

Trading and Enforcing Patent Rights¹

Alberto Galasso
University of Toronto

Mark Schankerman
London School of Economics and CEPR

Carlos J. Serrano
University of Toronto and NBER

August 24, 2012

¹We would like to thank two anonymous referees and the Editor for very constructive comments on an earlier version of the paper. We are also grateful to Victor Aguirregabiria, Ashish Arora, Pierre Azoulay, Christian Catalini, Dietmar Harhoff, Nico Lacetera, Josh Lerner, Megan MacGarvie, Matt Mitchell, Jeff Thurk, Dan Treffer and Heidi Williams for comments and suggestions on earlier drafts, and to Grid Thoma for helping with the matching algorithm. We also thank seminar and conference participants at Berkeley, Duke, Georgia Tech, Kellogg, Carlos III University Madrid, Max Planck Institute, Pompeu Fabra, SUNY Stony Brook, Toulouse, Toronto and the ZEW. Jelena Bozovic and Christina Kim provided excellent research assistance. We are grateful for financial support from the Centre for Economic Performance at the London School of Economics and the Social Sciences and Humanities Research Council of Canada.

Abstract

We study how the market for innovation affects enforcement of patent rights. Conventional wisdom associates the gains from trade with comparative advantage in manufacturing or marketing. We show that these gains imply that patent transactions should increase litigation risk. We identify a new source of gains from trade, comparative advantage in patent enforcement, and show that transactions driven by this motive should reduce litigation. Using data on trade and litigation of individually-owned patents in the U.S., we exploit variation in capital gains tax rates as an instrument to identify the causal effect of trade on litigation. We find that taxes strongly affect patent transactions, and that reallocation of patent rights reduces litigation risk on average, but the impact is heterogeneous. We show that patents with larger potential gains from trade are more likely to change ownership, suggesting that the market for innovation is efficient, and the impact of trade on litigation depends on characteristics of the transactions.

Keywords: patents, litigation, market for innovation, capital gains taxation.

JEL Codes: K41, H24, O32, O34.

1 Introduction

The ‘market for innovation’ – the licensing and sale of patents – is an important source of R&D incentives, especially for small firms and individual inventors for whom patents are often their critical asset. Transactions in patent rights are also important for developing efficient market structures in high-technology sectors. They do this by shaping the division of labor, and the nature of competition, between small firms (or individuals) who typically specialize in innovation but lack the capacity for large scale development, production and marketing, and large firms whose comparative advantage lies in the commercialisation of these inventions (Gans and Stern, 2000; Gans, Hsu and Stern, 2002). The key to realizing these social gains is efficient technology transfer.

Despite these private and social benefits, there is growing concern voiced in both academic and policy debates about the potentially deleterious effects of patent transactions. The modern innovation landscape is characterized by a large number of patents, with often fuzzy boundaries and fragmented ownership (Bessen and Meurer, 2008). The main concern is that, in this environment, patent transactions can deter innovation if they take place in order to extract rents through patent litigation, rather than to facilitate welfare enhancing technology transfers. This issue is at the center of a recent report by the U.S. Federal Trade Commission (2011), and the Supreme Court has raised similar concerns in a recent, prominent case (*MerckExchange, L.L.C. v. eBay, Inc.*, 126 S. Ct. 1837, 2006).¹ However, there is sharp disagreement among economic and legal scholars about the scope and severity of this problem. For example, Mann (2005) claims that the detrimental effects from patent transactions are minimal, while Lemley and Shapiro (2007), among others, argue that patent transactions constitute a serious threat of ex post hold-up for manufacturing firms, discouraging investment and innovation, and requiring policy intervention in the form of limits to patent enforcement for non-practicing entities.² Some have even gone so far as to recommend more draconian reductions in permissible patenting; especially in relation to software patents.³

Despite the importance of these issues, there are no empirical studies of the impact of

¹The FTC report highlights the risk associated with the activity of patent assertion entities (sometimes called patent trolls), which it defines as firms that obtain nearly all of their patents through acquisitions in order to assert them against manufacturing companies.

²The policy recommendation by Lemley and Shapiro is to limit the ability of non-practicing entities to obtain preliminary injunctions – specifically, to allow them only when the patent holder can claim actual lost profits (which requires that the patent holder to actually be working the patent), but not when only reasonable royalties are claimed. Other policy proposals are currently being examined by the federal government. For example, the U.S. House of Representatives is discussing the Saving High-Tech Innovators from Egregious Legal Disputes (SHIELD) Act (HR 6245), whose objective is to deter patent litigation by patent assertion entities.

³See Hall and McGarvie (2010) for an overview of the software patents debate.

the market for patents on patent litigation. Indeed, this lack of empirical evidence led the U.S. House Judiciary Committee, in April 2011, to amend the Patent Reform bill (H.R. 1249, “The America Invents Act”) to require the Comptroller General of the United States to study the impact of patent transactions and litigation on innovation.

In this paper we take a first step in this direction by studying how the market for patents affects the enforcement of patent rights. The economics and management literature typically associates the gains from trade in patent transactions with vertical specialization (Teece, 1986; Arora, Fosfuri and Gambardella, 2001) and comparative advantages in manufacturing or marketing (Arora and Ceccagnoli, 2006). By raising the potential profit from the innovations, these mechanisms imply that market reallocation of patent rights should increase the likelihood of litigation. In this paper we identify a novel source of private, and social, gains from trade – comparative advantage in patent enforcement. The market for innovation can reduce litigation if it reallocates patents to entities that are more effective at resolving disputes over these rights without resorting to the courts.⁴ A third, more controversial motivation for patent transactions is patent trolling – acquisition of patent rights for later use against existing manufacturing firms. If this is the driving force behind patent transactions, we would also expect to observe that a change of ownership raises the likelihood of litigation on the traded patent.

The main focus of this paper is to identify empirically the causal effect of trade on litigation, and to assess the relative importance of commercialization and enforcement gains from trading patent rights (we briefly explore the patent troll issue later in the paper). To do so, we construct a new, comprehensive data set that matches information on trades (Serrano, 2010) and litigation (Lanjouw and Schankerman 2001, 2004) involving patents owned by individual inventors in the United States during the period 1983-2000.

The empirical challenge in studying how reallocation of patent rights affects litigation is the endogeneity of patent trading. To address this concern, we exploit a provision in the U.S. tax law that allows us to use variation in capital gains tax rates across states and over time as an instrument to identify the causal effect of a change in patent ownership on litigation. Under U.S. law, for an individual patent-holder the profits from the sale of a patent are taxed as capital gains, whereas any damage awards from litigation are taxed as ordinary income. This means that capital gain tax rates affect the incentives to sell patents for individual owners, but not their incentives to undertake patent litigation, and are thus a suitable instrument for change in ownership in the patent litigation regression. This identification strategy means that we can only study patents that are originally owned by individual inventors in this paper.

⁴This may involve acquisition by firms to accumulate defensive patent portfolios for resolving dispute non-litigiously (Hall and Ziedonis, 2001) or some form of economies of scale in enforcement (Lanjouw and Schankerman, 2004). Defensive patenting is particularly prevalent in high technology sectors where there is widespread fragmentation of patent rights over important inputs used in the R&D and production processes.

The main empirical findings in the paper are as follows. First, we show that capital gains taxation strongly affects the decision to trade patent rights for individual inventors. This finding is consistent with recent literature on how taxation affects the frequency and timing of the sale of small businesses (Chari, Golosov, and Tsyvinski, 2005; Gentry, 2010). We conduct simulations using our parameter estimates that show that changes in capital gains taxation can have large effects on the frequency of patent transactions and litigation.

Second, we find that changes in patent ownership reduce the likelihood of litigation for patents originally owned by individual inventors, *on average*. This implies that enforcement gains dominate commercialization gains (and the effects of any patent trolling activity) in the market for such patents. This finding is consistent with our hypothesis that patent transactions exploit differences across firms in their ability to enforce these rights. However, the marginal treatment effect of an ownership change is highly heterogeneous and depends on the characteristics of the patent and the transacting parties. We also show that the market for innovation appears to be efficient in the sense that patents are more likely to be traded when the estimated private enforcement gains from doing so are larger.

Third, we unbundle the heterogeneous treatment effect of patent transactions on litigation by exploring how specific characteristics of the transaction influence this treatment. We show that the impact of trade on litigation depends on the size of the buyer's patent portfolio and the technological fit of the traded patent in that portfolio. Sales by individual inventors to other individuals or small firms are not associated with a decline in the (post-trade) probability of litigation. By contrast, sales to firms with larger patent portfolios significantly reduce litigation risk. This is consistent with the economies of scale in enforcement first documented by Lanjouw and Schankerman (2004).⁵ In addition we find that, holding the buyer's portfolio size constant, reallocation of patents increases litigation risk more when the traded patent is a better technological fit in the buyer's existing portfolio. This is what we expect since the potential commercialisation gains from the transfer are likely to be larger in such cases.

Finally, we examine whether this increase in litigation risk is due to patent assertion entities – firms that typically gather patents through acquisitions in order to assert them against manufacturing companies (often referred to as patent trolls). We do not find any evidence that patent trolls play a substantial role in our sample of transactions involving individually-owned patents during the period 1983-2000. Whether this conclusion would apply to corporate patent transactions, or the post-2000 period, is left for future research.

Taken together, our empirical findings indicate that a well-functioning market for innova-

⁵Lanjouw and Schankerman (2001, 2004) show that the litigation risk is systematically related to characteristics of the patent (including measures of value and the technology field) and of the patentholder. In particular, they find economies of scale in patent enforcement – firms with larger patent portfolios are more able to resolve disputes without resorting to the courts.

tion is important for allocating patent rights efficiently, and that taxation strongly affects this process. Moreover, as long as small innovators can appropriate part of the commercialization and enforcement gains generated by these transfers, this market increases their incentives to innovate.⁶

The paper is organized as follows. In Section 2 we present a simple model that highlights the commercialisation and enforcement gains from trade, the impact of trade on litigation, and the role of taxation. Section 3 describes the data. In Section 4 we develop the baseline econometric model for estimating the causal effect of trade on litigation, and present the results. In Section 5 we allow for heterogenous marginal treatment effects, and empirically link them to characteristics of the trade. In Section 6 we quantify the impact of taxes on patent trade and litigation by simulating changes in individual tax rates. Section 7 provides a brief discussion of the welfare implications of our findings. Brief concluding remarks close the paper.

2 An Illustrative Model of Patent Trade and Litigation

Consider an individual, A , owning a patent and a firm, B , willing to acquire the patent from the individual.⁷ If the individual does not sell the patent, he obtains product market profits from commercializing (licensing or using) the innovation equal to π^A . If the patent is acquired by the firm, it generates product market profits equal to π^B . For simplicity, we assume that the individual has all the bargaining power and extracts the entire surplus from the transaction (results are similar if there is Nash bargaining).

Both A and B face an infringing action by a third party, firm C with probability β . If the infringing action takes place, the patent owner chooses whether to litigate or settle the dispute. With litigation the patent owner $i = \{A, B\}$ sustains litigation costs l_i to secure product market profits. To settle the dispute, the owner gives up a fraction $(1 - \theta_i)$ of the profits to firm C . We also assume that there is a zero mean, random (monetary) component in the settlement payoff, ε . In this setup, there will be litigation if

$$\pi^i - l_i \geq \theta_i \pi^i + \varepsilon \tag{1}$$

which occurs with probability

$$\Pr \{ \varepsilon \leq \pi^i (1 - \theta_i) - l_i \}.$$

⁶Our paper is also connected to the growing literature on the interplay between innovation and the transactions across firm boundaries. For example, Azoulay (2004) studies how the nature of knowledge affects outsourcing, Cockburn, MacGarvie and Mueller (2010) examine the impact of the intellectual property (IP) landscape on licensing and Williams (2011) studies how IP affects cumulative innovation.

⁷In this paper we do not model the microfoundations of the search process through which matching occurs. The role of the model is simply to illustrate the two different sources of gains from trade, their impact on litigation and their interplay with income and capital gains taxes.

We refer to the vector $e^i = (l_i, \theta_i)$ as the enforcement vector of owner $i = \{A, B\}$. Litigation takes place with probability $\Omega(\pi^A, e^A) = \beta \Pr\{\varepsilon \leq \pi^A(1 - \theta_A) - l_A\}$ if the patent is owned by the individual and with probability $\Omega(\pi^B, e^B) = \beta \Pr\{\varepsilon \leq \pi^B(1 - \theta_B) - l_B\}$ if the patent is owned by the firm. Notice that $\partial\Omega(\pi^i, e^i)/\partial\pi^i > 0$ whereas $\partial\Omega(\pi^i, e^i)/\partial l_i < 0$ and $\partial\Omega(\pi^i, e^i)/\partial\theta_i < 0$.

To start, we consider the case in which there are no taxes. If the individual does not trade the patent, expected profits are

$$\begin{aligned} & (1 - \beta)\pi^A + \Omega(\pi^A, e^A)(\pi^A - l_A) + (\beta - \Omega(\pi^A, e^A))\theta_A\pi^A \\ &= (1 - \Delta_A)\pi^A - \Omega(\pi^A, e^A)l_A \end{aligned}$$

where the term $\Delta_A = (\beta - \Omega(\pi^A, e^A))(1 - \theta_A)$ captures the expected fraction of profits lost because of settlement between A and C, and $\Omega(\pi^A, e^A)l_A$ captures the expected litigation costs. Similarly, if the patent is owned by firm B, profits are $[(1 - \Delta_B)\pi^B - \Omega(\pi^B, e^B)l_B]$ where $\Delta_B = (\beta - \Omega(\pi^B, e^B))(1 - \theta_B)$.

The individual will sell the patent if

$$[(1 - \Delta_B)\pi^B - \Omega(\pi^B, e^B)l_B] \geq [(1 - \Delta_A)\pi^A - \Omega(\pi^A, e^A)l_A]$$

which can be re-written as

$$(\pi^B - \pi^A) + (\Delta_A\pi^A - \Delta_B\pi^B) + (\Omega(\pi^A, e^A)l_A - \Omega(\pi^B, e^B)l_B) \geq 0. \quad (2)$$

Condition (2) highlights three possible sources of gains from trade. The first term captures product market gains, i.e. the greater profits that firm B obtains from selling the product. The second and third terms capture the enforcement gains which take the form of losing less profit from settlement, $\Delta_A\pi^A - \Delta_B\pi^B$, and incurring lower expected litigation costs, $\Omega(\pi^A, e^A)l_A - \Omega(\pi^B, e^B)l_B$.

It is straightforward to introduce taxes into the analysis. If the individual owner commercializes the patent, the profits are taxed at the personal income tax rate τ^I . If the patent is traded to the firm, the product market profits are taxed at the corporate income tax rate τ^C . If the individual owner sells the patent, the gains from the transaction are taxed at the capital gains tax rate τ^G . This setup conforms to the U.S. tax code (see Section 4 for more details). With taxes, we get the following conditions for the decision to litigate and to trade the patent, respectively:

$$(\pi^i - l_i)(1 - \tau^i) \geq (\theta_i\pi^i + \varepsilon)(1 - \tau^i) \quad (3)$$

$$\begin{aligned} & [(1 - \Delta_B)\pi^B - \Omega(\pi^B, e^B)l_B](1 - \tau^C)(1 - \tau^G) \geq \\ & [(1 - \Delta_A)\pi^A - \Omega(\pi^A, e^A)l_A](1 - \tau^I) \end{aligned} \quad (4)$$

where $\tau^i = \tau^I$ if $i = A$ and $\tau^i = \tau^C$ if $i = B$.

Note that the capital gains tax rate does not enter the first inequality that governs the litigation decision. The second inequality, however, shows that the condition required to have trade becomes more stringent with an increase in τ^G , an increase in τ^C or a decrease in τ^I . Higher capital gains and corporate taxes reduce the likelihood that patent rights are reallocated, and higher (personal) income tax rates increase it. We test these predictions in the empirical analysis, and exploit the capital gains tax rate as an instrument for trade based on it being excluded from the condition for litigation.⁸

To see how litigation is affected by a change in patent ownership, let *NewOwner* be an indicator variable equal to one if the patent changes ownership and zero otherwise. If individual *A* does not sell the patent, the probability of litigation is $\Pr(\textit{Litigation}|\textit{NewOwner} = 0) = \Omega(\pi^A, e^A)$. If trade takes place, the probability is $\Pr(\textit{Litigation}|\textit{NewOwner} = 1) = \Omega(\pi^B, e^B)$. Thus the impact of trade on litigation is

$$\Pr(\textit{Litigation}|\textit{NewOwner} = 1) - \Pr(\textit{Litigation}|\textit{NewOwner} = 0) = \quad (5)$$

$$- [\Omega(\pi^A, e^A) - \Omega(\pi^B, e^B)].$$

This equation shows that the effect of trade on litigation depends on whether it reallocates the patent to an entity with greater product market gains and/or lower enforcement costs. The effect of trade can be either positive or negative, depending on the difference $\Omega(\pi^A, e^A) - \Omega(\pi^B, e^B)$.

