high number of claims per patent.*

Of course, these firms also would prefer to reduce their patenting costs by reducing the number of patent applications, but due to the low number of patents to begin with, reducing this number even further would likely lead to a too strong reduction in the level of protection of the complex invention. At least those firms with a large enough number of claims per patent should therefore choose to forego the reduction in costs in favour of keeping the target level of patent protection at the optimal level.

Hypothesis 5. *On aggregate, the set of firms reacting in claims contributes a greater reduction of the total number of patented claims than the set of firms reacting in patents.*

The above hypotheses are supposed to be valid for a "representative" firm. In any particular case the effects may differ considerably. Firms in environments (industries, technological areas, etc.) with few current patentees and low fragmentation of existing patent ownership ("easy-to-navigate" environments) may feel less need to use large patent portfolios for defensive reasons or holding-up of competitors and accordingly show a greater propensity to substitute some of their patent applications for secrecy.

Furthermore, the splitting or not of an invention into multiple patents and claims affects the observed reaction. The more similar are two patents, the less is aggregate patent protection reduced if one of them was missing. If claims within a patent are rather similar, the same logic applies to them.

A final remark needs to be made about the role of R&D spending in regressions. When explaining the number of patents, controlling for R&D is crucial, because increasing R&D spending, all else equal, yields *additional* inventions, which will *have to* be included in *additional* patents. When explaining the number of claims per patent, R&D is not of much help.

Even if strengthened secrecy law does increase the *amount* of R&D spending (Png, 2017a), this does not affect the number of independent claims per patent when the additional inventions are technologically similar to existing R&D projects, and only their number increases. There is no reason to expect a non-random change in the number of patented components per invention.

And even if strengthened secrecy law does affect the *nature* of (additional) inventions, this would be more of a categorical effect rather than a linear effect similar to that of R&D on the number of patents. It might be, e.g., that stronger secrecy protection makes firms shift their R&D activities to more complicated inventions which consist of a greater number of components (higher K). In such a case, a higher k^* could be expected. Similarly, firms could be induced to shift their R&D to inventions that rely more on secrecy for technological reasons, which would lead one to expect a lower k^*. Such changes in the nature of inventions cannot be controlled for by R&D spending.

If these changes entail a change in technology classes, a difference between the most common before-treatment and after-treatment technology class could be used as a dummy variable. If technology classes do not change, it seems hard to think of an alternative control.[17] In any case, such an effect would be ascribed to the legal (treatment) variable. Therefore:

[17] A potential way to approach this problem could be via searching for non-random changes in keywords in patent texts used to describe inventions. Such an analysis is beyond the scope of this project.

Hypothesis 6. *R&D expenditure is a crucial control in the regression of patents, but is not important in the regression of claims.*

4 Identification and Legal Framework

4.1 General Identification Strategy

Until the coming into force of the Defend Trade Secrets Act of 2016, trade secrets protection was in the exclusive domain of state law (Denicola, 2011), giving rise to a great variety of legal provisions, providing different levels of protection between states and changing at different points in time. Patent law, on the contrary, is exclusively federal law, making all US firms equally subject to its changes. This unique setting allows applying a difference-in-differences identification strategy with multiple treatment groups while controlling for patent law via time-fixed effects. Van Zeebroeck et al. (2009) note that applicants seeking patent protection at multiple patent offices across the world show the tendency to orient the design of patent documents towards what is common in the US, increasing the external validity of research findings obtained with US data.

4.2 US Trade Secret and Patent Law

The degree of protection of trade secrets and the availability of legal precedent differ(ed) greatly between US states (Denicola, 2011). The UTSA was drafted out of an initiative to change this and to harmonise trade secret protection across states. In many states, it extended the definition of trade secrets, increased their protection, and clarified available legal remedies, including punitive damages. Furthermore, it UTSA included regulations concerning TS litigation. This and the fact that legal precedent was replaced by codified law reduced uncertainty about the ability to defend trade secrets in any one state.

It was not mandatory for states to implement the UTSA, but a majority did. Tables 1.1-1.3 give an overview of the changes occurring to legal TS protection in the US during the sample period 1990–1999, as measured by Png's (2015) index, covering 39 of the 47 legal changes caused by state-wise implementation of the UTSA. This index considers six different components of TS law that are each assessed on a scale between 0 and 1 and then averaged to yield the index value for each state and year. The components are "(i) Whether a trade secret must be in continuous business use; (ii) Whether the owner must take reasonable efforts to protect the secret; and (iii) Whether mere acquisition of the secret is misappropriation; (iv) The limitation on the time for the owner to take legal action for misappropriation; (v) Whether an injunction is limited to eliminating the advantage from misappropriation; and (vi) the multiple of actual damages available in punitive damages" (Png, 2017b). To allow more precise estimates, the main treatment variable measures the protection due to the UTSA *on top of* the existing protection level via legal precedent (common law). The appendix to Png (2017a) contains a number of robustness checks indicating exogeneity of UTSA enactment in the states that decided to do so to firm lobbying activity and other state policies directed at R&D and the general business climate.

5 Empirical Setup

5.1 Main Regression Specification

The main specification used in estimations is a panel difference-in-differences setup. Unless stated otherwise, every regression is performed twice, once with the number of yearly (eventually granted) patent applications and once with the average number of independent claims in these applications (simply referred to as "patents" and "claims" in the following) of company i in state s and year t as the variable to be explained:

$$\ln(\text{patents})_{ist} = \lambda_{is} + \mu_t + \beta X_{ist} + \gamma \text{UTSA}_{st-1} + \varepsilon_{ist}. \tag{1.8}$$

$$\ln(\text{claims})_{ist} = \lambda_{is} + \mu_t + \beta X_{ist} + \gamma \text{UTSA}_{st-1} + \varepsilon_{ist}. \tag{1.9}$$

λ_{is} and μ_t are firm-per-state and year fixed effects, X_{it} is the vector containing the time-varying covariates at the firm level, and UTSA_{st-1} contains the "increase in the legal protection of trade secrets arising from the UTSA being in effect" (Png, 2017b) in state s in the respective previous year. This means that this variable captures solely the effect of a state adopting the UTSA, as the level of protection granted by common law or a statute different from the UTSA is represented by different variables in Png (2015). Alternatively, the index capturing total protection level can be used, but exogeneity of treatment would be much harder to argue for that case. γ is the coefficient of interest in this study. Standard errors will be estimated clustered at both the state and company-level, unless otherwise stated. The number of claims per patent are very close to normally distributed when taken the log of, while the (log of the) number of patent applications follows a heavily right-skewed distribution. For this reason, specification (1.8) is additionally estimated by Poisson regression.

5.2 Selection of Sample

The study covers the years 1979–1999. 1999 is chosen as the final year of observation due to the enactment of the American Inventors Protection Act in late 1999. Among other issues, this Act has patent applications published 18 months after their filing date, irrespective of whether they lead to a patent grant or not.[18] This potentially fundamentally changes the incentives to patent or to keep secret (Graham & Hegde, 2015). In fact, since now only those inventions will be patented for which the applicants are either certain enough that they satisfy the requirements of patentability or they are much harder to keep secret, the time after 1999 might make an even better study period. However, there is much more variability in the trade secrets protection index before 1999 since only four states enact the UTSA after that year.[19]

Patentees are defined by the [*gvkey*] firm identifier used in the Compustat company information database. The two identifiers are not perfectly congruent, one [*gvkey*] variable may cover more than one value of [*pdpass*]; [*pdpass*] therefore captures patentees

[18] This brought the American patent system closer in line with most other patent systems around the world. However, even after 1999 it is still possible to keep an application secret by opting out of the possibility of additionally seeking patent protection outside of the US.

[19] These four states being Pennsylvania (2004), Wyoming (2006), New Jersey (2012) and Texas (2013).

at the subsidiary level if such a level exists. As the setup exploits legal changes at the state level, a question is how to treat the fact that most firms have some sort of relation to more than one state. Png (2017b) aggregates patenting at the parent-company level and constructs company-specific treatment variables as the average of the UTSA index weighted by the share of employees a company has in any given state. Data on employee shares is taken from the Bowkers directory of private R&D facilities, which is a reliable source, but covers mostly very large firms, leading to a considerable reduction in sample size, which proves particularly problematic when additionally conditioning on pre-treatment patenting of firms.

A very straightforward alternative (followed by Dass et al., 2015) uses exclusively the trade secrets protection index (TSPI) of the state in which the headquarters of the firm are located. Assuming that headquarters must be to a certain extent be informed about trade secrets used in the company, there is certainly a strong correlation to be assumed between the local TSPI and the actual level of secrecy protection that the company enjoys. In the case of using patent data without Compustat company information, there is no information on the headquarter location, so a workaround would be to restrict the sample to firms patenting in a single state.

It is generally possible to use inventor addresses contained in patent documents to increase information about the exposition of a firm to trade secret regimes of states other than their home state. On the one hand it seems credible that an inventor residing in a different state is subject to the state law of its employer. On the other hand, around 90% of patentees have at least one inventor coming from a different state than the patentee, so this is not a feasible sample restriction. More feasible is restricting the sample to patentee-years in which all foreign-state inventors, in the present and all future years, come from states with at least as strong trade secrets protection as the patentee's home state, which still leaves about half of the sample of all US patentees.

5.3 Data and Main Variables

This study builds on the dataset of all granted patent applications to the US Patent and Trademark Office (USPTO) between 1976 and 2006, made available by the the NBER patent data project (Hall et al., 2001), which determines unique identifiers for all patent assignees, allowing empirical analysis at the assignee level without having further information on assignee characteristics. These data also include a match to the company identifiers of Compustat North America (Bessen, 2009), allowing to add fundamental financial data for publicly listed US companies. This matching yields richer data, but at the cost of heavily reducing size and representativeness of the sample (2,160 vs. 20,598 assignee-state locations).

The study uses for the main part the dataset including additional firm level information, but uses the larger patent data-only dataset as a robustness check. For the latter case, patent data has been restricted to non-government patent assignees with a US postal address at the time of application (state code "US", assignee code "2"), and with patent information aggregated at the level of state-locations per unique assignee code (variable [*pdpass*]). Ideally, patent data should be aggregated at the level of corporate headquarters when studying the protection decision if this decision is taken at the company level and the headquarters of a company are informed about about trade secrets held at the company's subsidiaries. However, grouping companies according to their previous (pre-treatment) patenting behaviour—explained in detail in section 6.2—requires information on the respective date of treatment, which is state-specific. There is no obvious way of aggregating the treatment date and pre-treatment data, and the weighting of states' TS laws allows for additional discretion. Section **??** contains a robustness check in which the sample is restricted to assignees that are active in a single state only.

As argued in section 3, the extent of a reaction to changes in secrecy law can be greater in complex product industries, while the nature of discrete products makes it more likely that—absent dramatic shifts in the legal regimes—the decision is restricted to including or not production methods and other "method" claims in a patent. Patentees therefore need to be assigned to one of the two industries. The classification of complex and discrete industries by Cohen et al. (2000) is based on the International Standard Industrial Classification (ISIC) and defines as complex any industry with an ISIC code of 2900 or greater. The subsample matched to Compustat data contains a unique industry code for each company. While for smaller companies this seems an appropriate description, larger companies may well be active in more than one six-digit NAICS industry. Furthermore, for the larger sample of patent data, only the technology classes (IPC) to which a patent is assigned by the patent office are available. Accordingly, the ISIC classes are matched to IPC classes via a concordance table (Lybbert & Zolas, 2014). A first indicator of the nature of the products commonly assigned to a class is therefore the share of ISIC codes greater or equal to 2900 out of all assigned ISIC codes, as defined in Cohen et al. (2000). A second indicator is the previous indicator multiplied by the probability weights (provided by Lybbert & Zolas, 2014) of all matches to complex product industry ISIC codes relative to the sum of all probability weights of that IPC class. The degree of complexity of an assignee's innovations is then determined as the respective average of each of these indicators. Alternatively, for every assignee the most commonly used four-digit IPC class is determined over the whole sample period. Both complexity indicators are matched to assignees based on this most common IPC class. For each of the two indicators a dummy variable is created ($[complex1]$ for the first and $[complex2]$ for the second, accounting for probability weights) that takes the value of one if the respective indicator is greater than 0.5, and zero otherwise.

The time dimension of the sample is restricted to patent applications from 1979 to 1999. Starting in 2001 there is a notable decrease in the number of patent applications contained in the NBER data which does not correspond to the development of real patent applications but results from truncation of the data on granted patent applications available at the time of creating the NBER dataset. Furthermore the entry into force of the American Inventors Protection Act at the end of the year 1999 brought with it a number of fundamental changes of patent law (Gallini, 2002) which impede identification of the effect of secrecy law on patenting behaviour (Johnson & Popp, 2003; Graham & Hegde, 2015).[20]

Further sources of data are the research datasets recently made available by the USPTO, containing the full text of patent claims (described in Marco et al., 2016) and

[20]The most important change for the topic of this paper is the introduction of mandatory publication of any patent application 18 months after filing, unless the applicant agrees not to seek patent application outside of the United States.

a summary of the patent examination process (described in Graham et al., 2015). This latter dataset contains such information as whether the application led to or stems from a divisional application, whether the patent applicant was eligible for "small entity" fee discounts, etc.

Information on the number of (unique) inventors and the state of their residential address is taken from data made available by Lai et al. (2011). Finally, data assigning patents to classes based on the semantic content of their abstracts is use to proxy for similarity between patents and claims and was made available by Bergeaud et al. (2017).

6 Results

6.1 Data description

Table 1.5 presents summary statistics of the dependent variables and the most important covariates for the main sample. The averages with regards to patents and claims between complex and discrete industries is negligible; if anything, discrete product patents contain fewer independent claims, which seems intuitive. It is noteworthy that firms in discrete product industries are significantly larger in terms of revenue and EBITDA, while average R&D spending is very similar. Both sub-samples feature comparable variability in the legal indexes. The row "assignees" below each industry category is the number of parent companies, as identified by the [*gvkey*] company identifier variable in the Compustat database. The row "assignee-locations" denotes the unit of analysis (the panel identifier) of this study and is created by the state identifier available for each patent.

Figures 1.5 and 1.6 plot the number of patents and claims/patent, averaged over assignees, separately for complex and discrete industries for every state and the District of Columbia and furthermore denote changes in the TS protection index by coloured bars (see description below the tables). Both patents and claims per assignee increase over time, but the trend in claims seems more stable and much less sensitive to extreme events.[21] This is mirrored in the summary tables: patents are characterised by much greater standard deviations than claims; furthermore—at least in the whole sample—within deviation is greater than between variation for patents, while the order is reversed for claims. This indicates an important role of firm and industry fixed effects in the explanation of claims, while strategic considerations besides invention protection are behind firms' decisions to apply for patents. In discrete product industries, in contrast, where the attribution of patents to inventions is much more straightforward, both types of standard deviation are approximately equal.

[21] Compared to the average diagram in each figure, "exceptional" development of claims over time can only be identified for Hawaii, Montana, Nebraska, Nevada and North Dakota (considering only the curves for complex product industries; these states account for less than 2% of all post-treatment observations in these industries), while for patents this is true for California, Delaware, the District of Columbia, Idaho, Illinois, New Jersey, New Mexiko, New York, North Carolina, Pennsylvania, South Dakota, Texas and Washington (accounting for almost exactly 50% of post-treatment observations).

6.2 Baseline Results

Whole sample Tables 1.10 presents results for the most basic regressions of patents and independent claims on the UTSA and common law indexes and a set of firm variables. The distinction between complex and discrete product industries is based on the SIC industry code contained with every firm record in the Compustat database, matched to ISIC industry codes according to a correspondence table by the US Office of Management and Budget (following Png 2017b). ISIC codes greater or equal 2900 are considered "complex product industries" following the definition by Cohen et al. (2000). The columns including "nn." are additionally restricted to non-negative financial values in the Compustat database (except for EBITDA).

None of the estimated coefficients are statistically significantly different from zero, though only the coefficients in complex product industries have negative signs. This finding is broadly in line with Png (2017b) who finds no statistically significant effect without instrumenting the UTSA.

Of the firm-level covariates, R&D expenditure is the only one with significant explanatory power in all samples, strongly determining the number of annual patent applications. The estimated coefficient seems too large to be taken literally. R&D in this (and all other currently included) specification is assigned to state-locations based on the current share of patent applications over all state-locations sharing the same [*gvkey*]. Using pre-treatment shares of patent applications to assign R&D instead roughly halves the coefficient of R&D expenditure, which arguable still seems to high. Size in that specification, however, has a slightly positive effect on patents, while the estimate of the UTSA index changes in the order of 1 decimal place. Regarding independent claims, the UTSA coefficient is virtually independent of the way of R&D allocation, but R&D itself contributes nothing to explain the number of claims per patent (coefficient 0.001, s.e. 0.02), which is in line with hypothesis 8.[22]

Conditioning on pre-treatment patent design According to hypotheses 3 and 4, the observed reaction to stronger trade secrets protection depends on the shape of the two cost functions, which also affect the optimal patent design in general. In an attempt to uncover the effect of existing patent design, the sample is split in various ways following pre-treatment values of patent and claims counts. As a first effort, claims are regressed separately for firms that have above-median and below-median number of patents before the legal change. The median used here is the one of the whole sample (i.e., over all states), but measured before the individual treatment date. The treatment year is set equal to the year of enactment of the UTSA. To include firms from states that never adopt the UTSA or that adopt it outside of the sample period (i.e., in 2000 or later), pre-treatment data for those states is defined as data from the year 1987 or earlier. Exceptions to the latter rule are made for states in which a notable change in secrecy protection occurred from common law or enactment of a statute different from the UTSA. This is the case for Alabama (1987),[23] New Jersey (1994), New York (1995), Pennsylvania (1996), Texas (1982) and Wisconsin (1985).

[22] Accordingly, a revised version of this paper will use this allocation of R&D spending as the baseline specification.

[23] Years in parentheses are the years in which the legal change happened.

Tables 1.11 and 1.12 report the results for six different ways of determining the pre-treatment number of patents for every state-location: (I) the number of patents in the most recent year before treatment with non-zero patents (i.e., years in which the assignee applied for at least one patent); (II) the average number of patents in the most recent four non-missing years before treatment;[24] (III) the average over all non-missing values in the four years immediately preceding the state's treatment year; (IV) the average over all non-missing years before treatment; (V) the average of the four years preceding the state's treatment year, counting years without patent applications as "zero"; and (VI) the average of the whole pre-treatment period, counting years without applications as "zero".

Grouping assignees based on their patenting in the single most recent pre-treatment year as in (I) may capture randomness because of the considerable amount of within-assignee variance in claims and (especially) patents. If assignees are grouped into the "wrong" side of the median based on an exceptional year of patenting, over time there is an effect of these wrongly classified assignees reverting to their average patent and claim numbers. This effect might yield positive estimates of the UTSA coefficient for low and negative estimates for high values of the respective grouping variable that are in fact not caused by the UTSA.

Results in the columns of table 1.11 roughly agree with each other, however, showing that for assignees with many patents per year before treatment, patents in states with higher values of the UTSA index have fewer claims per patent than those with lower UTSA index values. In sharp contrast are results of the same regression for assignees with few patents per year before treatment (table 1.12), for which all seven UTSA estimates are much smaller and less precise than the corresponding estimates in table 1.11.[25]

While for firms with a small number of parallel R&D projects the raw number of patents may be a good enough proxy of their patenting propensity, this is less and less so the greater the diversity of R&D. Selecting firms based on their patenting propensity—the number of annual patents divided by their deflated annual R&D spending— yields very comparable results, however, presented in table 1.13:[26] again, the UTSA is associated with considerably fewer claims per patent when patenting propensity was low rather than high.

Conditioning on *state-specific* pre-treatment patent design In order to generate even deeper insights, samples are not only divided by their overall medians, but rather by quintiles, and the quintiles differ across states. This seems appropriate, because otherwise—in an extreme case—firms in one state may apply for considerably more patents than firms in another state, and when selecting firms based on overall sample medians, the sample therefore might be split along state lines, counteracting the very

[24] If a firm has less than four pre-treatment years of data the average is computed over these fewer years. The same approach is followed for method III.

[25] Conditioning on a single pre-treatment year yields a highly right-skewed distribution of pre-treatment patent numbers, having a median of 1. For this reason, the regressions using this pre-treatment value as sample selector are additionally performed using the 75th percentile instead of the median.

[26] This and future tables only contain results using samples selected according to ways I–IV, since V and VI are not a useful categorisation for the number of claims per patent: this is by definition undefined when there are zero patents applied for.

idea of the identification approach. Independent computation of quintiles of patents and claims generates groups following a rule that is comparable across different samples, at the expense of resulting in an unequal allocation of observations to groups: only relatively few patentees apply for "many broad" patents.

The first such analysis is performed with the number of patents as the dependent variable (table 1.14), which only react negatively (or at all) to the UTSA when the number of claims per patent before treatment was in the lower quintiles of the distribution. The right-hand half of the table contains the same regressions that additionally include state-specific linear time trends. All estimates are greater in absolute value, but the fundamental differences between the cases of many and few independent claims per patent are preserved.

A couple of further tables report very similar results. Table 1.16 contrasts the results of table 1.14 with the case of few pre-treatment patents, showing that there is no decrease associated with the UTSA in these cases. Regressing claims instead (table 1.17) shows that claims react significantly only when their number is high pre-treatment, while there does not seem to be an effect of the number of pre-treatment patents. This changes only slightly when patent propensity is used instead (table 1.15).

To summarise this subsection, there seems to be a sizeable effect of the UTSA on both patents and claims, the size of which depends on the relative size of both variables before the change in secrecy law. Patents do indeed seem to follow quite strongly a pattern set by the number of independent claims included in each of them, pointing towards independent claims as the level of the patenting decision. The number of claims included in each patent decreases relatively independent of the number of patents, indicating that "topical grouping" (figure 1.3) seems to be the more prevalent way of assigning components into patents.

Robustness checks The regressions presented so far used samples selected based on pre-treatment data exclusively including firms that were assigned to "complex product" industries based on their Compustat industry code. To check whether the described pattern is indeed a result of the particular nature of complex products, table 1.18 compares identical regressions performed using complex or discrete product firms . The estimate for patents in column 1 has a not-significant discrete-industries counterpart in column 2. The estimate of claims in column 8 is still significantly different from zero, though smaller in absolute value than the complex-products estimate in column 7. Apart from imperfect categorization of—particularly large—firms into industries by a single three-digit industry code, there would indeed be a certain amount of reaction in discrete industries be expected, since production methods that otherwise would have been patented might instead be kept secret under the UTSA.

To understand by how much simple "reversion to the mean" is driving results, the same regressions (again restricted to complex-product industries) were run using the length of independent claims as a variable. Independent claim length was identified as a measure of patent breadth by Kuhn & Thompson (2017), but as argued in section 2, this definition of patent breadth is different to the one in the focus of the present study, therefore it should be independent of the patenting-secrecy decision. Indeed, no significant differences between the various samples (selected based on pre-treatment patents and claims) were found (table 1.19). Similarly, using pre-treatment independent claim length to select samples does not yield (highly) significantly different estimates, as apparent from comparing neighbouring columns of tables 1.20 and 1.21. These results allow to put additional trust in the previous split-sample results.

Third, the driver behind the differential effect of secrecy law on patents and claims, as argued in section 3 and summarised in hypothesis 1, is the difference in the necessary fees: adding an independent claim to an existing patent application is far less expensive than starting another patent application, even if it contains only one claim. To get some empirical support for this hypothesis, the USPTO fee schedule over the sample period was collected (summarised in table 1.4) and added to the regressions. As tables 1.22 and 1.23 show, the sign of the respective fee variables are negative as expected.[27]

Comparison to the results by Png (2017b) Png (2017b) does, as does the present study, not find a significant reduction of patents in reaction to the UTSA unless instrumenting the legal change by the enactment of other innovation-unrelated uniform state laws. Table 1.24 reports a replication of the baseline results of Png (2017b) in columns 1 and 4, as well as splitting the sample into complex and discrete industries using the same indicator as above. Column 2 shows that even without the instruments, the number of patents show a significant decrease in reaction to the UTSA, while no reaction could be identified in discrete industries. This difference in results is robust to the author's instrumenting strategy. Furthermore, the author performs a Hausman test by including the first-stage residuals of the IV Poisson regression as a regressor in the model. While the author reports a statistic of 3.82 ($p = 0.051$), this statistic is 4.67 ($p = 0.031$) in the complex-products sample, lowering the level of significance at which we can reject the null-hypothesis of a coefficient estimate of zero for the residuals.

A potential way of explaining this difference may be via the summary statistics of this dataset (table 1.6). Similar to the main dataset used above, firms assigned to discrete-product industries are considerably larger, making it seem possible that at least part of the lower endogeneity of the UTSA is explained by firm size. Larger firms may be better informed about IP legislation (in the making) and possess greater in-house expertise to quickly adapt to such changes. If true, this would make it even more appealing to extend analysis to smaller (in particular non-publicly traded) firms, which is at least partly achieved by analysing the whole NBER patent dataset in section 6.5.

6.3 A second dimension: how many patents are covering the same invention?

Using the number of patents divided by R&D expenditure can to some extent proxy for the number of patents applied for for each "unit of R&D output", albeit under the assumptions that units of R&D are homogeneous and—more importantly—R&D spending is a good-enough proxy of R&D output. Knowing about the number of patents protecting each "unit of invention" takes us back to the initial question (depicted in figure 1.1): are

[27] A further analysis could include *actual fees paid* in previous years as a sort of "importance-weighted" firm-specific fee variable.

all components of a complex invention included in a common patent, or is protection split over multiple patents?

In a first attempt to address this question, I use data provided by Bergeaud et al. (2017) who assign patents to "semantic classes" based on text analysis of patent abstracts. This pre-categorisation of patents allows to use a very simple measure of patent similarity, namely the ratio of the number of patents in a year to the number of distinct semantic classes in the same year. This measure is smallest when all patents share the same semantic classes, while it takes value 1 when they are all assigned to different semantic classes.[28] As summarised by table 1.25, the number of patents reacts to the UTSA only when this ratio is low (before treatment), while the number of claims decreases (more strongly) when the ratio is high. These results are entirely in line with the idea of "splitting" or combining components into patents that is underlying the model in section 3.

6.4 A third dimension: how many components of the complex product does a single patentee produce alone?

A final question addressed by the main body of analysis is the extent of the patenting decision that is left to a single patentee. In other words, the model in section 3 assumes that a single complex invention is developed by a single firm alone. An important feature of complex product industries (and one of major concern to patent policy, as it gives rise to all sorts of hold-up and related problems) is that commercialised products consist of components contributed by multiple firms. While there is no way of knowing which combinations of patents constitute a commercializable product, a less direct approach looks at the extent of vertical integration of a patentee. Vertical integration is measured by the variance in "upstreamness" of patents. Upstreamness is measured by values provided by Antràs et al. (2012) at the six-digit industry code of the 2002 US benchmark Input-Output Tables. These values are then translated to NAICS industry classes based on tables by the US Bureau of Economic Analysis.

Table 1.26 presents regressions on samples selected by above- and below-median values of the (pre-treatment) standard deviation of "upstreamness" assigned to all of a firm's patents. The results are intuitively appealing, showing a significant reaction of patenting only for those firms that have been assigned a greater variety of upstreamness values. In the interpretation intended here this means that these firms are covering a greater range of components of the complex product they are contributing to, increasing the scope of reducing patenting in favour of secrecy under the assumption that a minimum level of patent protection will be retained if alone for strategic purposes.

Every result based on correspondence tables requires a certain level of caution, even though the result does not change qualitatively with different weighting of IPC-NAICS assignments nor with different ways of weighting IPC classes themselves (e.g., weighting

[28]As the number of semantic classes is limited (to around 300), this measure is not entirely robust to a bias introduced by the size of the patent portfolio: when the number of annual patents is greater than the number of semantic classes, at least two patents will be sharing a semantic class. However, since most patents are assigned to multiple semantic classes, the measure should be able to provide a reasonably good first assessment of the effect of similarity. I am currently performing a more elaborate analysis of the effect of patent similarity that will be added to the paper as soon as possible.

them by relative occurence in each single patent vs. over all patents vs. counting them once only per assignee-year). In regressions not reported here, patenting was found to react only when the level of concentration of patent ownership in a technology field (as measured by the WIPO classification) was high enough, implying both an easier-to-navigate technology space (potentially making strategic patenting less important) and a greater share of the technology field occupied by each single patentee.

6.5 Extending the sample beyond the Compustat match

Matching patent data with firm information from Compustat increases the amount of data available on every assignee and improves precision of regression estimates. On the downside, however, it decreases sample size considerably: in complex product industries by a factor of roughly ten concerning the number of companies included, by a factor of five concerning the raw number of observations, and by a third concerning the number of patents covered by the sample (see table 1.7). This potentially becomes a problem when splitting the sample further, leading to some regressions involving less than 1,000 observations, or almost about as many companies as there are US states. Such a reduced sample size questions the validity of the identification approach which rests on the idea that there is a sufficient number of valid control states available for every state that experiences a legal change.

To address this problem, this section presents analyses performed on the whole NBER patent dataset.[29] In absence of Compustat [*gvkey*] company identifiers, analysis is based on the unique assignee identifier [*pdpass*], which itself is derived from disambiguated applicant names (Bessen, 2009). The sample is restricted to non-government patent assignees with a US postal address at the time of application (country code "US", assignee code "2"), and with patent information aggregated at the level of state-locations per unique assignee code.

As in the Compustat dataset above, no significant effect is found in the overall sample. Conditioning on pre-treatment patents and claims, as before assigned to state-specific quintiles, shows effects very similar to those derived above, if not even more clear-cut (table 1.28). Comparing means and quintile boundaries between groups, the group names seem to fit post-treatment data reasonably well: patent counts are higher for the "many" than for the "few" groups, while claim numbers are higher for "broad" than for "narrow" patent applicants (tables 1.8 and 1.9). A robustness check shows that—not surprisingly—the number of independent and dependent claims as well as their sum is moving simultaneously, though the estimated semi-elasticity is largest for independent claims (1.29).

To allow an assessment of economic significance, the estimated coefficients (semi-elasticities) translate into absolute changes[30] in patents (per assignee-year) of −6.49 (many narrow) and −4.24 (many broad), and in independent claims of −0.66 (few broad) and −0.49 (many broad). The average reduction in the number of independent claims patented that is due to increased secrecy protection is therefore 7,627 (few broad), 30,109 (many narrow), 17,591 (many broad, change in patents) and 18,025 (many broad, change

[29]This dataset was the only one used in the earlier phases of this research project. A number of analyses have yet to repeated with the Compustat-matched dataset analysed above.

[30]These numbers are back-of-the-envelope calculations using the average number of patents and independent claims post-treatment and the respective sample sizes.

in claims). Overall, ignoring the change in independent claims underestimates the effect of the UTSA by 1/3. In the sub-sample of many broad patents, which is close in summary statistics to the sample of patents matched to Compustat firms, the effect due to less patents being applied for captures only half of the total effect.[31]

Robustness: Does disclosure really decrease? This section provides additional evidence that the reduction in claims found so far is equivalent to a reduction in the amount of "inventive content" put into a patent. The first additional measure of patent content is the ratio of USPTO patents to Japanese (JPO) patent applications belonging to the same OECD "triadic patent family", i.e., patents filed at different patent offices that claim the same priority patent, which means that they originate from the same invention. Japanese patents traditionally contain fewer claims (Sakakibara & Branstette, 2001), therefore a US patent belonging to a patent family usually corresponds to two or even more Japanese patents. If the US patents start containing less inventive matter, the number of Japanese patents to cover that matter should decrease. Table 1.34 indicates that this is what happens: since the dependent variable has the number of Japanese patents in the denominator, a decrease in their number implies an increase in the dependent variable.

A second measure of the amount of inventive content is the number of divisional applications that are related to a patent application. Any application containing more than one independent and distinct invention needs to be split into multiple patent applications to be processed. More precisely, one invention needs to be selected for examination, while the remaining inventions can be included in separate applications called "divisional applications". If there are fewer inventions included in the average patent, the likelihood that the patent examiner requires restriction to one invention and accordingly filing of divisional applications decreases. Though only mildly significant, the relevant coefficient estimates (again table 1.34) point in the right direction.

Robustness: Product vs. process innovations. It seems arguably much easier to keep a production process secret that never leaves a firm's premises than a product that can be disassembled and potentially reverse-engineered. It might accordingly be that the observed reduction in patent claims is entirely due to keeping secret of method claims, while product claims are not affected. To find out whether the UTSA has a differential effect on different types of claims, I identified by keyword matching and some limited manual quality control all 75 million claims that are in the USPTO claims research dataset as either a process, a use , or a product.[32] Using numbers of these claims as

[31] There is also high similarity in terms of summary statistics between the "many narrow" sample and the dataset used by Png (2017b), which may explain why in that dataset a sizeable decrease in patents could be found, but not a similar reduction in claims.

[32] Claims that start with "A method" or "A process" are classified as *process claims*; claims that start with "The use" or "The application" are classified as *use claims* (this is by far the least common category). All remaining unclassified claims are considered product claims. This seems to be the most promising way, since product claims are the least structured claims: many begin with "An apparatus", but many much more specific nouns are often used to refer to the protected invention. It needs to be noted that there are instances of overlap between the product and process categories: some claims simultaneously protect, e.g., "An apparatus and a method to product it".

dependent variable instead shows that there is indeed a much stronger effect of secrecy law on process claims, but the effect on product claims is also remarkably strong (table 1.33). Patents under stronger secrecy law do indeed contain a lower number of product inventions.

Heterogeneous effects The dataset allows to answer a number of further questions regarding what factors are influencing the patenting-secrecy decision. What follows is a brief summary of the most important ones. Both large and small firms are decreasing their patenting activities (table 1.30), where assignee size is determined by whether the assignee was in a majority of years eligible for reduced USPTO patent application fees. State-locations belonging to larger companies that include patenting activities in multiple US states show a slightly weaker reaction to the UTSA, but the differences are statistically not significant (table 1.31). Assignees that share ownership of a considerable share of their patents with other assignees show no reaction to the UTSA in the number of patents, while there is no effect of shared patent ownership on the reduction in claims. This points towards the raw number of patents being a variable of importance in contractual relationships, beyond providing protection against imitation.

Extending the result of weaker impact of state law changes on companies that are active beyond state boundaries, table 1.32 looks at the effect of foreign inventors on the secrecy decision. Enforcement of even national law beyond the borders of a country is challenging, which is why changes in state law are expected to be of minor importance to firms that have to (potentially) protect internal R&D knowledge outside the US. In line with this reasoning, estimates are negative but not significant for firms that have at least one inventor with a non-US residential address in any of their assigned patents.

Finally, while R&D expenditures control for the magnitude R&D effort, they do not capture whether improved secrecy protection affects the (technological) nature of a firm's R&D projects. One way to account for changes in patented technologies is to identify the most important (modal) IPC class for each assignee separately for the time before and after treatment, and then check whether there are significant differences in the reaction to the UTSA when firms changed their modal IPC class after treatment. The right-hand half of table 1.32 shows that indeed firms that have a different modal IPC class after treatment do not significantly reduce their number of patents, even though the estimate for the UTSA is still negative and sizeable. This is also true for the estimate of the claims regression, but due to the small magnitude the difference is not statistically significant.

7 Implications of the Endogeneity of Patent Claims for the Measurement of Innovation

This section builds on the fact established above that patentees do have discretion in determining the extent of a complex innovation included in a patent. Patenting more, accordingly, does not always imply that a greater number of patents needs to be observed. Rather, within some legal and technological boundaries, an additional component can be included in a patent that would otherwise be applied for as well. Each patent can have a number of independent claims, each covering a different component of the innovation.

Defining the relationship $k = p{\cdot}c$, where k is the total number of patented components, p the patentee's number of patents [33], and c the average number of independent claims in those patents,[34] then the total amount of patented innovation is measured by k, not p. p will most probably be positively correlated with k, but to the extent that c is variable, using p instead of k may yield misleading results.

Consider the simple regressions equation as in 1.8, leaving out the treatment variable for simplicity:

$$p = \lambda_{is} + \mu_t + \beta X_{ist} + \varepsilon_{ist}. \tag{1.10}$$

Regressing p here means implicitly regressions k/c. If c is fixed, for whatever reason, all variation in k will necessarily show in p. If c is variable, p may—in an extreme case where all additional patenting only affects c—remain unchanged.

In the following, I replicate two prior studies. The first, Kogan et al. (2017), calculates patent value by stock market reactions to patent publications. I use the study to validate independent claim counts as a measure of patent content and the potentially different nature of independent claims in complex and discrete product industries. The second study, Hashmi (2013), is itself a replication attempt at Aghion et al. (2005). Using British patent data, the latter study finds that the nature of competition and innovation, measured by (citations-weighted) patent counts, takes the shape of an inverted U. Hashmi, using US patent data for the same time frame and selection of industries, finds a mildly but consistently negative relationship. I use this study to showcase how—both regarding the result and regarding the mechanism—using k instead of p affects findings.

7.1 Kogan, Papanikolaou, Seru, & Stoffman (2017).

The authors's main contribution is a dataset containing a value measure for 1,801,879 US patents, spanning a time period of almost 100 years. To measure patent value, previous studies have commonly resorted to the number of citations that a patent receives. While patent citations can occur for a variety of reasons (Jaffe & Rassenfosse, 2017), the consensus is that a highly-cited patent has proven more influential than the average patent, and this influence should be somehow positively affecting value. This can be, e.g., due to licensing opportunities to follow-on innovators, or by forward citations merely indicating a technology area with ample commercialisation opportunities, attracting a lot of activity.

Kogan et al. compare their measure to that of citations received, finding a strongly positive correlation (see figure II in their article), possibly with the exception of patents with very few citations. I use their dataset and add my independent claims count as well as an SIC industry variable from Compustat, which as before allows to distinguish complex- and discrete-technology main areas of firm activity. Panels (a) and (b) of figure 1.10 show that in both cases there is some sort of positive relationship between the log number of independent claims per patent and the associated log patent value, but the

[33] Again, in the data used I do not observe patent applications, but patent grants, but the logic applies to both settings.

[34] Patents differ regarding the number of independent claims they contain, but the present considerations are supposed to guide regression analysis, which considers means only.

relationship is much stronger in complex product industries. For those, and at least in the sample covered the dataset by Kogan et al., the number of independent claims work as well as citation numbers when proxying patent value.

Panels (c) and (d) repeat the exercise, but all measures have been first aggregated to the assignee (patent-portfolio) level, using the CRSP permno variable matched to patents by Kogan et al. At the portfolio level, citations and independent claim counts fare almost equally well regardless of the industry to which the assignee belongs.

Additionally, for a set of 1,102,098 patents[35] I use average-linkage cluster analysis with a Jaccard similarity measure on the cleaned and stemmed claim text to identify groups of very similar claims. Adding several similar claims may increase the breadth of protection of a patent, but arguably does not add a lot in terms of disclosure. Therefore, grouping similar claims may give a better claim-based measure of the innovative content of a patent. I create several measures with Jaccard similarity threshold values between 0.2 and 0.9. For illustration purposes, I added the number of claim clusters computed with a threshold of 0.5 to panels (c) and (d); their performance is very comparable to that of citations and claims counts.

7.2 Hashmi (2013).

The author compiles a dataset of patent and citation numbers at the industry-year level, using US patent data from 1976 to 2001. While the final dataset is available for download, in order to add claim and similarity measures, the dataset needs to be re-created starting at the patent level. I do so following the instructions provided by the author in his Appendix A, including the measure of competition based on firm financial data from Compustat North America and the instrumental variable based on dollar exchange rates with major trading destinations per industry.

In this section I focus on the author's first result, presented in his figure 1, panel (b): a negative relationship between innovation and competition. This needs to be contrasted with the finding by Aghion et al., replicated by the author in his panel (a), who find an inverted-U shaped relationship between innovation and competition, where citations-weighted patent output is maximised in industries with a medium degree of competition intensity.

Figure 1.11a shows the result of my replication attempt. The estimated relationship is negative as well, albeit with a much stronger estimated effect of competition on innovation. The levels of competition are on average higher than in Hashmi's data, which is likely due to differences in the Compustat variables used to calculate an industry's Lerner index. The distribution of competition levels is in fact closer to that of Aghion et al. Importantly, though, there is no sign of an inverted U-shaped relationship.

This changes when using counts of independent claims instead of citations, as shown in panel (b). Panels (c)-(f) use the number of claim clusters described in the preceding subsection. With decreasing similarity threshold, i.e., when grouping together increasingly less-similar claims, the maximum of the curve shifts further to the right. In all cases, however, the relationship is distinctly inverse-U shaped.

[35] 334,781 of those patents contain only a single independent claim, which by definition implies a single claim cluster.

In order to delve deeper into possible explanations of this difference, I calculate the pairwise similarity of independent claim text both between claims within a patent as well as across patents of the same patentee. Again, I use Jaccard similarity as the main measure, but similar results are obtained by using Sorensen-Dice and Cosine similarity. I then aggregate the claim-pair-wise similarity measures at the patentee-year level, computing averages and medians, alternatively using a 1-year or a 3-year window.[36]

This assignee-level similarity measure is similar in nature to the across-patent similarity measure computed by Bergeaud et al. (2017) used before. Following the same logic as the cluster measure, when claims within a patent are very similar, there is little to be gained in terms of lower disclosure by keeping one of those claims secret. Vice versa, if patents of the same patent applicant are similar, keeping one of them secret reduces disclosure less than in a case where patents are highly dis-similar. The underlying logic therefore again leads to the conjecture that a firm reacts to changes in the propensity to patent by changing p when claims are highly similar, while it rather changes c when claims are dis-similar and/or patents are similar.

Figure 1.12 applies this logic to the present dataset. Panel (a) is the same as in figure 1.11, while panel (c) shows c, the average number of claims per patent, for the same sample. The right column with panels (b) and (d), respectively, contains estimates of the same regression, but in each case the sample is restricted by leaving out those industries that feature the highest average degree of similarity in terms of the dependent variable. In (b), this implies leaving out the upper quintile of across-patent claim similarity, while in panel (d), industries with above-median similarity were excluded. In each case, an easily-distinguishable inverse-U shape results.

One question that remains is how such a difference between US and UK patent data may come about. To a large extent, this may be due to institutional differences between the respective patent offices. Both academic and practitioner literature reports that on average, US patents include more claims at application and at grant, the fees per claim are lower at the USPTO compared to the EPO, and on average, the USPTO is more lenient towards the concept of "unity of invention" during its patent examination (cf. van Pottelsberghe & François, 2009; de Rassenfosse & van Pottelsberghe, 2012; van Zeebroeck et al., 2009). This may indicate that there is less scope for the behaviour described in the present essay when applying to a patent in Europe. A more definitive assessment can only be made using European patent data.

8 Conclusion

This chapter makes two main contributions. The first is in estimating the total effect of UTSA enactment on patenting by considering independent claims in addition to patent counts, the second is in showcasing the implication of such endogeneity of claims.

A number of effects have been identified that are in line with the idea of independent claims as the patenting decision variable. The reaction to the UTSA shown by patents (claims) depends negatively on the pre-treatment value of claims (patents). Greater similarity between patents yields a reaction in patents, while lower similarity yields a reaction

[36]The use of a three-year window is motivated by the idea that complex innovations can be covered by more than a single patent, which may be applied for in different (but consecutive) calendar years.

in claims. This is in line with greater similarity between patents implying that similar patents are covering the same invention. Greater variety in the level of "upstreamness" of a firm's patents yields a stronger reaction to the UTSA, which would be expected of a firm covering a greater share of the K components of a complex product.

Further evidence reinforces the ideas underlying the theory section. Patenting reacts negatively to patent office fees, including the number of independent claims, agreeing with the idea of discretion at the claim level on the side of the patentee; both product and process claims are reduced; and the systematic reduction of claims versus patents is found only in complex product industries where patenting needs to be decided for multiple components per invention. This is in line with the reported strategy of licensing single independent claims, not always whole patents (Dahlin, 2015; Schmidt, 2015), treating independent claims as the units of innovation.

"Patent breadth" being a choice variable similar to—or instead of—patent count would be of importance to both patent policy and empirical patent research. First, this means that counting patents is not enough to get a complete picture of firms' inventive output. This fact is not new to the literature analysing patent data, but it is given a new dimension: not only is there R&D output left unpatented, but in certain cases additional claims represent additional R&D results. An example of an application would be the study of the effect of patent invalidation on follow-on innovation by Galasso & Schankerman (2015). They look at the effect of invalidation of *at least one claim* in a patent on the future accumulation of forward citations by the invalidated patent. For patents in complex product industries, this effect should be increasing in the number of invalidated independent claims.[37]

Simply controlling for patent attributes may as well not be a valid strategy since such attributes can be as endogenous as the patent count itself. When analysing whether the propensity to patent generally differs with firm size,[38] comparing patent counts might under- or overestimate the true extent of patenting in complex product industries. When, e.g., small firms apply for fewer patents, but these patents are on average broader, then controlling for the number of claims in the patent regression would no yield a meaningful correction of the estimated coefficient of firm size.

Certainly the validity to use the number of claims as a control variable depends on the aim of the particular study at hand. For reasons of strategic patenting, the sheer number of patents may indeed be the more important criterion. If interest lies instead on the production of knowledge, the informational content of patents is of major importance. In this respect, Tong & Frame (1994) seem to have been on the right track. In the context of complex product industries, the number of independent claims seems to behave like a good measure of the number of inventive pieces included in a patent.

This interpretation is, however, restricted to complex product industries.[39] Just as

[37]The authors only look at the effect of invalidation of claim 1, but not the number of claims. They however find a positive relationship between a judge's "propensity to invalidate" and the number of claims actually being invalidated.

[38]A question of considerable policy interest, see i.a. the summary in Kingston (2004).

[39]This remark requires acknowledging that distinguishing between complex and discrete industries is empirically challenging and most often incomplete. This is an even more common problem for patents, which are hard to unambiguously allocate to an industry in the first place. Most industries will locate somewhere between the endpoints of a continuous scale of product complexity.

innovation and patenting strategies in these industries differ considerably from those common in discrete products, so should the peculiarities of complementary innovation that characterises these industries be taken into account by patent policy. On one hand, an above-average number of claims per patent does not necessarily mean that this number is “excessive”; instead, in complex product industries the ‘throughput’ of patent examiners may be more appropriately measured based on claims rather than patents (cf. van Pottelsberghe & François, 2009). On the other hand, the large difference between the per-patent and the per-claim components of patent fees may not be equally appropriate in complex as in discrete product industries.

References

Patents cited

U.S. Patent No. 216,231: Improvement in velocipedes. Inventor: Shire, J. Application filed by J. Shire, 1879.

U.S. Patent No. 5,687,577: Apparatus and method for spray-cooling an electronic module. Inventors: Ballard, G. W., Turocy, J. W., & Beise, T. R. Application filed by Motorola Solutions Inc., 1996.

U.S. Patent No. 5,718,117: Apparatus and method for spray-cooling an electronic module. Inventors: Inventors: McDunn, K. J., Limper-Brenner, L., & Press, M. D. Application filed by Motorola Solutions Inc., 1996.

U.S. Patent No. 5,761,035: Circuit board apparatus and method for spray-cooling a circuit board. Inventor: Beise, T. R. Application filed by Motorola Solutions Inc., 1996.

U.S. Patent No. 5,768,103: Circuit board apparatus and apparatus and method for spray-cooling an electronic component. Inventors: Kobrinetz, A., & Bullock, M. K. Application filed by Motorola Solutions Inc. in 1996.

U.S. Patent No. 5,831,824: Apparatus for spray-cooling multiple electronic modules. Inventors: McDunn, K. J., Limper-Brenner, L., & Press, M. D. Application filed by Motorola Solutions Inc., 1996.

U.S. Patent No. 5,854,092: Method for spray-cooling a tunable semiconductor device. Inventors: Root, L. F., & McDunn, K. J. Application filed by Motorola Solutions Inc., 1997.

U.S. Patent No. 7,280,838: Paging transceivers and methods for selectively retrieving messages. Inventor: Helferich, R. J. Application filed by The Thompson Investment Group, LLC, 2005; assigned to Helferich Patent Licensing, LLC, 2008.

U.S. Patent No. 8,836,922: Devices and methods for a rotating LIDAR platform with a shared transmit/receive path. Inventors: Pennecot, G., Droz, P.-Y., Ulrich, D. E., Gruver, D., Morriss, Z., & Levandovski, A. Application filed by Google LLC, 2013.

U.S. Patent No. 9,123,979: Printed waveguide transmission line having layers with through-holes having alternating greater/lesser widths in adjacent layers. Inventor: Izadian, J. S. Application filed by Google LLC, 2013.

U.S. Patent No. 9,285,464: Devices and methods for a rotating LIDAR platform with a shared transmit/receive path. Inventors: Pennecot, G., Droz, P.-Y., Ulrich, D. E., Gruver, D., Morriss, Z., & Levandovski, A. Application filed by Google LLC, 2014.

U.S. Patent No. 9,935,514: Printed circuit board layout. Inventors: Lenius, S. W., Karplus, P., & Zbrozek, A. D. Application filed by Waymo LLC, 2014.

Literature

Aghion, P., Bloom, N., Blundell, R., Griffith, R., & Howitt, P. (2005). Competition and Innovation: An Inverted-U Relationship. *The Quarterly Journal of Economics, 120*(2), 701–728.

Arundel, A., & Kabla, I. (1998). What percentage of innovations are patented? empirical estimates for European firms. *Research Policy, 27*(2), 127–141.

Arundel, A. (2001). The relative effectiveness of patents and secrecy for appropriation. *Research Policy, 30*(4), 611–624.

