Estimating the Gains from Trade in the Market for Patent Rights

Carlos J. Serrano[1]

UPF, Barcelona GSE, University of Toronto, and NBER

This draft: January, 2013

[1] I would like to thank Thomas Holmes, Sam Kortum, Victor Aguirregabiria, Zvi Eckstein, Nicolas Figueroa, Suqin Ge, Flor Gragera de Leon, Hugo Hopenhayn, Chris Laincz, Katherine Lande, Rob McMillan, Matt Mitchell, Valeriu Omer, Ariel Pakes, and Andy Skrzypacz. In addition, I thank Concordia University, Universite de Montreal, University of Toronto, University of Iowa, Federal Reserve Bank of Minneapolis, University of British Columbia - Sauder Business School, University of Texas, Dallas - Management School, Federal Reserve Board, London School of Economics, Universitat Autonoma de Barcelona, Harvard, Wilfrid-Laurie, IESE, Duke, and the Stanford Institute for Economic Policy Research for their comments. In addition, we gladly acknowledge Ryan Lampe's generosity for providing access to the raw data of Ocean Tomo patent auctions. All errors are mine. An earlier draft of this paper went under the title "The Market for Intellectual Property: Evidence from the Transfer of Patents".

Carlos J. Serrano, Universitat Pompeu Fabra; Ramon Trias Fargas, 25-27; 08005 Barcelona; Spain. Tel: +34-93-542-1532; Email: carlos.serrano@upf.edu.

Abstract

The "market for patents" – the sale of patents– is an often discussed source of incentives to invest in R&D. This article presents and estimates a model of the transfer and renewal of patents that, under some assumptions, allows us to quantify the gains resulting from the transfer of patent rights. The gains from trade measure the benefits of reallocating the ownership of a patent from the original patentee to a new owner for whom the patent has a higher value. In addition, we study the effect that lowering transaction costs has on the proportion of patents traded and the gains from trade.

1 Introduction

The market for patent rights — the sale of patents — is an often discussed source of incentives to invest in R&D, especially for small firms being patents typically their critical assets. The market for patent rights can generate private and social gains from technological trade by facilitating the reallocation of innovation to firms that are more effective at commercializing the patented innovations (Arrow, 1962; Arrow, 1983; Teece; 1986; Arora, Fosfuri, and Gambardella, 2001a; Gans, Hsu, and Stern, 2008). Another source of private and social gains from trade is comparative advantages in patent enforcement. Patent transactions can reduce litigation costs if the market reallocates patent rights to entities that are more effective at resolving disputes over these rights without resorting to the courts (Galasso, Schankerman, and Serrano, 2011). The gains that original inventors can obtain from this market, and the corresponding incentives to invest in R&D, are greater the greater are the extra-profits that non-inventors can generate from the patented innovation.

At the same time, transactions in patent rights can adversely affect the rate and direction of innovation activity of firms, especially for those firms not directly involve in a transaction. This can occur if these patent rights get reallocated to entities that have the capacity and the incentives to exploit the patents to extract excessive royalty fees by holding up rivals and/or infringers (see e.g., Shapiro (2001) and Lemley and Shapiro (2007)).[1] This issue is at the center of recent reports by the U.S. Federal Trade Commission (FTC (2011, 2012)) analyzing the working of the market for patent rights.[2] The key to understand the overall impact of transactions in patent rights is to assess quantitatively the different ways the market for patents affects social welfare. While it is well known that many economic activities have seen the benefits of markets for tradable property rights, working examples are the market for taxi medallions, liquor licenses, pollution permits, among many others, little systematic empirical work has been undertaken examining the magnitude of the gains from trade in the market for patent rights.

This paper takes a first look at quantifying the private gains from trade in the market for patents. This is a first step towards understanding the broader policy question of whether patent transactions increase total welfare. If this private gains from trade were significant, they could potentially offset concerns about hold-up problems and excessive concentration in patent rights.

[1] Alternatively, patent transactions can also reduce the incentives to innovative if they lead to excessive concentration of patent rights (Hahn, 1984).

[2] See report on the "The Evolving IP Marketplace: Aligning Patent Notice and Remedies with Competition" by the Federal Trade Commission (March, 2011). More recently, the "Patent Assertion Entity Activities Workshop" jointly sponsored by the Department of Justice and the Federal Trade Commision (December, 2012) (http://www.ftc.gov/opp/workshops/pae/).

To do this, we develop and estimate a model of patent sales and renewals that, under some assumptions, allows us to recover the gains resulting from the sale of patents. We extend the Pakes and Schankerman's patent renewal model to allow the original patentee to sell patents to other firms (Pakes (1986) and Schankerman and Pakes (1986)). We combine data on patent sales and renewal to estimate the parameters of this model. Our method allow us to estimate the distribution of patent values and the distribution of gains from trade from selling one patent from one firm to another one. In addition, we estimate market transaction costs and study the effect that lowering transaction costs has on the proportion of patents traded and the gains from trade. Counterfactual experiments indicate that the private gains from trade account for about 10 percent of the value of the volume of the trade of patents, and that lowering transaction costs by fifty percent increases the gains from trade by an additional ten percent. Our findings also point out that the value of the volume of the trade of patents represents about 50% of the total value of all patents, and that the distribution of the gains from trade is highly skewed.

We make use of patent assignment data to identify changes in the ownership of patent rights. An interesting feature of the sale of patents —that distinguishes it from the licensing of patents— is the fact that the sale of patents is most often publicly recorded because of the legal requirement that all patent sales have to be filed with the United States Patent and Trademark Office (USPTO) in order to be legally binding (Dykeman and Kopko, 2004). This property of patent sales allows us to link these transactions to patent numbers, which can then be merged to the basic patent data that others have used: the payment of renewal fees, patent citations, the patent's issue date, a patentee identifier, etc. (Griliches, 1992; Pakes and Schankerman, 1984; Hall, Jaffe and Trajtenberg, 2001). A challenge in using patent assignment data is to distinguish changes in patent ownership from other events. To do this, we follow Serrano (2010) and conservatively drop all the assignments that appear not to be associated with an actual patent trade.[3] Serrano shows that on average 13.5 percent of all patents are traded at least once over their lifetime, and that this rate is much higher when weighted by patent citations received, and for patents of small firms.[4]

In our empirical analysis, we will focus on patents applied for and granted to small firms. We operationalize this focus by restricting attention to patents granted to firms with no more

[3] Serrano (2010) documents the dynamics of the sale and renewal of patents. He studies the effect of firm size, patent age, patent citations, patent generality, patent technology classes, whether a patent was previously traded, etc. on the decision to trade and renew a patent. For recent empirical work using patent assignment data, see Figueroa and Serrano (2011) and Galasso, Schankerman, and Serrano (2011).

[4] But the sale of patents is not important just solely for small firms. For instance, large firms such as AOL, Dell, IBM, Facebook, Google, Microsoft, etc. have recently made available for sale hundreds of their patents or acquired large bundles of patents.

than 500 employees as of the application date of the patent. There are a number of reasons that support our attention on these patents. The first one is that in the policy arena there exists special interest in the importance that the market for patents plays on the incentives to invest in R&D by small firms (FTC, 2011). Another reason is that by focusing on small firms the economic forces that we highlight will be more salient than in transactions involving the patents from larger firms.[5] At the same time, the focus on small firms allows us to parallel our empirical analysis with the existing theory on patent renewals, in which the decision making is at the patent level. Furthermore, small firms are interesting in their own right, given the importance they play in the innovation process (Arrow, 1983 and Acs and Audretsch, 1998).

We then develop a theory of patent sales and renewals. The starting point of our model is Pakes (1986) and Schankerman and Pakes (1986). They examine the problem of a patent owner deciding in each time period whether or not to pay a renewal fee, and thereby to extend the life of a patent, in a context with heterogeneity in the economic value of inventions. The distinct element of our theory is the introduction of the possibility of the arrival of opportunities for surplus-enhancing transfers, which may lead to a new owner for whom the patent has a higher value. However, to transfer the patent to a new owner involves a transaction cost. The gains from trade measure the net benefits of reallocating the ownership of a patent from the original patentee to a new owner. Therefore, whereas Pakes and Schankerman's framework has one margin —whether the patent owner should pay the fee for renewing the patent— our model has a second margin —whether a transaction cost should be covered to reallocate the patent to an alternative owner.

The parameters of the model are estimated using the simulated generalized method of moments (McFadden, 1989 and Pakes and Pollard, 1989). The empirical moments we use to identify these parameters describe both the patent life cycle properties of the trading and expiring decisions of patent owners. The first set of moments refers to the trading decision: (1) the probability that an active patent is traded at alternative ages conditional on having been previously traded; and (2) the probability that an active patent is traded at alternative ages conditional on having been previously untraded. The second set of moments relates to the expiring decision and how this decision interacts with past trading history of patents: (3) The probability that an active patent is allowed to expire at alternative renewal dates conditional on having been previously traded; and (4) the probability that an active patent is allowed to expire at alternative renewal dates conditional on having been previously untraded. The joint set of moments describes both the

[5] Serrano (2010) reports that patents granted to large corporations may not only be traded for the technology that they represent, but also as a result of large acquisitions that may be pursued to increase the buyer's market share in a particular product market, etc.

history of the trading and the renewal decision of patent owners over the lifetime of the patents.

The model evaluated at the parameter estimates fits the empirical moments reasonably well. It captures both the facts that the probability that a patent is traded is decreasing with age and that previously traded patents are more likely to be traded than the previously untraded ones. The model also reflects that previously traded patents are less likely to be allowed to expire than the untraded ones. Fit is also assessed by the mean-squared difference between the empirical and the simulated moments, as well as figures that depict both the empirical and simulated moments pooled across patent grant years. Both methods show that the model evaluated at the parameter estimates fits well the empirical moments.

Several results emerge from the estimation. First, the value of the volume of the trade of patents represents about 50% of the total value of all patents. The relative value of the volume of the trade of patents is obtained by multiplying the number of patents traded in our simulation times the average value at age one of a traded patent, and then divided by the total value of all patents. In our analysis, the mean value of a traded patent ($164,670) is about three times the mean value of a non-traded patent ($50,162) (all values are in 2003 US dollars), and about 23% of the patents are traded at least once over their lifetime. The median value of a traded patent and untraded patent are estimated at $58,320 and $10,605, respectively. Instead, the mean value of a patent was estimated at $76,598, and the median $16,184. These findings suggest that the gains from trade in the market for patents may be significant, but the large differences between the average value of the traded and untraded patents also suggest that there may be selection in the trade of patents.

Second, we find that the gains from trade account for about 10 percent of the value of the volume of the trade of patents. The gains from trade were obtained by comparing the value of a patent actually traded with the value that the same patent would have obtained had the option to sell it to a potential buyer not been allowed for over the patent's lifetime. An advantage of structurally estimating the model is that the counterfactual of shutting down the option of selling the patent allows us to quantify the gains from trade while accounting at the same time for selection in the transfer of patents. Our empirical results confirm the fact that patents with higher importance are more likely to be renewed and traded.