Previous literature associates the surplus generated by patent trades with gains from vertical specialization or comparative advantages in manufacturing or marketing. In our model, this commercialisation hypothesis corresponds to the case where $\pi^A < \pi^B$ and $e^A = e^B = e$. Because $\partial\Omega(\pi^i, e^i)/\partial\pi^i > 0$, in this case the change in patent ownership is unambiguously associated with an increase in patent litigation, since $\Omega(\pi^B, e) - \Omega(\pi^A, e) > 0$. Intuitively, in this scenario trade increases the product market profits generated by the patent but does not alter the enforcement capability of the owner. Because an increase in patent value increases the likelihood of patent litigation (Galasso and Schankerman, 2010), trade increases litigation rates if it is only motivated by product market gains.

By contrast, if the difference $\theta_B - \theta_A$ is positive and large enough to guarantee that $\Omega(\pi^A, e^A) > \Omega(\pi^B, e^B)$ trade is associated with a reduction in the level of patent litigation.

⁸The model assumes that the fee the company pays for the patent is not tax deductible. If we assume that a fraction g of the fee is deductible, the optimal fee becomes

$$\left[(1 - \Delta_B) \pi^B - \Omega(\pi^B, e^B) l_B \right] (1 - \tau^C) / (1 - g\tau^C)$$

which depends (negatively) on corporate taxes as long as $g < 1$. Incomplete deductibility is a plausible assumption because, under the current tax code, the cost of acquiring intellectual property must be capitalized (I.R.C. § 263) and are also subject to a variety of tax depreciation rules (Maine and Nguyen, 2010).

In the patent context there are two main reasons why patentees may vary in their likelihood to enforce the patent without filing a suit. First, patent owners may have different ability to exchange intellectual property through licensing or cross-licensing agreements (Hall and Ziedonis, 2001). Second, not all owners may be able to generate an expectation of repeated interaction large enough to sustain cooperation over time. Lanjouw and Schankerman (2004) provide evidence in support of these two mechanisms, showing that firms with large patent portfolios are less likely to file a suit on any individual patent in their portfolio (controlling for patent characteristics).

Discussion of Modeling Assumptions

There are two assumptions in the model that warrant additional discussion. First, we assumed that infringement does not occur between the seller and the buyer. However, it is possible that patent trades occur as the outcome of patent infringement or invalidity disputes. To accommodate this, in Appendix 1 we develop an extension to our model that includes the possibility of infringement between buyer and seller, and we show that our identification strategy still holds. The intuition behind this is that, also in this extended model, higher capital gains and corporate taxes reduce the payoff of a patent transaction and thus the likelihood that patent rights are reallocated, while higher (personal) income tax rates increase it. This extension of the model, however, introduces an additional mechanism by which the reallocation of patent ownership reduces litigation. There are now two distinct mechanisms: the first is the differential ability of the buyer and the seller to settle disputes with third parties (this was the original channel); the second, new channel is avoidance of litigation between the buyer and the seller involved in the patent dispute.

Second, we assumed that the probability of facing an infringing action by a third party, β , is exogenous and does not depend on the characteristics of the patent owner. Allowing for different values β_A and β_B does not affect the main predictions of the model. In particular, we can simply redefine $\Omega(\pi^i, e^i) = \beta_i \Pr\{\varepsilon \leq \pi^i(1 - \theta_i) - l_i\}$ and $\Delta_i = (\beta_i - \Omega(\pi^i, e^i))(1 - \theta_i)$ for $i \in \{A, B\}$, and then there is no change in equations (2), (3) and (4) above. Intuitively, even if the probability of filing a suit depends on a combination of ex ante, owner-specific characteristics that affect the likelihood of infringement action, β_i , and an ex post random shock that affects the likelihood of litigation, $\Pr\{\varepsilon \leq \pi^i(1 - \theta_i) - l_i\}$, we can still distinguish between gains from trade arising from commercialization and enforcement – i.e. , we still get equation (2). The only difference is that these enforcement gains now consist of an ex ante and an ex post component. Moreover, our result that the capital gains tax rate does not affect the individual’s litigation decision directly also holds in this generalized setup. In addition, as before, the decision to trade is more likely when capital gains and corporate tax rates are low,

and when income tax rates are high (i.e. there is no change in equations (3) and (4)).

In this generalized setup, equation (5) can now be re-written as:

$$\Pr(\textit{Litigation}|\textit{NewOwner} = 1) - \Pr(\textit{Litigation}|\textit{NewOwner} = 0) = -\beta_A (z^A - z^B) + (\beta_B - \beta_A) z^B$$

where $z^i = \Pr\{\varepsilon \leq \pi^i(1 - \theta_i) - l_i\}$. This decomposition confirms the result that the effect of trade on litigation can be positive or negative. As in our baseline model, the difference $z^A - z^B$ is positive when ex post enforcement gains dominate product market gains. It also shows that two types of enforcement gains reduce the level of litigation: 1) litigation declines because of a greater ability to settle a dispute (the ex post effect, shown in the first term in the equation when $z^A > z^B$) and 2) litigation declines because of a lower likelihood of infringing action (the ex ante effect, given by the second term in the equation above when $\beta_B < \beta_A$).⁹

3 Description of the Data and Motivating Evidence

Our starting point is the panel of patents granted in the period 1975-2000 that are either owned by the original inventor at the grant date or have been assigned to U.S. individuals by the grant date. Hall, Jaffe and Trajtenberg (2001) refer to the first group of patents as ‘Unassigned’ and to the second group of patents as ‘U.S. Individuals’ patents. The USPTO refers to both groups as ‘Individually Owned’ patents. For each of these patents we obtained information on the U.S. state of the primary (first listed) inventor, their reassignment and litigation history. We also collected information on the U.S. state and federal ordinary income taxes, capital gain taxes and corporate taxes during the sample period.

We now describe the main components of our data set.

Patent trade data: We follow Serrano (2010) and use re-assignment data to identify transfers of patents across owners. The source of these data is the USPTO Patent Assignment Database. When a U.S. patent is transferred, an assignment is recorded at the USPTO acknowledging the change in ownership. A typical re-assignment entry indicates the patent involved, the name of the buyer (assignee), the name of the seller (assignor), the date at which the re-assignment was recorded at the patent office, and the date at which the private agreement between the parties was signed. The data set covers the period 1983-2001.

⁹Endogenising the parameter β is beyond the scope of this paper. Nonetheless, notice that in our baseline model the infringer gets at most $\pi^i(1 - \theta_i) - l_i$, which is decreasing with the patent owner’s ability to settle disputes without litigation denoted by θ_i . Moreover, patent transactions characterized by large enforcement gains will reallocate the patent to owners with large θ_i . Thus, in a generalized version of our model, we would expect to find that ex ante infringement β_i is less likely when ex post enforcement gains are larger.

Under Section 261 of the U.S. Patent Act, recording the assignment protects the patent owner against previous unrecorded interests and subsequent assignments. If the patentee does not record the assignment, subsequent recorded assignments will take priority. For these reasons, patent owners have strong incentives to record assignments and patent attorneys strongly recommend this practice (Dykeman and Kopko, 2004).

A challenge in using re-assignment data is to distinguish changes in patents ownership from other events recorded in the USPTO assignment data. To this end, we use an algorithm developed in Serrano (2010) that conservatively drops all the assignments that appear not to be associated with an actual patent trade. Specifically, we drop assignments in which the buyer is the assignee at the grant date of the patent, and assignments recorded at the patent application date. We also dropped transfers to financial institutions to eliminate transactions (recorded in the USPTO Patent Assignment Database) in which a patent is used as collateral.¹⁰ Another concern is that the first assignment of an unassigned patent may not correspond to a trade but rather to the transfer of ownership from the inventor to the company in which the inventor works. To deal with this, we drop any transactions where there is evidence that the seller is an inventor working for the buyer.¹¹

Litigation data: The patent litigation data set was compiled by Lanjouw and Schankerman (2001, 2004). This data set matches litigated patents identified from the Lit-Alert database with information on the progress or resolution of suits from the court database organized by the Federal Judicial Center. The data set contains 14,169 patent cases filed during the period 1975-2000. For each of these case filings, the data set reports detailed information on the main patent litigated, the patentee, the infringer and the court dealing with the case. The data set contains information on patent cases filed in U.S. federal district courts (and not on appeal). For each patent in our data, we identify the suits in which the patent was involved and the year in which the case was filed.¹²

¹⁰We also dropped records in which the buyer and seller are the same entity and in which the execution date is either before the application date or after patent expiration. For additional details on the procedure, see Serrano (2010).

¹¹Specifically, for each transfer between a seller i and buyer j , we identified all the patents which list the seller i as the (primary) inventor and checked whether any of these patents was assigned to the buyer j at its grant date. We drop all such transactions.

¹²The use of re-assignment data as a proxy for activity in the market for innovation can be problematic because technology can be transferred through patent licensing without changes in ownership. This concern is less relevant in our study that focuses on patent litigation because typically it is the owner of the patent that brings patent infringement actions. Non-exclusive licensees do not have the right to sue for patent infringement (Textile Prods. v. Mead Corp., 134 F.3d 1481, 1485 Fed. Cir. 1998). An exclusive licensee may have standing to bring such a suit but only under some restrictive contract arrangements (Resonant Sensors Inc. v SRU Biosystems, Inc., No. 3:08-cv-1978-M).

Tax data: Information on state and federal income and capital gain taxes are obtained from the NBER Tax Rates data base. This contains marginal income tax rates by year and state for a representative household with \$500,000 of wage income.¹³ The data set also reports maximum federal and state long-term capital gains tax rate by year and state, computed using the NBER TAXIM model. We obtain information on the maximum federal and state corporate marginal tax rates, for each year and state, from two government publications: the Significant Features of Fiscal Federalism (available for the period 1982-1995) (ACIR, 1982-1995) and the Book of the States (for the period 1996-2000) (CSG, 1996-2000). For each assigned patent in our data set, we use the ordinary income and capital gains marginal tax rates in the state of the initial patent assignee. For unassigned patents, we used the state of the primary inventor as identified by the USPTO. To measure tax rates faced by potential corporate buyers, we construct a weighted average of state corporate taxes where state weights are determined by the fraction of state patent applications in the technology class of the patent.¹⁴

Matching data on income and capital gain taxes to patents is meaningful as long as the patent is owned by an individual at the time of the transaction. To ensure this, we focus our analysis on the first transfer of a patent. Subsequent owners are generally not individuals and thus are not subject to either personal income or capital gains taxation on the patent transaction. Focusing on the first transfer involves dropping very few patent trades. Most of the traded patents in our data are traded only once (94.9 percent) and only 0.15 percent of traded patents are traded more than three times.

The final data set is a panel with 299,356 patents and 2,436,649 patent-age observations. The main variables used in the empirical analysis are described below.

Litigation Dummy: dummy variable equal to 1 if at least one suit is filed in a federal court involving the patent in a given year

NewOwner: dummy variable equal to 1 for patent-ages in which the patent is no longer owned by the original individual assignee/inventor

Income Tax Rate: for each patent-age, the sum of the federal income tax rate and the state income tax rate for the state of the primary (first listed) inventor of the patent

Capital Gains Tax Rate: for each patent-age, the sum of the federal capital gain tax rate and the state capital gain tax rate for the state of the primary (first listed) inventor of the patent

¹³For details, see the description of the TAXSIM program in Feenberg and Coutts (1993). The simulation and the resulting data are available at <http://www.nber.org/~taxsim/state-rates/>

¹⁴All our results are robust to dropping corporate tax rates or to using corporate tax rates in the state of the inventor, which assumes that trading of patents occurs only within states.

Corporate Tax Rate: for each patent-age, the sum of the federal corporate tax rate and a weighted average of the state corporate tax rates. State weights are equal to the fractions of state patent applications in the technology class (USPTO nclass) of the patent in that calendar year.

In principle, exploiting the information contained in the USPTO assignment data it is possible to recover the patenting activity of the buyers in our sample. Unfortunately, the names of the buyer and seller in the Patent Assignment Database were never standardized by the USPTO. Therefore, to back out buyer patent portfolios we need to match each buyer name manually with a unique assignee identifier required to identify the buyer’s patents. Because of the large size of our sample (17,605 traded patents), we manually matched only patents that were both traded and litigated at least once in their lifetime (569 patents). In the empirical analysis below, we will focus on regression results for the entire data set (299,356 patents), but also show that the findings also hold for the smaller data set of traded and litigated patents, where we are able to investigate the role of buyer characteristics on the impact of trade.¹⁵

Table 1 reports summary statistics for the key variables. Panel A shows the fraction of sample patents involved in trade or litigation at least once in their life. Of the total sample, 4.55 percent of patents are traded and 0.69 percent are involved in at least one suit. These rates are low but it is worth noting that, for the later patents in the sample, data on trade and litigation are truncated and this biases downward litigation and trade rates.¹⁶ Moreover, patents that are traded or litigated are much more valuable than the those that are not (as measured by citations received).¹⁷ The striking fact from this table is the strong association between trading and litigation. Of patents that are traded, 4.2 percent are also litigated; for patents that are not traded, the litigation rate is only 0.51 percent. Of patents that are litigated, 27.9 percent are also traded; for patents that are not litigated, only 4.4 percent are traded.

The second panel of Table 1 illustrates the combined (state plus federal) individual and corporate tax rates averaged across states for four five-year time periods. There is a substantial decline in income tax rates in the late eighties and an increase in the early nineties. Conversely, there is an increase in capital gain tax rates in the late eighties and a decrease in the late nineties. The summary statistics show the range of variation across U.S. states. The difference between

¹⁵In Section 5 we also extend our analysis of the impact of buyer characteristics by using a much larger data set that includes corporate buyers identified by using a disambiguation algorithm developed by Thoma et. al. (2010).

¹⁶For patents where we have litigation and trade data during the first ten years of life (i.e. patents granted in 1983-1991), we find that 11.8 percent are traded and 2.2 percent are litigated.

¹⁷The mean number of citations for patents that are neither traded nor litigated is 6.1. The mean is 10.8 for traded patents and 16.5 for litigated patents. For those that are both traded and litigated, the average is 19.3.

the lowest and the highest capital gains tax rates across states ranges from 7-9 percentage points (depending on the year). The difference between the minimum and the maximum income tax rate across states is 6-16 percentage points. Corporate tax rates decline during the sample period and the difference between the lowest and highest rates is 12-15 percentage points.¹⁸ Analysis of variance shows that 89.4 percent of the overall variance in capital gains tax rates is variation over time and 8.7 percent is variation across states (the small remainder is residual). The breakdown for ordinary income tax rates is 92.9 and 6.8 percent; for corporate tax rates, 49.1 and 48.6 percent.

In Appendix Table A1 we provide a more detailed breakdown of the variation in capital gain tax rates. For the period 1982-2001, our data show 268 changes in capital gains tax rates at the state level. The marginal tax rate increases in 138 of these cases and decreases in 130 instances. The average increase in capital gains tax rates is 1.5 percentage points, equivalent to 54 percent of the rate in the year immediately before the tax change. The average decline in the capital gains tax rate is 0.6 percentage points, representing about a 9.2 percent tax cut. The table confirms that there is substantial variation in the rates across time and states. Only nine states (Florida, Texas, Washington, Tennessee, Nevada, New Hampshire, Wyoming, Alaska, and South Dakota) do not experience any change in state level capital gain tax rates during our sample period.

Individually owned patents represent 17.9 percent of the patents granted in the period 1975-2000 (about 19 percent if weighted by citations received). If we exclude patents granted to foreign entities and government agencies, individually owned patents account for about 22 percent of the remaining sample. In Table 2 we compare sample means of the number of citations received for individually owned and corporate patents, granted in the 15 year window 1985-2000 for which we also obtained data on litigation and reassignment of corporate patents. We distinguish between small corporate innovators (defined following Serrano, 2010) as entities applying for fewer than 5 patents in a calendar year) and other corporate innovators. On average, individually owned patents receive fewer citations than corporate patents. Nonetheless, if we focus on traded patents, we see only very minor differences in citations across the three ownership types (this is particularly important because traded patents are key for the identification of the effect of trade in our fixed effects regressions).¹⁹

Table 2 also examines the differences in the likelihood of trade and litigation. The

¹⁸Similar figures are observed if we restrict the analysis to the 20 states with the most individually owned patents.

¹⁹Because these samples are very large, the differences between these means are statistically significant. However, these differences are small when measured in terms of the percentile of the distributions of citations. Specifically, the sample means of citations for traded patents owned by individuals (10.34), small corporate (10.87) and other firms (9.28) all lie between the 77th and the 80th percentiles of the distribution.

fraction of corporate patents that are traded is three times as large as those of individually owned patents. However, there is essentially no difference between trade rates of individuals and small firms once we weight patents by citations (this is consistent with the evidence in Serrano, 2010 that show greater evidence of selection into trade for individuals and small firms). For each of the ownership types we also constructed the annual litigation rates. The litigation rates for individuals and small innovators are quite similar. At the same time, consistent with the findings in Lanjouw and Schankerman (2004), there is a substantially lower litigation rate for larger corporate innovators. Overall, these figures indicate that individually owned patents are not sharply different from corporate owned patents, especially those of small patenting firms.

