Property Rights and Frictions in the Sale of Patents

Jeffrey M. Kuhn*

November 6, 2016

Abstract

Patent scope is central to the sale of ideas, which can spur economic growth and provide significant gains from trade. Awarding an inventor a patent on a new idea partially solves a commitment problem that would otherwise prevent the inventor from selling the idea. (Arrow, 1962). In the absence of a patent, a prospective buyer cannot credibly promise not to steal the idea should the inventor reveal it, while the inventor cannot credibly promise to reveal the idea should the prospective buyer pay for it. A firm's ability to use a particular patent to overcome this transactional hurdle derives from two factors: (1) the scope of the patent's legal right to exclude and (2) the effectiveness of that legal right in providing market exclusivity. I first show that a broader patent is more likely to be sold by employing a causal instrument that provides a plausibly exogenous shock to the scope of a patent's legal right to exclude, holding fixed the underlying idea. I then examine variation in the effectiveness of the right by interacting the instrument with endogenous firm, industry, and market characteristics. These results shed light on how firms profit from innovation and also connect the important but understudied market for patents, widely believed to be illiquid and inefficient, with fundamental research about how markets function in other contexts.

*PhD Candidate, UC Berkeley Haas School of Business, jeffrey_kuhn@haas.berkeley.edu. The author thanks Rui de Figueiredo, Noam Yuchtman, Brian Wright, Kenneth Younge, Robert Merges, Jeffrey Furman, Deepak Somaya, Rosemarie Ziedonis, Neil Thompson, Bo Cowgill, and seminar participants at UC Berkeley, the Consortium on Competitiveness and Competition, the Academy of Management Dissertation Consortium, and the Institutions and Innovation Conference for helpful comments. The Hoover Institution Working Group on Intellectual Property, Innovation, and Prosperity at Stanford University provided generous financial support during the preparation of this paper.

1 Introduction

The market for ideas, which includes the sale and license of patents, is an important part of how the patent system provides incentives for innovation. Moreover, ideas are a key input in a firm's production function, enabling the productive and competitive use of land, labor, and capital. Transactions for patent rights can thus spur economic growth and provide significant gains from trade by enabling the division of innovative labor, facilitating the exploitation of complementary assets, and avoiding the misallocation of resources. (Akcigit, Celik, & Greenwood, 2016; Arora, Fosfuri, & Gambardella, 2004; Serrano, 2011). However, a long line of research finds that the market for ideas suffers from substantial imperfections, and the market for buying and selling patents is thought to be particularly illiquid and inefficient. (Agrawal, Cockburn, & Zhang, 2015; Arora et al., 2004; Teece, 1986).

The prevalence of market frictions raises the concern that exchanges efficient from an institutional perspective may fail to occur, potentially reducing social welfare. On the supply side, market failures in technology transfers may prevent the exploitation of patented technology that a firm develops but does not commercialize. On the demand side, ownership fragmentation of complementary patents is associated with a reduced likelihood of in-licensing and more aggressive patenting within the firm. (Cockburn, MacGarvie, & Mueller, 2010; Ziedonis, 2004). Thompson and Kuhn (2016) found that over 10% of patents duplicate contemporaneous efforts by other firms, which may reflect in part the difficulty in transacting for technology. Even worse, the expectation of transactional difficulties may deleteriously affect market entry decisions and market structure. Indeed, imperfections in the market for ideas have been found to drive both firm commercialization strategies and inter-industry differences in firm boundary decisions. (Gans & Stern, 2003, 2010).

Patent scope is central to the sale of ideas. Awarding an inventor a patent on a new idea partially solves a commitment problem that would otherwise prevent the inventor from selling the idea. (Arrow, 1962). In the absence of a patent, a prospective buyer cannot credibly promise not to steal the idea should the inventor reveal it, while the inventor cannot credibly promise to reveal the idea should the prospective buyer pay for it. A firm's ability to use a particular patent to overcome this transactional hurdle derives from two factors: (1) the scope of the patent's legal right to exclude and (2) the effectiveness of that legal right in providing market exclusivity. In this article I employ a causal instrument to investigate the effect of an exogenous shock to the scope of a patent's legal right to exclude on the likelihood that the patent is sold. I then examine variation in the effectiveness of the legal right by interacting the instrument with endogenous firm, industry, and market characteristics.

This article contributes to a well-established literature on the transfer of technology via the market for ideas, the study of which is crucial for understanding the effects of the patent system on firm strategy and for crafting public policy to improve economic efficiency and provide incentives for innovation. Both patent sales and patent licenses can be used to transfer rights to technology, and a long line of research investigates the determinants of patent licensing. (See, e.g., Agrawal et al., 2015; Arora et al., 2004; Arora & Ceccagnoli, 2006; Arora & Gambardella, 2010; Cockburn et al., 2010; Fosfuri et al., 2006; Mowery et al., 2001). However, because a patent license typically does not transfer to the recipient the right to sue those who infringe the patent, the determinants and consequences of sales and licenses often differ in various ways. Although the literature on the market for ideas provides important insight into transactions for patent rights, causal evidence is limited and the literature does not address the relationship between patent sales and a patent's legal right to exclude. This article investigates what is thus a fundamental and unanswered question—how, when, and under what conditions the scope of the right to exclude affects patent sales.

Patent sales have received less attention from scholars than patent licenses, but a bourgeoning literature highlights the importance of the secondary market for patents. Lamoreaux and Sokoloff (1999, 2001) showed that an active secondary market for patents in the late eighteenth and early nineteenth centuries gave rise to a large group of specialized independent inventors. Serrano (2010) presented descriptive evidence that patent sales are correlated with factors such as patent age, firm size, and forward citation counts. Hochberg, Serrano, and Ziedonis (2014) found a positive relationship between patent sales and the rate of lending to startups. Serrano (2011) quantified the gains to trade in the market for patents based on patent sales and renewal rates, while Akcigit et al. (2016) estimated a structural model and found that patent sales spur economic growth. Bessen, Meurer, and Ford (2011) argued that the explosion in patent litigation over previous decades is attributable in large part to non-practicing entities (NPEs) that purchase and monetize patents through licensing and litigation, but other research suggests that the patents acquired and asserted by NPEs are similar in quality and scope to other litigated patents. (Fischer & Henkel, 2012; Risch, 2011). Finally, Galasso, Schankerman, and Serrano (2013) showed that while patent sales driven by comparative advantage in commercialization increase litigation, sales resulting from advantages in patent enforcement reduce it.

To investigate the role of property rights and frictions in patent sales, this article employs a range of novel measures and data sets developed by the author. The identification strategy is based on work by Kuhn, Roin, and Thompson (2016), who developed a measure of a patent's legal right to exclude by analyzing the full text of the claims of more than 1 million patents. Because patent applications are quasi-randomly assigned to patent examiners within an area of technology and because patent examiners vary in their average effect on claim scope, a patent's examiner serves as a causal shock to the right to exclude provided by the patent relative to the underlying idea. To measure the effect of this shock, patent sales data were constructed by analyzing more than 4 million USPTO patent assignment agreements, which patent owners used to record changes to ownership rights. Measures of patent congestion, invention familiarity, and invention originality were constructed using patent-to-patent textual similarity data developed by Younge and Kuhn (2015, 2016), who employed a vector space similarity model to compare the full text of every patent to the full text of every other patent.

I show that receiving a broader patent increases the likelihood that the patent is sold, likely reflecting the fact that the sale of a broader exclusionary right provides greater gains to trade that are more likely to overcome the considerable transaction costs associated with patent sales.

However, the magnitude of the effect depends substantially upon firm, technology, and market characteristics, likely reflecting variation in the practical effectiveness of patent rights across these dimensions for providing market exclusivity. In particular, the scope of the right to exclude is more important for smaller firms and for patents in thicker and less congested markets. I also demonstrate that the between-industry variation in the effect of patent scope on patent sales aligns broadly with survey evidence described by Cohen, Nelson, and Walsh (2000) regarding industry-level patent importance and efficacy. Importantly, the likelihood that a patent is sold depends on the interaction of the realized scope of the legal right to exclude with the reported effectiveness of patents for exclusion in an industry. Likely for this reason, the right to exclude plays an outsized role in facilitating transfers in mechanical, chemical, and medical device industries and a relatively minor role in the semiconductor and electrical components industries.

This article makes several contributions to the literature on the market for ideas. First, it provides the first causal evidence of the relationship between patent sales and the scope of the legal right to exclude. Second, it demonstrates how the effectiveness of the right to exclude, and hence the strategic role patents play for firms, varies across firm and industry characteristics. Third, it links the market for patents with fundamental research on market design by illustrating interactions between market characteristics and the causal effect of the right to exclude. Fourth, it introduces new data and measures to provide current and comprehensive descriptive evidence of the extent and characteristics of the market for patents. Together these results have important positive implications for firm strategy since they illustrate ways in which a firm's ability to acquire, sell, and exploit innovation depend on market and intellectual property characteristics. At the same time, these findings have important normative implications for how the patent system may be modified to better incentivize innovation and allocate new ideas to where they may be most efficiently exploited.

Section 2 describes why understanding the determinants of patent sales is important for scholars of firm strategy and innovation economics. Section 3 develops testable hypotheses based on a discussion of the real-world implications of the right to exclude. Section 4 describes the empirical specification, including the econometric model, data and measures, and summary statistics. Section 4 also discusses the causal instrument used to investigate the effect of receiving a broader patent on the likelihood of a sale and introduces new measures of patent rights congestion, patent market thickness, technological originality, and technological familiarity. Section 5 presents the results from estimating the empirical model, which broadly support the hypotheses. Section 6 discusses the results and concludes.

2 The Importance of Patent Ownership

Ronald Coase argued that a production input is incorporated into a firm when the transaction cost associated with coordinating production in the market, under imperfect information, is greater than within the firm. (Coase, 1937). In many cases, economies of scope mean that multiple products may be produced more efficiently via joint production using shared inputs. (Williamson, 1975). These forces are particularly salient in the development of new ideas, the production and commercialization of which frequently requires an array of complementary assets. (Teece, 1986). Joint production is optimally organized within the same firm if the economies of scope are based on the ongoing use of proprietary knowhow. (Teece, 1980, 1982). The importance of organizational knowledge explains why a firm needs access to many ideas, but not why the firm needs to own the ideas outright. Indeed, a belief that contracting is sufficient to facilitate the efficient flow of rights seems implicit in much of the literature on the market for ideas, given that the literature to date has focused primarily on patent licensing. Nevertheless, several important distinctions between patent licensing and patent sales suggest that patent sales warrant independent investigation.

Most fundamentally, a license does not generally transfer the right to sue. Patent law is grounded primarily on negative, exclusionary rights rather than positive rights. For example, a patent does not provide a right to practice the claimed invention—a commercial product or process may be covered by potentially many different patents. Instead, a patent provides the right to exclude others from making, using, or selling the claimed invention. In the same way, a patent license transfers to the recipient not the right to sue, but rather the right to not be sued by the patent owner. However, patent law generally that any law suit alleging infringement of a patent be instigated by the patent's owner. A licensee only has standing to file a lawsuit without the patent owner's participation if the license is not only exclusive but also transfers "all substantive rights" associated with the patent to the license holder. Courts have interpreted these requirements broadly, holding that even field of use restrictions are enough to deny standing to the licensee. (*Luminara Worldwide, LLC. V. Liown Elecs. Co.*, Fed. Cir. Feb. 29, 2016). A licensee can in some circumstances initiate a lawsuit in conjunction with the patent's owner, but the difficulty of contractually resolving in advance the incentives and responsibilities inherent in a complex and prolonged patent infringement suit hinder such arrangements. For instance, the patent owner may have disincentives to incur costs, exert effort, sue its own actual or potential customers, or suffer negative publicity associated with a patent suit. Faced with a high and somewhat ambiguous bar to judicial standing, significant negotiation costs, and potentially conflicting incentives, a prospective licensee has significant motivation to instead acquire the patent outright if exercising the legal right to exclude is likely to be an important component of the licensee's strategy.

The United States Patent and Trademark Office ("USPTO") evaluates each patent application via a complex examination process, but rigorously examining each patent application to the point of certainty would be an inefficient use of patent office resources given that few patents are ever litigated. (Lemley, 2001). Patents are thus probabilistic in the sense that the validity, scope, and infringement of a patent are all unclear until each issue has been carefully litigated, and different actors may estimate these probabilities differently. (Lemley & Shapiro, 2005; Merges, 1999). The probabilistic nature of patents exacerbates the contracting costs of negotiating a license and underscores the importance of patent litigation. In the presence of uncertainty, rights holders may choose to establish a reputation for toughness by demonstrating the willingness and ability to engage in suit. (Agarwal, Ganco, & Ziedonis, 2009). A reputation for toughness can only be established on the basis of a meaningful threat.

In practice, very few patents are ever litigated, and scholars frequently study other uses of patents such as signaling, collaboration, and knowledge transfer. Nevertheless, the efficacy of these non-litigation uses for patents is grounded in the threat of litigation since patents encourage the free exchange of information by discouraging appropriation without consent. Acquiring the right to sue is an important motivator for patent sales because it facilitates a key strategy for many patent holders. Faced with a fragmented web of property rights spread across many different owners, a firm intending to commercialize in an area of technology may need to acquire patents for defensive purposes. The firm can then respond to litigation, actual or threatened, with a countersuit using its own patents. The threat of mutual holdup facilitates a faster resolution to the potential standoff. (Somaya, 2002, 2003, 2012).

The differences in the legal effect of patent sales and patent licensing drives a distinction in their strategic implications. Both are ways for the original patent owner to generate revenue. However, patent sales are in theory driven by the most efficient owner of the right, while patent licensing is in theory driven by the most efficient use of the right. For instance, licensing may be an appropriate strategy when a patent protects a general purpose technology with applications in a variety of non-competing contexts that are not embodied in a single firm. In contrast, if a patent protects an idea whose exclusive use would provide a significant competitive advantage but whose development would require extensive investment and a range of complementary assets, then the most efficient course of action may be to sell the patent to the firm best equipped to exploit it. I do not observe licensing and therefore cannot directly investigate the trade-offs between licensing and assignments in this article. Nevertheless, as discussed in the following section, the extensive literature on patent licensing provides important insights that inform my analysis of patent sales.

3 Institutions and Hypotheses

The patent system exists in part to solve a fundamental commitment problem that hinders the sale of new ideas. In the absence of patent protection, the prospective buyer cannot commit to paying for an idea if the idea is revealed before payment is received, while the prospective seller cannot commit to revealing the idea if payment is received before the idea is revealed. (Arrow, 1962). A patent reduces the danger of expropriation by granting the patent holder a tradable right to exclude others from using the idea. However, the patent system provides only an imperfect solution to the danger of expropriation, for two reasons. The first, addressed in Section 3.1, is that the scope of the legal right to exclude provided by the patent may fail to completely encompass the underlying idea. The second, addressed in Section 3.2, is that the practical effectiveness of the legal right to exclude varies with firm, industry, and market characteristics.

3.1 The delineation of the right to exclude

To secure a patent, an applicant files a patent application that includes both a detailed description of the invention and a set of claims that delineate the requested legal right to exclude. The USPTO examines each patent application both to weed out unwarranted applications and to ensure that patents are not granted with overly broad exclusionary rights. In practice the decision as to whether to grant a patent at all affects only patent applications on marginal ideas. After all, about 75% of patent applications ultimately issue as patents, and of the ones that do not issue nearly half are abandoned for non-technological reasons such as financial constraints. (Lemley & Sampat, 2008). In contrast, the delineation of the patent's right to exclude is the defining outcome of the patent examination process. (Kuhn et al., 2016). Indeed, for most patent applications the examination process takes the form of a negotiation between the patent applicant and patent examiner over the scope of the patent's right to exclude (i.e. the patent's claims).

In a typical patent examination, the patent examiner rejects the applicant's initial claims on the basis that the claimed idea is either not new or is obvious in view of existing inventions. Although the detailed description of the idea is fixed at the time of filing the patent application, the patent applicant can usually overcome the initial rejection by amending the claims to encompass a narrower exclusionary right. This process of rejection and response may be repeated indefinitely until the examiner agrees to issue the patent or the applicant abandons the applications. In practice, about 70% of all patents are granted with claims that are narrower than those originally filed with the patent application. (Kuhn et al., 2016). Importantly for this paper's identification strategy, patent applications are quasi-randomly assigned to patent examiners within a particular area of technology, and patent examiners vary considerably in the degree to which they cause applicants to narrow their claims. (Kuhn et al., 2016).