Third, we show that the distribution of the gains from trade is highly skewed. We find that about 50 percent of the patents had gains from trade below $3,416, and these patents accounted for only 3.7 percent of the total value of the gains from trade. At the same time, the top 10% of the patents with the highest gains from trade (gains from trade higher or equal than $30,970 per patent) accounted for about 67 percent of the total value of the gains from trade. Furthermore, we demonstrate that the top one percent of the patents represented about 25% of the total gains from trade. This finding is what we expect because patents with higher importance are more likely to be traded and the distribution of the value of patents is known to be skewed. The relative small proportion of patents with significant gains from trade may have important implications on the workings and institutional setting of the market for patent rights.

Fourth, the remaining of the findings refer to the effect that lowering transaction costs has on the market for patents. In a counterfactual experiment, we dropped the estimated sunk transaction cost (estimated at about $5,500) by fifty percent. As an outcome to the experiment, there are two results we want to highlight. First, we found that the proportion of traded patents increases by six percentage points (from 23.1% to 29.6%). Second, we determine that the gains from trade increase by about 10%. Moreover, the additional gains from trade come from patents with returns above the median and below the eightieth percentile of the distribution of initial per period returns, suggesting that these patents, and their owners, will be the ones that benefit the most from lowering transaction costs.

Taken together, our empirical findings indicate that the market for patents generates significant private gains from trade, but that only a small proportion of the traded patents accounts for a significant share of the benefits. Moreover, transaction costs affect this process by reducing the proportion of patents traded and by creating a selection in the transfer of patents. Our interpretation is that as long as small firms can appropriate part of the gains from trade, this market increases their incentives to innovate.

Our empirical strategy to estimate the gains from trade, and the corresponding returns to patent protection, is admittedly unique. We do not observe sale prices; instead, we use the decision to sell patents and the decision to renew patents, which requires a renewal fee, to learn about the distribution of patent values for the original inventor and potential buyers.[6] While the magnitude of the renewal fees increases significantly over the lifetime of patents, there may be insufficient variation in fees, which can limit what we can learn about the stochastic process generating the patent returns without imposing further functional form restrictions (Pakes, 1986). At the same time, recent evidence presented in Harhoff, Narin, Scherer, and Vopel (1999) reassures the functional form we consider for the distribution of initial patent returns to patent protection. We also did not find that alternative functional forms provided a better fit to the data. Another possibility to assess how reasonable are our estimates, and still identify important aspects of the stochastic process generating patents returns and gains from trade, is to limit the identification

[6] Pakes and Simpson (1989) show that with unlimited variation in the renewal fee schedule and under mild regularity conditions the stochastic process generating returns to patent protection can be non-parametrically identified. The actual fees are typically small during the early years of a patent's life and grow over time.

analysis of the renewal and trading decision choices in the observed range of renewal fee schedules (Pakes and Simpson, 1989). Following this strategy, we present evidence that indicates that our results on the gains from trade in patents are robust. Finally, we also confront our estimates of the value of traded patents with data we have obtained on actual sale prices of patent lots sold at auctions organized by patent broker ICAP Ocean Tomo. All in all, the findings suggest that our estimates of patent value are reasonable.

This paper contributes to two strands of the literature on the economics of innovation and technological change. We first contribute to the measurement of the market for innovation – the sale and licensing of patents—. Existing studies have used data on technology licensing to estimate the volume of the market for technology (see e.g., seminal work in Arora, Fosfuri, and Gambardella (2001a, 2001b), and a recent survery by Arora and Gambardella (2011)). Our work complements these studies not only in that we focus on quantifying the value of tradable patent rights (the sale of patent rights rather than the licensing of patents), but also in that we attempt to quantify the gains from trade in this market —the added value by potential buyers to the patents—. We also examine the impact of policy counterfactuals, such as the effect of lowering transaction costs on both the volume of the market for patents, the proportion of patents sold, and the gains from trade. This is the first paper providing estimates quantifying the gains from trade and transaction costs, and the effect that these costs have on the trading of patents and the corresponding gains from trade.

Our work also contributes to a well-known literature estimating the returns to patent protection (Pakes, 1986; Schankerman and Pakes, 1986; Putnam, 1996; Lanjouw, 1998; Schankerman, 1998; Deng, 2007; and Bessen, 2008). In this literature, scholars have used the schedule of renewal fees and the patent's renewal decisions to identify the distribution of the value of holding patents at alternative ages. Our work is distinct in that we use new data on the transfer of patents, together with the renewal fees and renewal decisions, in order to recover the distribution of the value of holding patents at alternatives ages and the gains from trade. The new data allow us to build on their work in two ways. First, we estimate how much of the value of holding a patent is due to the possibility of selling it to potential buyers. In other words, we can uncover, under some assumptions, an estimate of the gains from trade in the market for patent rights. Second, the identification of the parameters of the distribution of the initial per period returns is not uniquely based on observables from patents presumably in the left tail of this distribution (i.e., renewal decisions), but also on observables of what appears to be high value patents (i.e., traded patents) (Serrano, 2010). To the best of our knowledge, this is also the first attempt to estimate the value of holding a patent using information on patent sales.

The paper is organized as follows. Section 2 presents the data sources and the patterns of the transfer and renewal of patents. Section 3 presents a model of patent trading. Section 4 describes the estimation strategy. Section 5 presents the estimation results of the distribution of patent value for traded, untraded, and all patents. Section 6 presents results regarding the volume of the market for patents and the gains from trade. Section 7 examines the effect of lowering transaction costs on the proportion of patents traded and the gains from trade in patent markets. Section 8 concludes the paper. All proofs are included in the Appendix.

2 Data

Our starting point is a panel of patents granted in the period 1988-1997 that were applied for and granted to U.S small firms. Hall, Jaffe, and Trajtenberg (2001) refers to patents granted to U.S firms as "U.S Corporate Patents". The USPTO classifies small firms as business entities with no more than 500 employees (including all subsidiaries). The remaining of the firms are large businesses. About three quarters of all the patents in the period of analysis were granted to firms; the rest were granted to government agencies, individual inventors, or were unassigned as of their issue date. We can distinguish small and large business for patents granted in 1988 and thereafter.[7] Small businesses patents account for approximately 15 percent of all the patents granted to firms. For each of these patents, we obtained information on their reassignment and renewal history.

We now describe the main components of our dataset.

Patent renewal data Patent applications become patents as of their grant date.[8] Multiple patent renewal fees must be paid to extend the life of a patent until its maximum legal length. If a renewal fee is not paid at a renewal date, then the patent immediately expires.[9] In the US, patents are subject to renewal fees only if they were applied for after December 12, 1980. We have obtained these records, as of December 2001, for patents granted for the period 1988 to 1997. The patents granted in 1997 are the last cohort where we can assess their first renewal event. The amount of renewal fees is increasing with the age of the patent, with three renewal events four, eight, and twelve years following a patent's grant date.[10] Small entities are subject to application

[7] Patents granted during and after year 1988 face the first renewal fees starting on January 1, 1992 and thereafter.

[8] The maximum legal length of new patents applied for prior to 1995 was 17 years following their grant date. This maximum legal length was subsequently modified to 20 years following patent application date. During the period 1981-2002 patent applications were granted on average in 2.5 years.

[9] The USPTO began charging renewal fees in 1984 on patents applied for after December 12, 1980.

[10] Standard renewal fees were \$890, \$2,050, and \$3,150 as of January 1, 2003. These fees have increased in value over the years. As of January 1, 2012, the standard renewal fees were \$1,130, \$2,850, and \$4,730.

fees and renewal fees that are half of the amount than large entities pay. While our sample only includes patents originally applied for and granted to small businesses, subsequent patentees pay renewal fees depending on their entity status as of the time the patent's renewal fees are due.

Patent assignment data Another event that can happen during the lifetime of a patent is the transfer of the ownership of the patent rights to others. We use assignment data to identify changes in the ownership of patent rights. The source of this data is the USPTO Patent Assignments Database. When a U.S. patent is transferred, an assignment is recorded at the USPTO acknowledging a change in ownership. We have also obtained these records for all transfers in our sample for the period 1988 to 2001. Some of the assignments recorded with the patent office are administrative events, like a name change or a security interest, as opposed to a true economic transfer of the ownership between two distinct parties. A challenge in using assignment data is to distinguish changes in patent ownership from and other events. To do this, we follow Serrano (2010) and conservatively drop all the assignments that appear not to be associated with an actual patent trade. We 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).

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).

The remaining assignment records identify the sale of a patent and have information about patent numbers, making it possible to merge them at the patent level with information on the payment of renewal fees, the patentee's entity status, and other patent characteristics.

Estimation sample The final dataset is a panel with $54,840$ patents distributed across ten calendar years. Each patent cohort is defined as the grant year of a patent. There are $20,790$ patents granted for the years 1988-1992 and $34,050$ for the rest of the cohorts. Each patent is observed on average 8.5 years, but the minimum and maximum amount of time the patents can be observed is four and thirteen years respectively depending on the patent's grant year. The number of years a patent is observed during its life cycle differs across cohorts because our transfer and renewal data ends in 2001. Patents granted in 1988 are observed from their issue date to age thirteen (the last patent age a renewal decision is observed), and for the patents granted in 1997

Table 1: Percentage of Active Small Business Patents Traded and Expired

Age	All	Not Previously Traded	Previously Traded (Years since last trade) Any Year	One year
A. Probability that an active patent is traded				
2	2.99	2.85	7.47	7.47
7	2.81	2.46	4.79	6.63
11	2.51	2.13	3.77	2.55
B. Probability that an active patents is allowed to expire				
5	17.2	17.7	12.7	6.2
9	25.6	26.6	21.4	11.6
13	25.5	26.6	22.5	14.1

we only have data for the first four years (the first renewal decision is observed at the beginning of the fourth year since a patent's grant date.) For each of these granted patents we observe the trading and renewal decisions at alternative ages over their lifetime.

Table 1 presents a summary of the patent renewal and the transfer rates. The top part of the table presents the probability that an active patent is traded at alternatives ages for both the previously traded and the untraded patents. The probability of trade of previously traded patents is higher than the probability of trade for patents that were previously untraded. Furthermore, the table shows that the probability of trade, whether or not the patent was previously traded, decreases with age. The bottom part of the table presents the expiration rates of active patents at alternative renewal dates for previously traded and untraded patents. Expiration rates of previously traded patents are lower than the expiration rates of the previously untraded patents. In addition, the table shows that the expiration rates increase with patent age. In Serrano (2010) we showed that these patterns were consistent across type of patentees (e.g., individual inventors, and small and large innovators) and patent technology classes.[11] In Figueroa and Serrano (2012) we found that individual inventor and small business patents are more likely to be traded and being allowed to expire than patents of large businesses.