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

Anton, J. J., Greene, H., & Yao, D. A. (2006). Policy Implications of Weak Patent Rights. In A. B. Jaffe, J. Lerner & S. Stern (Eds.), *Innovation Policy and the Economy, Volume 6*. MIT Press.

Antràs, P., Chor, D., Fally, T., & Hillberry, R. (2012). Measuring the Upstreamness of Production and Trade Flows. *American Economic Review, 102*(3), 412–416.

Archontopoulos, E., Guellec, D., Stevnsborg, N., van Pottelsberghe de la Potterie, B., & van Zeebroeck, N. (2007). When small is beautiful: Measuring the evolution and consequences of the voluminosity of patent applications at the EPO. *Information Economics and Policy, 19*(2), 103–132.

Bartz, D. (2015, February). Patent licenser Helferich wins court round against New York Times, others. *Reuters*. Retrieved from https://www.reuters.com/article/helferich-new-york-times-patent/patent-licenser-

Belleflamme, P., & Bloch, F. (2013). Dynamic Protection of Innovations through Patents and Trade Secrets. *CESifo Working Paper No. 4486*.

Bergeaud, A., Potiron, Y., & Raimbault, J. (2017). Classifying patents based on their semantic content. *PLOS ONE, 12*(4), e0176310.

Bessen, J. (2009). NBER PDP Project User Documentation: Matching Patent Data to Compustat Firms. Available at http://users.nber.org/~jbessen/matchdoc.pdf

Breen, B., & Stockwell, D. (2015, February). Helferich Patent Licensing v. New York Times et al.: Federal Circuit Reverses District Court Summary Judgment Ruling Regarding Patent Exhaustion. *JDSUPRA*. Retrieved from https://www.jdsupra.com/legalnews/helferich-patent-licensing-v-new-york-t-22128.

Ciardullo, J.-P. (2015, March). United States: Federal Circuit Limits Patent Exhaustion Doctrine For Complementary Technology. *Mondaq*. Retrieved from http://www.mondaq.com/unitedstates/x/381770/Patent/Federal+Circuit+Limits+Patent+Exhaustion+Doctrine+for+Complementary+Technology.

Cohen, W. M., Nelson, R. R., & Walsh, J. P. (2000). Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not). *NBER Working Paper No. 7552*.

Crass, D., Garcia-Valero, F., Pitton, F., & Rammer, C. (2016). Protecting Innovation Through Patents and Trade Secrets: Determinants and Performance Impacts for Firms with a Single Innovation. *ZEW Discussion Paper No. 16-061.*

Dahlin, L. (2015). When Is a Patent Exhausted: Licensing Patents on a Claim-by-Claim Basis. *Chicago-Kent Law Review, 90*, 757–781.

Dass, N., Nanda, V. K., & Xiao, S. C. (2015). Intellectual Property Protection and Financial Markets: Patenting vs. Secrecy. *Mimeo.* doi:10.2139/ssrn.2517838

de Rassenfosse, G., & van Pottelsberghe de la Potterie, B. (2007). Per un pugno di dollari: a first look at the price elasticity of patents. *Oxford Review of Economic Policy, 23*(4), 588–604.

de Rassenfosse, G., & van Pottelsberghe de la Potterie, B. (2012). On the Price Elasticity of Demand for Patents. *Oxford Bulletin of Economics and Statistics, 74*(1), 58–77.

de Rassenfosse, G. & van Pottelsberghe de la Potterie, B. (2013). The Role of Fees in Patent Systems: Theory and Evidence. *Journal of Economic Surveys, 27*(4), 696–716.

Denicola, R. (2011). The Restatements, the Uniform Act and the status of American trade secret law. In R. C. Dreyfuss & K. J. Strandburg (Eds.), *The Law and Theory of Trade Secrecy: A Handbook of Contemporary Research* (pp. 18–45). Cheltenham, UK: Edward Elgar.

Denicolò, V. (1996). Patent races and optimal patent breadth and length. *The Journal of Industrial Economics, 44*(3), 249–265.

Denicolò, V. (2000). Two-stage patent races and patent policy. *The Rand Journal of Economics, 31*(3), 488–501.

Frakes, M. D., & Wasserman, M. F. (2017). Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Microlevel Application Data. *The Review of Economics and Statistics, 99*(3), 550–563.

Galasso, A., & Schankerman, M. (2015). Patents and Cumulative Innovation: Causal Evidence from the Courts. *The Quarterly Journal of Economics, 130*(1), 317–369.

Gallini, N. T. (1992). Patent policy and costly innovation. *The Rand Journal of Economics, 23*(1), 52–63.

Gallini, N. T. (2002). The Economics of Patents: Lessons from Recent US Patent Reform. *Journal of Economic Perspectives, 16*(2), 131–154.

Gamba, S. (2016). The effect of Intellectual Property Rights on domestic innovation in the pharmaceutical sector. *FBK-IRVAPP Working Paper No. 2016-05.* Retrieved from http://irvapp.fbk.eu/publications/effect-intellectual-property-rights-domestic-innovation-pharmaceutical-sector

Gilbert, R. J., & Shapiro, C., (1990). Optimal patent length and breadth. *The RAND Journal of Economics, 21*, 106–112.

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

Graham, S. J. H., & Somaya, D. (2006). Vermeers and Rembrandts in the Same Attic: Complementarity between Copyright and Trademark Leveraging Strategies in Software. *Georgia Institute of Technology TI:GER Working Paper.* doi:10.2139/ssrn.887484

Graham, S. J. H., Marco, A. C., & Miller, R. (2015). The USPTO Patent Examination Research Dataset: A Window on the Process of Patent Examination. *USPTO Economic Working Paper No. 2015-4.*

Hall, B. H., Jaffe, A. B., & Trajtenberg, M. (2001). The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. *NBER Working Paper 8498.*

Hall, B. H., Helmers, C., Rogers, M., & Sena, V. (2014). The Choice between Formal and Informal Intellectual Property: A Review. *Journal of Economic Literature, 52*(2), 375–423.

Hashmi, A. R. (2013). Competition and Innovation: The Inverted-U Relationship Revisited. *The Review of Economics and Statistics, 95*(5), 1653–1668.

Hong, S., Oxley, L., & McCann, P. (2012). A Survey of the Innovation Surveys. *Journal of Economic Surveys, 26*(3), 420–444.

Jaffe, A. B., & de Rassenfosse, G. (2017). Patent citation data in social science research: Overview and best practices. *Journal of the Association for Information Science and Technology, 68*(6), 1360–1374.

Johnson, D. K. N., & Popp, D. (2003). Forced out of the Closet: The Impact of the American Inventors Protection Act on the Timing of Patent Disclosure. *The RAND Journal of Economics, 34*(1), 96–112.

Kingston, W. (2004). Making patents useful to small firms. *Intellectual Property Quarterly, (4)*, 369–378.

Klemperer, P. (1990). How broad should the scope of patent protection be? *The Rand Journal of Economics, 21*, 113–130.

Kuhn, J. M., & Thompson, N. (2017). How to Measure and Draw Causal Inferences with Patent Scope. International Journal of the Economics of Business, Forthcoming.

Kwon, I. (2012a). Patent Races With Secrecy. *The Journal of Industrial Economics, LX*(3), 499-516.

Kwon, I. (2012b). Patent Thicket, Secrecy, and Licensing. *The Korean Economic Review, 28*(1), 27-49.

Lai, R., D'Amour, A., Yu, A., Sun, Y., & Fleming, L. (2011). Disambiguation and Co-authorship Networks of the U.S. Patent Inventor Database (1975 - 2010). *Harvard Dataverse, V5.* hdl:1902.1/15705

Lanjouw, J. O., & Schankerman, M. (2001). Characteristics of Patent Litigation: A Window on Competition. *The RAND Journal of Economics, 32*(1), 129–151.

Lee, J. F. (2017). Measuring Innovation with Patents when Patenting is Strategic. *Job Market Paper, Queen's University, Ontario.* Retrieved from https://sites.google.com/site/leejonf

Lemley, M. A. (2001). Rational ignorance at the patent office. *Northwestern University Law Review, 95*(4), 1495–1532.

Lemley, M. A., & Sampat, B. (2012). Examiner Characteristics and Patent Office Outcomes. *The Review of Economics and Statistics, 94*(3), 817–827.

Lemley, M. A., & Shapiro, C. (2005). Probabilistic patents. Journal of Economic Perspectives, 75–98.

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

Levin, R. C., Klevorick, A. K., Nelson, R. R., Winter, S. G., Gilbert, R., & Griliches, Z. (1987). Appropriating the Returns from Industrial Research and Development. *Brookings Papers on Economic Activity, 1987*(3), Special Issue on Microeconomics (1987), 783–831.

Levine, D. (2017, October 12). Exclusive: Alphabet's Waymo demanded $1 billion in settlement talks with Uber. Reuters. Retrieved from https://www.reuters.com/article/us-alphabet-uber-lawsuit/exclusive-alphabets-waymo-demanded-1-billion-in-settlement-talks-with-uber-sources-idUSKBN1CH0QC

Lybbert, T. J., & Zolas, N. J. (2014). Getting patents and economic data to speak to each other: An 'Algorithmic Links with Probabilities' approach for joint analyses of patenting and economic activity. *Research Policy, 43*(3), 530–542.

Marco, A. C., Sarnoff, J. D., & deGrazia, C. (2016). Patent Claims and Patent Scope. *USPTO Economic Working Paper No. 2016-04*.

Ohnsman, A. (2017, March 1). Waymo's Uber Lawsuit May Be Start Of A Google Patent War. Forbes. Retrieved from https://www.forbes.com/sites/alanohnsman/2017/03/01/waymos-trade-secret-suit-against-ub

Ottoz, E., & Cugno, F. (2008). Patent–Secret Mix in Complex Product Firms. *American Law and Economics Review, 10*(1), 142-158.

Ottoz, E., & Cugno, F. (2011). Choosing the scope of trade secret law when secrets complement patents. *International Review of Law and Economics, 31*, 219-227.

Park, W. G. (2008). International patent protection: 1960–2005. *Research Policy, 37*(4), 761–766.

Png, I. P. L. (2015). Replication Data for: Law and Innovation: Evidence from State Trade Secrets Laws. *Harvard Dataverse, V1.* doi:10.7910/DVN/BFP2IC

Png, I. P. L. (2017a). Law and Innovation - Evidence from the State Trade Secrets Laws. *Review of Economics and Statistics, 99*(1), 167–179.

Png, I. P. L. (2017b). Secrecy and Patents: Theory and Evidence from the Uniform Trade Secrets Act. *Strategy Science, 2*(3), 176–193.

Risch, M. (2015). Empirical Methods in Trade Secret Research. In P. S. Menell & D. L. Schwartz (Eds.), *Research Handbook on the Economics of Intellectual Property Law*, Vol. II. Cheltenham, UK: Edward Elgar (forthcoming 2017). Retrieved from http://papers.ssrn.com/abstract=2658685

Sakakibara, M., & Branstetter, L. (2001). Do Stronger Patents Induce More Innovation? Evidence from the 1988 Japanese Patent Law Reforms. *The RAND Journal of Economics, 32*(1), 77–100.

Schmidt, J. W. (2015). Authorized Personnel Only: The Patent Exhaustion Doctrine after Helferich. *Notre Dame Law Review, 91*, 2115–2154.

Scott, J. R. M., & de Jonge, S. (2008). When is a search not a search? Part 2 -? Non-unity, the EPO approach. *World Patent Information, 30*(3), 199–205.

The United States Patent and Trade Mark Office. (2017). *Manual of Patent Examination Procedure (MPEP)* (Ninth Edition, Revision 08.2017). Alexandria, VA. Available at https://www.uspto.gov/web/offices/pac/mpep/index.html.

Tong, X., & Frame, J. D. (1994). Measuring national technological performance with patent claims data. *Research Policy, 23*(2), 133–141.

van Pottelsberghe de la Potterie, B., & François, D. (2009). The Cost Factor in Patent Systems. *Journal of Industry, Competition and Trade, 9*(4), 329–355.

van Zeebroeck, N., Stevnsborg, N., van Pottelsberghe de la Potterie, B., Guellec, D., & Archontopoulos, E. (2008). Patent inflation in Europe. *World Patent Information, 30*(1), 43–52.

van Zeebroeck, N., van Pottelsberghe de la Potterie, B., & Guellec, D. (2009). Claiming more: the Increased Voluminosity of Patent Applications and its Determinants. *Research Policy, 38*(6), 1006–1020.

Yiannaka, A., & Fulton, M. (2006). Strategic patent breadth and entry deterrence with drastic product innovations. *International Journal of Industrial Organization, 24*(1), 177–202.

Yiannaka, A., & Fulton, M. (2011). Getting Away with Robbery? Patenting Behavior with the Threat of Infringement. *Journal of Economics & Management Strategy, 20*(2), 625–648.

A Figures and Tables

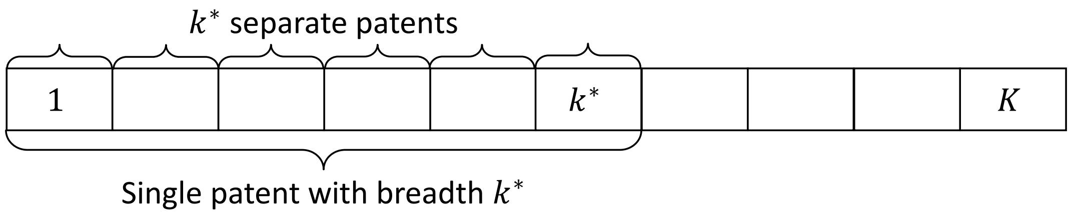

Figure 1.1: Patent protection of a complex product

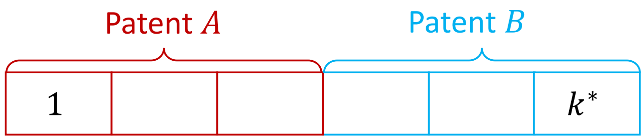

Figure 1.2: Allocation of components to patents *sequentially*

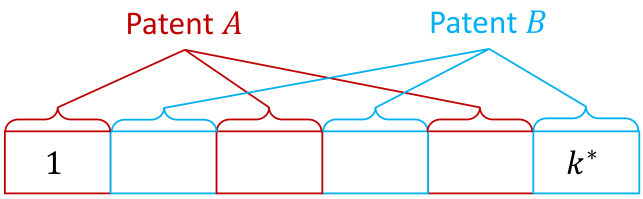

Figure 1.3: Allocation of components to patents *topically*

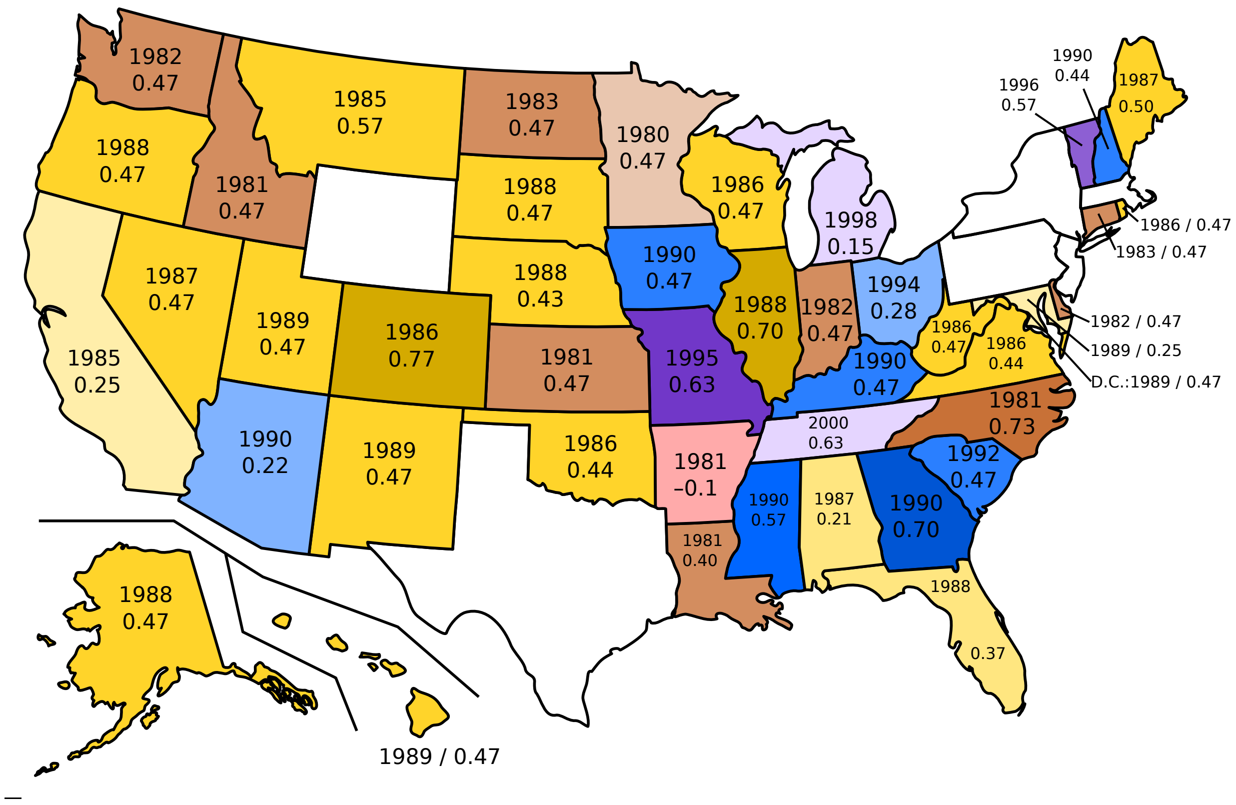

Figure 1.4: Map of enactment years and changes in the states' trade secret protection indexes (Png, 2015) due to enactment of the UTSA.

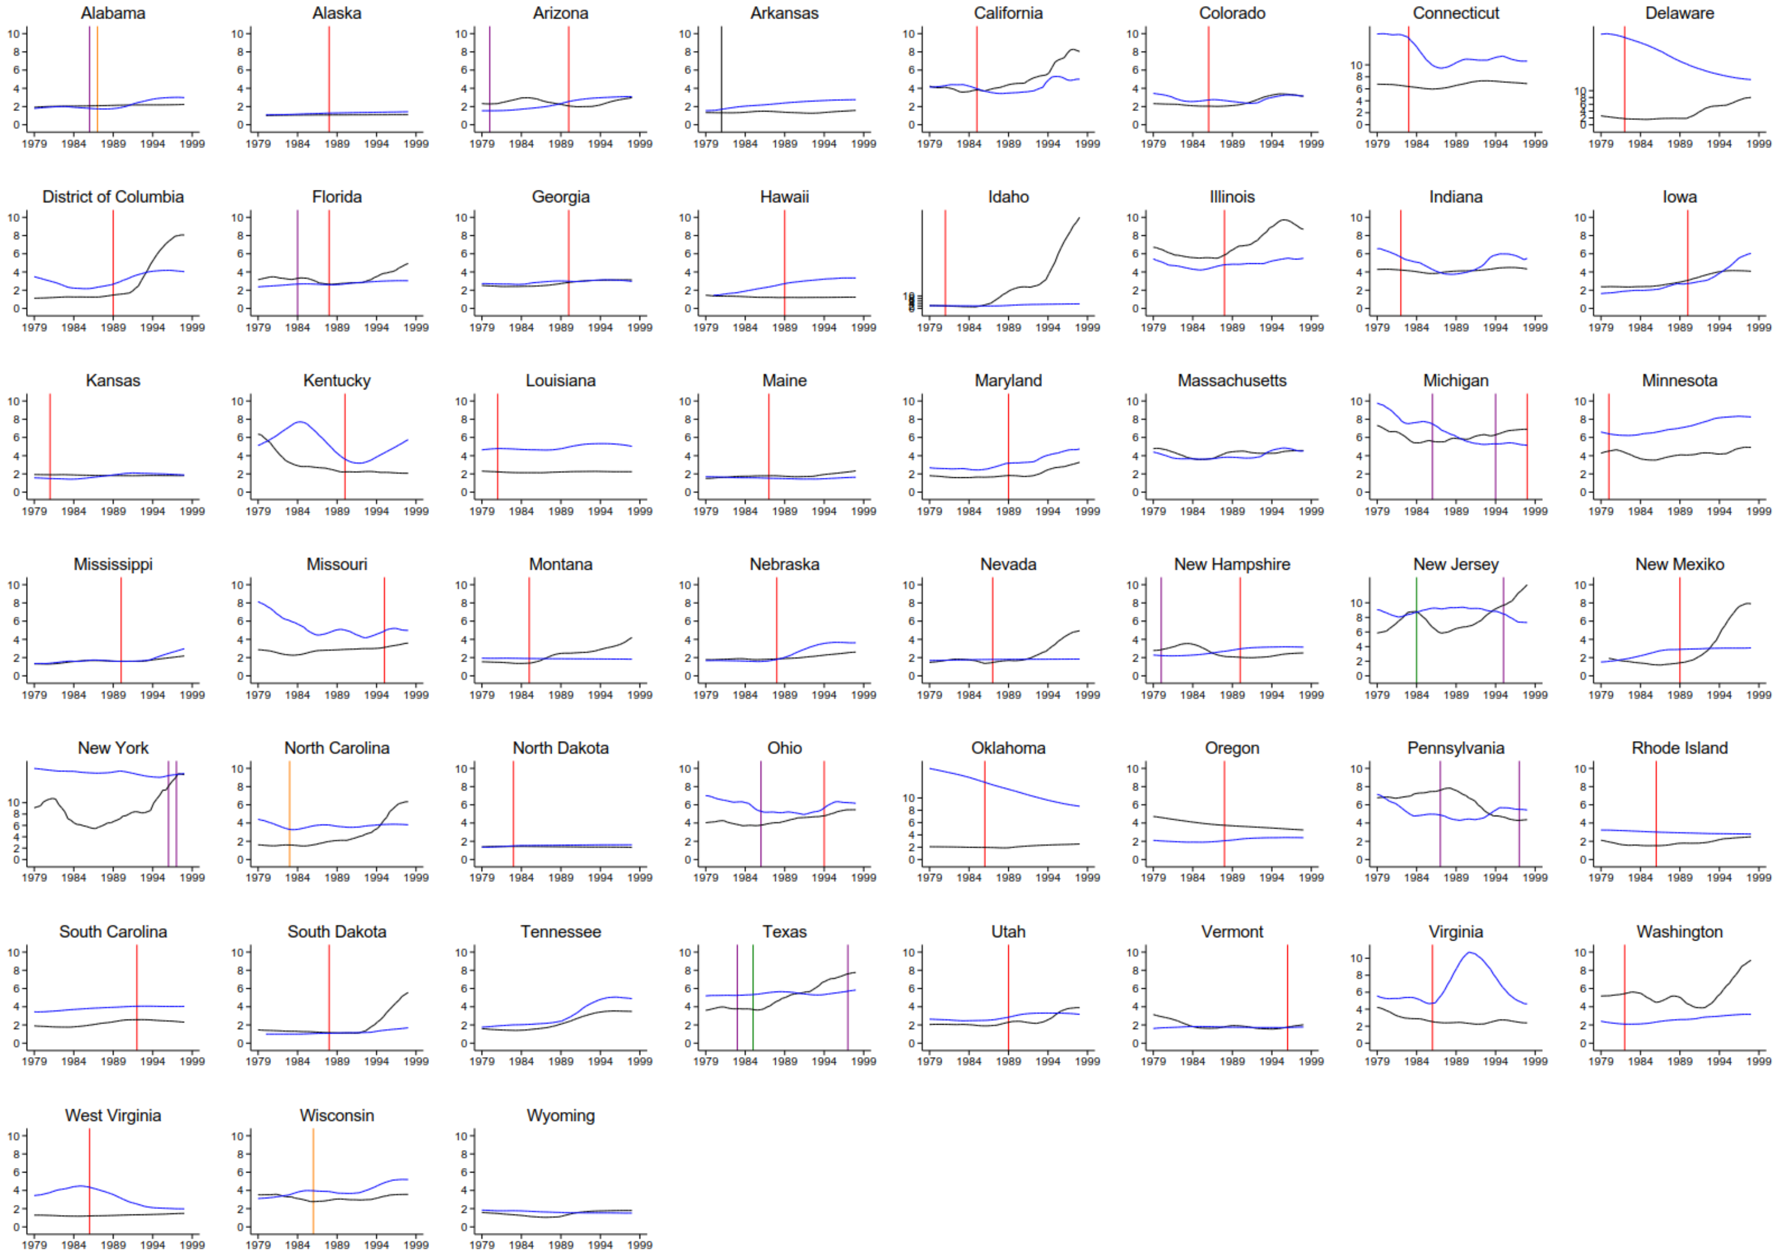

Curves are local polynomial smooths. A red vertical line depicts the year of UTSA implementation, an orange line enactment of a TS statute different from the UTSA, and a purple line important changes in the common law. If a legal change decreases the total level of TS protection, this is depicted by a black (UTSA) or green line (common law). If possible, the vertical axis is scaled from 0 to 10 to facilitate comparison between states.

Figure 1.5: Patent applications in complex (black) and discrete (blue) product industries over time

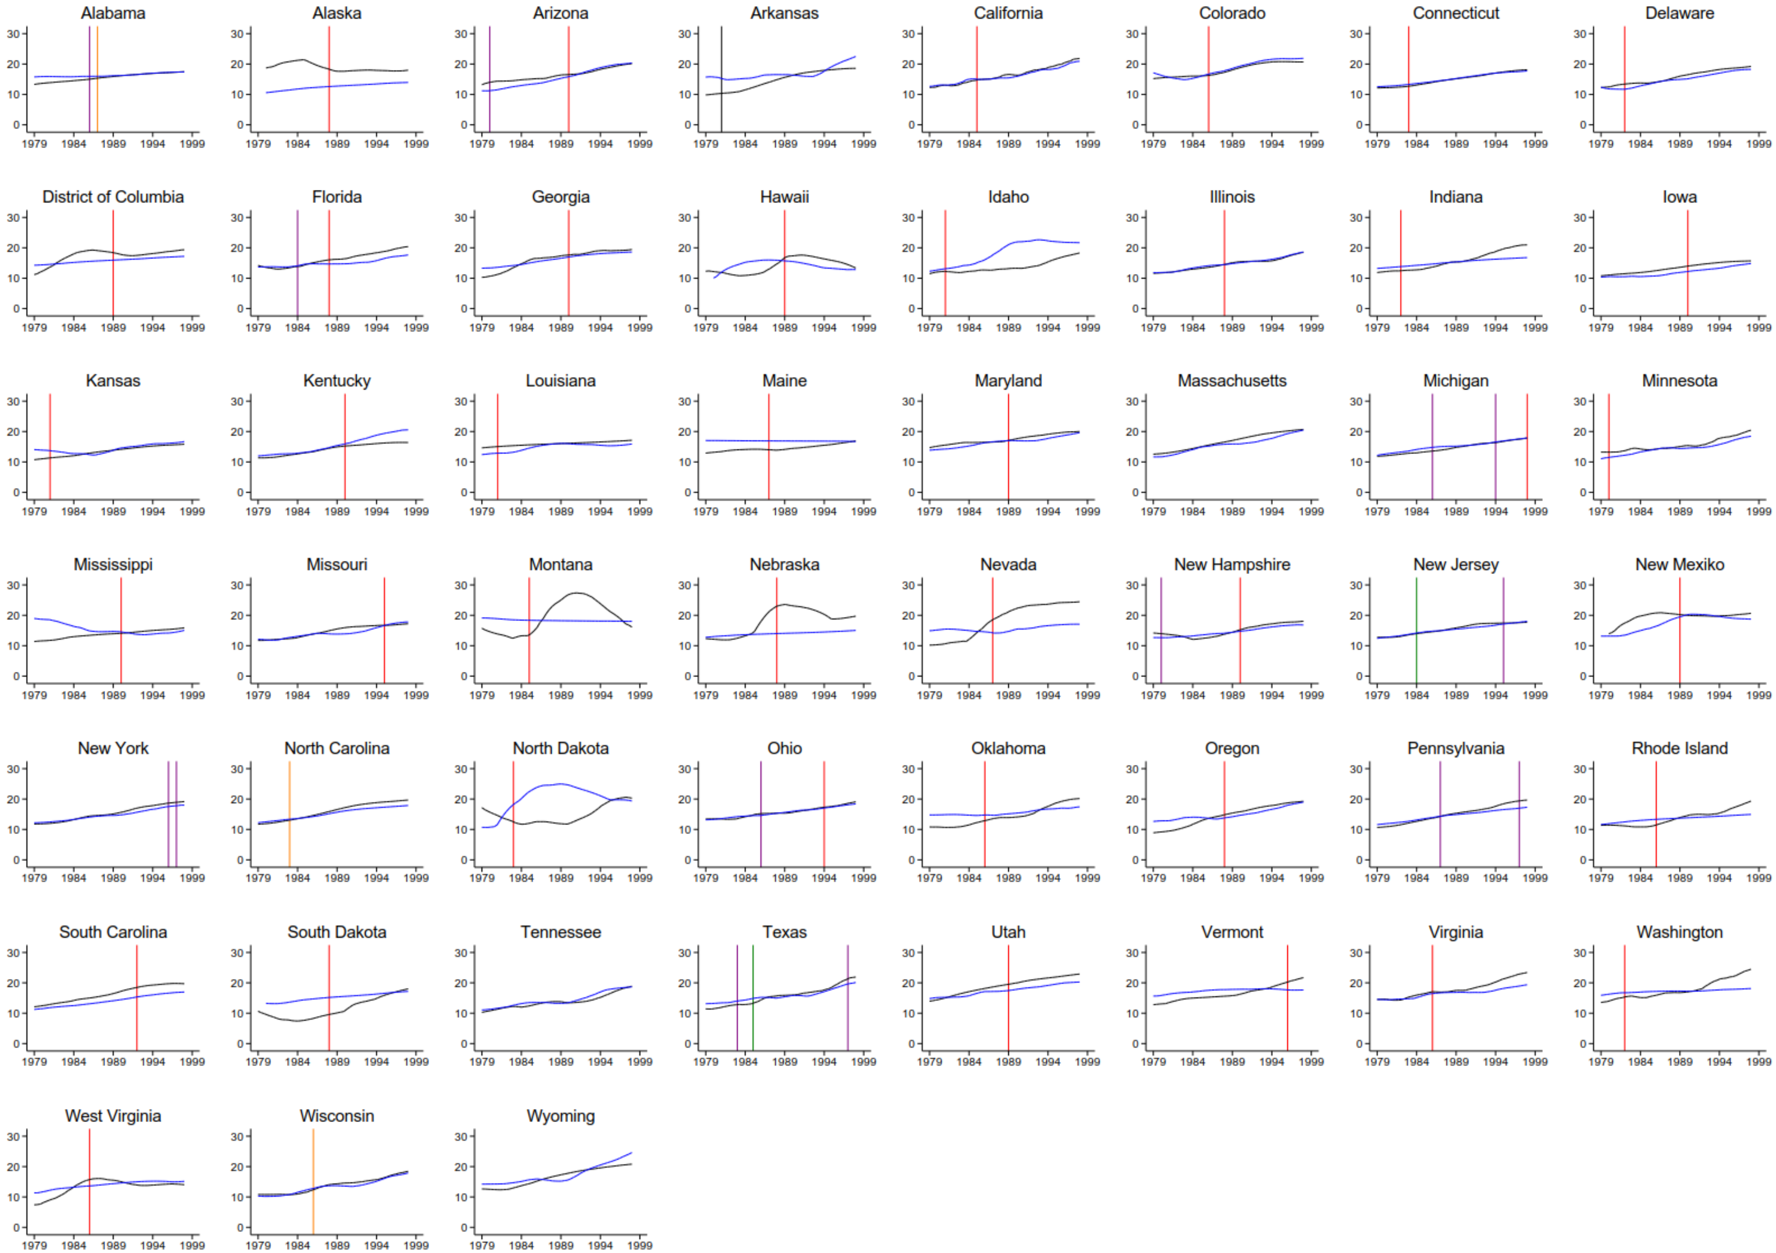

Curves are local polynomial smooths. A red vertical line depicts the year of UTSA implementation, an orange line enactment of a TS statute different from the UTSA, and a purple line important changes in the common law. If a legal change decreases the total level of TS protection, this is depicted by a black (UTSA) or green line (common law).

Figure 1.6: Claims per patent in complex (black) and discrete (blue) product industries over time

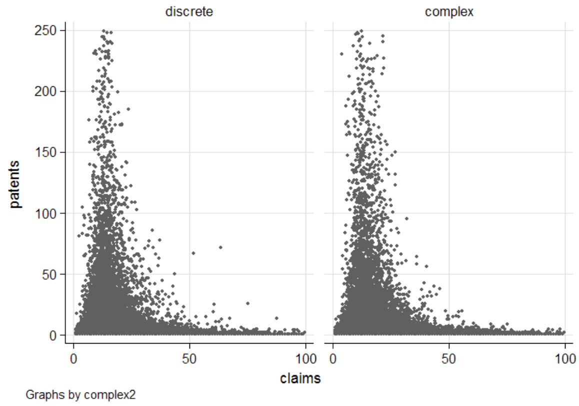

(a) Whole sample period

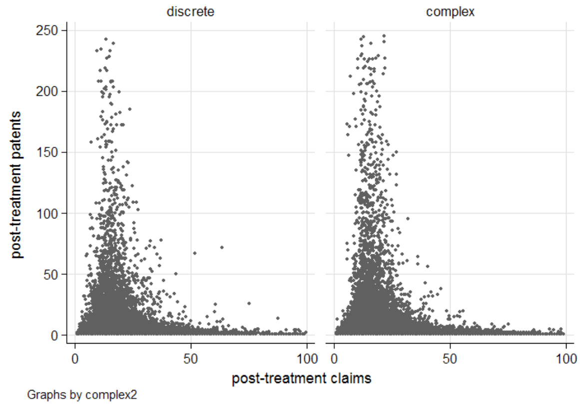

(b) Post-treatment only

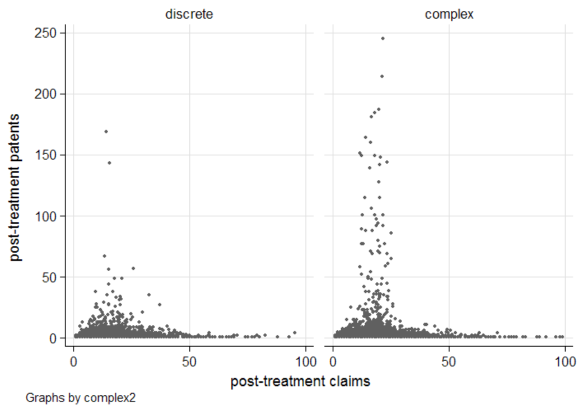

(c) Few narrow patents

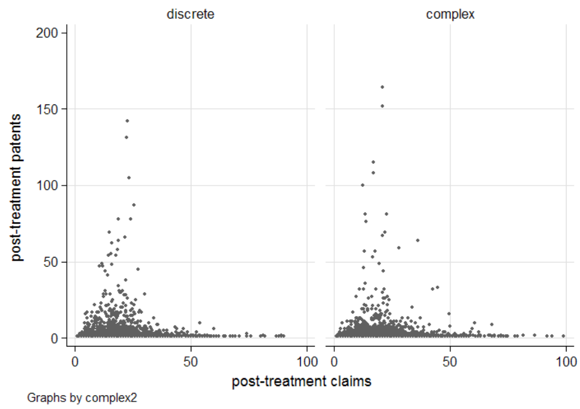

(d) Few broad patents

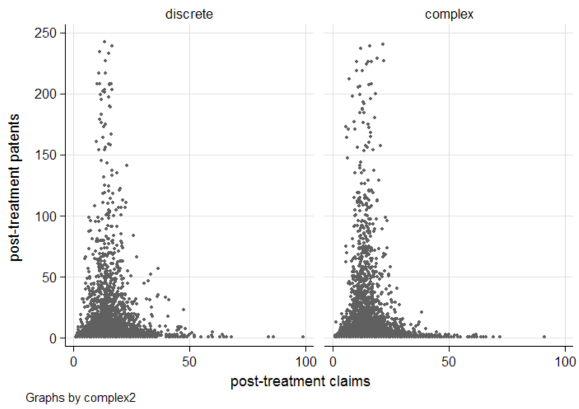

(e) Many narrow patents

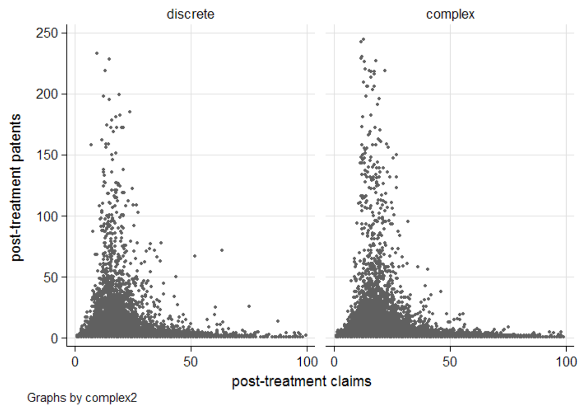

(f) Many broad patents

Figure 1.7: Scatter plots of (post-treatment) patents versus pre-treatment claims

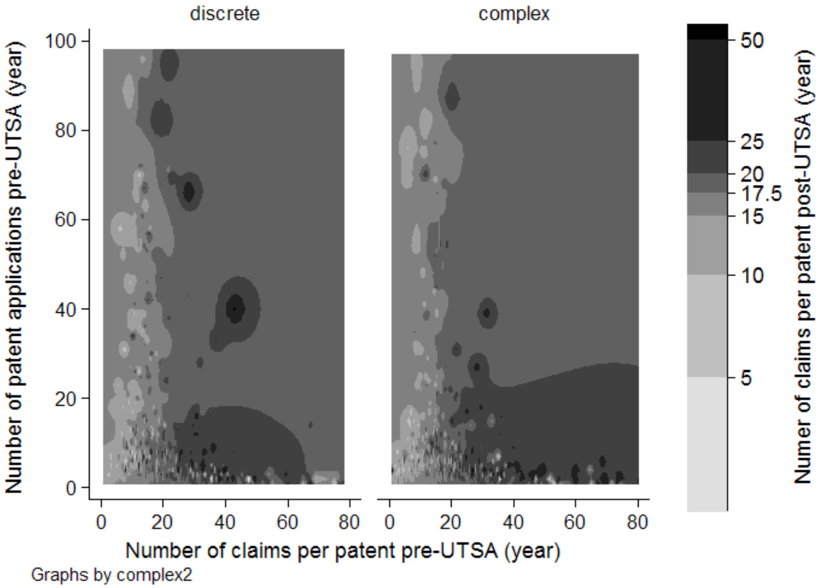

(a) Claims: Shepard interpoliation

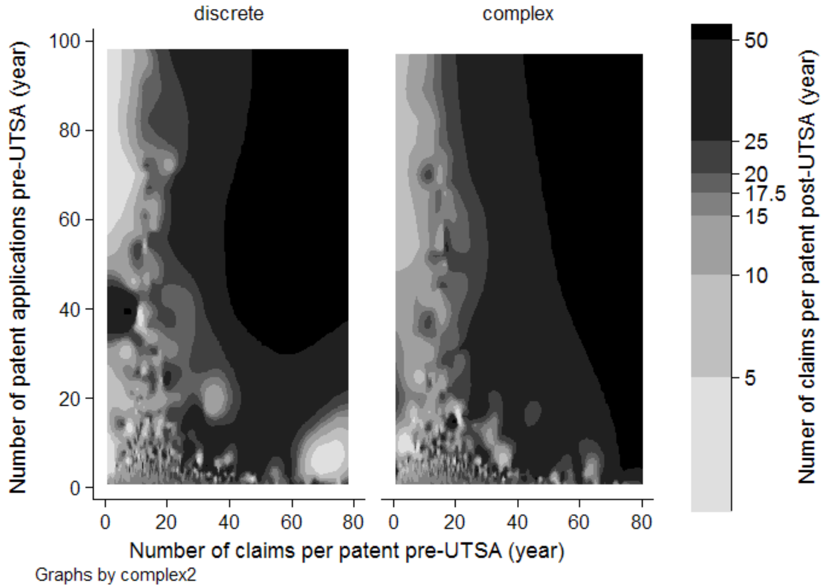

(b) Claims: Thin-plate-spline interpolation

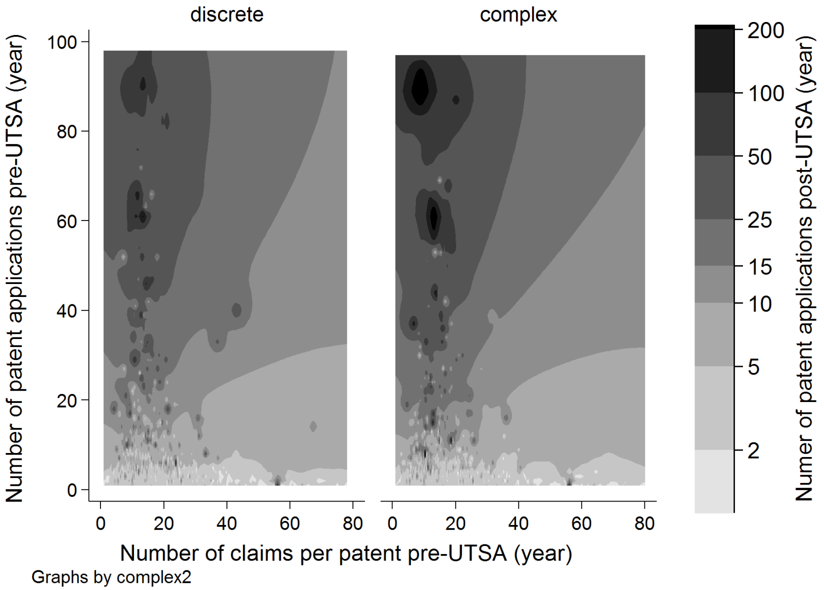

(c) Patents: Shepard interpolation

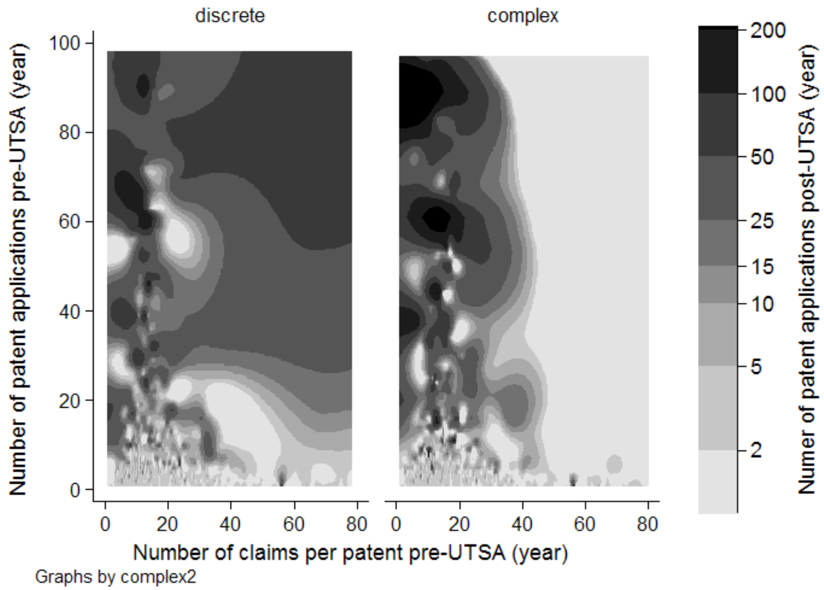

(d) Patents: Thin-plate-spline interpolation

Figure 1.8: Contour plots of post-treatment claims and patents with respect to pre-treatment claims and patents (single year)

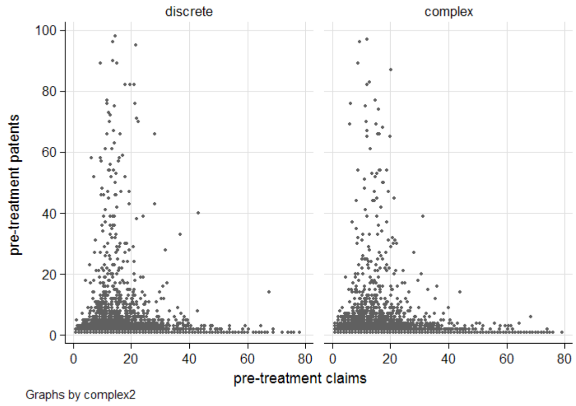

(a) Single year

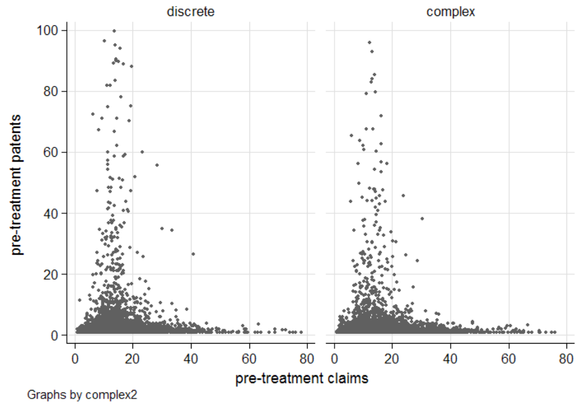

(b) Average 1

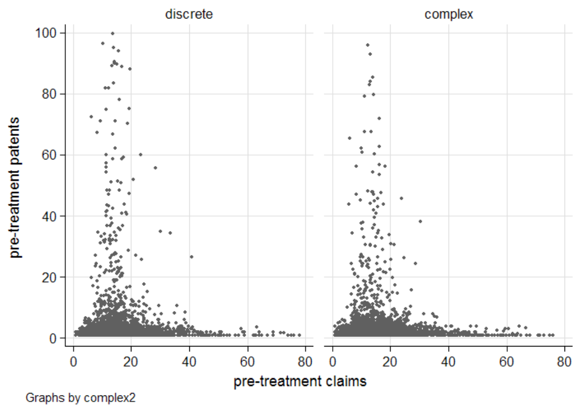

(c) Average 2

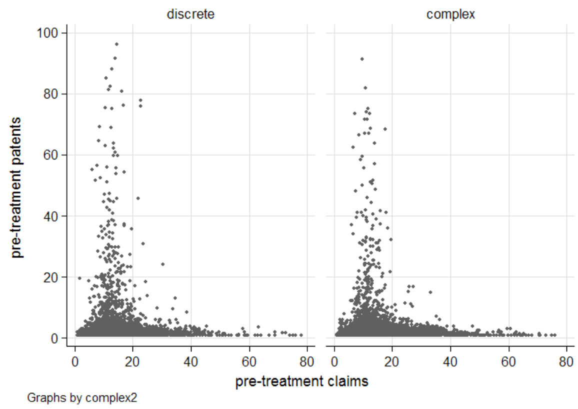

(d) Average 3

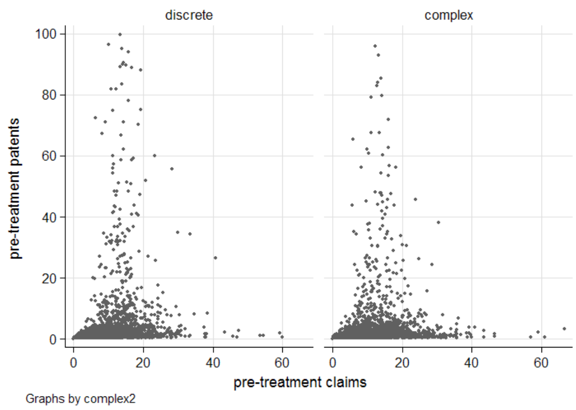

(e) Average 4

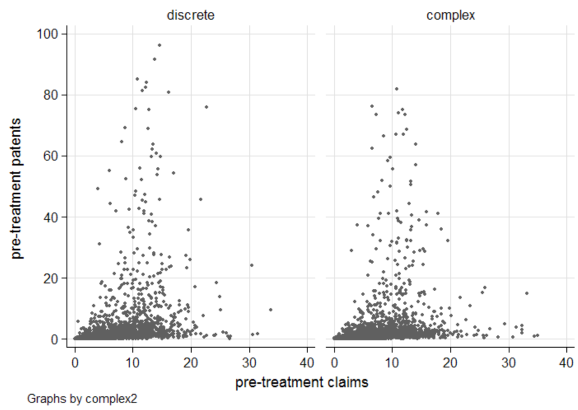

(f) Average 5

Figure 1.9: Scatter plots of pre-treatment patents versus pre-treatment claims

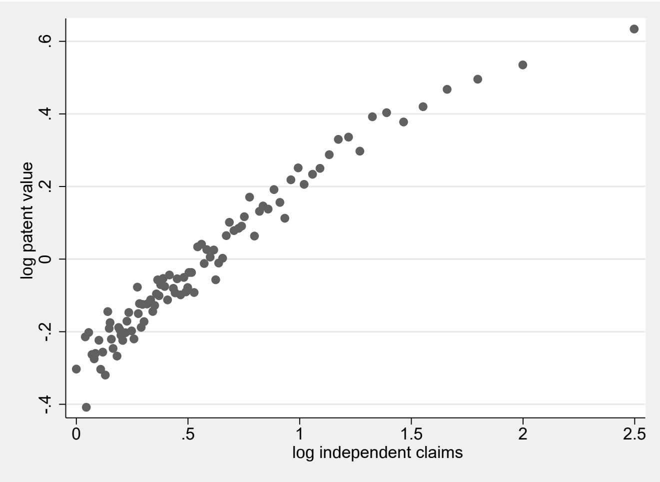

(a) Complex-product industries

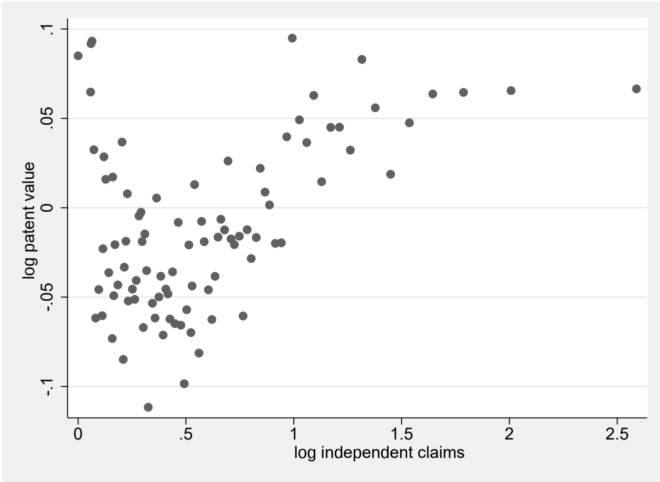

(b) Discrete-product industries

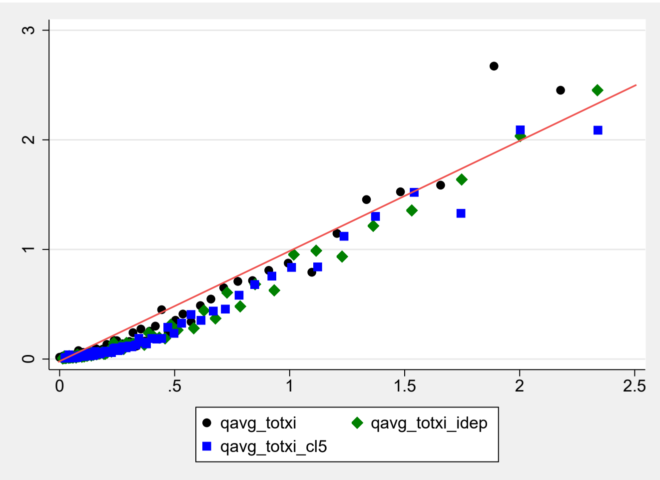

(c) Linear values on the horizontal axis

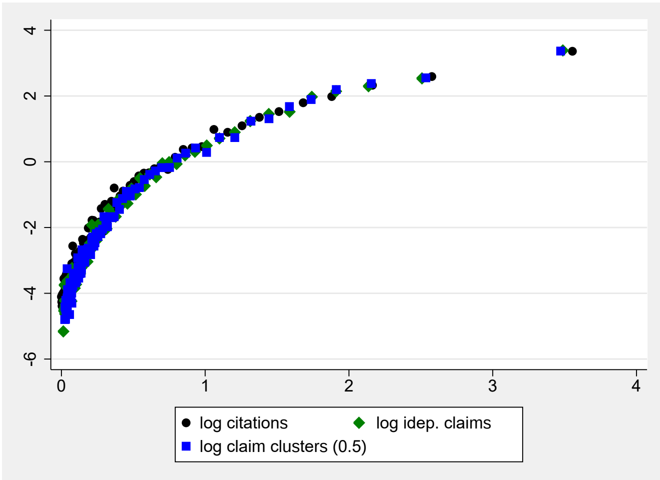

(d) Log values on the horizontal axis

Figure 1.10: Replication of Figure II of Kogan et at. (2017). Panels (a) and (b) are at the patent level, panels (c) and (d) at the patent-portfolio (firm) level. The vertical axis shows the log of patent value (as calculated by Kogan et al.), scaled by mean value of patents granted in the same year.