The literature offers few clear theoretical predictions about the effects of receiving a broader or narrower right to exclude. Merges and Nelson (1994) argue that awarding broad patent claims may hinder technological progress, but that the effects will vary based on underlying features of the technology such as the extent to which rapid innovation requires diverse information inputs. However, they do not predict the effect of scope on the likelihood that a patent is sold. Empirically, Gans et al. (2008) show that resolving uncertainty about the scope of a patent under examination increases the likelihood that the patent will be licensed, but whether these licenses resulted from a reduction in information asymmetries or a higher-than-expected realized patent scope is unclear.[1] Further, sample size constraints and fundamental differences between licensing and sales limit the

[1] It should be noted that patent claims are subject to several types of uncertainty. Gans et al. (2008) investigate the effect of resolving one type—what the claims are—which is not determined until the patent office indicates to the applicant that patent application is ready to be issued. Of course, even after a patent is issued, uncertainty remains regarding both whether a court would invalidate the claims (e.g., if new prior art is unearthed) and whether the claims are being infringed. Applicants deal with this second type of uncertainty in part by including in the patent "dependent" claims that are narrower than the patent's main, independent claims and that serve as fallback positions in the event that the independent claims are invalidated. Because of these fallback positions, an applicant granted a broader patent generally does not give up the narrower protection, which means that a broader patent is strictly better for the patent owner than a narrower one even if broader claims are at greater risk of being invalidated. (Kuhn et al., 2016).

applicability of their result to the sale of patents across firms, time, and industries.

Prospective buyers and sellers must overcome significant transaction costs to achieve a patent sale. In addition to the transaction costs associated with incomplete and uncertain appropriability, patent transactions also involve non-trivial search, bargaining, and enforcement costs. A prospective buyer and seller must find each other, convey the idea, ascertain the characteristics of the patent that protects the idea, and negotiate the purchase of the patent rights. The dollar value of these costs vary widely across deals, firms, technologies and industries, but Teece (1977) estimated the total cost of an inter-firm technology transfer at between 2% and 59% of the value of the technology in an early study of 27 projects.

The key to understanding the effect of an increase in a patent's scope on the likelihood that the patent is sold lies in a return to the commitment problem the patent system is in part designed to solve. As part of the random process of discovery, a firm produces a variety of ideas. The firm may decide to commercialize some of these ideas but not others. Importantly, the lengthy time period required to secure patent protection and the difficulty in measuring the right to exclude mean that the decision to commercialize often must be made independently of realized USPTO outcomes. If the firm has decided to commercialize the idea, then presumably it will retain the patent rights regardless of the specific right to exclude afforded by the patent, although the firm would certainly prefer to receive a broader exclusionary right. If instead the firm has decided not to commercialize the idea, then the patent that protects the idea likely has limited value to the firm and may be a candidate for sale.

Of course, a firm can employ a patent offensively even in the absence of commercialization, but in practice most firms focus on competing in the product market rather than litigating patents on ideas they have chosen not to pursue. Firms that litigate must not only incur large up front costs but also risk retaliatory countersuits based on the defendant's own patents. Such costs and risks are likely unwarranted without the added benefit of product-market exclusivity that can only derive from actually practicing the patented invention. Some non-practicing firms exclusively pursue strategies focused on licensing and litigation, but these firms tend to purchase rather than sell patents.

If patent's original owner would like to sell it, then the likelihood of that sale occurring may depend in part upon how effectively the right to exclude afforded by the patent resolves the commitment problem. If the right to exclude is narrow, a prospective purchaser could commercialize the idea without actually purchasing the patent by tailoring the commercial product to avoid infringing the patent's narrow exclusionary right (i.e. "designing around" the patent). Alternately, the prospective purchaser of a narrow exclusionary right might worry that purchasing the patent would secure insufficient protection to prevent competitors from designing around the patent. In either case, gains to trade provided by the sale of a broader patent would be more likely to exceed the considerable transaction costs typically associated with patent sales. These observations lead to the first hypothesis.

Hypothesis 1. *Holding fixed the underlying idea, a patent with a broader exclusionary right is more likely to be sold.*

3.2 The effectiveness of the right to exclude

The practical effectiveness of patents for different purposes such as market exclusion or cross-licensing varies widely between industries (Cohen et al., 2000), and patent licensing rates reflect this variation. For instance, licensing rates are correlated with industry-reported patent effectiveness. (Anand & Khanna, 2000; Arora & Ceccagnoli, 2006). Industry-level variation in patent effectiveness may in part reflect fundamental differences in technology. Industries differ in terms of product development lifecycles, technological opportunities to "design around" a given invention, the number of distinct ideas that tend to be embodied in a given product or process, the ways in which inventions are translated from ideas to language, and many other characteristics. At the same time, variation in patent effectiveness may also be driven by industry structure, which itself may be endogenously related to technological characteristics. Regardless of the cause of this differential effectiveness, the result is likely the same. In industries where patents are reported as being more effective for exclusion, the precise delineation of the right to exclude is likely more important than in other industries for providing market safety and thereby facilitating sales. These observations lead to the second hypothesis.

Hypothesis 2. *Holding fixed the underlying idea, the effect of the scope of a patent's exclusionary right on the likelihood that the patent is sold is greater in industries where patents are reported as being more effective for exclusion.*

The likelihood that a good is sold depends not only upon the characteristics of the good itself, but also on the market in which the prospective sale would occur. According to the market design literature, an efficient market must exhibit thickness, a lack of congestion, and market safety (Roth, 2008), all of which have been shown to correlate with an increased likelihood of patent licensing (Agrawal et al., 2015). Indeed, as discussed in this section, patent markets seem particularly at risk of market design problems. Market safety is the degree of risk incurred by participating in a market. The risk incurred by participating in the market for ideas is that of expropriation, the prevention of which is the very purpose behind the right to exclude provided by the patent system. In line with Hypothesis 1, the right to exclude thus serves to provide market safety, ensuring that the buyer actually receives the right to exclude purportedly provided by the patent.

Under a classical definition, market thickness is directly proportional to both the number of buyers and the number of sellers in the market. (Roth, 2008). However, the market for patents functions quite differently than conventional markets. (See, e.g., Niioka, 2006; Richardson, Oliver, & Costa, 2016). The supply of patent rights is relatively high, since the random nature of technological discovery means that many firms have produced patents that are of limited value to the firm and that they would be happy to sell. On the other hand, the demand for patent rights is naturally low, in part because only a limited number of firms are equipped to make use of patents in a particular area of technology and in part because of the imperfect exclusion provided by patent protection. Finally, transactions are often conducted through brokers since search and negotiation costs are often relatively high and since prospective buyers prefer anonymity to avoid strategic reactions by other firms. For these reasons, market thickness for patents depends not on the number of sellers but rather on the number of active buyers who have demonstrated a willingness to purchase patent rights.

A thin market may be related in both cause and consequence to the effect of the right to exclude on the likelihood that a patent is sold. In a thin market, the dearth of active buyers makes asset pricing unstable, and patents are difficult assets to price in the best of markets. Thus, in a thin market, variation in the right to exclude may have little impact on the patent's value to a prospective purchaser. At the same time, a lack of market thickness may indicate that patents are ineffective at exclusion regardless of how the right to exclude is delineated. In either case, the effect is the same and leads to the third hypothesis.

> **Hypothesis 3.** *Holding fixed the underlying idea, the effect of the scope of a patent's exclusionary right on the likelihood that the patent is sold increases with market thickness.*

Market congestion arises from buyers' inability to evaluate (and hence to price) a surfeit of offers. When evaluating an offer requires a non-trivial amount of time, too many competing offers can destablize asset prices. (Roth, 2008). Under this classical definition, market congestion should not be problematic in the secondary market for patents since the market does not feature offers to sell and is instead driven primarily by the demand side. However, in practice the market for patent rights suffers from congestion of a different sort.

As the number of patents issued has exploded, scholars have become increasingly concerned about patent thickets, which are generally defined as a web of overlapping patent rights in an area of technology. Patent thickets are associated with a decreased frequency of in-licensing (Ziedonis, 2004) but an increase in in-licensing expenditure (Cockburn et al., 2010). Congestion also speeds litigation settlement (Galasso & Schankerman, 2008) and may deter product market entry (Hall et al., 2015). In a finding that implicates both congestion and thickness, a firm's in-licensing rate exhibits an inverted U-shaped relationship with the number of potential suppliers. (Fosfuri, 2006).

In a market with congested rights, an individual patent's legal right to exclude is likely less effective for product market exclusion. For instance, a firm seeking to exclude a competitor from the product market by way of a lawsuit would likely face a retaliatory countersuit that relies on the competitor's own patent with a similar exclusionary right. Further, a mass of overlapping patent rights by different firms may raise the likelihood that in litigation a court would invalidate a broad patent on the grounds that it was improvidently granted by the USPTO. When the right to exclude is ineffective for excluding competitors from the product market, it provides less market safety for overcoming the commitment problem inherent to the sale of ideas. These observations lead to the fourth hypothesis.

> **Hypothesis 4.** *Holding fixed the underlying idea, the effect of the scope of a patent's exclusionary right on the likelihood that the patent is sold decreases with patent congestion.*

Aside from technology and market characteristics, the role of the right to exclude in protecting an idea is also likely to vary across firms. Larger firms with more established patent portfolios may be able to rely partially on repeated interactions and a network of related patents to protect ideas. For instance, larger firms may simply cross-license all patents necessary to produce a product rather than sharply negotiating usage rights for individual patents. In contrast, the precise contours of a single patent's legal right to exclude may be quite important for a smaller firm with fewer ideas and a shorter history of interactions with other firms. These observations lead to the fifth and final hypothesis.

> **Hypothesis 5.** *Holding fixed the underlying idea, the effect of the scope of a patent's exclusionary right on the likelihood that the patent is sold decreases with firm size.*

4 Methodology

Section 4.1 describes the sample selection and variables. Section 4.2 introduces the identification strategy. Section 4.3 presents descriptive statistics on patent sales, and Section 4.4 describes the econometric specification.

4.1 Sample selection and data

Table 1 reports summary statistics for variables. The data are drawn from several sources. Data on patent assignments were constructed based on raw transfer data downloaded from the USPTO. The identification strategy, a plausibly exogenous shock to patent scope described in Section 4.2, was developed in related work based on a textual analysis of patent claims. (Kuhn et al. 2016). Patent bibliographic data was parsed from XML files provided by the USPTO as described by Thompson and Kuhn (2016). Measures of technological originality, technological familiarity, and patent rights congestion were constructed based on patent-to-patent similarity data described and validated by Younge and Kuhn (2015).

— Insert Table 1 about here. —

The unit of observation is the individual patent, and the sample includes 346,802 patents filed between January 1, 2004 and December 31, 2006. The sample includes all patents filed within the indicated time range with several exceptions. Most notably, the sample excludes biotechnology patents because the causal instrument is not valid in biotechnology. Additional details regarding sample selection are described in Appendix B.1.

I constructed assignment-related variables by analyzing 4,188,639 assignments from the USPTO patent assignment database. Each assignment agreement transfers ownership of a patent from the existing owner to a new owner. Although patent owners are not required by law to record assignments at the USPTO, they face strong incentives to do so because an earlier assignment can be legally voided by a subsequent assignment unless the earlier assignment is recorded. In the data, an assignment record identifies the patent number whose ownership is transferred, the names of the assignor (seller) and assignee (buyer), and the date the assignment was recorded. Instead, the parties to an assignment are identified via self-provided name and address information, not unique identifiers. Because this identification information is replete with typographical errors and alternative spellings, disambiguation was required in order to assign a unique firm identifier to each party to an assignment. The process for disambiguating the assignment parties is discussed in Appendix B.2.

Under U.S. law, every patent is initially owned in undivided and equal part by its inventors. In practice, inventors are typically required to assign the rights of their inventions to their employers under an employment agreement. I refer to the initial assignment transferring rights from the inventor or inventors to a firm as the original assignment. This article introduces a novel classification scheme in which each transfer after the original assignment is classified, based on the logic represented in Figure 1, as a sale, a spinoff, a firm exit, or a firm renaming by analyzing the full population of assignment data. Specifically, I first determine for each assignment whether the buyer ever previously received ownership of another patent. If not, I treat the focal assignment as the seller's entry into the patent market. I also determine for each assignment whether the buyer ever subsequently received ownership of another patent. If not, I treat the focal assignment as the buyer's exit from the patent market. These entry and exit variables allow me to classify each assignment. If an assignment is neither a seller exit nor a buyer entry, then I classify it as a sale. If an assignment is a seller exit but not a buyer entry, I classify it as a firm exit by the seller. If an assignment is a buyer entry but not a seller exit, I classify it as a spinoff that creates the buyer. If an assignment is both an exit by the seller and an entry by the buyer, I classify it as a firm renaming and exclude the assignment from subsequent analysis. Appendix B.2 provides more detail regarding the construction and validation of these classification categories.

— Insert Figure 1 about here. —

The first variable in Table 1 indicates whether a patent was reassigned to another firm after it was issued to an initial firm. The second through fourth variables indicate whether the patent was reassigned post-grant in particular types of assignments, including sales, exits, and spinoffs. The fifth through seventh variables are similar but only count those transfers that include 20 or fewer patents between the buyer and seller in a given year. The cutoff of 20 is chosen arbitrarily, but conservatively, to focus attention on potentially strategic patent sales rather than the wholesale transfer of entire portfolios.

The data on assignments are analyzed subject to an important caveat. Multinational corporations sometimes engage in intra-firm patent transfers such as between domestic headquarters and foreign subsidiaries, often for organizational or tax reasons. (Arora, Belenzon & Rios, 2011). The current version of the data set does not identify such transfers. However, many such transfers occur before the patent is granted and as such do not affect the analysis of this paper. Further, transfers between firm subsidiaries are not arms-length market transactions and therefore should bias the estimates in the results toward zero.

The thirteenth and fourteenth variables in Table 1 were constructed based on patent-to-patent textual similarity data. Familiarity measures the extent to which the focal patent is similar to past work by the same firm. Originality measures the extent to which the focal patent is different than past work by other firms. Appendix B.2 provides further details regarding the construction of the familiarity and originality measures.

The fifteenth and sixteenth variables in Table 1 measure market thickness and rights congestion—characteristics of the local technology market in which a patent is situated. I define the local technology market based on the USPTO examining division (i.e. art unit) that examined the patent. Patent applications at the USPTO are assigned for examination to one of 531 art units based on the technology described and claimed in the application. The art units, which are split among eight Technology Centers, each includes a number of patent examiners skilled in examining a particular type of technology, such as mechanical presses (art unit 3725) or furnaces (art unit 3743).

Conventionally, market thickness is defined based on the number of buyers and sellers making bids and offers. However, since the market for patents is driven primarily by demand rather than supply, as discussed in Section 3, I instead defined market thickness as the ratio of buyers to patents within the local technology market (i.e. art unit) of the focal patent. The raw value was logged and normalized to produce a z-score that provides a sense of the number of buyers to which a potential seller of a focal patent would have access. The interpretation of the thickness variable is that a patent having a thickness level of 1 is in an art unit that is 1 standard deviation thicker than the mean for all patents.

Conventionally, market congestion is a feature of markets in which offers must be evaluated in serial rather than in parallel and in which evaluating offers requires non-trivial time. However, because the market for patents features a surplus of rights rather than a surplus of offers, I instead measure the related concept of patent thickets. Scholars have proposed several measures of patent thickets, but all are based on patent citations or patent office technological classification. (Cockburn et al., 2010; Von Graevenitz, Wagner, & Harhoff, 2011; Ziedonis, 2004). In order to avoid many of the concerns about bias and selection inherent in patent citations and manual technological classification, I instead constructed a measure based on patent-to-patent textual similarity data. Specifically, the patent congestion variable is constructed by comparing the mean originality of all patents in the local technology market (i.e. art unit) to the originality of other patents in the sample. The interpretation of the congestion variable is that a patent having a congestion level of 1 is in an art unit that is 1 standard deviation more congested than the mean for all patents.

Appendix B.2 provides further details regarding the construction of the congestion measure.