Table 2 shows summary statistics for patents traded and untraded. At the top of the table we present the mean number of patent citations received for patents that had paid none, one, two, or three renewal fees. In the previous literature, patent citations received are have typically been used as a proxy of patent importance. Patents for which renewal fees were never paid received the lowest number of patent citations received. The table also shows that the number

[11] Serrano (2010) defines firm size (small vs. large innovator) based on the size of the patentee's patent portfolio. Figueroa and Serrano (2012) and the current paper define firm size based on the number of employees.

Table 2: Mean Number of Patent Citations Received of Traded, Untraded, Expired, and Renewed Patents

A. Mean number of patent citations received by renewal decision

Age	All	Renewed	Expired
4	2.79	2.95	2.01
8	7.36	8.04	5.37
12	11.90	12.92	8.95
		Always renewed	Never renewed
Lifetime	9.45	15.69	8.28

B. Mean number of patent citations received by trading decision

Age	All	Traded	Not Traded
2	0.68	0.83	0.67
7	6.49	8.25	6.43
11	10.93	13.09	10.88
		Eventually traded	Never Traded
Lifetime	9.45	12.41	8.84

Note: Mean number of citations received reported for expired, renewed, traded, and not traded columns includes all active patents as of a given age. Mean of citations received reported for eventually traded and never traded columns includes all patents.

of patent citations received increases with the number of renewal fees being paid, indicating a positive correlation between patent citations received and patent value (at least as seen by the patentee and reflected by the payment of the renewal fees), that is higher valued patents received on average higher number of patent citations. The bottom part of the table presents the mean number of patent citations received as of the year an active patent is up for trade. We show that for any given age, the patents that were traded at that age had received higher number of patent citations received (as of the same age) than the patents that had not been traded. We also show that patents that are eventually traded received higher number of patent citations received than the rest of the patents. These results confirm the finding that patents with higher number of patent citations received prior to either being traded or being up for renewal are more likely to be both traded and renewed, respectively (Serrano, 2010).

In the estimation of the structural parameters of the model that follows, we maximize a criterion function that measures the goodness of fit of the model to the following stylized facts:

1. The probability that an active patent is allowed to expire at a renewal date for previously traded patents;

2. The probability that an active patent is allowed to expire at a renewal date for previously untraded patents;

3. The probability that an active patent is being traded at a given age for previously traded patents;

4. The probability that an active patent is being traded at a given age for previously untraded patents.

The dimension of the vector of moment conditions will be $m = 186$, which is the sum of 36 conditional probabilities based on the renewal dates and 150 conditional probabilities based on the transfer dates.[12]

In addition, our model's predictions will also be consistent with the facts presented in Table 2. Namely, that high value patents (i.e., patents with higher number of patent citations received) are more likely to be traded (selection effect) and less likely to be allowed to expire; and that the probability that an active patent is traded is decreasing with patent age (horizon effect).

3 A Model of Patent Trading

This section presents a model of the transfer and renewal of patents to be estimated. We first describe the decision making problem that a patent holder faces, and then present new theoretical results that indicate how observations on the transfer and renewal decisions can provide information on the gains from trade, transaction costs, and the distribution of the value of holding a patent. Patent returns illustrate the effects of patent protection.

The starting point for our theory is Pakes (1986) and Schankerman and Pakes (1986). They examine the problem of a patent owner deciding in each period whether or not to pay a renewal fee, and thereby extend the life of a patent, in a context with heterogeneity in the economic value of inventions. Our contribution is adding the possibility of the arrival of opportunities for surplus-enhancing patent transfer, which may lead to alternative potential owners having a greater valuation for a patent right than the current owner. In addition, we consider that to transfer a patent to a new owner involves a transaction cost.

Alternative owners could generate a greater valuation than the current owner for a number of reasons. The economics of innovation literature has traditionally associated this additional valuation with vertical specialization (Teece, 1986; Arora, Fosfuri, Gambardella, 2001a), complementarities in assets (Teece, 1986 and Gans and Stern, 2003), comparative advantages in

[12] The 1988 patent cohort will count towards twelve probabilities in the trading decision (i.e., from age two to thirteen), and three in the renewal dates for each the previously traded and the untraded patents. The 1997 patent cohort, instead, will count towards three probabilities in the trading decision (i.e., from age 2 to four) and one in the renewal decision (i.e., the first renewal date).

marketing or manufacturing (Arora and Ceccagnoli, 2006), and more recently with comparative advantages in patent enforcement (Galasso, Schankerman, and Serrano, 2011). As for the transaction costs, previous scholars have identified them with technology adoption, expropriation risk, and transaction intermediaries such as patent brokers (Arrow, 1962; Teece, 1977; and Astebro, 2002).

Assuming that both the transfer and renewal of patents are based on rational decision making at the patent level, patent owners will only sell their patents whenever the price obtained is higher than the value of retaining the ownership until the next period. Since transaction costs to reallocate the patent to the potential buyer are also a factor, the patent will only be sold if the added valuation is greater than the transaction cost. Similarly, patent owners would renew their patents if the value of keeping the patent until the next renewal date is higher than the renewal fee to be paid. If the renewal fees are not paid, the patent expires. If the patent owner pays the renewal fees and keep the ownership of the patent, the patent holder would face a similar problem next period. Finally, if a patent reaches its maximum legal length of patent protection, the patent will expire next period.

Patent owners maximize the expected discounted value of per period returns and the price they would obtain from selling their patent. Patent owners know the sequence of renewal fees and the transaction costs, but are uncertain about both the arrival of opportunities for surplus-enhancing transfer and the sequence of per period returns that would be earned from retaining the patent. There is evidence pointing at both the large degree of uncertainty about the commercial significance of the invention as well as about the validity and scope of the legal right being granted (see e.g., Pakes (1986) and Lemley and Shapiro (2005)). In this context, patent owners will not necessarily sell their patent to a buyer with higher returns because there is a positive probability that they will match in a future period with another buyer with much higher returns. Similarly, patent owners will pay the renewal fees and thereby retain for an extra period the patent ownership not only because the per period returns from retaining the patent may be high, but also because there is a positive probability that a future potential buyer will have much higher returns than them. The same decision problem will be faced in future periods with the exception that the patent horizon is shorter. An implication of the shorter horizon is that the patent holder will have less time to meet a buyer with high returns, and the potential buyer will have less time to amortize the transaction costs.

More formally, let $V_a(x, y)$ be the expected discounted value of patent protection prior to the transfer and renewal decision from a patent of age a with per period return x if kept by current owner and y if sold to the best potential buyer. If the renewal fee is not paid, the patent expires and $V_a(x, y) = 0$. If the patent is not sold and the renewal fee paid, the owner earns the per period return x and keeps the patent until the next renewal date. If the patent is sold, the buyer earns per period return y and pays the renewal fee. For simplicity, we will assume that the current owner obtains all the surplus in the sale (results are similar if there is Nash bargaining).[13] The value of holding a patent prior to the renewal and selling decision is:

$$V_a(x, y) = \max\{0, V_a^K(x, y), V_a^S(x, y)\}$$

where the age of the patent is $a \in \{1, ..., L\}$, L is the maximum legal length of patent protection, and $V_a^K(x, y)$ and $V_a^S(x, y)$ are the values of keeping and selling the patent, respectively. In particular, $V_a^K(x, y)$, is the sum of the per period patent return x and the option value of the patent minus the renewal fee c_a. $V_a^S(x, y)$ is the price the current owner would obtain from selling the patent, which given the previous assumptions, is equal to the sum of the per period return that the buyer would earn y and the option value of the patent minus the renewal fee and the transaction cost. Therefore we have

$$\begin{aligned} V_a^K(x, y) &= x + \beta E[V_{a+1}(x', y')|\Lambda_a] - c_a \\ V_a^S(x, y) &= y + \beta E[V_{a+1}(x', y')|\Lambda_a] - c_a - \tau \end{aligned}$$

where $\beta \in (0, 1)$ is a discount factor, c_a is the renewal fee at age a, τ is the transaction cost, Λ_a is the information set of the patent owner prior to the transfer and renewal decision in the a^{th} year of the patent, and $E[.|\Lambda_a]$ is an expectation operator over the sequence of per period returns from internal (x') and external (y') opportunities conditional on the set of information Λ_a. Patent owners pay renewal fees c_a based on the patent's owner entity status as of the time a renewal fee is due.

The following assumptions simplify the description of the decision making problem of a patent holder. We have divided these assumptions into three groups. The first group of them allows us to complete the description of the option value of the patent.

Assumption 1 (A1) $Pr(x', y'|\Lambda_a) = F_a(x', y'|x, y, d, \theta)$, where $Pr(.|.)$ denotes a conditional probability, $d \in \{K, S, 0\}$ is the decision to keep the patent, sell it, or let it expire, and

[13] The results may differ if the bargaining process is not efficient. Inefficiency could arise due to asymmetric information regarding x or y. Asymmetric information is less likely to be source of concern in the transfer of patents than in their licensing. For instance, if there was significant asymmetric information about x, the current owner and potential buyer could write a licensing agreement with an option to buy the patent. The option of licensing should decrease the degree of asymmetric information as of the time the patent is sold. In any case, if there was asymmetric information, the decision whether or not a patent had been previously traded would also be a state variable. The analysis of that decision problem, while interesting, is beyond the scope of this paper.

θ is a vector of parameters.

Assumption 2 (A2) Patent holders and potential buyers know the sequence of renewal fees c_a that will be required to be paid in order to keep the patent active. Furthermore, the renewal fee schedule in every renewal date is increasing with the age of the patent.

The option value of a patent now depends only on the current per period return of the patent owner (x), the return of the potential buyer (y), and the decision to keep or sell the patent (d). Assumption A1 implies that the past history of per period returns does not affect the future returns. Assumption A2 is consistent with the fact that renewal fees are increasing with the age of the patent at alternatives renewal dates. Therefore, under the above assumptions, the option value of keeping and selling the patent becomes

$$E[V_{a+1}(x', y') \mid x, y, d] = \int \int V_{a+1}(x', y') \mathrm{d}F_a(x', y'|x, y, d)$$

The second group of assumptions allows us to simplify the decision problem of a patent holder. We will first motivate these assumptions, and then group them into assumptions A3-A5. Assumption A3 considers that the return under the best outside opportunity, y', depends on the current internal return, x'. This is motivated by the fact that while per period returns may be firm specific, it is also likely that they will be patent specific. Assumption A4 considers that the internal return x' depends on the previous realized return (x if $d = K$ or y if $d = S$), but it does not depend on the alternative return that was not realized (y if $d = K$ and x if $d = S$). This assumption is consistent with a framework where the arrival of new inventions depend on the inventions that have been adopted but not of those that were not. In assumption A5, we consider as true that the process of internal returns is the same for all patents with the same period return independently of whether or not the patent had previously changed ownership. This premise allows us to simplify the decision problem of patent owner, excluding the possibility that, conditional on patent returns, future patent returns may depend on the past history of trading decisions. An implication of the previous assumptions is that the law of motion $F_a(x', y'|x, y, d)$ can be rewritten as the product of the transition of external and internal returns $F_a(x', y'|x, y, d) = P_a^y(y'|x')[1_{\{d=K\}} P_a^x(x'|x) + 1_{\{d=S\}} P_a^x(x'|y)]$. These assumptions are summarized into Assumptions A3-A5.

