

« Strategies of search and patenting under different IPR regimes »

Auteurs


Robin Cowan, Nicolas Jonard, Ruth Samson

Document de Travail n° 2024 – 20

Avril 2024

Bureau d'Économie
Théorique et Appliquée
BETA

www.beta-economics.fr

 @beta_economics

Contact :
jaoulgrammare@beta-cnrs.unistra.fr

Strategies of search and patenting under different IPR regimes

Robin Cowan

BETA, Université de Strasbourg; 67085 Strasbourg, France
UNU-MERIT, Maastricht University; 6221AX Maastricht, Netherlands
`cowan@unistra.fr`

Nicolas Jonard

DEM, University of Luxembourg; L1359 Luxembourg, Luxembourg
`nicolas.jonard@uni.lu`

Ruth Samson

BSE, Université de Bordeaux; 33608 Pessac, France
`ruth.samson@u-bordeaux.fr`

Abstract

Many scholars observed changes in the intellectual property rights systems in the 1980s and 1990s throughout the world. Patent systems in particular seemed to be expanding their scope, and the legal system seemed to be changing its attitudes towards intellectual property rights. At the same time, and probably in response, firms started to change their patenting behaviour — treating patents as tools of competition and bargaining rather than as a means to protect the fruits of intellectual labour. In this paper we present a simulation model that can be used to discuss that shift. Firms search for new technologies and patent what they find. But different firms have different strategies: one is to protect an invention; a second is to protect a technology space; the third is to attack others' technology spaces. In the literature the latter two have been described as different types of blocking. We examine different IPR regimes, characterized by who is able to infringe whose patent rights. This is an extreme case of who is able to extract rents from a given configuration of patent rights.

Keywords: Innovation, Patents, Knowledge network, Blocking strategies

JEL numbers: O31; O34; C6; L5

This paper is also circulating as a UNU-MERIT Working paper.

1 Introduction

In March 2006, Research in Motion, maker of the Blackberry mobile phone and email device, paid 613 million dollars to NTP to settle a patent infringement case. NTP was a firm that had never made anything, and had no plans to do so in the future. It has been described as a patent troll. The case was settled and the payment made, in spite of the fact that during the litigation the patents were being re-examined on the basis of “overlooked prior art” and were in fact finally rejected as invalid. This case is often seen as one of the paradigm examples illustrating the problems emerging in US, and later world-wide, patent systems. And it was a general issue that exercised Paul David very much at the time. He was very concerned that the then current rapid evolution of patent systems was in the process of destroying the innovation systems they were designed to foster. He spent considerable energy during those years arguing about changes not only to the patent system but to systems of intellectual property rights more generally. He feared that the apparent broadening and strengthening of the system and attendant changes in patenting behaviour of firms in it, was threatening the ecosystem that supported innovation and thus economic growth.

Paul David was not the only person thinking about these things; several people noticed an increase in patenting activity, and an increase in the ratio of patent applications to R&D expenditure (Blind et al. 2006, Hall and Ziedonis, 2001 for example) in the late 1980s and 1990s. While there are several candidate explanations for these trends, one in particular raised concerns. This is that firms move away from the traditional motivation for patenting, that is to secure temporary protection of their intellectual property, to strategic motives. Firms developing a product or technological line want to protect possible future technological developments, and so patent pre-emptively to permit continued development of a current trajectory. On the other hand, they may patent with competitors’ trajectories in mind, thus attempting to prevent competitors’ future developments. Both of these strategies were considered troubling, as rather than protecting inventions and innovations, they are walling off pieces of the technological space. If successful, of course, this can impede innovation. A third motive, with less clear cut implications, is the acquisition of patents as bargaining chips, for use in case the firm is accused of infringement. A firm with a large patent portfolio is more likely to be able to defend such a suit by a counter-suit. Less litigiously, firms may simply be able to trade patents, if there is a double coincidence of wants. (See Blind et al. 2006 for further discussion of other patenting motives.) All of these motives raise concerns that the initial rationale for the IPR system has been lost, and that the IPR no longer supports but rather (in several important industries at least) impedes innovation.

This paper takes up that concern, and presents a simple model of patenting behaviour, asking about the interaction between firm strategies and micro and macro performance

in the face of changes to the structures (and strengths) of the patent system.

There is now considerable literature in Management on the implications of patent protection for innovation and R&D investment. Recent studies show that firms patent strategically with motives other than the classical goal of appropriating returns from R&D (Hall and Ziedonis, 2001; Blind et al., 2006; Harhoff et al., 2007; Holgersson, 2013; Holgersson and Granstrand, 2017; Blind et al., 2018). The implications of strategic patenting (for instance, on firms' R&D investment and market value) have received much attention empirically (Grimpe and Hussinger, 2014; Clancy, 2018; Blind et al., 2009; Helmers and Rogers, 2011; Panagopoulos and Park, 2018) and theoretically (Schneider, 2008; Chu, 2009; Krasteva, 2014; Mihm et al., 2015). Yet, we are still without a nuanced understanding of how firms should search a technology landscape and how optimal search strategies may differ under different IPR regimes.

This paper presents a simple model of firms discovering and patenting technologies in the development of different products. We model innovation as a process of building technological bridges, where new technological inventions are combined with existing technologies to advance existing products. This takes place under different IPR regimes, and with a population of firms having idiosyncratic R&D and patenting strategies. Essentially a “product” begins with a pioneering patent, and then is developed as new technologies are discovered, possibly patented, and added to the product (Cohen and Tripsas, 2018). Thus a product or technology evolves through the accretion of new, improvement discoveries. Using a network representation, we capture the structure and evolution of the technology landscape, while focusing on the search and patenting strategies firms might follow.

1.1 Patent protection and strategic patenting

In principle the patent system has the primary goal of creating incentives for innovation, which it does by offering limited-time monopoly rights in exchange for full technological disclosure (Drivas et al., 2016; Erkal, 2005; Archibugi and Planta, 1996; Dunlavey, 1955; Crass et al., 2019). However, evidence on whether formal disclosure through patenting encourages or discourages idea generation and innovation is mixed (Boldrin and Levine, 2009; Suzuki, 2015; Ouellette, 2012). On the one hand, studies show that imperfect patent protection, allowing a certain level of imitation, stimulates investment in R&D and innovation (Krasteva (2014); Hellmann and Perotti (2011); Klein (2020)). On the other hand, there are numerous arguments against patents: disclosure through patents hurts innovation as it discourages agents from engaging in creative endeavours, and the requirements for patent grants involve so much disclosure that imitation is made too easy (Anton and Yao, 2004; Teece, 1986; Boldrin and Levine, 2013; Horstmann et al., 1985; Chu et al., 2012; Suzuki, 2015; Gans et al., 2017; Cockburn et al., 2010).

Despite weaknesses in the patent system, firms still do patent their inventions. Indeed, there has been a tremendous increase in patent applications and grants (especially in the US) since the 1980s (Gallini, 2002; Hall and Ziedonis, 2001; Hall, 2004). Many of these commentators have commented on a “regime change” taking place, starting in the US in the 1980s. Among others, Lerner and colleagues¹ observed a “pro-patent shift” in the legal environment, particularly with the establishment of the Court of Appeals for the Federal Circuit, which focused its attention on patent issues.² One might expect that such a regime change would induce a change in patenting behaviour and in particular a change in the strategic uses to which patents are put.

In light of this increase in patenting (and in particular the ratio of patents to R&D expenditure), researchers have established that firms’ patenting decisions vary across industries and types of inventions and are not entirely motivated by the protection of proprietary know-how. Instead, firms’ decisions to file for a patent may encompass a combination of strategic acts to sustain a competitive edge. Firms use patents for strategic reasons, including building a shield around a particular technology to push away rivals, or pre-empting competitors on advances in a particular technology — fencing or offensive blocking (Walsh et al., 2016; Grimpe and Hussinger, 2014; Czarnitzki et al., 2020; Jell et al., 2017). Also, firms use patents to force entry into a particular technology or to block another firm’s downstream invention in a particular direction — defensive blocking (Blind et al., 2006; Hall et al., 2021; Harhoff et al., 2016). Phenomena such as “patent thickets” may arise from fragmentation of the technology landscape when an inventive entity holds several interlocking patents around a core invention to prevent rivals from designing around or developing substitutes (Shapiro, 2000; Cockburn et al., 2010; von Graevenitz et al., 2013). It is also a mechanism that inventors use to block competitors from entering a particular technology and(or) to avoid being blocked by competitors (Mihm et al., 2015; Desyllas et al., 2018).

1.2 The cumulative nature of innovation and product development

The notion of “standing on the shoulders of giants” implies that a firm’s quest to innovate does require not only the firm’s “own” stock of knowledge but also the knowledge stock of “others” (Scotchmer, 1991). However, in the context of intellectual property protection, especially using patents, this also implies “standing on the shoulders of rights holders”

¹Kortum and Lerner (1998) or Lerner (1995) for example. See Hall and Ziedonis (2001), page 105 for a discussion.

²There was also the discovery by, patent owners, of the Eastern District of Texas, which was particularly favourable towards plaintiffs in patent infringement suits, and especially so towards entities sometimes known as patent trolls. This stopped in 2017 when the US Supreme court limited the ability of plaintiffs to choose the jurisdiction in which they filed suit. See Mullin, 2017.

(Duffy, 2008; Ziedonis, 2004). The cumulative nature of innovation engendered the notion of core or pioneering inventions (protected by basic patents) and improvement inventions (protected by improvement patents) (Silverman, 1995; UpCounsel, 2020). Core or pioneering inventions are inventions upon which many other technological developments can be built. The US supreme court defines a basic patent (and invention) as

“a patent covering a function never before performed, a wholly novel device, or one of such novelty and importance as to mark a distinct step in the progress of the art” (Merges and Nelson, 1990, pg. 854).