Table 2 reports summary statistics by broad industry category. As discussed earlier in this section, each patent is examined by an art unit at the patent office. The art units are divided into eight broad technology centers. In addition, I list as a separate category those patents issued by art units responsible for medical devices. As is shown in Table 2, the frequency of sales varies considerably by industry, with sales being most common in computer, networking, and electrical technologies and least common in chemical and mechanical technologies. The medical device, chemical, and computer industries are all thick and uncongested, suggesting that the right to exclude may be particularly effective in these broad categories. For the other industry categories, the mean congestion and thickness values are more ambiguous. For instance, mechanical category is moderately thick but also quite congested, while the electrical components industry is relatively uncongested but also quite thin.

— Insert Table 2 about here. —

4.2 Identification strategy

This paper employs as an identification strategy an instrumental variable that provides a plausibly exogenous shock to the final scope of the patent relative to the underlying idea. Patent applications are quasi-randomly assigned to patent examiners within an art unit, and patent examiners vary in the stringency with which they examine patent applications. (Cockburn, Kortum, and Stern (2002); Kuhn et al. (2016); Lemley and Sampat (2012); Sampat and Williams (2015)). The assignment of patent applications to patent examiners thus provides a natural experiment—patents examined by less stringent examiners provide on average a broader exclusive right relative to the underlying idea than those examined by more stringent examiners. Kuhn et al. (2016) show that the number of words in a patent's first independent claim is a good proxy for the patent's scope, and that the mean number of words added to the first independent claim in an examiner's patents (excluding the focal patent) is a good measure of examiner stringency.

Table 3 presents the first stage for all patents and for industry sub-samples. Were all examiners identical, a patent's scope would be determined entirely by factors such as the underlying technology and would not depend on examiner characteristics. However, Table 3 reveals considerable variation between examiners. A one standard deviation increase in examiner score predicts a .18 ($p<.01$) standard deviation increase in patent scope for all patents. Although the strength of the instrument varies somewhat by industry, examiner score remains a positive and highly significant predictor of patent scope.

— Insert Table 3 about here. —

A natural concern with any instrument is that although institutional factors suggest quasi-random assignment, the practical realities are such that assignment is not in fact random. Balance statistics indicate this is not the case. Kuhn et al. (2016) show that initial applications examined by lenient and stringent examiners are nearly identical on characteristics observable at the time of filing, and that patent examiner stringency varies little over time and by experience. Table 4 presents balance statistics comparing the applications examined by the upper and lower quartiles of examiners by stringency for the specific sample analyzed in this article. The statistics show a slight selection effect due to fact that more stringent examiners are more likely to cause applicants to abandon low-quality applications. All of these differences are statistically significant due to the large sample size, but the differences are not economically significant. The means differ by at most .033 standard deviations, well below the threshold of .25 that would indicate a substantial difference. (See, e.g., Imbens & Wooldridge 2009). Further, because higher-quality applications are more likely to be transferred and because the small selection effect removes more low-quality applications from the set of applications examined by stringent examiners, the selection effect should slightly bias the results toward zero. Indeed, in unreported results, controlling for the variables included in Table 4 slightly increase the point estimates relative to the uncontrolled results.

— Insert Table 4 about here. —

Another concern with any instrument is violation of the exclusion restriction. Although patent examiners appear to be quasi-randomly assigned, a more stringent patent examiner may affect patents in ways other than the scope of the exclusive right. The highly-regulated, arms-length, and quality-controlled nature of the procedural rules governing patent examination should allay many of these concerns. However, tougher examiners may cause patents to issue more slowly, which could deter subsequent transfers. (Gans et al., 2008). In practice, the difference is relatively small—patents examined by a patent examiner one standard deviation tougher than the mean are delayed on average by about 2 months (Kuhn et al., 2016)—but this concern cannot be ruled out completely. Another possibility, which again cannot be ruled out, is that examiner stringency may cause patent owners to perform increased or decreased follow-on innovation, which may affect whether or not the patent is sold. In part for this reason, the results include reduced form versions of all regression tables.

4.3 Descriptive statistics

Conventional wisdom suggests that patents are sold almost exclusively in large portfolios, which may render irrelevant the characteristics of any one patent. However, the data indicates that this is not the case. Figure 2 shows a patent-level histogram of first post-grant transfers, binned by the number of other patents included in the transfer. Although many transfers include 50 or more patents, the majority of patent transfers include 25 or fewer patents. Figure 3 plots the average examiner score of patents transferred post-grant as a function of the number of other patents included in the transfer. As would be expected, scope appears to plays a larger role in the transfer of small numbers of patents than in transactions for large portfolios. Figure 4 shows the frequency of different types of post-grant transfers by patent. Clearly, most patents remain with the firm that produces them. However, about 5.7% of patents in the sample are transferred post-grant. Many of these transfers occur when large portfolios change hands, but about 1.5% of patents are transferred in small batches of 20 or fewer patents. To identify the role of patent scope as a driver of patent transfers, the results in the rest of the paper focus on transfers of patents that occur in groups of 20 or fewer. Although some patents are transferred more than once, such patents are quite rare, and subsequent analysis in this paper focuses on the first post-grant transfer for analytic simplicity.

— Insert Figure 2, Figure 3, and Figure 4 about here. —

Most patent transfers are likely motivated by idiosyncratic characteristics of the underlying technology that are difficult to measure. Nevertheless, the data reveals patterns that may shed light on the motivations behind some transfers. Table 5 presents the results of four models estimating the relationship between patent sales and patent citations made to the focal patent. Consistent with the findings of Serrano (2010), columns 1 and 2 show that patents that are sold receive 31% ($p<.01$) more total forward citations than patents that are not sold, but they receive 14% ($p<.01$) fewer citations by the firm that originally owned the patent. Thus, patents that are sold have on average a greater impact than those that are not, but unsurprisingly are less important to their original owners. Columns 3 and 4 show that patents that are sold are much more likely to be cited by patent examiners to reject future patent applications. Indeed, a patent that is sold is 2.3 ($p<.01$) times more likely to be cited in a novelty rejection than a patent that is not sold. Thus, patents that are sold are on average not only more impactful in a general sense, but also more strategically important in blocking future patent rights by other firms.

— Insert Table 5 about here. —

Another way to understand which patents are more likely to be transferred is to compare a focal patent's similarity to both prior work by the same firm and prior work by other firms. Table 6 presents the results of four models using a logistic regression to estimate the relationship between a patent's technological positioning using patent-to-patent textual similarity data and the odds ratio that a patent is transferred. The models split the sample by firm size into large firms and small firms.

— Insert Table 6 about here. —

A large firm seeking to establish and maintain a competitive advantage should retain the patents that contribute to such an advantage and sell the patents that do not. One way to establish a competitive advantage is to develop a technological competence that distinguishes the firm from its competitors. Accordingly, a large firm should be less likely to sell patents that are more central to the firm's area of competence and that are more economically original, since those are precisely the exclusionary rights likely to provide the greatest advantage. Columns 1 and 2 in Table 6 support such a story, suggesting that large firms are less likely to sell patents on technology familiar to the firm and less likely to sell patents on technology more original to the world outside the firm. Specifically, a large-firm patent one standard deviation more familiar to the firm is .84 ($p<.01$) times as likely to be sold, while a large-firm patent one standard deviation more original to the world is .89 ($p<.01$) times as likely to be sold.

Patent sales by small firms are likely to be driven by different considerations. As noted in columns 1 and 2, patents on technology original to the world appear more valuable to large firms than patents on less original technology and thus in greater demand. For a small firm that seeks to monetize its patent portfolio but has relatively few patents to sell, the likelihood of a sale may be likely to be driven by demand-side considerations of precisely which ideas are valuable on the market. However, the extent to which a technology is unfamiliar to the firm may be less likely to influence whether a patent is sold for a small firm since small firms may be more willing to expand into unfamiliar technologies. Columns 3 and 4 in Table 6 support such a story, with small firms being much more likely to sell patents that are more original and with technological familiarity playing a minor role. Specifically, a small-firm patent one standard deviation more familiar to the firm is .95 ($p<.1$) times as likely to be sold, while a small-firm patent one standard deviation more original to the world is 1.136 ($p<.01$) times as likely to be sold.

4.4 Econometric specification

Equation 1 provides the primary econometric specification. The likelihood that a patent is sold at any time after it is issued is estimated using a two-stage least squares regression. Each observation i represents a single patent, while firms are indexed with j. In the first stage, a patent's scope is predicted based on the score of the examiner responsible for examining that patent. In the second stage, the estimated scope is used to predict the likelihood that the patent is sold. Dummy variables are included for each art unit during each patent application filing year to capture technology-level and year-level variation in the rate of sales.

$$\begin{aligned} Scope_i &= \eta * ExaminerScore_i + \rho * ArtUnit_i * AppYear_i + \mu_i \\ Sale_i &= \beta * \widehat{Scope}_i + \gamma * \widehat{Scope}_i * C_i + \delta * C_i + \tau * ArtUnit_i * AppYear_i + \epsilon_{ij} \end{aligned} \tag{1}$$

The main specification does not include any additional control variables, but other specifications include a control variable C_i that controls for firm size, market thickness, rights congestion, or patent effectiveness. In some secondary specifications, the dependent variable is replaced with a dummy indicating not whether the patent was sold, but rather whether the patent was transferred in a firm exit or spinoff transaction.

Patent scope (i.e. $Scope_i$) measures the breadth of the focal patent's right to exclude relative to other patents in the same area of technology. Similarly, $ExaminerScore_i$ measures the extent to which the focal patent's examiner tends to issue broad patent rights relative to the examiner's peers. Following Kuhn et al. (2016), $Scope_i$ and $ExaminerScore_i$ are defined on the level of the patent as described in Equations 2 and 3. The variable $Words$ measures the number of words in the first claim of the focal patent ($Words_i$), other patents in the art unit of the focal patent ($Words_a$), or other patents examined by the examiner of the focal patent ($Words_{ei}$). Importantly, Equation 3 is calculated for each patent with the focal patent excluded from the calculation to avoid introducing endogeneity into the instrument.

$$Scope_i = \frac{Words_i - mean\left(Words_a\right)}{StdDev\left(Words_a\right)} \tag{2}$$

$$ExaminerScore_i = \frac{mean\,(WordsAdded_a) - mean\,(WordsAdded_{ei})}{StdDev\,(WordsAdded_a)} \tag{3}$$

Autocorrelation of the error term does not seem to be problematic since in unreported results the Huber-White robust standard errors are quite close to classical standard errors across all models. However, the error term in the regression could be correlated within any of various groupings. For example, firms could vary in their sensitivity to changes in patent scope. To address this concern, standard errors in ordinary least squares and two-stage least squares regressions were clustered at the level of the firm that originally owned the focal patent. Errors could be also correlated within patents examined by the same examiner or within patents examined by the same art unit. In unreported results, clustering at the level of art unit or patent examiner yields standard errors similar to clustering at the level of the firm.

As discussed in Section 4.2, violation of the exclusion restriction cannot completely be ruled out. For this reason, and because the dependent variable is binary with a relatively low mean (between 1% and 2%), I also estimated each model using a conditional logistic regression model in which the likelihood of a sale is regressed directly on examiner score. Equation 4 presents the specification for the conditional logistic regression. Estimating a logistic regression with fixed effects (i.e. a conditional logistic regression) can render estimates less consistent due to the incidental parameters problem. However, the number of fixed effects included Equation 3 is very small in comparison with the number of observations, and the number of observations within each group is quite large. Together these factors should minimize any bias introduced by including the incidental parameters in the model.

$$\log\left(\frac{Sale_i}{1 - Sale_i}\right) = \beta * ExaminerScore_i + \gamma * ExaminerScore_i * C_i + \tau * TechCenter_i * AppYear_i + \mu_i \tag{4}$$

5 Results

5.1 Main results

Table 7 presents estimates for naive models in which the likelihood that a patent is sold is regressed directly on patent scope. Model 1 presents the results for all patents, while models 2 through 7 present the results for industry subsamples. Broader patents are more likely to be sold, with a one standard deviation increase in patent scope associated with a .085 percentage point increase ($p<.05$) in the likelihood of a sale across the entire sample. The effect is positive but insignificant for all industries and is larger within the computer, chemical, and mechanical industries. However, patent scope is an endogenous result of the underlying invention, so the effects shown in Table 7 may be driven by characteristics of the underlying idea or the patent applicant rather than the scope of the exclusionary right.

— Insert Table 7 about here. —

To address endogeneity concerns, Table 8 presents estimates for reduced form ordinary least squares fixed effects models in which the likelihood that a patent is sold is regressed on the examiner score instrument. Examiner score has a larger effect than patent scope in Table 7, with a one standard deviation increase in examiner score increasing the likelihood that the patent is sold by .17 percentage points ($p<.01$). The largest effect is observed in the medical device industry, where a one standard deviation increase in examiner score increases the likelihood of a sale by 1.4 percentage points ($p<.01$). Moderate and statistically significant increases are also observed in the computer, chemical, and mechanical industries, while patent scope seems to exert little to no effect on sales in the electrical and semiconductor industries.

— Insert Table 8 about here. —

Table 9 presents the main results, a set of two-stage least squares regressions of the likelihood that a patent is sold on patent scope, with patent scope instrumented by the examiner score. The main effect is shown in Model 1, where a one standard deviation increase in patent scope increases the likelihood that a patent is sold by .95 percentage point ($p<.01$), a result which directly supports the main hypothesis.

The largest effect is observed in the medical device industry, where a one standard deviation increase in patent scope increases the likelihood of a sale by 9.1 percentage points ($p<.01$). Large and statistically significant effects are also observed in the computer, chemical, and mechanical industries. As noted in Table 2, the patent markets in chemical, computer, and medical device technologies are all thick and uncongested. While the mechanical sector is highly congested, it is also rather thick, which suggests a mature market in which buyers know how to value patents.

In contrast, patent scope has a negligible effect in the two thinnest industries: electrical component and semiconductor industries. Important semiconductor products are typically covered by many hundreds, and sometimes thousands, of different patents. Because of this fragmentation, the importance of the scope of any particular patent is likely to be negligible. The electrical component industry, including technologies such as television and telephony, is well-established and highly consolidated. In a market where the large majority of patents are owned by very large companies and are never transferred, the scope of a patent is unlikely to have much effect on the likelihood of sale.

The estimates presented in Table 8 are much larger and more significant than those presented in the endogenous regressions in Table 7, likely due to error in measuring patent scope. In the endogenous regression, the likelihood of a sale is regressed on the normalized number of words in its first independent claim, which is a statistically significant but very noisy measure of patent scope (Kuhn et al., 2016). The examiner score provides a much less noisy measure by aggregating the patent examiner's impact on patent scope across all other patents issued by the examiner.

When interpreting the estimates in Table 8 and other instrumental variable regression results in this article, the reader should keep in mind the magnitude of the effects in the first stage. Because a one standard deviation increase in examiner score increases a patent's scope by about .2 standard deviations, a variation of one standard deviation in patent scope as indicated in the estimates in Table 8 is larger than most between-examiner variation observed in practice. For this reason, reduced form estimates regressing directly on examiner score are somewhat more interpretable than the instrumental variable results, so Table 11, Table 12, Table 13, and Table 14 present estimates of reduced form logistic regression models.

— Insert Table 9 about here. —

Table 10 repeats the main results controlling for firm size. Because patents are randomly assigned to patent examiners and because firm size is demeaned, the main effect of patent scope on the likelihood that a patent is sold is not appreciably altered by including firm size as a covariate. The results show that in the unrestricted sample, patents by larger firms are less likely to be sold—a 100% increase in firm size decreases the likelihood that a patent is sold by .27 percentage points ($p<.01$). Further, the effect of an increase in patent scope on the likelihood of a sale decreases with firm size–a 100% increase in firm size decreases the effect of patent scope by .15 percentage points ($p<.05$). However, on the industry level the negative interaction is significant only for the computer and mechanical subsamples. Indeed, for chemical patents the effect is reversed—scope has a greater effect on patent sales for larger companies. Together these results provide qualified support for Hypothesis 5, suggesting that patent scope is in general more important for facilitating sales by smaller firms but that the relationship between patent strategy and firm size is industry-specific.