Assumption 3 (A3) Let $F_a(x', y'|x, y, d, w) = P_a^y(y'|x', x, y, d) * P_a^x(x'|x, y, d)$ where P_a^y and P_a^x are probability functions. The return under the best outside opportunity, y', depends on the current internal return, x', but not on the previous history of returns or decisions:

$P_a^y(y'|x', x, y, d) = P_a^y(y'|x')$.

Assumption 4 (A4) The internal return x' depends on the previous period realized return (x if $d = K$ or y if $d = S$) but not on the alternative return that was not realized (y if $d = K$ or x if $d = S$). That is:

$$P_a^x(x'|x, y, d) = \left\{ \begin{array}{lll} P_a^{x,K}(x'|x) & \text{if} & d = K \\ P_a^{x,S}(x'|y) & \text{if} & d = S \end{array} \right\}$$

Assumption 5 (A5) Conditional on the realized return at previous period, the internal return does not depend on the decision of keeping or selling the patent. That is:

$$\begin{array}{rcl} P_a^{x,K}(x'|x) & = & P_a^x(x'|x) \\ P_a^{x,S}(x'|y) & = & P_a^x(x'|y) \end{array}$$

The next two assumptions allow us to guarantee that there exists a cut-off rule solution to the agent's decision problem. Assumption A6 considers that the transaction costs are sunk.[14] This assumption is motivated by the fact that previous scholars have identified that part of these costs with large fixed noncapital and capital costs of adoption, organizational changes, and intermediaries such as patent lawyers and brokers (Teece, 1977 and Åstebro, 2002). Assumption A7 describes the transition probabilities of internal and external returns, P_a^x and P_a^y. We consider a continuous transition probability of internal returns such that the probability that next period return's are larger than any given number is larger the higher is the current patent return. More formally, $P_a^x(x'|x)$ is continuous in x' at every x and decreasing in x, and $P_a^y(y'|x')$ is continuous in y' at every x' and decreasing in x'. This is reasonable: it states that patents with higher per period returns today are more likely to have higher returns tomorrow than patent with lower returns. At the same time, for the solution of the decision problem to be well defined, we also need that patent returns do not increase too fast. That is, we need that for a given patent return the probability that next period returns are higher than a given number is not increasing with patent age. This requires that for a given x and any z the probability $P_a^x(x' > z|x)$ not increases with a; and that the probability $P_a^y(x' > z|x)$ not increases with a. These assumptions are summarized into A6-A7.

Assumption 6 (A6) Patent holders and the best potential buyer know the transaction costs τ

[14] Below we will discussed the extent to which this assumption can be relaxed. We will also relax this assumption later in the empirical section.

that will be required to be paid in order to transfer the patent. Furthermore, the transaction costs are sunk and constant with the age of the patent.

Assumption 7 (A7) The function $P_a^x(x'|x)$ is continuous in x' at every x and decreasing in x. The function $P_a^y(y'|x')$ is continuous in y' at every x' and decreasing in x' At the same time, for a given x and z, the probability $P_a^x(x' > z|x)$ decreases with a; and for given x' and z, the probability $P_a^y(y' > z|x')$ decreases with a.

Finally, the remaining assumption is made for convenience. The assumption allow us to characterize the solution of the decision making problem of the patent holder. A sufficient condition is that the transition probabilities described in A7 come from a random walk process. That is, the transition probability of internal returns, $P_a^x(x'|x)$ (if $d = K$) or $P_a^x(x'|y)$ (if $d = S$) comes from a random walk process $x' = g_a^i x$ (if $d = K$) and $x' = g_a^i y$ (if $d = S$). We model the uncertainty concerning experimentation outcomes at the firm by assuming that the growth of internal returns $g_a^i \in [0, \infty)$ is drawn independently of x and is non-increasing with a as the patented technology matures (A8.1). Note that the next period return x', that depends on x, can take any value higher or equal than zero. This assumption guarantees that the probability that next period internal return's are larger than any given number is larger the higher is the current patent return. This process reflects the dynamics of per period returns over the lifetime of a patent within the firm; as new technologies age we expect the patent per period return to decrease on average too. As for the process of external returns, we consider a transition probability of external returns such that the probability that a potential buyer return is higher than any given number is higher the larger is the current patent return. This process captures the possibility that potential buyers may generate higher value the than current owner, from the same patent right. As we argued above, this may be due to comparative advantages in marketing, production, complementary assets, new technologies that fall in the same patent rights' scope, and comparative advantages in the ability to enforce patent rights without resorting to courts (patent enforcement). A sufficient condition is that the process of external returns $P_a^y(y'|x')$ is a random walk process such that $y' = g_a^{e\prime} x'$. We model the uncertainty concerning the source of external returns by assuming that $g_a^{e\prime} \in [0, \infty)$ is a random variable independent of the return of the current owner x' and non-increasing with a (A8.2). This is consistent with a number of economic studies on the process of R&D activity and product innovation (see e.g., Grossman and Helpmann (1991) and Bentall and Peled (2002)). Note that y' depends on the patent per period return of the current owner (x'), and can take any value higher or equal than zero. Moreover, we think of the improvement factor g_a^e as the arrival of opportunities for surplus-enhancing via the sale of patents, and thus the outright transfer of all patent rights to other firms, relative to the in-house use and the potential licensing proceeds of the current owner.[15] The approach we follow to model the uncertainty about the intertemporal transition of patent per period internal and external returns is consistent with the work of Grossman and Helpmann (1991). They view the process of innovation as the action of generating an ever-expanding range of horizontally differentiated products. These assumptions are summarized in A8.

Assumption 8 (A8) The transition probabilities $P_a^y(y'|x')$ and $P_a^x(x'|x)$ come from random walks processes. That is:

$$y' = g_a^{e\prime} x'$$

where $g^{e\prime}$ is a random variable non-increasing in a and independent of x' with cdf F^{g^e} that is continuous in y' at every x' such that, for a given x' and return z, the $P_a^y(y' > z|x')$ decreases with a. Similarly:

$$x' = \left\{ \begin{array}{ll} g_a^{i\prime} x & \text{if} \quad d = K \\ g_a^{i\prime} y & \text{if} \quad d = S \end{array} \right\}$$

where $g^{i\prime}$ is a random variable non-increasing in a independent of x with cdf F^{g^i} that is continuous in x' at every x such that, for a given x and z, the $P_a^x(x' > z|x)$ decreases with a.

To characterize when a patent holder will decide to sell, to keep, or let its patent expire, we must describe the properties of the value function and option value of a patent. The following Lemma shows that the value function of a patent is continuous, weakly increasing in the returns of the patent, weakly increasing in the return of potential buyers, and weakly decreasing in patent age.

Lemma 1. *The value function $V_a(x, y)$ is continuous and weakly increasing in the current return of the holder of the patent, x, and the return factor of the potential buyer, y. The option value $EV_{a+1}(x', y'|x, y, d)$ is weakly decreasing in a.*

Proof. See appendix. ■

Based on the results from this Lemma, the following proposition shows that the solution to the problem of the firm can be summarized into two cutoff rules – $\widehat{x}_a(w)$ and $\widehat{g}_a^e(x, \theta)$ – that divide the state space into three areas (keep, sell or do not renew).

[15] There are no systematic data on patent licensing revenue, but there is anecdotal evidence. IBM's licensing revenue was \$1.6 billion in the year 2000 (Berman (2002) as reported in Merrill, Levin and Myers (2004)). In 1996, U.S. corporations received \$66 billion in income from royalties of unaffiliated entities (Degnan, 1998). Texas Instruments reported to have obtained \$1.6 billion in licensing royalties from 1996 to 2003 (Grindley and Teece, 1997).

Figure 1: Optimal choices of a Patent Holder

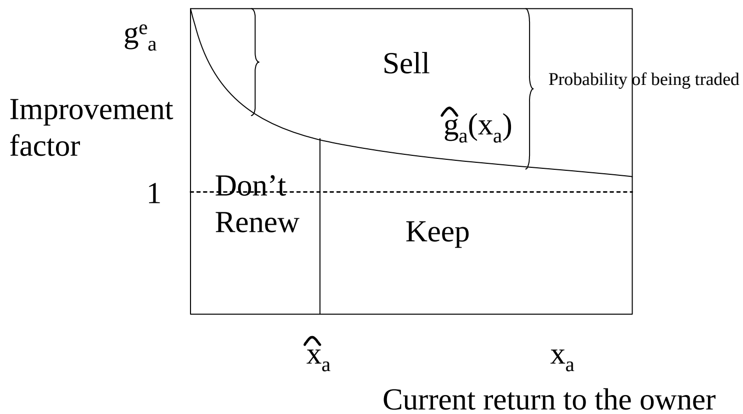

Proposition 1. *For each age, there exist cutoff rules $\widehat{x}_a(w)$ and $\widehat{g}_a^e(x,\theta)$ such that patents with an improvement factor above $\widehat{g}_a^e(x,\theta)$ will be sold. Among the patent holders that met a potential buyer with improvement factor below $\widehat{g}_a^e(x,\theta)$, the patents with per period returns x above $\widehat{x}_a(\theta)$ will be renewed, and those with returns x below $\widehat{x}_a(\theta)$ will be retained by their current owners.*

Proof. See Appendix. ■

Figure 1 presents the results of this proposition. In the vertical axis we have the improvement factor of the best potential buyer while in the horizontal axis we have the per period return of the current patent owner prior to the trading and renewal decision. The first cutoff is the per period return $\widehat{x}_a(\theta)$ that makes the owner indifferent between keeping the patent or letting it expire. This is a vertical line that separates the regions "Don't renew" and "Keep." This is the well known Pakes and Schankerman's renewal rule. The second cutoff $\widehat{g}_a^e(x,\theta)$ is new. This is the improvement factor that makes a patent owner indifferent between selling or not selling a patent with per period return x and age a. This is the curve that separates the region "Sell" from the regions "Don't renew" and "Keep." If a patent holder faces a potential buyer with an improvement factor above the cutoff $\widehat{g}_a^e(x,\theta)$, the patent will be sold. These patents belong to the "Sell" region. Moreover, as long as the improvement factor is lower than $\widehat{g}_a^e(x,\theta)$, patents with lower per period returns than $\widehat{x}_a(\theta)$ will be allowed to expire, this is the "Don't renew" region, and the ones with per period return higher than $\widehat{x}_a(\theta)$ will be renewed and retained by the current owner. This is the "Keep" region. The following proposition characterizes the properties of the function $\widehat{g}_a^e(x,\theta)$ and $\widehat{x}_a(\theta)$ with x and a.