Motivating Evidence

Panel A in Table 1 shows that trade and litigation are associated, but it does not reveal how litigation rates differ before and after a trade occurs. To show this, we focus on patents that are eventually traded (in our sample period). In Figure 1 we compare the probability of being involved in at least one suit prior to and after the date at which trade occurs. In aggregate, a patent that has not been traded but that will be traded in its lifetime is involved in at least one suit in that year with probability 0.61 percent. A patent that has already changed ownership is involved in at least one suit with probability 0.48 percent. The post-trade litigation probability is lower after trade even after we condition on age. For example, a patent that has not yet been traded at age 7 is involved in at least one dispute with probability 0.76 percent, whereas for a patent of the same age that has been already traded, the litigation rate is about that level (0.43 percent). In short, Figure 1 suggests that the reallocation of patent rights is temporally related to lower litigation risk. In the econometric analysis we exploit capital gains tax rates as an instrument to pin down the causal relationship.

The large number of changes in state level capital gains and corporate tax rates that are present in our data provide potentially rich variation for identification of the tax effect on trade. Specifically, the mean number of years between two tax changes in a state is about two years, with 50 percent of the tax changes in a state followed by another tax change in the next year. Similarly, the average number of years between a tax hike and tax cut is about three years (on average tax cuts are observed 3.7 years after a tax hike and tax hikes are registered 2.7 years after a tax cut).

To provide preliminary evidence of the effect of taxes on patent trade, we focus on eight illustrative tax events (4 tax cuts and 4 tax hikes). Specifically, we searched for tax changes that were both sizeable (i.e. tax cuts higher than 1 percentage point and tax hikes larger than 2 percentage points, which approximately correspond to the top quartile of the distributions) and that were not confounded by other capital gains and corporate tax changes within a six

year window.²⁰

We exploit these events to compare the changes in trade rates between individually owned patents and corporate owned patents, which are used as the control group (since trade for them is unaffected by capital gains taxes), before and after a tax change. We do this separately for tax cuts and tax hikes. This leads to the familiar difference-in-differences estimator. In these regressions we control for additive fixed effects for patent age, years, technology sub-categories and states, as well as for the level of corporate and income taxes (sum of state and federal rates). Unreported regressions show that the trading of individually owned patents, relative to the corporate patent control group, decreases when there is a capital gains tax hike, and increases with a tax cut, and these effects are strongly statistically significant. The estimated treatment effects are 0.009 (with std. error equal to 0.002) for a capital gains tax cut and -0.010 (with standard error equal to 0.002) for a tax hike. These effects are large, corresponding to about 80 percent of the mean probability of trading for individually owned patents.

Figure 2 depicts the point estimates from a more general empirical specification that allow the treatment effect to vary for each year before and after the tax event. We normalize the coefficient to zero for the year preceding the tax event, as is common practice. We also depict the estimated 95 percent confidence bands. In the first panel, which examines the effects of cuts in the state capital gains tax rate, we see no statistically significant “treatment effects” in the years prior to the date of the tax changes. The estimated coefficients are not individually statistically different from zero (even though the point estimates rise somewhat), and we also do not reject the hypothesis that they are jointly equal to zero (p-value=0.19). This result indicates that there is a common trend for individual and corporate patents prior to the treatment. However, the estimated treatment effects for the two years immediately after the tax cuts are positive and statistically significant – there is a sharp increase in the relative trading rate for individually owned patents. The second panel presents the effects of the tax hikes. While the point estimates of the “treatment effect” decline before the tax increase, these are not statistically different from zero separately or jointly (p-value=0.25 for the joint test), but after the tax hike there is a statistically significant decline in the trading rate for

²⁰In our sample, 52 out of 269 events involve a sufficiently large tax change, with 19 tax cuts and 33 tax hikes. Among these tax events, we can construct a six-year window only for events in the period 1986-1998, which leads us to 46 tax changes out of 52 (14 tax cuts and 32 tax hikes). Moreover, we focus on capital gains tax events with no contemporaneous changes of the state-level corporate tax rate in the six year window. This reduces the number of tax events to 26 (12 tax cuts and 14 tax hikes). Finally, we consider tax events that involve a tax cut with no confounded tax hike in the six year window, and similarly for tax hikes. This reduces the sample to eight “clean” tax events (4 tax cuts and 4 tax hikes). These tax cuts take place in Connecticut, Maryland, New Jersey, and Virginia; the tax hikes in Connecticut, South Carolina, and Wisconsin (in two different years).

individually owned patents relative to the corporate patent control group.²¹

4 Estimating the Effect of Trade on Litigation

Baseline Econometric Model

Let L_{it} denote an indicator variable that is equal to 1 if at age t (period τ) at least one patent case is filed involving patent i . We assume that patents are litigated according to the following linear probability model:

$$L_{it} = \beta X_{ot} + \mu_i + \lambda_\tau + a_t + u_{it}$$

where X_{ot} are the characteristics of the owner at age t and u_{it} is the residual component. The terms μ_i , λ_τ and a_t capture patent fixed effects, period effects and age effects.²²

Letting j denote the initial owner of the patent and k the buyer of the patent, we can write

$$X_{ot} = (1 - \text{NewOwner}_{it})X_j + \text{NewOwner}_{it}X_k$$

where NewOwner_{it} is a dummy that equals one from the date the patent changes ownership, and X_j and X_k are owners characteristics that we assume are constant over time for simplicity. Then the litigation model can be expressed as

$$L_{it} = \beta X_j + \text{NewOwner}_{it}\beta(X_k - X_j) + \mu_i + \lambda_\tau + a_t + u_{it}. \quad (6)$$

Equation (6) provides useful guidance in interpreting our empirical results. In the next section we will regress litigation on trade in panel regressions of the form

$$L_{it} = \alpha \text{NewOwner}_{it} + \varpi_i + \lambda_\tau + a_t + u_{it}. \quad (7)$$

In light of equation (6), the patent fixed effects, ϖ_i , will capture the combined effect of the characteristics of the initial owner, βX_j , and the patent characteristics, μ_i (i.e., $\varpi_i = \beta X_j + \mu_i$). More importantly, the coefficient on the traded dummy, α , can be rewritten as $\beta(X_k - X_j)$. This has two implications. First, we can interpret the coefficient on trade as the impact that

²¹As a robustness check on the difference-in-differences identification strategy, we follow the suggestion of Angrist and Pishke (2009) to add a time trend interacted with the treated group (individually-owned patents). This allows for the time trends for the treated and control (corporate owned patents) groups to be different. To do this, in each of the diff-in-diff regressions we include both an intercept dummy for individually-owned patents as before and an interactive dummy between individually owned patents and a time trend variable. We find that the estimated effect of tax changes is robust to the inclusion of the differential trends.

We also tried an even more flexible specification which allows for an interaction between the dummy for individually owned patents and a piece-wise linear trend (with the same four sub-periods as used in the paper). Again the results confirm the estimated effects of tax cuts and hikes on patent trading.

²²We cannot include year dummies because the patent grant year is absorbed by the patent fixed effect. We include dummies for four time periods: before 1986, 1986-90, 1991-95, and after 1995.

the change in ownership characteristics (if unobservable to the econometrician) has on patent litigation. If we were able to control for all the owner characteristics that affect litigation risk, the coefficient α should be zero. The second implication is that α will differ from zero only if two conditions hold: first, there are unobservable owner characteristics that affect litigation outcomes (i.e., $\beta \neq 0$) and, secondly, the market for innovation reallocates patents to entities that differ substantially in these characteristics (i.e., $X_k \neq X_j$).²³ Previous literature on patent litigation confirms that owner characteristics substantially affect litigation risk (e.g., Lanjouw and Schankerman, 2001 and 2004), but there is no existing research on the link between the reallocation of patent rights and litigation risk. To our knowledge, this is the first paper that studies this link and the sorting which the market for innovation induces.

Identifying the Causal Effect of Trade on Litigation

To identify the causal effect of trade on litigation, we need to address the potential bias arising from correlation between $NewOwner_{it}$ and u_{it} . This can arise in a variety of ways. A positive shock to the value of the technology covered by the patent may lead to an increase both in the likelihood of trade and litigation, inducing positive correlation. Another possibility is that litigation may increase because firms acquire patents strategically with the purpose of blocking competitors through patent litigation. Negative correlation can arise if a cash constrained innovator is more likely to sell its patent and less likely to enforce it aggressively.

To address potential endogeneity, we need an instrument that affects the likelihood of trading a patent but does not belong directly in the litigation equation. We exploit a feature of the U.S. tax code that allows us to use the capital gain tax rates as an instrument. In the United States Internal Revenue Code, individuals face a lower tax rate on capital gains (from sales of assets) than on ordinary ('earned') income. U.S. corporations do not receive this preferential tax rate on capital gains (Desai and Gentry, 2004). According to section 1235 of the Internal Revenue Code, the transfer of a patent by an individual is treated as the sale of a capital asset and is subject to capital gain taxes. On the other hand, patent litigation damages (and licensing royalties) are taxed as ordinary income. This treatment of litigation damages is acknowledged in a number of tax court decisions (Maine and Nguyen, 2003). This means that

²³It is easy to extend the model to introduce observable characteristics of the owner. Consider the model $L_{it} = \beta X_{ot} + \gamma \tilde{X}_{ot} + \mu_i + \lambda_\tau + a_t + u_{it}$ where X_{ot} are the unobservable characteristics of the owner at age t and \tilde{X}_{ot} are the observable characteristics of the owner. Because \tilde{X}_{ot} are observed, we can estimate $L_{it} = \alpha NewOwner_{it} + \varpi_i + \gamma \tilde{X}_{ot} + \lambda_\tau + a_t + u_{it}$. In this extended model, the patent fixed effects, ϖ_i still captures the combined effect of the time invariant *unobservable* characteristics of the initial owner, βX_j , and the time invariant patent characteristics, μ_i (i.e., $\varpi_i = \beta X_j + \mu_i$). The coefficient on the traded dummy, α , can be still be rewritten as $\beta(X_k - X_j)$ and measures the impact that the change in *unobservable* ownership characteristics has on patent litigation. This extension also implies that if the econometrician is able to observe all the patent characteristics that have an impact on litigation (i. e. X_{ot} is empty) then α would be equal to zero.

the decision to trade a patent will be affected by the capital gains tax rate, but the decision to litigate will not. We limit the analysis in this paper to trades of individually owned patents because this tax distinction does not apply to patent sales by corporations, so unfortunately we cannot use this instrumental variable for transfers of company owned patents.

We start by specifying a probit equation that determines how taxes affect the probability that a patent is traded by the original assignee at age t . To do this, we generate a dummy variable, $Trade_{it}$, that equals one only in the year in which the patent changes ownership. We drop all the observations that follow the first change in ownership and estimate the following probit regression

$$Trade_{it} = \begin{cases} 0 & \text{if } p(Z_{it}, X_{it}) \leq \varepsilon_{it} \\ 1 & \text{if } p(Z_{it}, X_{it}) > \varepsilon_{it} \end{cases} \quad (8)$$

where Z_{it} is the capital gains tax rate for individuals in the state of the inventor and X_{it} is a vector of patent characteristics and additional controls. Given the probability of being traded at age t , p_{it} , we can compute the probability that the patent is not owned by the initial owner at age t as

$$P_{it} = P_{it-1} + (1 - P_{it-1})p_{it}$$

with $P_{i1} = p_{i1}$. Intuitively, the probability of not being owned by the original inventor at age t is equal to the probability of having changed ownership in the previous periods plus the probability of not having changed ownership up to age t and being traded at age t .

Denote the predicted probability from the probit model (8) as \hat{p}_{it} . We use \hat{p}_{it} to construct an estimate of the probability of not having changed ownership up to age t , \hat{P}_{it} , for the *entire sample* of observations. This estimate satisfies two important properties. First, \hat{P}_{it} depends on capital gain tax rates Z_{it} which are assumed to be uncorrelated with the likelihood of patent litigation (except through changes in ownership). Specifically, we expect larger capital gain tax rates to reduce the probability of a change in ownership.²⁴ Second, \hat{P}_{it} is equal to expected value of $NewOwner_{it}$, conditional on the current value and past sequences of Z_{it} and X_{it} . These properties allow us to exploit \hat{P}_{it} as an instrumental variable to estimate the effect of trade on litigation (Angrist, 2001; Angrist and Pishke, 2009). Newey (1990) and Wooldridge (2002) argue that this IV estimator is asymptotically efficient in the class of estimators where the IVs are functions of Z_{it} and X_{it} . Moreover, because we are using \hat{P}_{it} as an instrument for $NewOwner_{it}$, the model for P_{it} does not need to be correctly specified as long as the linear projection of $NewOwner_{it}$ on P_{it} actually depends on P_{it} (Wooldridge, 2002).

²⁴Because \hat{P}_{it} is a non-linear function of $\hat{p}_{it}, \hat{p}_{it-1}, \dots, \hat{p}_{i1}$, it depends on the entire vector of current and past capital gain taxation rates $\bar{Z}_{it} = (Z_{i1}, Z_{i2}, \dots, Z_{it})$. Nonetheless, the relevant thing in our setting is that, conditional on a patent not having been traded, a change in ownership only depends on Z_{it} .

In short, we estimate the following 2SLS model. The first stage regression is

$$NewOwner_{it} = \beta \widehat{P}_{it} + \varpi_i + \lambda_\tau + a_t + \xi_{it}$$

and the second stage is the same as (7) except that $NewOwner_{it}$ is replaced by the fitted values from the first stage

$$L_{it} = \alpha \widehat{NewOwner}_{it} + \varpi_i + \lambda_\tau + a_t + u_{it}. \quad (9)$$

Empirical Results

Trade and Litigation: Correlations

In Table 3, we begin by presenting OLS estimates of our baseline econometric model (7). The first three columns present estimates using the full sample (including patents that are not traded and/or litigated). In column 1, where we do not include any controls, the coefficient on the NewOwner dummy is positive and significant. However, this result is likely due to selection into trading, since more valuable patents are both more likely to be traded and litigated.²⁵ In column 2 we include patent fixed effects in the specification to control for this possibility. This specification makes use only of within-patent litigation variation. Once fixed effects are included, the coefficient becomes negative and significant, indicating that a patent is *less likely* to be involved in a suit after it changes ownership. A Hausman test strongly rejects the null hypothesis that the patent effects are random.²⁶ The negative correlation between change of ownership and litigation is robust when we introduce age effects and time period dummies, in column 3.²⁷ Finally, in column 4 we present a similar regression using the much smaller subsample of patents that are both traded and litigated at least once in their life. Also in this smaller sample we find a negative correlation between trading and litigation, but not statistically significant.²⁸

²⁵A similar positive correlation is found by Chien (2011) in a small random sample of patents granted in 1990.

²⁶To perform the test, we run a random effect panel regression with additional covariates. The additional controls are the number of citations made by the patent, the number of citations received, the number of claims, and the technology class of the patent. The random effect coefficient on NewOwner is positive and significant in the random effect specification ($\hat{\beta} = 0.003$, p -value < 0.01). The estimate from the fixed effect specification is negative and significant ($\hat{\beta} = -0.002$, p -value < 0.01). We strongly reject that the two estimates are equal ($\chi^2_{20} = 1878.19$, p -value < 0.01).

²⁷In this regression we add dummies for four sub-periods: before 1986, 1986-90, 1991-95, and after 1995. In a more general specification with a dummy for each year and no age dummies, we do not reject the joint hypothesis that the individual year coefficients can be summarised by these four period dummies. Results are very similar if we drop the period dummies and control for only for age and patent fixed effect.

²⁸We observe a similar negative correlation even when we drop the unassigned patents and focus on patents that are assigned to a US individual at their grant date (14,576 patents). The correlation between trade and

We also looked at the relationship between trade and litigation for small corporate innovators. As for individual inventors, we find a positive (but insignificant) correlation between trade and litigation when we do not include patent fixed effects, but the correlation becomes negative once we control for patent fixed effects. The magnitude of the correlation is very similar to the one estimated for individually owned patents (-0.016 and $p < 0.01$), further suggesting that the two samples are quite similar.

The results in Table 3 are to be interpreted as correlations between litigation rates and changes in ownership, not causal impacts. As we argued above, there are a number of reasons why we should expect unobservable factors to affect both the trading and litigation decisions. This intuition is confirmed by a Rivers-Vuong test that provides strong evidence against the exogeneity of patent trade.²⁹ To address this endogeneity we will now construct an instrument that relies on the effect of capital gains taxes on patent trading.

Impact of Taxes on Patent Trading

Table 4 presents estimates of the impact of taxes on changes in patent ownership. The dependent variable is an indicator variable, *Trade*, that equals one only in the year in which the patent changes ownership. Because tax rates affect the initial owner incentives to sell the patent up to the time at which the patent is sold, we estimate these regressions dropping all observations that follow the first change in ownership. In all the regressions we control for a range of observable patent characteristics, including the number of citations received by the patent, a measure of the patent generality, technology class (36 two digit sub-categories), plus year and age fixed effects. Because the main variable of interest (capital gain tax rates) varies at the state-year level, we report state-level block bootstrapped standard errors (Bertrand et. al., 2004).³⁰

Column 1 presents the estimates of the probit model (8). The regression confirms that

litigation is negative and significant at the 5 percent level. The confidence interval of the coefficient on traded overlaps with the one estimated in the large sample.