— Insert Table 10 about here. —

Table 11 presents estimates of conditional logistic regressions corresponding to those in Table 8, with the logit of post-grant sale regressed directly on examiner score. The coefficients are presented as log odds, but to aid in interpretability Figure 5 plots the probability of a sale for different industries as curves predicted based on these estimates. The results are broadly consistent with those shown in Tables 8 and 9. In the unrestricted sample, a one standard deviation increase in examiner score increases the odds that the patent is sold by a factor of 1.22 ($p<.01$). The effect is larger and statistically significant in the computer, chemical, and mechanical industries and largest in the medical device industry, where a one standard deviation increase in examiner score increases the odds that the patent is sold by a factor of 2.28 ($p<.01$). As is shown in Figure 5(a), industries that are largely similar in the overall likelihood that a patent is sold differ dramatically in the effect of examiner score on the likelihood of a sale. Figure 5(b), which compares the effect of examiner score in the medical device industry to the effect for all patents, highlights the outsized role that the right to exclude plays in certain industries.

— Insert Table 11 and Figure 5 about here. —

In order to address concerns related to possible violation of the exclusion restriction for the instrumental variable and concerns about model fit with a low-probability binary dependent variable, the remainder of the results tables estimate reduced form conditional logistic regression models rather than two-stage least squares regression models. However, two-stage least squares regression tables corresponding with the tables in Section 5.2 are included in Appendix A for robustness.

5.2 Industry and market characteristics

Although the effect of patent scope exhibits considerable heterogeneity across industry divisions, this cross-classifications does not precisely reflect reports of where and how market participants believe that patents are effective. Table 12 and Table 13 present the results of eight models examining how the effect of patent scope on the likelihood of a sale interacts with the effectiveness of patents in different industries. As discussed in Appendix B.1, survey measures of patent effectiveness provided by Cohen et al. (2000) were converted to z-scores by matching industries to art units and normalizing the measures across the sample of patents. In each table, models 1 and 2 present the results for product patent effectiveness, while models 3 and 4 present the results for process patent effectiveness. The coefficients are presented as log odds, but to aid in interpretation Figure 6(a) and Figure 6(b) graphically present the estimates of model 1 in Table 12 and Table 13.

The models presented in Table 12 include interaction terms for overall patent effectiveness (models 1 and 3) and patent effectiveness for market exclusion (models 2 and 4).[2] Table 12 shows that the effect of examiner score on the likelihood of a sale increases as the reported effectiveness of patents increases overall (models 1 and 3) and specifically for exclusion (models 2 and 4). Indeed, these interaction effects are in most instances (models 1, 2, and 4) larger than the main effects, which are even negative for exclusion (models 2 and 4). As shown in Figure 6a, in industries where product patents are reported as being ineffective for market exclusion, patents that vary widely in terms of their scope differ little in terms of the probability that they will be sold. In contrast, in industries where patents are reported as being effective for market exclusion, the probability that a patent with an examiner score one standard deviation greater than the mean is sold is over twice that for a patent with an examiner score one standard deviation narrower than the mean. Thus, patent effectiveness in general and for exclusion in particular as reported by industry participants may in part reflect the extent to which a patent's scope provides market safety for the purpose of facilitating sales.

— Insert Table 12 and Figure 6 about here. —

One concern with these results might be that any increase in patent importance or effectiveness, even for non-exclusionary purposes, may be associated with an increased effect of patent scope on the likelihood of a sale. Table 13 reveals that this is not the case. Using patents for cross-licensing or negotiating with rivals reflects a portfolio-based strategy in which a single patent is less likely to be strategically important. In such a context, the scope of a patent is unlikely to have a particularly strong effect on the likelihood of a sale. The models presented in Table 13 include interaction terms for patent effectiveness for cross-licensing (models 1 and 3), and patent effectiveness for negotiating with rivals (models 2 and 4).[3] The main effects for each of these models are positive—an increase in patent effectiveness for negotiation or cross-licensing is associated with an increased likelihood that a given patent is sold. However, the effect of patent scope on the likelihood of a sale does

[2] The categories "Overall" and "Exclusion" correspond to the labels "Patents" and "Fences" in Cohen et al. (2000).

[3] The categories "Cross-licensing" and "Negotiation" correspond to the labels "Cross-licensing" and "Player" in Cohen et al. (2000).

not increase with the reported effectiveness of product patents for cross-licensing (model 1) or negotiation (model 3). In fact, the effect decreases as the reported effectiveness of process patents for cross-licensing (model 2) or negotiation (model 4) increases.

— Insert Table 13 about here. —

Overall, the results presented in Table 12 and Table 13 are consistent with Hypothesis 2. The likelihood that a patent is sold increases not only with the patent's scope, but also with the interaction of patent scope with the industry-reported effectiveness of the legal right to exclude. In contrast, the likelihood that a patent is sold does not increase with the interaction between the scope of the right to exclude and the effectiveness of patents for negotiation and cross-licensing.

The realized effect of patent scope on the likelihood of a sale thus reflects industry reports about where patents are effective for exclusion as opposed to more aggregate purposes such as cross-licensing or negotiation. While this differential effectiveness may in part reflect fundamental aspects of the underlying technology, it may also relate to characteristics of the market for technology. Table 14 presents the results of nine models examining how market characteristics moderate the relationship between the examiner score and different types of transfers. The dependent variable in the first three models is the logit of whether the patent is sold, and all coefficients are statistically significant for these models with $p<.01$. In model 2, the examiner score remains a positive and statistically significant predictor of a sale even when controlling for market thickness and the interaction of scope and thickness. Moreover, the effect of examiner score increases substantially in thicker markets. Figure 7(a) presents this result graphically by plotting the probability of a sale as a function of market thickness for patents one standard deviation above and below the mean examiner score. As is shown in Figure 7(a), broad and narrow patents have nearly the same probability of being sold in very thin markets, but broader patents are over three times more likely to be sold in very thick markets. Model 3 presents a similar result for patent congestion. Patents are much less likely to be sold in more congested markets, and the effect of patent scope on the likelihood of a sale is substantially reduced in congested markets. Figure 7(b) presents this result graphically. Again, broad and narrow patents have nearly the same probability of being sold in very congested markets, but broader patents are much more likely to be sold in very uncongested markets. The results presented in Table 14 are thus broadly consistent with Hypotheses 3 and 4.

— Insert Table 14 and Figure 7 about here. —

The models presented in Table 14 include fixed effects for broad industry categories. Nevertheless, one concern with these results is that market characteristics are highly endogenous, so an unobserved variable correlated with market characteristics rather than the market characteristics themselves may be responsible for the change in the effect of patent scope. For instance, industry structure may both give rise to market characteristics and influence the effect of patent scope. This concern cannot be completely ruled out, but one way of addressing it is to restrict the sample to a specific industry that nevertheless contains subindustries that vary in market characteristics. The medical device industry works for this purpose since it is composed of a number of different technologies that vary in their patent market characteristics, with large firms buying patents from different areas of medical device technology. Figure 7(c) and Figure 7(d) plot the effect of plus or minus one standard deviation in examiner score on the likelihood of a sale as a function of market thickness and market congestion for the subsample patents in the medical device industry, with the x-axis limited to where the data provides full support. The interactions between examiner score and market characteristics for the medical device subsample are consistent with the results for the unrestricted sample, with an increase in examiner score having virtually no effect in thin or congested markets and a very large effect in thick or uncongested markets.

Patents are transferred not only in sales between firms, but also via firm exits (e.g., bankruptcy, acquisition, merger) and firm spinoffs. A patent transfer via a firm exit may be in many cases a form of market transaction—this is clear in the case of bankruptcy, but even in mergers and acquisitions a large part of a firm's value may depend on its patent portfolio. Thus, patent transfers via firm exits should exhibit many of the same characteristics as patent sales. Models 4, 5, and 6 in Table 12, in which the dependent variable is the logit of whether the patent is transferred in a firm exit from the patent market, are consistent with this story. Across each of these models, all coefficients are statistically significant with $p<.01$. Examiner score has a positive effect on the logit of post-grant exit transfers. Examiner score interacts positively with market thickness and negatively with market congestion. Figure 8(a) and Figure 8(b) present the results of models 5 and 6 graphically. As shown in Figure 8(a), in a very thick market the probability of a firm exit transfer is over three times greater for a patent having a score one standard deviation greater than the mean than for a patent having an examiner score one standard deviation narrower than the mean, while the difference between high and low examiner scores is negligible in a thin market. Similarly, Figure 8(b) shows that patents in congested markets are almost never transferred via firm exits, and that the examiner score has little effect on the likelihood of such a transfer. However, the examiner score exerts a large and positive effect on firm exit transfers in very uncongested markets.

— Insert Figure 8 about here. —

In contrast to firm exits, spinoff transfers are not conducted in a marketplace. Because a spun off firm typically has a close relationship to the parent firm and because the spinoff is unlikely to involve a sharp negotiation over patent rights, the likelihood that a patent is transferred via a spinoff should not depend very much on the scope of the patent. Models 7, 8, and 9 in Table 12, in which the dependent variable is the logit of whether the patent is transferred in a firm spinoff, are consistent with this story. An increase in examiner score slightly but significantly increases the likelihood of a spinoff transfer—a one standard deviation increase in examiner score increases the odds of such a transfer by a factor of 1.13 ($p<.01$). However, the interactions between examiner score and market characteristics are small in magnitude and insignificant. Figure 8(c) and Figure 8(d) present the results of models 8 and 9 graphically.

6 Conclusion

The results demonstrate that an exogenous increase in the scope of a patent's legal right to exclude, holding fixed the underlying idea, increases the likelihood that the patent is sold. A patent having a broader exclusionary right is likely more valuable, so the sale of that patent is more likely to overcome the considerable transaction costs involved in the transfer of intellectual property. The effect is stronger for smaller firms and in general holds not only for patent sales but also for the transfer of patents upon firm exit. Broadly speaking, this finding highlights the importance of the legal right to exclude as a strategic factor for firms seeking to profit from innovation by transferring intellectual property rights. (Gans et al., 2008; Teece, 1986).

The results also show that the effect of patent scope on the likelihood that the patent is sold varies considerably across industries. The effect is large in the medical device, chemical, mechanical, and computer industries, but is small and not statistically significant in the electronic component and semiconductor industries. The results suggest that these differences may reflect in part differences in the effectiveness of the legal right to exclude for providing market exclusivity, which at the industry level is distinct from both the overall rate of patent sales and the effectiveness of patents for purposes other than exclusion. (Cohen et al., 2000) For example, patents are no more likely to be sold in industries where they are reported as being more effective for exclusion, but the effect of patent scope on the likelihood of a sale increases with the effectiveness of patents for exclusion. However, although patents are more likely to be sold in industries where they are more effective for cross-licensing, the effect of patent scope on the likelihood of a sale is unrelated to the effectiveness of patents for cross-licensing.

At a more detailed level, the results highlight how the relationship between patent scope and patent sales depends substantially on characteristics of the market in which the patent is situated, even within the broad industry categories noted above. Viewed through the lens of the market design literature (see, e.g., Roth, 2008) and consistent with Arrow (1962), the legal right to exclude provides market safety by ensuring that the purchaser of an idea actually receives ownership of the idea being purchased. As in all markets, the results suggest that patent sales are inhibited by congestion, a lack of market safety, and a lack of market thickness. (Roth (2008)). In a finding specific to patent markets, the results show that patent scope has less of an effect on the likelihood of a sale in a thin or congested market, suggesting that patent scope is less effective at providing market safety in these contexts. In particular, congested patent rights seem to render individual patents less effective, consistent with the findings of Ziedonis (2004) and Hall et al. (2015). Together these demonstrate that just as the market for ideas through patent licensing is plagued with imperfections (see, e.g., Agrawal et al., 2015; Arora et al., 2004; Gans & Stern, 2010), so too is the free flow of rights on the secondary market for patents.

Although only about 5.7% of patents are transferred post-grant, the data suggests that many of these transfers are strategic and important. Patents that are transferred are more highly cited and are more likely to block other patents via novelty or non-obviousness rejections. Patent transfers also seem to reflect choices by large firms about firm boundaries, since such firms are more likely to transfer patents on technology that is unfamiliar to them and less likely to transfer patents on technology that is more original. The secondary market for patents may thus be important for a range of strategic considerations such as establishing a technological competitive advantage (Teece, 1982), reducing IP ownership fragmentation (Ziedonis, 2004), and allowing a firm to profit from innovation that it is poorly equipped to commercialize (Teece, 1986). These strategic considerations may explain in part the economic growth and gains to trade that recent studies have found to be associated with patent sales. (Akcigit et al., 2016; Serrano, 2011). In addition, strategic transfers may spur further development of the underlying technology rather than allowing it to languish unexploited, for instance when technologies radical to the firm are orphaned when an intrapreneur leaves the organization. (See, e,g. Pinchot, 1987).

The results also highlight the complex tradeoffs inherent in deciding issues related to patent grant rates and patent scope determinations. (Merges & Nelson, 1990). The promise of a patent is intended to provide incentives for innovation, and ownership of a patent can also provide incentives for further development and commercialization. Indeed, awarding broad rights for a particular patent may facilitate the patent's transfer to a firm better able to develop the underlying ideas. However, consistent with the findings of Ziedonis (2004), a surfeit of overlapping and uncertain rights can create market congestion that inhibits the free flow of rights between firms.

I argue that this article supports three main conclusions. First, holding fixed the underlying idea, a broader patent is more likely to be sold. Second, patent sales often represent strategic decisions for firms that stem in part from the strategic importance of the underlying ideas, in part from the legal characteristics of the intellectual property right, and in part from the market characteristics of the technology in which the patent is situated. Third, the effect of the scope of the legal right to exclude on firm decisions differs substantially across firms, markets, and industries, highlighting both variation in the effectiveness of patent rights and the importance of within-industry analysis for better understanding the strategic implications of the patent system for firms. Together these conclusions suggest that both patent scope and patent sales are important components of the market for ideas and may play an important role in determining why some firms succeed while others fail.

References

Agarwal, R., Ganco, M., & Ziedonis, R. H. (2009). Reputations for toughness in patent enforcement: Implications for knowledge spillovers via inventor mobility. *Strategic management journal, 30*(13), 1349-1374.

Agrawal, A., Cockburn, I., & Zhang, L. (2015). Deals not done: Sources of failure in the market for ideas. *Strategic management journal, 36*(7), 976-986.

Akcigit, U., Celik, M. A., & Greenwood, J. (2016). Buy, Keep, or Sell: Economic Growth and the Market for Ideas. *Econometrica, 84*(3), 943-984.

Anand, B. N., & Khanna, T. (2000). The structure of licensing contracts. *The Journal of Industrial Economics, 48*(1), 103-135.

Anton, J. J., & Yao, D. A. (2002). The sale of ideas: Strategic disclosure, property rights, and contracting. *The Review of Economic Studies, 69*(3), 513-531.

Arora, A., Belenzon, S., & Rios, L. A. (2011). *The organization of R&D in American corporations: the determinants and consequences of decentralization*. National Bureau of Economic Research.

Arora, A., & Ceccagnoli, M. (2006). Patent protection, complementary assets, and firms' incentives for technology licensing. *Management Science, 52*(2), 293-308.

Arora, A., Fosfuri, A., & Gambardella, A. (2004). *Markets for technology: The economics of innovation and corporate strategy*: MIT press.

Arora, A., & Gambardella, A. (2010). Ideas for rent: an overview of markets for technology. *Industrial and Corporate Change, 19*(3), 775-803.

Arrow, K. (1962). Economic welfare and the allocation of resources for invention *The rate and direction of inventive activity: Economic and social factors* (pp. 609-626): Princeton University Press.

Aspex Eyewear, Inc. v. Miracle Optics, Inc, No. No. 04-1265, 434 1336 (Court of Appeals, Federal Circuit 2006).

Bessen, J. E., Meurer, M. J., & Ford, J. L. (2011). The private and social costs of patent trolls. *Boston Univ. School of Law, Law and Economics Research Paper*(11-45).

Bloom, N., Schankerman, M., & Van Reenen, J. (2013). Identifying technology spillovers and product market rivalry. *Econometrica, 81*(4), 1347-1393.

Burk, D. L., & Lemley, M. A. (2003). Policy levers in patent law. *Virginia Law Review*, 1575-1696.

Coase, R. H. (1937). The nature of the firm. *Economica, 4*(16), 386-405.

Cockburn, I. M., Kortum, S., & Stern, S. (2002). *Are all patent examiners equal?: The impact of characteristics on patent statistics and litigation outcomes*: National Bureau of Economic Research.