In contrast, improvement inventions build directly over another invention, the core or pioneering invention. Without the latter, while the former may exist as a piece of knowledge (invention), it cannot be exploited or marketed (as an innovation). In principle improvements can be “original” if they provide the same functions but with technical modification to the core invention (Merges and Nelson, 1990; Silverman, 1995; Scotchmer, 1991). But, an improvement invention can also be additional, when it adds features, or value, to existing technology.³ Inherently, an improvement patent (protecting an improvement invention that meets the patentability criteria) builds directly upon a basic patent (protecting the pioneering invention). In this paper we conceive of improvement patents as adding features to an existing product or technology. Such improvements could be (discovered or) made by the owner of the pioneering patent, or by other firms. Who discovers the improvement patent will have different effects under different IPR regimes.

1.3 Firm strategies

The strategic use of patents must be coupled with a relevant search strategy. Where in the technological landscape should a firm focus its R& activities? For fencing, search around one’s own technologies would be appropriate; for defensive blocking, search around a rival’s. If the traditional patent motive is in play the search wherever there might be “good stuff”.⁴ We model firms’ search on the landscape (movement on the landscape) as a deliberate decision to invest in research and development based on various search motives, which are as well the motives behind patenting the resulting inventions. Further, the unpatented inventions represent the uncharted territory of the landscape that no one knows about. In our model, the knowledge discovered and disclosed by other inventive entities act as streetlights illuminating the way. Thus, the landscape is uncharted unless a firm discovers something and discloses it through patenting (pioneer or improvement patent).

³The current technology landscape has evolved to the extent that most inventions seeking protection through patents are improvement technologies, providing meaningful novel, non-obvious, and useful improvements to existing technological inventions (Justia, 2022).

⁴We have obviously abstracted from issues of path dependence at the level of a firm’s knowledge or competence, so a richer discussion would ask for more nuance in these strategies.

A firm’s ultimate payoff from innovation activities depends on its R&D strategy — deliberate movement in a particular direction on the technology landscape — contingent on different external and internal factors such as industry characteristics and appropriability regimes (Teece, 1986; Granstrand, 1999; Pisano, 2006; Laursen and Salter, 2014; Holgersson et al., 2018; Cohen et al., 2000). We follow Granstrand’s (1999) argument that firms’ R&D strategies and patenting strategies are tightly coupled. Mihm et al. (2015) separated firms’ patenting decision (strategy) from R&D strategy and found that firms’ R&D strategy (leader or follower) is the main determinant of patenting strategy (patent all, close competitor, high quality, or nothing). We do not focus on how or what firms decide to patent. Rather, our interest is in the extent to which products are developed, and whether different patenting strategies pay off differently in different IPR regimes, and under different population ecologies.

In what follows we develop a simple agent-based model in which firms explore a technological space in order to improve existing products. A product begins with a “pioneering” patent which can be further developed by technological discoveries. Thus firms explore the technology space, hoping to find a technology that can be linked to an existing product, patented and thus exploited (or to find a new pioneering patent). Different firms follow one of three different strategies, roughly speaking: the classic “appropriate rents from this discovery”; patent to protect my current products; and patent to stop competitors from developing their products. We use the model to explore the value of these different strategies under different IPR regimes and different distributions of the three strategies over the population of firms.

2 Ideas and products

There are three parts to the basic structure of the model. We can think of this as a multiplex network of three layers, one lying above the other, and each layer being a reduction of the layer below. That is, the bottom layer represents the technology space, showing how technologies are located relative to each other. The middle layer connects technologies or groups of technologies into potential products. The top layer is a subgraph of the middle layer, showing which technologies have actually been discovered and which products are brought to market.

2.1 Ideas

The most basic, foundational, part of the model are ideas or technologies. There are many ideas in the world which might or might not be discovered. Technologically or scientifically, ideas can be close to, or distant from, each other. Thus in the model we

assume that the world contains a fixed number of ideas which might or might not be discovered at future some time. But any pair of ideas has a distance in knowledge space indicating how (dis-)similar they are to each other. For simplicity, we represent knowledge space with a finite two-dimensional lattice.⁵ Each node represents a potentially patentable idea or technology, and nodes close to each other in the lattice-space are considered similar technologically. This structure is unchanging over time.

2.2 Potential products

Ideas or technologies can be combined in various ways into products. One can think of a product as a collection of attributes, each attribute arising from one idea or technology. Some technologies can be sensibly combined into a product (camera and mobile telephony) some cannot (milking machines and asphalt paving perhaps). Thus a product is a set of ideas that can potentially be combined (in a specific order) into a something with value. Many sets are supersets of others: today’s automobile contains all of the ideas of older automobiles with the addition of bluetooth communication for example. All are feasible products, though you might expect the “superset product” to cannibalize one or other, or both, of the two individual products.

The middle layer of the structure, then, consists of a list of sets of ideas, each set defining one product. Initially, this structure indicates only potential: initially no products exist, but they could be created as time passes and ideas are discovered. It defines (small) sub-graphs of the underlying lattice, with each sub-graph representing a potential product. The middle part contains all possible technology combinations that could have value if brought to the market. This structure is also unchanging over time. Figure 1 represents the two bottom layers in the network.

2.3 Patents and products

The top layer contains information about what ideas have been discovered, patented, exploited, and joined together into products. Ideas once discovered are patented, so the top part of the structure represents, at each point in time, the current state of knowledge — which of the ideas has been discovered. Further, though, given the constraints of the product space, ideas can be exploited, if they are used in products that are brought to market. That is, if in this top layer all the ideas of a set in the middle layer have been discovered, they can be exploited as that product (if such exploitation is profitable to the owners of the ideas). We display an illustration of this third layer later in the paper.

⁵In point of fact the lattice structure is not necessary; what is necessary is some structure that provides a distance measure between ideas, where “closer” ideas are technologically similar to each other, and thus “discoverable” from each other. We use the lattice structure just to facilitate a graphical illustration of the model.




Figure 1: The idea and potential product spaces. The left panel is the bottom layer of technologies or ideas and connections between them, following which firms discover new technologies or ideas; an edge between two ideas expresses that the idea at one end of the edge is discoverable through local search from the idea at the other end. The right panel is the middle layer of potential products. A single product is identified by one shade of grey, with the pioneering idea shown as a solid-grey node, subsequent nodes in the product identified by edges of that shade of grey. Disconnected, unframed vertices represent ideas that belong to no potential product. Observe that some ideas (4, 5 or 77 in this case) can be used in more than one product.

3 Patents, products, and strategies

In this section we develop the structures under which firms search for new ideas, patent inventions, and develop products.

3.1 Pioneering ideas, improvement patents and intellectual property

We think of R&D as a search process on the idea landscape. Firms search the idea space in different manners, and the search style determines how they interact with one another (see the next subsection). For each firm, history starts with the ownership of one random pioneering idea. A pioneering idea is the starting point of a potential product. Without it, the product cannot exist, but as soon as it is discovered and patented it offers a possible source of income as the patented invention can be exploited. From here on, further improvement patents (corresponding to ideas in the same potential products) can be discovered. Every idea is patentable, i.e., offers sufficient novelty to meet the patentability requirements. However, an improvement patent cannot be exploited without accessing the rights to the basic, pioneering patent, and all the subsequent patents to which it constitutes an improvement. Therefore, when the protection conferred by patents is perfect, as soon as different firms are involved in the ownership of patents constituting a

product, improvements will be blocked.

Not any improvement is possible at any time for a given product, and we assume that improvements must follow a unique, well-defined sequence. Imposing a sequence between the various improvements made to a product captures well the features of industrial R&D, but within the context of the model, it is also a way of creating the conditions for effective blocking by firms. In reality, though all improvements are not possible at all times, there are probably several ways in which the full potential product can be developed and so blocking strategies are less effective as perhaps some more advanced product variants can still be brought to the market. Our assumptions thus make the conditions and consequences of firms employing blocking strategies more salient.

To illustrate, consider a potential product consisting of the sequence $(29 - 19 - 30 - 9)$ and a single inventor searching the idea space. In Figure 1, this product is in the bottom right corner of the right panel. In the sequence constituting the potential product, the pioneering idea is 29, which can be patented and exploited by its inventor (through a license or direct development and commercialisation) yielding an initial set of patents $\{29\}$ and a product (29) , which is the first generation of the (improvable) product. If patent 19 is discovered (recall there is only one inventing firm so far, so no issue of conflicting property rights), augmenting the set of patented ideas to $\{29, 19\}$ then the product can be improved into $(29 - 19)$ and exploited. The demand side will always prefer the improved version of the product, and switch entirely to the most recent variant of the product. If however patent 30 is discovered before the discovery of 19, the set of patented ideas becomes $\{29, 30\}$ and 30 is valueless as the sequence $(29 - 30)$ is not part of the product. Nothing can happen on the product front (and the value of 30 as an improvement cannot be realized) until 19 is discovered, in which case the set of patented ideas becomes $\{29, 30, 19\}$ and the product offered on the market is $(29 - 19 - 30)$. There is no point in offering $(29 - 19)$ as consumers will unambiguously prefer $(29 - 19 - 30)$. This process continues until the inventor has discovered the entire sequence constituting the product, in which case search “around” this product stops.