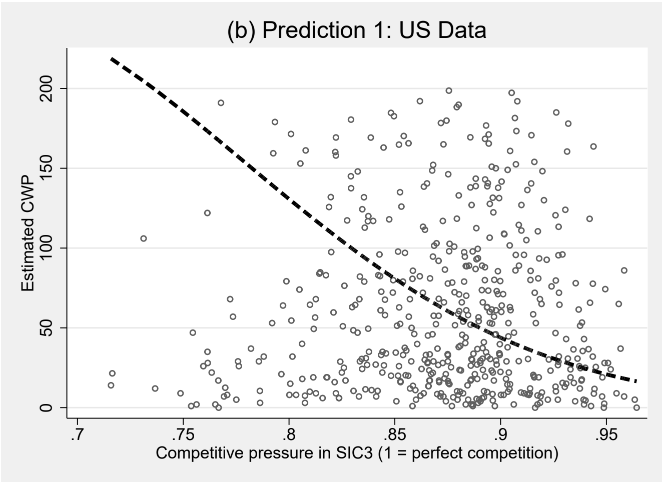

(a) Citations-weighted patents

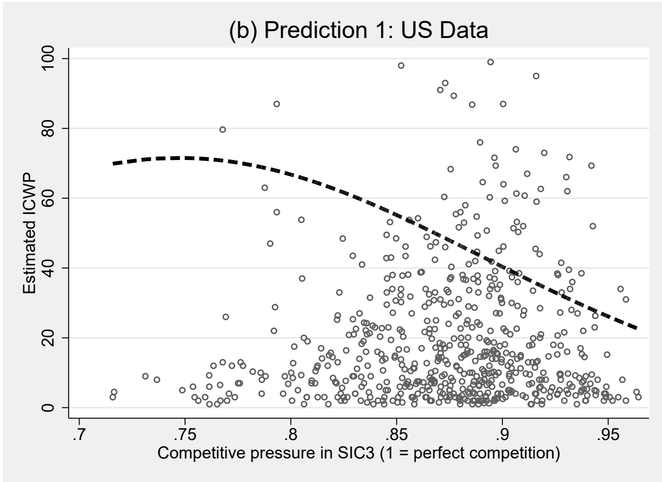

(b) Patents weighted by independent claims

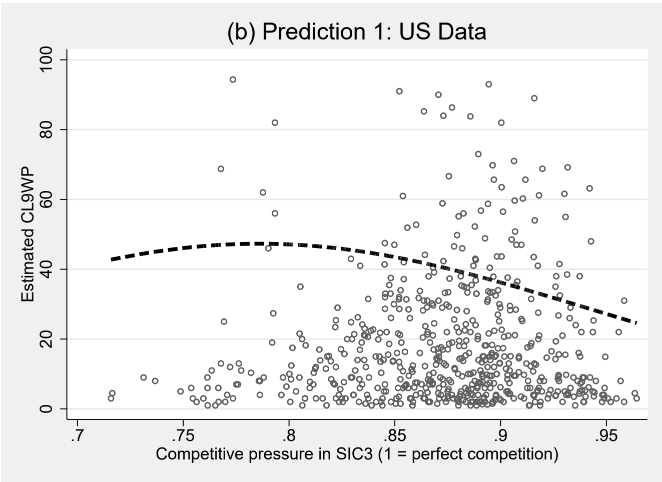

(c) Patents weighted by independent-claim clusters (0.9)

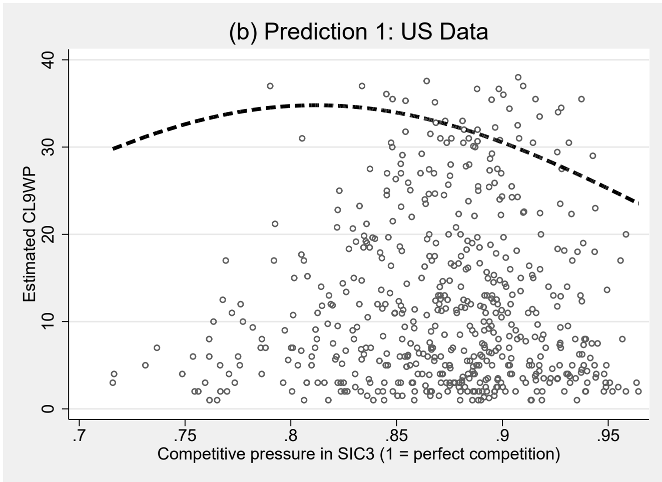

(d) Patents weighted by independent-claim clusters (0.7)

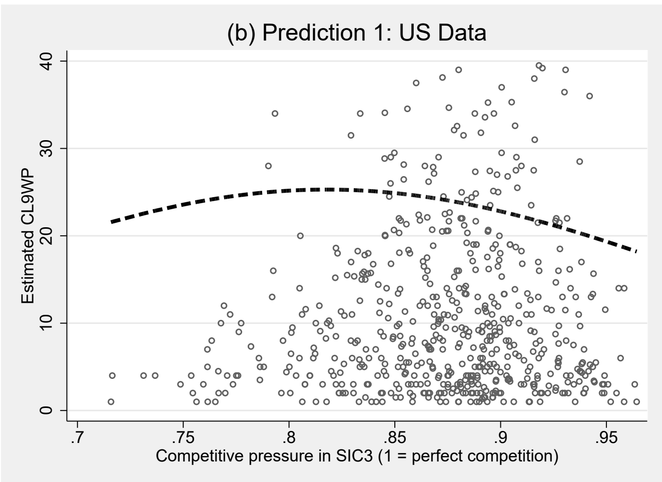

(e) Patents weighted by independent-claim clusters (0.5)

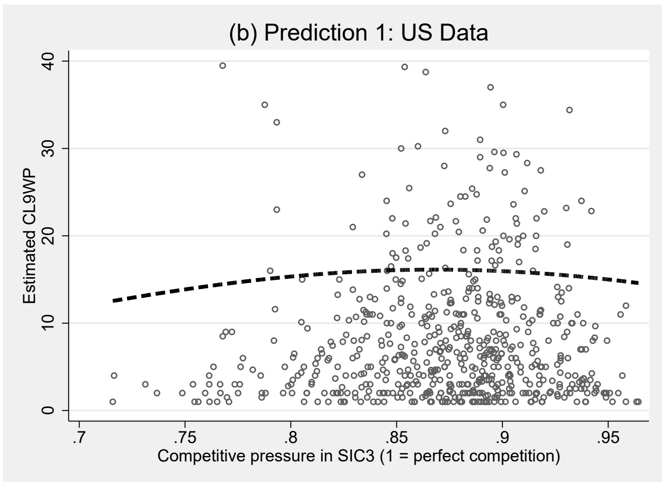

(f) Patents weighted by independent-claim clusters (0.3)

Figure 1.11: Replication of Figure 1(b) of Hashmi (2013). His prediction 1 reads: "there is an inverted-U relationship between product market competition and innovation." The dotted curve is the fitted line from the negative binomial regression performed in each case, the hollow circles are the predicted values.

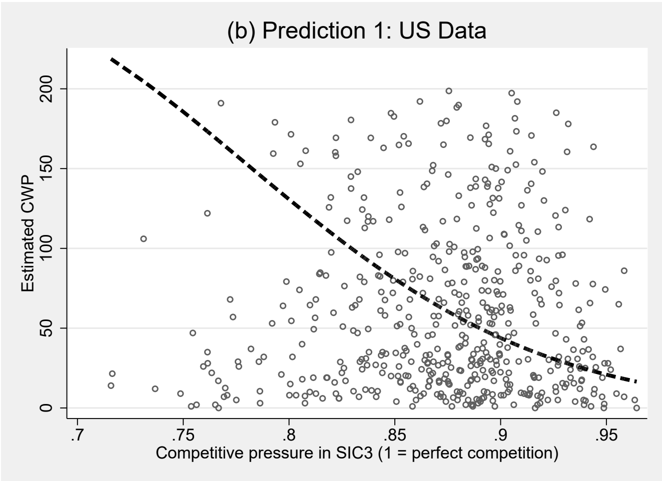

(a) Citations-weighted patents

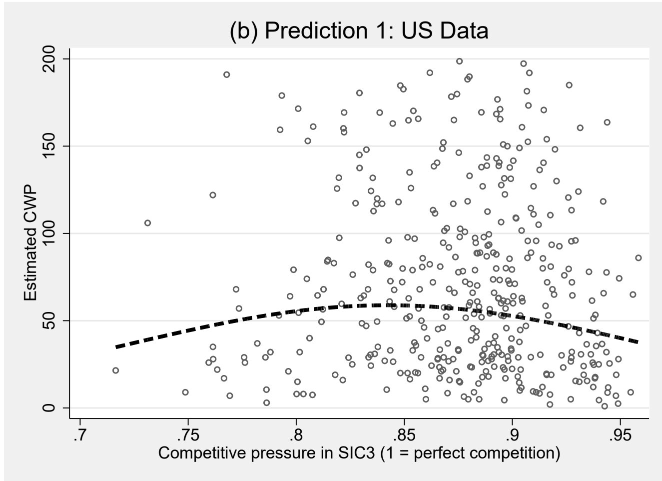

(b) Citations-weighted patents as in (a), but industries with highest across-patent claim text similarity excluded

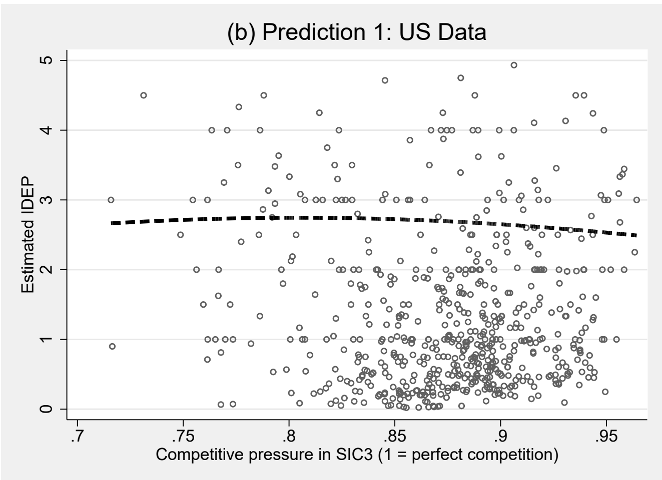

(c) Number of independent claims per patent

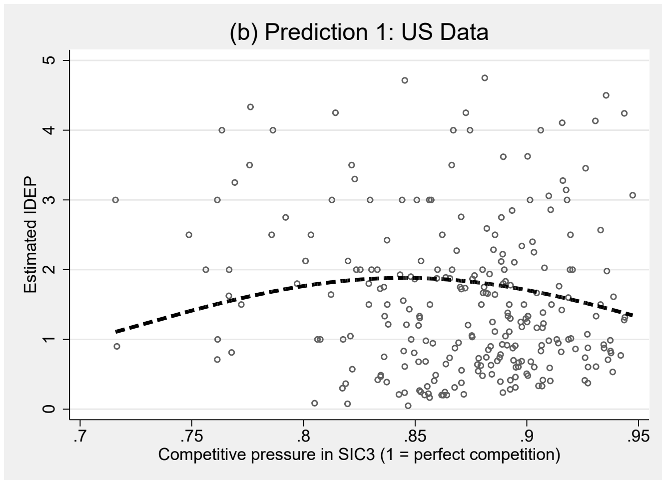

(d) Number of independent claims per patent as in (c), but industries with highest within-patent claim text similarity excluded

Figure 1.12: Replication of Figure 1(b) of Hashmi (2013). Panel (a) is equal to panel (a) in figure 1.11.

State	Initial level (Common law)	Year	Change	Year	New level
Alabama	0.025	1987*	0.21		
Alaska	0	1988	0.47		
Arizona	0.25	1990	0.22		
Arkansas	0.5	1981	−0.1		
California	0.22	1985	0.25		
Colorado	0	1986	0.77		
Connecticut	0	1983	0.47		
Delaware	0	1982	0.47		
District of Columbia	0	1989	0.47		
Florida	0.1	1988	0.37		
Georgia	0	1990	0.70		
Hawaii	0	1989	0.47		
Idaho	0	1981	0.47		
Illinois	0	1988	0.70		
Indiana	0	1982	0.47		
Iowa	0	1990	0.47		
Kansas	0	1981	0.47		
Kentucky	0	1990	0.47		
Louisiana	0	1981	0.40		
Maine	0	1987	0.50		
Maryland	0.22	1989	0.25		
Michigan	0.25	1998	0.15		
Minnesota	0	1980	0.47		
Mississippi	0	1990	0.57		
Missouri	0	1995	0.63		
Montana	0	1985	0.57		
Nebraska	0	1988	0.43		
Nevada	0	1987	0.47		
New Hampshire	0.025	1990	0.44		
New Mexico	0	1989	0.47		
North Carolina	0	1981*	0.73		
North Dakota	0	1983	0.47		
Ohio	0.25	1994	0.28		
Oklahoma	0.025	1986	0.44		
Oregon	0	1988	0.47		
Rhode Island	0	1986	0.47		
South Carolina	0	1992	0.47	1997*	0.17
South Dakota	0	1988	0.47		
Tennessee	0	2000	0.63		
Utah	0	1989	0.47		
Vermont	0	1996	0.57		
Virginia	0.025	1986	0.44		
Washington	0	1982	0.47		
West Virginia	0	1986	0.47		
Wisconsin	0	1986*	0.47		

* marks changes due to statutes that are not following the UTSA. Year: effective year of enactment; Initial level: Strength of legal protection of TS prior to enactment; Change: Change in legal protection of TS due to enactment.

Table 1.1: Changes in Trade Secret Protection due to Legal Statutes

State	Initial level	Year	Change	Year	Change	Year	Change
Alabama	0	1986	0.025				
Arizona	0	1980	0.25				
California	0.025	1978	0.22				
Florida	0	1984	0.1				
Michigan	0.02	1986	0.083	1994	0.167		
New Hampshire	0	1980	0.025				
New Jersey	0.25	1984	−0.05	1995	0.13		
New York	0.25	1979	0.1	1996	0.15	1997	0.16
Ohio	0	1986	0.25				
Oklahoma	0	1986	0.025				
Pennsylvania	0	1987	0.05	1997	0.19		
Texas	0.05	1983	0.19	1985	−0.01	1997	0.04

Table 1.2: Changes in Trade Secret Protection due to Common Law

State	Level
Massachusetts	0.27
Wyoming	0.5

Table 1.3: States without Changes in Trade Secret Protection

Component	24 Oct 1965	24 Jan 1978	12 Dec 1980	1 Oct 1982	10 Dec 1991	1 Oct 1998	29 Dec 1999
Application							
fee	65	65	65	300	500	760	690
independent claim > 1	10	10	10				
independent claim > 3				30	52	78	78
any claim > 10	2	2	2				
any claim > 20				10	14	18	18
multiple dependent claim*				100	160	260	260
International application**							
Search & Examination by USPTO					450	670	670
Only Search by USPTO					500	760	690
Neither Search nor Examination by USPTO					670	970	970
Issuing	100	100	100	500	820	1210	1210
Application for Reissue							
fee	65	65	65	300	500	760	690
any claim > 10 *and* > #claims in orig patent	2	2	2				
any claim > 20 *and* > #claims in orig patent				10	14	18	18
any claim > #indep. claims in orig patent	10	10	10				
any indep. claim > #indep. claims in orig patent				30	52	78	78
Maintenance							
After 3.5 years			***	400	650	940	830
After 7.5 years			***	800	1310	1900	1900
After 11.5 years			***	1200	1980	2910	2910
Remedies							
Disclaimer	15	15	15	50	78	110	110
Appeal to Board of Appeals	50	50	50	115	190	300	300
Brief in support of appeal				115	190	300	300
Request of oral hearing				100	160	260	260

Dates denote entry into force of the respective act. All values in nominal US$. *Fees for inclusion of multiple dependent claims need only be paid once per patent. If the total # claims in the patent is > 20, the multiple dependent claim is subject to the standard excessive claim fee times the number of references to independent claims it contains. **The Patent Cooperation Treaty entered into force on 24 January 1978, but separate fees for the national stage of international applications were not mentioned before 1991. ***Maintenance fees were first mentioned in the 1980 revision, but their value fixed by the law only in 1982.

Table 1.4: Most important components of the USPTO Patent Fee Schedule over time

	obs.	mean	standard deviation overall	between	within	P25	median	P75
Complex products								
Patent applications	12,144	19.28	75.98	35.37	46.64	1	3	9
Average no. claims	12,144	16.76	10.44	8.74	7.79	10.17	15.00	20.67
Avg. no. idep. claims	12,144	3.17	1.87	1.48	1.46	2.00	3.00	3.94
Revenue (real)	12,144	4,457.02	15,923.68	13,857.08	3,517.91	88.30	363.37	2372.85
EBITDA (real)	12,124	577.59	2,291.91	1,913.21	931.44	7.38	46.24	322.69
R&D (real)	12,144	82.15	382.06	234.31	177.52	2.85	8.43	28.84
Common law index	12,144	0.13	0.11	0.11	0.04	0	0.10	0.22
UTSA index	12,144	0.19	0.22	0.20	0.11	0	0	0.28
Assignees	1,484							
Assignees-locations	2,160							
Observations	12,144							
Discrete products								
Patent applications	7,143	18.31	48.28	25.78	22.22	1	3	12
Average no. claims	7,143	16.37	10.20	8.59	7.80	10.21	14.47	20.00
Avg. no. idep. claims	7,143	2.73	1.70	1.41	1.32	1.85	2.33	3.14
Revenue (real)	7,143	7,163.40	18,401.53	13,721.75	4,611.81	105.03	1597.09	6711.70
EBITDA (real)	7,126	1197.28	2,803.91	2,181.26	916.40	8.66	201.55	1101.90
R&D (real)	7,143	83.07	270.23	216.94	125.02	3.09	12.30	42.30
Common law index	7,143	0.12	0.12	0.11	0.04	0	0.10	0.22
UTSA index	7,143	0.17	0.23	0.21	0.11	0	0	0.37
Assignees	783							
Assignees-locations	1,175							
Observations	7,143							

Table 1.5: Descriptive statistics of the main dataset — Patent data matched to Compustat financial data — at the company-state level — Complex and discrete technology industries according to SIC-ISIC correspondence for the sample period 1979–2000

	obs.	mean	standard deviation overall	between	within	P25	median	P75
Complex products								
Including years with zero patents:								
Patent applications:	4,097	25.07	88.47	71.12	45.39	1	3	14
Employees	4,097	9.39	18.71	15.69	4.59	0.68	1.98	8.47
Revenue (real)	4,097	1716.95	3715.95	3421.39	989.12	104.72	301.03	1418.42
PPE (real)	4,097	832.01	1974.40	1710.95	529.45	41.06	141.70	671.67
R&D exp. (real)	4,097	67.92	172.25	151.53	63.37	2.65	10.53	45.56
No R&D exp. (dummy)	4,097	0.07	0.25	0.22	0.10	0	0	0
Facility-weighted:								
Common law index	4,097	0.13	0.10	0.10	0.04	0	0.11	0.22
UTSA index	4,097	0.20	0.20	0.19	0.09	0	0.21	0.35
Headquarter state:								
Common law index	2,924	0.13	0.12	0.11	0.04	0	0.10	0.23
UTSA index	2,924	0.24	0.24	0.22	0.06	0	0.25	0.47
Excluding years with zero patents:								
Patent applications:	3,133	33.39	99.90	71.76	51.84	2	6	25
Average no. claims	3,133	16.73	8.30	6.95	6.42	11.80	15.625	20.00
Avg. no. idep. cl.	3,133	3.09	1.43	1.09	1.13	2.23	2.91	3.61
Assignees	469							
Observations	4,097							
Discrete products								
Including years with zero patents:								
Patent applications:	3,194	19.96	51.08	40.70	20.96	0	3	13
Employees	3,194	13.65	18.89	17.60	5.67	1.90	5.80	18.50
Revenue (real)	3,194	3337.31	5196.84	4803.34	1407.30	424.31	1315.40	4129.76
PPE (real)	3,194	2613.06	5124.26	4792.61	1079.57	191.94	692.84	2956.69
R&D exp. (real)	3,194	75.32	210.52	188.33	79.46	0	7.18	42.85
No R&D exp. (dummy)	3,194	0.28	0.45	0.44	0.12	0	0	1
Facility-weighted:								
Common law index	3,194	0.11	0.10	0.10	0.04	0	0.08	0.20
UTSA index	3,194	0.19	0.22	0.20	0.12	0	0.11	0.36
Headquarter state:								
Common law index	2,025	0.11	0.12	0.11	0.04	0	0.05	0.23
UTSA index	2,025	0.24	0.27	0.26	0.08	0	0	0.47
Excluding years with zero patents:								
Patent applications:	2,266	27.63	58.71	40.91	24.69	2	6	24
Average no. claims	2,266	16.84	9.44	7.00	7.57	11.67	15.10	20.00
Avg. no. idep. cl.	2,266	2.74	1.43	1.10	1.14	2.00	2.42	3.14
Assignees	324							
Observations	3,194							

Table 1.6: Descriptive statistics of the dataset used in Png (2017b) — Patent data matched to Compustat financial data and Bowker R&D data - at the parent-company level — Complex and discrete technology industries according to SIC-ISIC correspondence for the sample period 1984–1999

	obs.	mean	standard deviation overall	between	within	P25	median	P75
Complex products								
Patent applications	67,272	4.86	35.82	12.73	22.78	1	1	3
Average no. indep. claims	66,552	3.17	2.30	2.08	1.65	2.00	3.00	4.00
Average no. total claims	67,272	16.62	12.63	12.28	8.54	9.00	14.00	20.58
Average no. IPC subclasses	67,271	1.34	0.58	0.46	0.45	1.00	1.00	1.50
Assignees	15,708							
Assignees-locations	20,598							
Observations	67,272							
Discrete products								
Patent applications	75,387	6.27	32.21	13.11	16.28	1	1	3
Average no. indep. claims	74,354	2.87	1.98	1.62	1.50	1.76	2.40	3.50
Average no. total claims	75,387	15.92	11.23	9.52	8.18	9.00	14.00	20.00
Average no. IPC subclasses	75,387	1.58	0.79	0.65	0.58	1.00	1.25	2.00
Assignees	16,125							
Assignees-locations	18,465							
Observations	75,387							

Table 1.7: Descriptive statistics of the full NBER patent dataset — Patent data only at the company-state level — Complex and discrete technology industries according to IPC-SIC-ISIC correspondence for the sample period 1979–1998

	obs.	mean	standard deviation overall	between	within	P25	median	P75
few narrow patents								
Patent applications	23,046	2.70	12.51	4.86	9.59	1	1	2
Average no. indep. claims	22,709	2.52	1.72	1.03	1.40	1.33	2.00	3.00
Average no. total claims	23,046	13.59	9.73	6.86	7.54	7.00	11.91	18.00
few broad patents								
Patent applications	11,324	2.32	6.96	3.28	5.12	1	1	2
Average no. indep. claims	11,114	3.78	2.48	1.81	1.98	2.00	3.00	4.67
Average no. total claims	11,324	17.84	12.25	9.97	9.24	10.00	15.33	22.25
many narrow patents								
Patent applications	3,629	24.66	66.77	39.37	39.48	2	6	18
Average no. indep. claims	3,606	2.60	1.30	0.77	1.12	1.91	2.40	3.00
Average no. total claims	3,629	13.43	7.38	5.39	6.03	9.00	12.42	16.42
many broad patents								
Patent applications	2,931	25.17	143.54	81.69	93.20	2	5	14
Average no. indep. claims	2,915	3.49	1.78	1.39	1.51	2.45	3.16	4.00
Average no. total claims	2,931	17.07	9.18	6.35	7.58	11.32	15.50	20.08

Values are for current-year numbers of patents and claims, including **both pre- and post-treatment periods**.

Table 1.8: Descriptive statistics of the four groups of pre-treatment patent-claims combinations in the NBER-only dataset

	obs.	mean	standard deviation overall	between	within	P25	median	P75
few narrow patents								
Patent applications	8,498	3.71	17.10	5.49	12.53	1	1	2
Average no. indep. claims	8,450	2.93	1.98	8.58	7.38	2.00	2.50	3.64
Average no. total claims	8,498	15.54	10.91	8.58	7.38	8.25	13.67	20.00
few broad patents								
Patent applications	4,686	2.45	3.25	3.51	6.94	1	1	3
Average no. indep. claims	4,658	3.49	2.23	8.58	7.38	2.00	3.00	4.00
Average no. total claims	4,686	17.63	11.47	10.51	8.22	10.00	15.79	22.18
many narrow patents								
Patent applications	1,594	27.24	76.10	211.60	64.73	2	7	22
Average no. indep. claims	1,591	2.91	1.43	8.58	7.38	2.00	2.69	3.45
Average no. total claims	1,594	15.39	7.97	5.21	5.53	10.67	14.31	18.67
many broad patents								
Patent applications	1,170	31.73	186.11	85.67	27.15	2	6	20
Average no. indep. claims	1,168	3.55	1.72	8.58	7.38	2.51	3.20	4.00
Average no. total claims	1,170	18.19	8.89	6.41	6.59	13.00	16.83	21.00

Values include **post-treatment years only**.

Table 1.9: Descriptive statistics of the four groups of pre-treatment patent-claims combinations in the NBER-only dataset

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent var.:	Patents				Claims			
Products:	all	complex	complex nn.	discrete	all	complex	complex nn.	discrete
UTSA index	0.077	-0.048	-0.022	0.263	0.017	-0.058	-0.066	0.035
	(0.115)	(0.123)	(0.113)	(0.184)	(0.033)	(0.048)	(0.051)	(0.044)
Common law index	-0.405	-0.561	-0.450	-0.399	0.030	0.107	0.133	-0.164
	(0.382)	(0.516)	(0.508)	(0.280)	(0.094)	(0.101)	(0.108)	(0.123)
Company age	-0.005	-0.021	0.001	-0.035	0.013	0.038	0.045	-0.010
	(0.025)	(0.037)	(0.038)	(0.046)	(0.019)	(0.034)	(0.034)	(0.038)
EBITDA (log)	-0.002	-0.009	-0.011	-0.087**	-0.002	0.011	0.013	-0.029
	(0.018)	(0.023)	(0.022)	(0.036)	(0.007)	(0.009)	(0.009)	(0.024)
Revenue (log)	-0.059	0.029	0.000	-0.070	0.014	0.010	0.022	-0.003
	(0.072)	(0.095)	(0.098)	(0.079)	(0.017)	(0.017)	(0.017)	(0.046)
R&D exp. (log)	0.806****	0.766****	0.771****	0.830****	0.025****	0.040****	0.031***	0.028**
	(0.022)	(0.030)	(0.033)	(0.034)	(0.007)	(0.010)	(0.011)	(0.014)
R^2	0.37	0.35	0.35	0.37	0.03	0.03	0.03	0.03
Observations	16,801	8,731	7,859	5,080	16,801	8,731	7,859	5,080
Assignees	1,691	952	887	432	1,691	952	887	432
Assignee locations	2,450	1,315	1,216	665	2,450	1,315	1,216	665

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of granted patents (columns 1–4) / average number of independent claims per patent (columns 5–8), all in logs. Samples restricted as stated. Columns (3) and (7) restricted to non-negative financial variables. Selection of complex vs discrete product industries is based on Compustat 3-digit SIC codes, translated into ISIC rev.3 codes from the US Office of Management and Budget. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.10: Regression of patents and independent claims in the whole sample (matched to Compustat data).

Average:	(1) 1 year	(2) 1 year p75	(3) 4 avlb. years	(4) 4 years	(5) all years	(6) 4 years, zero	(7) all years, zero
UTSA index	-0.248	-0.126*	-0.336***	-0.113	-0.287**	-0.332***	-0.696***
	(0.186)	(0.073)	(0.110)	(0.086)	(0.131)	(0.107)	(0.213)
Common law index	0.005	0.046	0.252	0.136	0.366	0.197	0.351
	(0.337)	(0.169)	(0.271)	(0.216)	(0.267)	(0.340)	(0.442)
Company age	-0.019	0.037	0.031	-0.019	0.031	0.021	0.014
	(0.060)	(0.084)	(0.063)	(0.072)	(0.062)	(0.080)	(0.049)
EBITDA (log)	-0.041*	-0.018	0.019	-0.006	0.019	0.003	0.012
	(0.023)	(0.021)	(0.018)	(0.022)	(0.018)	(0.022)	(0.019)
Revenue (log)	0.051	0.070*	0.017	0.012	0.017	-0.041	0.066
	(0.038)	(0.039)	(0.039)	(0.031)	(0.040)	(0.045)	(0.061)
R&D exp. (log)	0.071***	0.051**	0.042*	0.083**	0.054**	0.102****	-0.009
	(0.026)	(0.025)	(0.024)	(0.036)	(0.026)	(0.028)	(0.029)
R^2	0.05	0.04	0.02	0.04	0.03	0.04	0.05
Observations	1,440	2,735	1,813	1,841	1,809	1,818	1,316
Assignees	231	373	345	255	336	306	269
Assignee locations	262	453	398	296	390	365	312

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual average number of independent claims per patent, in logs. Samples restricted to **below-median** number of annual patents before treatment as follows: Columns (1) and (2) based on single most recent pre-treatment year (column (1) is below-median, column (2) is lower three quartiles); column (3) based on average of four most recent available pre-treatment years; column (4) based on average of four pre-treament years; column (5) based on average over all available pre-treatment years; column (6) based on average over four pre-treatment years, treating missing as zero; column (7) based on average over all pre-treatment years, treating missing as zero. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.11: Regression of independent claims when the pre-treatment number of patents is small.

Average:	(1) 1 year	(2) 1 year p75	(3) 4 avlb. years	(4) 4 years	(5) all years	(6) 4 years, zero	(7) all years, zero
UTSA index	-0.035	-0.048	-0.139**	-0.036	-0.176***	-0.005	-0.027
	(0.055)	(0.060)	(0.064)	(0.052)	(0.061)	(0.058)	(0.051)
Common law index	0.159	0.188	0.118	0.132	0.065	0.166	0.159
	(0.101)	(0.140)	(0.199)	(0.103)	(0.186)	(0.107)	(0.131)
Company age	0.021	-0.019	-0.005	-0.010	-0.008	0.004	0.037
	(0.060)	(0.045)	(0.041)	(0.057)	(0.041)	(0.051)	(0.034)
EBITDA (log)	0.032**	0.036**	0.019	0.016	0.023	0.005	0.016
	(0.014)	(0.015)	(0.020)	(0.015)	(0.019)	(0.015)	(0.013)
Revenue (log)	0.005	-0.025	-0.011	0.023	-0.014	0.054**	0.013
	(0.026)	(0.028)	(0.041)	(0.025)	(0.039)	(0.021)	(0.025)
R&D exp. (log)	0.034*	0.038**	0.033*	0.021	0.026	0.017	0.039**
	(0.018)	(0.019)	(0.018)	(0.016)	(0.019)	(0.017)	(0.015)
R^2	0.05	0.07	0.05	0.06	0.05	0.06	0.05
Observations	4,244	2,949	2,542	3,302	2,546	3,866	5,419
Assignees	353	210	281	231	286	275	481
Assignee locations	438	247	325	276	333	335	612

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual average number of independent claims per patent, in logs. Samples restricted to **below-median** number of annual patents before treatment as follows: Columns (1) and (2) based on single most recent pre-treatment year (column (1) is below-median, column (2) is lower three quartiles); column (3) based on average of four most recent available pre-treatment years; column (4) based on average of four pre-treament years; column (5) based on average over all available pre-treatment years; column (6) based on average over four pre-treatment years, treating missing as zero; column (7) based on average over all pre-treatment years, treating missing as zero. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.12: Regression of independent claims when the pre-treatment number of patents is large.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Restriction:	below median				below median			
Average:	1 year	4 avlb. years	4 years	all years	1 year	4 avlb. years	4 years	all years
UTSA index	-0.184**	-0.277**	-0.195***	-0.321***	0.050	-0.119*	0.117	-0.077
	(0.085)	(0.118)	(0.068)	(0.105)	(0.112)	(0.071)	(0.087)	(0.088)
Common law index	-0.085	0.139	0.032	0.174	0.364**	0.146	0.256	0.158
	(0.118)	(0.235)	(0.139)	(0.232)	(0.179)	(0.214)	(0.226)	(0.235)
Company age	0.011	0.023	-0.010	0.014	0.015	-0.060	-0.011	-0.035
	(0.057)	(0.033)	(0.053)	(0.034)	(0.060)	(0.100)	(0.056)	(0.105)
EBITDA (log)	-0.012	0.018	-0.006	0.029**	0.028	0.024	0.016	0.014
	(0.015)	(0.014)	(0.019)	(0.015)	(0.023)	(0.025)	(0.021)	(0.024)
Revenue (log)	0.067	0.040	0.068*	0.020	-0.036	-0.030	-0.021	-0.009
	(0.046)	(0.039)	(0.036)	(0.041)	(0.033)	(0.041)	(0.033)	(0.041)
R&D exp. (log)	0.044	0.025	0.038	0.030	0.041**	0.039	0.044**	0.029
	(0.028)	(0.022)	(0.028)	(0.021)	(0.016)	(0.025)	(0.021)	(0.022)
R^2	0.06	0.04	0.06	0.04	0.04	0.03	0.05	0.03
Observations	3,219	2,403	2,965	2,331	2,417	1,920	2,033	1,992
Assignees	280	327	237	317	259	267	206	268
Assignee locations	380	398	314	391	310	313	236	320

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual average number of independent claims per patent, in logs. Samples restricted to annual patenting propensity (annual patents divided by R&D expenditure) before treatment as follows: Columns (1) and (5) based on single most recent pre-treatment year; columns (2) and (6) based on average of four most recent available pre-treatment years; columns (3) and (7) based on average of four pre-treament years; columns (5) and (8) based on average over all available pre-treatment years. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.13: Regression of independent claims, conditional on pre-treatment “patenting propensity”.

Pre-treatment sample:	(1) many claims	(2) few claims	(3) many claims	(4) few claims	(1) many claims	(2) few claims	(3) many claims	(4) few claims
UTSA index	0.012	-0.581****	-0.013	-0.222	-0.069	-0.788***	-0.161	-0.600****
	(0.174)	(0.171)	(0.131)	(0.155)	(0.387)	(0.297)	(0.258)	(0.162)
Common law index	-0.168	-0.065	0.019	0.195	-0.142	-0.003	0.038	0.140
	(0.426)	(0.460)	(0.359)	(0.513)	(0.435)	(0.253)	(0.400)	(0.494)
Company age	-0.065	0.023	-0.126**	0.007	-0.061	-0.139	-0.142**	0.009
	(0.080)	(0.064)	(0.058)	(0.053)	(0.087)	(0.085)	(0.058)	(0.058)
EBITDA (log)	-0.111	-0.020	0.002	0.058	-0.121*	-0.058	0.017	0.032
	(0.068)	(0.087)	(0.071)	(0.069)	(0.065)	(0.082)	(0.073)	(0.067)
Revenue (log)	0.134	0.184	-0.022	0.022	0.180	0.528****	-0.078	0.087
	(0.207)	(0.215)	(0.147)	(0.151)	(0.233)	(0.123)	(0.162)	(0.149)
R&D exp. (log)	0.889****	0.705****	0.822****	0.725****	0.890****	0.599****	0.832****	0.687****
	(0.099)	(0.117)	(0.073)	(0.090)	(0.104)	(0.073)	(0.076)	(0.078)
R^2	0.50	0.40	0.41	0.39	0.52	0.48	0.43	0.43
Observations	865	792	1,717	1,510	865	792	1,717	1,510
Assignees	55	49	127	108	55	49	127	108
Assignee locations	59	53	139	120	59	53	139	120

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of granted patents, in logs. Samples restricted by annual patents and claims before treatment as follows: Columns (1) and (2) restricted to 5th quintile of patents; columns (3) and (4) restricted to 4th & 5th quintile of patents. Additionally: (1) and (3) restriected to 3rd–5th quintile of claims; (2) and (4) restricted to 1st–3rd quintile of claims. Columns (5)–(8) are equivalent to (1)–(4) but additionally contain state-specific linear time trends. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.14: Regression of patents, conditional on pre-treatment patents and claims.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent var.:	patents				claims			
Pre-treatment patent prop.:	high		low		high		low	
Pre-treatment claims:	many	few	many	few	many	few	many	few
UTSA index	-0.274	-0.389*	-0.047	0.268	-0.222**	0.106	-0.269***	-0.107
	(0.279)	(0.203)	(0.176)	(0.241)	(0.102)	(0.098)	(0.095)	(0.087)
Common law index	-0.728	0.626	-0.582	-0.466	0.412	0.433	0.345	-0.064
	(0.793)	(0.480)	(0.466)	(0.882)	(0.259)	(0.304)	(0.255)	(0.198)
Company age	-0.091	0.038	-0.130	-0.066	-0.042	0.093	-0.128*	0.074
	(0.104)	(0.072)	(0.100)	(0.050)	(0.090)	(0.096)	(0.072)	(0.074)
EBITDA (log)	-0.005	0.060	0.003	0.007	0.024	0.028	0.029	-0.028*
	(0.052)	(0.052)	(0.039)	(0.041)	(0.028)	(0.027)	(0.024)	(0.016)
Revenue (log)	-0.191*	-0.173	-0.011	0.104	-0.023	-0.026	-0.039	0.104**
	(0.106)	(0.154)	(0.141)	(0.169)	(0.053)	(0.034)	(0.049)	(0.052)
R&D exp. (log)	0.836****	0.884****	0.842****	0.843****	-0.004	0.061**	0.042	0.042
	(0.095)	(0.078)	(0.048)	(0.065)	(0.033)	(0.027)	(0.028)	(0.032)
R^2	0.30	0.36	0.37	0.47	0.06	0.04	0.06	0.07
Observations	1,246	1,873	1,358	2,412	1,242	1,870	1,352	2,392
Assignees	141	213	133	215	140	213	133	215
Assignee locations	156	249	156	280	155	249	156	278

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of granted patents or claims as stated, in logs. Samples restricted by annual patent propensity and claims before treatment as follows: Columns (1), (2), (5) and (6) restricted to the 3rd–5th quintiles of patent propensity; columns (3), (4), (7) and (8) restricted to the 1st–3rd quintiles. Additionally: (1), (3), (5) and (7) restriected to 4th–5th quintile of claims; (2), (4), (6) and (8) restricted to 1st–3rd quintile of claims. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.15: Regression of patents and claims, conditional on pre-treatment patent propensity and claims.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pre-treatment patents:	high				low			
Pre-treatment claims:	many	few	many	few	many	few	many	few
UTSA index	-0.152	-0.222	0.047	-0.581****	0.167	0.262	-0.003	0.227
	(0.224)	(0.155)	(0.182)	(0.171)	(0.308)	(0.230)	(0.189)	(0.216)
Common law index	-0.651	0.195	-1.092*	-0.065	-0.284	-0.656	-0.239	-0.028
	(0.519)	(0.513)	(0.660)	(0.460)	(0.741)	(1.159)	(0.449)	(0.856)
Age	-0.212**	0.007	-0.115	0.023	-0.021	-0.057	-0.116	-0.042
	(0.083)	(0.053)	(0.119)	(0.064)	(0.110)	(0.090)	(0.091)	(0.062)
EBITDA (log)	-0.040	0.058	-0.111	-0.020	0.062	-0.002	0.034	0.033
	(0.066)	(0.069)	(0.090)	(0.087)	(0.043)	(0.043)	(0.040)	(0.042)
Revenue (log)	0.003	0.022	0.085	0.184	-0.286***	-0.019	-0.175*	-0.060
	(0.142)	(0.151)	(0.256)	(0.215)	(0.098)	(0.140)	(0.102)	(0.145)
R&D exp. (log)	0.851****	0.725****	0.923****	0.705****	0.697****	0.900****	0.728****	0.871****
	(0.086)	(0.090)	(0.132)	(0.117)	(0.053)	(0.073)	(0.063)	(0.064)
R^2	0.40	0.39	0.51	0.40	0.23	0.42	0.25	0.41
Observations	1,071	1,510	484	792	1,113	2,020	1,700	2,738
Assignees	84	108	31	49	159	264	205	314
Assignee locations	91	120	34	53	171	321	228	388

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of granted patents, in logs. Samples restricted by annual patents and claims before treatment as follows: Columns (1) & (2) restricted to the 4th–5th quintiles of patents; columns (3) & (4) restricted to the 5th quintile; columns (5) & (6) restricted to the 1st–3rd quintiles of patents; (7) & (8) to the 1st–4th quintiles. Additionally: (1), (3), (5) and (7) restriected to 4th–5th quintile of claims; (2), (4), (6) and (8) restricted to 1st–3rd quintile of claims. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.16: Regression of patents, conditional on pre-treatment number of patents and claims.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pre-treatment patents:	high				low			
Pre-treatment claims:	many	few	many	few	many	few	many	few
UTSA index	-0.308****	0.063	-0.295**	-0.028	-0.363***	0.110	-0.273****	-0.022
	(0.078)	(0.072)	(0.132)	(0.122)	(0.140)	(0.091)	(0.075)	(0.103)
Common law index	0.316	0.051	0.294	-0.018	0.189	-0.119	0.383*	0.140
	(0.200)	(0.126)	(0.275)	(0.421)	(0.373)	(0.155)	(0.217)	(0.292)
Company age	0.005	-0.031	-0.123	0.123	-0.027	-0.059	-0.071	0.082
	(0.081)	(0.055)	(0.098)	(0.103)	(0.081)	(0.044)	(0.074)	(0.081)
EBITDA (log)	0.002	0.030	0.036	-0.032	-0.026	0.035	0.029	-0.016
	(0.019)	(0.019)	(0.035)	(0.022)	(0.025)	(0.025)	(0.027)	(0.019)
Revenue (log)	-0.020	0.040	-0.003	0.056	0.046	0.019	-0.028	0.056*
	(0.051)	(0.050)	(0.074)	(0.042)	(0.080)	(0.056)	(0.048)	(0.033)
R&D exp. (log)	0.016	0.019	0.033	0.073***	-0.002	0.056	0.030	0.051**
	(0.025)	(0.041)	(0.039)	(0.024)	(0.034)	(0.057)	(0.028)	(0.023)
R^2	0.11	0.11	0.04	0.04	0.17	0.19	0.03	0.05
Observations	1,068	1,509	1,107	2,000	482	793	1,693	2,716
Assignees	84	108	158	262	31	49	204	312
Assignee locations	91	120	170	319	34	53	227	386

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual average number of independent claims per patent, in logs. Samples restricted by annual patents and claims before treatment as follows: Columns (1) & (2) restricted to the 4th–5th quintiles of patents; columns (3) & (4) restricted to the 5th quintile; columns (5) & (6) restricted to the 1st–3rd quintiles of patents; (7) & (8) to the 1st–4th quintiles. Additionally: (1), (3), (5) and (7) restriected to 4th–5th quintile of claims; (2), (4), (6) and (8) restricted to 1st–3rd quintile of claims. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.17: Regression of claims, conditional on pre-treatment patents and claims.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pre-treatment sample:	many patents				many claims			
Dependent var.:	patents		claims		patents		claims	
Industries:	complex	discrete	complex	discrete	complex	discrete	complex	discrete
UTSA index	-0.425***	0.285	0.003	0.058	-0.053	0.132	-0.292****	-0.161**
	(0.144)	(0.251)	(0.062)	(0.079)	(0.141)	(0.148)	(0.064)	(0.066)
Common law index	-0.399	0.678	-0.023	-0.060	-0.508	0.289	0.285	0.069
	(0.339)	(0.768)	(0.124)	(0.188)	(0.357)	(0.396)	(0.210)	(0.067)
Company age	-0.046	-0.228***	-0.046	-0.022	-0.132**	-0.092	-0.062	0.132**
	(0.050)	(0.085)	(0.044)	(0.035)	(0.067)	(0.120)	(0.068)	(0.060)
EBITDA (log)	-0.056	0.046	0.011	-0.015	0.009	-0.012	0.021	-0.040
	(0.073)	(0.079)	(0.027)	(0.028)	(0.038)	(0.071)	(0.022)	(0.041)
Revenue (log)	0.127	0.025	0.027	0.058	-0.110	-0.157	-0.015	0.013
	(0.205)	(0.118)	(0.059)	(0.055)	(0.075)	(0.148)	(0.037)	(0.084)
R&D exp. (log)	0.800****	0.776****	0.035	0.007	0.789****	0.884****	0.022	0.032
	(0.102)	(0.103)	(0.028)	(0.040)	(0.059)	(0.048)	(0.022)	(0.029)
R^2	0.44	0.41	0.14	0.13	0.32	0.38	0.04	0.05
Observations	1,276	867	1,275	866	2,184	1,500	2,175	1,489
Assignees	77	52	77	52	226	137	225	137
Assignee locations	87	56	87	56	262	164	261	163

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of patents or independent claims as stated, in logs. Samples restricted by annual patents and claims before treatment as follows: Columns (1)–(4) restricted to the 5th quintile of patents; columns (5)–(8) restricted to the 4th–5th quintiles of claims. Samples are furthermore restricted to complex or discrete product industries as stated, and to non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.18: Robustness check: Complex vs. discrete product industries.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pre-treatment patents:	many				few			
Pre-treatment claims:	many	few	many	few	many	few	many	few
UTSA index	0.044	-0.066	0.058	0.013	-0.004	0.068	-0.090	0.049
	(0.081)	(0.046)	(0.060)	(0.052)	(0.076)	(0.081)	(0.089)	(0.088)
Common law index	0.019	-0.208	-0.237**	-0.372**	-0.276	-0.296	-0.190	-0.115
	(0.165)	(0.211)	(0.099)	(0.145)	(0.193)	(0.181)	(0.180)	(0.232)
Company age	-0.070	-0.004	-0.035	0.000	-0.001	-0.006	-0.002	-0.011
	(0.045)	(0.035)	(0.029)	(0.051)	(0.043)	(0.050)	(0.050)	(0.069)
EBITDA (log)	0.007	0.007	0.011	0.012	-0.000	0.011	-0.009	0.010
	(0.015)	(0.012)	(0.019)	(0.020)	(0.015)	(0.015)	(0.020)	(0.019)
Revenue (log)	-0.012	-0.028	-0.049	-0.075*	-0.004	0.031	0.041	0.081***
	(0.025)	(0.039)	(0.032)	(0.039)	(0.038)	(0.030)	(0.043)	(0.029)
R&D exp. (log)	-0.003	0.005	0.016	0.050	0.016	0.007	0.013	-0.017
	(0.023)	(0.026)	(0.019)	(0.031)	(0.025)	(0.022)	(0.030)	(0.020)
R^2	0.20	0.20	0.13	0.12	0.08	0.05	0.08	0.06
Observations	863	793	1,713	1,509	2,476	2,716	1,626	2,000
Assignees	55	49	127	108	289	312	227	262
Assignee locations	59	53	139	120	333	386	253	319

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: length of the shortest independent claim, in logs. Samples restricted by annual patents before treatment as follows: Columns (1) & (2) restricted to the 5th quintile of patents; (3) & (4) restricted to the 4th–5th quintiles of patents; (5) & (6) to the 1st–4th quintiles; and (7) & (8) to the 1st–3rd quintiles. Samples restricted by annual claims before treatment as follows: columns (1), (3), (5) and (7) are restricted to the 3rd–5th quintile of claims; columns (2), (4), (6) and (8) are restricted to the 1st–3rd quintiles. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.19: Robustness check: Independent claim length.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	patents				claims			
Pre-treatment patents (propensity):	many		few		many—		low	
Pre-treatment claim length:	long	short	long	short	long	short	long	short
UTSA index	-0.647****	-0.359**	0.224	0.124	-0.202	-0.142	0.085	0.056
	(0.189)	(0.162)	(0.137)	(0.262)	(0.146)	(0.095)	(0.056)	(0.052)
Common law index	-0.840	-0.439	-0.613	-0.345	0.240	-0.084	-0.013	-0.180
	(0.638)	(0.368)	(0.444)	(1.234)	(0.250)	(0.319)	(0.140)	(0.197)
Company age	-0.044	-0.033	-0.059	-0.082	0.040	0.013	-0.022***	-0.025*
	(0.080)	(0.078)	(0.082)	(0.100)	(0.041)	(0.033)	(0.008)	(0.015)
EBITDA (log)	-0.042	0.031	0.041	0.032	0.021	0.033	0.018	0.011
	(0.090)	(0.103)	(0.034)	(0.049)	(0.036)	(0.026)	(0.012)	(0.013)
Revenue (log)	0.093	0.010	-0.096	-0.150	-0.160	-0.084	-0.057	-0.037
	(0.344)	(0.236)	(0.107)	(0.154)	(0.101)	(0.076)	(0.049)	(0.051)
R&D exp. (log)	0.854****	0.721****	0.772****	0.909****	-0.184***	-0.248****	-0.057****	-0.034**
	(0.154)	(0.078)	(0.068)	(0.104)	(0.063)	(0.061)	(0.017)	(0.014)
R^2	0.35	0.35	0.30	0.31	0.17	0.14	0.05	0.02
Observations	840	925	1,669	1,927	1,186	1,177	2,144	2,460
Assignees	56	54	240	245	142	136	204	212
Assignee locations	60	59	276	287	161	150	255	266

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of patents or patent propensity, as stated, both in logs. Samples restricted by annual patents (or patent propensity) before treatment as follows: Columns (1) & (2) restricted to the 5th quintile of patents; (3) & (4) restricted to the 1st–3rd quintile of patents; (5) & (6) to the 4th–5th quintiles of patent propensity; and (7) & (8) to the 1st–3rd quintiles. Samples restricted by average length of the shortest independent claim before treatment as follows: columns (1), (3), (5) and (7) are restricted to the 3rd–5th quintile of claim length; columns (2), (4), (6) and (8) are restricted to the 1st–3rd quintiles. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.20: Robustness check: Regression of annual number of patents or patent propensity, conditional on pre-treamtment independent claim length.