Cockburn, I. M., MacGarvie, M. J., & Mueller, E. (2010). Patent thickets, licensing and innovative performance. *Industrial and Corporate Change*, 19(3), 899-925.

Cohen, W. M., Nelson, R. R., & Walsh, J. P. (2000). *Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not)*. National Bureau of Economic Research.

Dahlman, C. J. (1979). The problem of externality. *The Journal of Law & Economics, 22*(1), 141-162.

Fischer, T., & Henkel, J. (2012). Patent trolls on markets for technology–An empirical analysis of NPEs' patent acquisitions. *Research Policy, 41*(9), 1519-1533.

Fosfuri, A. (2006). The licensing dilemma: understanding the determinants of the rate of technology licensing. *Strategic management journal, 27*(12), 1141-1158.

Galasso, A., & Schankerman, M. (2008). *Patent thickets and the market for ideas: evidence from settlement of patent disputes*.

Galasso, A., Schankerman, M., & Serrano, C. J. (2013). Trading and enforcing patent rights. *The Rand Journal of Economics, 44*(2), 275-312.
Gallini, N. T., & Wright, B. D. (1990). Technology transfer under asymmetric information. *The Rand Journal of Economics*, 147-160.
Gambardella, A., Giuri, P., & Luzzi, A. (2007). The market for patents in Europe. *Research Policy, 36*(8), 1163-1183.
Gans, J. S., Hsu, D. H., & Stern, S. (2008). The impact of uncertain intellectual property rights on the market for ideas: Evidence from patent grant delays. *Management Science, 54*(5), 982-997.
Gans, J. S., & Stern, S. (2003). The product market and the market for "ideas": commercialization strategies for technology entrepreneurs. *Research Policy, 32*(2), 333-350.
Gans, J. S., & Stern, S. (2010). Is there a market for ideas? *Industrial and Corporate Change*, 19(3): 805-837.
Graham, S., & Hegde, D. (2015). Disclosing patents' secrets. *Science, 347*(6219), 236-237.
Hall, B. H., Helmers, C., & Von Graevenitz, G. (2015). *Technology entry in the presence of patent thickets*. National Bureau of Economic Research.
Hochberg, Y. V., Serrano, C. J., & Ziedonis, R. H. (2014). *Patent collateral, investor commitment, and the market for venture lending*. National Bureau of Economic Research.
Klepper, S., & Sleeper, S. (2005). Entry by spinoffs. *Management Science, 51*(8), 1291-1306.
Kuhn, J. M., Roin, B. N., & Thompson, N. C. (2016). *Causal Inference on Patent Protection*.
Kuhn, J. M., & Younge, K. A. (2016). *Patent Citations: An Examination of the Data Generating Process.*
Lamoreaux, N. R., & Sokoloff, K. L. (1999). Inventors, firms, and the market for technology in the late nineteenth and early twentieth centuries *Learning by doing in markets, firms, and countries* (pp. 19-60): University of Chicago Press.
Lamoreaux, N. R., & Sokoloff, K. L. (2001). Market trade in patents and the rise of a class of specialized inventors in the 19th-century United States. *The American Economic Review, 91*(2), 39-44.
Lemley, M. A. (2001). Rational ignorance at the patent office. *Northwestern University Law Review, 95*(4).
Lemley, M. A., & Sampat, B. (2012). Examiner characteristics and patent office outcomes. *Review of Economics and Statistics, 94*(3), 817-827.
Lemley, M. A., & Sampat, B. N. (2008). Is the Patent Office a Rubber Stamp? *Emory Law Journal, 58*, 181.
Lemley, M. A., & Shapiro, C. (2005). Probabilistic patents. *The Journal of Economic Perspectives, 19*(2), 75-98.
Manning, C., Raghavan, P., & Schütze, H. Introduction to Information Retrieval: Cambridge: 2008: Series.
Merges, R. P. (1999). As many as six impossible patents before breakfast: Property rights for business concepts and patent system reform. *Berkeley Technology Law Journal*, 577-615.
Merges, R. P., & Nelson, R. R. (1990). On the complex economics of patent scope. *Columbia Law Review, 90*(4), 839-916.
Merges, R. P., & Nelson, R. R. (1994). On limiting or encouraging rivalry in technical progress: the effect of patent scope decisions. *Journal of Economic Behavior & Organization, 25*(1), 1-24.
Mowery, D. C., Nelson, R. R., Sampat, B. N., & Ziedonis, A. A. (2001). The growth of patenting and licensing by US universities: an assessment of the effects of the Bayh–Dole act of 1980. Research policy, 30(1), 99-119.
Niioka, H. (2006). Patent auctions: business and investment strategy in IP commercialization. *Journal of Intellectual Property Law & Practice, 1*(11), 728-731.
Pinchot, G. (1987). Innovation through intrapreneuring. *Research management, 30*(2), 14-19.
Richardson, K., Oliver, E., & Costa, M. (2016). The 2015 Brokered Patent Market: A Good Year to be a Buyer.
Risch, M. (2011). Patent troll myths.

Roth, A. E. (2008). What have we learned from market design? *The Economic Journal, 118*(527), 285-310.
Salton, G., Wong, A., & Yang, C.-S. (1975). A vector space model for automatic indexing. *Communications of the ACM, 18*(11), 613-620.
Sampat, B., & Williams, H. L. (2015). *How do patents affect follow-on innovation? Evidence from the human genome*. National Bureau of Economic Research.
Serrano, C. J. (2010). The dynamics of the transfer and renewal of patents. *The Rand Journal of Economics, 41*(4), 686-708.
Serrano, C. J. (2011). *Estimating the gains from trade in the market for innovation: Evidence from the transfer of patents*. National Bureau of Economic Research.
Somaya, D. (2002). Theoretical perspectives on patent strategy. *Strategic management journal, 24*(1).
Somaya, D. (2003). Strategic determinants of decisions not to settle patent litigation. *Strategic management journal, 24*(1), 17-38.
Somaya, D. (2012). Patent strategy and management an integrative review and research agenda. *Journal of Management, 38*(4), 1084-1114.
Somaya, D., Kim, Y., & Vonortas, N. S. (2011). Exclusivity in licensing alliances: using hostages to support technology commercialization. *Strategic management journal, 32*(2), 159-186.
Teece, D. J. (1977). Technology transfer by multinational firms: The resource cost of transferring technological know-how. *The Economic Journal, 87*(346), 242-261.
Teece, D. J. (1980). Economies of scope and the scope of the enterprise. *Journal of Economic Behavior & Organization, 1*(3), 223-247.
Teece, D. J. (1982). Towards an economic theory of the multiproduct firm. *Journal of Economic Behavior & Organization, 3*(1), 39-63.
Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy, 15*(6), 285-305.
Thompson, N., & Kuhn, J. (2016). *Does Winning a Patent Race Lead to More Follow-on Innovation?*
Von Graevenitz, G., Wagner, S., & Harhoff, D. (2011). How to measure patent thickets—A novel approach. *Economics letters, 111*(1), 6-9.
Williamson, O. E. (1975). Markets and hierarchies. *New York*, 26-30.
Younge, K. A., & Kuhn, J. M. (2015). *Patent-to-Patent Similarity: A Vector Space Model*. Available at SSRN.
Ziedonis, R. H. (2004). Don't fence me in: Fragmented markets for technology and the patent acquisition strategies of firms. *Management Science, 50*(6), 804-820.

Tables and Figures

Table 1: Variable definitions and summary statistics

	Variable	Level	Definition	Mean	SD
1	Reassigned	Patent	Dummy = 1 if patent is transferred post-grant	0.057	0.231
2	Reassigned - any sale	Patent	Dummy = 1 if patent is transferred post-grant in a sale	0.041	0.198
3	Reassigned - any exit	Patent	Dummy = 1 if patent is transferred post-grant in a firm exit	0.007	0.083
4	Reassigned - any spinoff	Patent	Dummy = 1 if patent is transferred post-grant in a spinoff	0.015	0.121
5	Reassigned - small sale	Patent	Dummy = 1 if patent is transferred post-grant in a sale of 20 or fewer patents	0.015	0.122
6	Reassigned - small exit	Patent	Dummy = 1 if patent is transferred post-grant in a firm exit of 20 or fewer patents	0.005	0.069
7	Reassigned - small spinoff	Patent	Dummy = 1 if patent is transferred post-grant in a spinoff of 20 or fewer patents	0.007	0.083
8	Firm size	Patent	The log of the number of patents issued to the firm in the five years prior to the focal patent	4.492	3.236
9	Forward citations (total)	Patent	The number of forward citations to the focal patent	8.233	17.524
10	Forward citations (self)	Patent	The number of forward self-citations to the focal patent	0.809	3.166
11	Forward citations (102)	Patent	The number of forward blocking (novelty) citations to the focal patent	0.045	0.238
12	Forward citations (103)	Patent	The number of forward blocking (non-obviousness) citations to the focal patent	0.082	0.330
13	Familiarity	Patent	The normalized maximum textual similarity between the focal patent and any prior patent by the same firm	0.000	0.978
14	Originality	Patent	The normalized maximum textual similarity between the focal patent and any prior patent by a different firm, times -1.	0.000	0.999
15	Market congestion	Art unit	The normalized mean maximum textual similarity between the focal patent and any prior patent application	0.000	1.000
16	Market thickness	Art unit	The logged ratio of unique buyers to patents	0.000	1.000

Table 2: Summary statistics by industry

Industry	Observations	Sale	Spinoff	Exit	Congestion	Thickness
All	339,769	4.09%	1.49%	0.69%	0.00	-0.00
Electrical	53,368	5.45%	2.09%	0.69%	-0.18	-0.43
Semiconductors	123,404	3.49%	1.47%	0.53%	0.16	-0.06
Chemical	43,990	2.51%	1.30%	1.03%	-0.61	0.12
Computers	30,199	11.30%	1.26%	0.79%	-0.58	0.15
Miscellaneous	39,745	2.27%	1.37%	0.48%	0.40	0.26
Mechanical	41,605	2.27%	1.13%	0.49%	0.52	0.12
Medical devices	7,458	4.16%	2.15%	3.35%	-0.45	0.72

Table 3: First stage by industry

	Dependent variable:						
	Patent scope						
	All	Electrical	Semicond.	Computers	Chemical	Mechanical	Medical
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Examiner score	0.183***	0.274***	0.148***	0.246***	0.135***	0.162***	0.152***
	(0.004)	(0.007)	(0.004)	(0.014)	(0.007)	(0.007)	(0.017)
Observations	339,769	53,368	123,404	30,199	43,990	41,605	7,458
R^2	0.046	0.085	0.029	0.089	0.022	0.043	0.042
Adjusted R^2	0.043	0.083	0.028	0.083	0.017	0.040	0.036
Residual Std. Error	0.961	0.941	0.933	1.007	1.118	0.890	0.952

Note: *p<0.1; **p<0.05; ***p<0.01
All variables normalized as z-scores. All models run with application year by art unit fixed effects. Standard errors are clustered at the firm level.

Table 4: Balance statistics (t-tests).

Variable	Broad	Narrow	T-Stat	Normalized Difference
Number of inventors	2.544	2.554	−1.259	0.004
Number of words in patent	7,323.451	7,669.633	−9.707	0.033
Number of claims	20.717	20.794	−0.747	0.003
Number of words in first claim	132.472	135.838	−8.240	0.028

Observations: 169,880

Table 5: Sales and forward citations

	Dependent variable:			
	Total (logged)	Self (logged)	Blocks others (102)	Blocks others (103)
	OLS		Logistic (odds ratio)	
	(1)	(2)	(3)	(4)
Is sold	0.305***	−0.137***	2.312***	1.206***
	(0.015)	(0.008)	(0.515)	(0.171)
Observations	339,769	339,769	339,769	339,769
Note:	*p<0.1; **p<0.05; ***p<0.01 All models run with application year fixed effects.			

Table 6: Sales and technological relatedness

	Dependent variable:			
	Odds ratio of sale			
	Large firms		Small firms	
	(1)	(2)	(3)	(4)
Familiarity	0.840***		0.954*	
	(0.027)		(0.025)	
Originality		0.888***		1.136***
		(0.029)		(0.022)
Observations	123,807	123,963	80,462	110,609
Log Likelihood	-5,610.591	-5,634.137	-8,254.129	-12,675.660
Akaike Inf. Crit.	11,227.180	11,274.270	16,514.260	25,357.320
Note:	*p<0.1; **p<0.05; ***p<0.01 All models run with application year fixed effects.			

Table 7: Main results by industry, endogenous regression

	Dependent variable:						
	Likelihood of sale (percentage points)						
	All	Electrical	Semicond.	Computers	Chemical	Mechanical	Medical
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Patent scope	0.085**	0.093	0.029	0.142*	0.132	0.151	0.080
	(0.042)	(0.077)	(0.088)	(0.079)	(0.082)	(0.103)	(0.175)
Observations	339,769	53,368	123,404	30,199	43,990	41,605	7,458
R^2	0.076	0.013	0.012	0.016	0.220	0.080	0.334
Adjusted R^2	0.073	0.011	0.010	0.010	0.216	0.077	0.329
Residual Std. Error	11.445	10.578	11.598	9.672	12.115	11.984	15.050
Note:	*p<0.1; **p<0.05; ***p<0.01 Covariate normalized as a z-score. All models run with application year by art unit fixed effects. Standard errors are clustered at the firm level.						

Table 8: Main results by industry, reduced form

	Dependent variable:						
	Likelihood of sale (percentage points)						
	All	Electrical	Semicond.	Computers	Chemical	Mechanical	Medical
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Examiner score	0.174***	0.074	0.040	0.275**	0.291*	0.405***	1.384***
	(0.047)	(0.056)	(0.094)	(0.115)	(0.150)	(0.124)	(0.469)
Observations	339,769	53,368	123,404	30,199	43,990	41,605	7,458
R^2	0.076	0.013	0.012	0.017	0.220	0.081	0.338
Adjusted R^2	0.073	0.011	0.010	0.010	0.217	0.078	0.333
Residual Std. Error	11.444	10.578	11.598	9.669	12.113	11.979	15.004

Note: *p<0.1; **p<0.05; ***p<0.01
Covariate normalized as a z-score. All models run with application year by art unit fixed effects. Standard errors are clustered at the firm level.

Table 9: Main results by industry, instrumental variable

	Dependent variable:						
	Likelihood of sale (percentage points)						
	All	Electrical	Semicond.	Computers	Chemical	Mechanical	Medical
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Patent scope	0.951***	0.269	0.273	1.115**	2.155*	2.510***	9.115***
	(0.255)	(0.206)	(0.636)	(0.468)	(1.117)	(0.775)	(3.295)
Observations	339,769	53,368	123,404	30,199	43,990	41,605	7,458
R^2	0.071	0.013	0.012	0.006	0.193	0.051	0.112
Adjusted R^2	0.068	0.011	0.010	-0.001	0.189	0.048	0.106
Residual Std. Error	11.476	10.579	11.600	9.724	12.327	12.172	17.375

Note: *p<0.1; **p<0.05; ***p<0.01
Covariate normalized as a z-score. All models run with application year by art unit fixed effects. Standard errors are clustered at the firm level.
A one standard deviation increase in examiner score increases patent scope by about .2 standard deviations.

Table 10: Instrumental variable results by industry with firm size interactions

	Dependent variable:						
	Likelihood of sale (percentage points)						
	All	Electrical	Semicond.	Computers	Chemical	Mechanical	Medical
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Patent scope	0.949***	0.281	0.437	1.503**	2.480**	2.166***	9.014***
	(0.255)	(0.230)	(0.705)	(0.617)	(1.175)	(0.639)	(3.298)
Firm size	−0.272***	−0.317***	−0.304***	−0.350***	−0.087	−0.226***	−0.391**
	(0.019)	(0.050)	(0.031)	(0.039)	(0.070)	(0.039)	(0.172)
Scope * size	−0.145**	−0.025	−0.299	−0.352**	0.481*	−0.351*	0.096
	(0.071)	(0.072)	(0.217)	(0.139)	(0.256)	(0.203)	(0.815)
Observations	339,769	53,368	123,404	30,199	43,990	41,605	7,458
R^2	0.074	0.021	0.014	0.007	0.155	0.046	0.127
Adjusted R^2	0.071	0.019	0.012	-0.00005	0.151	0.043	0.121
Residual Std. Error	11.455	10.536	11.589	9.719	12.614	12.203	17.229

Note: *p<0.1; **p<0.05; ***p<0.01
Patent scope normalized as a z-score. All models run with application year by art unit fixed effects. Standard errors are clustered at the firm level.
Firm size operationalized as the log of patents issued in the past two years.
A one standard deviation increase in examiner score increases patent scope by about .2 standard deviations.