In the example above we have considered a single inventor who owns patent rights to all inventions. Thus everything potentially exploitable is eventually exploited. Consider now that there are several inventing firms. If ideas are discovered and patented by different inventors, then the patent system plays a role. Specifically, as Silvermann (1995) puts it: “A patent grants the patent owner a negative right (i.e., the right to prevent others from making, using, or selling the patented invention); it does not give the patentee the right to practice the invention. In many instances, a patented improvement cannot be made, used, or sold without infringing the basic patent.” Suppose in the previous example that ideas 29 and 19 are discovered by different firms. The owner of 29 (the pioneering idea) does not need the consent of the owner of 19 to exploit the pioneering idea, but the owner of 19, which is an improvement idea over 29, cannot use it without the consent

of the owner of patent 69. And the same issue would apply to the owner of 30, as it is an improvement over 29 and 19. Thus, although the set of patented ideas is $\{29, 30, 19\}$, if no agreement can be found, product improvement and value creation stop with (29). In principle, though, an agreement in the form of revenue sharing, or of patent trading could permit further development and exploitation. We develop this in the next sections.

3.2 R&D and patenting behavior: the candid inventor and the strategic patentees

Three patenting motives coexist in the industry. They represent archetypes identified in Cohen et al. (2000) and considered in a large literature on strategic patenting motives that emerged in the early 2000s (see section 1.1 above). In this paper, we consider the individual and aggregate performance implications of having these patenting motives interacting in a race for innovation. We do not claim firms do or should adhere strictly to one archetype over time, but rather that certain contingencies will make certain behavioral features superior. We will explore 3 such contingencies in detail: the possibility to trade patents, the properties of the IPR regime, and the ecology of behaviors in the industry.

The first behavior is the traditional ‘patent to protect to exploit’ motive. A firm searches for pioneering patents (in our model a firm needs at least one to generate income) and improvement patents to further develop its pioneering patents. A firm sees when further improvement of a product is blocked by the patent of another firm, and searches away from ideas in existing products, regardless of their stage of development. This view of R&D and IP is somewhat short-sighted (candid?) but we believe captures the essence of the balance between exploration and exploitation that is reflected in the traditional logic of innovation. Why would the firm waste resources searching for improvement patents that will be impossible to exploit if the firm does not control the pioneering idea or is blocked on the path of product improvement by another firm’s patent? The focal firm is better off looking for another pioneering idea, of which there are plenty in the idea space.

The second and third behaviors move away from the traditional ‘patent to protect to exploit’ logic, as they embed an element of strategic forward thinking. The second firm-type searches and patents around its own patents. In so doing it further develops its own products (those for which it owns the pioneering patent) but also patents around its ideas to forestall encroachment by competitors. These are firms that seek to create patent ‘fences’.

The third firm-type, though it owns its own product(s), also looks for opportunities to stall its competitors. It seeks around others’ ideas — if successful it either stops their product development or shares in the rents associated with the others’ products, using its patents as bargaining chips. These are firms that seek to create patent ‘thickets’.

These three behaviours reflect the main motives for patenting stated by firms in the Yale survey (Levin et al. 1987) and all its variant offspring (see e.g. Arundel and Kabla 1998; Cohen et al. 2000; Reitzig 2014 among others).

4 Model

We can describe the model informally as follows.⁶ The world consists of a set of discoverable ideas for which firms search, and which, if successful, they patent. Ordered sets of ideas constitute a product, each set based on one “pioneering” patent. Ideas are discovered one at a time, not necessarily in order, and when an idea is “the next one” in the development of a product, it might, depending on ownership, be added to the product and exploited. In each period a firm is chosen to act. In the base version of the model, patent rights are absolute and strong, and neither side-payments nor trading are possible. Thus, a firm can only develop products for which it owns the pioneering patent. However, if a different firm discovers “the next” idea in a product, its development stops. In later sections of the paper we modify the strength of patents and how they are used, to see how different strategies perform under different IPR regimes.

4.1 Experimental settings

The base case we examine has 2025 ideas arrayed on the two-dimensional non-periodic lattice, with 45 ideas along each dimension. There are 10 firms of each of the 3 types, and each firm owns one pioneering patent at the outset. There are other pioneering patents, undiscovered yet, in the idea space (2 times more than the number of initially available pioneering patents). Product length is equal to 10.⁷ A simulation runs until all profitable patents have been discovered (this can take up to $2025 - 30 = 1995$ periods, but often takes less time). We replicate each configuration 400 times, and report performance distributions for the three types of strategies in the form of box-plots.

ADD regimes (3), nb of pioneering ideas at the outset, in total, IP protection β , nb of replications, product length

⁶A full formal description is in the Appendix.

⁷Product length can be translated into many discussions of product complexity: a complex product is often described as having many features or involving many patents. In the context of our model, then, changing the length of the product strings could capture changes in product complexity. Doing so though, experimenting with products of length 5 and of length 20, we have found no qualitative differences with the results presented below.

Parameter	Value
Number of pioneering ideas	$(1 + 2) \times 3 \times 10 = 90$
Total number of ideas	$45 \times 45 = 2025$
Product length	10
Number of strategies	3
Number of firms	$3 \times 10 = 30$
Number of replications	400
Number of IPR regimes	3
IPR protection coefficient	$\beta \in \{0.5, 1\}$

Table 1: Parameter settings and ranges for the simulation experiment.

5 Results

We treat as the base case absolute, blocking patent rights, without patent trading. We present the results for this very strong case, and then examine the effects of introducing patent trading to the system while keeping absolute, blocking patent rights. Then we move to two regimes of imperfect protection, one favorable to incumbents but not to “interlopers”,⁸ the other more favorable to interlopers. In both cases we consider the possibility of patent trading.

5.1 Perfect IP protection

We start with a context of perfect IP protection. Perfect IP here means that blocking is fully effective. For any potential product, as soon as an improvement patent is obtained by a firm (say j) that is not the owner of the product’s pioneering patent (say i), further product development is impossible. This can be interpreted as j blocking i , in the sense that i cannot continue improving its product (by “its product”, we mean the product for which it owns the pioneering patent) because there is only one sequence of product improvement, and j occupies it. Here we take the blocking power of patents very literally, and assume that firms do not negotiate over revenue sharing details. We make this assumption in order to focus on blocking strategies. By contrast, if revenue sharing agreements were considered, one could imagine a split such that every additional patent is exploited, with revenue shared among the firms *pro rata* relative to the number of patents owned within the product, for instance. In such a case it would be beneficial for all firms to permit any patent to be included, and so all potential products eventually become fully developed. We abstract from such arrangements, assuming firms are interested only in revenue generated by the products they control, regardless of royalties or other payments they could get from holding a patent in products controlled by others.

Performance of the firms using the different strategies are shown in Figure 2. In an

⁸By “interloper” we mean a firm that patents a technology that is part of a different firm’s product, thus potentially blocking further development.

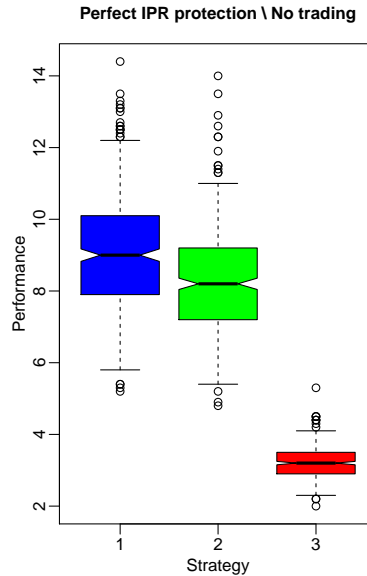


Figure 2: The distribution of average performance of firms of each of the three behavioral types over 400 replications; perfect IP protection and hard blocking.

IPR regime guaranteeing perfect protection, type 1 firms prevail, followed closely by type 2 firms, both far ahead of type 3 firms. A type 1 firm exploits its products through local search, but it also follows the most exploratory behaviour, therefore discovering more pioneering patents, and so controlling more products. A type 2 firms explores less than a type 1 firm, but by remaining in the vicinity of the pioneering patents it holds, also collects high revenues by developing the products it controls. Type 3 firms by contrast behave in a very unproductive way, only searching in the vicinity of other firms' patents; patents they obtain they cannot exploit themselves, unless by pure luck they happen to be part of a product whose pioneering patent they control (and at the right position in the sequence of improvements). Without trading, they have no means to extract rents from other firms owning pioneering ideas.

Given the behaviour of type 3 firms, which tends to block others, most products do not experience their full potential development: indeed, we find that only roughly 30% of the valuable ideas are exploited, which means that blocking occurs on average after the pioneering patent and the first 2 improvements have been found. Thus it is an advantage to own many pioneering patents, since when one product is blocked the firm can move to another. It is the strategy of type 1 firms that leads to the discovery of the most pioneering patents.

5.2 Trading IP

Suppose now that if two firms block each other they could trade patent ownership. The idea is that if firm R holds a blocking patent on a product of firm G and vice versa, R

and G could trade those patents (one-for-one) and both firms would be better off. This implies a slight change in the strategy of type-3 firms: they should not search randomly among others' technologies, but should concentrate search around firms that hold patents blocking their own products. That is, they search for patents that would be useful in making trades. Figure 3 illustrates a case of trading.