²⁹Following Rivers and Vuong (1998), we regressed *NewOwner* on capital gain taxes, age dummies and period dummies in a linear probability model with fixed patent effects. We constructed the residuals (\hat{v}) for this model and then regressed the litigated dummy on *NewOwner*, age, period dummies and \hat{v} . The coefficient on \hat{v} is positive and highly significant (point estimate of 0.054, $p - value < 0.01$).

³⁰We use the NBER data set for information on the number of citations, grant date and technology class for each patent. Since citation counts are inherently truncated, we use the truncation-adjusted citations counts contained in the NBER patent data (Hall et. al., 2001 for details). The NBER data also provides an index of patent “generality,” defined as one minus the Herfindahl index of the citations received by a patent across different technology classes. The measure is high if the patent is cited by a wide range of technology fields.

Our findings are robust to dropping patent citations, generality and including only patent characteristics that are fixed at the time of issue. Results are also robust to clustering standard errors at the state and state-year level.

higher capital gains and corporate tax rates reduce the likelihood that patent rights are traded, and higher income tax rates increase it. These results are consistent with the predictions of the model presented in Section 2. The estimated marginal effects imply large tax impacts. The elasticity of the probability of trade with respect to the capital gains tax rate is -1.62. The corresponding elasticities for the personal income tax rate and corporate tax rate are 1.22 and -0.77, respectively. Column 2 shows that results are similar if we use a linear probability model.

As a robustness check, in column 3 we present the estimates for a discrete time, proportional hazard model. This nonlinear model is more demanding in that it is designed to explain the timing of trade – i.e. the probability of being traded in a year, given that trade has not yet occurred. Following Jenkins (1998), we estimate a discrete specification of a Weibull proportional hazard model by maximum likelihood method. The parameter estimates from this hazard model also imply large tax effects on trading: the elasticity of the hazard of trade with respect to the capital gains tax rate is -1.64. The corresponding elasticities for the personal income tax rate and corporate tax rate are 1.33 and -0.72, respectively. Finally, in column 4 we focus on the sub-sample of patents that are litigated and traded at least once in their lifetime. Despite the huge reduction in sample size, we still find a negative and significant coefficient for capital gains tax rates. In this restricted sample the estimated coefficients on citations received and generality are not significant. This is not surprising because all the patents in this smaller sample are traded, and our time-invariant measures of patent characteristics have little explanatory power on the timing of trade.

There is a concern that the impact of taxes on trade may reflect omitted variables correlated with tax rates rather than incentives to trade generated by the asymmetric tax treatment of trade and litigation faced by individuals. To check this possibility, in Table 5 we explore the relationship between taxes and corporate patent transfers. Because capital gains and income tax rates should not affect corporate transfers, the analysis has the natural interpretation of a placebo test. Moreover, a simple extension of our model to corporate trades suggests that corporate income taxes should be negatively associated with the transfer of patents.³¹

Ideally, to estimate the effect of taxes on corporate transfers, we would like to use the effective marginal tax rate relevant to the company decision-making. This is very difficult to measure because in practice this rate depends on statutory rates, firm income, and dynamic features of the tax code such as the effects of net operating losses and projections of future

³¹If the gross profits for the corporate buyer are π^B and the seller, by keeping the ownership, obtains gross profits equal to π^S , trade occurs if $(1 - \tau^C)^2 \pi^B \geq (1 - \tau^C) \pi^S$, a condition that is more stringent with larger τ^C . The square on the left hand side of the inequality reflects the fact that the buyer's willingness to pay is only his net profit, and the seller then pays corporate income tax on the proceeds of the sale.

income. To proxy this tax rate, we exploit a measure developed by Graham (1996) – essentially it is a simulated marginal tax rate that approximates the “true” tax variable managers use in their decisions (Graham provides a detailed description). These tax rates, which vary across companies and time, have been computed by Graham for essentially all firms listed on Compustat.³² We obtained this simulated marginal tax rate for the period 1980-2010 and matched it with the sample of large corporate innovators described in Serrano (2010). The final sample contains 150,511 corporate patents.

Table 5 reports OLS regressions of corporate transfers against tax rates in this sample. Because all the patents in the sample are owned by Compustat firms, we were able to include additional, time-varying controls for employees, cash flow, assets and the portfolio size. We include firm fixed effects and cluster the standard errors at the firm level. The first column shows a strong negative correlation between the new corporate tax variable and trade. In columns 2-4 we add controls for individual income and capital gain taxes. The coefficient of income and capital gain taxes are statistically insignificant and have the opposite sign from the coefficients obtained for individuals in the OLS regression (2) of Table 4. For the estimates in column 2 the elasticities are substantially different than those for individuals. The elasticity of the probability of trade with respect to the capital gains tax rate is -1.6 and statistically significant for individually-owned patents (p-value<0.001), but only 0.19 and statistically insignificant for corporate patents (p-value= 0.94). The large differences in statistical significance, magnitude and sign of the two elasticities confirm that the strong relationship between capital gain taxes and patent trade that we uncovered for individual inventors is *not* present in company-owned patents.³³ In column 3 we show that results are robust to controlling for citations received. In column 4 we show that results are qualitatively similar when we control for income and capital gains taxes in the state of the company headquarters, rather than the state of the inventor.

In Table 6, we move to a stronger specification for trading by individuals that includes patent fixed effects. Column 1 shows a negative correlation between trade and capital gain taxes in a panel regression with patent fixed effects as well as controls for patent age and time period effect. In column 2 we add controls for income and corporate tax rates, and we still find the key negative (and significant) correlation between trade and capital gain taxes. The correlations with income and corporate taxes, despite having the sign predicted by our model,

³²Graham conducts a variety of validation exercises that show the simulated tax variables outperform alternative tax measures. For example, Graham and Mills (2008) show that these tax rates do a better job in explaining financial statement debt ratios than alternative tax measures.

³³We also ran a similar probit regression using the entire sample of patents that are assigned to U.S. corporations by the time they are granted and measuring the corporate tax rate faced by the firm owning the patent, using the sum of the federal and state top marginal rates for corporates in the state of the inventor. In this sample, we estimate an elasticity of the probability of trade with respect to the capital gain tax rate equal to -0.2 for corporate patents, which is substantially different from the one estimated for individual patents.

are no longer statistically significant in this more demanding specification.

In column 3 we re-run our placebo test in the specification that includes patent fixed effects. This regression uses the entire sample of corporate patents (not just Compustat firms), so the corporate tax variable we use here is constructed using the statutory rates. The coefficient on capital gain taxes is now much smaller and statistically insignificant (p-value = 0.99). The coefficient on the capital gain tax rate for trade of individually owned patents, estimated in column 2, implies an elasticity of -0.6 (s.e. = 0.1), whereas the estimate for corporate owned patents in column 3 corresponds to an elasticity of -0.01 (s.e. = 0.28). A formal test rejects the null hypothesis that the two elasticities are the same (p-value = 0.04). In column 4 we present the placebo test for the sub-set of large patenting firms for which we have the more precise measure of corporate tax rates described in Graham (1996). The estimated coefficient on the capital gain tax rate remains statistically insignificant, as does the coefficient on the personal income tax rate, consistent with the falsification test (moreover, their point estimates have the opposite sign than those obtained for the analysis of individual inventors). The estimated coefficient on the corporate tax rate using this preferred measure is negative, as one expects, and significant at the 10 percent level.³⁴

Overall, this evidence strongly supports the hypothesis that capital gains, ordinary income and corporate tax rates affect the likelihood that individual patent rights are traded. Since market based reallocations presumably increase the surplus generated by the patented innovations, the fact that taxes affect transactions in intangible assets is of independent interest, quite apart from the usefulness of capital gains taxes as an instrument for identifying the impact of such trade on litigation. Our finding adds to the recent literature that documents the impact of capital gains taxation on the frequency and timing of small business transfers involving tangible assets (Chari, Golosov and Tsyvinski, 2005; Gentry, 2010).

Causal Effect of Trade on Litigation

The parameter estimates from the regressions in Table 4 allow us to compute the probability that a patent is traded *at a specific age*, \hat{p}_{it} . The estimate of \hat{p}_{it} can be used to construct an estimate of the probability that $NewOwner_{it} = 1$, which we denote by \hat{P}_{it} . To estimate the causal effect of trade on litigation, we use \hat{P}_{it} as an instrument for the endogenous variable, $NewOwner_{it}$. Econometrically, the exogenous variation is derived from the capital gain tax rates, but any monotonic function of this variable can be used as an instrument and \hat{P}_{it} is a

³⁴In unreported fixed effect regressions, we also examined whether variation in corporate tax rates is associated with variation in *litigation* of corporate owned patents. We find no statistically significant correlation between litigation of corporate patents and (statutory) capital gains tax rates in a regression, both in the entire sample of corporate patents and in the sample of large corporate innovators for which we have the firm-specific data on marginal corporate tax rates.

typical choice when the endogenous variable is binary (Doyle, 2007; Lileeva and Treffer, 2010). In all the first stage regressions (reported in Appendix Table A2), \widehat{P}_{it} is strongly correlated with the indicator variable *NewOwner*, and the F-test of joint exclusion of the instruments does not indicate problems of weak instruments (p-values <0.01).

Table 7 presents the parameter estimates using this IV strategy. Column 1 reports estimates when *NewOwner* is instrumented by the \widehat{P}_{it} , constructed from the probit regression in column 1 of Table 4. Columns 2 and 3 show that the estimated coefficient is nearly identical if the instrument is obtained from the linear probability model or the proportional hazard model. In all regressions the estimated causal effect of a change in ownership on litigation is negative and significant, and the point estimate (in absolute value) is about six times larger than the simple OLS estimate in column 3 of Table 3. This result highlights the importance of controlling for the endogeneity of trade, and indicates a strong positive correlation between *NewOwner* and the disturbance in the litigation equation (inducing an upward bias, in absolute value, if we treat changes in ownership as exogenous).

In column 4 we present similar IV regressions using the sub-sample of traded and litigated patents. In this case too, the estimated coefficient on the change in ownership variable is negative and significant, and much larger in absolute value than the OLS estimates in Table 3.³⁵

For the full sample, the IV estimate indicates that a change in ownership reduces the annual litigation probability by about 1.2 percentage points. In the sub-sample of traded and litigated patents, the causal effect is an order of magnitude larger, corresponding to a reduction in litigation probability by about 16 percentage points. While the difference in the magnitude of the marginal effect across the two samples is very large, the implied elasticities are fairly similar (i.e., the differences in marginal effects are driven by differences in the mean litigation probability). Evaluated at sample means, the implied elasticity is -0.43 (std.dev.=0.10) in the full sample, and -0.91 (std.dev.=0.39) in the restricted sample.³⁶

The econometric model in equation (9) assumes a constant effect of trade on litigation. It is tempting to interpret these 2SLS estimates as weighted averages of heterogeneous responses. However, Angrist and Imbens (1995) show that estimates from constant effect models may differ substantially from average heterogeneous effects when there are continuous instruments

³⁵There is the concern that spurious correlation may arise between low capital gain taxes and litigation because of macro-economic variables. To address this concern we exploit U.S. Bureau of Economic Analysis data on gross state product per capita (GSP) in 1997 dollars. Controlling for GSP, we obtained results very similar to those in our baseline specification. GSP itself is not significantly correlated with patent trade or litigation.

³⁶In our empirical setting, these elasticities measure the expected drop in annual litigation rate for a patent that has not yet been traded. In Section 6, we provide a more complete quantification of the effect of trade on litigation by simulating changes in individual tax rates.

and covariates. To address this concern, we follow the procedure they suggest and convert the instrument into a dummy variable. Specifically, we generate the indicator variable *Large Tax Difference*_{it} = 1 if the difference between capital gains and income tax rates is above the 75th percentile of our dataset (19 percentage points). We exploit this binary instrument to estimate the local average treatment effect (LATE): the average effect of a change of ownership for those patents whose owners were induced to sell their patents by the substantial difference between capital gains and income tax rates (Imbens and Angrist, 1994). The average treatment effect is "local" because not all patent owners are induced by the instrument to sell the patent.

In column 5 we present the estimates from the Angrist-Imbens procedure.³⁷ The second-stage LATE estimates indicates that a change in ownership reduces the annual litigation probability by about 3.3 percentage points. While the magnitude of the effect is larger than the one obtained in the constant effects model, the confidence intervals of the two coefficients overlap. Thus, in our setting the estimates of the simple IV model do not differ substantially from the mean of the heterogeneous responses.

The estimated LATE measures the effect of trade on the unidentifiable sub-population of patents that change ownership because of a change in capital gains taxation. Thus it is difficult to map the coefficients estimated in Table 4 to reductions in litigation rates for the average patent in our sample. A plausible assumption is to consider those sample patents that are litigated but not traded as "at risk" of being affected by a change in taxation. For these patents, the estimated LATE from column 4 implies a reduction in the annual litigation rate of about 32 percent.³⁸

Finally, there is the concern that individual inventor mobility may affect our IV estimates. The first reason is a measurement issue. We measure the capital gains tax rate based on the state of residence of the inventor. If inventors move before trading (even if the decision to move is unrelated to the decision to trade), we will get random measurement error in the capital gains tax rate – we should use the new state of residence but we use the original one. This would lead to attenuation bias in the coefficient of the capital gains tax rate in the regression of trade on tax rates. If inventors move to lower capital gains tax states with the intention of trading,

³⁷The LATE is estimated as a two-stage least square regression with the endogenous variable instrumented by the dummy instrument and interactions between the instrument and covariates. For details on the procedure, see Angrist and Imbens (1995; Theorem 3) and Lileeva and Treffer (2010).

³⁸Specifically, the average annual probability of litigation among patents that are litigated (but not traded) at some point in their life is 0.103. Thus the percentage change implied by the estimate LATE is equal to $-0.033/0.103 = -0.32$. We obtain very similar results using the 90th and 95th percentiles, and the point estimates are larger than the one obtained using at the 75th percentile. However, while our model suggests that larger tax differences are likely to generate a larger amount of trade, the actual impact on the average litigation rate depends on the relative strengths of commercialization and enforcement gains for the patents that are traded because of the larger tax differences.

this will reinforce the attenuation bias, as we do not measure the lower tax rates (in the new state) that is relevant. However, while it affects the coefficients in the trade equation, there is no reason to think that inventor mobility would affect the consistency of our IV estimates. All that matters for consistency is the orthogonality condition that the disturbance in the litigation equation is uncorrelated with the relevant capital gains tax rate. If inventors move state to pay less capital gain taxes on the profits from the sale of the patent; this is unlikely to affect the new owner incentives to litigate.³⁹

5 Heterogeneous Effects of Trade on Litigation

Estimation and Results

The econometric model developed in the previous section assumes that the treatment effect of trade on litigation is identical across all patent transactions. However, the underlying motivation for transactions may vary, with the gains from trade coming from a variety of sources with different implications for post-trade litigation. Therefore, we now extend the model to allow for heterogeneous treatment effects as follows:

$$\begin{aligned}
 NewOwner_{it} &= \begin{cases} 0 & \text{if } P_{it} \leq v_{it} \text{ and } NewOwner_{it-1} = 0 \\ 1 & \text{otherwise} \end{cases} \\
 L_{it}(NewOwner_{it} = 1) &= \alpha_i + \varpi_i + \lambda_\tau + a_t + u_{it} \\
 L_{it}(NewOwner_{it} = 0) &= \varpi_i + \lambda_\tau + a_t + u_{it}
 \end{aligned} \tag{10}$$

where ϖ_i are patent fixed effects, a_t and λ_τ are patent age and year effects.⁴⁰ We assume that α_i can be decomposed into a common component ($\bar{\alpha}$) and a random component (ψ_i): $\alpha_i = \bar{\alpha} + \psi_i$. The heterogeneous effect of new ownership on litigation is

$$L_{it}(NewOwner_{it} = 1) - L_{it}(NewOwner_{it} = 0) = \bar{\alpha} + \psi_i.$$

Consider an increase in the value of the technology that makes the patent both more likely to be traded (small v_{it}) and more likely to be litigated after trade (high ψ_i). Together these imply that we should observe a negative correlation between ψ_i and v_{it} , and thus a

³⁹In any event, cross-state mobility of individual inventors is quite limited, in part because social connections affect access to finance for small ventures (Sorenson and Stuart, 2001; Breschi and Lissoni, 2009).

⁴⁰Notice that, conditioning on a patent not having changed ownership ($NewOwner_{it-1} = 0$), trade occurs if $p(Z_{it}, X_{it}) > \varepsilon_{it}$. Multiplying both sides of the inequality by $(1 - P_{it-1})$ and adding P_{it-1} , we obtain

$$P_{it} > v_{it} \equiv P_{it-1} + (1 - P_{it-1})\varepsilon_{it}$$

This is the relationship between $NewOwner$ and P described in the first formula of the econometric model. Even if the ε_{it} are assumed to be independent draws, the impact of v_{it} on $NewOwner$ depends on the entire sequence of past realizations of ε_{it} . The serial correlation in v_{it} is not a problem as long as v_{it} is uncorrelated with Z_{it} .

negative correlation between v_i and the effect of trade on litigation. More formally we should expect $E(\bar{\alpha} + \psi_i|v_{it})$ to be decreasing in v_{it} . Because v_{it} is not observed it is not possible to condition on it. Nonetheless, for an inventor that is just indifferent between trading and not trading, it must be that $P(Z_{it}) = v_{it}$. Exploiting this equality, we obtain the marginal treatment effect $E(\bar{\alpha} + \psi_i|P(Z_{it}))$, which corresponds to the (heterogenous) effect of trade on litigation for patents that are traded because of the instrument.