3.1 The selection and horizon effect

Proposition 2. *The following properties characterize the functions $\widehat{x}_a(\theta)$ and $\widehat{g}^e_a(x,\theta)$:*

- **(a)** *The function $\widehat{x}_a(\theta)$ is increasing with patent age a at the renewal dates.*
- **(b)** *If the transaction cost τ is positive, then*
  - **(b.1)** *For a fixed patent age a, the function $\widehat{g}^e_a(x,\theta)$ is decreasing with x;*
  - **(b.2)** *For a fixed patent return x, the function $\widehat{g}^e_a(x,\theta)$ is increasing with a.*

Proof. See Appendix. ■

The findings in Proposition 2 are intuitive. In part (a), the renewal cutoff $\widehat{x}_a(\theta)$ increases with patent age a because while the fee schedule increases with a, the per period returns do not on average increase with a. This result is consistent with both Pakes (1986) and Schankerman and Pakes (1986) and the rest of the literature estimating patent value.

The new results are in part (b). In part (b.1), we show that $\widehat{g}^e_a(x,\theta)$ is decreasing with patent returns (x). In particular, if for a given improvement factor g^e the gains from trade by potential buyers are scaled up by the patent returns x, the higher is x the lower is the necessary improvement factor to amortize the sunk transaction cost. This result hinges on the assumption that the improvement factor of the potential buyer (g^e) is complementary with the size of the invention (x), inducing that the value added of potential buyers to patents increases with the patent's per period return of the current owner. This is consistent with most theoretical settings in the economics of innovation, where acquired capital (e.g., patents) would become a complementary asset with the capabilities or capital of potential buyers (see e.g., Teece (1986), Jovanovic and Rousseau (2008), Silveira and Wright (2010), Galasso, Schankerman, and Serrano (2011)).

Lastly, part (b.2) of the proposition shows that for a fixed patent return x, the function $\widehat{g}^e_a(x,\theta)$ is monotonically increasing with patent age a. This is because for a given patent return x a shorter horizon implies less time to amortize the transaction cost. The result does not depend on whether the transaction cost is sunk; more generally, it requires that the transaction cost is not proportional to the difference between the value of selling and keeping a patent. Economic settings were transaction costs may depend on the per period return x or $V(x,y)$, which our estimation strategy implicitly will allow for, would not be proportional to the difference between the value of selling and keeping a patent.[16] The rest of the patents, which will be let to expire, will have the lowest returns.

[16] The key element that allows us to prove these results is that the transaction cost is not proportional the to the difference between the value of selling and keeping a patent (see Appendix). To the extent that a part of this cost is sunk, and thus does not fully internalize how the difference between the value of selling or not selling a patent changes as the revenue and age of the patent varies the result will hold. In other words, a part of the transaction

Clearly, the functions $\widehat{g}_a^e(x,\theta)$ and $\widehat{x}_a(\theta)$, together with the distribution of improvement factor and the distribution of the per period returns of a patent at age a and prior the decision of the patent holder, determine the probability that an active patent will be sold, kept, or let to expire. The probability that a patent with return x will be sold at age a is the likelihood that an improvement factor is above $\widehat{g}_a^e(x,\theta)$.

An immediate implication from Proposition 2 is that there is both a selection and a horizon effect in the trading of patents. There is a *selection effect* because the probability than an active patent is traded at age a increases with the patent return x. That is, patents with higher per period returns are more likely to be traded than the rest of the patents. Note that as x increases the cutoff rule $\widehat{g}_a^e(x,\theta)$ decreases and thus the probability that a patent is being traded increases. The selection effect implies that traded patents, and especially those that have been recently traded, have on average higher returns than the previously untraded ones as of the time of the trading decision. The results are consistent with the stylized facts on the transfer of patents presented in Table 1 and 2. In particular, we reported that patents with higher number of patent citations received are more likely to be traded; and that previously traded patents are more likely to be traded than the untraded patents.

There is a also *horizon effect* because, for a fixed patent return x, the probability that an active patent is being traded decreases with patent age. Note that as the patent age increases the cutoff rule $\widehat{g}_a^e(x,\theta)$ decreases with age and thus the probability that a patent is being traded decreases with age too because the random variable g_a^e is independent of x. The finding is also intuitive: a shorter horizon implies less time to amortize the transaction costs. The result is also consistent with the stylized facts shown in Table 1; the horizon effect accounts for why the proportion of active patents being traded decreases with patent age.

3.2 Discussion of the results

The previous theoretical results suggest that observations on the transfer and renewal decisions can provide information on the distribution of the value of holding a patent, the gains from trade, and the transaction costs. First, observations of the proportion of patents allowed to expire and their corresponding renewal fees provide information to identify the distribution of patent value at alternatives ages. This is insight was first shown in Pakes (1986) and Schankerman and Pakes (1986). Previous scholars have used the observed proportion of patents renewed at alternatives

cost could in principle depend on the current patent return x and the age of the patent a. On the other hand, if the transaction costs were proportional to the difference between the value of selling and the value of keeping the patent, then neither part (b.1) nor (b.2) would hold.

dates, together with the schedule of renewal fees, to estimate the distribution of the value of holding patents at alternatives ages.

Second, the selection into the trading decision suggests that observations on the trading of patents can provide additional information about the distribution of patent returns. At alternative ages we observe whether patents are traded, kept by their owners, or let to expire. Based on the selection effect, traded patents will have on average the highest returns of all. Patents retained by their owners will have, on average, lower returns than the traded patents. This is because patent owners holding patents with high per period returns will pay the renewal fees at alternative renewal dates not only because their current returns are high, but also because they have a higher probability than the rest of patent holders to find it optimal to sell their patent to a buyer with much higher returns. While the probability to meet a buyer is independent of the current patent returns, the patent returns potential buyers can potentially generate from a specific patent depend on the patent return the current owner is already generating. This hinges on the complementarity between current patent returns and the improvement factor of potential buyers upon acquiring the patent.[17] Therefore, observations on the trading decision are also informative about the distribution of patent returns.

Third, observations of patents being traded can provide information about the resulting gains from trade. For a given distribution of per period returns at age a, and prior to the transfer and renewal decision, the observed proportion of patents being traded at age a will be associated with the proportion of buyers with improvement factors above the selling cutoff $\widehat{g}_a^e(x,\theta)$. Assuming that the transfer of a patent is based on decision making at the patent level characterized by the function $\widehat{g}_a^e(x,\theta)$, conditional on patent age and the distribution of patent returns, higher proportions of patents being traded will be associated with higher gains from patent trading. Thus, the observed proportion of traded patents at alternatives ages will provide some information on the distribution of the process of external returns (improvement factor), which ultimately determines the gains from trade in patents. In addition, the fact that a large proportion of patents are not active immediately after a renewal date provides an important drastic and sudden variation in the distribution of the patent returns of active patents (at each renewal date), which together with the observed rate of trade in the same and the years immediately before and after a renewal date can help us learn something else about the surface of $\widehat{g}_a^e(x,\theta)$ and the stochastic process of

[17] The assumption underlying complementarity between acquired capital and the improvement factor of potential buyers provides enough flexibility to the model to account for the observed patterns of the transfer of patents. In fact, if alternatively the per period return of a potential buyer (y) was independent of the per period return of the current owner (x), a theoretical setting like ours would not be flexible enough to match the empirical facts presented in Table 1 and 2. In particular, the facts that previously traded patents are more likely to be traded later in their life; or that patents with higher per period returns are more likely to be traded than the rest of the patents.

external returns.

Fourth, observations of patents being traded at alternative ages can provide information about the transaction costs. Conditional on the distribution of per period returns at alternative ages and the properties of the process of external returns, the decreasing pattern with age of the observed proportion of active patents being traded will be associated with the magnitude of the transaction cost. Proposition 2 indicates that for a fixed patent return and a distribution of the process of external returns, the decreasing pattern of the observed proportion of patents being traded with age depends on the size of the transaction cost. In practical terms, higher transaction costs imply steeper decreasing patterns with age of the proportion of patents being traded.

Finally, the renewal rates of previously traded and untraded patents can provide additional information about both the gains from trade and the transaction costs. Conditional on the distribution of patent returns at alternative ages, the difference between the renewal rates of previously traded and untraded patents provides information about both the gains from trade of previously traded patents and transaction costs. In particular, higher transaction costs and higher improvement factors by potential buyers will imply a larger difference between the renewal rates of previously traded and the untraded patents. At the same time, while higher transaction costs imply that the proportion of patents being traded decreases with age (and become steeper); higher improvement factors will shift up the proportion of patents being traded at all ages. Thus, the observed renewal rates of previously traded and untraded patents can provide information on both the process that determines the gains from trade and the transaction costs.

3.3 Model stochastic specification

Initial per period patent returns Following the literature on the estimation of patent value we assume that initial patent returns are lognormally distributed. This is not an unreasonable assumption, empirical studies using survey data on commercialization outcomes have shown that lognormal distributions provide a reasonable fit to the observed distribution patent returns (Harhoff, Narin, Scherer, and Vopel, 1999). The distribution of the initial patent return is

$$\log(x_1) \sim F_{x_1}(\mu, \sigma_R)$$

where $F_{x_1}(.)$ is a normal distribution.

Transition of internal patent returns The transition of internal returns is defined as

$$x_{a+1} = \left\{ \begin{array}{ccc} 0 & & \text{with probability } (1-\gamma^i) \\ x_a\delta & \text{if } d=K & \multirow{2}{*}{\text{with probability } \gamma^i,} \\ y_a\delta & \text{if } d=S & \end{array} \right\}$$

where $d = K$ if the patent was retained by its current owner, $d = S$ if the patent was sold, and $\delta < 1$.

This stochastic process has an economic interpretation. We model the uncertainty concerning a firm's experimentation outcomes by assuming that the growth of internal returns $g_a^i \in [0, \infty)$ is independent of the size of the existing innovation x. In particular, we assume that at each age patent owners and potential buyers find out whether patent returns may become obsolete ($g_a^i = 0$) and thus the returns are zero thereafter; or alternatively, the returns depreciate at a rate $\delta < 1$ ($g_a^i = \delta$). This captures both the uncertainty about commercial significance of an innovation as well as the validity and scope of the legal right being granted. This specification follows Lanjouw (1998). In contrast, Schankerman and Pakes (1986) examined renewal decisions using the more stylized model of deterministic transition of patent returns (i.e., $g_a^i = \delta$). At the same time, Pakes (1986) allowed for uncertainty and age dependence in the transition of patent returns over the patent's life. Nonetheless, Pakes estimates imply that most age dependence uncertainty is resolved before the fifth year of the patent's life. While we observe the trading decision over the patent's entire lifetime, the first renewal fee in the U.S. schedule is not collected until four years after a patent's grant date, which limits what we can learn about age dependence early in a patent's life. In light of Pakes' result and the fact that we are using U.S. patent data, we have conservatively assumed a transition of internal patent returns that is independent of patent age while still stochastic over a patent's lifetime.[18]

Transition of external patent returns and transaction costs We model the uncertainty concerning external patent returns by assuming that the size of the improvement factor of potential buyers g_a^e is a random variable independent of the size of the preceding per period returns x. This is consistent with a number of economic settings. For instance, Bental and Peled (2002) describes R&D activity as a sequential search involving random draws, i.e., technology draws, from a population of existing technologies. Alternatively, Grossman and Helpmann (1991) view