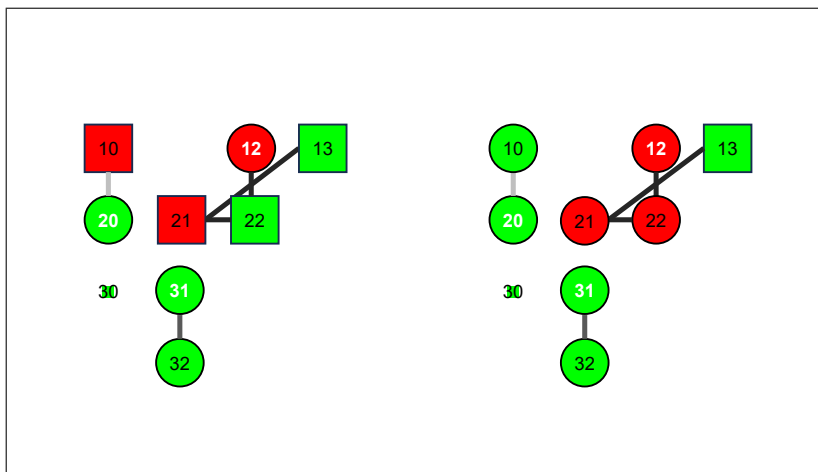


Figure 3: The current patent and product spaces before (left) and after (right) patent trading with 2 firms (red and green): vertex colors code for firm identity; circles indicate exploitation and squares indicate blockade; coloured edges have the same interpretation as in the potential product space; pioneering patents labeled in bold white font. Red originates a product with node 12, but is blocked by Green who owns node 22 (which renders 21 also unexploitable). Green originates a product with node 20, but is blocked by Red who owns 10. Green owns a second product, (31 – 32). Idea 30 has been discovered by Green, but is valueless as it belongs to no potential product. In the right panel, Red and Green have traded patents (10 for 22), and so have both extended their products: Green now offers (20 – 10) and Red offers (12 – 22 – 21). Red would like to include 13, currently owned by Green, but has at this point nothing to offer in trade, and thus (12 – 22 – 21) is blocked.

In Figure 4, we present two boxplots. On the left we repeat the base case; on the right we show the outcomes for the different strategies when trading is possible.

A first observation is that trading is beneficial to all, even if only slightly. This implies that more improvement patents are exploited, consumers enjoy better products with more features incorporated, and firms earn higher profits. In terms of relative performance, type 1 still prevails, but the benefits trading provides to type 3 behavior are clearly visible: trading is mostly beneficial to the firm that accumulated bargaining chips by searching where there are many products being developed by others. Nonetheless, accumulating trading opportunities (which might or might not be exploited) remains less profitable than the two other strategies which prioritize developing one's own pioneering ideas.

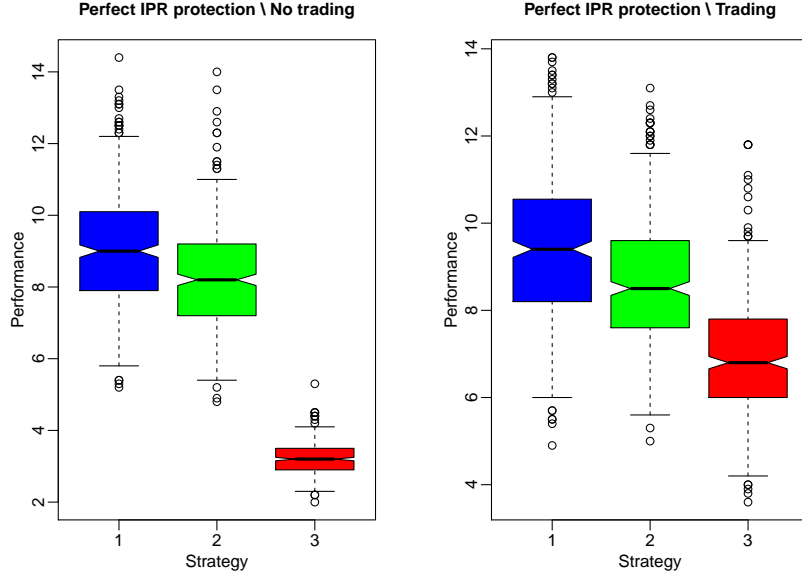


Figure 4: Box-plot representations of the average performance of the three behavioral types in the absence (left) and presence (right)) of patent trading over 400 replications.

5.3 Imperfect IP protection

We consider two forms that imperfect IPR protection can take. The first corresponds to IP infringement that goes unsanctioned when performed by the owner of the pioneering patent. Whether courts are failing at properly defending prior IP, whether they consider IP violation to be less detrimental to innovation than blocking product development would be, or whether the owners of the violated IP do not take the case to court, control over the product is maintained by the incumbent (the owner of the first, pioneering patent in the potential product). In effect, the IPR regime is favoring incumbents. We model this probabilistically, as the outcome of a Bernoulli trial. With probability $1 - \beta$, where β measures IP protection, when the owner of the latest improvement also owns the pioneering patent, the incumbent infringes the property rights of any other firm that owns a patent in that product sequence, and captures total product revenue, thereby appropriating the value created by the innovative efforts of others. Essentially, courts are reluctant to stop the development of a product because of a possible patent infringement.⁹ In the other outcome of the Bernoulli trial, occurring with probability β , strong property rights hold and a patent held by a non-incumbent firm blocks further development.

The second form of IP imperfection has a different focus. While IP infringement was a prerogative of incumbent firms, it now becomes possible for firms that do not own

⁹Of course this is an extreme implementation. More generally what we consider is a case in which it is not very costly for an incumbent firm to infringe another firm's patent. Or a case in which the holder of a single patent relevant to a product that involves many patents does not thereby have a large claim on rents that accrue to the overall product.

the pioneering patent. So in effect, the IPR regime here favours interlopers: the last firm to patent something in a product sequence extracts all the rents of that product, so effectively owns the product, and becomes the incumbent. But that firm can lose it all if another interloper (who might of course be the original pioneer, given the way we have modelled it) appears. With some probability, any firm holding the latest improvement patent infringes the IP of its predecessors in the product, and then behaves with subsequent improvements as if holding the pioneering idea — the owner of the latest improvement, when successful at infringing prior IP, controls the product. Again with probability $1 - \beta$, where β measures IP protection, the owner of the latest improvement infringes the property of all the other firms involved in the product, and captures total product revenue, thereby appropriating the value created by the innovative efforts of others.¹⁰

Figure 5 presents the relative performance of the three patenting behaviours, when protection is perfect (top row), imperfect and favorable to incumbents (middle row), and finally imperfect and favorable to interlopers (bottom row), in the absence (left column) and presence (right column) of patent trading.

Weakening patents to favour incumbents increases overall payoffs generally. It becomes harder to block an incumbent, so more improvement patents are exploited and more products are more fully developed. Aggregate “output” increases. Additionally, though, the dominant strategy changes: type 2 now dominates types 1 and 3, in this order. By acquiring patents in the neighbourhood of its own patents (which include pioneering patents), the type 2 firm is ideally placed to build up long sequences of improvement patents when infringement is “successful”. This applies to other firm types, but with less pronounced effect because they spend more time either exploring away from existing patents (type 1) or searching the neighbourhood of other firms’ patents (type 3). All firms thus perform better, but the increase in performance is maximal for type 2 firms.

When the weakening of patents favours interlopers, the performance ordering of type 2 and type 3 firms is reversed, while type 1 firms maintain an intermediate position. When infringement is possible half of the time, searching around the patents held by others rather than seeking to develop one’s own product pays off — type 3 firms are ideally placed to take advantage of imperfect IP protection. Other types also benefit from the failure of IP (profits are larger for all firms than they are under incumbent-favourable weakening), but to a lesser extent. It is quite remarkable that although at first sight the third patenting motive is somewhat unreasonable — rather than trying to develop

¹⁰Again this is an extreme. In principle the idea is that the holder of a single patent has “hold-up” power and can extract an unseemly proportion of the rents that accrue to the entire product (see the example of Research in Motion in the introductory paragraphs of this paper).

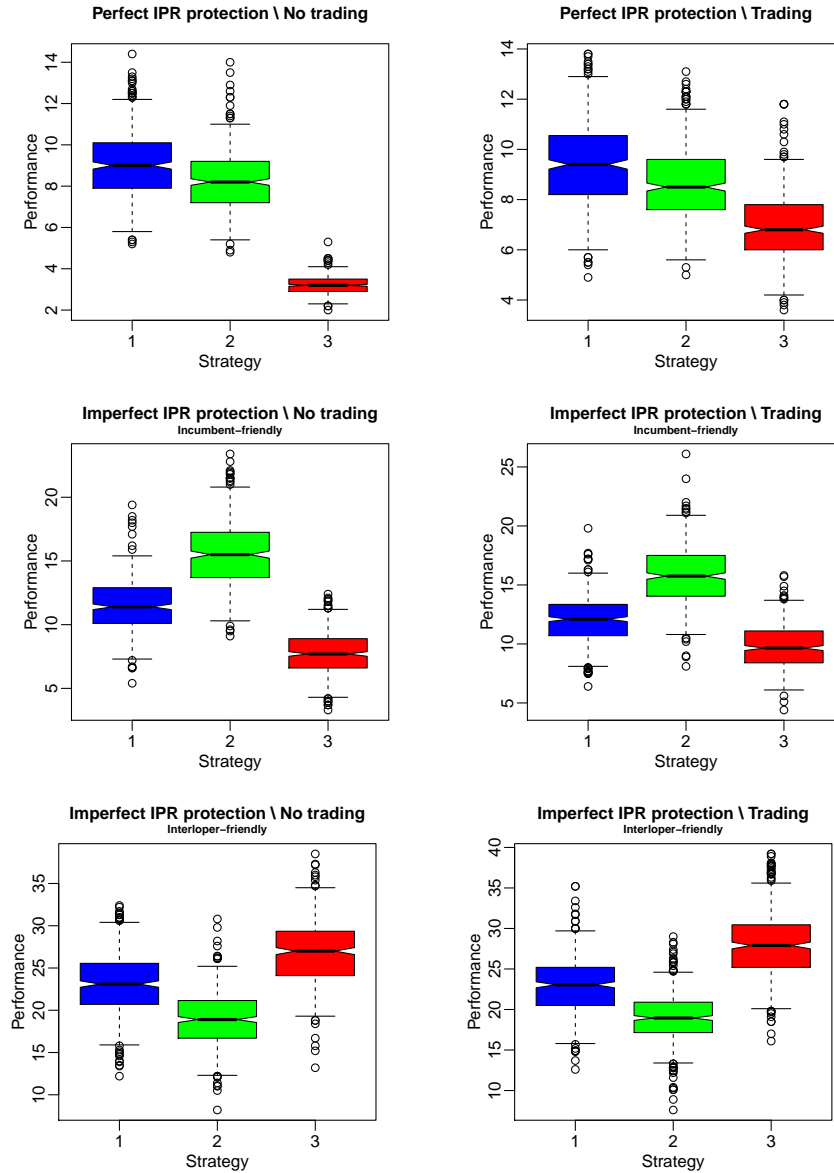


Figure 5: Box-plot representations of the average performance outcomes of the three behavioral types over 400 replications. Note that different panels have different scales on the ordinate axis.