Heckman and Vytlacil (1999) provide a formal treatment, where they show that

$$E(\bar{\alpha} + \psi_i|P = v_{it}) = \frac{\partial E(L_{it}|P)}{\partial P}\Big|_{P=v_{it}} \quad (11)$$

and establish identification of the marginal treatment effect. Specifically, the observed litigation is

$$\begin{aligned} L_{it} &= L_{it}(NewOwner_{it} = 1)NewOwner_{it} + L_{it}(NewOwner_{it} = 0)(1 - NewOwner_{it}) \\ &= (\bar{\alpha} + \psi_i)NewOwner_{it} + \varpi_i + \lambda_t + a_{it} + u_{it}. \end{aligned}$$

The marginal treatment effect (*MTE*) can be computed by estimating the expected litigation conditional on P , $E(L_{it}|P)$. Let \hat{P} be our estimate of the probability that a patent is not owned by the initial inventor. Substituting this into the observed litigation equation, we obtain a partially linear model

$$E[L_{it}|\hat{P}_{it}] = E[(\bar{\alpha} + \psi_i)NewOwner_{it}|\hat{P}_{it}] + \varpi_i + \lambda_t + a_{it}. \quad (12)$$

The derivative of (12) can be semi-parametrically and non-parametrically estimated in order to obtain the marginal treatment effect. For the semi-parametric estimation, we follow Carneiro, Heckman, and Vytlacil (2010) and approximate $E[(\bar{\alpha} + \psi_i)NewOwner_{it}|\hat{P}_{it}]$ with a third order degree polynomial, obtaining

$$E[L_{it}|\hat{P}_{it}] = c_1\hat{P}_{it} + c_2(\hat{P}_{it})^2 + c_3(\hat{P}_{it})^3 + \varpi_i + \lambda_t + a_{it}$$

which implies a *MTE* equal to $c_1 + 2\hat{P}_{it}c_2 + 3c_3(\hat{P}_{it})^2$. For the non-parametric approach, we follow the multistep procedure developed by Heckman, Ichimura, Smith and Todd (1998) and Carneiro, Heckman, and Vytlacil (2010).⁴¹

Figure 3 shows the semi-parametric estimates of the *MTE* for the entire sample.⁴² The horizontal axis depicts the estimated probability that the patent has changed ownership, \hat{P} .

⁴¹For details, see Appendix 2. The first part of the procedure involves estimating the litigation equation non-parametrically. This is a non-parametric counterpart of the *IV* estimates and it involves the use of local linear regressions. The second part of the procedure involves numerically differentiating the estimated $E[L|\hat{P}_{it}]$

A simple test of heterogeneity suggested by Carneiro et. al. (2010) involves testing the null hypothesis that the coefficients of the second and third order are jointly equal to zero. The *F*-statistic for $\hat{c}_2 = \hat{c}_3 = 0$ is 24.28 ($p < 0.01$) in the sample for litigated and traded patents and 13.61 ($p < 0.01$) in the full sample.

⁴²For the large sample, only the semi-parametric *MTE* could be estimated because running local linear regressions in a panel with more than two million observations is infeasible with the available computer hardware.

The vertical axis shows the effect of trade on litigation for different values of \hat{P} (dashed lines are 95 percent bootstrapped confidence intervals). The support for \hat{P} goes up to 0.15, which corresponds to the 99th percentile for the measure. The absolute value of the estimated marginal treatment effect is monotonically increasing in \hat{P} . Patents with low value of \hat{P} are those that, given their observables, are unlikely to have changed ownership (e.g. patents that are not highly cited or with low generality index). The small (or insignificant) values of the *MTE* in this range show that, if a change in capital gain taxes induced the owner of one of these patents to sell, the change in litigation risk would be negligible. Conversely, patents with high \hat{P} are those at high risk of being traded. For these patents the *MTE* is negative, indicating a substantial drop in the likelihood of litigation from transfer of ownership. The *MTE* becomes statistically significant for values of \hat{P} above 0.03 that roughly correspond to the median of the \hat{P} distribution.⁴³

Appendix Figure A1 reports the non-parametric estimates of the *MTE* for the subsample of patents that are traded and litigated. The support for \hat{P} differs from the one in the previous figure because the estimated probability of a change in ownership is greater in the sample where all patents are traded. Nonetheless, in this case too we find that the absolute value of the estimated effect is monotonically increasing in \hat{P} , and statistically significant for values of \hat{P} greater than 0.5. The figure looks similar if we estimate the *MTE* with the semi-parametric procedure employed for the larger sample.

These results point to two important conclusions. First, the main impact of trading in patent rights, over most of the range of \hat{P} , is to reduce litigation risk, suggesting that comparative advantage in patent enforcement may be more important than comparative advantage in commercialisation, at least for transfers involving individually owned patents. Second, the results show that patents with larger estimated (enforcement) gains from trade are in fact those with the highest predicted likelihood of changing ownership. This suggests that the market for innovation reallocates patent rights efficiently, at least in this sense.

⁴³Two points should be noted. First, to compare the *MTE*'s with the 2SLS estimate of the *LATE*, we split the distribution of \hat{P} into seven adjacent bins and computed a weighted average of the *MTE*s evaluated at the bin mid-points, with weights equal to the fraction of observations in each bin. The average is equal to -0.013, which is very close to our 2SLS estimates.

Second, to assess what fraction of the variation in \hat{P} is generated by the instrument, we re-estimated the probability of trade using only variation in capital gains taxes and fixing all other covariates at their sample means. The resulting distribution is approximately uniform, with a support ranging from zero to 0.08, which is approximately the 90th percentile of the distribution of \hat{P} . This confirms that the instrument generates substantial variation in the predicted change in ownership.

Unbundling the Marginal Treatment Effect

Baseline Analysis

We have shown that the effect of trade on litigation is heterogenous, and that the effect reduces litigation more strongly for patents at greater risk of being traded. This suggests that the nature of the transactions varies and that there is a particular type of sorting: patents that are less likely to be traded (low values of \hat{P}) are more likely involved in transactions based on commercialisation advantages, and patents with high values of \hat{P} are more likely to be in transactions driven by enforcement gains. To understand this sorting better, in this section we unbundle the marginal treatment effect and relate it to observable characteristics of the transaction.

To do this, we need information on patent buyers. The USPTO reassignment data contain non-standardized names of the buyers, so buyer characteristics must be manually recovered. We perform this manual match for the 569 patents that were both traded and litigated at least once in their lifetime. For each of these patents, we constructed the size of the portfolio of the buyer, defined as the number of patents obtained in the twenty years before the trade occurs. Our matching shows that most transactions involve trade from an individual owner to a firm (only 11.4 percent of cases involve two individuals). The distribution of buyer portfolio size is highly skewed. The median portfolio size for acquiring firms is one patent, the 75th percentile is three, and the mean is 106.1.

We use the buyer portfolio to construct two variables to capture the two basic motivations for transactions: enforcement gains and commercialisation (product market) gains. The first variable, *LargeBuyer*, is equal to one if the buyer’s portfolio includes at least eight patents at the time of the transaction (i.e., if the buyer had that number of patents granted in the preceding twenty years). which corresponds to the top decile of the portfolio size distribution. Lanjouw and Schankerman (2004) show that firms with large patent portfolios are less likely to file a suit on any individual patent in their portfolios, controlling for patent characteristics, and argue that this reflects their ability to resolve disputes ‘cooperatively’ without resorting to the courts. In addition, larger firms often use broad cross-licensing agreements to avoid costly litigation and preserve freedom to operate in their R&D activities (Galasso, 2012). Building on this idea, we expect “enforcement gains” (reduction in litigation) to be greater for patents acquired by large buyers.

The second variable is designed to capture transactions where the traded patent is a good match for the technology profile of the buyer, where comparative advantage in manufacturing or marketing are more likely to be realized. We define *TechFit* as a dummy variable equal to one if the acquired patent belongs to the technology area to which the plurality of the buyer’s

patents are assigned. To do this, we use the 36 technology sub-categories defined in Hall et. al. (2001). The *TechFit* measure is equal to one for 140 patents (about 25 percent of the sample). The hypothesis is that in such cases the product market gains from the transaction will be larger, and thus that such transactions tend to raise, not lower, litigation risk.

Table 8 presents instrumental variable regressions that examine how buyer portfolio size and technology fit affect the impact of trade on litigation. Column 1 confirms that patents traded to small entities, and that fit well in the buyer's portfolio, experience an *increase* in litigation after they are traded. These are transactions where we expect product market gains to be important and enforcement gains negligible. In sharp contrast, column 2 shows that the largest reduction in litigation rates occurs when patents are traded to large entities with low fit in the buyer patent portfolio, where enforcement gains are large and product market gains are small. Columns 3 and 4 show that trade is associated with a reduction in litigation of smaller magnitude for transactions where both sources of gains (or none) are present. In column 5 we confirm the results using the pooled sample and interacting the traded indicator with the dummies for large portfolio and patent fit.⁴⁴

These results are consistent with the idea that the market for patents generates efficiency gains by reallocating patents to large patentees that have an advantage in avoiding patent disputes and/or litigation arising out of disputes. We can illustrate this mechanism with a concrete example from our sample. Our data indicate that five individually owned patents were re-assigned to Sandisk Corporation in 1995. SEC annual report filings show that, during the period 1989-99, Sandisk signed cross-licensing deals with Intel, Sharp, Hitachi, Samsung, Silicon Storage Technology and Toshiba. This means that, for the five acquired patents, the risk of litigation with these companies disappeared once they were re-assigned to Sandisk. For a more recent example, on July 6, 2012 Facebook announced a patent portfolio cross-licensing arrangement with Yahoo. Before entering the deal, Facebook acquired a number of patents from AOL, Microsoft and IBM. Because of the cross-licensing deal, all these patents will face no litigation risk with Yahoo. More importantly, the widely discussed start-up acquisition strategy followed by Facebook (Cenicola, 2012) implies that patents initially owned by small innovators will not be involved in suits with Yahoo, once they are assigned to Facebook.

⁴⁴The regression in column 5 is a constrained version of those in columns 1-4, where period dummies are the same across the different samples. In a more general specification, we do not reject the hypothesis that the period dummies are the same across the four groups of transactions. In column 5 we allow the age dummies to differ across the samples because the data strongly indicated differences in the impact of these dummies in the first stage regression.

Robustness of Unbundling Analysis

In Appendix Table A3 we present the unbundling results for a number of alternative specifications. First, we increase the buyer threshold from 8 to 12 patents. Second, in the baseline analysis we set $TechFit = 1$ when the buyer is an individual (in such cases we cannot measure the portfolio size). Results hold up if we instead set $TechFit = 0$ in such cases. Third, we employ a $TechFit$ measure constructed using a finer technology classification (we move from 36 technology sub-categories to about 400 USPTO patent classes). Finally, we use a citations-based $TechFit$ measure (rather than one based on the patent classification). Specifically, we define $TechFit = 1$ if either the acquired patent cites one of the patents of the buyer or if the patents of the buyer cited the acquired patent. In all these regressions, our baseline unbundling results are robust: the interaction between the *NewOwner* indicator and $TechFit$ is positive and significant, and the interaction with the *LargeBuyer* dummy is negative and significant. These results are consistent with the theoretical framework developed in Section 2, where the relative magnitude of product market and enforcement gains determine whether a change in ownership has a positive or negative impact on patent litigation.

We also examine the possibility that changes in ownership may simply be the way patent disputes are settled, rather than reflecting an efficient reallocation to entities that are less likely to resort to courts. There are two reasons why we do not think this case is important in our analysis. First, it is rare in our sample that trade occurs as the settlement of disputes. To identify these events, we compared the names of the parties trading a patent with those involved in the litigation. There is very little overlap: in only 20 patent cases (3.5 percent of the sample of litigated and traded patents) does a patent transfer follow a suit filed by the same parties. This fact suggests that the main effect of trade on litigation operates through comparative advantage in enforcement, and not through facilitating settlement of an existing dispute. Second, the cases in which trade occurs as the outcome of a settlement are not associated with a larger impact of trade on litigation, as one would expect if avoiding litigation with the buyer is the main reason for trade. To check this, we re-estimated the unbundling regression introducing an interaction for trades occurring as settlement (i.e., where the names of the parties in the trade and litigation are the same). The point estimate on the interaction is very small and is not statistically significant. There is essentially no change in the other coefficients.

There is a concern that the heterogeneous effects that we uncover are peculiar to the small set of (manually matched) traded and litigated patents and cannot be generalized to the rest of the sample. To address this important concern, we employed the matching procedure described in Thoma et al. (2010) in order to standardized the names of patent buyers and

those of USPTO assignees and to match the two datasets.⁴⁵ The matching is performed by constructing a matching score (the Jaccard weighted distance) based on the division of name strings into sequences of characters ('tokens'). Specifically, the score measure is given by the ratio between the weighted sum of common tokens in the two names and the weighted sum of tokens in the two names. Each token has a weight that is inversely related to its frequency in the dataset. The algorithm classifies patent buyers as "unmatched" if there is no USPTO assignee name that can be matched to the buyer name with a score above 0.3. If there is at least one assignee with score above 0.3, the software matches the buyer with the assignee having the highest matching score. Additional details about the procedure can be found in Thoma et al. (2010). Exploiting the algorithm, we were able to identify an assignee number for 8,123 of the 13,607 traded patents in our sample (about 59 percent).⁴⁶

In Table 9 we use the matching algorithm to extend the heterogeneous effect analysis to this larger sample. For consistency, in each sample we construct the *Large Buyer* dummy to capture the upper decile in the buyer portfolio size distribution of the sample. In column 1 we perform the heterogeneous effects analysis looking at the 2570 traded patents that were matched with a score of 1. The regression confirms the qualitative results obtained with the manual matching: the interaction between the *NewOwner* indicator and *TechFit* is positive and significant, and the interaction with the *LargeBuyer* dummy is negative and significant. The magnitude of the coefficients is smaller: this is not surprising given that the mean litigation rate is lower in this larger sample. In columns 2-4 we extend the sample by including patents matched with lower precision. Column 2 uses matching scores above 0.9, column 3 above 0.5 and column 4 above 0.3. In this last regression, the sample size is 8123 patents, about 14 times larger than the one manually matched. Despite the large differences in sample size, our main results are robust. Being acquired by a large buyer amplifies the negative effect of trade on litigation, whereas high fit with the buyer portfolio reduces the negative effect of trade on litigation.

⁴⁵In Thoma et al. (2010), name standardization is carried out by developing a dictionary of cleaned names in the two datasets. For example, the software removes common company words (as INC or AB) and replaces spelling variations with their harmonized equivalent (as INTL for INTERNATIONAL and its variants).

⁴⁶Before studying this enlarged sample, we explored the validity of the matching procedure in the smaller sample of traded and litigated patents for which we can compare the manual matching procedure with the new automated matching procedure. Of the 569 traded and litigated patents used in the heterogeneous effects analysis, the software identified a buyer for 368 patents (about 65 percent of the sample). For 124 of these patents the match had a score equal to 1, and for about half of the matched patents the score was above 0.8. Roughly 10 percent of the matched patents had a score below 0.4. In unreported regressions we find that in the small sample we obtain qualitatively similar results if we replace the manually matched buyers with those matched by the algorithm. We also compare the performance of the regression using the matching algorithm with one in which we randomly select a fictitious buyer. The matching algorithm is informative in the sense that it performs substantially better than a random matching procedure.

These unbundling results provide more insight into the pattern of marginal treatment effects documented in Figures 2 and 3. Our estimates suggest that patents with low values of \hat{P} are more likely to be involved in transactions driven by product market gains, and patents with high \hat{P} are more likely to be involved in transactions driven by enforcement gains. To explore this idea further, we look at the types of transactions at each level of \hat{P} . Controlling for patent age, our data show that, as \hat{P} increases, there is a decline in the number of trades to small buyers with high *TechFit*, and a corresponding increase in the low *TechFit* trades. For example, for patent age 5, about 30 percent of patents in the first quartile of the \hat{P} distribution are involved in small buyer/high fit transactions. The fraction drops to 16 percent for patents in the fourth quartile of the \hat{P} distribution. Among those patents, the fraction of low fit trades is about 65 percent in the first quartile, but 82 percent in the fourth quartile of the \hat{P} distribution.

Product Market Gains or Patent Trolls?

The unbundling exercise shows a positive association between trade and litigation only for patents traded to small entities that fit well in the buyer’s portfolio. This finding is consistent with the model developed in Section 2, where higher litigation rates are generated by product market gains. A possible alternative explanation for this finding is that the patents in this sub-sample are acquired by small, specialized patent assertion entities (a.k.a. *patent trolls*).