[18] A more flexible stochastic process for g^i allowing for age dependence such as in Pakes (1986) can, in principle, be estimated using both the renewal and the transfer data. Serrano (2006) argued that a stochastic process where g^i depended on age could be identified with information of the observed trading decisions prior to the first renewal date. Serrano also used the flow of patent citations received by a patent at alternatives ages to estimate the value of holding a patent.

product innovation as a process of generating an ever-expanding range of horizontally differentiated products independent of the size of existing innovations and their vintage. Furthermore, the external patent returns may be the result of acquisition of firms to accumulate defensive patent portfolios for resolving legal disputes non-litigiously (Hall and Ziedonis, 2001) or due to economics of scale in enforcement (Lanjouw and Schankerman, 2004), and there is evidence indicating the probabilistic nature of patent rights (Lemley and Shapiro, 2005). To this end, we assume that there is a positive probability γ^e that the patent return earned by the potential buyer following a patent sale will be $y_a = 0$. Alternatively, there is a probability $(1 - \gamma^e)$ that the patent return earned by the potential buyer will be based on the improvement factor g^{*e} observed prior to the decision of selling the patent. In this case we will have that $y_a = g^{*e} x_a$, where the random variable g^{*e} is distributed exponentially with cdf $F^{g^{*e}}(g^{*e}) = 1 - \exp(-\frac{g^{*e}}{\sigma_a^e})$ and $\sigma_a^e = \sigma^e$. Consequently, the random variable g^e is distributed:

$$F^{g^e}(u) = \left\{ \begin{array}{cc} \gamma^e & \text{if} \quad u = 0 \\ (1-\gamma^e)[1 - \exp(\frac{u}{\sigma^e})] & \text{if} \quad u > 0 \end{array} \right\}$$

In light of the benchmark model analyzed in this section, another interpretation of this stochastic specification is that the transaction costs have both a fixed and a variable component. In particular, the transaction costs could now account for the sum of the sunk fixed cost (τ) and a variable component that equals to the parameter γ^e times the net value of selling the patent $\gamma^e y + \gamma^e \beta E[V_a(x', y')|x, y]$. In our setting, the variable component of the transaction costs naturally raises from the decision to acquire a patent under an uncertain outcome. Previous scholars have reported that the process of acquiring innovations and patent rights is typically characterized by risks associated with the adoption of technology, patent enforcement, or the match of the patent and innovation with potential buyer's complementary assets (see e.g., Saha, Love, and Schwart (1994), Rosenberg (1996), and Lemley and Shapiro (2005)).

Measurement error Because there may be measurement error in the patent transactions recorded at the USPTO, our specification allows for it in the trade of patents. The specification we consider is the following. Let d^* and d be the true and the observed decision variables, respectively. If $d = 0$, then $d^* = 0$; If $d = K$, then $d^* = K$; but if $d = S$, then there is a probability $\varepsilon > 0$ that $d^* \neq S$.

4 Estimation

The parameters of the model are characterized by the vector $\theta = (\mu, \sigma_R, \sigma^e, \gamma^e, \tau, \varepsilon, \gamma^i, \delta)$. The discount factor is set to $\beta = 0.9$ as in Pakes (1986). The rest of the parameters are jointly estimated using the simulated generalized method of moments. This method involves finding θ in order to minimize the distance between the empirical moments, defined as those from the data, and the simulated moments generated by the model. The moments generated by the model are simulated because they cannot be solved for analytically due to the fact that patent returns are unobserved and serially correlated over time.

4.1 Estimation algorithm

More formally, let the k parameter vector θ_0 be the unique solution to $G(\theta) = E[h - \eta(\theta)] = 0$, where the vector h is defined as the true hazard probabilities, and the vector $\eta(\theta)$ contains the probabilities predicted from the model given a parameter vector θ. Both h and $\eta(\theta)$ have dimension m, which is the number of moment conditions. Let N be the sample size. The sample moment condition is $G_N(\theta) = h_N - \eta_N(\theta)$. The vector $\eta_N(\theta)$ contains the hazards simulated by the model given the vector of parameters θ; and h_N represents the sample hazard proportions in the data.

The simulated minimum distance estimator, $\widehat{\theta}_N$, of the true parameter vector θ_0 is defined as

$$\widehat{\theta}_N = \arg\min_{\theta} [G_N(\theta, Z_N)]' W_N [G_N(\theta, Z_N)]$$

where W_N is a positive semi definite weighting matrix with dimension m by m defined as $W_N(\theta) = diag(\frac{1}{h_j^2}\sqrt{n_j/N})$, where n_j is the number of patents in which hazard j is conditioning on the relevant age/cohort, and h_j is the j^{th} component of the vector of the sample hazard proportions.[19]

Given that the conditions required to have the consistency and asymptotic normality of our estimator are satisfied (Pakes and Pollard, 1989), and that $W_N(\theta)$ converges in probability to a semi-definite matrix W_0, we have that $G_N(\theta)$ converges in distribution to a normal distribution,

$$\sqrt{N} G_N(\theta_0, Z_N) \to N(0, V)$$

[19] The optimal weighting matrix would involve using the inverse of the asymptotic variance-covariance matrix of the sample moment conditions. Since efficient GMM is computationally very costly in our framework, we follow Lanjouw (1998) and use a diagonal matrix that weights each moment according to the number of observations of each sample hazard. Because the absolute value of the sample renewal rates approximately are ten times larger than the transfer rates, our weighting matrix divide each element in the diagonal matrix by its corresponding squared sample hazard. Thus, our minimization criterion is equivalent to one that minimizes the relative differences between the sample and the simulated hazards.

and $\widehat{\theta}_N$ converges to θ_0, and $\sqrt{N}(\widehat{\theta}_N - \theta_0)$ satisfies a central limit theorem

$$\sqrt{N}(\widehat{\theta}_N - \theta_0) \to N(0, (\Gamma' W_0 \Gamma)^{-1} \Gamma' W_0 \Omega W_0 \Gamma (\Gamma' W_0 \Gamma)^{-1})$$

where

$$V = (\Gamma' W_0 \Gamma)^{-1} \Gamma' W_0 \Omega W_0 \Gamma (\Gamma' W_0 \Gamma)^{-1}$$
$$\Gamma = p \lim \frac{\partial G_N(\theta_0, Z_N)}{\partial \theta'}$$
$$\Omega = E([G_N(\theta_0, Z_N)][G_N(\theta_0, Z_N)]')$$

4.2 Obtaining $\widehat{\theta}$

There are a number of steps we follow to obtain the estimate $\widehat{\theta}$. First, we simulate the model in order to obtain the vector of hazard probabilities $\eta_N(\theta)$. This involves obtaining the cutoff rules $\widehat{x}_a$ and $\widehat{g}_a^e$ and the value function of keeping and selling the patent. The second step involves finding the value of θ that minimizes a metric of the distance between the simulated and the empirical hazard probabilities.

Given a vector of parameters θ, the hazard probabilities $\eta_N(\theta)$ are generated by first solving the model and then simulating it. To solve the model, we compute both the value of keeping a patent and the cutoff rules of renewing and selling it. With the purpose of doing so, we solve the model recursively, starting from the maximum legal length of patent protection and moving backwards until the patent's grant date. To obtain the value function of keeping the patent at alternative ages, we discretize the per period return into a number of grid points.[20] Solving the value function backwards involves constructing cutoff rules for selling and renewing a patent. At each patent age we use this cutoff rules in order to approximate the integral that defines the expected value of patent protection (i.e., the patent's option value).[21] The backward process ends when we reach the patent's age one. At this point, we have calculated the value of keeping a patent at alternative ages and the grid points of the per period patent returns.

The life cycle history of N patents is simulated S times using the previously calculated values of keeping and selling the patent. Simulating the life cycle of these patents consists of taking pseudo-random draws from the distribution of the initial returns, and then pass each initial patent

[20] We apply a grid of 10 points, but have used 100 and 300 grid points and found no significant differences. Rather than the number of points used in the grid, the important aspect is to chose their location based on the curvature of the value function. The location of the grid points was chosen to match a lognormal distribution with parameter similar to those of the current estimates of initial patent per period returns. This allowed us to ensure that the density of the grid points is higher where the curvature of the value function is steeper.

[21] We use quadrature methods to numerically approximate the value of the multi-dimensional integral that defines the patent's option value.

return through the stochastic processes of internal and external patent returns, including the measurement error in recorded transfers. History of trade and renewals decisions for each patent, as well as their patent returns and patent value, are obtained and recorded. The simulated hazard probabilities $\eta_N(\theta)$ are calculated by averaging out, over the number of patents and simulations, the patent renewal and transfer decisions.

Finally, an algorithm based on simulated annealing methods is used to obtain the simulated minimum distance estimate. Annealing methods are used because the objective function is likely to be non-smooth and have multiple local minima.[22] Parameter estimates and their respective standard errors were obtained by applying a Newton method algorithm once the simulated annealing method converged to the global minimum of the objective function of the estimator.

5 Estimation Results

Table 3 presents the parameter estimates and the corresponding standard errors. On the top part of this table, estimates for the initial distribution of patent returns as well as the processes of depreciation and obsolescence can be found. The bottom part shows the parameter estimates for the process of the gains from trade and the transaction cost. The parameter estimates are all positive and significant.

An indicator of how the estimated model fits the data is comparing the empirical moments and the simulated moments from the model. Figures 2 and 3 (see appendix) outline both the empirical and simulated moments pooled across patent grant years. The model fits the empirical moments used in the estimation algorithm reasonably well. Figure 2 presents both the simulated and empirical probabilities that an active patent is allowed to expire for previously traded and untraded patents. The probability of an active patent is being allowed to expire conditional on having been previously untraded is well fitted by our model simulations. For the untraded patents, however, the simulation generated by the model tends to predict less steepness than what is observed in the data, especially in the second and last renewal date. This feature may be due to unobserved patent heterogeneity in that some patents are more likely to be traded than others independently of their patent returns. Alternatively, it could be due to our rather simple specification of the process of the transition of internal returns. For another moment that sheds light on how the estimated model fits the data, Figure 3 presents the probability of that an active patent is traded for previously traded and untraded patents.[23] The model is able to capture well

[22] In particular, we use the simulated annealing algorithm developed and tested in Goffe, Ferrier and Rogers (1994).