	(1)	(2)	(3)	(4)
Pre-treatment claim number:	many		few	
Pre-treatment claim length:	long	short	long	short
UTSA index	-0.382****	-0.285****	0.161	-0.010
	(0.082)	(0.085)	(0.145)	(0.091)
Common law index	0.275	0.248	-0.326	-0.064
	(0.351)	(0.249)	(0.250)	(0.447)
Company age	0.057	-0.121	0.011	0.119
	(0.096)	(0.076)	(0.068)	(0.109)
EBITDA (log)	0.047*	0.012	-0.003	0.023
	(0.028)	(0.024)	(0.023)	(0.034)
Revenue (log)	0.007	-0.033	0.016	0.046
	(0.050)	(0.042)	(0.033)	(0.089)
R&D exp. (log)	0.024	0.016	0.085***	0.066**
	(0.034)	(0.023)	(0.027)	(0.033)
R^2	0.07	0.04	0.09	0.08
Observations	1,180	1,382	1,545	1,399
Assignees	129	149	170	140
Assignee locations	143	166	201	161

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of patents or patent propensity, as stated, both in logs. Samples restricted by annual patents (or patent propensity) before treatment as follows: Columns (1) & (2) restricted to the 5th quintile of patents; (3) & (4) restricted to the 1st–3rd quintile of patents; (5) & (6) to the 4th–5th quintiles of patent propensity; and (7) & (8) to the 1st–3rd quintiles. Samples restricted by average length of the shortest independent claim before treatment as follows: columns (1), (3), (5) and (7) are restricted to the 3rd–5th quintile of claim length; columns (2), (4), (6) and (8) are restricted to the 1st–3rd quintiles. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.21: Robustness check: Regression of claims, conditional on pre-treamtment independent claim length.

Fees:	(1) application	(2) issue	(3) maintenance	(4) remedy	(5) (1)+(2)	(6) (5)+(3)	(7) (6)+(4)
Patent fees (log)	-0.115****	-0.111****	-0.092*	-0.097****	-0.113****	-0.061****	-0.065****
	(0.023)	(0.023)	(0.056)	(0.020)	(0.023)	(0.014)	(0.015)
UTSA index	0.096	0.096	0.026	0.099	0.096	0.107	0.106
	(0.115)	(0.115)	(0.103)	(0.116)	(0.115)	(0.115)	(0.115)
Common law index	-0.220	-0.221	-0.586	-0.221	-0.220	-0.226	-0.226
	(0.382)	(0.382)	(0.490)	(0.383)	(0.382)	(0.386)	(0.386)
Company age	-0.002	-0.002	0.000	-0.002	-0.002	-0.004	-0.003
	(0.005)	(0.005)	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)
EBITDA (log)	0.011	0.010	0.015	0.010	0.011	0.009	0.009
	(0.018)	(0.018)	(0.019)	(0.018)	(0.018)	(0.018)	(0.018)
Revenue (log)	-0.062	-0.062	-0.091	-0.063	-0.062	-0.066	-0.065
	(0.074)	(0.074)	(0.086)	(0.074)	(0.074)	(0.073)	(0.073)
R&D exp. (log)	0.803****	0.803****	0.812****	0.803****	0.803****	0.803****	0.803****
	(0.023)	(0.023)	(0.024)	(0.023)	(0.023)	(0.023)	(0.023)
R^2	0.36	0.36	0.36	0.36	0.36	0.36	0.36
Observations	14,922	14,922	12,400	14,922	14,922	14,922	14,922
Assignees	47	47	47	47	47	47	47
Assignee locations	2,229	2,229	2,011	2,229	2,229	2,229	2,229

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of granted patents, in logs. Samples restricted to non-negative financial variables. Patent fees deflated by the US GPDI deflator. The fee variable includes: column (1): application fees; (2) issue fees; (3) maintenance fees (one-time); (4) remedy fees; (5) application + issue fees; (6) application + issue + maintenance fees; (7) application + issue + maintenance + remedy fees. Robust standard errors clustered by state and assignee. All regressions include assignee-state fixed effects.

Table 1.22: Regression of patents on patent fees.

Fees:	(1) application	(2) issue	(3) maintenance
Claim fees (log)	-0.073***	-0.116****	-0.067***
	(0.023)	(0.034)	(0.020)
UTSA index	0.011	0.010	0.012
	(0.038)	(0.038)	(0.038)
Common law index	0.043	0.045	0.051
	(0.094)	(0.096)	(0.095)
Company age	0.015****	0.014****	0.014****
	(0.002)	(0.001)	(0.001)
EBITDA (log)	-0.003	-0.003	-0.003
	(0.008)	(0.008)	(0.008)
Revenue (log)	0.012	0.013	0.013
	(0.019)	(0.019)	(0.019)
R&D exp. (log)	0.027***	0.027***	0.027***
	(0.008)	(0.008)	(0.008)
R^2	0.02	0.02	0.02
Observations	12,399	12,399	12,399
Assignees	47	47	47
Assignee locations	2,011	2,011	2,011

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of granted patents, in logs. Samples restricted to non-negative financial variables. Claim fees deflated by the US GPDI deflator. The fee variable includes: column (1): fees for each independent claim above the third; (2) fees for each claim above the 20th; (3) fees for existence of multiple-dependent claims. Robust standard errors clustered by state and assignee. All regressions include assignee-state fixed effects.

Table 1.23: Regression of independent claims on claim fees.

Estimation:	Poisson F.E.			Endog UTSA: GMM		
Products:	all	complex	discrete	all	complex	discrete
Effective UTSA (facility wgh)	-0.360	-0.917***	0.330	-2.448**	-3.743***	-1.090
	(0.246)	(0.328)	(0.331)	(1.224)	(1.391)	(1.694)
Employees (log)	0.729****	0.839****	0.460**	0.710****	0.862****	0.423*
	(0.150)	(0.150)	(0.218)	(0.143)	(0.148)	(0.223)
Revenue per employee (log)	0.164	0.401	-0.076	0.132	0.531**	-0.304
	(0.209)	(0.267)	(0.154)	(0.217)	(0.259)	(0.251)
PPE per employee (log)	0.251	0.302	-0.114	0.215	0.123	-0.057
	(0.267)	(0.298)	(0.271)	(0.232)	(0.275)	(0.277)
R&D per employee (log)	0.405****	0.346**	0.365***	0.382****	0.356**	0.363***
	(0.104)	(0.146)	(0.124)	(0.114)	(0.174)	(0.140)
No reported R&D	-0.355	0.343	-0.774	-0.447	0.157	-0.691
	(0.606)	(0.229)	(0.703)	(0.579)	(0.336)	(0.673)
Prior common law (facility wgh)	-1.608**	-1.914	0.120	-3.228***	-4.681****	-0.878
	(0.762)	(1.185)	(0.786)	(1.197)	(1.377)	(1.339)
Observations	7,291	4,097	3,194	7,291	4,097	3,194
Companies	793	469	324			
Hansen p-val	.	.	.	.	.	.
Hausman stat	0.64	0.64	0.64	0.64	0.64	0.64
Hausman p-val	0.425	0.425	0.425	0.425	0.425	0.425
Marg. effect	-0.069	-0.170	0.066	0.066	0.066	0.066
Marg. p-val	0.143	0.005	0.319	0.319	0.319	0.319
Companies				793	469	324
Hansen J stat				2.81	3.83	0.92

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$ Replication of the main result of Png (2017b) as well as splitting the sample by into complex and discrete product industries. Dependent variable: annual number of patents, in levels. Estimation by Poisson regression. Robust standard errors clustered by assignee. All regressions include assignee and year fixed effects.

Table 1.24: Robustness check: Replication of Png (2017b).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	patents				claims			
Pre-treatment semantic index:	low	high	low	high	low	high	low	high
Additional restriction:	many patents		-	-	many claims		-	-
UTSA index	-0.420**	-0.061	-0.226*	0.399	-0.152	-0.289***	0.076	-0.273**
	(0.183)	(0.375)	(0.132)	(0.305)	(0.160)	(0.102)	(0.080)	(0.108)
Common law index	-0.289	0.510	-0.294	-1.780	0.278	0.126	0.105	-0.186
	(0.463)	(0.737)	(0.450)	(1.428)	(0.354)	(0.417)	(0.148)	(0.191)
Company age	-0.025	-0.117**	-0.011	-0.189	0.086	-0.052	0.005	0.107
	(0.075)	(0.059)	(0.067)	(0.121)	(0.112)	(0.087)	(0.041)	(0.116)
EBITDA (log)	0.001	-0.073	-0.014	0.020	-0.040	0.042	-0.023	-0.010
	(0.064)	(0.083)	(0.046)	(0.062)	(0.029)	(0.033)	(0.030)	(0.027)
Revenue (log)	0.231	-0.379*	0.220	0.010	-0.037	-0.052	0.037	0.063
	(0.218)	(0.216)	(0.202)	(0.126)	(0.070)	(0.065)	(0.045)	(0.061)
R&D exp. (log)	0.767****	1.083****	0.835****	0.660****	0.070*	0.035	0.060**	0.050
	(0.129)	(0.107)	(0.080)	(0.082)	(0.040)	(0.037)	(0.030)	(0.036)
R^2	0.46	0.47	0.44	0.30	0.10	0.05	0.07	0.05
Observations	1,226	434	1,621	1,257	453	1,063	1,620	1,244
Assignees	79	39	129	201	41	140	130	199
Assignee locations	86	43	147	221	43	154	148	219

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of patents or independent claims, as stated, both in logs. Samples restricted by semantic similarity between patents before treatment as follows: Columns (1), (3), (5) & (7) restricted to 1st quintile of the similarity index; (2) & (4) restriected to 3rd–5th quintile; (6) & (8) restricted to 4th–5th quintile. Columns (1) & (2) additionally restricted to 4th–5th quintile of patents; (3) & (4) additionally restricted to 4th–5th quintile of claims. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.25: The effect of semantic patent similarity.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:	patents				claims			
Variance in upstreamness:	low	high	low	high	low	high	low	high
UTSA index	0.166	-0.291**	0.469	-0.452****	0.013	-0.135*	-0.112	-0.281****
	(0.264)	(0.142)	(0.349)	(0.134)	(0.065)	(0.075)	(0.103)	(0.077)
Common law index	0.230	-0.299	-0.089	-0.053	0.354	0.036	0.263	0.390*
	(0.822)	(0.381)	(0.879)	(0.463)	(0.282)	(0.171)	(0.215)	(0.201)
Company age	-0.142	-0.063	-0.093	-0.084*	0.093	0.007	-0.043	-0.034
	(0.114)	(0.048)	(0.111)	(0.048)	(0.063)	(0.032)	(0.082)	(0.051)
EBITDA (log)	0.071	0.007	0.059	-0.028	0.016	0.015	-0.023	0.027*
	(0.048)	(0.065)	(0.054)	(0.086)	(0.013)	(0.014)	(0.034)	(0.015)
Revenue (log)	-0.066	-0.028	-0.142	0.132	0.045*	0.003	0.019	0.027
	(0.128)	(0.208)	(0.113)	(0.264)	(0.027)	(0.020)	(0.062)	(0.027)
R&D exp. (log)	0.593****	0.867****	0.648****	0.788****	-0.013	0.048***	0.003	0.026
	(0.107)	(0.048)	(0.107)	(0.075)	(0.024)	(0.015)	(0.035)	(0.029)
R^2	0.22	0.46	0.27	0.45	0.04	0.05	0.04	0.08
Observations	1,212	2,231	758	1,823	3,564	4,284	1,391	1,940
Assignees	126	138	70	113	555	348	170	151
Assignee locations	127	190	70	141	618	588	181	208

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of patents or independent claims, as stated, both in logs. To identify the level of upstreamness, ENTER EXPLANATION HERE. Samples restricted by the standard deviation of the average level of upstreamness as above or below the median value. Furthermore, columns (1) & (2) are restricted to the 3rd–4th quintiles of pre-treatment patents; (3) & (4) to the 4th–5th quintiles of patents; and (7) & (8) to the 3rd–5th quintiles of claims. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.26: The effect of variance in the level upstreamness.

	(1)	(2)	(3)	(4)
Dependent variable:	patents		claims	
Number of inventors:	low	high	low	high
UTSA index	-0.438****	0.066	-0.357****	-0.088
	(0.098)	(0.182)	(0.100)	(0.091)
Common law index	0.501	-0.961	0.028	0.621***
	(0.317)	(0.611)	(0.229)	(0.222)
Company age	-0.106	-0.059	-0.056	-0.022
	(0.088)	(0.051)	(0.090)	(0.041)
EBITDA (log)	0.017	0.090*	0.020	-0.014
	(0.048)	(0.054)	(0.024)	(0.014)
Revenue (log)	-0.070	-0.191	0.021	0.064*
	(0.112)	(0.184)	(0.034)	(0.038)
R&D exp. (log)	0.684****	0.942****	0.019	0.001
	(0.046)	(0.062)	(0.028)	(0.026)
R^2	0.27	0.48	0.06	0.05
Observations	1,787	1,656	1,756	1,583
Assignees	160	103	186	134
Assignee locations	191	126	232	160

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of patents or independent claims, as stated, both in logs. Samples restricted by average number of inventors over whole sample period; low and high number are relative to the median value of 2 inventors per patent. Samples are furthermore restricted to complex product industries and non-negative financial variables. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.27: The effect of the number of inventors per patent.

Dependent variable:	patents				claims			
Patent applications size:	few narrow	few broad	many narrow	many broad	few narrow	few broad	many narrow	many broad
UTSA index	0.216	0.121	-0.474****	-0.332**	0.137**	-0.406****	-0.067	-0.197**
	(0.139)	(0.130)	(0.139)	(0.131)	(0.058)	(0.082)	(0.077)	(0.084)
R^2	0.06	0.04	0.09	0.11	0.04	0.04	0.12	0.08
Observations	9,876	5,184	3,019	2,477	9,876	5,184	3,019	2,477
Assignees	2,090	1,054	287	229	2,090	1,054	287	229
Assignee locations	2,138	1,061	296	234	2,138	1,061	296	234

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable in logs as stated. Estimated by OLS with robust standard errors clustered by state and assignee. State macro variables are per capita. All variables in logs, apart from fractions and dummies. All regressions include assignee-state and year fixed effects.

Table 1.28: Whole NBER dataset: Reaction conditional on pre-treatment patents and claims.

Dependent variable:	Total number of claims				Dependent claims			
Patent applications size:	few narrow	few broad	many narrow	many broad	few narrow	few broad	many narrow	many broad
UTSA index	-0.033	-0.312****	-0.102	-0.180***	-0.055	-0.311****	-0.137	-0.174***
	(0.063)	(0.055)	(0.090)	(0.057)	(0.077)	(0.059)	(0.098)	(0.066)
R^2	0.07	0.08	0.19	0.16	0.06	0.07	0.19	0.16
Observations	10,002	5,290	3,027	2,483	9,631	5,144	3,002	2,466
Assignees	2,116	1,078	287	229	2,045	1,057	285	228
Assignee locations	2,165	1,085	296	234	2,091	1,064	294	233

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable in logs as stated. Estimated by OLS with robust standard errors clustered by state and assignee. State macro variables are per capita. All variables in logs, apart from fractions and dummies. All regressions include assignee-state and year fixed effects.

Table 1.29: Whole NBER dataset: independent vs dependent claims..

Dependent variable:	patents				claims			
Additional restriction:			narrow patents				few patents	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
"Entity size" :	large	small	large	small	large	small	large	small
UTSA index	-0.299****	-0.411***	-0.519****	-0.718**	-0.389****	-0.294**	-0.515****	-0.536***
	(0.089)	(0.144)	(0.114)	(0.283)	(0.074)	(0.132)	(0.083)	(0.176)
R^2	0.11	0.13	0.11	0.15	0.11	0.10	0.10	0.11
Observations	6,928	2,711	3,894	1,537	7,051	3,352	4,750	2,513
Assignees	722	326	420	186	1,072	695	874	619
Assignees locations	746	327	428	186	1,099	696	885	620

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable in logs as stated. Assignees are classified by eligibility for the USPTO's small entitiy fee discount. Regressions of patents are restricted to the 4th and 5th quintile group of patents in all four regressions and additionally to the 1st-3rd quintile group of claims in (3) and (4); regressions of claims are restricted to the 4th and 5th quintile group of claims in all four regressions and in (7) and (8) additionally restricted to the 1st-4th quintile group of patents.

Table 1.30: Heterogeneous impact depending on firm size.

Restriction:	Assignee active in multiple states?					
Dependent variable:	patents		claims		patents	claims
	(1)	(2)	(3)	(4)	(5)	(6)
	no	yes	no	yes		
UTSA index	-0.442****	-0.381****	-0.346****	-0.300****	-0.339****	-0.190****
	(0.075)	(0.107)	(0.070)	(0.078)	(0.070)	(0.033)
UTSA × multi-assignment					0.594****	0.028
					(0.127)	(0.043)
R^2	0.07	0.10	0.06	0.11	0.08	0.07
Observations	13,247	7,823	15,298	6,536	25,655	32,134
Assignees	3,184	1,070	3,864	1,114	5,066	6,603
Assignees locations	3,184	1,195	3,864	1,237	5,227	6,844

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable in logs as stated. Regressions of patents are restricted to 4th-5th ((1)-(2)) and 2nd-5th (5) quintile group; regressions of claims are restricted to 4th-5th ((3)-(4)) and 2nd-5th (6) quintile group.

Table 1.31: Heterogneous impact depending on activity in multiple states and shared patent assignment.

Restriction:	Inventors' residential address				Assignee changes technological area?			
Dependent variable:	patents		claims		patents		claims	
	(1) only US	(2) non-US	(3) only US	(4) non-US	(5) no	(6) yes	(7) no	(8) yes
UTSA index	-0.289****	-0.215	-0.234****	-0.087	-0.486****	-0.162	-0.354****	-0.282***
	(0.083)	(0.187)	(0.041)	(0.070)	(0.141)	(0.146)	(0.079)	(0.087)
R^2	0.06	0.16	0.08	0.21	0.18	0.13	0.12	0.10
Observations	11,504	2,716	17,776	2,720	3,376	2,327	4,468	3,932
Assignees	1,635	239	3,144	262	303	226	667	662
Assignees locations	1,674	244	3,213	269	308	231	675	664

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable in logs as stated. Starting with a single non-US residential address of an inventor in any patent, firms are classified into the "non-US inventor" category. Change of main area of technological activity is defined as having a different most common (i.e., mode of) 4-digit IPC class before and after treatment. Regressions of patents are restricted to 2nd-5th ((1)-(2)) and 5th ((5)-(6)) quintile group; regressions of claims are restricted to 2nd-5th ((3)-(4)) and 4th-5th ((7)-(8)) quintile group.

Table 1.32: Heterogeneous impact depending on foreign inventors and changes to the main area of technological activity.

Claims:	all	only products	only processes
UTSA index	-0.115***	-0.140****	-0.324****
	(0.041)	(0.036)	(0.056)
R^2	0.07	0.04	0.04
Observations	30,235	29,310	19,402
Assignees	6,287	6,093	4,161
Assignee locations	6,542	6,335	4,294

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: annual number of claims as stated, in logs. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.33: Regression of total, product and process claims

	USPTO/JPO patent ratio	divisional application: "resulted in"	divisional application: "results from"
UTSA index	0.052***	-0.058*	-0.054*
	(0.019)	(0.034)	(0.029)
R^2	0.11	0.10	0.13
Observations	11,240	30,235	30,235
Assignees	2,377	6,287	6,287
Assignee locations	2,474	6,542	6,542

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$; **** $p < 0.001$

Dependent variable: (1) number of US patents divided by number of JPO patents in the same OECD patent family; (2) log of number of annual patents that resulted in a divisional application; (3) log of number of annual patents that stem from a divisional application. Regressions on the full sample of complex product industries via complex1 dummy. Robust standard errors clustered by state and assignee. All regressions include assignee-state and year fixed effects.

Table 1.34: Robustness check: alternative measures of disclosure per patent

B Dataset construction

This section contains further details on the different steps undertaken in construction of the datasets.

B.1 Application years

The NBER patent data project contains US patents granted between 1976 and 2006. The values in the corresponding "application year" variable reach as far back as 1901. This points towards some sort of error in the data. For example, US patents 5180907, 5312928 and have a recorded application date of 1901. In fact, these patents were applied for in 1988 and 1991, respectively. A similar problem with patent grant dates has previously been identified by Kogan et al. (2017). I am following their procedure to correct some of the wrongly coded application years by relying on the sequential nature of patent application numbers.

For this, I downloaded the application data file from the USPTO's patentsview.org website, sort the data with respect to application numbers, and for every patent application compute the median date of a 40-patent window around the application (where the date is given in year-month-day format). For every application that has a difference between recorded and median application date of more than 180 days I replace the recorded by the median application date. For those patents I then replace the recorded application year in the NBER patent dataset.

It has to be noted that this correction might potentially introduce new errors when there is a mistake in the application number included in the USPTO's bulk application data file. For example, when sorting those data by application number, the first observation reads 133949 but should read 10133949, and the last observation reads 8296623 but should read 29296623. Similar mistakes throughout the dataset are basically undetectable. As a robustness check, I therefore also estimate the main regressions with using the original application year variable and get almost unchanged coefficient estimates.

B.2 Independent claims

Distinguishing independent from dependent claims is of major importance for the present study. Independent claims are valid on their own and therefore must be satisfy the patentability requirement independent of what else is included in the patent. Dependent claims are valid only in reference to one or several independent claims and allow independent claims to be as broad as possible while still including particular specifications of the subject matter covered in the independent claim in the patent. This can make it easier to enforce the patent against infringers and to avoid patent invalidation lawsuits to invalidate the full scope of the patent. Dependent claims are, however, by definition subsets of independent claims and therefore should not be treated as components of invention of the type considered in this study.

The USPTO provides, as part of its "Research Datasets" series, a dataset containing the parsed fulltext of all US patents granted between 1976 and 2014. In addition, this dataset comes with a dummy variable provided by the USPTO's Office of the Chief Economist that is supposed to mark independent claims. It is set equal to 0 if the fulltext of one claim contains a reference to another claim, and set equal to 1 otherwise.

These data have three limitations. First, the fulltext contains quite a large number of misspellings, leading to a large number of very infrequent words that are excluded from text similarity analysis (see section B.4 below). Second, the algorithm used to detect claims in the patent fulltext occasionally bunched more than one claim together. This most frequently happened with claims containing a lot of numbers, such as those covering chemical formulas. The reverse of this issue also applies, leading to many observations not containing a full claim, but only a part of it that begins with a number. This mainly applies to claims containing enumerations. The only way to address these two limitations would be to perform a more careful optical character recognition (OCR) procedure on the original patent fulltext, and accordingly is beyond the scope of this study. Chemical patents are excluded from the greatest part of the analysis anyway since they are a prototype example of "discrete" technology that can be fully captured in a single claim.

Third, the dummy itself contains errors as well. To address this issue, I perform the following steps: 1) I create a list of the most important words occuring in claim text, based on the words suggested in Milanez et al. (2017), and complemented by the USPTO's list of stopwords used in its patent search engine. This list comprises the following words:

apparatus, accordance, according, another, capable, characteriz/se(d), claim, combination, combining, composition, compound, compris(ing), consist(ing), contain(ing), corresponding, device, embodient, expressing, generally, having, includ(ing), invention, means, method, preferably, preferred, process, provided, respectively, suitable, system, where(in), which, whose.

For each of these words I then search, using regular expressions language, for instances of added space in the middle of the word, missing space between a noun and its preceding article, missing single letters, single pairs of letters that switched positions, double instances of single letters, and common misspellings as provided by http://tools.seochat.com/tools/online-keyword-typo-generator/, provided that the instance of the word is enclosed by blank spaces.

I then change the claim fulltext to lower case and identify the observations including the word "claim(s)" in the fulltext, plus a variety of common misspellings or OCR erros identified by manual inspection of several thousand claim text variables covering the whole length of the sample period.[40] I furthermore identify observations containing any of the the stopwords mentioned in Milanez et al. (2017, table A.1):[41]

as set forth in claim, according to claim, as claimed, as defined in claim, as in claim, defined in claim, produced by, prepared by.

Finally, I define as independent claims those claims which neither contain the word "claim" or its variations nor one of the stopwords listed above, as well as all claims for which the "claim_no" variable provided by the USPTO is equal to 1. (I have not identified a claim miscoded as first claim that was indeed a dependent claim.) Furthermore, I drop from the analysis all claims that are not numbered as first claim and have a length of 25 or less characters[42] as well as claims of patents for which not a single claim has been assigned a claim number of 1 (these are relatively rare).

B.3 Product and process claims

A great share of independent claims can be classified by the first noun that they contain: claims referring to an "apparatus", a "device", a "system", a "composition" or a "compound" are in most cases protecting a product innovation, while claims referring to a "method" or a "process" are protecting a process innovation. Unfortunately, these are but rules of thumb. A claim protecting an "apparatus produced by the method" is protecting the process of production, not the apparatus itself,[43] while a claim covering "to implement a process for [...], an apparatus..." protects and discloses the subsequently described apparatus, even though its first noun is "process".

To address these issues, I refer to manual inspection of several thousand claims including more than one of the above keywords, as well as those containing none of them. I then set up a collection of finely defined rules that take care of the greatest share of ambiguous classifications. These rules range from rather general cases like "An improved method" to very specific cases like "A slip clutch device control method". Common examples of product claims that do not contain any of the product stopwords above are "A circuit arrangement for carrying out a method" and "kit for carrying out a method". The complete list is available on request.

A tough challenge is the classification of software claims. There are reasons to classify either as product or process. The finished computer code seems to be more like an actual product, while the idea behind it is more like describing a process. Claims reading like "a computer-readable storage medium containing a program" are rather work-around claims trying to avoid potential non-patentability, or concerns in that direction by the responsible examiner, of software code alone. While classifying them generally as process claims, I additionally create a variable capturing such "computer-related" claims and exclude them as a robustness check.

B.4 Text similarity between claims and patents

I use the dataset with misspellings corrected and independent claims identified as described above in B.2. For all independent claims, I convert the text to lower case, drop all occurrences of Arabic numbers, drop all stopwords as contained in the above list, and reduce the remaining words to their word stem using the Porter stemmer. I then compute word frequencies over the period from 1976–2000. I drop all words with a length of less than three (almost exclusively spelling mistakes) or more than 25 (exclusively fragments of chemical formulas). I then drop all words that appear in fewer than 10 claims in total (almost exclusively spelling mistakes again), as well as the 30 most common terms:

first, second, signal, portion, least, end, surfac, plural, member, post, between, data, control, form, connect, receiv, select, about, extend, oper, compisinga, output, coupl, layer, materi, through, support, side, circuit, input, open.

These words are so common that they provide almost no information on how similar two claims actually are, but extend time and computational resources required for the analysis. Finally, I drop duplicate occurrences of the same word within a single claim.

I then create two separate datasets: first, I create pairwise comparisons for all the claims contained in any one patent, and average the similarity values at the patent level. Second, I concatenate the text of all independent claims of any one patent (and again drop duplicate occurrences of words) and create pairwise comparisons for all the patents of a single

[40] These mistakes that are not included in the list of common misspellings mentioned above include *ofclaim, toclaim, inclaim, withclaim, in?claim, nclaim, clalm, innclaim.*

[41] I do not include the arabic numbers suggested by Milanez et al. as conditions since there are too many claims with issues here: OCR has produced illegible numbers or simply left them out, claims are referred to in roman numerals, etc.

[42] In the rare event that these are indeed independent claims, they are most likely containing chemical compounds, but the far greatest share of them are claim fragments incorrectly split by the USPTO's original algorithm.

[43] To be precise, what is legally protected is the resulting apparatus *when* it is produced by the described method. If it is produced using another way, or the process is used to produce something else (given that is possible), these cases do not constitute patent infringement. However, here I am interested in what is *disclosed*, and that is most certainly a process, not a product.

assignee within a three-year time window. In both datasets, I then compute for all those paris both the Jaccard and the Sørensen-Dice similarity, both unweighted as well as weighted by inverse document frequency. I then compute a variety of statistics at the assignee level, including mean, median, second and fourth quartile, standard deviation, skewness, and kurtosis, in order to capture the distribution of pairwise similarity across the various patents that are filed by an assignee in any one year.

Chapter 2. Fee Shifting, Firm Size, and the Incentive to Patent

1 Introduction

The traditional motivation of the patent system is to give inventors the opportunity to recover development costs by temporarily excluding competitors from the market for their invention. Not all inventors plan to commercialise their inventions themselves. This can well be the result of optimal division of labour, where specialised R&D labs sell patented innovations to manufacturing firms. In these cases, though, the buyer of the technology subsequently acts as if they were the inventor. A different case is when the buyer is a non-manufacturing firm, too. Such non-practising entities licensing out intellectual property to interested manufacturers do not have to be a problem per se. However, not all of those firms are actually interested in selling licensing agreements. A business model has developed that has those firms wait until other firms have infringed their patents to then sue them in court. As such court suits are costly for each party even in case they win, there is scope for settlement agreements that are mutually preferred by either side to the outlook of fighting the trial to the end. In order to create such a situation of holding up an alleged infringer, the patent that is claimed infringed does not even have to be a very strong one.

Growing attention is being paid to the activities of such "patent trolls" in the United States (Risch, 2011; Love, 2012; Vickery, 2015; Cohen et al., 2016). Their activities have become so widespread in recent years that policy proposals have been put forward that aim to make their business model less profitable. One feature that most of those proposals have in common is the idea of introducing legal fee shifting into the US civil litigation system, at least for the category of patent infringement suits. Love et al. (2016) argue that fee shifting is indeed the main reason why patent trolls seem to be less of a problem than in the US.[1]

This chapter is concerned with the question how fee-shifting would affect firms that are using their patents for meritorious reasons: to protect their innovations from imitation/misappropriation. After all, this is what the patent system was designed for. While there are authors arguing that patents are not playing an important role anymore (Boldrin & Levine, 2013), and there are certainly differences across industries, firms continue to make heavy use of the patent system (Hall, 2004), and the extent of patent litigation between manufacturing firms (Bessen & Meurer, 2013) indicates that they do attach value to their patents. There is evidence that small firms are per se dis-advataged in using the

[1]This does not mean they do not exist at all, see e.g. Ann (2009).

patent system, and that a great share of this disadvantage comes from the high cost of enforcement (Kingston, 2000, 2004).[2] If a firm risks having to pay not only their own legal costs, but also those of their opponent, this might make patenting prohibitively costly for some smaller firms.

How does this compare to the fact that many European jurisdictions do have fee-shifting regulations in force? (Cremers et al., 2017) First, it can be observed that the patenting propensity in Europe is lower than in the US; whether this is for reasons of more costly enforcement or rather for a more stringent examination process. Second, fee-shifting often is not complete, but rather only a share of the costs of the winning party are shifted; which still does increase the incentive to spend, but maybe less so than generally expected. Furthermore, patent examination in European patent offices is on average more stringent than in the US (this is a deliberate policy choice), so an average European patent may have greater chances to stand up in court, and accordingly the incentive of an infringer to ?give it a try? may be lower

Finally, civil procedure in most European countries differs from that in the US in more respects than the presence of fee-shifting. One difference that is particularly important in the context of fee-shifting is the general importance of parties? own efforts in determining the outcome of court trial: inquisitorial system in Europe: judge plays a far greater role in selecting and interpreting evidence; adversarial system in the US: presenting of evidence is entirely up to the parties, while the judge is mostly a neutral observer (Parisi, 2002). This is why the US civil procedure attaches great importance to the process of "legal discovery" which gives each party to a lawsuit the right to request and obtain information from their adversary, which itself may be very costly for the party that needs to provide the information (Cooter & Rubinfeld, 1994; Farmer & Pecorino, 2013).

All things considered, the impact of fee-shifting on US patent litigation incentives may not be as obvious as it is so-far hypothesised. This paper aims to generate some insight by constructing a model that entails the features mentioned above. A firm has achieved an innovation and can decide whether to protect it with a patent. Patent protection allows suing an imitator in court, with court judgement being a function of both parties' expenditures as well as a measure of "patent quality", or strength. The main finding is that in a litigation system with lower importance of spending and high average patent quality (such as many European jurisdictions), the extent of fee-shifting needs to be large to have any effect on litigation and patenting incentives. In a regime with high importance of spending and lower patent quality, a much lower extent of fee-shifting would be sufficient to generate the same mitigating effect on patent litigation.

[2]As anecdotal evidence, take the lawsuit of Nestlé vs. Dualit, a small British manufacturer of kitchen appliances who according to Nestlé was infringing their patents on coffee capsules. The MD of Dualit is quoted as saying "They had a whole football team or even rugby team of lawyers; they tended to outnumber us five-to-one." (Lucas, 2013) He also claimed the lawsuit did cost Dualit about £1m, compared to yearly revenues of about £15m, and around £70bn for Nestlé.

2 The Model

2.1 Introduction

The model developed in the following is loosely based on the one by Bessen (2005), but features an endogenous court outcome and thereby an endogenous advantage for larger patentees. An innovative firm (N, "she") has developed a completely new product. Absent imitation, N will become the monopolist on the newly established market. Before market introduction, N decides on patenting the product (or rather the underlying unique and novel technology). Patenting has two effects: it enables N to sue an imitator in court, allowing her to defend her monopoly position in case of imitation; at the same time, it requires disclosure of information about the innovative technology, facilitating imitation (by either lowering the cost of imitation or increasing its probability of success). A single imitating firm (M, "he") decides on developing and selling an exact imitation of N's technology. If imitation happens and the innovation was patented, N can sue M for patent infringement. If she wins this lawsuit, or if M did not (successfully) imitate, N becomes the monopolist on the product market. If she loses the lawsuit, or decides against suing, or if the innovation was not patented, N and M proceed to the product market as duopolists competing in quantities.

A core ingredient of the model is that firms can differ in size. The simplest and most widespread way of generating endogenous differences in firm size is to assume differences in production cost when firms compete in quantities (see i.a. Rosen, 1991; Poyago-Theotoky, 1996; Yin & Zuscovitch, 1998).[3] I focus on the simple case comparing monopoly quantity to that of a quantity-setting duopoly. This stems from the assumption that in case the patent is upheld in court, patent protection is sufficiently broad as to prevent any profit loss from imperfect imitations. Besides the simplicity in modelling this has the advantage of firm size and expected profits not being a function of the imitation game's parameters, making it easier to work out the effect of size on patenting.[4] For the same reason I disregard the possibility that the innovator engages in entry-deterring behaviour other than potentially applying for a patent, which in fact is a form of entry deterrence.

Lastly, I also disregard issues of timing. Imitation both happens and is detected instantaneously, and the same is true about the decision of the lawsuit. This allows us to ignore both market monitoring effort (as is the focus in Langinier, 2005) as well as incentives of litigants to speed up or delay the conclusion of the lawsuit, which in reality might drag on for years. These issues would influence the model's outcome in predictable

[3] The precise reason why different firm sizes come about is not important to the results of the paper. Differences in production costs are the textbook way of implementation. Alternatively, it could be that firms enter the industry sequentially, and/or that capacity in the present period is a function of last period's capacity when capacity investment per period is constrained.

[4] Firm size would be endogenous to the model when, e.g., the imitator can choose to provide an imperfect imitation, leading to the two firms providing differentiated products. Each firm's size would then depend on the chosen imitation quality. A previous version of this paper studied the possibility to circumvent this problem by additionally assuming firms to be sufficiently capacity constrained. This approach trades off greater generality in one dimension (allowing imperfect imitation) against lower generality in another dimension (restricting at least one firm's capacity to values below its best-response value). A similar case could arise when the innovation consists of a reduction in production costs, the typical way of modelling a process innovation.

ways. Patent infringement going undetected reduces the innovator's expected profit from patenting her innovation; patent litigation dragging on undecided for a longer period similarly reduces either the innovator's or the imitator's expected profit, depending on whether the imitator is allowed to continue marketing the imitation during the trial.

Fogire 2.1 shows the structure of the game. The focus of interest will be on how a firm's optimal decision at the game's initial node depends on its own size. The cost of innovating will be assumed sunk when the game starts.[5] In the first step of the game the innovator can decide between patenting her innovation or not (thus "keeping it secret"). It is assumed that the patent application is costless and successful.

In the next step, the potential imitator can decide to imitate N's innovation. It is assumed that, alternatively, this imitation activity is costly or that it is risky and succeeds only with a certain probability. Importantly, the probability with which the imitation attempt is successful (or the expenditure necessary for imitation) depends on N's choice at stage 1: probability (cost) will generally be higher (lower) when the invention is protected by a patent because M will have access to the the publicly accessible patent specification which requires the patentee to disclose a considerable amount of technical information about the patented invention. This "disclosure effect" of patents is analysed in detail in Zaby (2010).

If M decides against imitation, or if his imitation attempt fails, N will sell the monopoly quantity, and M receives zero profit (outcomes f) and h)). If M imitates, and if N has not patented her product at stage 1, the firms compete in quantities (outcome g)). If N has patented, though, at stage 3 she can react to M's successful imitation by either accepting it (outcome e)) or by going to court and suing M for patent infringement. As an immediate reaction to filing suit, M can decide to concede to N's demand and exit the market, leaving the market to N again (outcome d)). If M decides to defend her case, the suit moves towards trial. Before it reaches trial, though, N can choose to propose a settlement contract to M, which involves licensing the innovation against payment of a fee. M can accept or reject this offer; if he accepts, the outcome will be a duopoly again, but a part of M's profit is shifted to N (outcome a)). If M rejects the settlement offer, the case will be decided by court, which either finds patent infringement (resulting in monopoly, outcome b)) or no patent infringement (duopoly, outcome c)). Court proceedings are modelled as a Tullock contest following Plott (1987), the outcome of which is influenced by the litigation expenditures of the two firms. The firm investing more effort has a greater chance to win the case. The model allows for a continuous fraction of the winner's litigation expenses be shifted to the loser. Other than that, the court system can be used without cost.

A larger firm will have an advantage in court, compared to a smaller firm in the same role, because its stakes in litigation (i. e., the difference in profits between winning and losing the court suit) are higher. Thus, when maximising its payoff from litigation, as long as the difference in sizes is large enough, the larger firm will choose a higher amount of litigation expenditure than the smaller firm would in the same situation.We assume that firms are not budget-constrained concerning their litigation expenditures. This can be due to, e.g., profits accruing from activity in other markets, or a budget set aside for legal

[5]The game could be extended to study the incentive to invest in developing the innovation, which here is a simple function of expected profit from the patenting game.

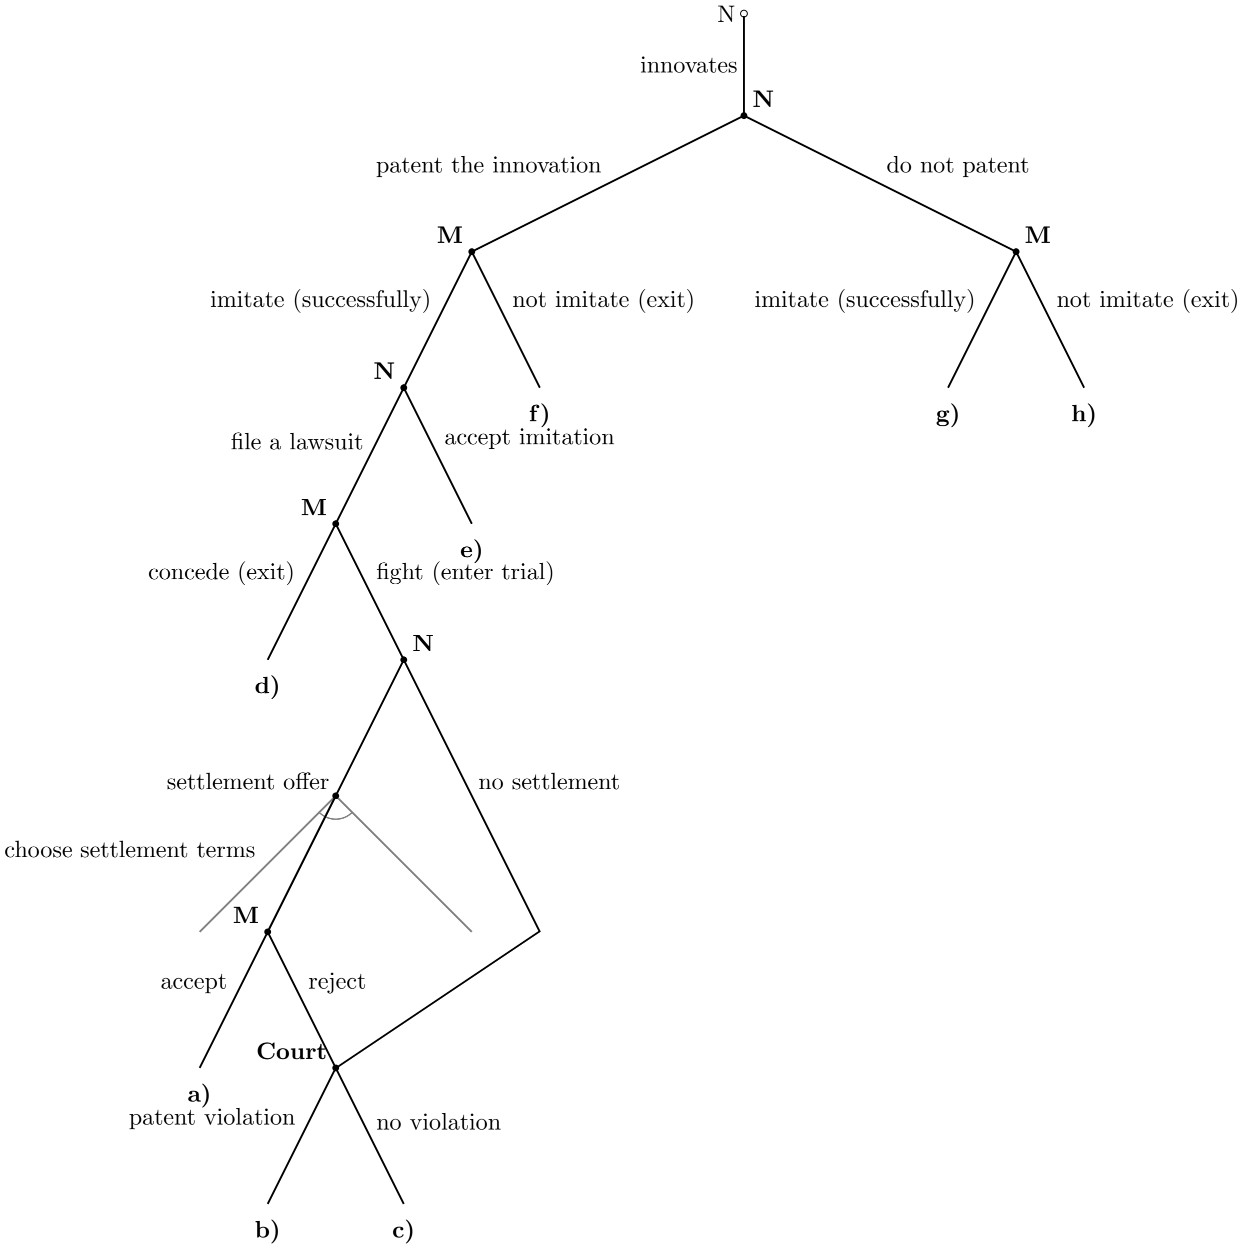

Figure 2.1: Extensive Form of the Game

activity from market income in previous periods or from the initial capital endowment of the firm.[6]

2.2 Market Outcomes

The final stage of the game consists of one of two possible situations: firm N supplying the monopoly quantity, or firms N and M producing their mutual duopoly-best response quantities. The precise values are not crucial to the result of the subsequent patenting-imitation-litigation game. The following profit levels are obtained from standard textbook calculations (Tirole, 1988) and can be thought of as the easiest of several ways of endogenizing firm size.

Assume demand for N's product (technology) given by the linear inverse demand function $P(Q) = a - b\,Q$, where $Q = q_N + q_M$ is the sum of firms' individual production quantity, and $a, b > 0$. Firms face individual linear production cost functions $C_i(q_i) = c_i q_i$. There is no restriction on which firm is more efficient, and costs of either firm can be chosen freely from within the below interval. In the end, it is only of interest how a change in the innovator's cost, given the imitator's cost, affects her patenting decision.

Without imitation or after winning court litigation, N is the only firm on the market. As a monopolist, she chooses q_N such as to maximise her profit $\pi_N^m = (P(q_N) - c_N)q_N^m$. Standard profit maximization yields a quantity of $q_N^m(c_N) = (a - c_N)/(2b)$ and a profit of

$$\pi_N^m(c_N) = \frac{(a - c_N)^2}{4b}. \tag{2.1}$$

With imitation, and after either foregoing or losing court litigation in case of patent protection, N and M act as standard Cournot duopolists, choosing their respective quantity as $q_i^d(c_i, c_j) = (a - 2c_i + c_j)/(3b)$, with $i \neq j$, and obtaining profits of

$$\pi_i^d(c_i, c_j) = \frac{(a - 2c_i + c_j)^2}{9b}. \tag{2.2}$$

Firm i will profitably supply a non-zero quantity as long as $c_i < (a + c_j)/2$.

Throughout the subsequent analysis of the game's various stages, general profit levels are assumed to be simply $\pi_N^m > \pi_N^d > 0$ and $\pi_M^d > 0$.[7] In order to more easily capture asymmetric size, we can write $\pi_M^d = \theta \pi_N^d$, where $\theta \in \mathbb{R}^+$ defines M's size relative to that of N.

[6] Furthermore, in a setting of vertical differentiation with capacity-constrained firms, without fee-shifting, firms' profit-maximizing expenditures are smaller than their initial profit levels for "empirically relevant" parameter levels (especially $r < 1$). Budget constraints would become crucial in a model studying process innovation, which is typically modelled as a reduction in production costs. This way it is possible for the initially small firm to become the larger firm in the market after reducing its costs. Accordingly, its stakes in litigation would be greater than those of the initially large firm, and so would be its willingness to spend in litigation. In that case, budget (or rather, liquidity) constraints would become crucial in order to incorporate a "disadvantage" of some sort of the small firm. Given our modelling choice of litigation as a Tullock contest, the literature on contests with budget constraints would apply, see i.a. Che & Gale (1997) and Yamazaki (2008) for an analysis and the difficulties associated with this approach.