its own product, the firm will only try block other firms in the industry — imperfect IP protection will make it most profitable.

It is worth pointing out that the dominant strategies do not emerge immediately.¹¹ Regardless of the regime, in the short run type 1 firms report the highest payoffs. They patent the most and exploit the highest number of patented technologies. The discussion above of the long run results of the base case of perfect IPR applies as well in the short run: when blocking is abrupt, it is advantageous to have a strategy that finds many pioneering patents. In the second and third regimes, all firms, but those of types two

¹¹We observed the system after 10 innovative attempts per firm.

and three in particular, can gain value by taking over products of other firms. This adds value most clearly to type three firms in the third regime, but also to type two firms in the second (even if someone has “stolen” their product, by continuing to search around their own patents, they can “steal it back”). This sort of theft results in the thieving firm acquiring all the value of the product. But the value of any product will increase with time, so the value of these thefts will also increase with time. Thus early on, developing one’s own products is of more value than stealing others’. So again type one firms dominate in the short run. So in both regimes of imperfect protection, it takes time for the dominant strategy to bear fruit.

6 Dominant strategies in heterogeneous populations

In this section, we relax the assumption of an equal number of firms following each type of behaviour, and ask how relative performance changes in population of firms characterized by different prevalence of each behavioral type. We consider a total of 30 firms, and vary the number of firms of each type by increments of 2. To represent the results, we use ternary plots. A ternary, or simplex plot depicts the ratios of three variables that sum to a constant (the shares of each type of behaviour sum to 1) as positions in a triangle. Each point in the triangle represents one population of firms defined by its distribution over the three strategies. At each position then, one reports the value of a 4th variable which is related to the three others. This way, one is able to represent 4 variables (the 3 variables summing to a constant, and the response variable) in a two-dimensional graph.

Reading a ternary plot requires a little practice. To illustrate, label the corners $\{1, 2, 3\}$. In any corner, all firms have the same strategy, strategy 1 in corner 1, strategy 2 in corner 2 and so on. Moving away from that corner (corner 1 for example), the proportion of firms using strategy 1 falls linearly, until reaching the edge (between the corners 2 and 3) opposite corner 1. Along that edge no firms use strategy 1, and are thus split between strategies 2 and 3. On any locus of points parallel to edge (2,3), all populations have the same share of firms adopting strategy 1.

The outcome variable we report is the identity of the type that most often maximizes performance (over the set of replications run for a given population mix). Each type has a color (blue, green, red). Color opacity increases with the difference between the best performing strategy and the second best performing strategy (in a way, this constitutes a 5th variable). Figure 6 presents 6 panels corresponding to three different IP regimes, crossed with the (im)possibility of trade.

Several things are apparent from this figure. The first is to observe the corners of the plot. Here there is little insight: when all firms are of the same type, then naturally this type is recorded as the dominant strategy. That said, on the interior of the simplex, we observe a clear dominant strategy that depends on the IPR regime: in the base, strong

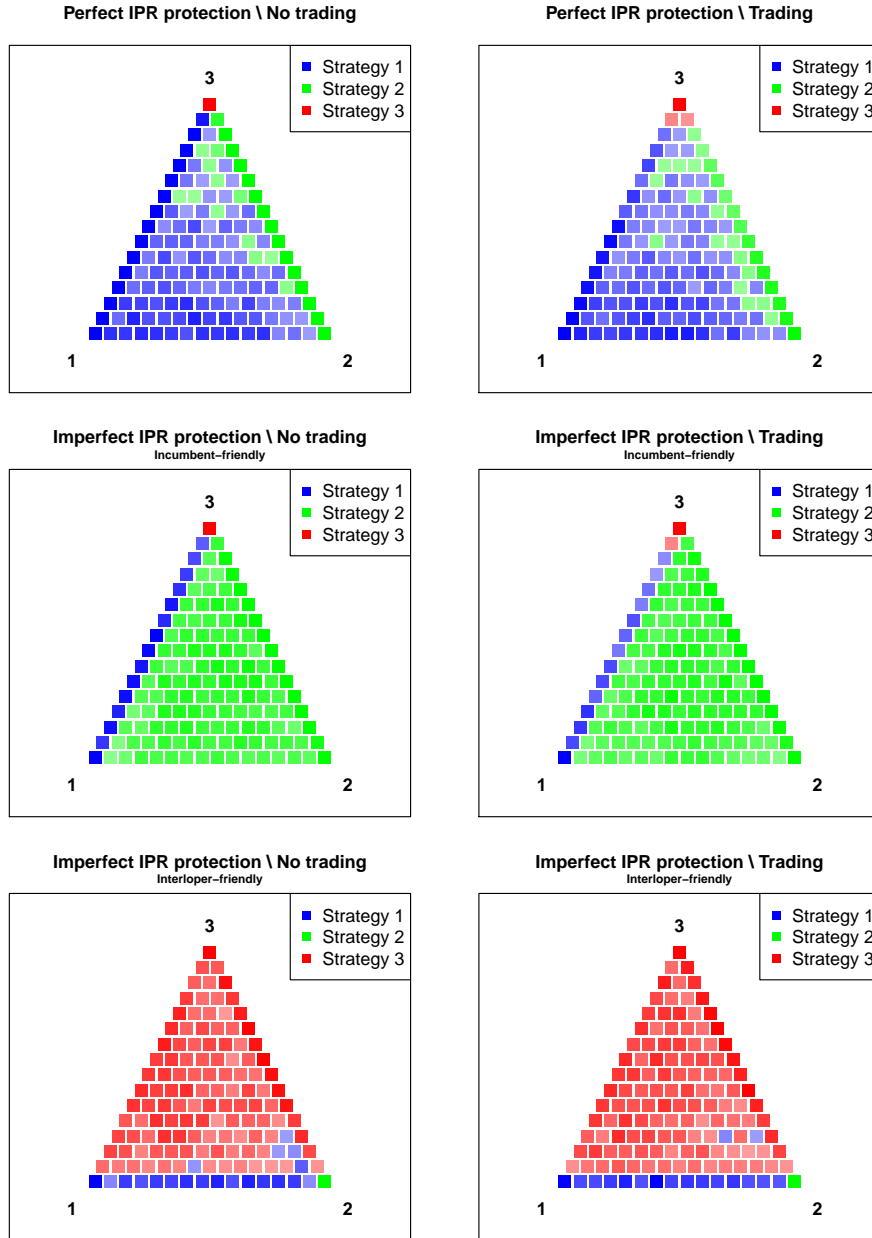


Figure 6: Ternary representations. Each triangle plot presents the simplex of proportions of firms of each type — a point defines three values that sum to 1. Colours indicate which strategy has the highest payoff, averaged over 100 replications; the opacity indicates the ratio of best to second best payoff, lighter shades being closer to one.

blocking case, strategy one dominates; when patent imperfections favour the incumbent strategy two is dominant; and when imperfections favour interlopers it is strategy three. Not only are the different strategies dominant in these cases, but, though we did not structure the experiment specifically to get at it, we can make a strong conjecture that they are also evolutionarily stable. In the middle row, for example, to look at the clearest case, even starting from zero firms of type two, if even one firm changes strategy to type one, it will have the highest payoffs. Assuming firms (eventually) imitate good

performers, in time all firms will switch to strategy one. Similarly, though sometimes not quite as stark, in the other cases the dominant strategy yields the highest payoff even if there are very few firms of that type in the population, leading us to conclude that the dominant strategy is also evolutionarily stable.

As discussed above, when incumbents can ignore patents that might block the development of their products, it is effective to search near one’s own patents, thereby discovering more improvement patents. Searching around others’ patents may prove futile if the incumbent is “lucky” and can ignore a potentially blocking patent. By contrast in the third regime, where an interloper can grab (in our case) all the rents, an easy way to make large profits is to find a technology that improves an existing product (your rivals having done all the previous hard work, you swoop in at the last minute and grab everything). We can also observe that trading makes little difference to which strategy dominates in general. The one possible exception is in the base case with strong blocking and trading, where if there are few firms of types one and two, type three can still be attractive. This follows from the discussion of trading in section 5.2.