We conduct a series of empirical tests to distinguish between these two competing explanations. Business press, legal studies and anecdotal evidence suggest that trolls tend to be common in a few industries, specifically those with complex technologies involving many patented inputs: computers, communication and electronics. We check whether these industries are over-represented in the high-fit/small-buyer sub-sample by testing whether there are differences between the technology field composition of these patents and the composition of the other traded and litigated patents. Mean comparison tests provide no evidence of industry specialization in the high-fit/small-buyer sub-sample. This result is confirmed by a multinomial logit regressions relating the four categories of patent transactions to technology field dummies.⁴⁷

Second, we examine whether the increase in litigation observed in the small-buyer/high-fit sub-sample is driven by serial buyers. From the USPTO re-assignment data, we retrieve the number of patents acquired by each buyer in this group of trades during the sample period.

⁴⁷The fraction of patents in computers and communication (NBER category 2) is 0.11 in the high-fit/small-buyer sub-sample and 0.09 for the other traded and litigated patents, and the difference is not statistically significant (p-value=0.52). Similarly, the fraction of patents in electrical and electronics (NBER category 4) is 0.08 in the high-fit/small-buyer sub-sample and 0.10 for the other patents, and again the difference is not statistically significant (p-value=0.62).

The distribution is highly skewed (the median and mean numbers of acquired patents are 2 and 9, respectively). We generate a dummy variable capturing serial buyers (the top decile of the sample). OLS and two-stage least squares regressions where the serial buyer dummy is interacted with NewOwner show no evidence that the increase in litigation is driven by serial buyers.⁴⁸

Finally, we examine whether the increase in litigation is driven by a few *serial litigants*. To do this, we compute the number of patent cases in which each buyer is involved as plaintiff during the sample period. Of the 116 buyers, 76 are involved in no patent cases, and only five buyers are involved in more than three cases. OLS and two-stage least squares regressions that interact the number of cases with NewOwner show no evidence that the increase in litigation is driven by serial litigants.

Overall, these exercises suggest that the increase in litigation rates in the high-fit/small-buyer transactions is more likely to be driven by product market gains than by the activity of patent trolls. Together with our finding that trade reduces litigation risk for all the other transaction types, this result indicates that during our sample period trolls do not play a substantial role in the market for individually owned patents. However, we emphasize that this finding does not rule out the possibility that trolls are important in the patent marketplace. The reason is that our data includes only patents owned by individual inventors and covers the period 1983-2000. It may be that patent assertion entities are more active in the market for company owned patents. Moreover, while there is documentary evidence of patent trolls throughout the 20th century (Resis, 2006), their activity may have intensified in the post-2000 period in patent activity has been growing rapidly.⁴⁹ We plan to investigate these issues in future research.

6 Simulating Tax Effects on Trading and Litigation

We have shown that capital gains taxes affect patent trading and that these transactions affect post-trade litigation risk. In this section we use our parameter estimates to simulate the impact of various tax scenarios on the frequency of patent transactions and litigation.

Let τ^G denote the capital gain tax rate, which we assume is constant for the entire life of a patent. Let $P_t(\tau^G)$ denote the probability that the patent has been traded by age t , L_{0t}

⁴⁸Similar results are obtained with less restrictive definitions of serial buyers. For example, we generate a dummy capturing buyers that acquire more patents than those granted to them (77.6 percent of buyers in the sub-sample) and also in this case, we find no evidence that the increase in litigation is driven by serial buyers.

⁴⁹There are no econometric studies of the impact of patent trolls on innovation (or even a good measure of the extent of their activity). A recent attempt to assess the private costs of trolls can be found in Bessen et al. (2011), who study the stock market reaction to lawsuits by non-practicing entities to compute the lost wealth to defendants.

be the likelihood of being involved in at least one dispute at age t if the patent is owned by the initial inventor, and L_{1t} be the likelihood of being involved in at least one dispute at age t if ownership has changed. Then the expected number of years in which at least one dispute is filed is

$$\begin{aligned} E(L(\tau^G)) &= \sum_{t=1}^T (1 - P_t(\tau^G))L_{0t} + \sum_{t=1}^T P_t(\tau^G)L_{1t} \\ &= \sum_{t=1}^T L_{0t} + \sum_{t=1}^T P_t(\tau^G)(L_{0t} - L_{1t}). \end{aligned}$$

The estimates in Tables 4 and 7 can be used to compute $E(L(\tau^G))$ for the average patent in our sample. We measure L_{0t} as the predicted litigation probability for a patent of age t that has not changed ownership, using column 1 from Table 7. $L_{0t} - L_{1t}$ is computed with the IV estimate in the first column of Table 7. $P_t(\tau^G)$ is the predicted probability of trade for different levels of capital gain taxes τ^G , constructed using the estimates in column 1 from Table 4, and evaluated at sample means.

We compute $E(L(\tau^G))$ for different tax scenarios. In the baseline scenario we assume $\tau^G = 29.2$ (percent), which is the average value in our sample and similar to the combined (state plus federal) tax rate faced by an individual in Texas in 1995. In the second, low tax scenario we set $\tau^G = 20$, which is the lowest rate in our sample and is the combined rate faced in Florida in 1985. In the third, high tax scenario $\tau^G = 40$, which is close to the highest rate in our sample, which prevailed in California in 1997. In the last two scenarios we study the impact of removing the differential tax treatment of capital gains. First, we increase the capital gains rate to be equal to the (personal) income tax rate. Second, we equate the two rates at the lower, capital gains rate. In all these exercises, corporate tax rates are kept at their sample mean.

Table 10 summarizes the results. In the baseline scenario $E(L(\tau))$ is 0.013. Multiplying this number by the average number of disputes filed in each year in which the patent is litigated (1.2 in our sample) and adjusting for litigation under-reporting using the weights in Lanjouw and Schankerman (2001), we can translate $E(L(\tau))$ into a number of predicted disputes. Our computations predicts about 36 disputes every 1,000 patents. This estimate, computed using an entirely independent method, is very similar to the litigation level estimated in Lanjouw and Schankerman (2001), which for individuals is 35 disputes per 1,000 patents.

In the low tax scenario – representing a reduction in the capital gains rate by 9.2 percentage points – the number of traded patents nearly doubles and this generates a 36 percent reduction in the number of disputes (to about 23 per 1,000 patents). In the high tax scenario – an increase in the capital gains rate by 11 percentage points – there is a 45 percent reduction in the number of traded patents with an associated 22 percent increase in the number of disputes.

Equalizing capital gains and income tax rates is associated with a contraction in the frequency of trade and an increase in litigation. The magnitude of the effect depends on whether the equality is reached by an increase in capital gains rate or a reduction in the income tax rate. The increase in litigation rates is stronger when capital gains rates are increased to the average level of income rates in our sample (42.6 percent). These computation exercises confirm that capital gains taxes have a powerful impact on the market for patents, and thereby on the litigation risk associated with the enforcement of those patent rights.

7 Discussion of Welfare Implications

We have shown that patent transactions are strongly affected by tax policy, and that this reallocation reduces the litigation risk after trade, on average, though there is heterogeneity in the effects of trade which depends on the characteristics of the patent and the transacting parties. This finding is consistent with our hypothesis that patent transactions effectively exploit differences across firms in their ability to enforce these rights. But the broader policy question motivating the analysis in this paper is whether patent transactions increase total social welfare. To shed light on this question, we first briefly discuss the different ways in which the market for patents can affect welfare and then identify how our empirical findings contribute to this broader welfare assessment.

The market for patents can generate welfare gains of three main types. First, patent transactions facilitate the reallocation of innovations to firms that are more efficient in *commercializing* patented innovation. These gains from trade arise from differences across firms in their access to complementary assets, vertical specialization, and comparative advantages in development, manufacturing and marketing of innovations. Such patent transactions generate both private benefits and social welfare through lower commercialization costs (Teece, 1986; Arora, Fosfuri and Gambardella, 2001). In addition, these transactions promote the emergence of efficient market structures in dynamic, high-technology sectors. In particular, they facilitate a more efficient division of labor in innovation activity between small firms (or individuals) who specialize in early stage innovation and large firms whose comparative advantage lies in the later development and commercialisation of these inventions (Gans and Stern, 2000; Gans, Hsu and Stern, 2002).

Second, as we show in this paper, patent transactions can generate welfare gains by exploiting comparative advantages in *patent enforcement*. This occurs when the market for innovation reallocates patents to entities that are more effective at resolving disputes over these rights without resorting to the courts, which reduces litigation costs associated with disputes.

Third, the two channels above – which relate to *ex-post* commercialization and enforce-

ment of innovations – also increase the *ex ante* incentives for R&D investment, especially for individuals and small firms, as long as the gains from trade are shared by the inventors. Given that innovation is the main engine of economic growth, and the consensus among economists, beginning with Arrow (1962), that the positive externalities from R&D imply under-investment relative to the socially optimal level, we expect this indirect increase in *ex ante* incentives for R&D to enhance social welfare.⁵⁰

Despite these social benefits, patent transactions may also generate social costs. The modern patent landscape is characterized by a large number of patents, with fuzzy boundaries and fragmented ownership (Bessen and Meurer, 2008). A number of scholars have argued that this environment enables patent holders to obtain excessive royalties by holding up infringers after they have already invested heavily to design, manufacture, market, and sell the product with the allegedly infringing feature (Shapiro, 2001). Such *ex post* hold-up may impose an ‘innovation tax’ which reduces the *ex post* profits of innovators who use the patent by more than would have been the case if the parties would have contracted *ex ante* before the sunk investment was made. This implies that *ex post* hold-up leads to greater dilution of the *ex ante* incentives of the patent user to undertake capital investments and R&D in the first place. Of course, this argument is not as clear-cut as it is sometimes presented because patent trolls increase the returns to the patents that they acquire which, if shared by the original inventor in the patent transaction, also enhances the R&D incentives of those inventors.⁵¹ Finally, patent transactions can reduce consumer welfare if patent accumulation, or concentration of patent ownership, leads to an increase in market power.

What our empirical analysis does is to provide new, causal evidence on the functioning of the market for patents, in particular by identifying a previously unnoticed source of gains from trade – comparative advantage in the enforcement of patent rights – and showing that this is an empirically important consequence of the observed patent transactions by individual inventors. We also show that the market for innovation appears to be efficient in the sense that patents are more likely to be traded when the estimated private enforcement gains from doing so are larger. From a welfare perspective, our findings imply that the market for innovation reduces litigation

⁵⁰In a recent paper, Bloom, Schankerman and Van Reenen (2012) show that, when R&D generates ‘business stealing’ effects that dominate technology spillovers, the conventional wisdom that there is under-investment in R&D can be overturned. Nonetheless, their empirical analysis shows that there is under-investment in R&D, with the social return roughly twice the private rate.

⁵¹Theoretical analysis of trolling activity is very limited (Turner, 2011 is an exception). It is not clear to us whether, in a dynamic contest, excessive hold up can be sustainable because it reduces R&D and thus the opportunities for future rent extractions. At a minimum, we would expect this dynamic link between rent extraction and R&D must limit the scope for profitable hold-up. In addition, the degree of ownership fragmentation, and the limits to *ex ante* contracting for serial trolls whose identity would be known by potentially investing parties, may play an important role in such a context. We leave these important issues for future research.

by reallocating patents to entities that are more effective at resolving disputes over these rights without resorting to the courts, and this represents a source of both private benefits and social welfare gains. In short, our analysis characterizes and quantifies the enforcement gains from trade from the market for patents, measured in terms of reduction in litigation risk. This is an important part of the benefit side of the welfare assessment, but it does not encompass the potential cost of a more fluid market for patent transactions.

As described in Section 5, we conducted a number of empirical exercises to investigate the role of patent trolls in our assessment of patent trading and litigation. While we found that, during our sample period, trolls do not play a substantial role in the market for individually owned patents, we stress that a verdict on the extent and impact of trolls is premature. Our data only covers patents owned by individual inventors and their trade history during the period 1983-2000. It is possible that trolls are more active in the market for company owned patents, and that their activity may have intensified in the post-2000 period. If this turns out to be the case, the results in our paper are still important because they show that the market for patent rights generates real and substantial private, and social, benefits. In doing so, these results show the importance of designing remedial policies that directly target non-practicing entities and hold-up behavior (such as those recommended by Lemley and Shapiro, 2007), rather than adopting policies that broadly restrict the level of patent transactions.

8 Conclusions

In this paper we study how the market for patents affects the enforcement of patent rights. Conventional wisdom associates the reallocation of patent rights through trade with comparative advantages in commercializing the innovation. The associated product market gains from trade should increase litigation risk for traded patents. We identify a new source of gains from trade, comparative advantage in patent enforcement, and show that this mechanism reduces patent litigation. Using data on trade and litigation of individually owned patents, and exploiting variation in capital gains tax rates across states and over time as an instrumental variable, we identify the causal effect of changes in patent ownership on litigation rates.

There are three key empirical findings. *First*, capital gains taxes strongly affect market transactions in patent rights for individual inventors. *Second*, the reallocation of these patent rights reduces litigation risk for individually owned patents, on average, indicating that enforcement gains are more important than product market gains for such patents. *Third*, the marginal treatment effect of trade on litigation is heterogeneous. Patents with larger potential gains from trading are those with the highest estimated probability of changing ownership, suggesting that market reallocates patent rights efficiently. Moreover, the impact of trade is related to transaction characteristics – specifically, the size of the buyer’s patent portfolio and

the technological fit of the patent in that portfolio. Finally, we find no evidence that patent trolls play a significant role in our sample, but whether this conclusion also holds for corporate patent transactions, or the last decade, remains an open question.

The findings in this paper demonstrate that a well-functioning market for innovation is important for allocating patent rights efficiently *ex post*, and that taxation affects this process. As long as small innovators can appropriate part of the gains from patent trading, this market also increases their *ex ante* incentives to innovate.

References

- [1] ACIR American Council on Intergovernmental Relations (1982-1995). *Significant Features of Fiscal Federalism, Budget processes and Tax systems* (Washington, D.C.: Government Printing Office)
- [2] Angrist, Joshua (2001), "Estimation of Limited Dependent Variable Models with Dummy Endogenous Regressors: Simple Strategies for Empirical Practice," *Journal of Business & Economics Statistics* 19: 2-16
- [3] Angrist, Joshua and Guido Imbens (1995), "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity," *Journal of the American Statistical Association* 90: 431-442
- [4] Angrist, Joshua D. and Jörn-Steffen Pischke (2009), *Mostly Harmless Econometrics: An Empiricist's Companion*, (Princeton: Princeton University Press)
- [5] Arora, Ashish, Andrea Fosfuri and Alfonso Gambardella (2001), *Markets for Technology: The Economics of Innovation and Corporate Strategy* (Cambridge, MA: MIT Press)
- [6] Arora, Ashish and Marco Ceccagnoli (2006), "Patent Protection, Complementary Assets and Firms' Incentives for Technology Licensing," *Management Science* 52: 293-308
- [7] Arrow, Kenneth (1962), "Economic Welfare and the Allocation of Resources for Inventions," in Richard Nelson, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors* (Princeton: Princeton University Press): 609-625
- [8] Azoulay, Pierre (2004), "Capturing Knowledge Within and Across Firm Boundaries: Evidence from Clinical Development," *American Economic Review* 94: 1591-1612.
- [9] Bertrand, Marianne, Ester Duflo and Sendhil Mullainathan (2004), "How Much Should We Trust Differences-in-Differences Estimates? *The Quarterly Journal of Economics* 119: 249-275.
- [10] Bessen, James and Michael J. Meurer (2008), *Patent Failure* (Princeton: Princeton University Press)
- [11] Bessen, James, Michael J. Meurer and Jennifer Ford (2011), "The Private and Social Costs of Patent Trolls," Working Paper
- [12] Breschi Stefano and Francesco Lissoni (2009), "Mobility of Skilled Workers and Co-invention Networks: An Anatomy of Localized Knowledge Flows", *Journal of Economic Geography* 9: 439-468.