[7] Is it important to assume the efficiency effect, i.e., $\pi_N^m > 2\pi_N^d$? What about the relationship between π_N^m and π_M^d?

3 The Patenting-Imitation-Litigation Game

We now proceed to the analysis of the stages prior to market competition. Firm size will play a crucial role because—*if* N has patented her innovation *and* M has successfully imitated it—N can file a patent violation suit against M. In court, the larger firm has the advantage that it has a greater incentive to win the case and accordingly spend on legal aid, thereby positively influencing the probability that the court comes to a decision in its favour.

3.1 Stage 5: The Litigation Subgame

Proceeding backwards we begin formal analysis at the very last stage (before firms compete in quantities) where both firms decide on their optimal litigation expenditure. This situation arises when N has patented, M has successfully developed an imitation, N has decided to sue M for this, and firms did not agree to settle litigation. Firms decide on expenditure to influence the outcome of patent litigation, which is modelled as a Tullock contest (Tullock, 1975, 1980). The innovator's probability of winning is given by:

$$P(x_N, x_M) = r \, \frac{x_N}{x_N + x_M} + (1 - r)z, \tag{2.3}$$

(first suggested by Plott, 1987) where x_N and x_M are the legal expenditures of N and M, respectively. $r \in [0, 1]$ determines the relative productivity of legal expenditure, i. e., the degree of influence of expenditure on the outcome of the lawsuit. Parisi (2002) uses the same specification to analyse legal systems that differ in the extent to which "the court" is in control of the proceedings. The higher is r, the less active is the court's role in evidence collection and evaluation, and the more influential are parties' efforts in doing so. A high level of r therefore more appropriately captures the "adversarial system" of civil litigation found in the United States, while Europe's "inquisitorial system" is associated with a lower value of r.

$z \in [0, 1]$ has different interpretations throughout the literature. Often it is termed "merits of the case" (Farmer & Pecorino, 1999; Hirshleifer & Osborne, 2001); individual beliefs about the probability of winning the case and a measure of available evidence are alternative interpretations. In any case z represents the part of the winning probability that parties cannot influence with their spending. It thereby captures several features of patent litigation in a single parameter. Winning a patent suit becomes easier the broader the patent is, making it less likely that M found a way to "invent around" the patent.[8] Inventing around is often possible when multiple alternative technologies exists that can be used in producing a product (patents are covering technologies, not products per se).[9] Patents in different technologies differ in their general "ease of interpretation" and the strictness of the delineation of the protected technology area. Patents covering

[8] The breadth of a patent to a large extent depends on the patent examiner and the definition of patentable subject matter and is thereby out of full control of the applicant.

[9] Different technologies may imply different levels of production costs, but production cost plays a role only after litigation, and therefore such cost differences can be captured by the imitator's expected profit term.

mechanical engineering are comparably easier to interpret than patents in (previously-)nascent technologies such as software or biotechnology, affecting the amount of chance inherent in patent litigation. Inventions also differ in the extent to which they satisfy the patentability requirements (novelty, non-obviousness, and usefulness), and the closer to the "inventive threshold" set by patent law an invention is, the more likely will it be invalidated in litigation.[10] Finally, some patent applications are simply more carefully drafted and accordingly stronger and more likely to survive a court challenge. (Maybe add some references?) z therefore can be summarized as "patent quality".

The two firms are thus maximising the following objective functions:

$$\max_{x_N} \quad \pi_N^{lit} = P\left(\pi_N^m - (1-s)\, x_N\right) + (1-P)\left(\pi_N^d - x_N - s\; x_M\right) \tag{2.4}$$

$$\max_{x_M} \quad \pi_M^{lit} = P\left(0 - x_M - s\; x_N\right) + (1-P)\left(\pi_M^d - (1-s)\, x_M\right), \tag{2.5}$$

where s is the share of legal costs of the succeeding party that will be shifted to the losing party. In the baseline model without fee-shifting, $s = 0$. The major driving force of firms' decision on litigation spending is the difference between the profit levels a firm gets when winning and when losing the case (the "*stakes of litigation*", Farmer & Pecorino, 1999). If N wins the case (outcome b)) M will not be allowed to enter the market, and accordingly his profit level is zero; if she loses (outcome c)) each firm will earn their duopoly profit.

Interior equilibrium spending in litigation. Interior equilibrium values of litigation spending x_N^*, x_M^* are obtained by joint maximisation of equations (2.4) and (2.5):

$$x_N^* = r\frac{(1 - s(1-(1-r)z))\pi_M^d(\pi_N^M - \pi_N^d)^2}{\left((1 - s(r(1-z)+z)\pi_M^d + (1 - s(1-(1-r)z))(\pi_N^m - \pi_N^d)\right)^2}, \tag{2.6}$$

$$x_M^* = r\frac{(1 - s(r+(1-r)z)(\pi_M^d)^2(\pi_N^m - \pi_N^d)}{\left((1 - s(r(1-z)+z)\pi_M^d + (1 - s(1-(1-r)z))(\pi_N^m - \pi_N^d)\right)^2}. \tag{2.7}$$

Both firms are motivated to expend resources in litigation by the possibility of higher profits when winning the court case. It will be convenient for the remainder of the model to define the following variables as the two parties' "stakes" of litigation:

$$J_N = \pi_N^m - \pi_N^d \qquad \text{and} \qquad J_M = \pi_M^d = \theta\pi_N^d. \tag{2.8}$$

To more easily compare relative firm size, we can write $J_N = qJ_M$, with $q \in \mathbb{R}^+$. Using the explicit expressions for firm profit from section 2.2 yields $q = 5/4$ for firms of equal size, i.e., if firms have symmetric production costs. Then, with $s = 0$, as presently (approximately) the case in US litigation, these expressions reduce to

$$x_N^*|_{s=0} = r\frac{J_M {J_N}^2}{(J_M + J_N)^2} \qquad \text{and} \qquad x_M^*|_{s=0} = r\frac{{J_M}^2 J_N}{(J_M + J_N)^2},$$

which are simple functions of the two firms' "stakes" in litigation.

[10] Note that in the present model, the outcome of finding "patent invalidity" and "no infringement of a valid patent" are equivalent in that they lead to a duopoly market outcome.

Litigation spending for various cases is depicted in figures 2.2-2.4. Generally, spending increases in own firm size. N's spending increases in J_M as long as $J_N > J_M \dfrac{1-s(r+(1-r)z)}{1-s(1-(1-r)z)}$. With $r=1$, $s=0$, or $z=1/2$, the fraction is equal to 1, and therefore spending increases in competitor's size as long as the competitor is smaller, while it decreases when the competitor becomes bigger than the own firm. Spending is accordingly highest for large and symmetric firms. Furthermore, without fee sifting, optimal spending is independent of the quality parameter z and monotonically increasing in the importance of spending r. Introducing a positive share of fee shifting complicates matters.

If currently $s=0$, then introducing a positive share of fee shifting increases spending unless z and r are sufficiently low and N is sufficiently small compared to M.[11] Furthermore, if $r=1$, spending monotonically increases in s, so any non-positive sign must come from the influence of patent quality z. All cases for which s has a negative effect on N's spending require a very low z; $z<1/3$ is required for low levels of s, while with already higher levels of s z can be up to around 0.41. For M's spending, analogous lower bounds for z exist. This points towards fee shifting possessing some of the alleged features of mitigating the litigation incentive for "unmeritorious", or weak, cases. It needs to be kept in mind that these cases in which s attenuates litigation spending also require the innovator size not to be too large compared to the imitator size. For low values of s this size relationship is always less than one; only for higher values of s does it affect innovators that are larger than the respective imitator. It does not keep x_i from increasing considerably with s as long as r and z are high enough.

Equilibrium litigation outcome. Expected litigation outcome with interior solutions is derived by substituting (2.6) and (2.7) into the litigation success function (2.3):

$$P(x_N^*, x_M^*) = r\left(\frac{(1-s(1-(1-r)z))\left(\pi_N^m - \pi_N^d\right)}{(1-s(r+(1-r)z))\,\pi_M^d + (1-s(1-(1-r)z))\left(\pi_N^m - \pi_N^d\right)}\right) + (1-r)z. \tag{2.9}$$

Without fee shifting, the "Tullock part", hereafter referred to as simply T, generates a simple relation of firms' stakes: $P^* = rJ_N/(J_N+J_M) + (1-r)z$. Obviously then, with $0<r<1$, P is a linear combination between the relative stakes and z. Again the direction of the effect of s crucially depends on z.

Figures 2.5 and 2.6 show P^* as a function of firm size and fee-shifting. With $z=1/2$, P^* is independent of s; if $z<1/2$, then P^* decreases in s, with $z>1/2$ it increases. This implies that with $s=0$, P^* is indeed bounded by the values of its two components, but with increasing s, P^* is pushed towards the direction in which z differs from $1/2$. If the Tullock part is equal to $1/2$, then P^* approaches z for $s=1$ for any $0<r<1$. If both parts of P^* are "on the same side" (i.e., both $\lesseqgtr 1/2$), then s close enough to 1 can push P^* beyond the more extreme of the two summands. Such a case is depicted in figure 2.6d. E.g., with symmetric firms ($J_N=5, J_M=4$) and $r=3/4$, then $P^*=0.825$ for $z=0.8$, while it is $P^*=0.228571$ for $z=0.2$. The greater the difference between both firms's stakes, the more advantaged is the firm with the greater stakes: if the innovator's stakes are greater than the imitator's, she benefits from high z, while if the imitator's stakes

[11] Precisely, increasing s from zero decreases x_N if $J_N < J_M(1-2r-3(1-r)z)/(1-(1-r)z)$.

are greater than the innovator's, he benefits from low z. The bottom line is that while s pushes P^* into the direction indicated by z, it does so more strongly for larger firms, and might yield a P^* more extreme than z if firms are sufficiently different in size.[12]

Corner solutions in litigation spending. Once the parties have reached trial, the only decision to be made is on how much to spend on convincing the court of their case. The interior maxima[13] obtained above are only an equilibrium as long as it is not a profitable deviation for either firm to spend zero instead. Unilaterally setting $x_i = 0$ while the competitor keeps spending $x_j = x_j^*$ yields expected profits as

$$\begin{aligned}
\pi_N^{lit}|_{x_N=0} &= (P|_{x_N=0})\pi_N^m + (1 - P|_{x_N=0})\left(\pi_N^d - s\ x_M^*\right) \ , \\
\pi_M^{lit}|_{x_M=0} &= (P|_{x_M=0})\left(-s\ x_N^*\right) + (1 - P|_{x_M=0})\,\pi_M^d \ ,
\end{aligned}$$

where $P|_{x_N=0} = (1-r)z$ and $P|_{x_M=0} = r + (1-r)z$. Considering $\pi_i^{lit^*} \geq \pi_i^{lit}|_{x_i=0}$ we can derive the following inequalities

$$\begin{aligned}
\pi_N^{lit^*} \geq \pi_N^{lit}|_{x_N=0} \quad &\Leftrightarrow \quad r\frac{x_N^*}{x_N^* + x_M^*}\left(J_N + s(x_N^* + x_M^*)\right) \geq x_N^*\left(1 - s(1-r)z\right) \ , \\
\pi_M^{lit^*} \geq \pi_M^{lit}|_{x_M=0} \quad &\Leftrightarrow \quad r\left(J_M + sx_N^*\right) \geq x_M^*\left((1 - s(1-(1-r)z)) - r\frac{x_N^*}{x_N^* + x_M^*}\right)\left(J_M + s(x_N^* + x_M^*)\right) \ .
\end{aligned}$$

Substituting (2.9), (2.6) and (2.4), we obtain

$$\begin{aligned}
\pi_N^{lit^*} \geq \pi_N^{lit}|_{x_N=0} \text{ if } \qquad & x_N^*\left(1 - s(1-r)z\right)J_N^2 \geq 0 \qquad \text{and} \\
\pi_M^{lit^*} \geq \pi_M^{lit}|_{x_M=0} \text{ if } \qquad & x_M^*\left(1 - s(r + (1-r)z)\right)J_M^2 \geq 0 \ .
\end{aligned}$$

These inequalities hold as long as either $s < 1$ or $r < 1$.[14] These two cases correspond to the extremes of the parameter space with either full fee shifting or complete irrelevance of parameter z. As those cases arguably have low empirical relevance we will in the following analysis restrict attention to $s, r \in [0, 1)$.[15]

3.2 Stage 4: The Settlement Offer

More than 90% of filed patent infringement suits in the US settle before trial, therefore it is crucial to incorporate this possibility in the model. If the case moves to trial, each party

[12] It needs to be kept in mind that firms cannot exceed a certain threshold in terms of their relation between production cost and market size. With the explicit functions from section 2.2, $c_i < (a + c_j)/2$.

[13] They are indeed maxima; the second derivatives are $-(2rx_mJ_N)/(x_n + x_m)^3$ and $-(2rx_nJ_M)/(x_N + x_M)^3$.

[14] A weak inequality holds when additionally $z > 0$ if $s = 1$.

[15] If we included those cases in the analysis where firm i optimally spends zero at trial we would have to face the problem of determining firm j's optimal response. In a standard Tullock contest, this would give rise to a situation in which no pure strategy equilibria exist, as the competitor would always find it beneficial to lower their own spending to a point in which a spending of more than zero would again be optimal from the original firm's point of view (Linster, 1993). Farmer & Pecorino (1999) solve this problem by assuming that this party will decide on the lowest expenditure level that will still make the opponent spend zero at trial. In the modified Tullock contest with the fixed component z used here, this case does not occur unless we make the LSF equivalent to that used by Farmer & Pecorino, which is the case precisely with $r, s = 1$.

expects a profit according to equations (2.4) and (2.5), respectively, with expenditure substituted by equilibrium values of (2.6) and (2.7). In the context of bargaining over the settlement terms, this pair of profits (π_N^{lit}, π_M^{lit} is the disagreement point. In case N refrained from bringing suit, each party would earn their respective duopoly profit. The feasibility set therefore contains all allocations of $(\pi_N^d + \lambda, \pi_M^- \lambda)$ which entail some positive transfer of λ from the imitator to the innovator.[16]

Generally, a mutually profitable allocation of duopoly profits exists if $\pi_N^{lit} + \pi_N^{lit} < pi_N^d + pi_M^d$. The left-and side of this inequality is equivalent to $P\left(J_N - J_M\right) + \pi_N^d + \pi_M^d - x_N - x_M$, which implies that settlement is profitable if $P\left(J_N - J_M\right) < x_N + x_M$. A settlement offer becomes more profitable the larger is the imitator.

Any allocation that satisfies the above inequality is a Nash equilibrium. Without further going into detail of the settlement bargaining process, I assume that parties will agree on *some* allocation. In the figures I restrict attention to the case involving an equal split of the bargaining surplus, and to the case in which A makes a take-it-or-leave-it offer involving the transfer of the whole surplus from M to N. A risk-neutral M would be indifferent between accepting this offer and moving to trial, but any arbitrary small deviation from a full transfer would be optimally accepted.

With trial, N can expect profit to be

$$\pi_N^{lit} = \pi_N^d + PJ_N - x_N(1 - Ps) - x_M s(1 - P),$$

whereas with a settlement payment of $\lambda \pi_M^d$ (note the definition of λ as a fraction of M's duopoly profit here) she would get

$$\pi_N^{set} = \pi_N^{lit} + \lambda\left(P(J_M - J_N) + x_N + x_M\right).$$

The precise allocation of surplus is not of crucial importance in the present model, as is rather the decision to file and/or defend a suit. I assume this decision to be independent of whether a settlement will be reached or not, as there will be no incentive for N to propose a settlement not for M to accept such an offer if the opponent would not be willing to move to trial should a settlement not be reached. This is the focus of the next section.

3.3 Stage 3: The Litigation Decision

There are three outcomes that can arise from this model when N patents her innovation: either the firms meet in court, spending their equilibrium amounts; they agree on a

[16] An alternative would be to allow the innovator to suggest a payment to the imitator in return for not entering the market. Besides potential antitrust concerns, settlement of patent litigation more often involves a payment from the alleged infringer to the allegedly infringed party. Giving the patentee the option to suggest sharing of the monopoly rent furthermore would move the focus of the model to one in which the optimal strategy of a small firm could be that of licensing the patent to a larger manufacturing firm, which then might or might not stick to the agreed licensing terms, or alternatively start a trial for patent invalidity. In such a case, the small firm's stakes would be determined by its licensing revenue, which would remove the endogenous disadvantage at trial from the model. To replace it, I would need to introduce some alternative way to capture the empirical fact that small firms are having a harder time in court, for example by introducing "liquidity constraints" on litigation spending, or a litigation cost function that depends on the number of employees.

settlement involving market entry and a transfer from M to N; or at least one of the firms' finds it more profitable to avoid litigation and accept the respective outside option. These three outcomes are the drivers behind the imitation and protection decisions made at preceding steps.

Anticipating the litigation expenditure levels determined above, N decides whether to file a suit against M for patent violation. In case she does, M can then decide whether to enter litigation, in which case the game proceeds to stage 4 and possibly 5, or whether to "concede" to N's demand and abstain from entering the market (outcome d)). If N decides against filing a suit, the game ends at outcome *e)* and the firms earn duopoly profits. Proceeding backwards again, we first analyse M's litigation decision given that N has filed a suit, and then move to N's litigation decision anticipating M's reaction.

M's litigation decision

Faced with a lawsuit filed by N, M compares his expected outcome from entering litigation to that of conceding to N's demand. As shown in the preceding section, entering litigation will entail litigation spending by the two parties given by (x_N^*, x_M^*) in (2.6) and (2.7) above. N's goal in filing the suit is to prevent M from entering the market,[17] so M's alternative to entering litigation is to stay out of the market and earn zero profit. Formally, M decides to enter litigation if he expects to earn a positive net profit from doing so, i.e., if $\pi_M^{lit^*} > 0$, or equivalently

$$(1 - P^*)J_M > P^* s\,(x_N^* + x_M^*) + (1 - s)\,x_M^*. \tag{2.10}$$

Substituting in equilibrium expenditure from the preceding section yields an expression with precisely one root (ϕ) within the permitted parameter ranges. M's expected profit from litigation is positive if $J_N < \phi$, with $\phi = \text{Max}\{0, \phi_1\}$, and ϕ_1 given by

$$\phi_1 = J_M \frac{\Big[\big(r(r+s+(1-r)(1-s)z-1)(1-s(r(1-z)+z))^2(1-s(1+(1-r)z))\big)^{1/2} + (1-r)(1-z)(1-s(r(1-z)+z))(1-s(1+(1-r)z))\Big]}{(1-s(1+(1-r)z))(r+s+z-1-(1-r)rs-(r+2(1-r)^2 s)z-(1-r)^2 sz^2)} \tag{2.11}$$

In general terms, we obtain the result that **conditional on a suit being filed by N, M will go to trial when $J_N \in [0, \phi)$, and avoids litigation by not entering the market otherwise**. The larger is the patent holder, the more likely it is that M will prefer conceding to fighting his case. Two results are immediate: the more trivial one is that litigation being unprofitable requires positive stakes for the innovator (who otherwise would be spending zero at trial); the less obvious one is that with positive stakes for either

[17] Often the goal in patent litigation is also to receive damages payments from the alleged infringer, but in the present static model—unless we explicitly consider punitive damages—a payment of damages amounting to lost profits of the patent owner is equivalent (for N) to "injunctive relief" preventing M from entering the market. For M, being liable to paying damages would unambiguously increase the stakes of litigation if it came in addition to losing market access. If damages in form of lost profits had to be paid *instead* of being prevented market access, patent infringement could become a profitable strategy if M is much larger than N.

party, $\phi > 0$, i.e., litigation is always profitable if $s = 0$ or $r = 1$. Since his expenditure at trial is a function of his and the plaintiff's stakes, without fee shifting there are no cases that the defendant—absent a settlement offer—will not find profitable to actually defend. Similarly, as found above, fee-shifting affects litigation behavior only as far as the value of z differs from that of T, and z playing a role requires $r < 1$.

This result of course also depends on the assumption of absence of fixed costs of litigation in general and of using the court system in particular, and is in line with literature findings comparing cost allocation rules: if each party has to bear their own costs of trial ("*American rule*" of cost allocation) and fixed costs of trial are sufficiently low, (almost) every possible dispute will go to trial (Farmer & Pecorino, 1999; Coughlan & Plott, 1997).

N's litigation decision

M's outside option was walking away from the market entirely. N's outside option, in contrast, depends on M's litigation decision. Anticipating M's decision to go to trial or to concede, N decides between filing a suit and accepting M's market entry. In the latter case, she will earn duopoly profits. If she files a suit and M concedes, she will earn monopoly profits, making her file a suit in this range of parameter values with certainty since $\pi_N^m > \pi_N^d$.[18] If M defends his case, N will file a suit as long as $\pi_N^{lit^*} > \pi_N^d$.

$$P^* J_N > (1 - P^* s)\, x_N^* + (1 - P^*)\, s x_M^* \tag{2.12}$$

Substituting in equilibrium litigation spending now yields a function with the following root, denoted by $\psi = \text{Max}\{0, \psi_1\}$, and ψ_1 given by:

$$\psi_1 = J_M \frac{\Big[\big(r((rs + (1-r)(1-s)z)(1 - s(r(1-z)+z))(1 - s(1+(1-r)z))^2\big)^{1/2} - (1-r)z(1 - s(r(1-z)-z))(1 - s(1-(1-r)z)) \Big]}{(r(1-z)+z)(1 - s(1-(1-r)z))^2} \tag{2.13}$$

Generally, N's reaction to s is qualitatively equivalent to that of M, just that the sign of comparative statics in own and competitor's size as well as in z are reversed: larger own stakes and lower competitor's stakes make trial more likely, as does a favourable (i.e., high) value of z. r works in the same direction as for M, higher importance of spending making litigation less likely. Due to the assumption of sequential decision-making in joining litigation, N enjoys some first-mover advantage, allowing her to prevent M from entering the market by filing a suit even when she wouldn't want it to proceed to trial. This case happens when both firms' litigation constraints are violated, in which case N anticipates M conceding in order to avoid expectedly unprofitable litigation.

N's equilibrium behaviour is therefore characterised as follows:

1. If $\psi < 0$ or not defined for a given combination of parameter values, then N will file a lawsuit for any value of J_N;

[18] A small but positive fixed cost of going to court could change this, but would complicate the model and not change the results apart from making litigation constraints being violated more often, reducing the incidence of low-value trials. Since value here follows directly from firm size, this would make small firms use the court system even less.

(a) if $J_N < \phi$, this lawsuit will go to trial;[19]

(b) if $J_N \geq \phi$, M will concede to N's demand of leaving the market.

2. If $0 < J_N \leq \text{Min}\{\psi, \phi\}$, then N will find it unprofitable to file a lawsuit, and rather accepts M's market entry;

3. if $0 < \psi < J_N < \phi$, then N will file a suite that goes to trial;

4. if $0 < \phi \leq J_N$, then N will file a lawsuit which however will only cause M to concede and leave the market.

Due to her first-mover position, in the last case N's decision does not depend on ψ.

Figures 2.7-2.10 illustrate the litigation decision. The first two sets of diagrams (figures 2.7 and 2.8) plot ϕ and ψ as a function of s, where as in the diagrams before, black curves belong to N (ψ) and orange (gray) curves to M (ϕ). The second two sets are plots of litigation profits, the two core variables of this model.

This manifests a second disadvantage of smaller innovators: not only do they have a lower likelihood of winning court litigation due to their smaller legal expenditures, but they also find litigation to be profitable less often. Most importantly, a large imitator will find it more profitable to defend the smaller is the patent holder, so that a smaller innovator cannot benefit to the same extent from the "costless enforcement" that happens when both firms' litigation constraints are violated.

3.4 Stage 2: The Imitation Decision

Imitation when the innovation is secret

Starting with the simpler case of secrecy, M decides to produce an imitation if the expected profit of doing so exceeds the (certain) cost, or $\rho^s \pi_M^d > \chi^s$, where superscript s denotes the case of secrecy. Alternatively, the model can be solved using only one of the two elements, success probability ρ^s or fixed imitation cost χ^s. With the former, imitation will always be attempted, but the expected profit scaled down, which can influence N's patenting decision at the preceding stage. With the latter, the decision is as simple as that an imitator with a high expected duopoly profit, i.e., a large imitator, will imitate. While very simple and to an extent probably realistic, this brings in yet another disadvantage for small innovators, as M's duopoly profit will be greater the smaller is N. For this reason it can be worthwhile to consider both cases, the success probability and the fixed cost, separately.

A possibility would be to endogenise ρ. If we assume that M can choose imitation effort y^s facing an increasing concave function $\rho^s(y^s)$ and an increasing convex function $\chi(y^s)$, then given the shape of those functions his optimal imitation effort depends only on π_M^d. Accordingly, imitation by a large firm would become more likely, but in a less step-wise manner than with fixed cost χ^s.

[19] I assume that N decides against the "hassle" of filing a lawsuit when the expected profit from doing so is zero.

Imitation when the innovation is patented

As patent protection opens the possibility of litigation, the expected litigation outcome needs to be taken into account. The simplest case is that in which M concedes to N, leaving him with zero expected profit. Accordingly, there is no point in investing in imitation, and we assume that imitation will not happen in this case.[20]

In the cases in which N prefers accepting market entry to fighting out the trial, without the threat of trial M faces a situation analogous to that of secrecy and will imitate if $\rho^p \pi_M^d > \chi^p$. The only difference is that the superscript now reads p for patent. Applying for patent protection requires the applicant to disclose information about the innovation allowing anyone "skilled in the art" to understand the innovation. The fundamental idea behind the patent system, besides providing an incentive to invest in innovation, is that technical knowledge that would otherwise be kept within the boundaries of the firm is disclosed to the public. For fixed exogenous values of either ρ or χ it seems reasonable therefore to assume that $\rho^p > \rho^s$ and $\chi^p < \chi^s$: patenting makes imitation easier, but potentially illegal. In the case of endogenous imitation effort, assuming one of the two differences is enough to obtain $\rho^p > \rho^s$. It is useful to assume that cost is low enough for *any* imitation to occur, as the present case is the one most favourable to the imitator (no litigation and higher success probability). If M decides against imitation here, he will never imitate, a trivial result.

The final case to be considered has both parties move to trial, conditional on imitation. M imitates if $\rho^p \pi_M^{lit} > \chi^p$. π_M^{lit} is the expected profit as given by equation (2.10). If $\chi^p = 0$, the decision is no different from the litigation participation decision, and the only effect of $\rho < 1$ is again on N's patenting decision. If imitation effort is endogenous, assuming that $\chi(y^p)$ is a smooth and continuous function starting at $\chi(0) = 0$ suffices to obtain the same result. Only with $\chi(0) > 0$ or an exogenously fixed χ^p it is that M may decide against imitation. The comparative statics are exactly identical to those at the litigation stage, implying that for some J_M sufficiently large compared to J_N imitation will happen, while no imitation will take place for smaller J_M.

$\pi_M^{set} \geq \pi_M^{lit}$, accordingly imitation can be more profitable if M anticipates a settlement offer by N. N's incentive to choose maximum λ therefore is not only due to maximisation of profit at the litigation stage, but to make imitation less profitable to begin with.

3.5 Stage 1: The Patenting Decision

The final section covers the defining decision of this paper. First I define the alternatives. If N keeps her innovation secret, she can do so at zero additional cost. This can be justified by defining the secrecy option as the mere "absence" of a patent, but also as an existing regime of protection of business information that was costly to set up, but the costs of which are sunk at the point of the protection decision. Depending on whether imitation happens, the game then ends with monopoly or duopoly profits.

If she patents her innovation, she will have to incur costs of f. These costs include both various fees charged by the patent office as well as those charged by patent attorneys. It is possible to argue that f varies depending on firm size, but since this would be a

[20] This is an entirely innocuous assumption.

rather obvious effect of firm size on the patenting propensity, and in order to keep focus on fee shifting as a policy measure, f is assumed to be fixed and equal for all firms.[21]

Just as for M, N's decision is made in anticipation of the litigation stage. To cover the trivial case first, if χ is prohibitively high and M does not imitate in either case, there is no point in applying for a costly patent. Similar reasoning applies if N will not defend her patent, as even with $f = 0$ it increases the occurrence of imitation. The contrasting case is that of M not defending her case at trial. Now, a patent will provide monopoly profits with certainty, while secrecy will do so only when χ^s is prohibitively high. Otherwise, it does so only with a probability $(1 - \rho^s) < 1$, making patenting the optimal choice if $\pi_N^m - f > \rho^s \pi_N^d + (1 - \rho^s)\pi_N^m$, or $\rho^s J_N > f$. Generally, the left-hand side is larger for larger M, while at the same time larger M will be less likely in the situation of not defending their court case.

Finally, when imitation leads to a trial, N's patenting decision $\pi_N^{pat} > \pi_N^{sec}$ is described by

$$\rho^p \pi_N^{lit} + (1 - \rho^p)\, \pi_N^m - f > \rho^s \pi_N^d + (1 - \rho^s)\, \pi_N^m \ , \text{ or}$$

$$J_N \big(\underbrace{\rho^p P^*}_{\text{benefit}} - \underbrace{(\rho^p - \rho^s)}_{\text{costs 1}} \big) > \underbrace{f + \rho^p \big(x_N^* \,(1 - P^* s) + x_M^* s\, (1 - P^*) \big)}_{\text{costs 2: monetary costs}} \ . \tag{2.14}$$

The left-hand side of the inequality consists of the changes in payoffs, while the right-hand side contains the monetary costs of patenting and litigation. The term labelled "benefit" is the payoff received by N in case of imitation and successful ensuing litigation. "Costs 1" is the increase in imitation probability associated with patent disclosure, a non-monetary cost of patenting that could alternatively be written on the right-hand side.

The patenting decision is illustrated in figures 2.11d as well as 2.12 and 2.13; the latter two additionally accounting for the case in which litigation is settled. The (monetary and non-monetary) costs of patenting make the patenting constraint more easily bind than the litigation decision, i.e., there are cases in which N would defend a patent in court but would not find it profitable to apply for one. In the "standard" diagram with J_N on the vertical axis, whenever patenting is costly or the imitation conditions differ compared to secrecy, then the patenting constraint is strictly above the litigation constraint.

This is not true for the patenting constraint with settlement, though when interpreting the diagrams it needs to be kept in mind that in this model, settlement only occurs when the litigation constraints are satisfied. Whenever the curve representing the patenting constraint under settlement is below the litigation constraint, than the latter becomes the effective patenting constraint for those cases. As long as settlement is possible, therefore, patenting becomes considerably more attractive to smaller firms. This of course requires that smaller firms in fact manage to obtain the same settlement terms as larger firms, something that is not captured by this model.

An interesting case is depicted by figures 2.13c and 2.13c. In those two cases, secrecy is much more effective in protecting against imitation, making patenting a very unattractive

[21] The USPTO charges firms that qualify as "small entities" or "micro entities" a lower rate of most fees. In contrast, some sources argue that legal advice is more costly for smaller firms that don't usually employ attorneys in-house and cannot to the same extent make use of economies of scale in the use of legal advice.

choice for firms interested in monopolising their innovation. In the examples drawn, the patenting constraint is so high that it is outside of the range of the vertical axis. A new case arises under settlement, however, where beyond a certain difference $\rho^s \ll \rho^p$ innovators with stakes *below* some threshold prefer patenting. Those small innovators prefer the duopoly arising from "being infringed and settling" to the much more probable monopoly from secrecy. Within the present model assuming production firms, this is the closest depiction of a "patent troll" that is possible. For this incentive to occur the imitator must be above some size threshold which still needs to be determined, as needs to be the degree of influence that fee-shifting has on this behaviour.

4 Conclusion

The model developed in this chapter predicts that the incentive to patent decreases with the extent of fee-shifting, and that the smallest firms are the ones that stop patenting first. The effect of increasing the extent of fee-shifting depends on characteristics of the litigation system, though: the more important is spending and the lower the quality of the litigated patents, the stronger is the negative effect of fee-shifting on the patenting incentive. When introducing fee-shifting to US patent suits, great care therefore needs to be taken in order to not render the patent system entirely unattractive to smaller firms. To achieve the same deterrent effect on patent trolls that is attributed to fee-shifting in Europe, a considerably lower extent of fee-shifting may be well sufficient.

In order to allow more detailed predictions, the model could be calibrated using US patent data. Parameter z could, e.g., be set according to the product patent effectiveness measure from Cohen et al. (2000), the probability (or cost) of imitation could be based on the measure of technological complexity (or rather "complicatedness") developed by Naghavi (2015)[22], while the difference between patenting and secrecy could come from the difference in Cohen et al.'s patent and secrecy effectiveness measures again. The USPTO has recently made available a dataset on patent litigation (Marco et al., 2017), which could be used to obtain industry-specific win rates of patent infringement litigation. Based on these and data on firm size, the parameter r could then be calibrated at the litigation stage (or alternatively the litigation decision stage, which would not even require determining the outcome of litigation). Firm size could be based either on Compustat data matched to the USPTO litigation data, or on industry-level data. A decision with some discretion would be which value to use for the size of the imitator, whether to use the average, median, largest firm, or rather some average of the top 5%.

Feng & Jaravel argue that the best policy tool against patent trolls would be an increase in patent quality. In the model, an increase in patent quality makes litigation more attractive to the patentee, but the model assumes a plaintiff who is practising their patent. To the extent that the business model of patent trolls requires the existence of low-quality patents that are commercially unattractive to their original owners, an increase in the quality of patent examination could reduce the number of such valueless patents in the population of US patents. This would have the advantage of avoiding the detrimental effect of introducing fee-shifting with the share of expenditure that is shifted

[22] This measure was used by Fernandez Donoso (2014) to control for the risk of imitation.

set too high. The precise effect of z on the existence of the "troll-like" behaviour of preferring infringement and settlement when secrecy offers comparably strong protection still needs to be determined.

Another effect that is not currently captured by the model is that of a portfolio of existing patents. Such a portfolio can have a strong mitigating effect on the incentive to file patent litigation between manufacturing firms, as the defendant to a suit could threaten to retaliate by counter-suing the plaintiff for infringement of one of his existing patents. A related problem is that of complementary innovation where the patent that is potentially litigated is one of several building blocks to a complex product.

References

Ann, C. (2009). Patent Trolls – Menace or Myth? In W. P. zu W. und Pyrmont, M. J. Adelman, R. Brauneis, J. Drexl, & R. Nack (Eds.), *Patents and Technological Progress in a Globalized World: Liber Amicorum Joseph Straus* (pp. 355–364).

Bessen, J. (2005). Patents and the diffusion of technical information. *Economics Letters, 86*(1), 121–128.

Bessen, J., & Meurer, M. J. (2013). The Patent Litigation Explosion. *Loyola University Chicago Law Journal, 45*, 401–440.

Boldrin, M., & Levine, D. K. (2013). The Case against Patents. *Journal of Economic Perspectives, 27*(1), 3–22.

Che, Y.-K., & Gale, I. (1997). Rent dissipation when rent seekers are budget constrained. *Public Choice, 92*(1), 109–126.

Cohen, W. M., Nelson, R. R., & Walsh, J. P. (2000). Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not). *NBER Working Paper 7552.*

Cohen, L., Gurun, U. G., & Kominers, S. D. (2016). The growing problem of patent trolling. *Science, 352*(6285), 521–522.

Cooter, R. D., & Rubinfeld, D. L. (1994). An economic model of legal discovery. *The Journal of Legal Studies*, 435–463.

Coughlan, P.J., & Plott, C.R. (1997). An experimental analysis of the structure of legal fees: American vs. English Rule. *CalTech Social Science Working Paper 1025.*

Cremers, K., Ernicke, M., Gaessler, F., Harhoff, D., Helmers, C., McDonagh, L., Schliessler, P., & van Zeebroeck, N. (2017). Patent litigation in Europe. *European Journal of Law and Economics, 44*(1), 1–44.

Farmer, A., & Pecorino, P. (1999). Legal expenditure as a rent-seeking game. *Public Choice, 100*, 271–288.

Farmer, A., & Pecorino, P. (2013). Discovery and Disclosure with Asymmetric Information and Endogenous Expenditure at Trial. *The Journal of Legal Studies, 42*(1), 223–247.

Fernandez Donoso, J. (2014). Do complex inventions need less international patent protection? *Economics Letters, 125*(2), 278–281.

Hall, B.H. (2004). Exploring the Patent Explosion. *The Journal of Technology Transfer, 30*(1?2), 35–48.

Hirshleifer, J., & Osborne, E. (2001). Truth, effort, and the legal battle. *Public Choice, 108,* 169–195.

Kingston, W. (2000). Enforcing Small Firms' Patent Rights. Publications Office of the Commission of the European Communities, Luxembourg. Available at: http://cordis.europa.eu/innovation-policy/studies/im_study3.htm

Kingston, W. (2004). Making patents useful to small firms. *Intellectual Property Quarterly*, 369–378.

Langinier, C. (2005). Using patents to mislead rivals. *Canadian Journal of Economics, 38*(2), 520–545.

Linster, B. G. (1993). Stackelberg rent-seeking. *Public Choice, 77*(2), 307–321.

Love, B.J. (2012). An Empirical Study of Patent Litigation Timing: Could a Patent Term Reduction Decimate Trolls without Harming Innovators. *University of Pennsylvania Law Review, 161*, 1309–1360.

Love, B.J., Helmers, C., Gaessler, F., & Ernicke, M. (2016). Patent Assertion Entities in Europe. In D. D. Sokol (Ed.), Patent Assertion Entities and Competition Policy.

Lucas, L. (2013, April 24). Dualit wins battle over Nespresso patent. *Financial Times.* Retrieved from https://www.ft.com/content/fab6a180-aceb-11e2-b27f-00144feabdc0

Marco, A. C., Tesfayesus, A., & Toole, A. A. (2017). Patent Litigation Data from US District Court Electronic Records (1963-2015). *USPTO Economic Working Paper No. 2017-06.*

Naghavi, A., Spies, J., & Toubal, F. (2015). Intellectual property rights, product complexity and the organization of multinational firms. *Canadian Journal of Economics, 48*(3), 881–902.

Parisi, F. (2002). Rent-seeking through litigation: Adversarial and inquisitorial systems compared. *International Review of Law and Economics, 22*(2), 193–216.

Plott, C.R. (1987). Legal Fees: A Comparison of the American and English Rules. *Journal of Law, Economics, & Organization, 3*(2), 185–192.

Poyago-Theotoky, J. (1996). R&D competition with asymmetric firms. *Scottish Journal of Political Economy, 43*(3), 334–342.

Risch, M. (2011). Patent troll myths. *Seton Hall Law Review, 42*, 457–499.

Rosen, R.J. (1991). Research and development with asymmetric firm sizes. *RAND Journal of Economics, 22*(3), 411–429.

Tirole, J. (1988). *The Theory of Industrial Organization.* Cambridge, MA; London, England: The MIT Press.

Tullock, G. (1975). On the Efficient Organization of Trials. *Kyklos, 28*(4), 745–762.

Tullock, G. (1980). *Trials on Trial: The Pure Theory of Legal Procedure.* New York: Columbia University Press.

Vickery, N. S. (2015). Don't Forget about the Little Guys: Trolls, Startups, and Fee Shifting. *Colorado Technology Law Journal, 13*, 171.

Yamazaki, T. (2008). On the Existence and Uniqueness of Pure-Strategy Nash Equilibrium in Asymmetric Rent-Seeking Contests. *Journal of Public Economic Theory, 10*(2), 317–327.

Yin, X., & Zuscovitch, E. (1998). Is firm size conducive to R&D choice? A strategic analysis of product and process innovations. *Journal of Economic Behavior & Organization, 35*, 243–262.

Zaby, A.K. (2010). Losing the Lead: The Patenting Decision in the Light of the Disclosure Requirement. *Economics of Innovation and New Technology, 19*(2), 147–164.

A Figures

In all of the following figures, black curves represent N and orange curves (printed as gray) represent M.

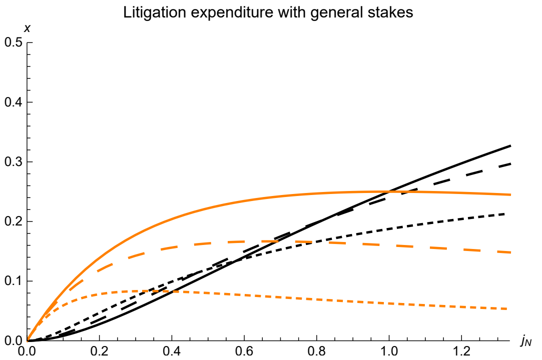

(a)

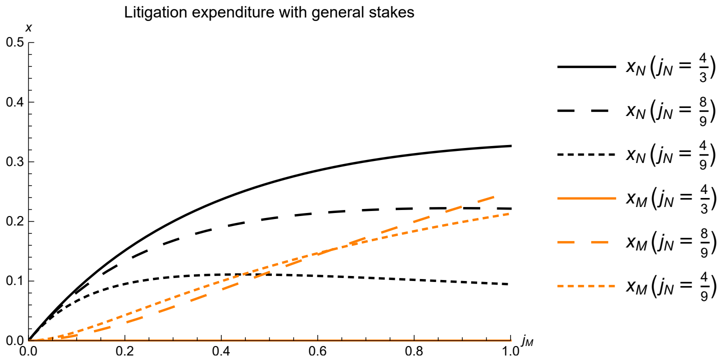

(b)

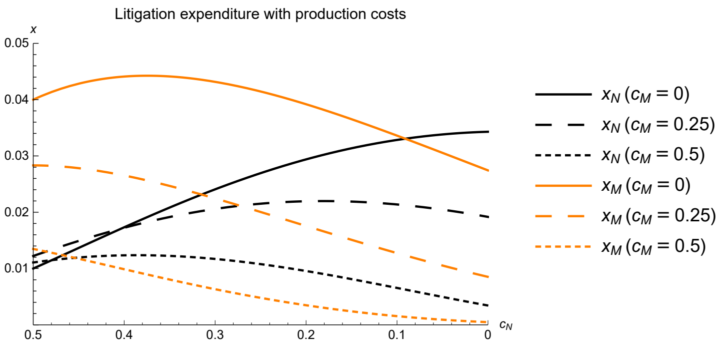

(c)

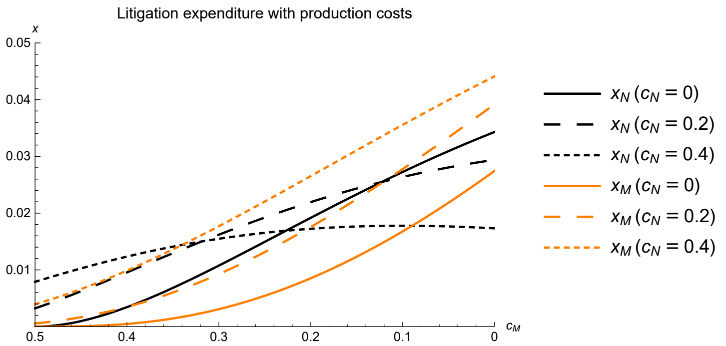

(d)

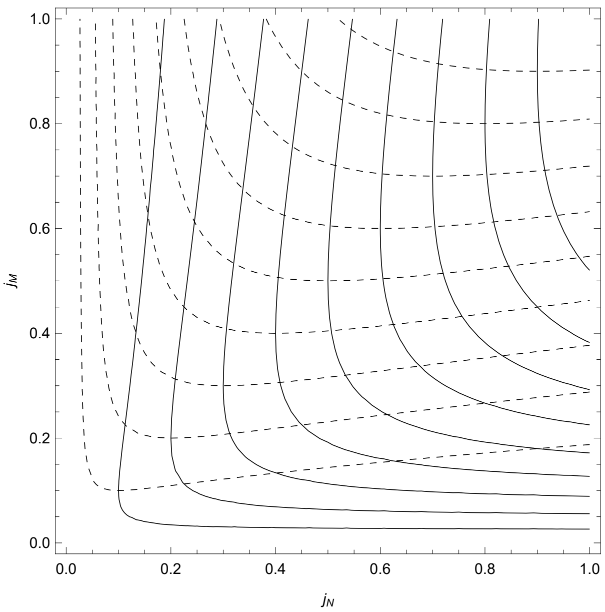

(e) Contour plot with general stakes

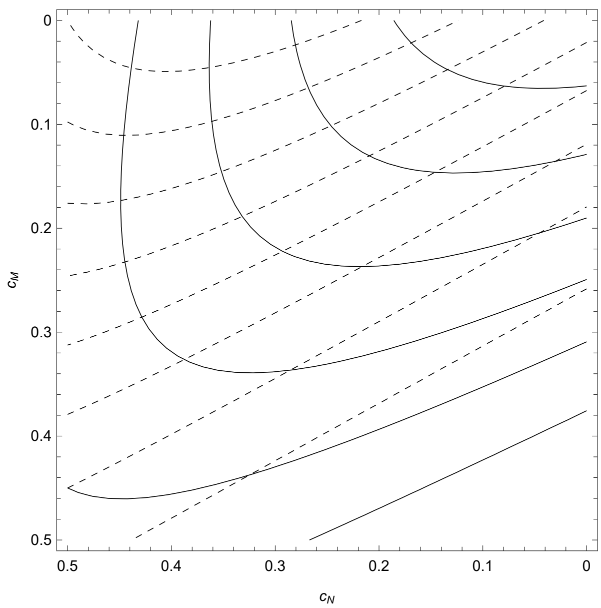

(f) Contour plot with explicit costs

Figure 2.2: Litigation expenditure by N and M; comparing the case with general stakes j_N and j_M to the case with explicit production costs c_N and c_M. In order to allow for easier comparison, the direction of the horizontal axis is reversed when plotting the case with explicit costs.

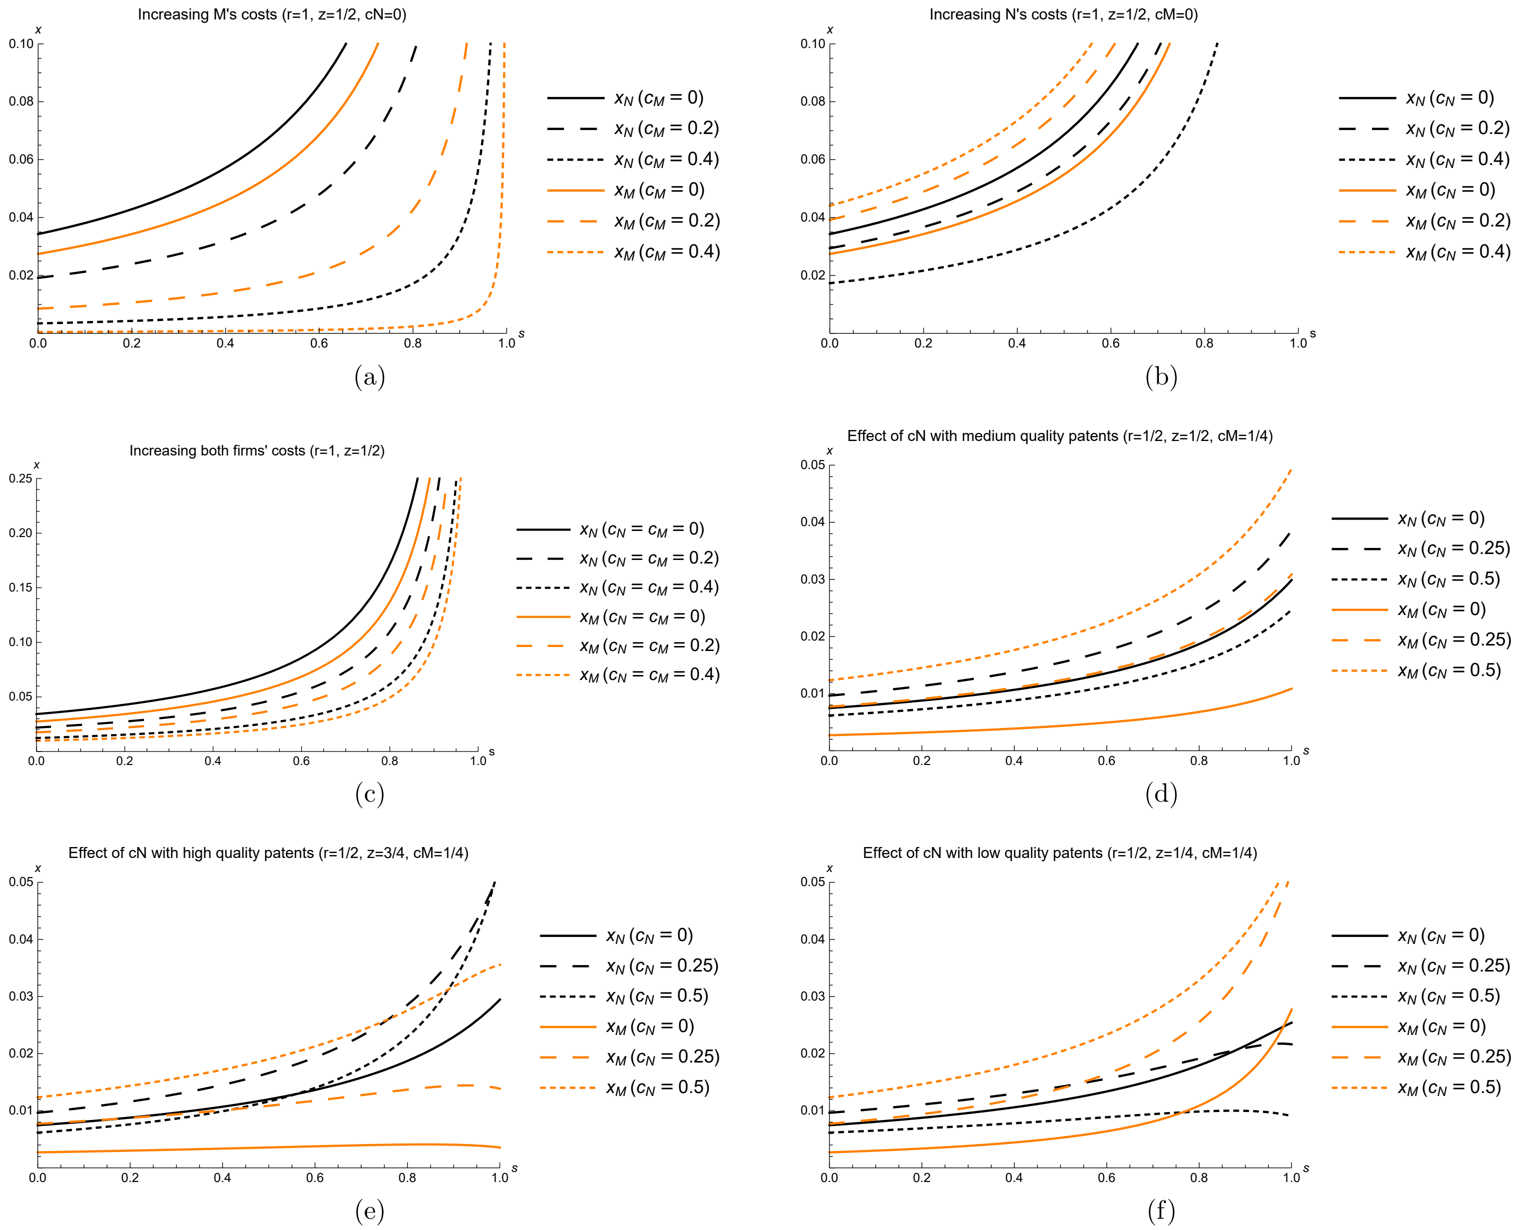

Figure 2.3: Comparative statics of litigation expenditure with explicit production costs (I).