A final observation regarding the two imperfect protection regimes: when the dominant firm-type is absent (type 2 in the middle row; type 3 in the bottom row), type one firms achieve the highest payoffs (see the 1 – 3 edge in the middle row; the 1 – 2 edge in the bottom row). The difference between the type-one strategy and the others is that type one firms explore the space much more than the other types do. Thus here we observe that in the absence of the strategy that is finely tuned to the regime, exploration, which is effective in discovering new pioneering patents and thus new products, is particularly valuable.

6.1 Aggregate performance of heterogeneous populations

Above we observed that each IPR regime has a different dominant strategy. In Figure 7 we show a measure of aggregate payoffs in a similar ternary representation as in Figure 6. Each point in the simplex represents one population distribution, and the grey scales indicate aggregate payoffs, with darker shades representing higher payoffs. As a measure of performance we use the total number of patents exploited at the end of the simulation, that is, when all technologies have been found. One striking pattern stands out: the more type three firms there are in the population, the lower the aggregate payoffs. The reason is clear: the type three strategy is to block others’ product development. Unless this can be circumvented by trading or IP violation, the more of this activity, the less product development there will be.¹² While it appears that overall, the third IPR regime has highest payoffs (and even the worst population there has higher payoffs than the best

¹²Recall that even in the second and third regime, $\beta = 50\%$ percent of the time patents are enforced, and an interloping patent blocks product development rather than permitting it through IPR violation.

population under other regimes), this should be treated with caution. The extent to which patents block depends heavily on the parameter β , the strength of the IPR enforcement regime. We have set $\beta = 0.5$ for the second and third regimes, implying that patent rights can be violated very frequently. These violations are what permits further product development. If violations were less common (higher β) this would happen less frequently, and more products would be blocked, lowering payoffs in the second and especially the third, regime.

In essence what we observe here is a conflict between optimal and evolutionarily stable individual strategies on the one hand, and social welfare on the other. Again this must be interpreted carefully, because in some sense this is a static model. Firms change neither their search strategies nor their R&D strategies. Considered more generally, blocking is a failure of two firms to come to some licensing agreement, and our regimes represent different distributions of bargaining power over the terms of that agreement. In principle, provided property rights are well established, a bargain should always be struck (even if it leaves the original owner with epsilon). However the predicted nature of these bargains will have an impact on the incentives for firms to become product originators (incumbents in our terms), and clearly the more power held by the interlopers, the less incentive there is to become an incumbent. This will slow down innovation overall, and points to the concerns held by Paul David at the time.

7 Discussion and Conclusion

We have presented a simple model in which firms search a technology space for innovations that could create new products or add value to existing products. Search strategies, where in the technology space firms look, are linked to underlying patenting strategies. In the model we have three types of firms: those who simply search for interesting technologies; those who protect their own piece of the space; and those who attack the space of other firms. the first type of firm is simply engaged in non-strategic search for good, that is to say patentable, ideas. In the literature the second type of firm is described as engaged in offensive blocking — its aim is to prevent other firms from encroaching on its activities; the third type of firm employs defensive blocking — it uses patents as bargaining chips either to extract rents from others' existing products, or for use in patent trading. In a population ecology where all firm types are present, which is the preferred strategy depends on structures of the IPR system. If the system is very strong, and rigid, meaning that a single “interloping” patent will block all further development of a product, type 1 firms dominate. They acquire the largest number of exploitable patents. The introduction of patent trading (so firms who block each other can swap ownership of the critical patents) improves performance of all types, but most dramatically type three firms. Nonetheless, even with trading, simple search for good ideas is the dominant

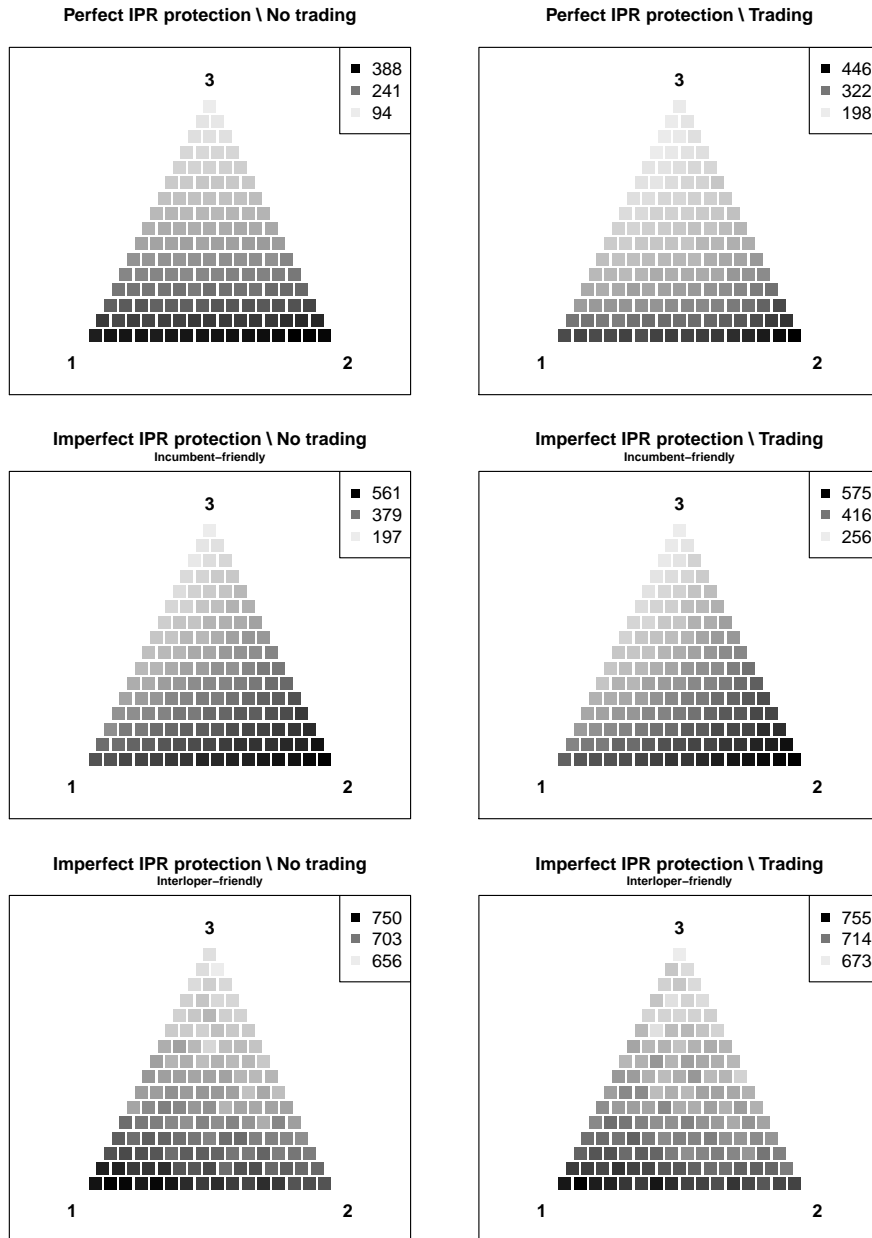


Figure 7: Ternary representations. Each triangle plot presents the simplex of proportions of firms of each type — the gray hue indicates aggregate performance (total payoff).

strategy.

Which strategy dominates, however, changes in response to changes in the patent system. We have avoided discussion of weak and strong patents and instead returned to basics, focusing on how the system rewards patent holders. In the model a product develops from a “pioneering” patent by the addition, in a pre-determined sequence, of improvement patents (characterised as adding features to the product). But what happens when two firms own patents in the same product sequence?

Above we referred to discussions of a change in regime taking place in the 1980s and 1990s, particularly, or perhaps starting, in the US. There was an increase in the ability of

the owner of a single patent to extract rents from the producer of a (complex) product. We try to capture that in the two “imperfect” regimes we introduce here.

In the first, perhaps representing an extreme version of the world prior to the 1980s, the owner of the pioneering patent can (with some probability) violate the rights of others who own patents in the sequence. The system¹³ favours the incumbent firm — the firm that owns the pioneering patent and thus initiates development of the product — making it difficult for owners of a single patent to impeded further development. Here, type two firms dominate: the best strategy is to protect one’s own patents. At the other extreme, representing the world after 1990, a firm that holds a patent relevant to another firm’s product has hold-up power and so can extract considerable (in our model all) rents accruing to the product. Courts became much more sympathetic to holders of single patents and granted them very strong rights. Here the dominant strategy is to invade the technology space of other firms hoping either for a trading opportunity, or to capture rents of existing products.

In the second part of the paper we look at population ecology, using different distributions of firms over types. The results are consistent, with the dominant strategies identified above remaining dominant regardless of the population composition. It remains the case, though, that in either “imperfect IP protection” regime, if no firm pursues the dominating strategy, the naive (non-strategic?) strategy of simply looking for good ideas pays off better than the other.

In the literature, “weak” has been used as an adjective both for patents and for patent enforcement. In the former, a “weak patent” is one that might be successfully challenged, and thus denied, or withdrawn. In the latter, “weak enforcement” refers to the willingness of courts to uphold patent rights. Here we are mostly concerned with the latter. However what we observe in the model is that “weak” or “strong” enforcement is not a subtle enough distinction to understand how firms respond to different IPR regimes. What matters is how weakness is effected in practice. Specifically, the issue turns on how rents are distributed when several firms own patents relevant to a single product. This makes a large difference to how firms understand and implement IPR strategies. One could characterise this in terms of bargaining power. Details of the IPR system will give different firms more or less bargaining power in different situations. Though we have modelled very extreme versions in our first case, all patent holders have equally, very strong bargaining power (and use it aggressively) and cannot come to a rent-sharing agreement. In the second case, the system creates bargaining power for the incumbent.