Table 11: Main results by industry, logistic regression

	Dependent variable:						
	Logit of sale						
	All	Electrical	Semicond.	Computers	Chemical	Mechanical	Medical
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Examiner score	0.200***	0.067	0.031	0.261***	0.215***	0.307***	0.824***
	(0.015)	(0.042)	(0.027)	(0.061)	(0.044)	(0.046)	(0.106)
Observations	339,769	53,368	123,404	30,199	43,990	41,605	7,458
Log Likelihood	-25,080.010	-3,108.598	-8,378.250	-1,455.206	-3,018.093	-2,872.838	-648.316
Akaike Inf. Crit.	50,198.010	6,437.195	17,130.500	3,318.412	6,446.185	6,013.676	1,400.632

Note: *p<0.1; **p<0.05; ***p<0.01
Covariate normalized as a z-score. All models run with application year by art unit fixed effects.

Table 12: Logistic regressions by transfer type and industry-reported patent effectiveness

	Dependent variable:			
	Logit of sale			
	Product		Process	
	(1)	(2)	(3)	(4)
Examiner score	0.221***	0.206***	0.227***	0.231***
	(0.020)	(0.020)	(0.020)	(0.020)
CNW Overall	0.163***		0.347***	
	(0.023)		(0.024)	
CNW Exclusion		−0.026		−0.254***
		(0.031)		(0.029)
Score x CNW Overall	0.177***		0.119***	
	(0.018)		(0.019)	
Score x CNW Exclusion		0.256***		0.237***
		(0.019)		(0.020)
Observations	207,037	207,037	207,037	207,037
R^2	0.031	0.031	0.035	0.031
Note:	*p<0.1; **p<0.05; ***p<0.01 All covariates normalized as z-scores. All models run with application year by technology center fixed effects.			

Table 13: Logistic regressions by transfer type and industry-reported patent effectiveness

	Dependent variable:			
	Logit of sale			
	Product		Process	
	(1)	(2)	(3)	(4)
Examiner score	0.249***	0.248***	0.259***	0.259***
	(0.020)	(0.020)	(0.020)	(0.020)
CNW Cross-licensing	0.195***		0.331***	
	(0.026)		(0.024)	
CNW Negotiation		0.188***		0.211***
		(0.027)		(0.023)
Score x CNW Cross-licensing	−0.007		−0.073***	
	(0.019)		(0.019)	
Score x CNW Negotiation		−0.013		−0.128***
		(0.019)		(0.019)
Observations	207,037	207,037	207,037	207,037
R^2	0.027	0.026	0.031	0.029
Note:	*p<0.1; **p<0.05; ***p<0.01 All covariates normalized as z-scores. All models run with application year by technology center fixed effects.			

Table 14: Logistic regressions by transfer type and market characteristics

	Dependent variable:								
	Logit of sale			Logit of exit			Logit of spinoff		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Examiner score	0.200***	0.177***	0.189***	0.219***	0.184***	0.118***	0.121***	0.126***	0.121***
	(0.015)	(0.017)	(0.016)	(0.027)	(0.028)	(0.033)	(0.022)	(0.023)	(0.022)
Market thickness		0.454***			0.291***			0.330***	
		(0.016)			(0.028)			(0.022)	
Patent congestion			−0.298***			−0.832***			−0.174***
			(0.017)			(0.026)			(0.024)
Score x thickness		0.062***			0.122***			−0.022	
		(0.016)			(0.027)			(0.022)	
Score x congestion			−0.070***			−0.153***			−0.023
			(0.016)			(0.025)			(0.023)
Observations	339,769	339,769	339,769	339,769	339,769	339,769	339,769	339,769	339,769
R^2	0.018	0.037	0.026	0.053	0.061	0.105	0.015	0.023	0.017

Note: *p<0.1; **p<0.05; ***p<0.01
All covariates normalized as z-scores. All models run with application year by technology center fixed effects.

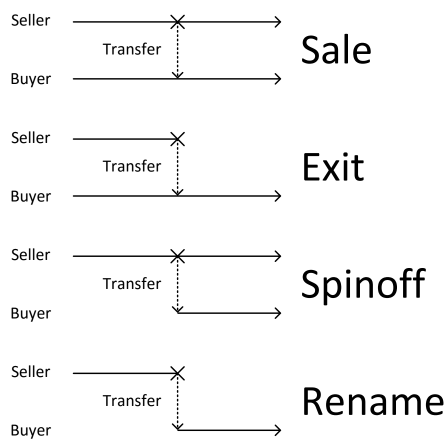

Figure 1: Representation of the logic used to identify patent transfer type.

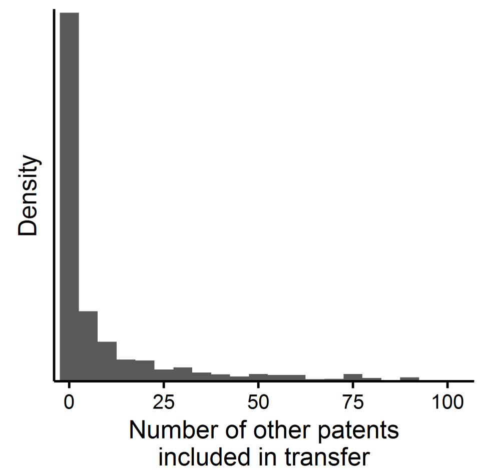

Figure 2: The density of patent transfers by the number of patents included in the transfer.

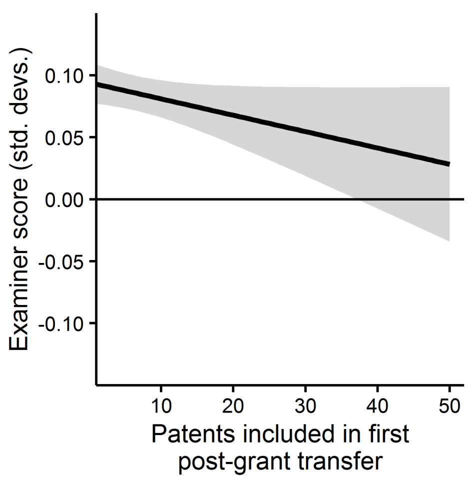

Figure 3: Linear relationship between examiner score and the number of patents in a transfer.

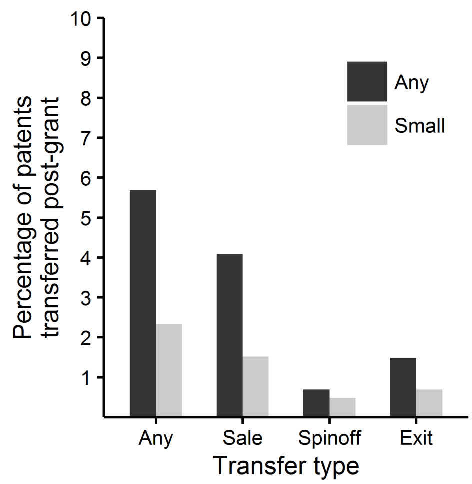

Figure 4: Percentage of patents transferred. Small transfers include 25 or fewer patents.

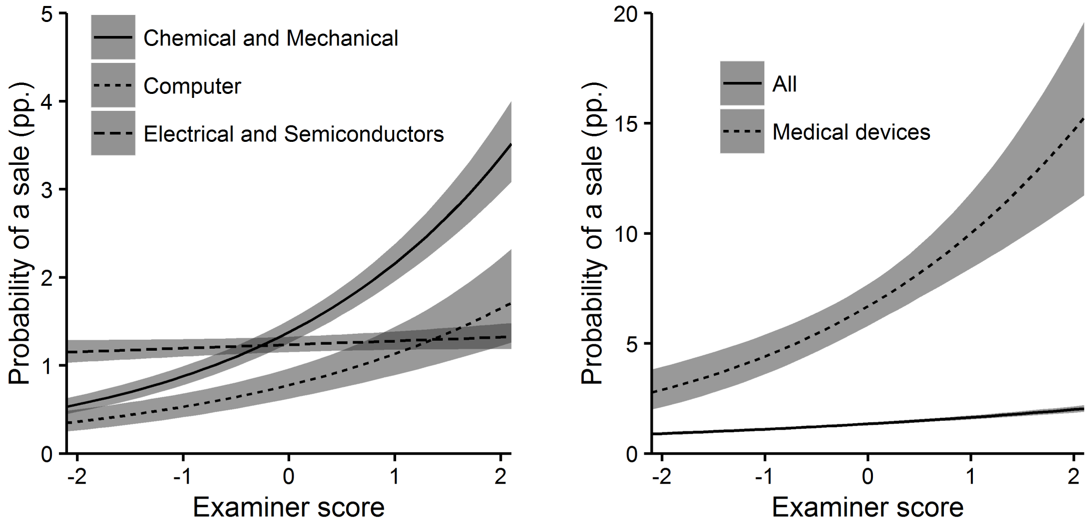

(a) Sales by industry **(b)** Sales of all patents vs. medical device patents

Figure 5: Probability that a patent is sold as a function of examiner score, by industry, predicted based on logistic regression estimates.

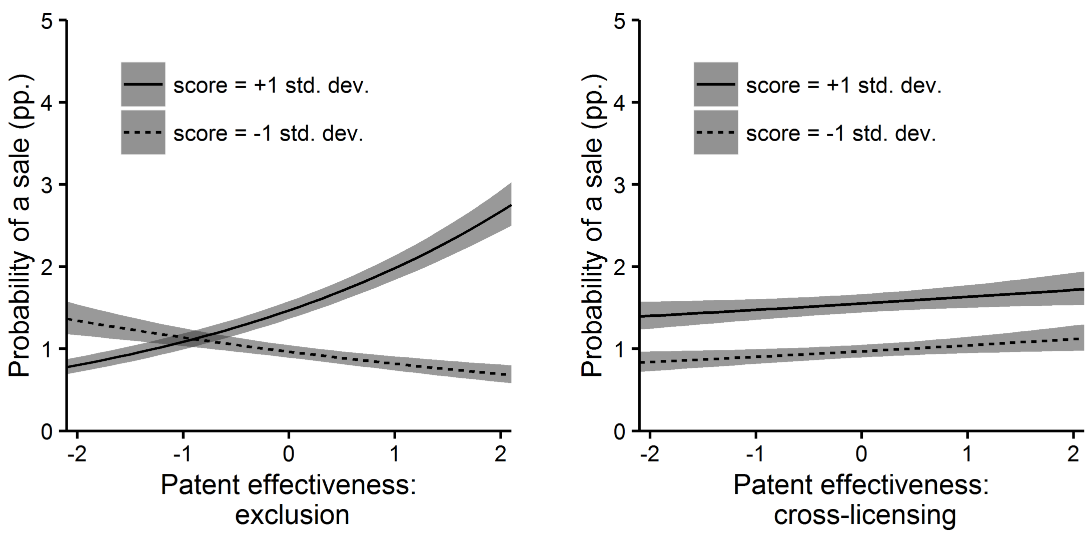

(a) Effectiveness for exclusion **(b)** Effectiveness for cross-licensing

Figure 6: Probability that a patent is sold as a function of industry-reported product patent effectiveness described by Cohen, Nelson & Walsh (2000), by examiner score, predicted based on logistic regression estimates.

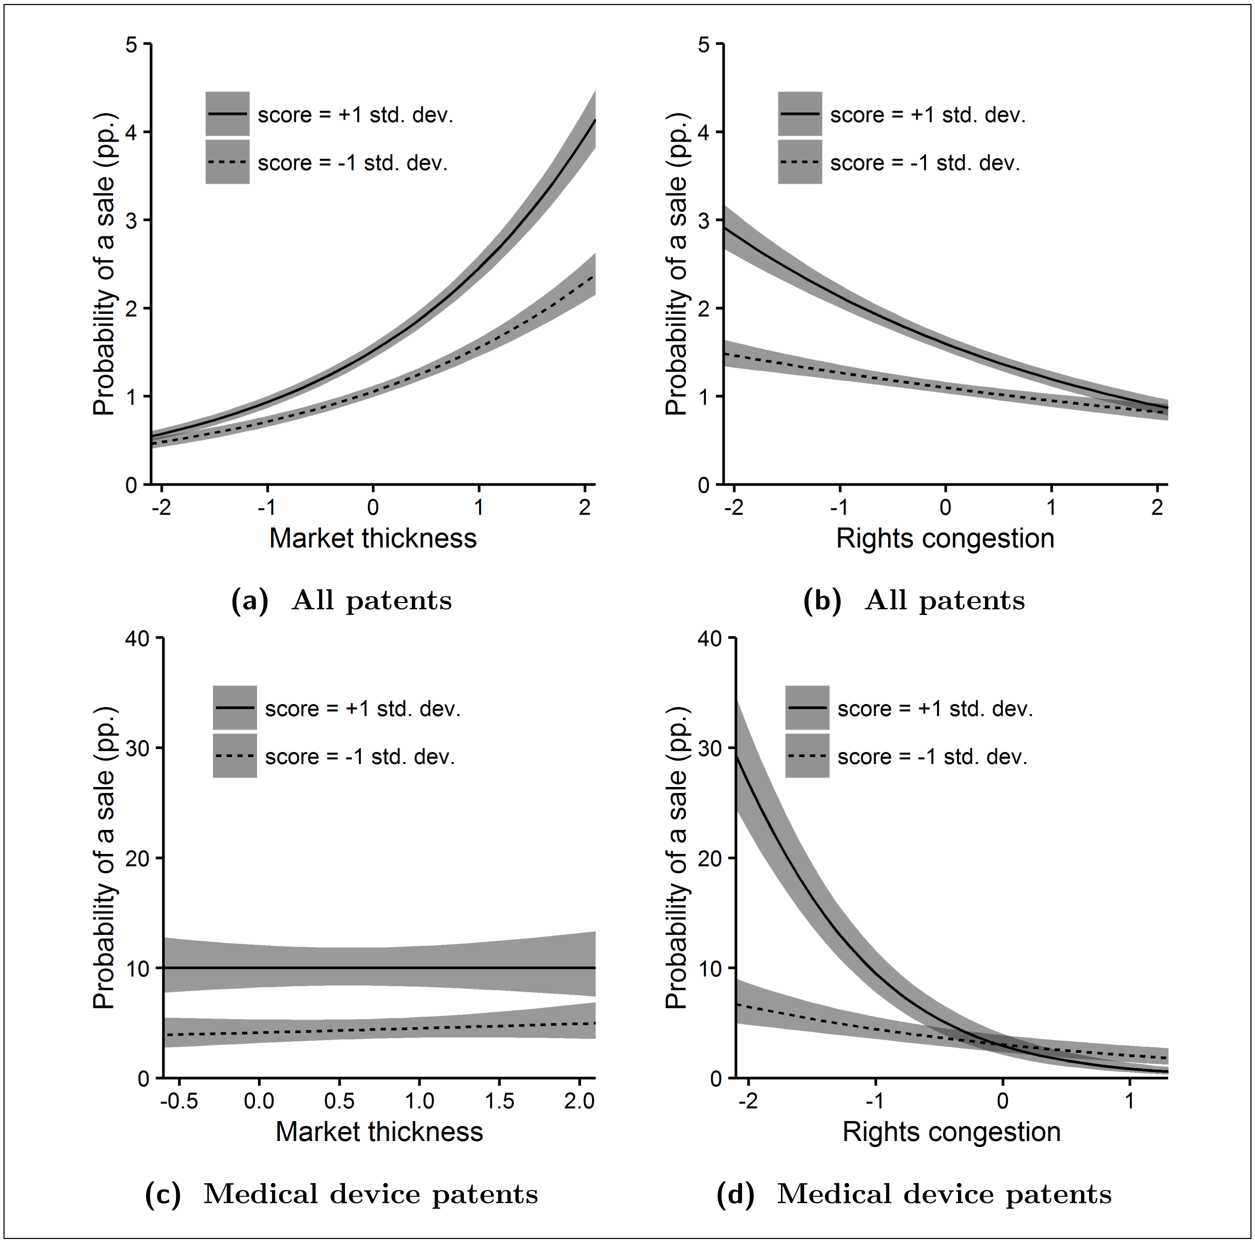

(a) All patents

(b) All patents

(c) Medical device patents

(d) Medical device patents

Figure 7: Probability that a patent is sold as a function of market characteristics, by examiner score, predicted based on logistic regression estimates.