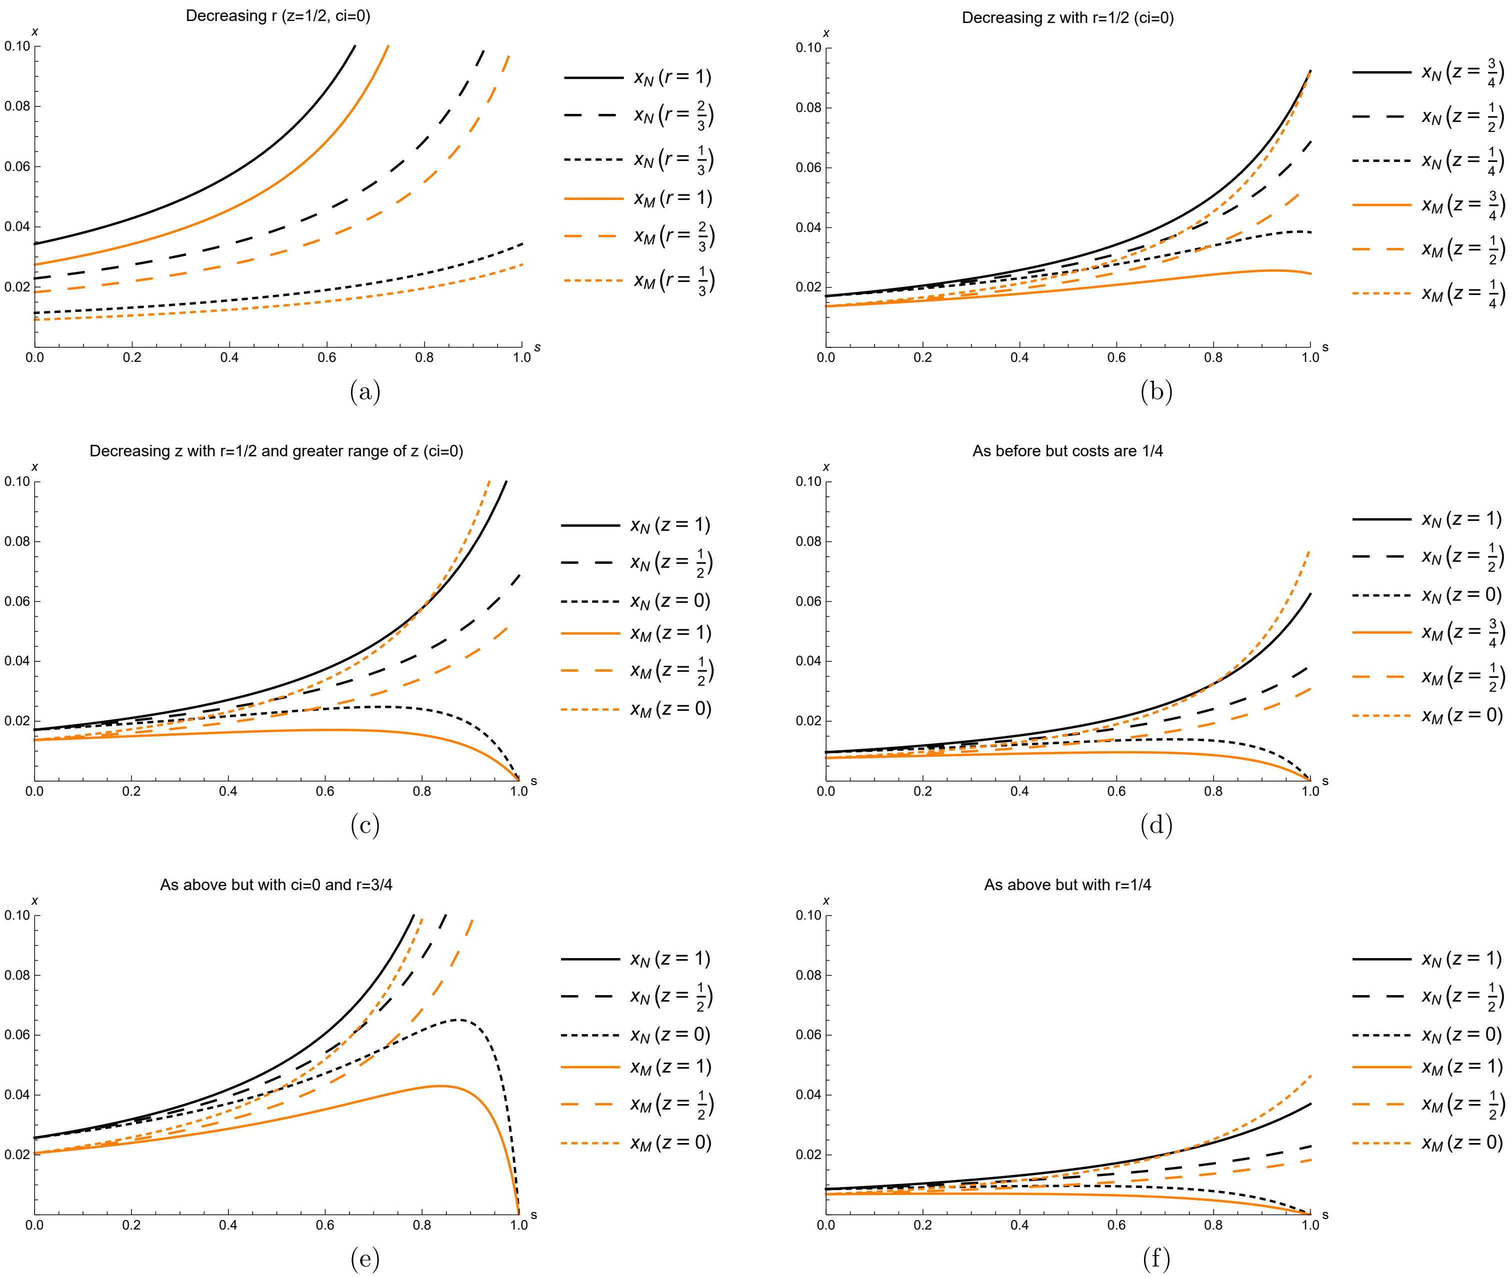

Figure 2.4: Comparative statics of litigation expenditure with explicit production costs (II).

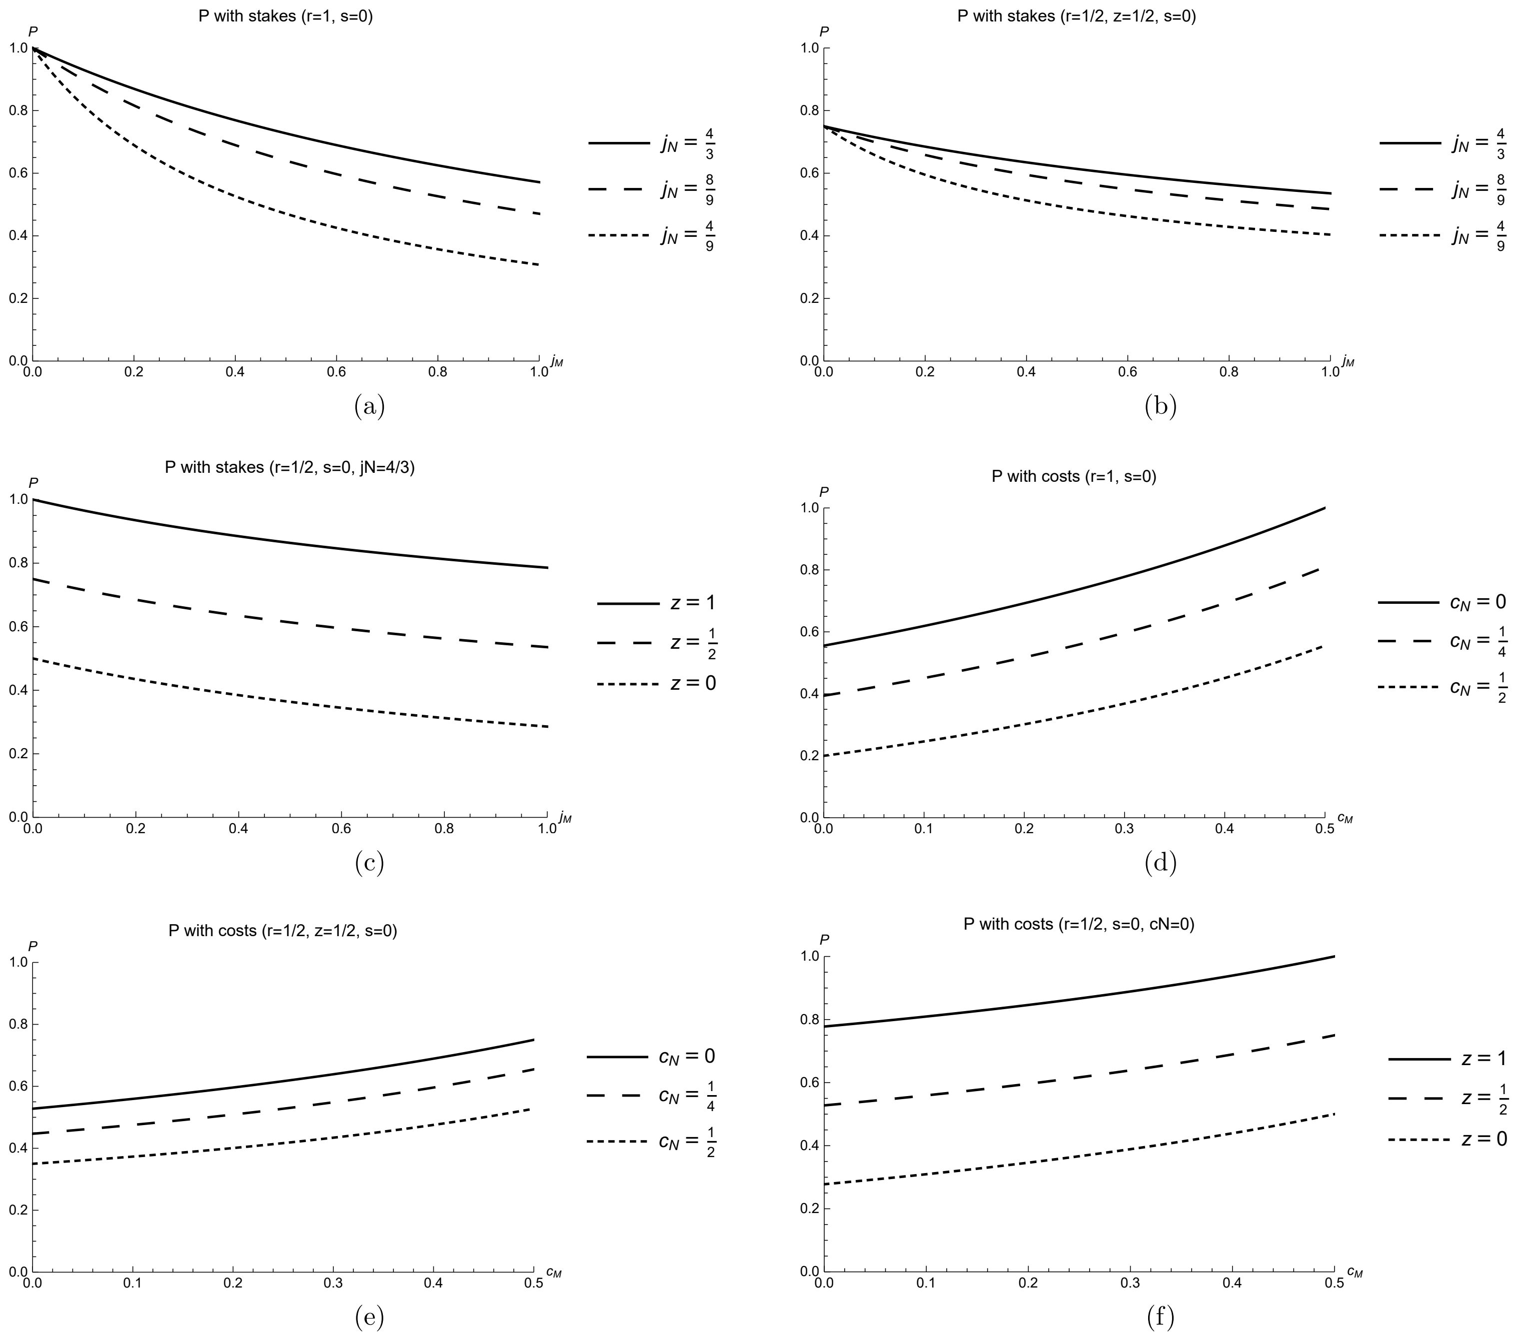

Figure 2.5: Expected litigation outcome P^* as a function of imitator's size.

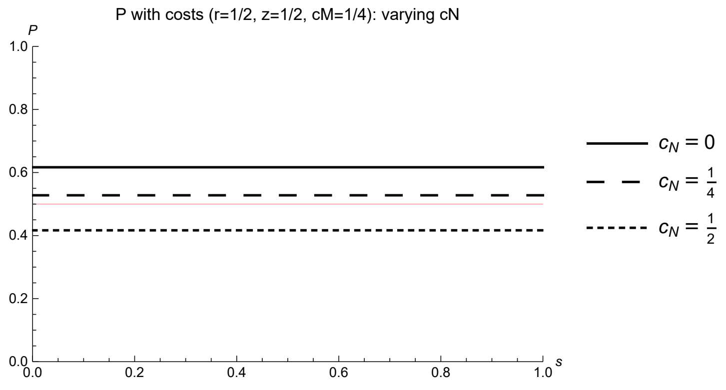

(a)

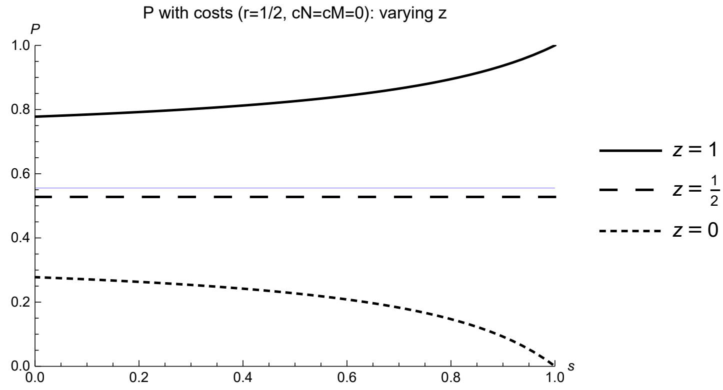

(b)

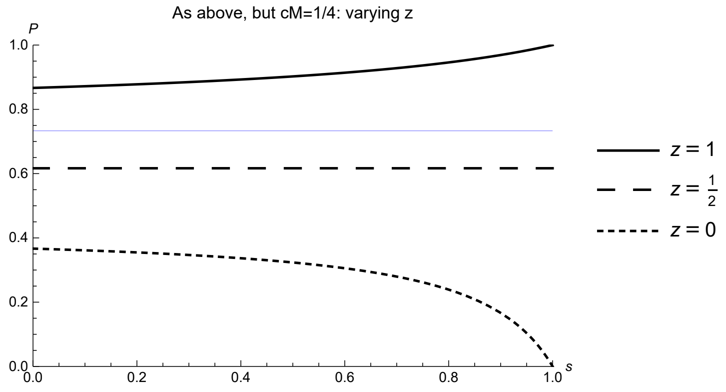

(c)

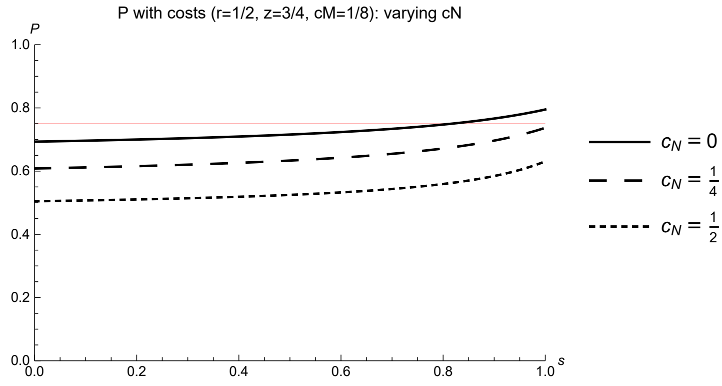

(d)

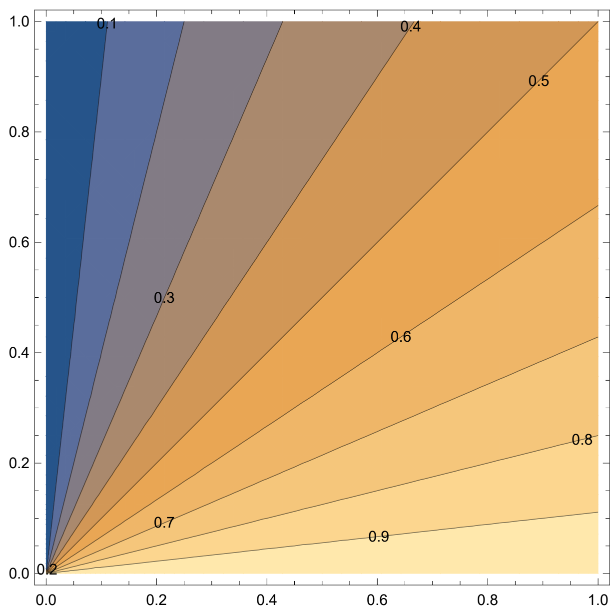

(e) Contour plot with general stakes

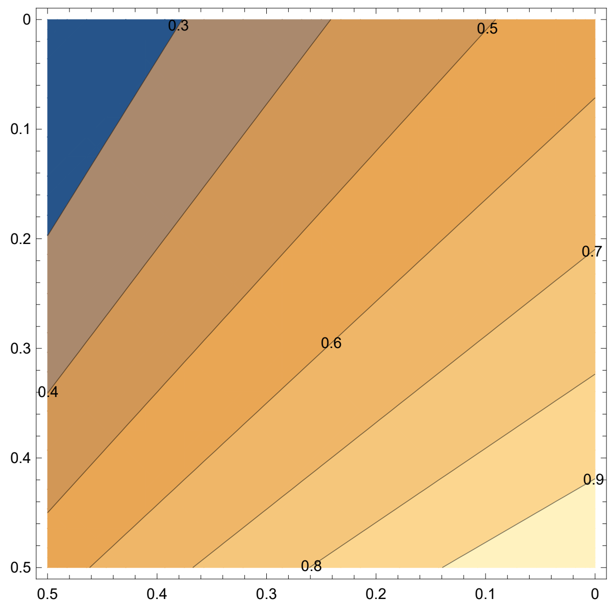

(f) Contour plot with explicit costs

Figure 2.6: Expected litigation outcome P^* as a function of the fee-shifting parameter s.

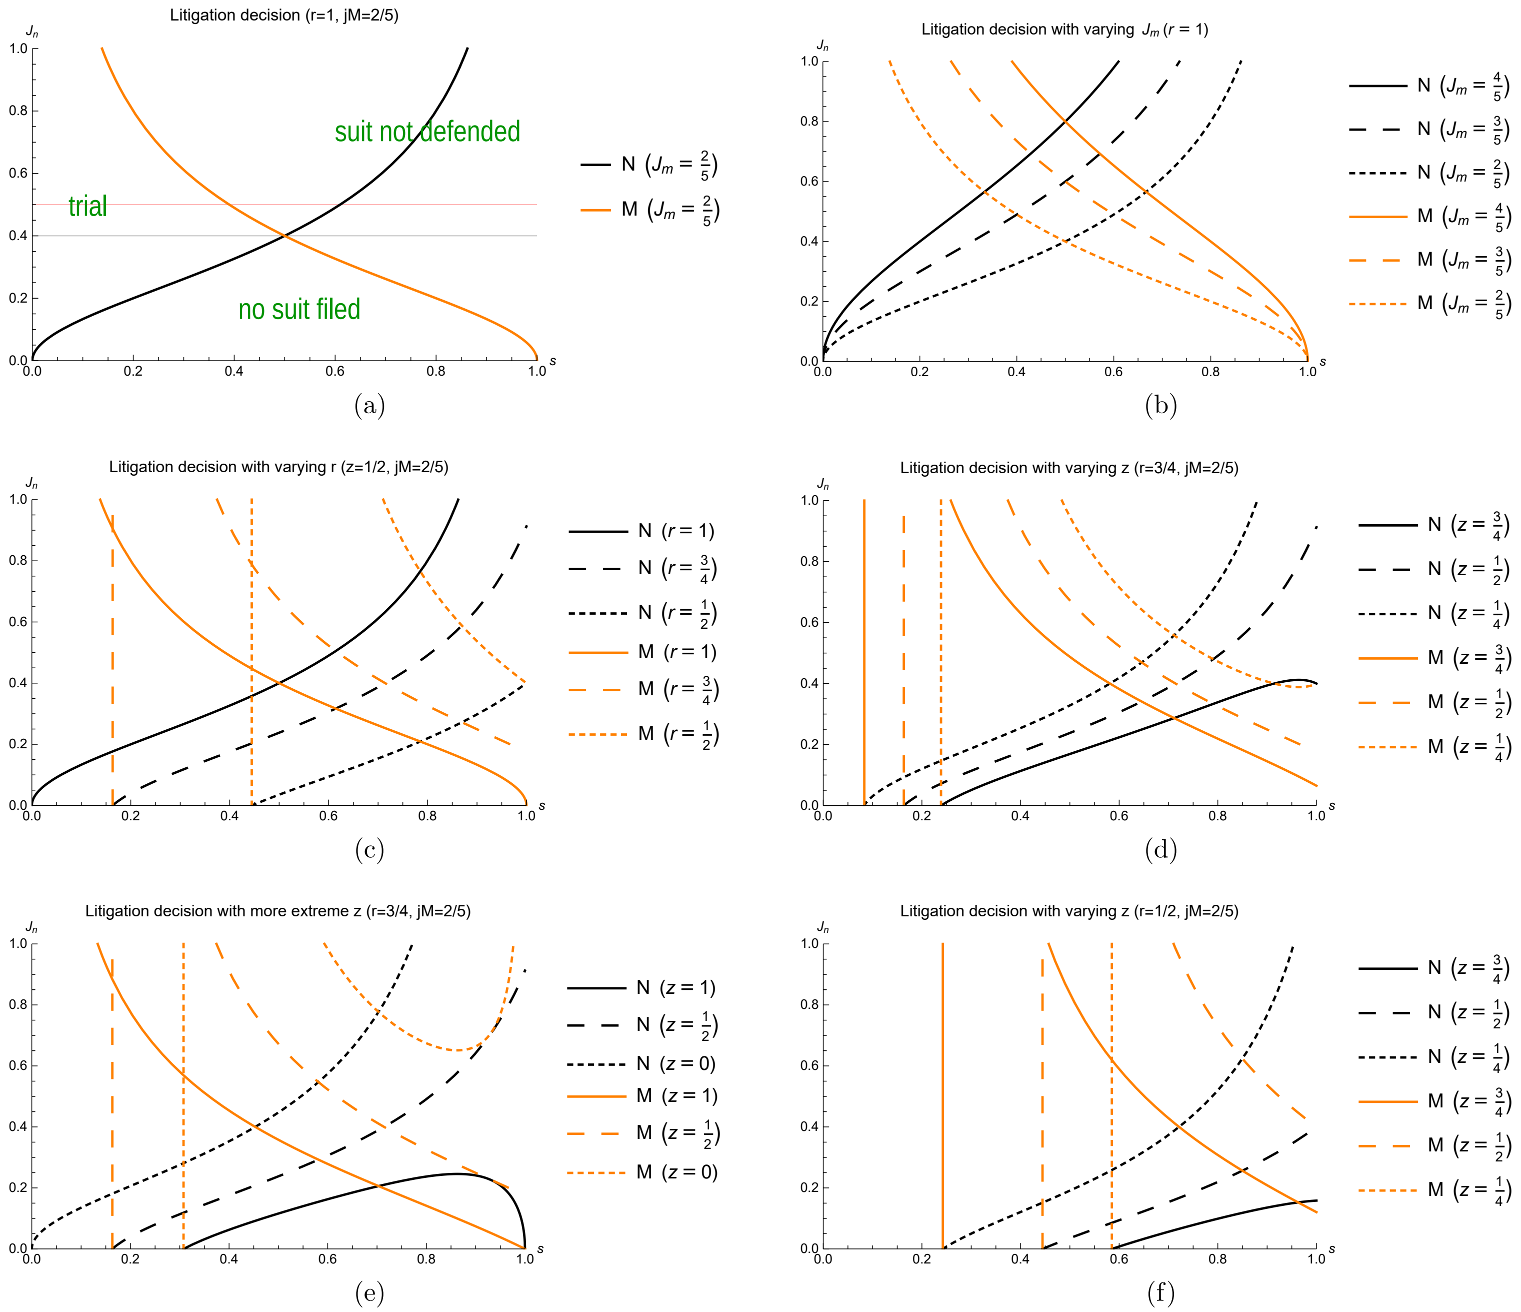

Figure 2.7: Litigation decision with general stakes. Panel a) shows the general structure that applies to all other pairs of litigation participation curves: In the area above N's and below M's participation constraint, a suit moves to trial; in the area below both firm's participation constraints, N will not file a suit; and above M's participation constraint, N files a suit that M will not defend. If this is below N's participation constraint, N does so by exploiting her first-mover advantage.

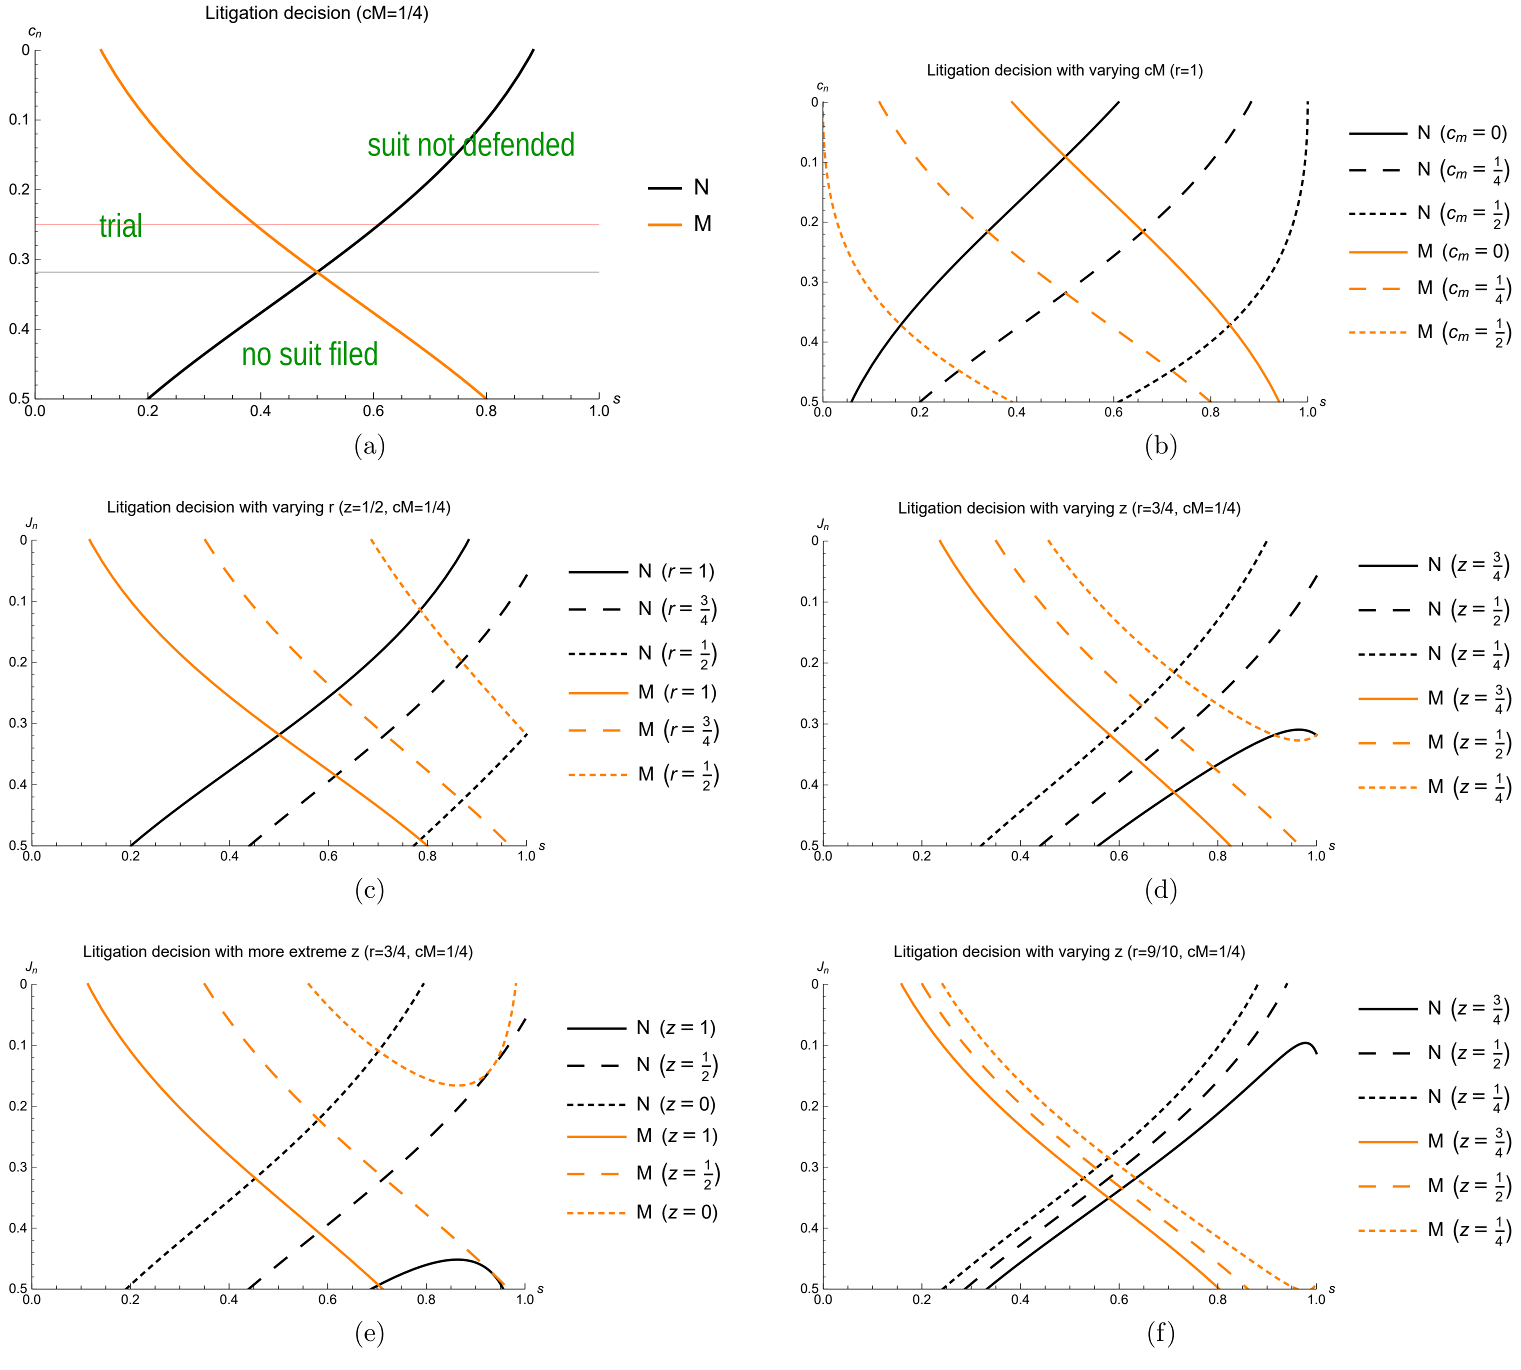

Figure 2.8: Litigation decision with explicit costs.

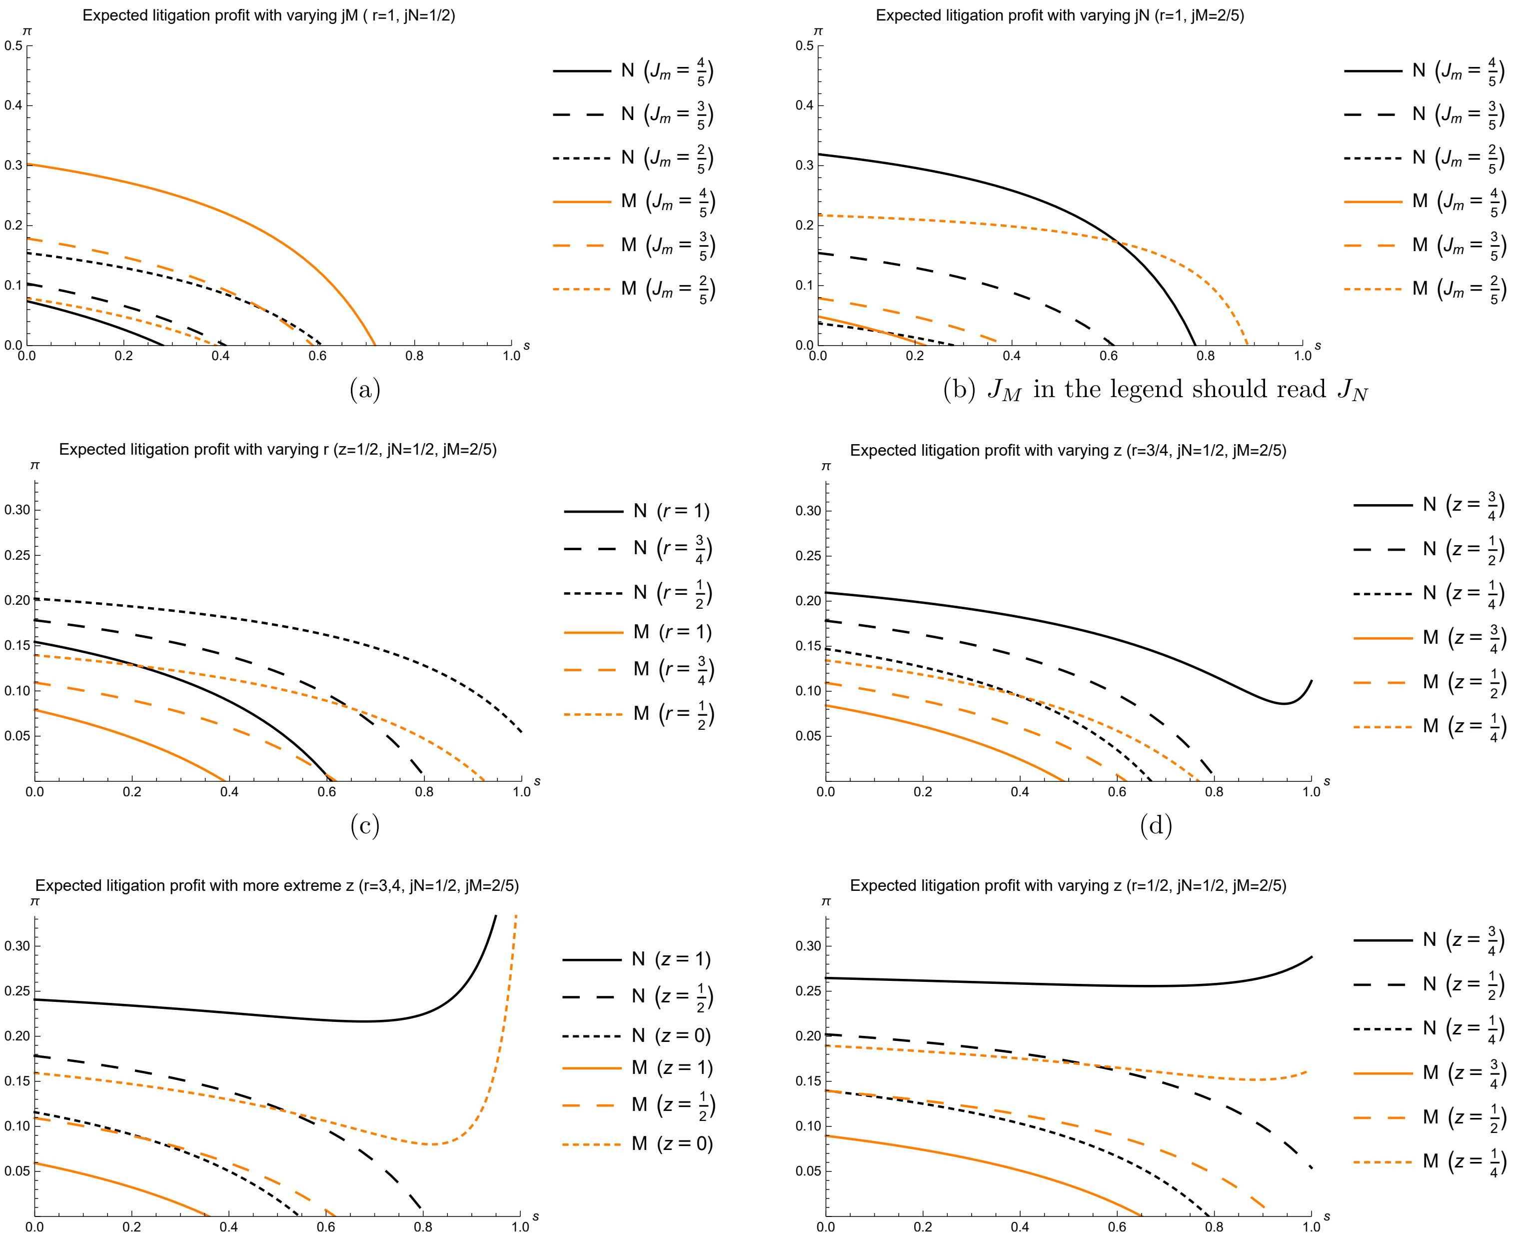

(a)

(b) J_M in the legend should read J_N

(c)

(d)

(e)

(f)

Figure 2.9: Expected profit from litigation with general stakes.

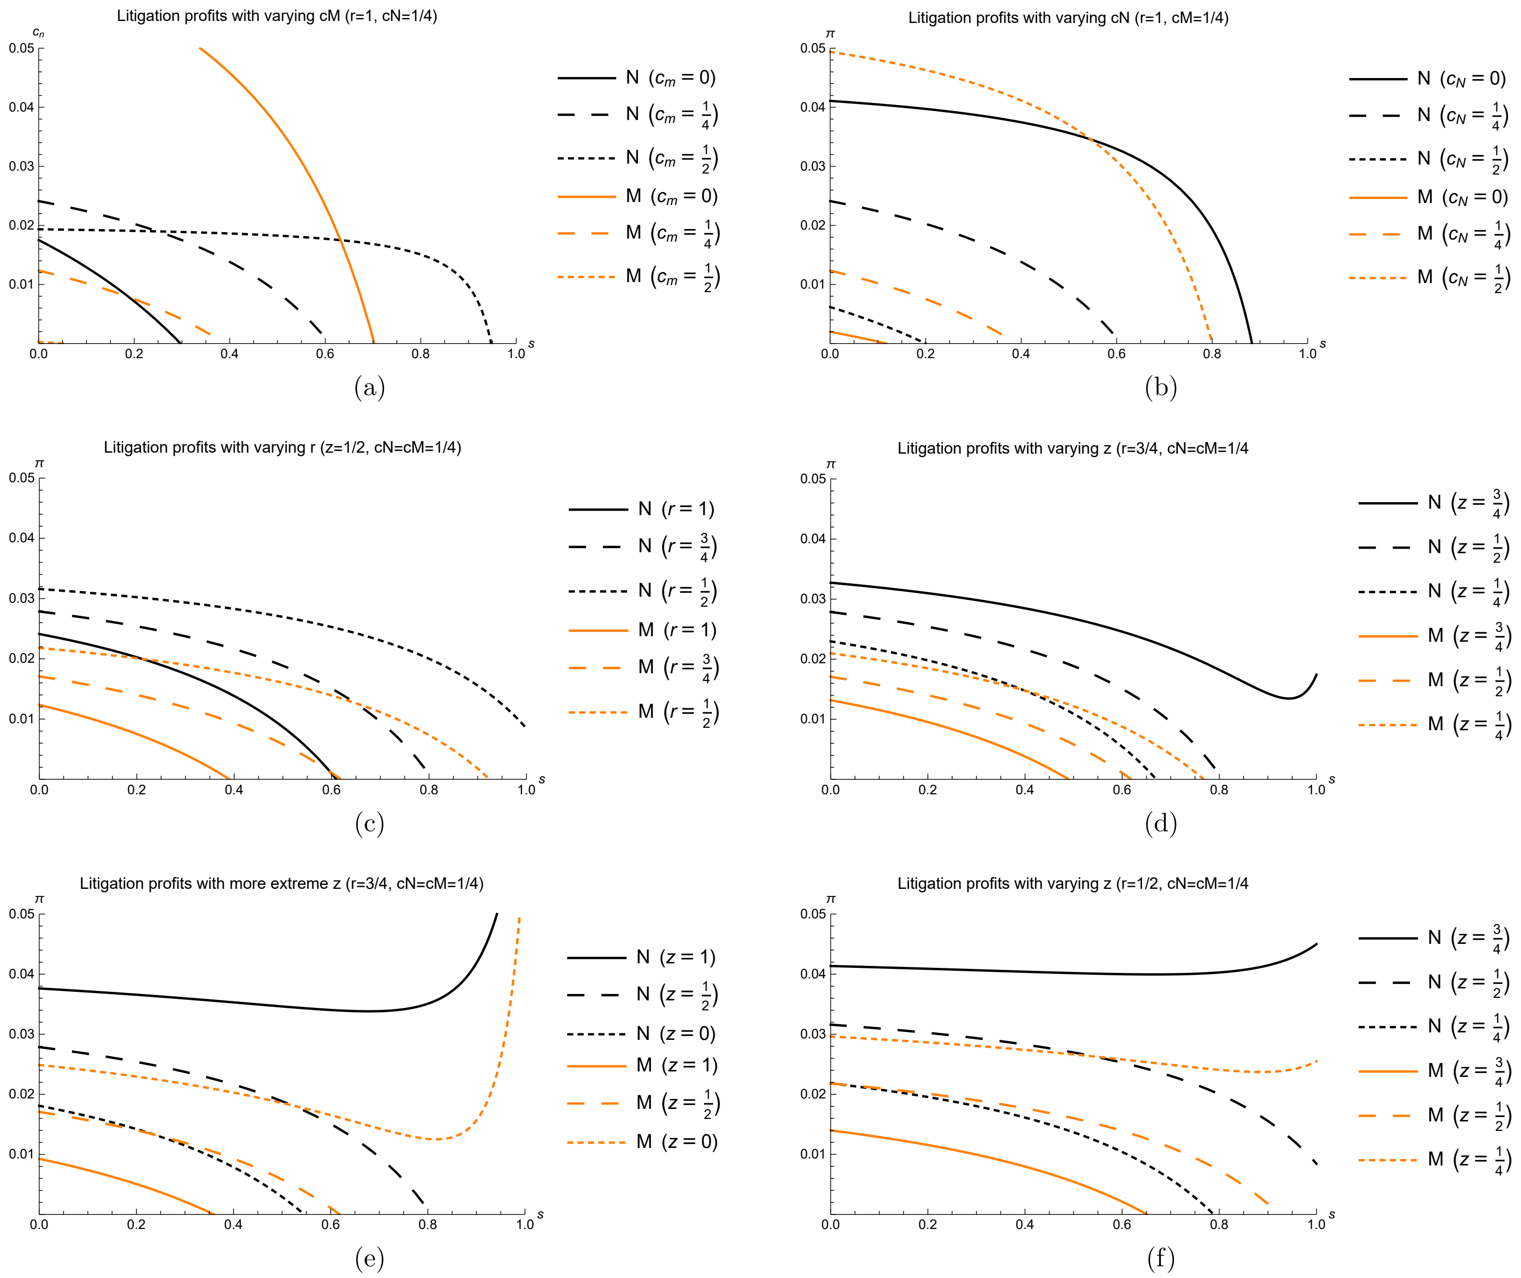

Figure 2.10: Expected profit from litigation with explicit costs.

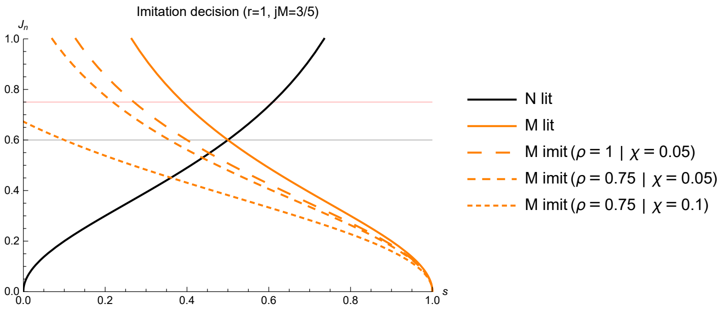

(a) Imitation decision with general stakes

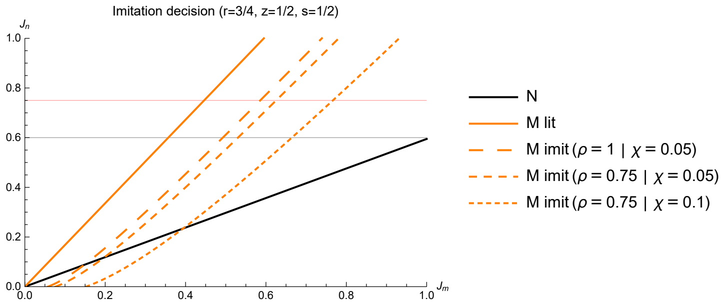

(b) Imitation descision as a function of j_M

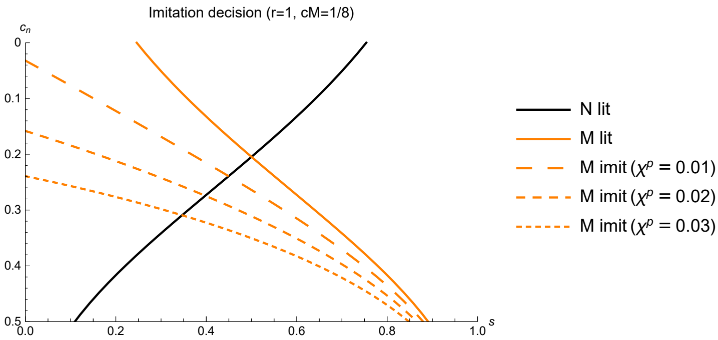

(c) Imitation decision with explicit costs

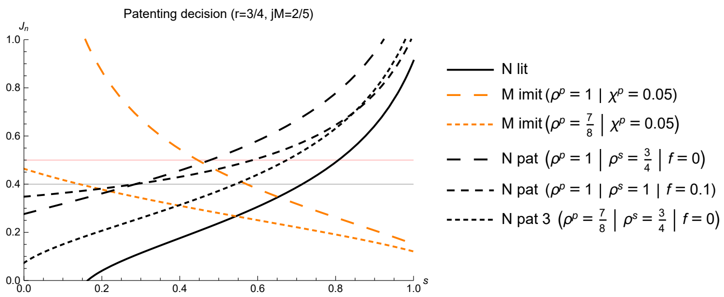

(d) Patenting decision

Figure 2.11: Imitation and patenting decisions.