¹³ We speak deliberately of “the system” so as to avoid the issue of why or how the rents are distributed in a particular way. Rent distribution will be affected by the legal structure, the way laws are written, enforcement mechanisms, attitudes of the courts, willingness of firms to fight or to litigate and so on. Which of these mechanisms drives the reward structure is beyond our scope, though definitely an interesting issue.

Here one can interpret this as one firm owning many patents pertinent to a product, and the system not giving hold-up power to a firm owning a few, possibly even unimportant, patents. In our third case, that sort of hold-up power is granted, and interlopers are able to extract large rents. What this draws attention to is the fact that “weak” or “strong” is not a simple, fixed property of a patent but rather the outcome of a combination of things involving many aspects of the intellectual property rights system, as pointed out in footnote 13.

Apart from our base case of extremely strong (not to say rigid) patents and enforcement, both regimes we model have weak patent enforcement. But the weakness affects different firms (incumbent versus interloper as we have styled them) differently in the two cases. And as one might expect, optimal strategies are different in the two cases.

A Appendix

In this appendix, we provide the details of search strategies and trading in the model.

Each firm $i = 1, 2, \dots, n$ starts the search process from a random initial pioneering idea s_i^0 , which also corresponds to the initial set of patents held by that firm, $\mathcal{P}_i^0 = \{s_i^0\}$. After this initialization, there is an exogenous random order in which firms act sequentially — a sequence of firm identities. Each firm appears the same number of times in the sequence, and a replication is a randomly generated sequence. When called to act, a firm explores the space of ideas by nearest neighbour moves on the two dimensional lattice (see Figure 1, left panel), starting from some discovered idea. The firm patents any new (in the sense of yet undiscovered) idea it finds. A discovered idea can be connected to an existing product if a) it is part of the product as defined in the middle layer, and b) the idea preceding it in the product sequence is part of the product. Any idea can belong to none, one or more than one product. Since there is one-to-one relationship between products and pioneering ideas, it is equivalent to saying that an idea can be related to none, one or more than one pioneering idea. The identity of the pioneering idea(s) to which an idea s is related is denoted $\mathcal{A}(s)$, where $\mathcal{A}(s) = \{\}$ if the idea is not part of any product, and might return more than one element for ideas with multiple product membership.

Consider now an arbitrary period $t \geq 1$. By that time, firms have undertaken R&D activities which have led to discoveries, so that by the end of period t , ideas $\{s^1, \dots, s^t\}$ have been discovered and patented, in addition to the initial pioneering patents. Firm i has patented a set \mathcal{P}_i^t of ideas, which contains at least s_i^0 and might contain other ideas if the firm has been called to act during the first t periods. The set of all patented ideas is $\mathcal{P}^t = \bigcup_{i=1}^n \mathcal{P}_i^t = \{s_1^0, \dots, s_i^0, \dots, s_n^0, s^1, \dots, s^t\}$, with $\mathcal{P}_i^t \cap \mathcal{P}_j^t = \{\}$ for any pair i, j (patents are uniquely owned).

The three archetypal R&D strategies we consider are represented by different search heuristics. Consider the beginning of the period $t + 1$, when the firm needs to decide about its starting point for the period's R&D efforts.

Starting points for search

All firms explore the idea space through local search, i.e., by considering the neighbourhood of patents that are already discovered. Local search is natural — inventors mostly search in the neighbourhood of what they know (which is not necessarily what they have discovered) — and it is effective in the model because products are made of locally clustered ideas. As firms can identify the product membership of an idea (provided this idea is connected to the pioneering ideas), searching iteratively around the currently identified

ideas making up a product will (eventually) uncover all the product's constitutive ideas (see Figure 1, right panel and the discussion in Subsection 3.1 about the development of product (29 – 19 – 30 – 9)). However, in the case that searching around the currently identified ideas making up a product cannot lead to an improvement, the firm searches elsewhere (for instance, when the owners of ideas 29 and 19 differ in Figure 1, the owner of 29 is blocked).

Three strategies

Type one A type 1 firm i first attempts to improve its products (that is products for which it owns the pioneering patent). To find a starting point for search it considers the set $\mathcal{S}_i = \{s \in \mathcal{P}^t : \mathcal{A}(s) \in \mathcal{P}_i^t\}$. For any idea $s \in \mathcal{S}_i$, i assesses whether improvement of the corresponding product is possible: it checks whether product $\mathcal{A}(s)$ is blocked by another firm $j (\neq i)$. If \mathcal{A} is blocked, s is removed from S which becomes $\mathcal{S}'_i = \mathcal{S}_i - \{s\}$. Idea s is also removed if \mathcal{A} has been fully developed (which firms can identify). When all candidate ideas have been considered and the final $\mathcal{S}'_i \neq \{\}$, the firm draws an idea in the nearest-neighbourhood of \mathcal{S}'_i . If this neighbourhood only contains already discovered patents, then the nearest-neighbourhood of this set is sampled, and so on if necessary.

If the final $\mathcal{S}'_i = \{\}$, i.e., there are no product improvement opportunities for the firm, it moves away from ideas belonging to products and focuses on exploration. It does so by removing from \mathcal{P}^t all ideas currently exploited in products owned by all firms, i.e., it determines $\mathcal{P}^t - \bigcup_{i=1}^n \mathcal{S}_i$ and draws an idea in the nearest-neighbourhood of this set.

Lastly, if the set $\mathcal{P}^t - \bigcup_{i=1}^n \mathcal{S}_i$ is empty, firm i simply draws an idea in the nearest-neighbourhood of \mathcal{P}^t , and hopes for the best.

Type two A type 2 firm i searches around its own patent set, i.e., it considers the set $\mathcal{S}_i = \mathcal{P}_i^t$. It does so regardless of any consideration of possible product improvement, but removes any idea s that belongs to any product that has been maximally developed, so that \mathcal{S}_i becomes $\mathcal{S}'_i = \mathcal{S}_i - \{s\}$. When all candidate ideas have been considered and the final $\mathcal{S}'_i \neq \{\}$, the firm draws an idea in the nearest-neighbourhood of \mathcal{S}'_i . If this neighbourhood only contains already discovered patents, then the nearest-neighbourhood of this set is sampled, and so on if necessary.

If the final $\mathcal{S}'_i = \{\}$, i.e., there are no more ideas to discover in the neighbourhood of i 's ideas, it moves away from ideas belonging to its products and focuses on exploration. It does so by removing from \mathcal{P}_i^t all the ideas it owns, i.e., it determines $\mathcal{P}^t - \mathcal{P}_i^t$, removing ideas which belong to maximally developed products, and draws an idea in the nearest-neighbourhood of this set.

Lastly, if there are no more improvement ideas to discover in the neighborhood of any firm's ideas, firm i simply draws an idea in the nearest-neighbourhood of \mathcal{P}^t , and hopes for the best.

Type three A type 3 firm i does the opposite of the type 2 firm, and searches around other firms' patents, i.e., it considers first the set $\mathcal{S}_i = \mathcal{P}^t - \mathcal{P}_i^t$. It does so regardless of any consideration of possible product improvement, but again removes any idea s that belongs to a product that has been maximally developed, so that \mathcal{S}_i becomes $\mathcal{S}'_i = \mathcal{S}_i - \{s\}$. When all candidate ideas have been considered and the final $\mathcal{S}'_i \neq \{\}$, the firm draws an idea in the nearest-neighbourhood of \mathcal{S}'_i . If this neighbourhood only contains already discovered patents, then the nearest-neighbourhood of this set is sampled, and so on if necessary.

If the final $\mathcal{S}'_i = \{\}$, i.e., there are no more ideas to discover in the neighbourhood of other firms' ideas, it focuses on exploration. It does so by considering its own patent set, setting $\mathcal{S}_i = \mathcal{P}_i^t$, removing ideas which belong to maximally developed products, and draws an idea in the nearest-neighbourhood of this set.

Lastly, if there are no more improvement ideas to discover in the neighborhood of any firm's ideas, firm i simply draws an idea in the nearest-neighbourhood of \mathcal{P}^t , and hopes for the best.

Trading

When patent trading is allowed, at the end of each period all blocking patents are identified. A patent s owned by j is "*blocking to i* " if $\mathcal{A}(s) \in \mathcal{P}_i^t$ and all the patents in the sequence running from $A(s)$ to s are exploited (except s). Then, mutually beneficial trades are identified. A mutually beneficial trade exists when there is a patent s owned by j that is blocking to i , and a patent owned by i that is blocking to j . If several such mutually beneficial trades exist, one is randomly selected and patent ownership is swapped, so that i becomes the owner of s and j becomes the owner of s' .

A replication lasts until all profitable ideas (ideas that are part of products) are discovered. Then the race stops and firm profits are computed.