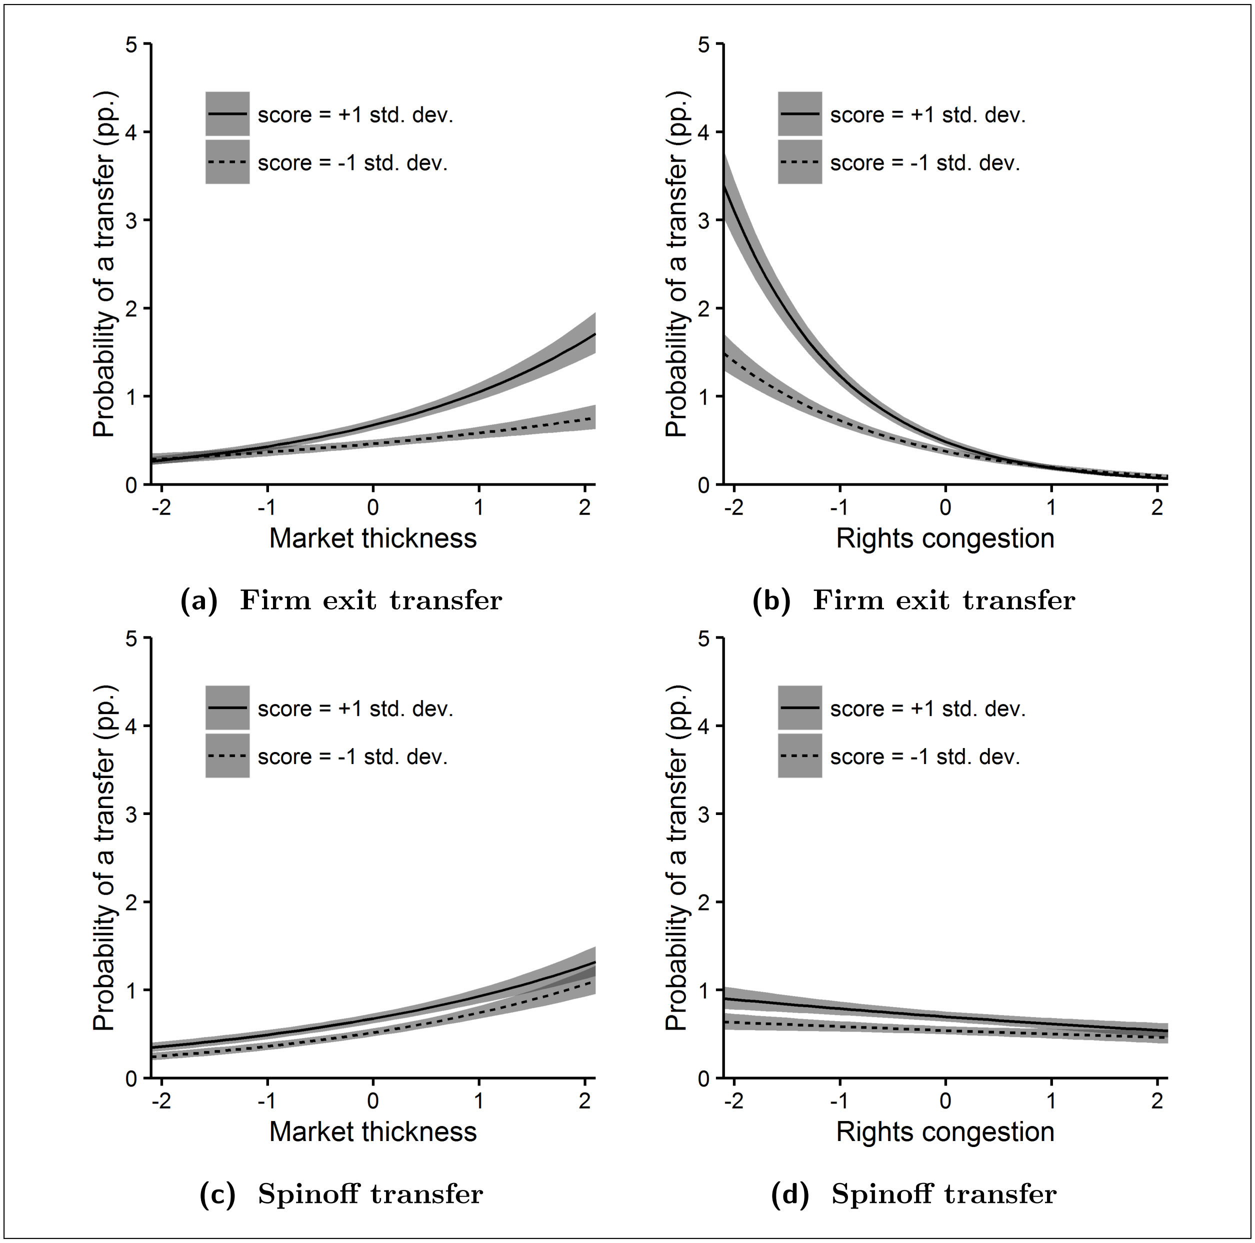

(a) Firm exit transfer (b) Firm exit transfer

(c) Spinoff transfer (d) Spinoff transfer

Figure 8: Probability that a patent is transferred via a firm exit or a spinoff as a function of market characteristics, by examiner score, predicted based on logistic regression estimates.

Appendix A: Robustness checks

Tables A.1-A.3 present the results of estimating models similar to those estimated in Tables 12-14, but with using instrumental variable regression instead of logistic regression. Each table includes both heteroscedasticity-consistent standard errors (top) and standard errors clustered at the firm level (bottom). In all models, the estimated coefficients for interaction terms are the same sign as those presented in the main results. With heteroscedasticity-consistent standard errors, significance is maintained throughout. Moreover, the heteroscedasticity-consistent are quite similar to the (unreported) classical standard errors, which provides some evidence that the model is not misspecified. (King & Roberts 2015).

The reported significance levels are based on the clustered standard errors. The significance of some of the interaction terms is not maintained under clustering. This fact, combined with the low likelihood (1-2%) that any particular patent is sold, suggests that the reduced form conditional logistic regression model for which results are reported in the main text provides a better fit for the data.

Table A1: Instrumental variable regressions by CNW-reported patent effectiveness overall and for exclusion

	Dependent variable:			
	Likelihood of sale (percentage points)			
	Product		Process	
	(1)	(2)	(3)	(4)
Patent scope	1.684*	1.661*	1.707*	1.651*
	(0.137)	(0.140)	(0.136)	(0.138)
	(1.006)	(0.986)	(1.023)	(0.990)
CNW Overall	0.315		0.666**	
	(0.035)		(0.039)	
	(0.283)		(0.313)	
CNW Exclusion		−0.004		−0.355
		(0.051)		(0.047)
		(0.314)		(0.283)
Scope x CNW Overall	1.706**		1.202	
	(0.145)		(0.127)	
	(0.859)		(0.934)	
Scope x CNW Exclusion		2.745		2.042
		(0.170)		(0.159)
		(2.720)		(1.844)
Observations	207,037	207,037	207,037	207,037
R^2	-0.038	-0.082	-0.023	-0.055

Note: *p<0.1; **p<0.05; ***p<0.01
All covariates normalized as z-scores. All models run with application year by technology center fixed effects. Top standard errors are heteroskedasticity-robust. Bottom standard errors are clustered by firm. Significance levels based on clustered standard errors.

Table A2: Instrumental variable regressions by CNW-reported patent effectiveness for cross-licensing and negotiation

	Dependent variable:			
	Likelihood of sale (percentage points)			
	Product		Process	
	(1)	(2)	(3)	(4)
Patent scope	1.678*	1.674*	1.694*	1.669*
	(0.136)	(0.136)	(0.136)	(0.136)
	(1.008)	(1.008)	(1.006)	(0.994)
CNW Cross-licensing	0.332		0.604**	
	(0.039)		(0.041)	
	(0.257)		(0.270)	
CNW Negotiation		0.320		0.354**
		(0.041)		(0.037)
		(0.255)		(0.179)
Scope x CNW Cross-licensing	−0.004		−0.347	
	(0.170)		(0.155)	
	(0.639)		(0.402)	
Scope x CNW Negotiation		−0.034		−0.907
		(0.159)		(0.154)
		(0.511)		(0.638)
Observations	207,037	207,037	207,037	207,037
R^2	-0.014	-0.014	-0.015	-0.022
Note:	*p<0.1; **p<0.05; ***p<0.01 All covariates normalized as z-scores. All models run with application year by technology center fixed effects. Top standard errors are heteroskedasticity-robust. Bottom standard errors are clustered by firm. Significance levels based on clustered standard errors.			

Table A3: Instrumental variable regressions by transfer type and market characteristics, clustered standard errors

	Dependent variable:								
	Likelihood of transfer (percentage points)								
	Sale			Exit			Spinoff		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Patent scope	1.461**	1.498**	1.454**	0.540	0.551	0.534	0.426***	0.434***	0.424***
	(0.114)	(0.114)	(0.114)	(0.066)	(0.066)	(0.066)	(0.078)	(0.078)	(0.078)
	(0.694)	(0.701)	(0.689)	(0.679)	(0.685)	(0.674)	(0.142)	(0.143)	(0.142)
Market thickness		0.564***			0.120***			0.209***	
		(0.021)			(0.012)			(0.015)	
		(0.065)			(0.046)			(0.025)	
Market congestion			−0.428**			−0.415*			−0.127***
			(0.024)			(0.014)			(0.017)
			(0.210)			(0.223)			(0.034)
Scope x thickness		0.971*			0.367			0.093	
		(0.101)			(0.058)			(0.069)	
		(0.504)			(0.373)			(0.079)	
Scope x congestion			−0.663			−0.536			−0.138
			(0.077)			(0.045)			(0.053)
			(0.535)			(0.579)			(0.085)
Observations	339,769	339,769	339,769	339,769	339,769	339,769	339,769	339,769	339,769
R^2	-0.011	-0.014	-0.012	-0.001	-0.004	-0.004	-0.001	-0.001	-0.001

Note: *p<0.1; **p<0.05; ***p<0.01
All covariates normalized as z-scores. All models run with application year by technology center fixed effects. Top standard errors are heteroskedasticity-robust. Bottom standard errors are clustered by firm. Significance levels based on clustered standard errors.

Appendix B: Data Construction

B.1 Supplemental discussion of variable definitions and sample selection

The sample includes all patents filed 2004-2006 with four exceptions. Although these restrictions impose some limitations on the external validity of the findings in this article and although not all of these restrictions are strictly necessary, the following exclusions aid in interpretability and in making the strongest possible case for valid causal inference. First, the sample excludes patents that were not published prior to issuance. Patent applicants can elect not to publish a patent application prior to its issuance as a patent, although in practice applicants choose non-publication for only about 7% of patents. (Graham & Hegde, 2015). These patents are excluded from the sample due to the inability to observe their features at the time they were filed, which as discussed in Section 4.2 is important for the identification strategy. Second, the sample excludes patents examined within USPTO Technology Center 1600, responsible for examining biotechnology inventions, because the instrument has not been validated for use with biotechnology patents. (Kuhn et al. 2016). Third, the sample also excludes all patents that claim priority to a previously-filed patent application in order to avoid introducing endogeneity into the instrument. (Kuhn et al. 2016). Fourth, because this paper examines the effect of technology transfer between firms, the sample excludes any patent not already assigned to a firm at the time of issuance. The conclusions in this article are thus limited to original patents filed by firms in technologies other than biotechnology and published pre-grant.

The initial date of the sample was selected to capitalize on a change in U.S. patent law. Patents filed in or after 2001 are published by default after 18 months, allowing me to observe the claims of those patents not only at the time of issuance, but also at the time of filing. This in turn allows the construction of a measure of the causal instrument. Publicly available patent metadata is more complete for patents issued in 2005 or later, so the initial date of the sample is restricted to patents filed during or after the beginning of 2004.

The terminal date of the sample was selected to mitigate rightward truncation, allowing sufficient time for patents to issue and potentially be sold. Figure B1 shows the density of the lag between patent filing and the first post-grant patent transfers for patents filed in 2001 or earlier. Because the patent assignment data is complete through 2015, each patent in this sub-sample has at least 14 years to potentially be transferred, and most have even more time. Few patents are transferred so long after filing. Nevertheless, the long lag for some initial post-grant transfers illustrates that any analysis will be subject to at least some rightward truncation since data availability limits the identification strategy to patents filed in 2001 or later. In order to minimize rightward truncation as much as possible, the sample is limited to patents filed 2006 or earlier. Although the assignment data includes assignments filed in 2015, the final year of data is used for transfer-type classification purposes only. Thus, each patent in the sample is at hazard of being assigned up until the end of 2014, giving each patent at least 8 years to be transferred.

— Insert Figure B1 and Figure B2 about here. —

Firm size—variable 9 in Table 1—is operationalized as the log of the number of patents the firm received in the five years preceding the filing of the focal patent. Patent counts are used to measure firm size for several reasons. Although firm size can be measured in different ways, the relevant aspect of firm size for this analysis is not the firm's headcount or financial footprint but rather the extent of the firm's participation in the patent system. Also, systematic data measuring the size of small, private firms is generally lacking, so measuring a firm's size based on its footprint in the patent system provides consistency across observations. Firm size is calculated as of the filing year of each patent to capture changes over time. Figure B2 shows a histogram of firm size on the patent level.

The tenth, eleventh, twelfth, and thirteenth variables in Table 1 measure counts of different types of forward citations received by the focal patent. Patents in the sample receive on average 8.4 total citations and about .8 citations by the firm that owned the patent at the time of issuance. Total and self citations were identified from USPTO XML data. A blocking citation (102 or 103) indicates an instance where a patent examiner cited the focal patent to reject the claims of a later-filed patent. As described in Section 4.2, the scope of most patent applications is narrowed during examination at the patent office by a process of negotiation between the patent applicant and the patent examiner. Typically, the patent examiner rejects the patent's initial claims as being too broad, and the applicant responds by amending those claims to encompass a narrower invention. Blocking citations are the evidence the examiner provides to support such claim rejections. A claim is rejected under 35 U.S.C. 102 for failing the novelty requirement if the claimed invention is entirely described in a single prior art reference, while a rejection under 35 U.S.C. 103 is made when the claimed invention is obvious in view of some combination of prior art references. Blocking citations were identified by converting more than 50 million pages of USPTO correspondence from image to text and then using regular expressions to extract the patent numbers for the blocking citations from the resulting text files. Thompson and Kuhn (2016), who use blocking citations as a source of exogeneity to analyze the effects of winning a patent race, provide more detail on both the identification and significance of blocking citations.

The sixteenth and seventeenth variables in Table 1 describe characteristics of the local market for technology in which the patent is situated. A market definition based on a company-level indicator, such as an SIC code, would be too coarse to examine heterogeneity across markets for different types of technology because many large firms span technological boundaries. Instead, I rely on internal classification of patents by the USPTO. When the USPTO initially processes a patent application, a reviewer assigns the application for examination to one of 531 art units based on the technology described and claimed in the application. The art units, which are split among eight Technology Centers, each include a number of patent examiners skilled in examining a particular type of technology, such as "Internal Combustion Engines" or "Wells and Earth Boring." I define each patent's local technology market as the USPTO art unit which issued the patent. For the analysis in Table 12, Table A3, and Table A4, I linked many of the art units to the industry categories identified by Cohen et al. (2000) ("CNW") to these USPTO art units by manual inspection. Some

USPTO art units do not correspond to any category in CNW and are therefore are not linked. Similarly, some categories in CNW are either too broad (e.g., "2411: Basic Chemicals") or too specific (e.g., "3212: Precision Instruments" to be linked to a distinct USPTO art unit or group of art units. Patents in art units not linked to an industry category were omitted from regressions using CNW variables. The linkages, along with the number of observations and the percentage of patents transferred post-grant via sale, spinoff, or exit are shown in Table B2.