- [13] Bloom, Nick, Mark Schankerman and John Van Reenen (2012), “Identifying Technology Spillovers and Product Market Rivalry,” LSE Working Paper, revised
- [14] Carneiro, Pedro, James Heckman and Edward Vytlacil (2010) “Evaluating Marginal Policy Changes and the Average Effect of Treatment for Individuals at the Margin,” *Econometrica* 78: 377-394
- [15] Cenicola Mark (2012), “Lessons in Lead Valuation from Facebook’s Acquisition Strategy,” Forbes.com
- [16] Chari, Varadarajan, Mikhail Golosov and Aleh Tsyvinski (2005), “Business Start-ups, the Lock-in Effect, and Capital Gains Taxation,” Yale University Working Paper
- [17] Chien Colleen (2011), “Predicting Patent Litigation,” *Texas Law Review* 90: 283-295
- [18] Cockburn, Iain, Megan MacGarvie and Elisabeth Mueller (2010), “Patent Thickets, Licensing and Innovative Performance,” *Industrial and Corporate Change* 19: 899-925
- [19] Cohen, Julie and Mark Lemley (2001), “Patent Scope and Innovation in the Software Industry,” *California Law Review* 89: 1-57
- [20] CSG Council of State Governments (1996-2000), *Book of the States*, Washington DC.
- [21] Desai, Mihir and William Gentry (2004), “The Character and Determinants of Corporate Capital Gains,” in James Poterba, ed., *Tax Policy and the Economy* (Cambridge, MA: MIT Press): 1-36
- [22] Doyle John, (2007), “ Child Protection and Child Outcomes: Measuring the Effects of Foster Care,” *American Economic Review* 97: 1583-1610.
- [23] Dykeman, David and Daniel Kopko (2004), “Recording Patent License Agreements in the USPTO,” *Intellectual Property Today*, August Volume:18-19
- [24] Federal Trade Commission (2003), “The Evolving IP Marketplace: Aligning Patent Notice and Remedies with Competition,” (Washington D.C.: Government Printing Office)
- [25] Feenberg, Daniel Richard and Elizabeth Coutts (1993), “An Introduction to the TAXSIM Model,” *Journal of Policy Analysis and Management* 12: 189-194
- [26] Galasso Alberto (2012), “Broad Cross-License Negotiations,” *Journal of Economics and Management Strategy* 21: 873-911.
- [27] Galasso, Alberto and Mark Schankerman (2010), “Patent Thickets, Courts and the Market for Innovation”, *RAND Journal of Economics* 41: 472-503

- [28] Gans, Joshua and Scott Stern (2000), "Incumbency and R&D Incentives: Licensing the Gale of Creative Destruction" *Journal of Economics and Management Strategy* 9: 485-511
- [29] Gans, Joshua, David Hsu and Scott Stern (2002), "When Does Start-Up Innovation Spur the Gale of Creative Destruction?" *RAND Journal of Economics* 33: 571-586
- [30] Gentry, William (2010), "Capital Gains Taxation and Entrepreneurship," Working paper, Williams College
- [31] Graham, John R., (1996), "Proxies for the Corporate Marginal Tax Rate," *Journal of Financial Economics* 42: 187-221.
- [32] Graham, John R. and Lillian Mills, (2008), "Simulating Marginal Tax Rates Using Tax Return Data," *Journal of Accounting and Economics* 46: 366-388.
- [33] Hall, Bronwyn, Adam Jaffe and Manuel Trajtenberg (2001), "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," NBER Working Paper 8498
- [34] Hall, Bronwyn and Megan Mcgarvie (2010), "The Private Value of Software Patents," *Research Policy* 39: 994-1009.
- [35] Hall, Bronwyn and Rosemarie Ziedonis (2001), "The Patent Paradox Revisited: An Empirical Study of Patenting in the US Semiconductor Industry, 1979-95," *RAND Journal of Economics*, 32(1) : 101-128
- [36] Heckman, James, Hidehiko Ichimura, Jeffrey Smith and Petra Todd (1998), "Characterizing Selection Bias Using Experimental Data," *Econometrica* 66: 1017-1098
- [37] Heckman, James and Edward Vytlacil (1999), "Local Instrumental Variables and Latent Variable Models for Identifying and Bounding Treatment Effects," *Proceedings of the National Academy of Sciences* 96: 4730-4734
- [38] Imbens, Guido and Angrist Joshua (1994), "Identification and Estimation of Local Average Treatment Effects," *Econometrica* 62: 467-476
- [39] Jenkins Stephen P.(1998), "Discrete time proportional hazards regression," *Stata Technical Bulletin* 7: 22-32.
- [40] Lanjouw, Jean O. and Josh Lerner (2001), "Tilting the Table? The Predatory Use of Preliminary Injunctions," *The Journal of Law and Economics*, 44(2): 573-603
- [41] Lanjouw, Jean and Mark Schankerman (2001), "Characteristics of Patent Litigation: A Window on Competition," *RAND Journal of Economics*, 32: 129-151

- [42] Lanjouw, Jean and Mark Schankerman (2004), "Protecting Intellectual Property Rights: Are Small Firms Handicapped?" *Journal of Law and Economics*, 47: 45-74
- [43] Lemley, Mark and Carl Shapiro (2007), "Patent Hold-Up and Royalty Stacking," *Texas Law Review* 85: 1991-2049
- [44] Lileeva, Alla and Daniel Trefler (2010), "Improved Access to Foreign Markets Raises Plant-Level Productivity ... for Some Plants", *Quarterly Journal of Economics* 125: 1051-1099
- [45] Mann, Ronald J. (2005), "Do Patents Facilitate Financing in the Software Industry?" *Texas Law Review* 83: 961-973
- [46] Maine, Jeffrey and Xuan-Thao Nguyen (2003), *Intellectual Property Taxation: Transaction and Litigation Issues* (Washington D.C.: The Bureau of National Affairs Inc.)
- [47] Maine, Jeffrey and Xuan-Thao Nguyen (2010), "Equity and Efficiency in Intellectual Property Taxation," *Brooklyn Law Review* 76: 1-63
- [48] Newey, Whitney K. (1990), "Efficient Instrumental Variables Estimation of Nonlinear Models," *Econometrica* 58: 809-837.
- [49] Resis, Robert (2006), "History of the Patent Troll and Lessons Learned," *Intellectual Property Litigation* 17: 1-6
- [50] Rivers, R. Douglas and Quang H. Vuong (1998), "Limited Information Estimators and Exogeneity Tests for Simultaneous Probit Model," *Journal of Econometrics*, 39: 347-65
- [51] Shapiro, Carl (2001), "Navigating the Patent Thicket: Cross Licenses, Patent Pools and Standard Setting," *Innovation Policy and the Economy*, volume 1, Adam Jaffe, Joshua Lerner, and Scott Stern, eds. (Boston: MIT Press)
- [52] Serrano, Carlos J. (2010), "The Dynamics of the Transfer and Renewal of Patents," *RAND Journal of Economics* 41: 686-708
- [53] Sorenson, Olav and Toby E. Stuart (2001), "Syndication Networks and the Spatial Distribution of Venture Capital Investments," *American Journal of Sociology*, 106: 1546-1588
- [54] Thoma, Grid, Salvatore Torrisi, Alfonso Gambardella, Dominique Guellec, Bronwyn Hall and Dietmar Harhoff (2010), "Harmonizing and Combining Large Datasets - An Application to Firm-Level Patent and Accounting Data," NBER Working Paper 15851.
- [55] Teece, David (1986), "Profiting from Technological Innovation: Implications for Integration, Collaboration, Licensing, and Public Policy," *Research Policy*, 15: 285-305

- [56] Turner, John (2011), “Patent Thickets, Trolls and Unproductive Entrepreneurship,” working paper
- [57] Williams, Heidi (2010), “Intellectual Property Rights and Innovation: Evidence from the Human Genome,” NBER working paper 16213
- [58] Wooldridge, Jeffrey M. (2002), *Econometric Analysis of Cross Section and Panel Data* (Boston: MIT Press)

TABLE 1. Summary Statistics

Panel A. Patent Trade and Litigation							
	Patents Not Traded			Patents Traded		Total	
	N	col. perc.	N	col. perc.	N	col. perc.	
Patents Not Litigated	284,281	99.49	13,038	95.82	297,319	99.31	
	row perc.	95.61	4.39				
Patents Litigated	1,468	0.51	569	4.18	2,037	0.69	
	row perc.	72.07	27.93				
Total	285,749		13,607		299,356		
	row perc.	95.45	4.55				

Panel B. Capital Gains and Income Tax Rates

Period	Capital Gains Tax Rates				Income Tax Rates				Corporate Tax Rates			
	Mean	Std. Dev.	Min	Max	Mean	Std. Dev.	Min	Max	Mean	Std. Dev.	Min	Max
1982-1986	21.4	1.2	20	27	52.7	1.9	50	56.7	52.6	3.1	46	58
1987-1991	31.6	2.1	28	37	34.4	4.5	28	44.6	46.5	3.3	39	54.5
1992-1996	32.4	1.9	28.9	37	42.4	3.9	31.9	48.1	45.8	3.1	39	51.2
1997-2001	26.9	5.6	21.2	40.3	43.9	1.8	40.3	46.9	45.8	2.9	39	51

NOTES: Capital gains tax rate is the sum of federal and state capital gain tax rates in state of first inventor. Income and Corporate tax rates are defined analogously.

TABLE 2. Patent Citations and Trade and Litigation Rates

	Individually Owned	Corporate Patents	
		Small Firms	All Other Firms
Panel A. Patent Citations Received			
All Patents	5.9	7.3	7.3
Traded Patents	10.3	10.8	9.3
Panel B. Rate of Trade and Litigation (in percentage)			
Trade rate	4.7	16.1	12.2
Trade rate weighted by cites	32.2	35.4	18.3
Annual Litigation rate	1.2	1.6	0.3
Number of Patents	204,592	236,776	496,284

Notes: Our sample includes all patents granted to U.S. individuals and corporations for the period 1985-2000. Patent Citations Received: Mean number of the truncation-adjusted forward cites (see Hall et al., 2001). The trade rate is the proportion of patents that were traded at least once in the sample. The trade weighted by cites is a sample mean computed assigning weights to patents based on their number of patent citations received. The annual litigation rate is the predicted probability that a patent will be involved in litigation in a given year. This probability was obtained using the coefficients of a Probit model of the decision to litigate a patent controlling for age dummies, technology dummies, year dummies, and individual and type of innovator dummies. patent, and evaluated at the sample means of the covariates. Technology dummies represent the 36 technology sub-categories defined in Hall et al. (2001). Small Firms are those with five or fewer patents in a given year (see Serrano, 2010 for details). All Other Innovator is the residual category.

TABLE 3: Trade and Litigation- Correlations

Estimation Method	1	2	3	4
Dependent Variable	OLS Litigation Dummy	OLS Litigation Dummy	OLS Litigation Dummy	OLS Litigation Dummy
NewOwner x 10	0.039*** (0.003)	-0.025*** (0.004)	-0.019*** (0.004)	-0.056 (0.153)
Patent Fixed Effects	NO	YES	YES	YES
Age Dummies	NO	NO	YES	YES
Time Period Dummies	NO	NO	YES	YES
Sample	Entire Sample	Entire Sample	Entire Sample	Litigated and Traded Patents
Patents	299,356	299,356	299,356	569
Observations	2,436,649	2,436,649	2,436,649	6,810

NOTES: Standard errors are clustered at patent level and are reported in parentheses. Statistical significance: * 10 percent, ** 5 percent, *** 1 percent. Litigation Dummy = 1 if the patent is involved in at least one case at that age; NewOwner = 1 when the patent changes ownership for the first time and remains equal to one for the remaining life of the patent. Time Period Dummies: before 1986, 1986-1990, 1991-1995, after 1995.

TABLE 4. Impact of Taxes on Patent Trading

Estimation Method Dependent Variable	1	2	3	4
	Probit Trade	OLS Trade	Proportional Hazard Trade	Probit Trade
	Mar. Eff. x 10 ³	Coefficients x 10 ³	Hazard Ratios	Mar. Eff. x 10 ²
Capital Gains Tax Rate	-0.204** (0.090)	-0.313*** (0.121)	0.953*** (0.014)	-0.590** (0.247)
Income Tax Rate	0.133** (0.066)	0.196** (0.090)	1.032** (0.013)	0.196 (0.214)
Corporate Tax Rate	-0.063** (0.031)	-0.147*** (0.046)	0.984*** (0.006)	-1.013* (0.593)
Patent Citations Received	0.061*** (0.001)	0.187*** (0.042)	1.011*** (0.001)	0.015 (0.019)
Patent Generality	0.193 (0.150)	0.052 (0.220)	1.071 (0.047)	5.576*** (1.811)
Sample	Entire Sample until Traded	Entire Sample until Traded	Entire Sample until Traded	Litigated and Traded Patents until Traded
Patents	299,356	299,356	299,356	569
Observations	2112507	2112507	2112507	3025

NOTES: All regressions include age, year and technology field dummies. Column 1 and 4 are Probit models. Column 2 is a linear probability. Column 3 is a proportional hazard model (estimates < 1 indicate a negative effect). Standard errors in parenthesis are block bootstrapped at the state level. Statistical significance: *10 percent, **5 percent, *** 1 percent. Trade= 1 when the patent changes ownership for the first time. Capital Gains Tax Rate: sum of federal and state capital gain tax rates in the state of first inventor. Income Tax Rate: sum of federal and state income tax rates in the state of first inventor. Corporate Tax Rate: weighted average of state corporate taxes with weights constructed using share of patenting in the technology area. Patent Citations Received: truncation-adjusted forward cites. Technology Dummies represent the 36 technology sub-categories defined in Hall et al. (2001).

TABLE 5. Impact of Taxes on Corporate Patent Trading

Estimation Method Dependent Variable	1	2	3	4
	OLS Trade	OLS Trade	OLS Trade	OLS Trade
	Coefficients x 10 ³	Coefficients x 10 ³	Coefficients x 10 ³	Coefficients x 10 ³
Corporate Tax Rate	-0.882** (0.340)	-0.883** (0.340)	-0.883** (0.344)	-0.875** (0.340)
Capital Gains Tax Rate		0.107 (1.512)	0.104 (1.514)	0.269 (0.712)
Income Tax Rate		-0.565 (1.540)	-0.565 (1.540)	-0.110 (0.824)
Patent Citations Received			0.019 (0.033)	
Patents	150,511	150,511	150,511	150,511
Observations	806,366	806,366	806,366	806,366

NOTES: All regressions include year, age, technology field and firm fixed effects and controls for employees, patent portfolio, cash and assets. Standard errors are clustered at firm level. Statistical significance: *10 percent, **5 percent, *** 1 percent. Trade= 1 when the patent changes ownership for the first time. Capital Gains Tax Rate: sum of federal and state capital gain tax rates in the state of first inventor (columns 2 and 3) and in the state of corporate headquarter in column 4. Income Tax Rate: sum of federal and state income tax rates in the state of first inventor (columns 2 and 3) and in the state of corporate headquarter in column 4. Corporate Tax Rate: see Graham (1996). Patent Citations Received: truncation-adjusted forward cites. Technology dummies represent the 36 technology sub-categories defined in Hall et al. (2001).

TABLE 6. Taxes and Patent Trade - Fixed Effect Regressions

Estimation Method Dependent Variable	1	2	3	4
	OLS Trade	OLS Trade	OLS Trade	OLS Trade
	Coefficients x 10 ³	Coefficients x 10 ³	Coefficients x 10 ³	Coefficients x 10 ³
Capital Gains Tax Rate	-0.117*** (0.027)	-0.124*** (0.038)	-0.013 (0.210)	0.090 (1.523)
Corporate Tax Rate		-0.053 (0.116)	-0.021 (0.743)	-0.564* (0.340)
Income Tax Rate		0.010 (0.027)	0.016 (0.196)	-0.493 (2.290)
Age dummies	YES	YES	YES	YES
Time Period Dummies	YES	YES	YES	YES
Patent Fixed Effects	YES	YES	YES	YES
Sample	Individual Patents	Individual Patents	All Corporate Patents	Large Corporation Patents
Patents	299,356	299,356	733,060	150,118
Observations	2,112,507	2,112,507	4,281,779	806,366

NOTES: Standard errors are block boot-strapped at state level. Statistical significance: *10 percent, **5 percent, *** 1 percent. Trade= 1 when the patent changes ownership for the first time. Capital Gains Tax Rate: sum of federal and state capital gain tax rates in the state of first inventor. Income Tax Rate: sum of federal and state income tax rates in the state of first inventor. Corporate Tax Rate: in columns (2)-(3) weighted average of state corporate taxes with weights constructed using share of patenting in the technology area. Firm specific marginal tax rate in column (4) described in Graham (1996).

TABLE 7. Impact of Trade on Litigation - Instrumental Variable Estimation

Estimation Method	1	2	3	4	5
Dependent Variable	2SLS Litigation Dummy	2SLS Litigation Dummy	2SLS Litigation Dummy	2SLS Litigation Dummy	2SLS Litigation Dummy
NewOwner (Instrumented)	-0.012*** (0.005)	-0.011*** (0.004)	-0.013*** (0.004)	-0.159*** (0.061)	-0.033*** (0.011)
Sample	Entire Sample	Entire Sample	Entire Sample	Traded and Litigated Patents	Entire Sample
Patents	299356	299356	299356	569	299356
Observations	2436649	2436649	2436649	6810	2436649
INSTRUMENT	\hat{P} estimated with probit	\hat{P} estimated with OLS	Predicted Survival from Hazard	\hat{P} estimated with probit	Large Tax Difference Dummy

NOTES: All regressions include age, period and patent fixed effects. Standard errors in parenthesis are clustered at the patent level. Statistical significance: * 10 percent, ** 5 percent, *** 1 percent. Litigation Dummy = 1 if the patent is involved in at least one case at that age; NewOwner = 1 when the patent changes ownership for the first time and remains equal to one for the remaining life of the patent. Time Period Dummies: before 1986, 1986-1990, 1991-1995, after 1995. \hat{P} is the estimated probability of not being owned by the original inventor. Large Tax Difference Dummy = 1 if difference between income tax rates and capital gain tax rates is above the 75th percentile.