References

- [1] Anton, J. J. and Yao, D. A. (2004). Little Patents and Big Secrets: Managing Intellectual Property. *RAND J. Econ.*, 35(1):1.
- [2] Archibugi, D. and Planta, M. (1996). Measuring technological change through patents and innovation surveys. *Technovation*, 16(9):451–519.
- [3] Arundel, A. and I. Kabla, (1998) What percentage of innovations are patented? empirical estimates for European firms, *Res. Policy*, Vol. 27(2): 127-141,
- [4] Blind, K., Cremers, K., and Mueller, E. (2009). The influence of strategic patenting on companies' patent portfolios. *Res. Policy*, 38(2):428–436.
- [5] Blind, K., Edler, J., Frietsch, R., and Schmoch, U. (2006). Motives to patent: Empirical evidence from Germany. *Res. Policy*, 35(5):655–672.
- [6] Blind, K., Pohlisch, J., and Zi, A. (2018). Publishing, patenting, and standardization: Motives and barriers of scientists. *Res. Policy*, 47(7):1185–1197
- [7] Boldrin, M. and Levine, D. K. (2013). The case against patents. *J. Econ. Perspect.*, 27(1):3–22.
- [8] Chu, A. C. (2009). Effects of blocking patents on R&D: a quantitative DGE analysis. *J. Econ. Growth*, 14(1):55–78.
- [9] Chu, A. C., Cozzi, G., and Galli, S. (2012). Does intellectual monopoly stimulate or stifle innovation? *Eur. Econ. Rev.*, 56(4):727–746.
- [10] Clancy, M. S. (2018). Inventing by combining pre-existing technologies: Patent evidence on learning and fishing out. *Res. Policy*, 47(1):252–265.
- [11] Cockburn, I. M., MacGarvie, M. J., and Müller, E. (2010). Patent thickets, licensing and innovative performance. *Ind. Corp. Chang.*, 19(3):899–925.
- [12] Cohen, S. L. and Tripsas, M. (2018). Managing technological transitions by building bridges. *Acad. Manag. J.*, 61(6):2319–2342.
- [13] Cohen, W., Nelson, R., and Walsh, J. (2000). Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not). Technical report, National Bureau of Economic Research, Cambridge, MA.
- [14] Crass, D., Garcia Valero, F., Pitton, F., and Rammer, C. (2019). Protecting Innovation Through Patents and Trade Secrets: Evidence for Firms with a Single Innovation. *Int. J. Econ. Bus.*, 26(1):117–156.

- [15] Czarnitzki, D., Hussinger, K., and Leten, B. (2020). How Valuable are Patent Blocking Strategies? *Rev. Ind. Organ.*, 56(3):409–434.
- [16] Desyllas, P., Miozzo, M., fen Lee, H., and Miles, I. (2018). Capturing Value from Innovation in Knowledge-Intensive Business Service Firms: The Role of Competitive Strategy. *Br. J. Manag.*, 29(4):769–795.
- [17] Drivas, K., Economidou, C., Karamanis, D., and Zank, A. (2016). Academic patents and technology transfer. *J. Eng. Technol. Manag. - JET-M*, 40:45–63.
- [18] Duffy, J. F. (2008). A timing approach to patentability. *Lewis Clark law Rev.*
- [19] Dunlavey, D. C. (1955). Protection of the Inventor outside the Patent System. *Calif. Law Rev.*, 43(3):457.
- [20] Erkal, N. (2005). The decision to patent, cumulative innovation, and optimal policy. *Int. J. Ind. Organ.*, 23(7-8):535–562.
- [21] Gallini, N. T. (2002). The economics of patents: Lessons from recent U.S. patent reform. *J. Econ. Perspect.*, 16(2):131–154.
- [22] Gans, J. S., Murray, F. E., and Stern, S. (2017). Contracting over the disclosure of scientific knowledge: Intellectual property and academic publication. *Res. Policy*, 46(4):820–835.
- [23] Granstrand, O. (1999). *The Economics and Management of Intellectual Property — Towards Intellectual Capitalism*. Edward Elgar.
- [24] Grimpe, C. and Hussinger, K. (2014). Pre-empted patents, infringed patents and firms’ participation in markets for technology. *Res. Policy*, 43(3):543–554.
- [25] Hall, B. H. (2004). Exploring the Patent Explosion. *J. Technol. Transf.*, 30(1-2):35–48.
- [26] Hall, B. H. and Ziedonis, R. H. (2001). The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995. *RAND J. Econ.*, 32(1):101.
- [27] Hall, B. H., von Graevenitz, G., and Helmers, C. (2021). Technology entry in the presence of patent thickets. *Oxf. Econ. Pap.*, 73(2):903–926.
- [28] Harhoff, D., Hall, B. H., von Graevenitz, G., Hoisl, K., Wagner, S., Gambardella, A., and Giuri, P. (2007). The strategic use of patents and its implications for enterprise and competition policies. Technical report, European Commission, Brussels.

- [29] Harhoff, D., von Graevenitz, G., and Wagner, S. (2016). Conflict Resolution, Public Goods, and Patent Thickets. *Manage. Sci.*, 62(3):704–721.
- [30] Hellmann, T. and Perotti, E. (2011). The circulation of ideas in firms and markets. *Manage. Sci.*, 57(10):1813–1826.
- [31] Helmers, C. and Rogers, M. (2011). Does patenting help high-tech start-ups? *Res. Policy*, 40(7):1016–1027.
- [32] Holgersson, M. (2013). Patent management in entrepreneurial SMEs: a literature review and an empirical study of innovation appropriation, patent propensity, and motives. *R&D Manag.*, 43(1):21–36.
- [33] Holgersson, M. and Granstrand, O. (2017). Patenting motives, technology strategies, and open innovation. *Manag. Decis.*, 55(6):1265–1284.
- [34] Holgersson, M., Granstrand, O., and Bogers, M. (2018). The evolution of intellectual property strategy in innovation ecosystems: Uncovering complementary and substitute appropriability regimes. *Long Range Plann.*, 51(2):303–319.
- [35] Horstmann, I., MacDonald, G. M., and Slivinski, A. (1985). Patents as Information Transfer Mechanisms: To Patent or (Maybe) Not to Patent. *J. Polit. Econ.*, 93(5):837–858.
- [36] Jell, F., Henkel, J., and Wallin, M. W. (2017). Offensive Patent Portfolio Races. *Long Range Plann.*, 50(5):531–549.
- [37] Justia (2022). Improvement and New Use Patents. <https://www.justia.com/intellectual-property/patents/types-of-patents/improvement-and-new-use-patents/>.
- [38] Klein, M. A. (2020). Secrecy, the patent puzzle and endogenous growth. *Eur. Econ. Rev.*, 126:103445.
- [39] Kortum, S, and J. Lerner (1998). Stronger Protection or Technological Revolution: What Is Behind the Recent Surge in Patenting?” *Carnegie-Rochester Conference Series on Public Policy*, Vol. 48:247- 304.
- [40] Krasteva, S. (2014). Imperfect Patent Protection and Innovation. *J. Ind. Econ.*, 62(4):682–708.
- [41] Laursen, K. and Salter, A. J. (2014). The paradox of openness: Appropriability, external search and collaboration. *Res. Policy*, 43(5):867–878.

- [42] Lerner, J. (1995) Patenting in the Shadow of Competitors. *Journal of Law and Economics*, Vol. 38:463-495.
- [43] Levin, R., Klevorick, A., Nelson, R.R., and Winter, S.G. (1987) "Appropriating the returns from industrial R&D", *Brookings Papers on Economic Activity*, 783-820.
- [44] Merges, R. P. and Nelson, R. R. (1990). On the Complex Economics of Patent Scope. *Columbia Law Rev.*, 90(4):839.
- [45] Mihm, J., Sting, F. J., and Wang, T. (2015). On the effectiveness of patenting strategies in innovation races. *Manage. Sci.*, 61(11):2662–2684.
- [46] Mullin, J. (2017) "Many patent-holders stop looking to East Texas following Supreme Court ruling" - 10/12/2017 *Ars Technica* <https://arstechnica.com/tech-policy/2017/10/patent-cases-in-east-texas-plunge-more-than-60-percent/>
- [47] Panagopoulos, A. and Park, I. U. (2018). Patents As Negotiating Assets: Patenting Versus Secrecy For Startups. *Econ. J.*, 128(615):2876–2894.
- [48] Pisano, G. (2006). Profiting from innovation and the intellectual property revolution. *Res. Policy*, 35(8 SPEC. ISS.):1122–1130.
- [49] Reitzig, M. (2004). The private values of 'thickets' and 'fences': towards an updated picture of the use of patents across industries, *Economics of Innovation and New Technology*, 13:457–476.
- [50] Schneider, C. (2008). Fences and competition in patent races. *Int. J. Ind. Organ.*, 26(6):1348–1364.
- [51] Scotchmer, S. (1991). Standing on the Shoulders of Giants: Cumulative Research and the Patent Law. *J. Econ. Perspect.*, 5(1):29–41.
- [52] Shapiro, C. (2000). Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting. *Innov. Policy Econ.*, 1:119–150.
- [53] Silverman, A. B. (1995). The relationship between basic and improvement patents. *JOM*, 47(1):50– 50.
- [54] Suzuki, K. (2015). Economic growth under two forms of intellectual property rights protection: patents and trade secrets. *J. Econ.*, 115(1):49–71.
- [55] Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Res. Policy*, 15(6):285–305.

- [56] Ouellette, L. L. (2012). Do Patents Disclose Useful Information? *Harv. J.L. Tech.*, 25(2):545.
- [57] UpCounsel (2020). Improvement Patents : Everything You Need to Know. <https://www.upcounsel.com/improvement-patents>.
- [58] von Graevenitz, G., Wagner, S., and Harhoff, D. (2013). Incidence and Growth of Patent Thickets: The Impact of Technological Opportunities and Complexity. *J. Ind. Econ.*, 61(3):521–563.
- [59] Walsh, J. P., Lee, Y. N., and Jung, T. (2016). Win, lose or draw? the fate of patented inventions. *Res. Policy*, 45(7):1362–1373.
- [60] Ziedonis, R. H. (2004). Don't fence me in: Fragmented markets for technology and the patent acquisition strategies of firms. *Manage. Sci.*, 50(6):804–820.