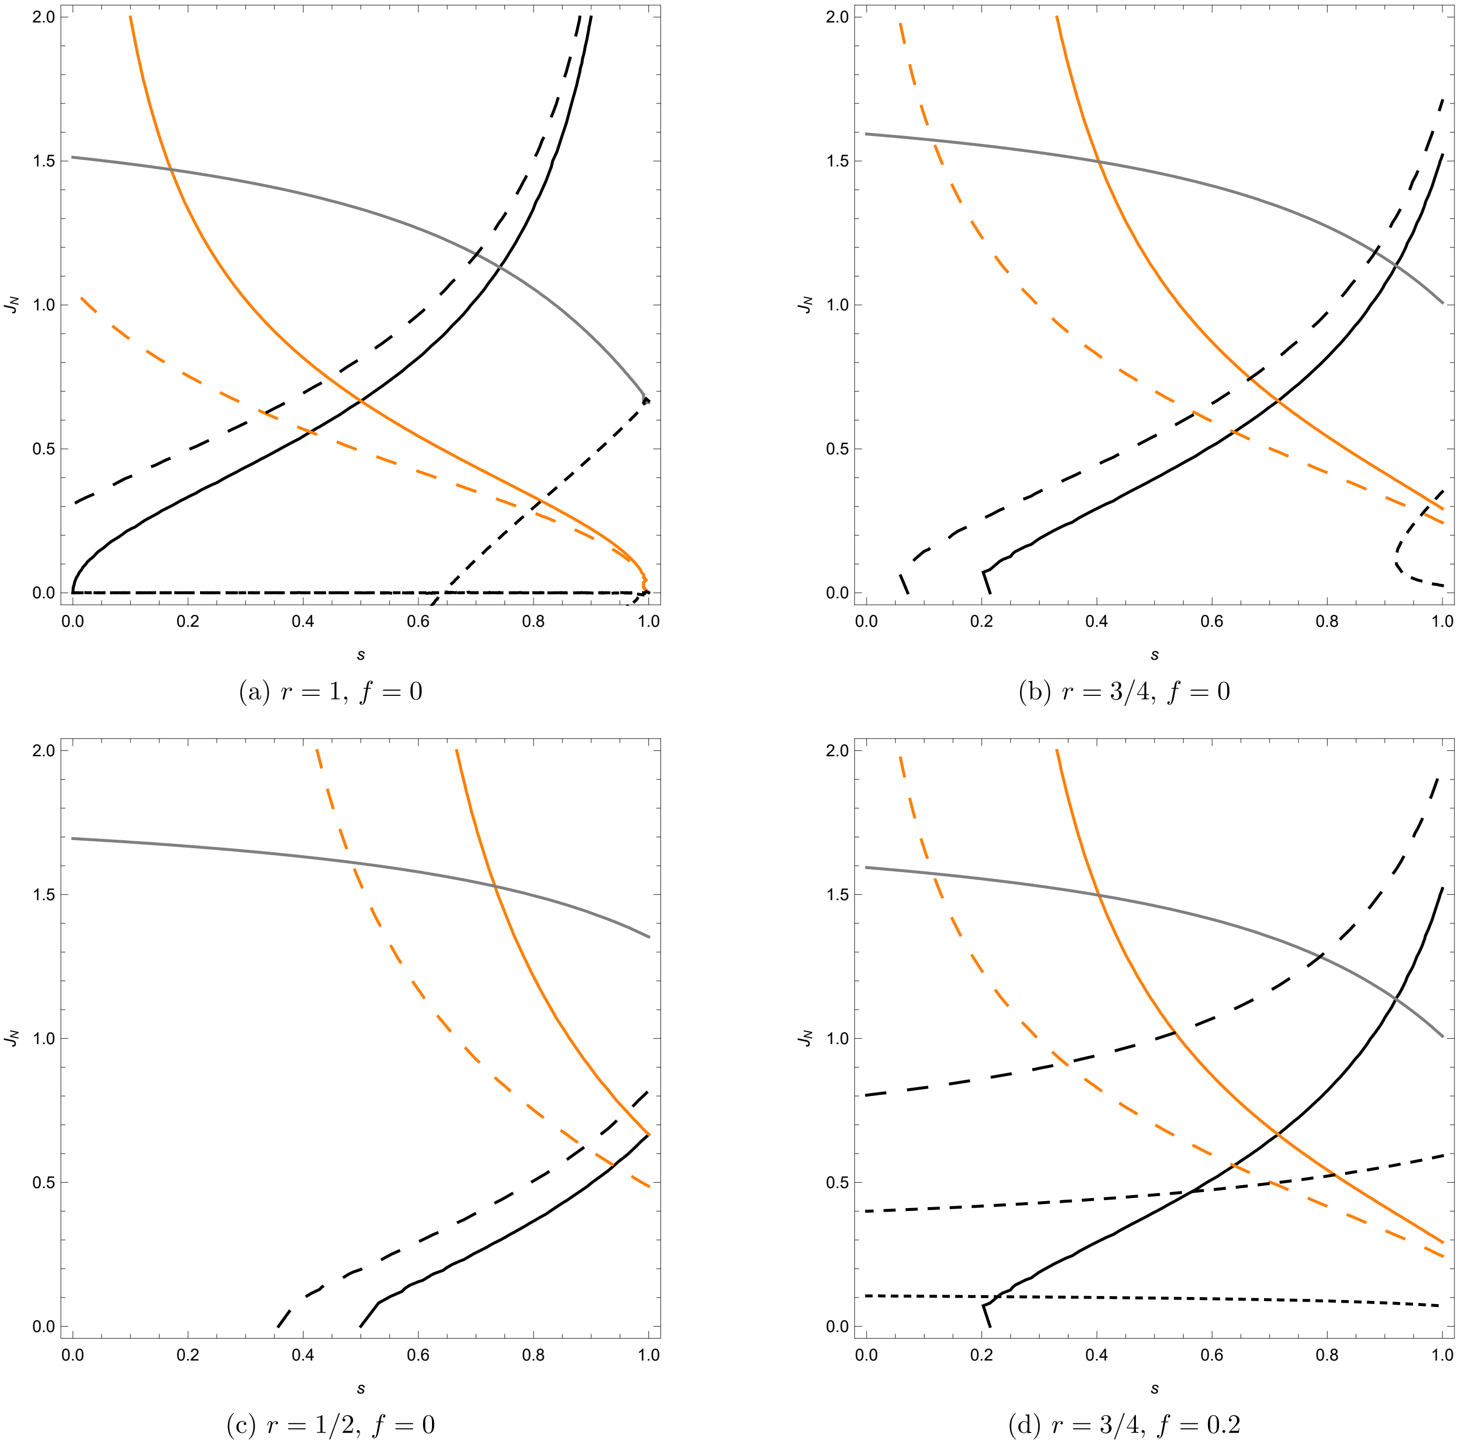

(a) $r = 1,\ f = 0$

(b) $r = 3/4,\ f = 0$

(c) $r = 1/2,\ f = 0$

(d) $r = 3/4,\ f = 0.2$

Figure 2.12: Patenting decisions with settlement (I).
Legend: solid black: N's litigation decision; solid orange: M's litigation decision; dashed orange: M's imitation decision (when suit leads to trial); solid gray: profitability of settlement (if > 0, settlement is profitable); long-dashed black: N's patenting decision with trial; medium-dashed black: N's patenting decision when firms share bargaining surplus equally; short-dashed black: N's patenting decision when the full surplus is transferred to N.
Parameter values: $J_M = 2/3,\ z = 1/2,\ \rho^p = 1,\ \chi^p = 0.1,\ \rho^s = 0.9.$

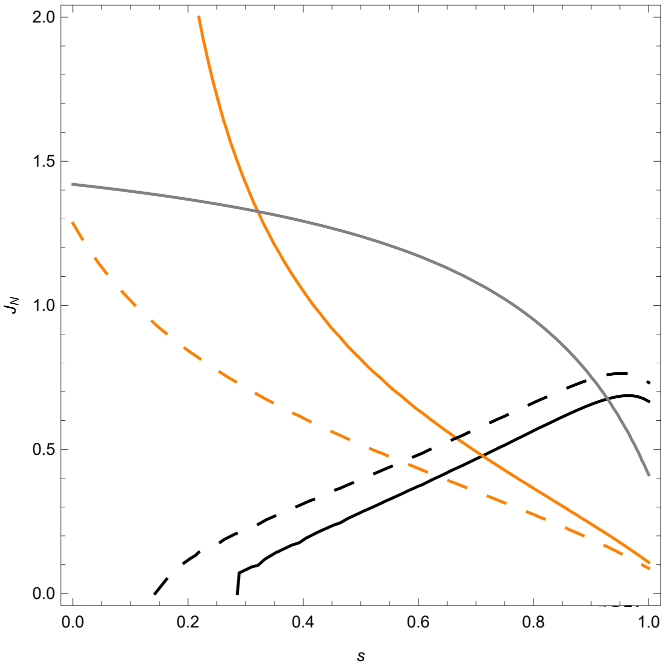

(a) $J_M = 2/3,\ r = 3/4,\ z = 3/4,\ \rho^s = 0.9$

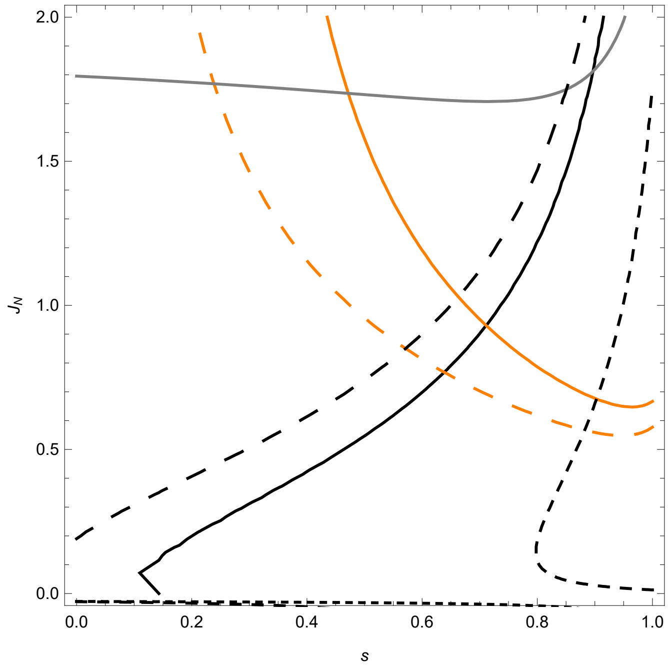

(b) $J_M = 2/3,\ r = 3/4,\ z = 1/4,\ \rho^s = 0.9$

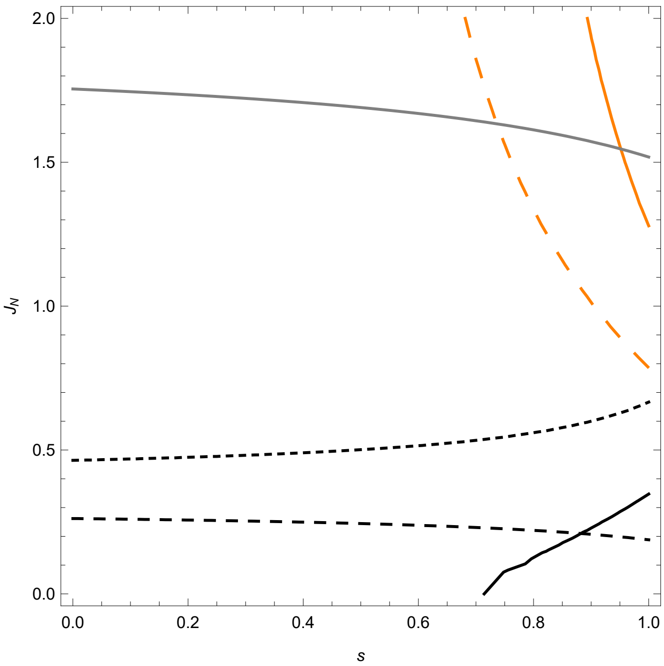

(c) $J_M = 2/3,\ r = 3/8,\ z = 1/2,\ \rho^s = 0.2$

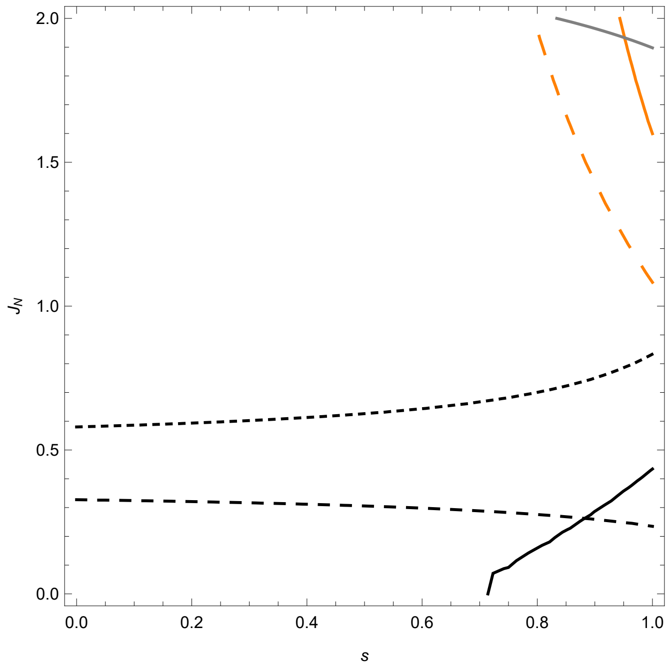

(d) $J_M = 5/6,\ r = 3/8,\ z = 1/2,\ \rho^s = 0.2$

Figure 2.13: Patenting decisions with settlement (II). Legend as in figure 2.12. Note that in panels c) and d) the black dashed lines depict upper bounds for patenting, not lower bounds as in all previous panels.
Parameter values: $\rho^p = 1,\ \chi^p = 0.1,\ f = 0.$

Chapter 3. As Soon as Optimal: Delayed Publication of Patent Applications

1 The issue

Patent applications are commonly published by the receiving patent office 18 months after the filing date. Yet, a common assumption made throughout the literature on studying firm competition for patents is that patent applications are immediately observable to all competitors. In the context of a patent race, firms decide to spend resources on a research project until the first participant succeeds. The winner of the race then files a patent application and all spending by any party on that particular project stops. A comparable assumption is at times made even when research output is not immediately patented, implying that while the technical knowledge related to a successful innovation is in possession of one firm only, the information about the existence of that innovation spreads immediately. Such an assumption simplifies the analysis of research competition, avoiding the necessity of firms to form beliefs about whether the race they are participating in is indeed still ongoing, and fits the similarly common assumption of perfect information between contestants.

Since at least Horstmann et al. (1985) it is acknowledged that it may indeed be in firms best interest not to patent each and every patentable research and development (R&D) output that they come up with precisely because a patent application itself is not only a government-issued license for a temporary monopoly in using a certain technology. Part of the bargain behind the patent system is that inventors' technical knowledge is disclosed to the public in order to let others benefit from it, by knowing what technical solutions have worked, what is currently protected, and how to imitate a patented innovation once the patent expires. This publicity, combined with a standardised format and a finely-grained technological classification of every submitted invention, makes it seem credible that both filed applications and granted patents are indeed a valuable source of information not only about the state of the art of a technological field, but also about the (patented) output of competitors' R&D departments.

Filing a patent application, however, does disclose more information than details about the technical nature of the innovation. Observers may use the patent application to infer all sorts of things about the innovation and its developing firm. In the mentioned Horstmann et al. it is in fact signalling the value of the innovation that might attract imitators that keeps innovators from patenting. A much more straightforward piece of information disclosed by a patent application is that the applicant has developed the underlying technology to the state of patentability. As mentioned before, in a standard patent race model, filing of the first patent ends the race for all participants. In all major patent systems today, a patent application is however only published 18 months after filing. Comparing this with evidence on the average length of R&D projects leading to patentable output of 6-7 years, an additional period of 18-months in which the losers of the race potentially keep on spending could possibly have serious consequences on the welfare effects of the patent system.

Interestingly, this time delay between filing and publication of a patent application has not received attention in the academic literature.[1] It might have been taken to be of little importance as the winner of a race could have been assumed to have an incentive to publicly announce this fact in order to drive competitors away from the patented innovation and towards the next yet-unpatented opportunity. A quite extensive literature studying the decision to patent—discussed in the next section—has found that there may be various reasons for firms to forgo patent protection entirely if the associated disclosure of innovation, be it on technology or on issues of more general business strategy, was perceived detrimental. All the findings of this literature could be expected to also apply in the context of the decision to announce or not a pending patent application, in particular since the decision to not apply for patent protection can come at the cost of being unable to prevent imitation, while this cost does not arise when in fact a patent application has already been filed.[2]

The concern of the present paper is therefore to ask what is the effect of the common statutory delay in patent publication on R&D competition between firms, and what are the incentives of the winner of a patent race to disclose or not the existence of an application. Based on this it might then be asked whether the publication delay merits further academic scrutiny.

After discussing the literature in the next section and introducing a very simple model of duopoly R&D competition in the third section, I proceed in two main steps. Section four introduces a statutory delay period in a simple race for a single innovation, a situation in which the winning firm is effectively indifferent between disclosing or not her pending patent. There, I ask the question how, conditional on the leading firm not disclosing her patent application, the existence of this statutory delay affects firms' participation decisions. The delay increases expected costs from participating over the whole expected duration of the race, but the effects are different from simply increasing the cost of performing R&D. While increased cost of R&D increase the efficiency level required of firms to profitably participate in the race at all, by construction an increased publication delay does not have an effect on firms' incentives to start the race in the beginning. It is only as the race progresses that the expected cost of participation reach a prohibitive level.

In sections five and six, instead, I take the patent race out of isolation and move attention to some fundamental cases in which firms may have an incentive to disclose their patent applications. Such cases involve the existence of a second patent race, which may either sequentially follow the first one (section five) or be connected to the first one via commercial complementarity between the resulting innovations (section six). Generally, I

[1] To the best of my knowledge, the only exception is an unpublished working paper by Ganglmair & Oh (2014).

[2] Of course, a patent application is only helpful in preventing imitation if the patent will be eventually granted, but this problem is true in both the decision to patent and the decision to disclose a pending patent application.

find that weaker firms may have an incentive to disclose a patent application to shift their competitor's attention to the other, missing innovation. I extend the analysis to a case in which firms have incomplete information about patent strength. I end section 5.2 by showing that in a setting with a sufficiently large number of complementary innovations, there could even be an incentive to delay patent application if statutory delay was set to zero.

The final section summarises and concludes that a more precise quantification of the net welfare effect of patent publication delay would be helpful in understanding the importance of this overlooked tool of patent policy.

2 The State of the Literature

Patent race models over the past decades have become some the workhorse models of theoretical economics literature on innovation, originating from works by Loury (1979) and Lee & Wilde (1980). A sizeable literature has built on these, reviewed first in Reinganum (1989).

Many features of the patent system have been studied with the help of patent races, such as patent duration and the breadth of protection (Denicolò, 1996; 1999), whether the patent is awarded to the first inventor or the first applicant (Scotchmer & Green, 1990), if complementary innovations should be patentable separately or only as a bundle (Ménière, 2008), if the award should only go to the single winner or if it should be shared to some extent with subsequent duplicators (Denicolò & Franzoni, 2010), or the size of the "non-obviousness requirement" in patent law that controls that governs how far an innovator must go beyond an existing patent to be able to receive patent protection herself (reviewed in Denicolò, 2008).

Regarding the delayed publication of patent applications not much has been written by economists. The legal literature has provided comments on the "bringing in line" of the US patent system with most other patent systems by the American Inventor Protection Act (AIPA) which came into force in 2001 (e.g., Watase, 2002; Ergenzinger, 2006). Up until then, US patent applications themselves were in fact never published, so a pending patent application was in fact secret to the outside world until a patent was eventually granted. Some economists and innovation researchers used the AIPA as a sort of natural experiment to identify the effect of having any disclosure at all (after 18 months) of patent applications—compared to having no disclosure—on matters such as patent grants or the diffusion of knowledge. Hegde & Luo (2017), for instance, find that patent licensing agreements are entered into considerably earlier post-AIPA, interpreting disclosure via the patent office as reducing information cost for both sides of a licensing transaction.

The theoretical literature went on as before though and extended the scope of patent races to the analysis of issues such as cumulative innovation (Scotchmer & Green, 1990; van Dijk, 1996; Scotchmer, 1996; O'Donoghue, 1998; Denicolò, 2000; Hunt, 2004), complementary innovations (Fershtman & Kamien, 1992; Green & Scotchmer, 1995; Denicolò & Halmenschlager, 2012; Biagi & Denicolò, 2014; Denicolò & Zanchettin, 2018).

Patent races have also been modified to allow the decision to keep an innovation secret (Denicolò & Franzoni, 2004; Kwon 2012a,b) or the decision about timing of the patent application as such (i.a., Langinier, 2005; Hopenhayn & Squintani, 2016). Somewhat related to the present paper is Akcigit & Liu's (2016) study of the inverse problem by allowing for "terminal failure" of R&D projects, i.e., an R&D project can turn out to be a dead end with no chance of eventual success. While every patent application is again immediately observable, non-observation of a patent application can now either mean that the competitor simply did not yet succeed, or that he realised that the project was a dead end and accordingly moved on. Both the works by Akcigit & Liu as well as by Hopenhayn & Squintani stand out by including a robustness check to verify that firms indeed prefer filing a patent application on a successful innovation instead of keeping the successful innovation secret. The authors acknowledge that in the latter case it would seem unrealistic to argue that the non-innovating competitors would still immediately learn about the existence of the new innovation. Fortunately for the authors, in the context of their models the option of immediate patent application is strictly preferred. It is however easy to show that such a strict relation would not hold for immediate public disclosure of an otherwise secret patent application.

Only few theoretical contributions specifically acknowledge this institutional setup of the patent system. Aoki & Spiegel (2009) compare firms' incentives to patent an intermediate innovation when patent applications are published or not before they are granted. The only model in which a patenting firm can choose to voluntarily disclose an otherwise secret patent application is a so-far unpublished paper by Ganglmair & Oh (2014). A successful innovator may decide to not disclose a pending patent if this affects a follower firm's incentive to invest in innovation themselves, as well as the intensity of market competition. Their model belongs to the class of non-race models in which an innovator has to choose what to do with an existing innovation. Most commonly this is about the decision to patent, such as in the fore-mentioned Horstmann et al. (1985), Anton & Yao (2004), Zaby (2010, or the preceding chapter of this dissertation, and Ottoz & Cugno (2008) for the case of complementary innovations. In contrast, the present model starts in a situation of on-going R&D competition; and unlike in Ganglmair & Oh (2014), incomplete information is not a strictly required ingredient to generate an incentive to (not) disclose a pending patent application.

Beyond this one paper, a number of previous studies have analysed the use of patents to transfer information, as opposed to protecting an invention (Crampes & Langinier, 1998; Long, 2002; Conti et al., 2013; Asay, 2015; Comino & Graziono, 2015). Very recently these considerations have been underpinned by some empirical evidence (Hottenrott et al., 2016; Saidi & Zaldokas, 2016). Finally, there is some ambiguity regarding optimal timing of decisions beyond patent application. Firms might decide to engage too early or too late in innovation (Barzel, 1968), patent litigation (Choi, 1998), imitation (Ponce & Henry, 2011), and commercialisation of patented ideas (Rousakis, 2014).

Do complex product firms engage in patent races? Kultti et al. (2006) report that in many industries, particularly network industries of which some are commonly regarded as producing complex products, e.g., consumer electronics and telecommunications, standardisation has narrowed down the set of possible paths for innovation so that most firms are in fact simultaneously working on the same inventions. While Sternitzke et al. (2013) contribute evidence of "patent-race like competition" in pharmaceuticals, Thompson & Kuhn (2017) study the occurrence of patent races in a multitude of technology fields and find that they are most common in information-technology fields, which again are characterised by complex products. Kim et al. (2017) exploit a change in the legal patent term to generate evidence allowing interpretation of increasing size of inventor teams to be driven by firms trying to finish R&D races more quickly. If this is true, R&D races recently have indeed become more intense.

Do firms strategically choose the timing of patenting-related decisions? The theoretical literature has identified a number of reasons why innovative firms might find it beneficial to delay filing of patent applications for otherwise patentable inventions. Firstly, in a setting of discrete and independent inventions, Langinier (2005) shows that patentees have an incentive to delay patenting when they are leading an R&D race of perfect information. When information is asymmetric, successful innovators optimally patent randomly and therefore more often, preventing their competitors to interpret patents as signals of completed R&D projects that should be abandoned by lagging firms. Ganglmair & Oh (2014) analyse circumstances in which it is optimal for an applicant to publicly announce their recently filed patent applications and those in which it is rather optimal to wait and use the time of secrecy until the application is inevitably disclosed by the patent office. In Hopenhayn & Squintani (2016a), innovators may patent too early relative to the social optimum duration of a patent race out of a "fear of preemption" by a rival. Stronger patents can in fact weaken this incentive.

Hall & Harhoff's (2012) survey of "Recent Research on the Economics of Patents" includes a sub-section entitled "timing", but that timing is about covers the extent to which patent applicants can and do influence the beginning of the examination of a pending patent. For example, Palangkaraya et al. (2008) argue that applicants delay patent *examination* in order to create investment uncertainty for competitors since a pending patent signals that the technology under development might already be owned. Henkel & Jell (2010) find that about 50% of applications to the German PTO include delays caused by the applicants themselves.

Similarly, a literature on "submarine patents" studied the incentives in the US patent system to delay the final granting decision of a patent during the time when patent applications would not be published, but this strategy is considered a general tool to hold up competitors, and the focus is not on how it affects innovation incentives.

Finally, Kim et al. (2016) provide evidence of an individually beneficial effect of waiting with patent application in settings with high uncertainty about future industry standards. Their analysis is based on the DVD standard, affecting mostly complex-product industries.

Most relevant to the present chapter is the explorative study by Graham & Hegde (2015) who find that 85% of those US patent applications that were eligible for application secrecy up until grant[3] opted for voluntary disclosure after 18 months; the authors conjecture that this is because a published patent may give credibility to the existence of the innovation, and that voluntary non-disclosure would lead to a delay of cumula-

[3] After the coming in force of the AIPA, a patent applicant can still choose to have their patent application not published by the US patent office if they announce to refrain from seeking patent protection of the same invention outside of the US.

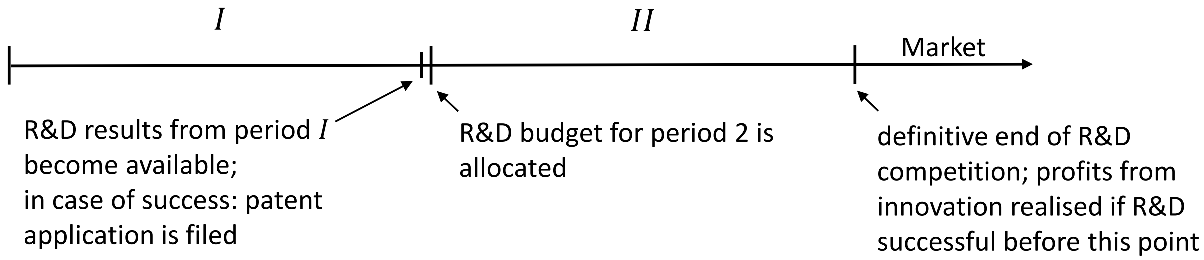

Figure 3.1: The basic structure of the two-period R&D model.

tive innovation and to duplicative R&D investments that is not in the interest of the applicant.

3 The Basic Model of R&D Competition

The analysis throughout this paper is based on variations of essentially the same simple model of R&D competition. Two firms A, B compete for one or two innovations with market profits π_1, π_2. The case of a single innovation arises in all variations by setting $\pi_2 = 0$. R&D activities are possible during two distinct periods, with results achieved during a period becoming available exclusively at the end of that period. There are therefore two distinct points in time at which R&D results can become available to a firm, and at which the firm—in case of R&D success—can then file a patent application. Figure 3.1 shows this basic structure.

Once a firm has successfully developed an innovation, it will immediately apply for a patent. It is well knownt that not all innovations are patented (see, e.g., the second chapter of this dissertation) but the present chapter restricts attention precisely to those innovations that will be patented. With the exception of section 5.1, the patent will be granted with certainty and be valid from the moment of application. Patent application is costless (or alternatively, in most cases, market profits can be understood as net of patenting costs). Due to the structure with discrete periods, there is the possibility that firms working on the same innovation during the same period will both be successful.[4] In such a case, each applicant will be awarded the patent with probability 1/2. The patent owner becomes monopolist on the market for the innovation, appropriating pi_i.

The probability with which a firm will be successful in a period is given by a fixed probability parameter α for firm A and β for firm B, with $0 < \alpha, \beta < 1$. The assumption on costs of R&D differs across sections, with one of two possible assumptions made: either the budget of R&D is fixed and exogenously given, and would expire if it wasn't used in the period it was intended for; or the budget could be used for other purposes, in which the decision to undertake R&D in any period comes at a cost c, assumed fixed and equal between firms and innovations. In other words, the opportunity cost of performing R&D are either zero or c.

[4]The majority of patent race models is constructed in a setting of continuous time, but assumes perfect information, at least regarding the patenting decision. Introducing imperfect information about the patenting decision that depends on a fixed time length shorter than the length of the race greatly complicates resulting expressions and is left for future analysis.

In cases in which $c > 0$, I assume that it is a flow cost and therefore could be stopped from occurring by stopping at any point during the period. To further simplify the analysis however I restrict attention to firms staying in the race for the whole duration of the period, *unless* they observe a competitor's patent at some point during that period. Generally, attention is paid to cases in which c is low enough to make it profitable for both firms to be on the market when the statutory delay is zero:

$$c < \alpha\pi \left(1 - \frac{b}{2}\right). \tag{3.1}$$

A patent application is published by the patent office after a statutory delay $0 < \ell_s < 1$, which in this discrete period setting implies that the maximum delay is one period. ℓ_s is an institutional variable chosen by policy. Starting in section 5, firms can choose to credibly announce a pending patent at some time before the patent office will officially publish it; in these sections therefore, the delay experienced by firm i is given by $\ell_{j \neq i}$, the delay with which her competitor decides to publish his pending patent, with $0 \leq \ell_j < \ell_s$.

4 Competition for a Single Innovation with Delayed Patent Publication

This section focuses on the competition for a single innovation which allows the patent holder to obtain monopoly market profits of π. R&D costs are assumed positive but low enough as above. A patent application filed at the end of period 2 is published by the patent office at a point ℓ_s during period 2. The decision that firms need to make is therefore simply whether to participate in each period or not. Since the patent race has a fixed and known start date and since R&D and spending in period 1 are entirely independent of that in period 2, ℓ_s does not affect a firm's decision to engage in R&D in period 1. A firm will participate if the costs are as specified in (3.1). The focus of the analysis therefore is on firms' participation decision in period 2. The participation decision would only need to be made in case the firm did not herself invent the innovation during period 1. In what follows, I restrict my attention to the decision made by firm A ("she"), but firm B's ("he") decision problem is entirely analogous.

At the beginning of period 2, if A was to observe a patent application by B, she knew the race was over and would cease to spend any further R&D resources on that specific innovation. If statutory delay was zero and she does not observe a patent, she would know that B did not invent the innovation and yet and accordingly the (per-period) profitability of the project did not change compared to one period before. If she does not observe a patent application in a patent system where $\ell_s > 0$, though, she needs to consider her expected profit. With probability β, firm B has been successful with his innovation attempt in period 1, and the associated patent application will be disclosed at point ℓ_s during period 2. Any resources spent on R&D in that case would have been essentially wasted. With probability $1 - \beta$, however, firm B is in the same situation as A, considering his participation in period 2. Now, should firm B find it profitable to participate in period 2, A's expected profit of participation is

$$\beta\left(-c \cdot \ell\right) + (1 - \beta)\left(\alpha\pi(1 - \beta/2) - c\right). \tag{3.2}$$

If instead firm B leaves the race after period 1, firm A has the opportunity to develop the innovation alone during period 2, should it not yet exist. Her expected profit in this case is

$$\beta\left(-c \cdot \ell\right)+(1-\beta)\left(\alpha \pi-c\right). \tag{3.3}$$

Now we can identify cases in which both, one, or no firm finds it profitable to participate in period 2. Combining this with the firms' decision to participate at the start of the race, we can find the equilibrium outcome of the game. The various parameter ranges are collected in table 3.2. The thresholds in period 2 are generally higher than in the corresponding case in period 1, so that we have cases in not all firms that start the race find it profitable to continue to period 2.

The cases are illustrated in figures 3.2 and 3.3. Each figure plots the participation decisions for each firm for different combinations of c and ℓ_s. This also allows to identify the different nature of the effects of c and ℓ_s. At the outset, both are increasing the expected cost of participating in the whole race. Increasing c, however, implies that low-ability firms will find it generally unprofitable to participate in any period, while there is another ability-threshold above which at least one firm will find it profitable to participate in period 2.

Increasing ℓ_s, on the other hand, leaves period-1 participation incentives unchanged, but it adds a second threshold (as a function of the competitor's R&D ability) above which even high-ability firms will not find it profitable anymore to compete in period 2. Participation constraints becomes such that they can intersect in an α, β diagram.

One immediate welfare effect of fewer firms participating in the race is that the probability of eventual R&D success is reduced. Such a result would likely also stem from a model that allowed choice of optimal R&D intensity with a convex R&D cost function.

Welfare properties of this effect that go beyond the probability would depend on which assumption we make about the population of firms in the economy. If the allocation of R&D ability is fixed, then completely excluding low-ability firms may not be optimal. A non-zero ℓ_s limits the amount of time that low-ability firms participate in the race, though.

Regarding the upper limit of ability beyond which it is not profitable for both firms to participate, it depends on the precise level. A not-too-long delay will still allow one firm to participate, while a too-high delay implies that neither firm finds it profitable to participate, even though each of them has a very high chance of success. The too-long delay may lead to innovations ending up not being developed even though there are firms that would be well equipped to do so.

5 Incentives to Disclose Pending Patents: Strategic Interaction

This section differs from the preceding one in that firms are now assumed to be able to choose the length of the period after which they make their pending patent application known to their competitor. The maximum delay that any firm can choose is given by the statutory delay ℓ_s, a policy variable, which here is assumed to be equal to 1. This means that unless a firm decides otherwise, a patent application will become publicly observable only with one period of delay, which in the present two-period setup is identical with the end of the game. Any individual $\ell_i < \ell_s$ chosen by firm i means that firm i's patent application will become publicly observable after a share ℓ_i of the period following patent application filing.

It needs to be noted that the analyses in this chapter depend on the assumption that there are no such thing as "fake" patent applications that either have firms submit applications that do not correspond to actual innovation results, or simply falsely claim that an application was submitted. Regarding the latter, it indeed carries a penalty to claim a pending patent if this is not in fact the case. The former problem is not an issue discussed in the practitioners' or academic literature, so it does not seem to be a widespread practice.

5.1 Sequential Innovation

The general case

In this subsection there are two innovations which, for technological reasons, can be only achieved in the pre-specified order that R&D on innovation 2 can not begin before innovation 1 has been invented. This resembles the basic case studied in Scotchmer & Green (1990). Other than that, the present model is unchanged compared to the preceding section with two discrete time periods with R&D results becoming available only at the end of each period, this implies that for innovation 2 to become available, at least one of the two participating firms must have achieved innovation 1 during period 1. If neither firm's R&D efforts were successful in period 1, they can still earn a positive profit from developing innovation 1 during period 2. If at the end of the game innovations 1 and 2 are owned by different firms, I assume that patent protection of the first innovation is sufficiently broad as to give the owner of that patent the chance to demand licensing fees from the producer of the second innovation. The broader is patent protection on the first patent, the greater is the profit share of innovation 2 that goes to the first innovator. It is precisely this potential for licensing revenue that will give rise to the incentive to disclose a pending patent.

Now, when there was no innovation success during period 1, the model in this section is equivalent to the model presented in the previous section: a firm that did not invent yet and that does not observe a competitor's patent when deciding whether or not to engage in R&D during the coming period will have to make the same participate-or-leave decision as it would need to make in absence of an innovation 2. In order to shift attention exclusively to the incentive to disclose that will become apparent in the subsequent sections, I assume for now that opportunity costs of R&D are zero, and that a firm therefore always participates in R&D.

This has the following consequences: if firm A did not invent innovation 1 during period 1, then she will try again in period 2, unless she observes a patent by firm B, in which case she will shift her resources to innovation 2. The assumption here is that the patent holder discloses enough information to enable the competitor to start working on innovation 2. This may be either by disclosing the technical information via a very thorough description section in the patent application document, or by directly making technical documents available. After all, enabling the competitor to work on the second innovation is the very reason the patent applicant discloses the information in this case. Since such disclosure would happen after the filing date of the patent application, it would not adversely affect the chances of the patent being granted due to lack of novelty compared to the "prior art".

Re-attempting innovation 1 in period 2 yields an expected profit of $(1-\beta)\alpha\pi_1$, while switching to innovation 2 with the help of information disclosed by the competitor yields $\alpha(1-\beta/2)\pi_2(1-\lambda)$. $0<\lambda<1$ denotes the share of monopoly profits that need to be paid to the patent holder of innovation 1. I will now give each firm the possibility to react to the disclosure of a patent application by their competitor by abandoning their running R&D projects that are affected by that patent. For this reason I introduce the ability of shifting resources during a period: if after period ℓ_B firm A observes a patent application by B, she will then shift her R&D efforts to innovation 2 for the remainder $1-\ell_B$ of the period. Her expected profit then is $(1-\ell_B)\alpha(1-\beta/2)\pi_2(1-\lambda)$. The decision making of firm B is exactly analogous.

While R&D success remains to become available at fixed discrete moments in time, the patent disclosure and R&D investment decisions are made in continuous time. In absence of a Poisson arrival process, this setting might give rise to time inconsistency where firms would prefer to change their R&D decision before the end of a period. The assumption of a small switching cost would be sufficient to prevent this from happening, while being small enough compared to the value of the innovation itself to be disregarded in modelling. Furthermore, results obtained in this chapter do not depend on this particular setting as the models have exclusively corner solutions where $\ell=0$ or $\ell=1$. I have chosen this approach because it yields objective functions that explicitly depend on the disclosure delay.

In case firm A did invent innovation 1 during period 1, it will become monopolist in the market for innovation 1 unless the patent will be awarded to firm B. The granting of the patent is entirely independent of the fact whether or not the applicant firm made the application publicly available before it was published by the patent office.

An important observation is the following: at the beginning of period 2, each firm can be in precisely two situations: having developed innovation 1, or having developed nothing. The only decision it has to make is about when to announce its pending patent application. This decision is only meaningful when it in fact has developed innovation 1 so far. A firm's R&D allocation for period 2, on the other hand, can only be influenced by its competitor's patent announcement when that firm has not yet invented innovation 1. Therefore, the only case in which a decision has to be made is when the competitor does not have to make any decision, and the equilibrium will be one in dominant strategies, where firm A's choice of ℓ_A does not affect firm B's choice of ℓ_B, and vice versa.

Firm A therefore chooses optimal publication delay ℓ_A by maximising the following expected profit equation:

$$\begin{aligned} E(\pi^A) = {} & \frac{\beta}{2}\left(\pi_1+\alpha\left(1-\frac{\beta}{2}\right)\pi_2(1-\lambda)+(1-\alpha)\,\beta\pi_2\lambda\right) \\ & +(1-\beta)\left(\pi_1+\alpha\left(1-\frac{\beta(1-\ell_A)}{2}\right)\pi_2+(1-\alpha)\beta(1-\ell_A)\pi_2\lambda\right) \end{aligned} \tag{3.4}$$

The first line gives the expected profit in case B invented during period 1, which happens with probability β. In that case, both firms compete for the first patent, the profit of which is therefore divided by 2, as is any profit from innovation 2 which equally depends on the allocation of the first patent. This part however is independent of the choice of patent disclosure. The second line is expected profit when B did not invent yet, in which case he will "waste" ℓ_A of his per-period budget on already-invented innovation 1. The first derivative with respect to ℓ_A is

$$(1-\beta)\beta\pi_2\left(\frac{\alpha}{2}-(1-\alpha)\lambda\right),$$

which has a negative sign if $\alpha < 2\lambda/(1+2\lambda)$, and a positive one otherwise. A negative sign implies an optimal choice of $\ell_A = 0$, which means that A will disclose her pending patent only if her R&D ability is low enough in relation to possible licensing revenues. The higher are those latter, the greater is the range of α for which A will find it optimal to disclose and thereby get B's "help" in obtaining at least a share of π_2.

Discounting

Every "period" in the present game is supposed to represent a sizeable amount of time, enough for an R&D project to have a reasonable chance of completion. As such, a firm that was successful in period 1 may not want to wait with commercialisation of innovation 2 until the end of one more period. It seems reasonable to believe that any market activity will be observable by the competitor and have an effect equivalent to that of patent publication. This means that to keep a pending patent undisclosed, a firm has to abstain from market activity involving the patented innovation. Delaying disclosure therefore implies delaying revenue. After all, commercialisation could be regarded as a particularly credible way to reveal innovation success.

To capture this greater incentive to disclose, equation (3.4) will be modified regarding the two terms of π_1: its second occurrence (when only A invented so far, in line 2) will be divided by $(1+r)^{\ell_A}$, the first one (when both firms invented in period 1) will be divided by $(1+r)^{\mathrm{Min}\{\ell_A,\ell_B\}}$. This latter piecewise definition of the profit function captures the fact that A can safely stop "hiding" her innovation when she observes a patent application by B on the same innovation.

The derivative of this modified profit function is

$$\frac{\partial\pi_A}{\partial\ell_A} = \begin{cases} \dfrac{\mathrm{Log}\left[\dfrac{(2-\beta)\pi_1\mathrm{Log}[1+r]}{(1-\beta)\beta(\alpha-2\lambda(1-\alpha))\pi_2}\right]}{\mathrm{Log}[1+r]} & \text{if } \ell_A \leq \ell_B \\[2ex] \dfrac{\mathrm{Log}\left[\dfrac{2\pi_1\mathrm{Log}[1+r]}{\beta(\alpha-2\lambda(1-\alpha))\pi_2}\right]}{\mathrm{Log}[1+r]} & \text{if } \ell_A > \ell_B. \end{cases} \tag{3.5}$$

While now the derivative does include ℓ_A and it is possible to solve for it, the sign of the second derivative is strictly negative, and accordingly only one of the two two corner solutions involving maximum or minimum publication delay can be profit-maximising.

The sign of the derivatives is positive, and delay therefore maximal at $\ell_A = 1$, if, as without discounting, $\alpha > 2\lambda/(1+2\lambda)$, and additionally

$$\pi_2 > \begin{cases} \dfrac{(2-\beta)r\pi_1}{(1-\beta)\beta(1+r)(\alpha-2(1-\alpha)\lambda)} & \text{if } \ell_A \leq \ell_B \\[2ex] \dfrac{2r\pi_1}{\beta(1+r)(\alpha-2(1-\alpha)\lambda)} & \text{if } \ell_A > \ell_B. \end{cases}$$

With discounting, there will additionally be disclosure at the beginning of period 2 if the profit from innovation 2 is not "worth waiting for". Disclosure instead is immediate, $\ell_A == 0$, when either condition above is violated. Without licensing revenue ($\lambda = 0$) there was previously no incentive for a firm to disclose a patent early; with discounting it is now optimal to disclose the patent at the beginning of period 2 when the profit from innovation 2 is low enough, i.e., if

$$\pi_2 < \begin{cases} \dfrac{(2-\beta)r\pi_1}{\alpha(1-\beta)\beta(1+r)} & \text{if } \ell_A \leq \ell_B \\[2ex] \dfrac{2r\pi_1}{\alpha\beta(1+r)} & \text{if } \ell_A > \ell_B. \end{cases}$$

Complementarity between the first and the second innovation

While there was a sort of "one-sided" complementarity in the preceding section by which possession of innovation 1 allowed to receive a share of innovation 2's profit via licensing, this section is dealing with "proper" complementarity in that combining both innovations allows reaping an additional profit that makes total profits greater than the simple sum of the stand-alone profits of each innovation. Following Denicolò & Zanchettin (2018) I am using a continuous parameter to measure the extent of complementarity. Total profit is given by π, of which a share $1-\gamma$ comes from stand-alone profits and a share γ requires both innovations to be available. Of the stand-alone profits, a share δ is contributed by innovation 1 and the remainder $1-\delta$ by innovation 2. When the two innovations are contributed by different firms, I assume that each receives half of the complementarity profit share γ. The greater is γ, the stronger is the degree of complementarity between the two innovations.

A's expected profit in this case is given by

$$\begin{aligned} E(\pi^A) = \beta\pi \Bigg[&(1-\gamma)\left(\frac{\delta}{2} + (1-\delta)\alpha\left(1-\frac{\beta}{2}\right)\right) \\ &+ \gamma\left(\frac{1}{2\cdot 2}(1-(1-\alpha)(1-\beta)) + \frac{1}{2}\alpha\left(1-\frac{\beta}{2}\right)\right)\Bigg] \\ &+ (1-\beta)\Bigg[(1-\gamma)\left(\delta + (1-\delta)\left(1-\frac{\beta(1-\ell_A)}{2}\right)\right) \\ &+ \gamma\left(\frac{1}{2}(1-(1-\alpha)(1-\beta(1-\ell_A)) + \frac{\alpha}{2}\left(1-\frac{\beta(1-\ell_A)}{2}\right)\right)\Bigg], \end{aligned} \tag{3.6}$$

The square brackets distinguish the two basic situations, the one in which B also developed innovation 1 (lines 1 and 2) from that were B was not successful in innovation (lines 3 and 4). One half of the complementarity profit γ is obtained by each patent-holding firm, and non-zero profit from innovation 1 crucially requires the development of innovation 2 to be successful.

The first derivative is

$$\frac{(1-\beta)\pi}{2}\left((1-\gamma)(1-\delta)\alpha\beta+\gamma\left(\frac{\alpha\beta}{2}-(1-\alpha)\beta\right)\right).$$

This derivative has a negative sign, and A accordingly discloses her patent at the beginning of period 2, if $\alpha < 2\gamma/(2+\gamma-2\delta(1-\gamma))$. The greater is the degree of complementarity and/or the more important is the not-yet-invented second innovation, the smaller is the range of α for which A will disclose her patent. This result once more has relatively weaker firms disclose their pending patents, and relatively stronger firms not disclose.

Incomplete information

So far the disclosure decision was a function of the R&D strength of the disclosing firm. This makes it seem as if it should also be possible to infer a firm's type from observing their patenting behaviour, conditional on having innovated in period 1. This is the subject of the present section.

I will revert to the basic sequential model without complementarity or discounting, where with certain patent rights a firm never has an incentive to disclose a pending patent as the only effect this would have is to allow an unsuccessful competitor to potentially take away the patent rights to innovation 2. However, in the above model with zero opportunity costs of R&D, the information about the competitor's R&D ability is not affecting decisions. As shown in section 4, even in a very simple setting the existence of non-zero opportunity costs gives rise to relatively complex continuation decisions. In this section, therefore, there is still complete information about each firm's R&D ability parameter. What is uncertain however is the quality of the filed patent application.

Up until now, each patent application was granted with certainty. In the following, however, each patent filed can be of high (HQ) or of low quality (LQ). With probability p, a firm's patent application is of low quality, and with $1-p$ it is of high quality. The difference is in terms of granting probability: a HQ patent will be granted with certainty, while a LQ patent will be granted only with probability $0 < g < 1$. The difference in patent quality can be thought of as a technical feature of the innovation that can only be observed once the research process is complete.[5] As before, in case of grant the decision of the patent office is final and protection is perfect, i.e., imitation is not possible. A non-zero chance of rejection by the patent office, however, can make it worthwhile for the so-far unsuccessful competitor to imitate innovation 1 in period 2 even though innovation 2 has become feasible to develop via patent disclosure. It is precisely this case that will give rise to a non-obvious disclosure decision.

[5] Alternatively, it would probably also be possible to give firms the choice of choosing between a costly HQ application and a cost-free (or less costly) LQ application, but a firm might be able to infer this decision when knowing α and β.

I make the following six additional assumptions, mostly for simplicity. The result as such should also arise when dropping them, though the cases would be less obvious. First, in case of disclosure the type of patent can be observed, so that uncertainty about the potential patent type is restricted to the case of non-disclosure; and second, patent quality will be the same for both firms in case they are competing for the patent on innovation 1. Dropping this assumption would require introducing an assumption about what happens if firm B applies for a HQ patent on top of an (already granted or rejected) LQ patent by firm A.

Third, by disclosing a patent the competitor's cost of imitating the described innovation are zero. This effect is independent of the patent's quality. Fourth, when imitation occurs and the LQ patent is rejected, the two firms become duopolists in the market for innovation 1, each getting a profit share of πd, where $0 < d < 1/2$. $d = 0$ represents perfect (price) competition, while with $d = 1/2$ the firms manage to collude in sharing the monopoly profit. Fifth, the patent for innovation 2 is always of high quality and will be granted with certainty. Sixth, after period 2 the game ends and no imitation can occur even if a patent was not granted, i.e., if firm B does not develop innovation during period 2, then firm A will be the monopolist on the market for innovation 1 even if her LQ patent is not granted.

Similar to Ganglmair & Oh (2014), by disclosing a pending patent application a firm can influence its competitor's decision by affecting its belief about the type of the patenting firm. Ganglmair & Oh look at the case of a single process innovation in a setting with a clear leader-follower distinction where uncertainty on the side of the follower firm is about whether the leader has in fact innovated, which would affect the follower's expected profit from innovating himself. In the present model, uncertainty is not only about existence of the innovation, but additionally about its "quality", while the competitor firm is not restricted to being a follower but can have innovated himself at the time the disclosure decision is made.