TABLE 8. The Roles of Buyer Portfolio Size and Patent Fit - Instrumental Variable Estimation

Estimation Method	1	2	3	4	5
Dependent Variable	2SLS Litigation Dummy	2SLS Litigation Dummy	2SLS Litigation Dummy	2SLS Litigation Dummy	2SLS Litigation Dummy
NewOwner	0.338*** (0.120)	-0.429* (0.236)	-0.278*** (0.081)	-0.383 (0.262)	-0.238*** (0.081)
NewOwner x LargeBuyer					-0.365* (0.196)
NewOwner x TechFit					0.461*** (0.137)
Sample	Trades to small buyers and high patent fit	Trades to large buyers and low patent fit	Trades to small buyers and low patent fit	Trades to large buyers and high patent fit	All traded and Litigated Patents
Observations	1585	507	4361	357	6810
Patents	116	47	382	24	569

NOTES: All regressions include age, period and patent fixed effects. Standard errors in parenthesis are clustered at the patent level. Statistical significance: * 10 percent, ** 5 percent, *** 1 percent. Litigation Dummy = 1 if the patent is involved in at least one case at that age; NewOwner = 1 when the patent changes ownership for the first time and remains equal to one for the remaining life of the patent. Time Period Dummies: before 1986, 1986-1990, 1991-1995, after 1995. LargeBuyer=1 if acquirer obtained more than 8 patents in the 20 years before trade. TechFit=1 if acquired patent belongs to technology sub-category in which buyer has more patents. NewOwner and its interactions are instrumented by the Probit estimates of the probability of not being owned by the original inventor.

TABLE 9. Portfolio Size and Patent Fit - Larger Samples

	1	2	3	4
Estimation Method	2SLS	2SLS	2SLS	2SLS
Dependent Variable	Litigation Dummy	Litigation Dummy	Litigation Dummy	Litigation Dummy
NewOwner	-0.040*** (0.013)	-0.021*** (0.008)	-0.012*** (0.005)	-0.009** (0.004)
NewOwner x LargeBuyer	-0.075** (0.036)	-0.079** (0.032)	-0.086*** (0.033)	-0.092*** (0.033)
NewOwner x TechFit	0.024*** (0.006)	0.013*** (0.004)	0.007** (0.003)	0.006** (0.003)
Sample	Matching Score=100%	Matching Score>75%	Matching Score> 50 %	Matching Score> 30 %
Observations	29161	46333	74457	89585
Patents	2570	4139	6654	8123

NOTES: Standard errors in parenthesis are clustered at the patent level. All regressions include age, period and patent fixed effects. Statistical significance: * 10 percent, ** 5 percent, *** 1 percent. Litigation Dummy = 1 if the patent is involved in at least one case at that age; NewOwner = 1 when the patent changes ownership for the first time and remains equal to one for the remaining life of the patent. Time Period Dummies: before 1986, 1986-1990, 1991-1995, after 1995. LargeBuyer=1 if acquirer was issued in the 20 years before trade more patents than the upper decile in the buyer portfolio size distribution at each sample. TechFit=1 if acquired patent belongs to technology sub-category in which buyer has more patents. NewOwner and its interactions are instrumented by the Probit estimates of the probability of not being owned by the original inventor.

TABLE 10: Simulated Effects of Capital Gains Taxes on Patent Trade and Litigation

Scenario	Capital Gains Taxes (in percent)	Traded Patents per 1000 patents	Predicted Suits per 1000 patents
Baseline	29.2	56.9	35.8
Low Tax	20.0	92.5	23.1
High Tax	40.0	30.9	45.5
Capital Gains Tax = Income Tax	42.6	26.4	47.1
Capital Gains Tax = Income Tax	29.2	35.1	44.1

FIGURE 1. Trade and Likelihood of Litigation







FIGURE 2. Impact of Tax Changes on Patent Trade



NOTES: These figures plot coefficient estimates (and 95 percent confidence intervals) from the diff-in-diff analysis described in Section 3. The x-axis plots years relative to a tax change (being zero the year of the tax change). The coefficients are estimates from OLS regressions with patent age effects, calendar year effects, technology sub-categories effects, an individually owned patent dummy, a dummy indicating whether a patent is subject to renewal fees, a set of U.S. state dummies, the level of corporate and income taxes (sum of state and federal rates), and a set of dummy variables for the length of the lag before and after the tax change (the year prior to the tax change is taken as the reference year). The standard errors are robust to heteroskedasticity. The dependent variable is whether or not a patent was traded in a given year, and the unit of analysis is patent-year.

FIGURE 3. Marginal Treatment Effect of Trade on Litigation



Appendices for "Trading and Enforcing Patent Rights"

Alberto Galasso
University of Toronto

Mark Schankerman
London School of Economics and CEPR

Carlos Serrano
University of Toronto and NBER

Appendix 1: Extended Model with Infringing Buyer

In this appendix we extend the model developed in Section 2 by allowing the buyer to be a potential infringer. The individual, A , owns the patent and a firm, B , is willing to acquire the patent from the individual. If the individual does not sell the patent, he may obtain product market profits equal to π^A . If the patent is acquired by the firm, it generates product market profits equal to π^B . Both A and B may face an infringing action by a third party, firm C . The infringement action occurs with probability β . If the infringement action takes place, the patent owner chooses whether to litigate or settle the dispute. With litigation the patent owner $i = \{A, B\}$ sustains litigation costs L_i to secure product market profits. To settle the owner gives up a fraction $(1 - \theta_{iC})$ of the profits to firm C . There will be litigation if

$$\pi^i - L_i \geq \theta_{iC}\pi^i + \varepsilon$$

that occurs with probability

$$\Pr \{ \varepsilon \leq \pi^i(1 - \theta_{iC}) - L_i \}.$$

At the time of the negotiation, firm B may also be potentially infringing the patent. To capture this possibility, we assume that a dispute between the two parties arises with probability ρ . Settlement of the dispute with B involves a loss of a fraction of profits equal to $(1 - \theta_{AB})$.

We refer to the vector $e^{iC} = (L_i, \theta_{iC})$ as the "enforcement" vector of owner $i = \{A, B\}$. Litigation between i and C takes place with probability $\Omega(\pi^i, e^{iC}) = \beta \Pr \{ \varepsilon \leq \pi^i(1 - \theta_{iC}) - L_i \}$. We define $e^{AB} = (L_A, \theta_{AB})$ and indicate with $\Omega(\pi^A, e^{AB})$ the probability that litigation is observed between A and B . In this extended setting, trade between A and B occurs if the following condition is satisfied:

$$\begin{aligned} & [(1 - \Delta_{BC})\pi^B - \Omega(\pi^B, e^{BC})L_B] (1 - \tau^C) (1 - \tau^G) \geq \\ & [(1 - \Delta_{AB} - \Delta_{AC} - \Delta_{ABC})\pi^A - (\Omega(\pi^A, e^{AC}) + \Omega(\pi^A, e^{AB}))L_A] (1 - \tau^I) \end{aligned}$$

where Δ_{AB} is the fraction of profits that are expected to be lost in the event of settlement between A and B (but not C) and Δ_{ABC} is the fraction lost if A settles with both parties. This inequality shows that also in this extended model the condition required to have trade becomes less stringent with an increase in τ^I and a decrease in τ^G .

We now investigate how the level of litigation is affected by a change in patent ownership. If individual A does not sell the patent, the expected number of suits in which the patent is involved is equal to $E(\text{Litigation}|T=0) = \Omega(\pi^A, e^{AB}) + \Omega(\pi^A, e^{AC})$. If trade takes place the expected number of suits is $E(\text{Litigation}|T=1) = \Omega(\pi^B, e^{BC})$. Therefore:

$$E(\text{Litigation}|T=1) - E(\text{Litigation}|T=0) = \tag{1}$$

$$-\Omega(\pi^A, e^{AB}) - [(\Omega(\pi^A, e^{AC}) - \Omega(\pi^B, e^{BC}))]. \tag{2}$$

Formula (1) shows how changes in patent ownership affect litigation rates through two different channels. First, as in the baseline model trade reallocates the patent to an entity that has a different likelihood to resort to courts. This "reallocation effect" is positive or negative depending on the difference $\Omega(\pi^A, e^{AC}) - \Omega(\pi^B, e^{BC})$ that in turns depends on magnitude of the "product market" and the "enforcement gains". Second, in this extended model by trading with the firm the individual substitutes litigation with trade. This "replacement" effect is captured by the term $-\Omega(\pi^A, e^{AB})$ that is always negative.

Appendix 2: Details on Non-Parametric Estimation

This section describes the details of the non-parametric estimation of the marginal treatment effect. Our approach is based on a multistep, non-parametric procedure of Heckman, Ichimura, Smith and Todd (1998) and Carneiro, Heckman, and Vytlacil (2010). We extend their procedure to a panel data setting in order to account for individual/patent fixed effects.

The first part of the procedure involves estimating the litigation equation non-parametrically. This is a non-parametric counterpart of the IV estimates in Table 4.

1. STEP 1. Regress each of the variables in the vector of covariates X on \hat{P} using local linear regression. In our setting this involves running multiple regressions. In particular, we run a regression for each age dummy on \hat{P} , a regression for each calendar dummy on \hat{P} , and finally another regression of income tax rates on \hat{P} . The regressions were run in STATA 10 using the command *lpoly*.
2. STEP 2. Let $\hat{\varepsilon}_X$ be the residuals and vector of residuals from the regression in step 1. Regress L on $\hat{\varepsilon}_X$ using OLS with patent fixed effects in order to obtain an estimate of the vector β_0 .
3. STEP 3. Let $\hat{\varepsilon}$ be an estimate of the residual from the previous OLS regression (accounting for the patent fixed effects). This is an estimate of $\beta_1\hat{P} + E[(\bar{\alpha} + \psi_i)T|\hat{P}]$. Regressing $\hat{\varepsilon}$ on \hat{P} using local linear regression allow us to obtain a non-parametric estimate of $\hat{\varepsilon}(\hat{P})$.

Putting all this together, we construct an estimate of $E[L|\hat{P}]$, $\widehat{E[L|\hat{P}]} = \widehat{\beta_0}X + \widehat{\mu}_i + \widehat{\varepsilon}(\hat{P})$.

The second part of the procedure involves numerically differentiating $E[L|\hat{P}]$. To do so, we divide observations into groups, based either on the deciles of the distribution of \hat{P} or the absolute value of \hat{P} . Recall that the variable component of $\widehat{E[L|\hat{P}]}$ with respect to \hat{P} is $\widehat{\varepsilon}(\hat{P})$. The mean of $\widehat{\varepsilon}(\hat{P})$ was calculated for each of these groups. The derivative of $\widehat{\varepsilon}(\hat{P})$ were obtained by finite differencing across neighboring groups. The confident intervals of the marginal treatment effects were obtained using 50 bootstrap iterations (seed = 123 in STATA 10).

Table A1. Capital Gains Tax Rate Policy Changes in US States

US State	Capital Gains Tax Rate (US States)					
	Tax hike			Tax cut		
	Changes in Taxes	Tax Change	Tax Change	Changes in Tax Rates	Tax Change	Tax Change
		(Percentage points)	(%)		(Percentage points)	(%)
California	2	2.25	0.2772	1	1.8	0.1544
New York	3	0.69	0.0948	6	0.36	0.0463
Florida	0	-	-	0	-	-
Texas	0	-	-	0	-	-
Illinois	3	0.33	0.1303	2	0.25	0.0871
Michigan	2	1.63	0.8725	6	0.18	0.0657
Pennsylvania	3	0.37	0.1621	5	0.1	0.0401
New Jersey	3	2	0.9667	3	0.21	0.0308
Ohio	8	0.83	0.2272	5	0.3	0.0518
Massachusetts	3	0.82	0.1757	3	0.59	0.1026
Washington	0	-	-	0	-	-
Maryland	3	1.36	0.4631	4	0.3	0.0511
Minnesota	4	1.69	0.5215	8	0.44	0.0788
Arizona	3	1.74	0.932	6	0.3	0.0491
Connecticut	1	4.2	1.5	2	1.25	0.1936
Colorado	3	1.16	0.6436	2	2.5	0.5109
Virginia	2	1.76	0.757	0	-	-
Wisconsin	0	-	-	6	1.88	0.1974
North Carolina	1	0.75	0.1026	0	-	-
Georgia	2	1.71	0.6942	0	-	-
Oregon	5	1.2	0.2276	2	0.6	0.0871
Indiana	3	0.88	0.7694	0	-	-
Missouri	3	1.61	0.9826	3	0.05	0.0114
Louisiana	4	0.99	0.7	2	0.16	0.0387
Oklahoma	5	1.17	0.3798	1	0.48	0.0713
Tennessee	0	-	-	0	-	-
Utah	4	1.73	0.985	4	0.46	0.0642
South Carolina	2	2.1	0.7512	4	0.92	0.1762
Iowa	5	1.06	0.3545	3	0.21	0.0276
Kansas	6	1.42	0.6113	3	1.29	0.2568
Nevada	0	-	-	0	-	-
Alabama	1	0.67	0.2	4	0.19	0.0483
Kentucky	2	2.48	1.486	1	0.03	0.0068
Nebraska	7	0.7	0.1648	3	0.22	0.0455
New Hampshire	0	-	-	0	-	-
New Mexico	3	1.94	0.6729	1	0.26	0.0322
Idaho	2	2.64	0.8721	2	0.2	0.0243
Montana	5	0.73	0.266	5	0.26	0.0672
Arkansas	0	-	-	2	0.97	0.1498
Mississippi	2	0.58	0.1452	1	0.07	0.0147
Rhode Island	7	1.41	0.2834	7	0.62	0.0859
Hawaii	2	2.85	0.6796	2	0.65	0.0694
Maine	4	2.05	0.4406	4	0.86	0.0856
West Virginia	3	1.16	0.2785	2	0.24	0.0484
Washington DC	2	2.84	0.6411	2	0.35	0.0354
North Dakota	7	0.59	0.2602	2	0.88	0.1932
Delaware	2	2.57	0.6489	7	0.63	0.0964
Vermont	6	1.41	0.2205	4	1.14	0.1101
Wyoming	0	-	-	0	-	-
Alaska	0	-	-	0	-	-
South Dakota	0	-	-	0	-	-
Average	2.71	1.50	0.5385	2.55	0.58	0.0923

NOTE: Capital gains tax rates for the period 1982-2001 are obtained from the NBER TAXIM dataset. Changes in Tax Rates is equal to the number of policy changes in the capital gains tax rate for the period 1982-2001. Tax changes is equal to the average tax change in percentage points (or growth rates) among all the changes in the capital gains tax rate in a State. US states with zero tax changes in the state level capital gains tax rate are states with no state level capital gains tax rate (e.g., Florida). The row labeled "Average" is the average across US states. U.S. States are ranked according to the number of individually-owned patents granted in the state.

Table A2. First Stage Regressions

	1	2	3
Dependent Variable	NewOwner	NewOwner	NewOwner
\hat{P} from probit	0.833*** (0.030)		
\hat{P} from OLS		0.816*** (0.025)	
Predicted survival from hazard			-0.875*** (0.030)
Weak Identification Test	F=3327.8 p<0.001	F= 4958.2 p<0.001	F= 3608.6 p<0.001
Sample	Entire Sample	Entire Sample	Entire Sample
Patents	299356	299356	299356
Observations	2436649	2436649	2436649

NOTES: All regressions include age , period and patent fixed effects. Standard errors in parentheses are clustered at the patent level. Statistical significance: * 10 percent, ** 5 percent, *** 1 percent. NewOwner = 1 when the patent changes ownership for the first time and remains equal to one for the remaining life of the patent. Time Period Dummies: before 1986, 1986-1990, 1991-1995, after 1995. P^ is the estimated probability of not being owned by the original inventor.

TABLE A3. The Roles of Buyer Portfolio Size and Patent Fit - Robustness

	1	2	3	4
Estimation Method	2SLS	2SLS	2SLS	2SLS
Dependent Variable	Litigation Dummy	Litigation Dummy	Litigation Dummy	Litigation Dummy
	<i>LargeBuyer =1 if more than 12 patents</i>	<i>TechFit=0 if trade among individuals</i>	<i>TechFit constructed with narrow technology classes</i>	<i>TechFit defined using patent citations</i>
NewOwner	-0.221*** (0.08)	-0.130** (0.06)	-0.190** (0.08)	-0.227*** (0.08)
NewOwner x LargeBuyer	-0.394* (0.22)	-0.338* (0.19)	-0.343* (0.19)	-0.401** (0.20)
NewOwner x TechFit	0.450*** (0.13)	0.133*** (0.04)	0.383*** (0.14)	0.484*** (0.13)
Sample	Traded and Litigated Patents	Traded and Litigated Patents	Traded and Litigated Patents	Traded and Litigated Patents
Observations	6810	6810	6810	6810
Patents	569	569	569	569

NOTES: Standard errors in parentheses are clustered at patent level. All regressions include age, period and patent fixed effects. Statistical significance: * 10 percent, ** 5 percent, *** 1 percent. Litigation Dummy = 1 if the patent is involved in at least one case at that age; NewOwner = 1 when the patent changes ownership for the first time and remains equal to one for the remaining life of the patent. Time Period Dummies: before 1986, 1986-1990, 1991-1995, after 1995. In columns 2 to 4 LargeBuyer=1 if acquirer obtained more than 8 patents in the 20 years before trade. In columns 1 and 2 TechFit=1 if acquired patent belongs to technology sub-category in which buyer has more patents. In column 3 TechFit constructed using USPTO patent nclasses. In column 4 TechFit=1 if either the acquired patent cites one of the patents of the buyer or if the patents of the buyer cite the acquired patent. NewOwner and its interactions are instrumented by the Probit estimates of the probability of not being owned by the original inventor.

FIGURE A1. Marginal Treatment Effect – Litigated and Traded Patents