[23] Note that the age profiles we examine are right censored at age thirteen since this is the longest age profile of

Table 3: Parameter Estimates

Description (Parameter)	Estimate[a]
A. Patent initial returns, depreciation, and obsolescence	
Depreciation factor (δ)	0.9003 (4.35*10^{-2})
Not obsolescence (γ^i)	0.9702 (7.60*10^{-3})
Mean parameter of the Lognormal Initial Distribution (μ)	8.3115 (4.50*10^{-3})
Std. Deviation parameter of the Lognormal Initial Distribution (σ_R)	1.7202 (4.30*10^{-3})
B. Market for patents and transaction costs	
Transaction cost (τ)	5,494.8 (63.46)
Mean External Growth of Returns (σ^e)	0.3775 (2.70*10^{-3})
Proportion of unsuccessful transfers (γ^e)	0.0386 (2.86*10^{-4})
Random transfers (ε)	0.0053 (8.15*10^{-4})
Size of sample	54,840
Simulations in the estimation	164,520
MSE[b]	4.2447*10^{-4}

[a]Estimated standard errors in parenthesis.
[b]MSE is the sum of squared residuals of the difference between the estimated and the actual moment conditions divided by 186, the number of moment conditions.

the decreasing shape of both probabilities with age.[24] Another indicator of the fit of the model is the sum of squared residuals of the difference between the estimated and the actual moment conditions divided by the number of moment conditions (MSE). At the bottom of Table 3, we show that the MSE is low ($4.2447 * 10^{-4}$).[25] Putting these findings together, both indicators suggest that our estimated model fits the data well.

Parameters μ and σ_R determine the initial distribution of the per period returns as of the grant date of patents. A high σ_R implies a high degree of heterogeneity in initial patent returns. A low μ implies that the initial per period returns to patents are low. The estimates of μ and σ_R indicate that a large proportion of the patents in our data will start with very low per period returns. Parameters δ and γ^i represent depreciation and obsolescence of patent returns, respectively. The estimate of γ^i implies that the patent per period return of about three percent of the active patents in a given year become zero and thus obsolete.[26] The estimate of δ indicates that the patent returns of the non-obsolete patents in a given year decrease by ten percent within a year.[27]

the oldest patent cohort in our estimation sample.

[24]In out of sample predictions, we also found that the probabilities sharply decrease when patents get closer to the maximum legal length of patent protection (age 17 for patents applied for before 1997), which is consistent with the evidence on the transfer of patents presented in section 2.

[25]This magnitude of the MSE is comparable to the reported MSE in Pakes (1986).

[26]The obsolescence rate is $1 - \gamma^i = .03$.

[27]The ten percent "depreciation" rate of patent returns over the patent age is computed in the following way: $(1 - .9003) = .10$.

Both estimates together imply that on average the per period returns of a patent decrease at the rate of approximately thirteen percent a year.[28] These parameters together account for the distribution of initial returns and its evolution over the age of patents, which we examine below.

Table A-1 presents a summary of the distribution of per period returns at ages one, four, and seven. This table and the remaining ones are in the appendix. The descriptive statistics provided in this and the subsequent tables were obtained from a simulation run of 90,000 patent draws using the estimates in $\widehat{\theta}$ and the renewal fee schedule of the year 1988 (in 2003 dollars).[29] There are two results in this table that are worth highlighting. First, the distribution of per period returns is skew. The table shows that while about twenty percent of the patents have returns at age one below or equal \$1,000, only the top 3.5 percent had returns above \$100,000. This result is associated with the magnitude of the estimated parameters (μ and σ_R), and is consistent with the work Pakes (1986). The second result to note is that the left tail of the distribution of per period returns becomes thicker with patent age. Per period returns decrease and a significant proportion of patents become obsolete between age one, four, and seven. In particular, eleven percent of the patents at age one had per period returns below \$500; this proportion doubled (22.6 percent) by age four, and was 33.2 at age seven. The median of the distribution of per period returns at age one (\$4,080) also dropped by approximately half (\$2,447) by age four. The median per period return dropped to \$1,404 by age seven. This result reflects the importance of the estimated effects of obsolescence and depreciation factor characterized by the parameters γ^i and δ, respectively. These findings suggests a highly skewed distribution of patent value that depreciates significantly with patent age.

In Table A-2, we provide the percentiles of the distribution of patent values at age one, four and seven. The value of a patent is defined as the discounted present value of the per period returns (minus the renewal fees if the patent is renewed, and minus the transaction cost if the patent is sold) from age one, four or seven to the maximum legal length of patent protection. Column (1) presents the patent value as of age one. Twenty-five percent of the patents in a patent cohort had values below or equal to \$4,584. These patents accounted for 0.7 percent of the total value of patents in a cohort, while the bottom fifty-percent of the patents contributed a slightly less than 4 percent of the total value in a cohort. The value of the median patent at age one was \$16,184 and the mean \$76,598. The top 10 percent of the highest valued patents accounted for approximately 67 percent of the total value of a cohort. These findings confirm

[28] In particular, $(1-.97)*0+.97*(.90)=.873$, indicating that the per period patent returns of untraded patents in a given period decrease on average about 12.7 percent.

[29] Alternatively we could have used the average of the renewal fee schedule across all the year in the sample. We do not expect the results to be different.

a very skewed distribution of patent values. The rest of the columns in the table present the distribution of patent value at ages four and seven. A last result we want to highlight is that the distribution of patent value appears to become more skewed with patent age.

Parameters γ^e and σ^e jointly determine the potential benefits of transferring patents. Recall that there is a probability γ^e that after a patent trade the returns are zero. Alternatively, with probability $(1-\gamma^e)$ the patent returns are based on the improvement factor observed prior to the transfer of the patent. The low estimate of γ^e (about 0.0386) indicates that in most of the transfers the realized returns depend on the observed improvement factor of the potential buyer. In other words, the uncertainty of the potential benefits of acquiring a patent involves about a four percent premium.[30] The improvement factor that determines the returns of a buyer is a function of the parameter σ^e. A high σ^e is associated with higher expected improvement factors and thus higher patent returns. The estimate of σ^e implies that with probability 0.068 the patent returns of the buyer in a given period are higher than the ones of the current owner.[31] Furthermore, conditional on a buyer having higher returns than a current owner, the buyer's per period returns are 38 percent higher than the ones of the current owner. These findings suggest important benefits of reallocating the ownership of a patent from the original patentee to a new owner for whom the patent has a higher value.

Furthermore, the low estimate of ε implies that about half of a percentage point of the simulated proportion of active patents can be accounted for our specification of the measurement error in the trade of patents.[32]

Another important point to note is that the estimate of transaction costs τ is significant (about $\$5,500$). The estimated sunk transaction costs are equivalent to a third of the estimated average patent return at age one. In other words, the transaction costs represent about eight percent of the average present value of a patent at age one. Alternatively, the estimate of output uncertainty ($\gamma^e = .0386$) can also be interpreted as a variable component of the transaction cost. According to this interpretation, the parameter γ^e implies that that a variable component of transaction costs represents about four percent of the sale price of a patent. For the average value of a traded

[30] Given our model specification, an alternative interpretation could be that the variable component of the transaction cost represent about a four percent of the sale price of a patent.

[31] The unconditional probability that the patent per period returns of the buyer are higher than the ones of the current owners is computed $(1 - .0386) * .071 = .068$. The term $(1 - .0386)$ refers to the probability that the acquisition would be successful (i.e., the realized improvement factor of the patent owner is based on the improvement factor g^{*e} observed prior to the decision of selling the patent). The term .071 refers to the probability that, conditional on a successful acquisition, the per period return of a potential buyer will be higher than the per period return of current owner, which $Pr(g^{*e} > 1) = .071$. Recall that g^{*e} is exponentially distributed with parameter σ^e.

[32] An alternative is that half of a percentage point of the simulated proportion of the traded patents can be accounted for a transfer process that does not involve neither sunk nor variable costs of technology transfer.

patent, the variable component of the transaction costs is estimated at about $\$6,350$.[33] The magnitude of the transaction costs points at the possibility of significant selection effects into the trading of patents.

5.1 Remarks and identification robustness checks

As in most papers estimating patent value from patent renewal data we use the decision to renew patents, which requires a fee, to learn about the distribution of patent values. Renewal fees schedule are set by the U.S. Congress every calendar year. Our estimation strategy uses ten distinct schedules of patent renewal fees faced by the patents granted for the period 1988-1997. In each cohort, the fees are small during the early years of a patent's life and grow over time. Pakes and Simpson (1989) show that with unlimited variation in the renewal fee schedule and under mild regularity conditions the stochastic process generating returns to patent protection can be non-parametrically identified. While the magnitude of the fees tends to increase over the lifetime of patents, there may be insufficient variation in the renewal fees, which can limit what we can learn can be about the stochastic process generating the patent returns without imposing further functional form restrictions. In our sample, approximately fifty-five percent of the patents are let to expire before their maximum statutory legal length, suggesting that, at least for the large majority of patents, and those with not too distant returns, the surface of the distribution of initial returns is non-pararametrically identified. The recent evidence presented in Harhoff, Narin, Scherer, and Vopel (1999) also reassures our functional form choice to model the initial patent returns to patent protection. They use survey data on patented commercialization outcomes to show that a lognormal distribution provides a reasonable fit to the observed distribution of patent returns. Moreover, we did not find that a Pareto distribution, a typical alternative candidate, provided a better fit to the data than the lognormal distribution.

Nonetheless, we can limit the identification analysis of the renewal and trading decision choices to the observed range of renewal fee schedules and still identify important aspects of the stochastic process generating patents returns and gains from trade (Pakes and Simpson, 1989). To this end, our counterfactual experiments in the next section put special emphasize on the patents with initial returns somewhat above the 50-60 percentile of the distribution of initial returns, which is the proportion of patents that are not renewed to their maximum statutory legal length of patent protection. In particular, the next section shows that the ratio of the computed gains from trade relative to the patent's value for these patents is similar to the one obtained using all the

[33] We computed the variable component of the transaction costs using the average value of traded patents times 0386, which is $\$164,670 * .0386 = \$6,356$.

distribution of initial patent returns. We also plan to confront our estimates of the value of traded patents with data we have obtained on actual sale prices of patents sold at public auctions. We hope the evidence will provide reassurance that our results on the gains from trade are robust.

6 The market for patent rights and gains from trade

This section documents our empirical results for the volume of the market for patent rights and the gains from trade. We first compare the value of the traded and untraded patents. We then estimate the volume of the market for patent rights, which is the total value of the traded patents. We end the section by quantifying how much of this volume is gains from trading patents.

In Table A-3A, we present the distribution of patent value at age one for both the traded and the untraded patents. Traded patents are those that were sold at least once over their life cycle; untraded patents are the rest. The first result we want to highlight is that the mean of the distribution of the value of patents eventually traded is about three times larger than the mean of the distribution of value of the patents untraded ($\$164,670$ vs. $\$50,162$). The finding that traded patents have a higher value than the untraded ones is consistent across the percentiles of the distribution of the value of the traded and untraded patents. Furthermore, the difference between the value of the traded and the untraded patents increases over the percentiles of the distributions. Differences range from few thousand dollars at the fifth percentile of each distribution ($\$655$ and $\$4,881$ for untraded and traded patents, respectively) to more than a million dollars above the 99th percentile.

Table A-3B quantifies the total value from patent ownership changes. We consider two indicator variables for transactions in the market for patent rights: the proportion of patents traded and the total value of the traded patents. In the first row we report that 23.1 percent of the all the patents are traded at least once over their lifetime.[34] The next row shows that traded patents represent almost half (49.6%) of the total value of all patents as of age one. According to our estimates, the volume of the market for U.S small firm granted patents in between 1988-1997 is then approximately $2.08 billion.[35]

In Table A-4, we begin tackling the issue of selection in the trade of patents. The findings in section 2 suggest that the gross differences between the value of the traded and the untraded patents (and thus the proportion of the total value of all patents that traded patents represent)

[34] The number of traded patents is equal to the proportion of traded patents (0.23) times the number of patents in our estimation sample ($54,840$).