— Insert Table B2 about here. —

Table B1 presents the Pearson correlation between each variable. Due to the large sample size, nearly all correlations are statistically significant. The correlation matrix reveals several meaningful descriptive relationships. First, areas of higher congestion also have on average lower market thickness. Second, patents on technology that is more familiar to a firm are also less original on average when compared to patents outside the firm. Third, patents originally owned by larger firms are less likely to be transferred.

— Insert Table B1 about here. —

B.2 Patent reassignment data

Under U.S. law, every patent is by default owned as common property equally and indivisibly by its inventors. However, most inventors agree to assign their inventions to their employers as a provision of their employment agreement. Both this initial assignment and subsequent purchases of a patent are formalized through an assignment agreement. An assignment agreement is a simple legal document that lists the patent right that is transferred, the assignor (seller), and the assignee (buyer). The assignee can, but is not required, to record the transfer with the USPTO, which logs the details of the assignment in a publicly available database when and if the patent is published.

In practice, patent purchasers have strong incentives to record assignments since failing to do so creates a risk that they will lose title to their property. Patent assignments are legal instruments similar in some respects to deeds to real property and certain forms of personal property such as automobiles. Failing to record a transfer with a government agency creates ambiguity as to the legal owner of the asset. For example, suppose that an entity A transferred a patent to entity B, but that entity B failed to record the assignment. If entity A then transferred the patent to entity C, which did record the assignment, then entity C becomes the legal owner of the patent. Entity B could potentially recover the patent through litigation but would need to overcome the basic presumption of ownership created by the valid assignment entity C recorded.

I downloaded the complete database of assignment data from 1975-2015 from the USPTO. Each assignment includes a text field that describes the type of assignment. Following Serrano (2010), I first employed textual analysis to exclude from consideration assignments indicative of non-substantive transfers such as security agreements, firm renamings, and bequests. For example, I excluded assignments for which the text field included words such as "lien," "security," "mortgage," "satisfaction," or "reversion."

Next, I developed a novel disambiguation algorithm to create a unique identifier for each assignee and assignor. In the raw data, the parties are identified only by text name and address fields, both of which are subject to typographical errors or differences in how firms record their information. For instance, the same firm may record different addresses for different patents or may use an abbreviation for its name in one assignment but not in another. To address these imperfections in the data, I divided assignees by location—either U.S. state or a foreign country. I cleaned each name and address, correcting common misspellings and creating a canonical version of commonly abbreviated terms. I then grouped identical records together and gave each record within the same group the same firm identifier. Finally, I ran a clustering algorithm to coalesce groups where appropriate. Starting with groups having a small number of members, the clustering algorithm compared the focal group to all groups having a larger number of members. I used a combination of Levenshtein distance and Levenshtein ratio, depending on the length of the string, to compare both the name and address. If the match exceeded a threshold level, I coalesced the two groups by assigning the members of the smaller group the firm identifier of the larger group.

The precise criteria used to identify a match are complex and were constructed via many iterative passes of the algorithm in conjunction with manual review to identify the rates of Type I and Type II errors. However, both *ex ante* and *ex post* validation suggest that the algorithm identifies firms with a high degree of accuracy, with both types of error rates in the low single digits.

Figure B3 shows the different types of patent transfers as a proportion of all transfers. Original transfers from the inventor to a firm comprise about 88% of all transfers since most patents are initially assigned to firms and not transferred again. Sales between ongoing companies make up the bulk of the secondary market, accounting for 6% of all transfers. Less common, but still sizeable, are transfers indicative of spinoffs and firm exits, accounting for 3.8% and .8% of all transfers. Finally, .5% of transfers are indicative of renaming and are omitted from subsequent analysis. The transfers shown in Figure B3 include those that occur both before and after patent grant. However, this article focuses on the relationship between patent scope and patent sales. Since a patent's scope is not determined until shortly before the patent is granted, the analysis in this article is limited to post-grant transfers.

— Figure B3 about here. —

B.3 Patent congestion measure

The literature has not yet settled on an accepted technique for measuring patent rights congestion (i.e. patent thickets). In existing work, congestion-related concepts frequently turn on relationships between patent citations and patent ownership. For example, Cockburn et al. (2010) measure patent fragmentation based on the rate and incidence of citations made between patents owned by different firms. Von Graevenitz, Wagner, and Harhoff (2011) recently provided a novel technique for detecting the presence of patent thickets. They argued that the frequency of citation triangles—instances in which three firms all cite each other's patents—indicate the likelihood of overlap between patent rights in a particular technology. However, Kuhn and Younge (2016) demonstrate that backward patent citations suffer from substantial selection effects and exhibit a secular time trend indicating near-exponential growth, raising questions of bias for most citation-based measures built using recent data.

I measured patent market congestion using patent-to-patent textual similarity data developed and validated by Younge and Kuhn (2015). In that project, Younge and Kuhn applied a vector space similarity model (VSM) based on term-frequency inverse-document frequency (tf-idf) weighting to calculate the textual similarity to every possible patent-to-patent comparison, involving over 14 trillion calculations and producing approximately 9 terabytes of data. Each calculation yields a cosine similarity value, normalized to a value between 0 and 100, that indicates the degree of textual similarity between the full technical description of pair of patents. The vector space model ("VSM") is one of the most common and robust models in the field of Information Retrieval Theory. (Manning, Raghavan, & Schtze). Originally developed for the SMART information retrieval system (Salton, Wong, & Yang, 1975), vector space models are used to construct a high-dimensional space based on the text of a given sample (or "corpus") of documents and now motivate the operation of nearly every modern-day search engine. The dimensionality of the vector space model was constructed from the originating vocabulary used in patent documents, not the assignment of patent classifications. Doing so allows the capture of similarity between patents whose patent classifications do not overlap, a key concern raised by Bloom, Schankerman, and Van Reenen (2013).

Younge and Kuhn (2015) validate the patent-to-patent similarity measure using a range of approaches. The patent-to-patent similarity measure was an extremely significant predictor of external ratings provided by a patent attorney, by a subject matter expert, and by the average response from a crowd of lay-person reviewers. The patent-to-patent similarity measure also predicted whether two patents share a patent class and subclass better than pre-existing techniques. Finally, relative patent-to-patent similarity values correspond intuitively with other constructs such as patent citations, patent families, and the degree of technological focus of a firm.

Since a randomly selected pair of patents is unlikely to have much in common, the resulting distribution is highly skewed. Figure B4, copied from Younge and Kuhn (2015), plots a histogram of the full population of 14.0 trillion values with a with a sharp spike at zero related to the 288.2 TB of calculated results that fell below the threshold of 0.10 and were therefore assumed to be zero, as well as the 8.1 TB portion of the distribution that fell above the threshold of 0.10 and was therefore retained. Figure B5 plots the distribution above 0.10 after the raw cosine similarity score was transformed by the arcsine of the square-root to improve interpretability.

— Figure B4 and Figure B5 about here. —

The fourteenth and fifteenth variables in Table 1 were constructed based on this patent-to-patent textual similarity data. Familiarity measures the extent to which the focal patent is similar to past work by the same firm. Familiarity is operationalized by finding, for every patent, the maximum textual similarity to any other patent previously filed by the same firm. The resulting raw values are normalized to produce a z-score having mean zero and standard deviation one. Originality measures the extent to which the focal patent is different than past work by other firms. Originality is operationalized by finding, for every patent, the maximum textual similarity to any other patent previously-filed by a different same firm. The resulting raw values are normalized and multiplied by negative one. Figure B6 plots the patent-level density of the maximum similarity to any earlier-filed patent, both within firm and between firms. As would be expected, same-firm similarity has a higher variance than between-firm similarity—firms sometimes branch out into technology new to the firm, but most patents tend to be somewhat similar to at least some prior work by other firms.

— Figure B6 about here. —

I developed a measure of congestion based on the likelihood of overlapping patent rights between different firms in the spirit of Cockburn et al. (2010) and Von Graevenitz et al. (2011). For each patent in the sample, I first found the maximum patent similarity value between the focal patent and any other patent that is filed before the focal patent and is assigned to a different firm as the focal patent. This value proxies for the extent to which the focal patent is likely to conflict with a pre-existing patent right. I then determined the mean of these raw values for each art unit. Each focal patent received the raw value corresponding to its art unit, and the raw values were normalized across the entire distribution of patents to produce a z-score.

B.4 Tables and Figures

Table B1: Correlation Matrix

	1	2	3	4	5	6	7	8	9	10	11	12
Sale (1)	1.00											
Exit (2)	0.21	1.00										
Spinoff (3)	0.10	0.14	1.00									
Firm size (4)	-0.07	-0.04	-0.07	1.00								
Cites total (5)	0.04	0.05	0.01	0.01	1.00							
Cites self (6)	-0.02	0.00	-0.01	0.06	0.45	1.00						
Cites 102 (7)	0.03	0.03	0.00	0.03	0.15	0.07	1.00					
Cites 103 (8)	0.01	0.00	0.01	0.03	0.18	0.07	0.37	1.00				
Familiarity (9)	-0.01	-0.01	-0.01	0.00	0.04	0.08	0.02	0.01	1.00			
Originality (10)	0.01	0.02	0.00	0.01	-0.03	-0.03	-0.01	-0.01	-0.36	1.00		
Thickness (11)	0.05	0.02	0.03	-0.16	0.05	-0.02	0.00	0.02	-0.05	0.00	1.00	
Congestion (12)	-0.02	-0.05	-0.01	0.02	-0.01	0.02	0.03	0.02	0.09	0.00	-0.26	1.00

Table B2: Summary statistics by CNW industry

Industry	Obs.	Sale	Spinoff	Exit	Art Units
1500:Food	4,888	1.29%	0.72%	0.16%	1791, 1792
1700:Textiles	1,502	1.33%	3.66%	2.66%	3765
2100:Paper	742	31.27%	0.54%	30.46%	1741, 3782
2200:Printing/Publishing	1,761	2.16%	1.76%	0.62%	3725
2320:Petroleum	4,795	3.59%	2.63%	0.42%	1797
2400:Chemicals, nec	25,781	1.77%	1.34%	0.55%	1700
2411:Basic Chemicals	0				
2413:Plastic Resins	0				
2423:Drugs	0				
2429:Miscellaneous Chemicals	0				
2500:Rubber/Plastic	3,310	2.05%	1.30%	0.39%	1711, 1712, 1732, 1734, 1742, 1743, 1747, 1761, 1762, 1763
2600:Mineral Products	0				
2610:Glass	0				
2695:Concrete, Cement, Lime	256	0.39%	1.95%	0.00%	1731
2700:Metal, nec	4,579	2.64%	0.39%	0.98%	1726, 1733, 1735, 1736, 1793
2710:Steel	0				
2800:Metal Products	2,360	1.69%	1.10%	0.25%	3726, 3729
2910:General Purpose Mach.	0				
2920:Special Purpose Mach.	0				
2922:Machine Tools	4,484	1.27%	0.83%	0.33%	3721, 3722, 3723, 3724
3010:Computers	30,199	11.30%	1.26%	0.79%	2100, 2400
3100:Electrical Equip.	0				
3110:Motor/Generator	10,492	2.57%	0.66%	0.30%	3745, 3746, 3747, 3748
3210:Electronic Components	24,249	2.95%	1.41%	0.64%	2827, 2831, 2832, 2833, 2834, 2835, 2836, 2837, 2838, 2842, 2843, 2844, 2845, 2847, 2848, 2859
3211:Semiconductors and Equip.	17,741	3.46%	0.83%	0.32%	2812, 2813, 2814, 1818, 2822, 2823, 2825, 2826, 2891, 2892, 2893, 2894, 2895
3220:Communications Equip.	45,297	5.87%	2.06%	0.72%	2600
3230:TV/Radio	8,071	3.09%	2.30%	0.57%	2422, 2486, 2612, 2622, 2651, 2656, 2661
3311:Medical Equip.	7,458	4.16%	2.15%	3.35%	3731, 3732, 3733, 3735, 3736, 3737, 3738, 3739, 3766, 3768, 3769, 3773, 3774, 3775, 3776, 3777, 3779
3312:Precision Instruments	0				
3314:Search/Navigational Equip.	3,005	2.70%	0.87%	0.00%	3646, 3655, 3661
3410:Car/Truck	0				
3430:Autoparts	4,201	1.98%	0.71%	0.02%	3611, 3616, 3618
3530:Aerospace	1,866	1.66%	1.23%	2.52%	3641, 3644
3600:Other Manufacturing	0				

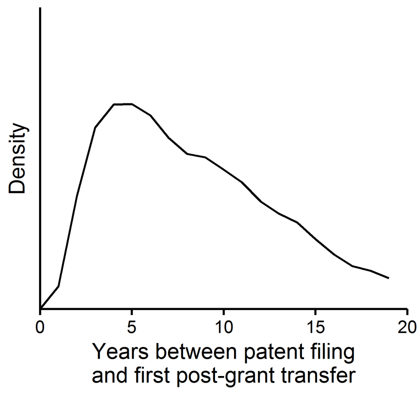

Figure B1: Density of years between patent filing and first post-grant transfer.

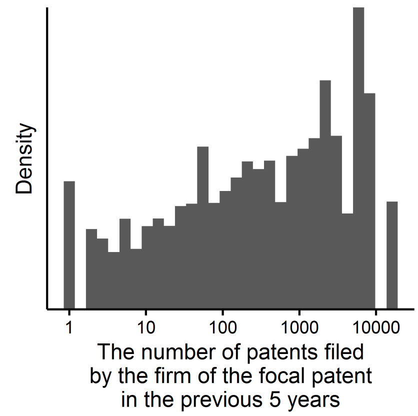

Figure B2: Density of firm size on the patent level.

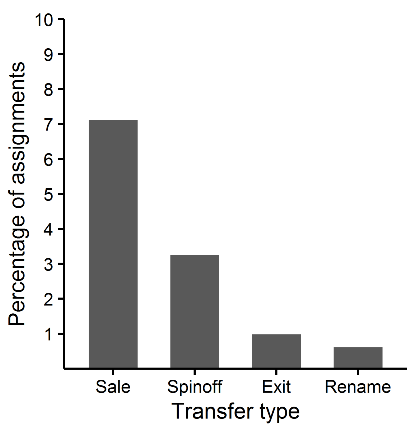

Figure B3: Patent transfers by type as a percentage of all transfers, with initial transfers from inventors to firms omitted.

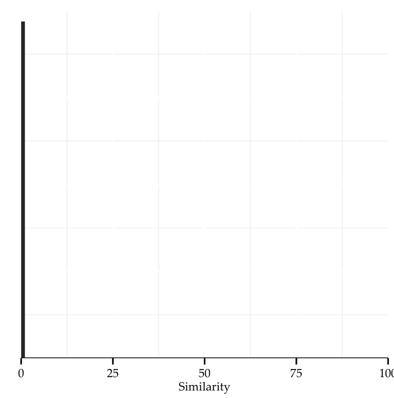

Figure B4: Density plot of raw similarity values, from Younge and Kuhn (2015).

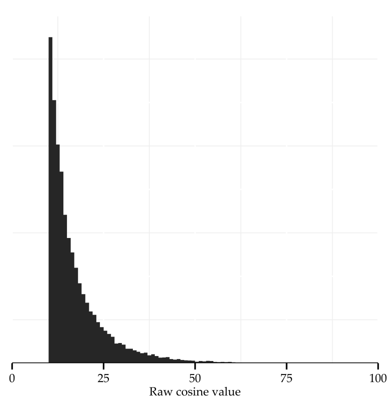

Figure B5: Density plot of transformed raw similarity values above 10, from Younge and Kuhn (2015).

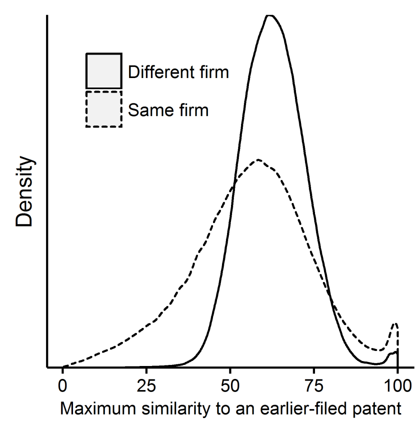

Figure B6: Patent-level density of plot of the maximum similarity value to a previously-filed patent.