Zero opportunity costs of R&D I start analysis by examining the case of zero opportunity costs in which the competitor never exits. In this case, when B was successful in period 1, he will in any case use period 2 to attempt and develop innovation 2. When B was not successful, his decision depends on whether and what type of patent A discloses. If A discloses a HQ patent, B's only chance for non-zero profit will be innovation 2. If A discloses a LQ patent, B can choose between imitating at zero cost and getting an expected profit of $\pi_1 d(1-g)$, and attempting development of innovation 2 for an expected profit of $\pi_2\beta(1-\alpha/2)$. When A does not disclose a patent, B's only chance of positive profits is to use period 2 to try again and develop innovation 1. If A did in fact not invent yet or if her LQ patent is not granted protection, then this possibility is still open for B. In absence of opportunity costs, even with the lowest non-zero chance of success he will therefore go and work on innovation 1 during period 2.

In the case on non-disclosure, B can in fact update his belief about the type of A's potentially undisclosed innovation, but since this will not affect his R&D decision for period 2, this is not further studied here. To sum up, given B was unsuccessful during period 1, then when disclosing a HQ patent he will work on innovation 2 during period 2, when not disclosing he will work on innovation 1, and when disclosing a LQ patent he will work on innovation 2 only if $\pi_1 < (\pi_2(2-\alpha)\beta)/(2d(1-g))$, that is, when innovation 2 is relatively more attractive given the context of the other parameters.

Turning to A, she will get the same expected profit independent of her disclosure decision if B has in fact already developed innovation 1:

$$\beta\left[\pi_1\left(\frac{g}{2}+(1-g)d\right)+\pi_2\alpha\left(1-\frac{\beta}{2}\right)\right].$$

The difference in expected profits therefore once again stems from the case in which B is without innovation success yet. When A has a HQ patent, we are effectively in the case of section 5.1 with $\lambda = 1$, so that A has no incentive to disclose her pending patent as this would only reduce her chance of getting innovation 2 without affecting her expected profit from innovation 1. When A has a LQ patent instead, expected profits differ. When she does not disclose it, she anticipates B working on innovation 1 during period 2, giving her an expected profit of

$$(1-\beta)\Big(\beta(\pi_1(g+1-g)d)+\pi_2\alpha\Big)+(1-\beta)\Big(\pi_1+\pi_2\alpha\Big).$$

Her expected profit from disclosing depends on B's reaction. If B subsequently imitates innovation 1, A expects

$$(1-\beta)\Big(\pi_1(g+(1-g)d)\Big),$$

while if B takes the opportunity to move on to innovation 2 A expects

$$(1-\beta)\left(\pi_1+\pi_2\alpha\left(1-\frac{\beta}{2}\right)\right).$$

Comparing the above profit equations, it becomes obvious that A does not have an incentive to disclose her LQ patent if in consequence B will still go for innovation 1, as thereby she effectively is only supporting him in his imitation attempt. When disclosure instead makes B go for innovation 2, A will find it worthwhile to disclose if $\pi_1 > (\pi_2\alpha)/(2(1-d)(1-g))$. Combining this threshold with the one affecting B's decision, I find that A will disclose her LQ patent if

$$\frac{\pi_2\alpha}{2(1-d)(1-g)} < \pi_1 < \frac{\pi_2(2-\alpha)\beta}{2d(1-g)}.$$

The value of innovation 1 must not be too low to give firm A an incentive to protect it, but it must also not be too high as otherwise firm B will not be willing to forgo the chance of duopoly profits. Since innovation values are variables with relatively arbitrary values, if we assume both values to be 1, the above condition becomes $g < (2(1-d)-\alpha)/(2(1-d))$ and $\beta > (2d(1-g)/(2-\alpha)$. Disclosure happens more often the greater is β (because B will be more likely successful when attempting innovation 2), but becomes happens less often when α is high (because A then puts more emphasis on innovation 2) or when d is high (because then B's imitation hurts less). The grant probability may neither be too high (since this would reduce the imitation threat to firm A) nor too low (as this would make imitation too attractive for B).

Positive opportunity costs of R&D The preceding section did have a more nuanced result than the complete-information cases studied before, but it did not depend on B's belief about A's innovation type. This changes once we (re-)introduce positive opportunity costs of R&D c. Again it is no longer true that B will always continue with R&D, since when the success probability is too low it will not be worthwhile to engage in R&D at all. What remains true is that B's decision is unaffected by A's disclosure decision when B has innovated himself, and that there is no case in which A would disclose a HQ patent application since there is no threat from an imitation attempt by B. The conditional probability of disclosure when A has a HQ patent is therefore zero: $Prob(\ell_A = 0|HQ) = 0$.

B's optimal decision changed slightly for each case. When A discloses a LQ patent, imitation is still costless, but attempted development of innovation 2 is not, which now incurs a cost of c. The condition on the maximum profitability of innovation 1 is unchanged at $\pi_1 < (\pi_2(2-\alpha)\beta)/(2d(1-g))$, but additionally now, costs may not be too high to make innovation 2 an ex-ante profitable project: $c < ((2-a)\beta\pi_2 - 2d(1-g)\pi_1)/2$.

When A does not disclose a patent, now B's belief about A's state becomes important, as a too high probability of a non-disclosed HQ patent will make it unprofitable to spend c on innovation 1. B therefore forms his belief about whether A has in fact filed a LQ patent application by Baye's rule as follows:

$$Prob(LQ|\ell_a = 1) = \frac{Prob(\ell_A = 1|LQ) \cdot Prob(LQ)}{\begin{aligned}&Prob(\ell_A = 1|LQ) \cdot Prob(LQ)\\ &\quad + Prob(\ell_A = 1|HQ) \cdot Prob(HQ) + Prob(\ell_A = 1|NI) \cdot Prob(NI)\end{aligned}},$$

where NI designates the case of no innovation by A in period 1. Obviously, $Prob(\ell_A = 1|NI) = 1$, as absent of innovation there is nothing to be disclosed. I denote the third conditional probability $Prob(\ell_A = 1|LQ)$ by q. The unconditional probabilities are defined as before with $Prob(NI) = 1-\alpha$, $Prob(LQ) = \alpha p$ and $Prob(HQ) = \alpha(1-p)$. B's posterior probability that A has filed a LQ patent application is therefore

$$\widehat{\alpha p} = \frac{q\alpha p}{q\alpha p + 1 \cdot \alpha(1-p) + 1 \cdot (1-\alpha)} = \frac{q\alpha p}{1 - \alpha p(1-q)}. \tag{3.7}$$

Posterior probabilities $\widehat{\alpha(1-p}$ and $\widehat{1-\alpha}$ are computed analogously.

Based on these posterior probabilities, B makes his decision on what to do in period 2 in case there was no disclosure. Innovation 2 remains unachievable, but the case for re-attempting innovation 1 looks less profitable now as there is a certain probability that A has already applied for a patent on it. When this probability is above some threshold, B now prefers to exit the race. This is the case when

$$c < \frac{\pi_1(\beta(2-\alpha)(1-\alpha) + 2\alpha d(1-g)pq}{2(1-ap(1-q))},$$

which may or may not be a stricter condition than the one in the case of disclosure.

Anticipating the effect of her disclosure choice on B, A decides between disclosure and non-disclosure. In fact, she could decide on a mixed strategy involving $0 < q < 1$, but as before such a strategy will always be dominated by one of the two pure strategies $q = 0$ and $q = 1$. She can accordingly choose between two alternatives, each of which can have two consequences (see table 3.1). Again, there is a component of expected profit that corresponds to the case in which B has innovated already and that is equal across cases and therefore ignored. The expressions in the third column of the table correspond to the case of so-far unsuccessful B and therefore would need to be multiplied by $1-\beta$.

denoted by	description	expression
π_A^I	disclose and B will work on innovation 1	$\pi_1(g+(1-g)d)+\pi_2\alpha-c$
π_A^{II}	disclose and B will work on innovation 2	$\pi_1+\pi_2(\alpha(1-\beta/2)-c$
π_A^{III}	do not disclose and B will work on innovation 1	$\pi_1(\beta(g+(1-g)d)+(1-\beta))$ $+\pi_2\alpha-c$
π_A^{IV}	do not disclose and B will exit	$\pi_1+\pi_2\alpha-c$

Table 3.1: A's choices and consequences under incomplete information with non-zero c

2×2 cases give rise to four comparisons. First, comparing π_A^I to π_A^{III} it is obvious that when disclosure is without effect, A never prefers disclosure. Second, whenever non-disclosure leads to exit (i.e., comparing π_A^I and π_A^{II} to π_A^{IV}), A will again never choose disclosure. The only case in which A potentially prefers disclosure is when disclosure induces B to attempt innovation 2 instead of innovation 1. Disclosure is preferred when $\pi_1>\pi_2\alpha/(2(1-g)(1-d))$, which is unchanged compared to the previous case with $c=0$.

We need to compare this condition with the ones describing the case in which B will indeed switch from innovation 1 to innovation 2 if a LQ patent is disclosed. In order to reduce the degrees of freedom, I consider results when $\pi_1=\pi_2=1$. Now, A prefers disclosure when either $g\leq 1-\alpha$, or when $d<(2(1-g)-a)/(2(1-g))$. The results, including the conditions required to ensure that B does in fact switch to innovation B when A discloses a LQ patent, are presented in table 3.3. The first condition relates α and g: one of the two variables needs not be too high compared to the other one, implying that either A will find it difficult to get π_2 by own innovation success, or she will unlikely get monopoly profits from innovation 1. Either way, protecting innovation 1 becomes important.

The second condition concerns d, which, if low, implies that imitation is particularly costly for A. The third condition is regarding the value of p, B's prior about whether A has a LQ innovation. If p is high, then B will have a particularly strong incentive to attempt innovation 1 in case of non-disclosure. The final condition relates β and c, where a high value of β increases the probability that B would be able to develop innovation 1 under non-disclosure, or is even already in possession of innovation 1, rendering the disclosure decision unimportant.

5.2 Complementary Simultaneous Innovations

In contrast to section 5.1, the innovations now do not technically build on each other, meaning that a firm can start working on either one as well as decide to work on both innovations simultaneously. This is similar to the case studied by Ménière (2008). With decreasing returns to R&D and absent the opportunity to collude in R&D, it is in a firm's best interest to start working on both innovations simultaneously.

With statutory delay: optimally choosing disclosure

To reduce notation clutter I assume that $\pi_1 = \pi_2 = 1$, i.e., the total profit by commercialising the two complementary innovations is 2, and each firm receives 1 if each innovation is contributed by a different firm. Furthermore I assume zero stand-alone profit, which corresponds to the case in section 5.1 when $\gamma = 1$, and that $c = 0$, implying here that a firm optimally does R&D on those components that are not yet patented. This allows a firm to have a clear idea of what their competitor is currently most likely working on without having to consider the optimal allocation of R&D first.

Splitting a budget of fixed size on two projects instead of focusing efforts on a single project, as before, makes it seem reasonable that the individual success probability of each of the two parallel projects is lower. This is implemented by assuming that each firm has a per-period success probability of α, β when doing a single project, but a per-period-and-project success probability of $\alpha v, \beta v$ when working on both innovations during the same period, where $1/2 < v < 1$ implies decreasing returns to R&D, a reasonable assumption at the project level.

As before, when A does not have an innovation ready yet, there is no meaningful decision to be made. When A has invented both innovations after period 1 already, she will be indifferent between disclosing or not those two patent applications since the R&D race itself is over.[6] The interesting case once more is when A has invented one of the two innovations.

Her competitor B can now be in four different situations: i) he has invented both components, which will have happened with probability $(\beta v)^2$; ii) he has invented a single component, which is the same component as A's (probability $\beta v(1 - \beta v)$); iii) his single component is the other one of the two (probability $\beta v(1 - \beta v)$); and iv) he has not developed any innovation yet (probability $(1 - \beta v)^2$).

In each case, A expects a different profit: in i), the game will have ended, and A will compete with B for the patent on one of the two components ($E(\pi_A) = 1/2$); in ii), the two firms also will compete for the same patent on the first component, but due to the innovations' complementarity will not obtain a positive profit unless the other component has been developed as well, which will happen with probability $(1-(1-\alpha)(1-\beta))$ during period 2 ($E(\pi_A) = (1 - (1 - \alpha)(1 - \beta))/2 + \alpha(1 - \beta/2)$); and in iii) the game is also over and each firm gets the profit share allocated to the innovation that they contribute ($E(\pi_a) = 1$).

In iv), finally, A's disclosure decision is playing a role. If she discloses her single patent application, this will make B optimally focus his R&D efforts on the remaining innovation, while if she does not disclose it, B will continue to spend on both innovations during period 2, reducing his success probability on innovation 2 by the factor v. A needs to decide whether she prefers a stronger or a weaker firm B working on innovation 2. During the time interval $[0, \ell_A]$, B will have a success probability of βv, while during the remainder of $(\ell_A, 1]$ he will succeed with probability β. His average success probability is accordingly $\beta(1 - \ell_A(1 - v))$. A chooses once again the optimal value of $0 \leq \ell_A \leq 1$ to

[6] Of course, we can assume that the firms will move on to the next comparable R&D race, in which A may have an incentive to not disclose her patents if this gives her a head-start in the subsequent race.

maximise the expression

$$E(\pi_A) = (1 - (1-\alpha)(1 - \beta(1 - \ell_A(1-v))) + \alpha\,(1 - \beta(1 - \ell_A(1-v))/2)$$

The derivative with respect to ℓ_A is

$$(1-\beta v)^2\,[\alpha\beta(1-v)/2 - (1-\alpha)\beta(1-v)]\,,$$

which has a negative sign if $\alpha < 2/3$, and a positive sign otherwise. Unless A is very strong, she prefers to disclose her pending patent immediately in order to make a potentially unsuccessful rival assist her in securing positive profit on her existing first innovation.

6 Conclusion

The goal of the present chapter is to draw attention to the potential impact of the statutory delay in patent publication on R&D competition, and to competing firms' incentives to voluntarily disclose pending patents. Generally, delayed notification about the end of a patent race yields to a period of R&D spending by all non-winners that has zero return and is therefore considered not only duplicative, but wasted. Increased expected participation cost raise the participation constraint and prevent participation by firms with R&D ability below some threshold of success probability. If policy aims at increasing the ability of participating firms, the publication delay is much more easily adjusted than the variable cost of performing R&D.

Section 4 shows that the precise effect of an increase in publication delay can however affect firms' participation decisions differently compared to an increase in per-period cost. Increasing the delay leaves period-1 participation unchanged, while restricting participation by a firm if the competitor's success probability is *above* some threshold. In an extreme case, this may lead to two high-ability firms mutually preferring to exit the race if the probability that the respective competitor already invented is too high, potentially leaving the innovation undeveloped despite having high-ability firms equipped to do so. Outside of this extreme case, participation is restricted to a single high-ability firm only.

In any case, though, the total probability of innovation success is reduced (or, alternatively, the time until expected development is increased) compared to the case with a zero delay in publication. A non-zero delay may still be chosen by policy for other reasons, e.g., if individually optimal R&D investment in a patent race is too high compared to the social optimum and the alternative policy response of lowering the strength of patent protection is too hard to affect. In a model with free entry and exit of firms both into the race and into the economy, the location of participation constraints may also be used to affect the distribution of R&D strength among firms in the economy. In a model of purely cumulative both otherwise independent innovation, a certain delay in patent publication may also be optimal if this yields the innovator a certain lead time in development of the subsequent innovation stage, potentially overcoming the "Arrow effect" (Arrow, 1962; Reinganum, 1983) by which a monopolist incumbent has a lower incentive to engage in innovation than a potential entrant.

In sections 5.1 and 5.2 I look into examples of ways in which the existence of a statutory delay gives rise—or not—to incentives to use the delay in order to influence competitors'

R&D behaviour. In the models presented there is a tradeoff between receiving profit due to own innovation success and profit due to the competitor's innovation success. In all cases, there exists a threshold value of R&D success probability below which a firm finds it optimal to announce the fact that it has been successful with a certain innovation already. Such an incentive to disclose could become even stronger when combined with weaker firms' lower incentive to enter the second period of R&D, as identified in section 4. It needs to be kept in mind that these strategic considerations would not be possible if patent publications were immediately observable. In that case, there could an incentive to delay the filing of the patent application itself, even though this incentive is arguable much weaker. I am currently working on showing the existence of this incentive. Table 3.4 collects all cases for which I found that disclosure is preferred that are not yet included in table 3.3.

Finally, the existence of a statutory delay also gives rise to signalling opportunities. I look at the case of differences in the probability of a patent grant, but one can think of ample other cases in which incomplete information can affect the decision to disclose or not a pending patent. One example is studied in Ganglmair & Oh (2014), but one might think that any case of signalling via the patenting decision itself (with the alternative of keeping an innovation secret indefinitely) would also apply in the situation in which a patent application itself can be kept secret for a while.

References

Akcigit, U., & Liu, Q. (2016). The Role of Information in Innovation and Competition. *Journal of the European Economic Association, 14*(4), 828–870.

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

Aoki, R., & Spiegel, Y. (2009). Pre-grant patent publication and cumulative innovation. *International Journal of Industrial Organization, 27*(3), 333–345.

Arrow, K. J. (1962). Economic Welfare and the Allocation of Resources to Innovation. In R. Nelson (Ed.), *The Rate and Direction of Economic Activity.* New York, USA: National Bureau of Economic Research.

Asay, C. D. (2015). The Informational Effects of Patent Pledges. In Jorge L. Contreras & Meredith Jacob (Eds.), *PATENT PLEDGES - GLOBAL PERSPECTIVES ON PATENT LAW'S PRIVATE ORDERING FRONTIER.* New York: Edward Elgar.

Barzel, Y. (1968). Optimal Timing of Innovations. *The Review of Economics and Statistics, 50*(3), 348–355.

Biagi, A., & Denicolò, V. (2014). Timing of discovery and the division of profit with complementary innovations. *Journal of Economics & Management Strategy, 23*(1), 89–102.

Choi, J. P. (1998). Patent Litigation as an Information-Transmission Mechanism. *The American Economic Review, 88*(5), 1249–1263.

Clark, D. J., & Konrad, K. A. (2008). Fragmented Property Rights and Incentives for R&D. *Management Science, 54*(5), 969–981.

Comino, S., & Graziano, C. (2015). How many patents does it take to signal innovation quality? *International Journal of Industrial Organization, 43*, 66–79.

Conti, A., Thursby, J., & Thursby, M. (2013). Patents as Signals for Startup Financing. *The Journal of Industrial Economics, 61*(3), 592–622.

Crampes, C., & Langinier, C. (1998). Information Disclosure in the Renewal of Patents. *Annales d'Économie et de Statistique*, (49/50), 265–288.

Denicolò, V. (1996). Patent Races and Optimal Patent Breadth and Length. *The Journal of Industrial Economics, 44*(3), 249–265.

Denicolò, V. (1999). The optimal life of a patent when the timing of innovation is stochastic. *International Journal of Industrial Organization, 17*(6), 827–846.

Denicolò, V. (2000). Two-Stage Patent Races and Patent Policy. *The RAND Journal of Economics, 31*(3), 488.

Denicolò, V. (2008). Economic Theories of the Nonobviousness Requirement for Patentability: A Survey. *Lewis & Clark Law Review, 12*, 443–459.

Denicolò, V., & Franzoni, L. A. (2004). Patents, Secrets, and the First-Inventor Defense. *Journal of Economics & Management Strategy, 13*(3), 517–538.

Denicolò, V., & Franzoni, L. A. (2010). On the Winner-Take-All Principle in Innovation Races. *Journal of the European Economic Association, 8*(5), 1133–1158.

Denicolò, V., & Halmenschlager, C. (2012). Optimal patentability requirements with complementary innovations. *European Economic Review, 56*(2), 190–204.

Denicolò, V., & Zanchettin, P. (2018). Some Simple Economics of Patent Protection for Complex Technologies. *CEPR Discussion Paper No. DP13087.*

Ergenzinger, E. R. Jr. (2006). The American Inventor?s Protection Act: A Legislative History. *Wake Forest Intellectual Property Law Journal, 7*, 145–173.

Fershtman, C., & Kamien, M. I. (1992). Cross licensing of complementary technologies. *International Journal of Industrial Organization, 10*(3), 329–348.

Graham, S., & Hegde, D. (2015). Disclosing patents? secrets. *Science, 347*(6219), 236–237.

Ganglmair, B., & Oh, J.-M. (2014). Strategic Secrecy of Pending Patents. *Mimeo*, The University of Texas at Dallas. Retrieved from http://ftp.cemfi.es/ftp/pdf/papers/wshop/PendingPatents_Ganglmair_Oh.pdf

Green, J., & Scotchmer, S. (1995). On the Division of Profit in Sequential Innovation. *RAND Journal of Economics, 26*(1), 20–33.

Hall, B. H., & Harhoff, D. (2012). Recent Research on the Economics of Patents. *Annual Review of Economics, 4*(1), 541–565.

Harris, C. J., & Vickers, J. S. (1987). Racing with Uncertainty. *Review of Economic Studies, 54*, 1–21.

Hegde, D., & Luo, H. (2017). Patent Publication and the Market for Ideas. *Management Science, 64*(2), 652–672.

Henkel, J., & Jell, F. (2010). Patent Pending – Why Faster Isn't Always Better. *Mimeo*, Technical University Munich. Retrieved from https://papers.ssrn.com/abstract=1738912

Hopenhayn, H. A., & Squintani, F. (2016). Patent Rights and Innovation Disclosure. *The Review of Economic Studies, 83*(1), 199–230.

Horstmann, I., MacDonald, G. M., & Slivinski, A. (1985). Patents as Information Transfer Mechanism: To Patent or (Maybe) Not to Patent. *Journal of Political Economy, 93*(5), 837–858.

Hottenrott, H., Hall, B. H., & Czarnitzki, D. (2016). Patents as quality signals? The implications for financing constraints on R&D. *Economics of Innovation and New Technology, 25*(3), 197–217.

Hunt, R. M. (2004). Patentability, Industry Structure, and Innovation. *The Journal of Industrial Economics, 52*(3), 401–425.

Kim, B., Kim, E., Miller, D. J., & Mahoney, J. T. (2016). The impact of the timing of patents on innovation performance. *Research Policy, 45*(4), 914–928.

Kim, J. (2017). Racing against Time in Research: A Study of the 1995 U.S. Patent Law Amendment. *IZA Discussion Paper No. 10815.*

Kultti, K., Takalo, T., & Toikka, J. (2006). Simultaneous Model of Innovation, Secrecy and Patent Policy. *American Economic Review, 96*(2), 82–86.

Kwon, I. (2012a). Patent Races With Secrecy. *The Journal of Industrial Economics, LX*(3), 499–516.

Kwon, I. (2012b). Patent Thicket, Secrecy, and Licensing. *The Korean Economic Review, 28*(1), 27–49.

Langinier, C. (2005). Using patents to mislead rivals. *Canadian Journal of Economics/Revue Canadienne d?économique, 38*(2), 520–545.

Lee, T., & Wilde, L. L. (1980). Market Structure and Innovation: A Reformulation. *The Quarterly Journal of Economics, 94*(2), 429–436.

Long, C. (2002). Patent Signals. *The University of Chicago Law Review, 69*(2), 625–679.

Loury, G. C. (1979). Market Structure and Innovation. *The Quarterly Journal of Economics, 93*(3), 395–410.

Ménière, Y. (2008). Patent law and complementary innovations. *European Economic Review, 52*(7), 1125–1139.

O?Donoghue, T., Scotchmer, S., & Thisse, J.-F. (1998). Patent Breadth, Patent Life, and the Pace of Technological Progress. *Journal of Economics & Management Strategy, 7*(1), 1–32.

Ottoz, E., & Cugno, F. (2008). Patent–Secret Mix in Complex Product Firms. *American Law and Economics Review, 10*(1), 142–158.

Palangkaraya, A., Jensen, P. H., & Webster, E. (2008). Applicant behaviour in patent examination request lags. *Economics Letters, 101*(3), 243?245.

Ponce, C. J., & Henry, E. (2011). Waiting to Imitate: On the Dynamic Pricing of Knowledge. *Mimeo*, Sciences Po Paris and Universidad Carlos III Madrid. Retrieved from https://ideas.repec.org/p/spo/wpmain/infohdl2441-eu4vqp9ompqllr09iatr32p81.html

Reinganum, J. F. (1983). Uncertain Innovation and the Persistence of Monopoly. *The American Economic Review, 73*(4), 741–748.

Rousakis, M. (2014). The length of patents and the timing of innovation. *EUI Working Paper MWP 2014/24*.

Saidi, F., & Zaldokas, A. (2016). Patents as Substitutes for Relationships. *CEPR Discussion Paper No. DP11580*.

Scotchmer, S. (1996). Protecting Early Innovators: Should Second-Generation Products be Patentable? *The RAND Journal of Economics, 27*(2), 322–331.

Scotchmer, S., & Green, J. (1990). Novelty and Disclosure in Patent Law. *The RAND Journal of Economics, 21*(1), 131–146.

Sternitzke, C. (2013). Patent race-like competition in drug lifecycle management. *Presented at the 35th DRUID celebration Conference, Barcelona*. Retrieved from http://druid8.sit.aau.dk/druid/acc_papers/0gek8c2qgbq3q6ncgsr77f6mu8gs.pdf

Thompson, N., & Kuhn, J. M. (2017). Does Winning a Patent Race Lead to More Follow-On Innovation? *Mimeo*, MIT & UC Berkeley. Retrieved from http://dx.doi.org/10.2139/ssrn.2899088

van Dijk, T. (1996). Patent Height and Competition in Product Improvements. *The Journal of Industrial Economics, 44*(2), 151–167.

Watase, R. (2002). The American Inventors Protection Act of 1999: An Analysis of the New Eighteen-Month Publication Provision Note. *Cardozo Arts & Entertainment Law Journal, 20*, 649–684.

Zaby, A. K. (2010). Losing the lead: The patenting decision in the light of the disclosure requirement. *Economics of Innovation and New Technology, 19*(2), 147–164.

A Tables

	participation by	condition 1	condition 2
Period 1			
	both firms	$c < \text{Min}\{\pi\alpha(1-\beta/2),\ \pi\beta(1-\alpha/2)\}$	
	only firm A	$\pi\beta(1-\alpha/2) < c < \pi\alpha(1-\beta/2)$ [1]	
		$\vee \quad \pi\beta < c < \pi\alpha$ [2]	
	only firm B	$\pi\alpha(1-\beta/2) < c < \pi\beta(1-\alpha/2)$	
		$\vee \quad \pi\alpha < c < \pi\beta$	
	only one of the firms[3]	$\pi\beta(1-\alpha/2) < c < \alpha\pi$	$2\beta/(2+\beta) < \alpha \leq \beta$
		$\pi\alpha(1-\beta/2) < c < \beta\pi$	$\beta < \alpha < 2\beta/(2-\beta)$
Period 2			
	both firms	$c < \text{Min}\left\{\frac{\alpha(1-\beta)(2-\beta)}{2(1-\beta(1-\ell))}, \frac{\beta(1-\alpha)(2-\alpha)}{2(1-\alpha(1-\ell))}\right\}$	
	only firm A	$\frac{\beta(1-\alpha)(2-\alpha)}{2(1-\alpha(1-\ell))} < c < \frac{\alpha(1-\beta)(2-\beta)}{2(1-\beta(1-\ell))}$ [1]	
		$\vee \quad \frac{\beta(1-\alpha)}{1-\alpha(1-\ell)} < c < \alpha(1-b)1-b(1-\ell)$ [2]	
	only firm B	$\frac{\alpha(1-\beta)(2-\beta)}{2(1-\beta(1-\ell))} < c < \frac{\beta(1-\alpha)(2-\alpha)}{2(1-\alpha(1-\ell))}$	
		$\vee \quad \frac{\alpha(1-\beta)}{1-\beta(1-\ell)} < c < \frac{\beta(1-\alpha)}{1-\alpha(1-\ell)}$	
	only one of the firms[3]	TBA (also π needs to be added, or removed from P1 terms)	

Table 3.2: Participation decision of firms in the competition for a single innovation.

[1] Firm A finds it profitable to enter independent of whether firm B enters, while firm B would not find it profitable to enter together.
[2] Firm B generally finds it unprofitable to enter.
[3] Equilibrium requires either some sort of coordination or mixed strategies.

No firm will find it profitable to enter if c is greater than the highest threshold.

condition 1	condition 2	condition 3	condition 4
$\alpha < 1 - \sqrt{g}$	$d < a(2-a)/(2(1-g))$	$p < \dfrac{2d(1-g)-(2-\alpha)\alpha}{\alpha(2d(1-g)-(2-\alpha))}$	$\beta \leq \dfrac{2d(1-g)(1-\alpha p)}{(2-\alpha)\alpha(1-p)}$ & $c < \beta(1-\alpha/2) - d(1-g)$
			$\beta > \dfrac{2d(1-g)(1-\alpha p)}{(2-\alpha)\alpha(1-p)}$ & $c < \dfrac{(1-\alpha)(2-\alpha)\beta}{2(1-\alpha p)}$
		$p \geq \dfrac{2d(1-g)-(2-\alpha)\alpha}{\alpha(2d(1-g)-(2-\alpha))}$	$c < \beta(1-\alpha/2) - d(1-g)$
	$d \geq a(2-a)/(2(1-g))$		$c < \beta(1-\alpha/2) - d(1-g)$
$1 - \sqrt{g} \leq \alpha \leq 1 - g$		$p < \dfrac{2d(1-g)-(2-\alpha)\alpha}{\alpha(2d(1-g)-(2-\alpha))}$	$\beta \leq \dfrac{2d(1-g)(1-\alpha p)}{(2-\alpha)\alpha(1-p)}$ & $c < \beta(1-\alpha/2) - d(1-g)$
			$\beta > \dfrac{2d(1-g)(1-\alpha p)}{(2-\alpha)\alpha(1-p)}$ & $c < \dfrac{(1-\alpha)(2-\alpha)\beta}{2(1-\alpha p)}$
		$p \geq \dfrac{2d(1-g)-(2-\alpha)\alpha}{\alpha(2d(1-g)-(2-\alpha))}$	$c < \beta(1-\alpha/2) - d(1-g)$
$1 - g < \alpha < \mathrm{Min}\,\{1, 2(1-g)\}$	$d < a(2-a)/(2(1-g))$	$p < \dfrac{2d(1-g)-(2-\alpha)\alpha}{\alpha(2d(1-g)-(2-\alpha))}$	$\beta \leq \dfrac{2d(1-g)(1-\alpha p)}{(2-\alpha)\alpha(1-p)}$ & $c < \beta(1-\alpha/2) - d(1-g)$
			$\beta > \dfrac{2d(1-g)(1-\alpha p)}{(2-\alpha)\alpha(1-p)}$ & $c < \dfrac{(1-\alpha)(2-\alpha)\beta}{2(1-\alpha p)}$
		$p \geq \dfrac{2d(1-g)-(2-\alpha)\alpha}{\alpha(2d(1-g)-(2-\alpha))}$	$c < \beta(1-\alpha/2) - d(1-g)$

Table 3.3: Cases in which A prefers disclosure because this makes B switch to innovation 2 under incomplete information when $\pi_1 = \pi_2 = 1$.

section	environment	condition 1	condition 2
4	single independent innovation	–	
5.1	sequential innovation	–	
5.1	– with licensing revenue	$\alpha < \dfrac{2\lambda}{1+2\lambda}$	
5.1	– with discounting (assuming $\ell_B = 1$)	$r < \dfrac{\sqrt{2}-1}{2}$	$\beta \leq 1 - \dfrac{1+\sqrt{1-4r(1+r)}}{2(1+r)}$
			$1 - \dfrac{1+\sqrt{1-4r(1+r)}}{2(1+r)} < \beta < 1 - \dfrac{1-\sqrt{1-4r(1+r)}}{2(1+r)}$ & $\alpha < \dfrac{\dfrac{(2-\beta)r}{(1-\beta)\beta(1+r)} + 2\lambda}{1+2\lambda}$
			$\beta \geq 1 - \dfrac{1-\sqrt{1-4r(1+r)}}{2(1+r)}$
		$r \geq \dfrac{\sqrt{2}-1}{2}$	
5.1	– with complementarity	$\alpha < \dfrac{2\gamma}{2+\gamma-\delta(1-\gamma)}$	
5.1	– with incomplete information and $c = 0$	$\alpha < \text{Min}\{1, (2(1-d)(1-g))\}$	$\beta > \dfrac{2d(1-g)}{2-\alpha}$
5.2	simultaneous complementary innovation	$\alpha < \dfrac{2}{3}$	

Table 3.4: Cases in which A prefers disclosure under complete information when $\pi_1 = \pi_2 = 1$.

In cases marked with "–" there is no endogenous incentive to disclose a pending patent application.

B Figures

In all of the following figures:

- black curves represent the participation constraint of firm A;
- orange (gray) curves represent those of firm B.
- Solid curves: participation constraint in period 1 if the competitor does not participate;
- medium-dashed curve: participation constraint in period 1 if the competitor does participate;
- short-dashed curve: participation constraint in period 2 if the competitor does not participate;
- long-dashed curve: participation constraint for period 2 if the competitor does participate.

With $\ell = 0$ the constraints for periods 1 and 2 coincide.

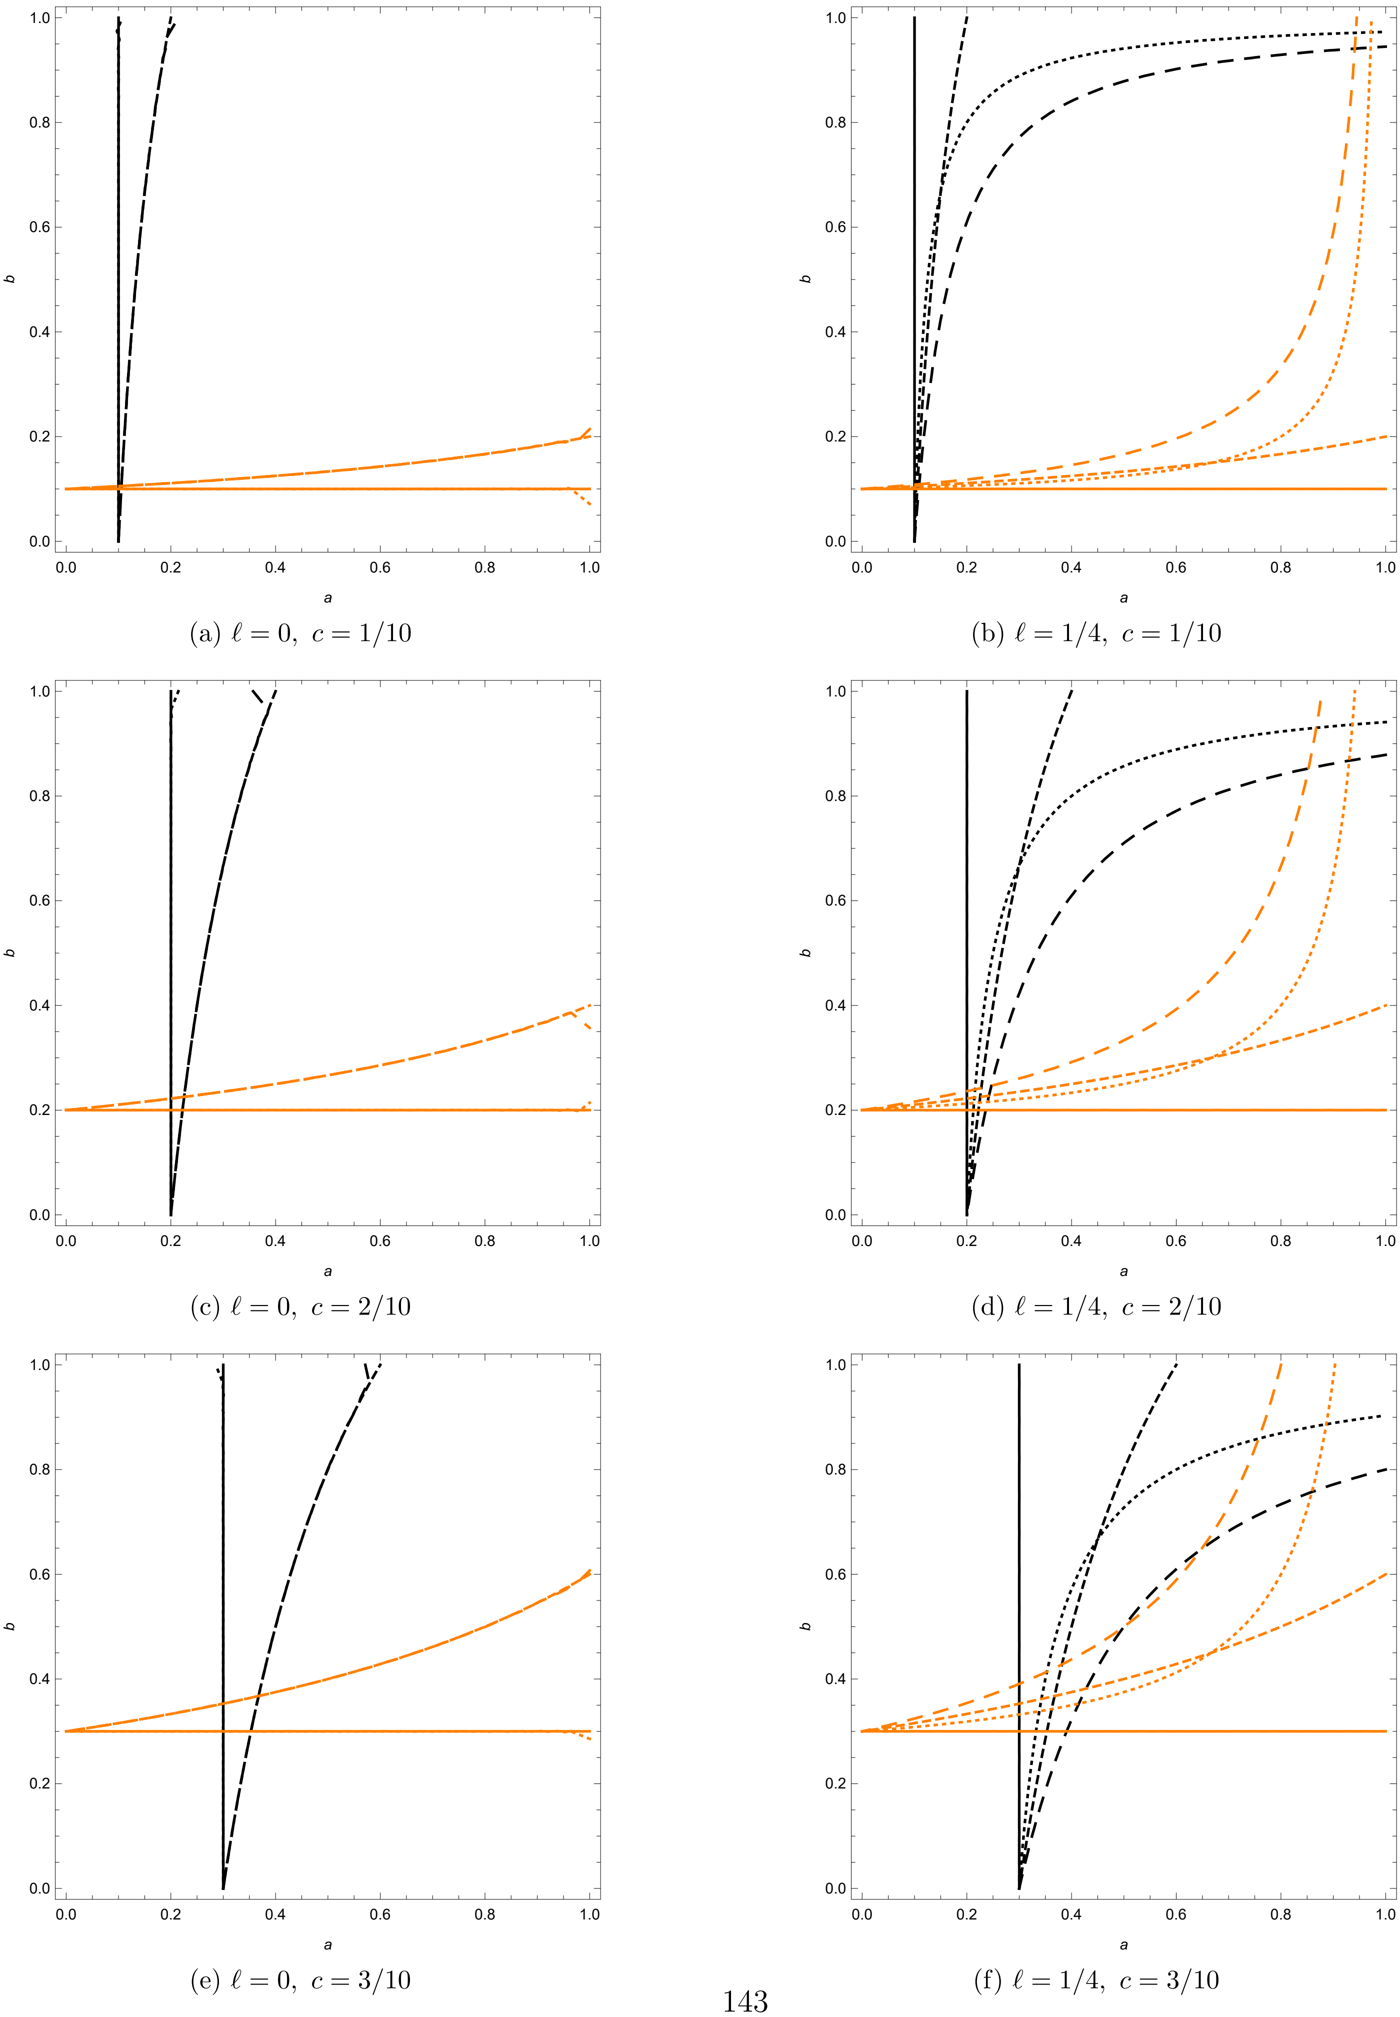

(a) $\ell = 0,\ c = 1/10$ (b) $\ell = 1/4,\ c = 1/10$

(c) $\ell = 0,\ c = 2/10$ (d) $\ell = 1/4,\ c = 2/10$

(e) $\ell = 0,\ c = 3/10$ (f) $\ell = 1/4,\ c = 3/10$

Figure 3.2: Participation constraints in the race for a single innovation (I).

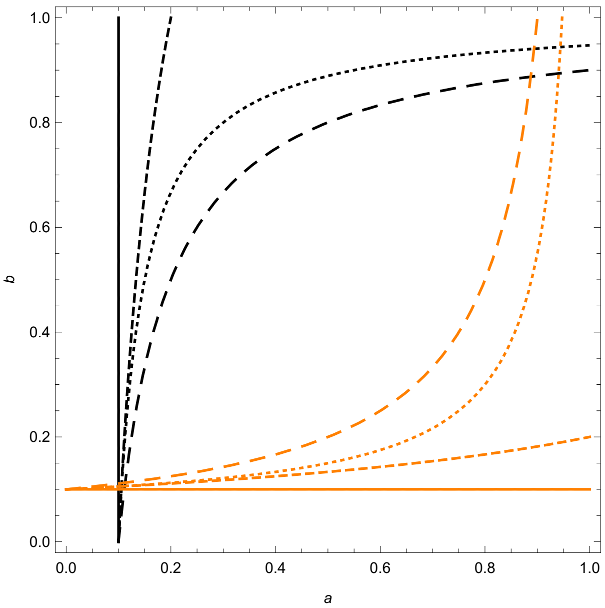

(a) $\ell = 1/2,\ c = 1/10$

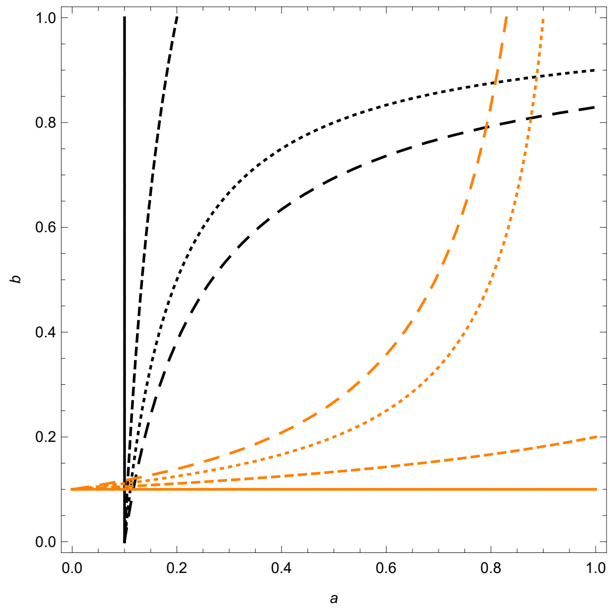

(b) $\ell = 1,\ c = 1/10$

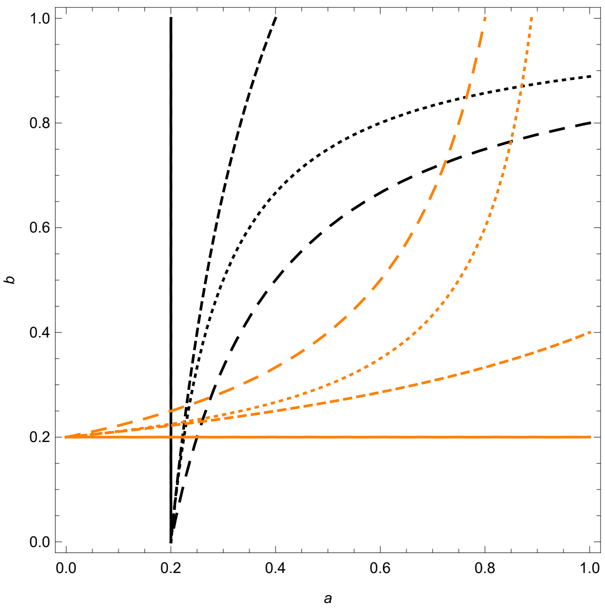

(c) $\ell = 1/2,\ c = 2/10$

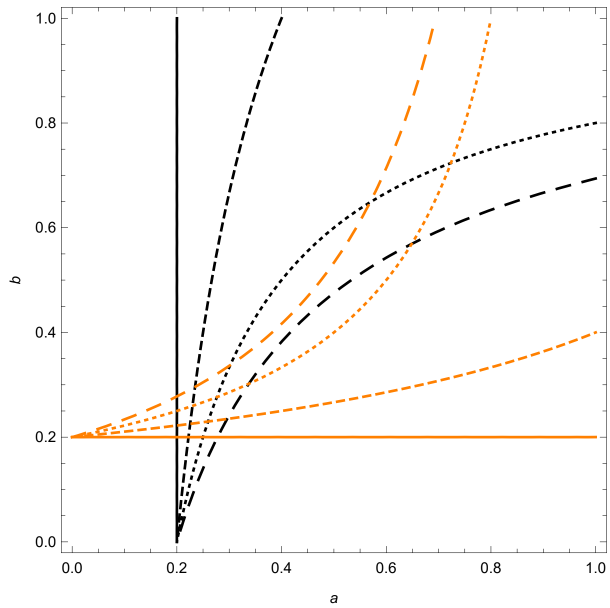

(d) $\ell = 1,\ c = 2/10$

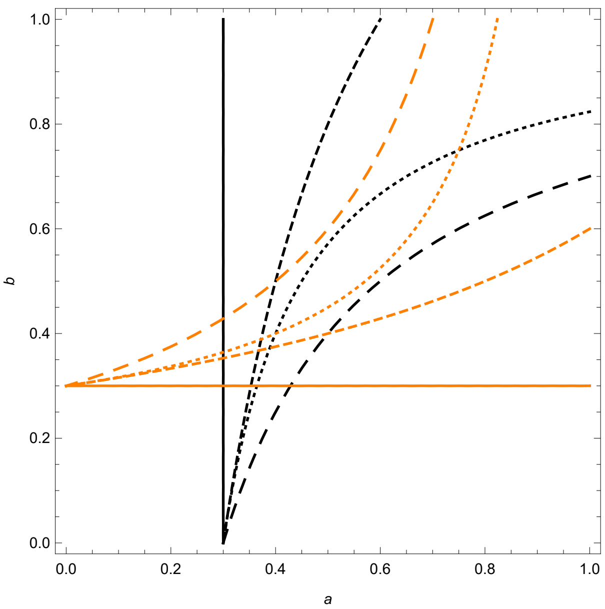

(e) $\ell = 1/2,\ c = 3/10$

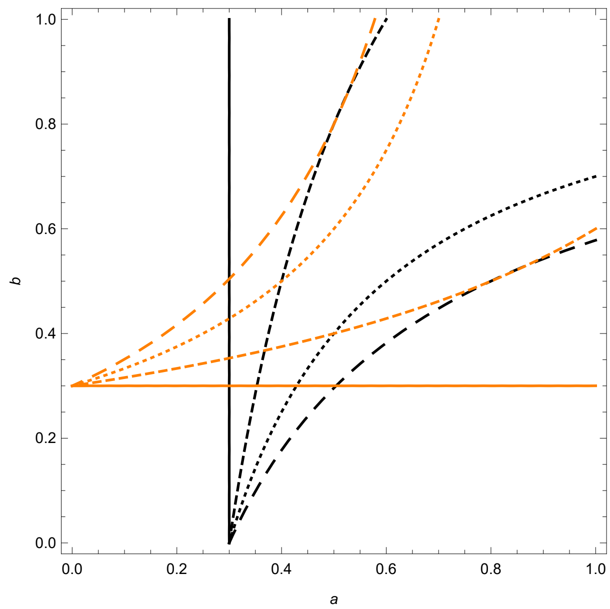

(f) $\ell = 1,\ c = 3/10$

Figure 3.3: Participation constraints in the race for a single innovation (II).