[35] This is the number of patents ($54,840$) $\times$ the proportion of patents traded (0.231) $\times$ the value at age one of a patent eventually traded ($\$164,670$).

can be due to potential gains from trade, but also because patents with higher per period returns are more likely to be traded. To address this issue, Column (1) presents the probability that a patent is traded obtained from simulations based on the estimated transaction cost. The bottom row shows the sample probability that a patent is traded over the life cycle, while the rest of the rows report the same probability across groups of patents formed based on the percentiles of the distribution of per period returns at age one. The sample probability that a patent is traded (0.231) is more than four times larger than the probability of patents with per period returns at age one in between the 10^{th} and 25^{th} percentile (0.05). Furthermore, patents with per period returns at age one below the 10^{th} percentile are virtually not traded. In contrast, the trade of patents is much more common at the highest percentiles of the distribution of per period returns at age one. For instance, about forty percent of the patents in between the 75^{th} and 80^{th} percentiles are traded over their life cycle. These results confirm the extent of the importance of selection in the trade of patents, as indicated in our motivating evidence presented in the data section. The results also suggest that the selection in trade can account for a sizable part of the estimated total value of patent ownership changes (i.e., total value of traded patents).

Table A-6A quantifies how much of the total value of patent ownership changes is due to gains from trade. The gains from trade were obtained by comparing the value of the patents actually traded with the value that the same patents would have obtained had the option to sell it to a potential buyer not been allowed for during a patent's lifetime. Column (1) and (2) report the average patent value at age one of traded patents with and without the possibility to transfer the patent rights to other firms. Column (3) we present the average realized gain from trade in patent markets as a percentage of the value of a traded patent across the percentiles of the distribution of patent returns at age one. Gains from trade for the average traded patent represent an additional 10 percent to the value that the original patentee could have earned by keeping the patent. We also show that the gains from trade as a percentage of patent value are heterogeneous across patents based on their initial returns at age one. The average gains from trade range from as low as zero, for traded patents with per period returns below the 10^{th} percentile, to as high as 11 percent, for patents with initial per period returns above the 95^{th} percentile. An advantage of structurally estimating the parameters of the model is that simple counterfactual experiments, such as the one of restricting the possibility to sell patents to other firms, allow us quantify the gains from trade in patent markets while at the same time account for selection into trade.

Interestingly, we have obtained that the value of the gains from trade as a proportion of the value of a traded patent is similar for traded patents with initial per period returns at levels where the identification analysis of the renewal and trading decision choices is based on the observed range of renewal fee schedule. In particular, for patents with per initial returns in between the 50^{th} and 65^{th} percentile, we find that about 9 percent of their value was due to gains from trade, which is comparable to the sample mean of the ratio of gains from trade relative to value of patents in the sample of all the patents. This finding provides some reassurance for our results on the relative magnitude of the gains from trade in patent markets.

In Table A-5, we assess in more detail the skewness and the distribution of the absolute value of the gains from trade. This table provides percentiles and cumulative percent of the distribution of the gains from trade of the traded patents. The result we want to highlight here is that the distribution of the gains from reallocating patent rights to new owners is highly skewed. The table reports that fifty percent of the traded patents had gains from trade below or equal $\$3,416$ (the mean of the value of the trade of patents is about $\$15,008$). These patents account for only 3.7 percent of the total value of the patents that experienced positive gains from trade in our simulations. On the other hand, the top ten percent of the patents with the largest gains from trade accounted for about 67 percent of the realized gains from trade. The findings indicate that only a small share of the traded patent contributes significantly to the realized gains from trade in patent markets.

The findings presented in this section indicate that the market for patent rights generates significant gains from trade for small firms, but that only a small proportion of the traded patents account for a significant share of the realized gains from trade in patent markets.

6.1 Remarks: Evidence from actual data on patent sales

It has been suggested to us that it would be nice to confront our results on the value of traded patents with some actual data on patent value from patent sales. It is challenging to find systematic data on the value of patents because most transactions are private. Moreover, our work suggests that the value of traded and untraded patents may quite different due to selection into trade.

To address this issue, we document the prices patent sales from the recent patent auctions organized and conducted by patent broker ICAP Ocean Tomo.[36] Ocean Tomo defines their business as matching buyers and sellers for the sale of patents and other intellectual property assets through the platforms of private brokerage and live auction.[37] The patent broker claims hundreds of successful transactions and over $150 million in sales over the last five years.

[36]We gladly acknowledge Ryan Lampe's generosity for providing access to the raw data of these auctions.

[37]Ocean Tomo was acquired by ICAP on June 16, 2009. See http://icappatentbrokerage.com/ and http://www.oceantomo.com/auctions.html for more details.

Ocean Tomo organized ten auctions for the period between April 2006 and July 2009. During this period, there were 672 lots including in total 1,867 patents up for sale. Approximately 40 percent of the sample lots were sold at the auctions, with an average of 2.4 patents in each of these lots. In total, there were 658 patents sold. We were able to obtain individual and small and large business entity status data for 500 of these patents. Among these patents, 73.6 percent had been originally issued to firms and 26.4 to individual inventors.[38] In addition, to parallel our analysis of the Ocean Tomo data with the sample used in our structural estimation, we focus on the sub-sample of these patents that were applied for and granted to U.S small firms. This represents 18.8 percent of the patents sold.[39] The average patent sale value was computed averaging the price paid for a lot over the number of patents included in the same lot. The mean value of the distribution of patent prices in the auction data is $140,845 (std. err. 31,070) and the median is about $66,935 (all values in 2003 U.S. dollars).[40] The distribution of patent value of the traded patents is skew, with a maximum patent value of $2.54 million.[41] In contrast, in the structural estimation, the mean and median value of a traded patent is $164,670 and $58,320, respectively.[42] Our interpretation is that the findings in the auction data suggest that our estimated results on the value of traded patents are reasonable.

7 Effects of transaction costs on the volume of patent trading and the gains from trade in patent markets

The remaining of our analysis studies the effect of lowering the sunk transaction costs on both the proportion of patents traded and the gains from trade in patent markets. To do this, we decrease the transaction costs by fifty percent, i.e., from $5,495 to $2,747, while keeping the rest of the parameters at their original estimated levels.

Several results emerge from this counterfactual experiment. First, we study the effects of reducing transaction costs on the probability that a patent is traded. Column (2) in Table A-4 reports the probability that a patent is traded across groups of patents based on their initial returns at age one, as well as the overall sample probability that a patent is traded. There are

[38] See Ughetto and Odasso (2010) for an earlier descriptive analysis of the first eight auctions organized by Ocean Tomo.

[39] Foreign and US small firm patents account for about 25 percent of the sample of patents sold. The rest of the patents are either individually owned patent or large firm patents

[40] We use the historical CPI index from the BLS to transfer current US dollars to US 2003 dollars.

[41] Comparable results were obtained in the whole sample of 658 patents, which includes both individual inventor and corporate patents. For this sample, The mean value of the distribution of patent prices is now $116,270 (std. err. 9,970) and the median is about $68,785. The maximum patent value was $4.34 million. All values in 2003 U.S. dollars.

[42] The confidence interval of the two means overlap.

two aspects worth highlighting here. The first one is that the probability that a patent is traded increases significantly, by six percentage points (from 0.231 to 0.296). The second one is that we show that the change in the probability that a patent is traded is heterogeneous across patents based on their returns at age one. Patents with returns in between the percentiles around the median of the distribution of initial returns at age one (25^{th} and 75^{th}) increased their probability of trade by about ten percentage points, while the rest of the patents were much less affected by the reduction in the transaction costs. In particular, for patents with per period returns at age one in the $50-65^{th}$ percentile, which are those with renewal and trading decision choices based on the observed range of renewal fee schedule, the reduction of transaction costs rises their probability of trade from 0.25 to 0.36. These findings provide some guidance to policy makers concerning the potential deleterious effects of transaction costs on the reallocation of patent rights.

Second, in Table A-7, we present the effects of lowering transaction costs on the gains resulting from the trading of patents. Column (1) depicts the percentage change in the value of the gains from trade for all patents and across groups of patents based on the initial returns at age one. In Column (2), we present the cumulative distribution of the additional gains from trade of the traded patents across groups of patents based on their initial returns at age one, and in Column (3) we show the cumulative distribution of the gains from trade across the same groups of patents in our baseline specification (i.e., prior to the reduction of the transaction costs.) There are two findings we especially want to highlight. First, in Column (1) we show that the gains from trading patents increase by about 10%. That is, decreasing the sunk component of the transaction cost by fifty percent has a limited increase in the total value of the gains from trade. Second, in Column (2) we show that the source of the additional gains from trade are patents with returns in between the median and the seventy-five percentile of the distribution of initial returns at age one. While the bottom seventy five percent of the traded patents with the lowest returns at age one account for 38% of the additional gains from trade, the patents with returns at age one in between the fifty and seventy percentile account for about 24% to the additional gains from trade. In contrast, in Column (3) we show that the patents with the lowest seventy-five percent of initial returns at age one accounted for just 6% of the total gains from trade in our benchmark specification. The findings suggest that the potential benefits associated with reductions in the transaction costs are limited, and that the patents that will benefit the most are those with per period returns somewhat above the median of the distribution of initial returns at age one.

Taken together, the results of the counterfactual of lowering transaction costs indicates that while the proportion of patents traded increases significantly, the realized gains from trade in patent markets associated with the reduction of transaction costs appear to be limited.

8 Conclusion

This paper develops and estimates a model of the transfer and renewal of patents that, under some assumptions, allows us to quantify the gains from trade and the transaction costs in the market for patents. The structural nature of the estimation strategy also help us assess the impact of lowering transaction costs on the proportion of patent traded and the gains from trade. The model is estimated using new data on the transfer of patents in a sample of patents applied for and granted to U.S. small firms.

Several findings emerge from our analysis. First, traded patents represent almost half (49.6%) of the total value of all patents. We showed that this magnitude is due to selection and gains from trade. Second, the gains from the trade account for 10 percent of the total value from patent ownership changes (i.e., in patent markets). Third, we showed that the distribution of the gains from trade was is skew, and that a small fraction of the traded patents accounted for a large part of the realized gains from trade. Fourth, we found that lowering transaction costs by fifty percent increased both the gains from trade in patent markets (10%) and the volume of trading (six percentage points).

The findings in this paper indicate that the market for patent rights generates significant gains from trade for small firms, but that only a small proportion of the traded patents accounts for a large share of the benefits. Our interpretation is that as long as small firms can appropriate part of the gains from patent in patent markets ex-post, this market increases the incentives to innovate ex-ante. The result that only a small proportion of patents will generate significant gains from trade may have important implications in the institutional setting and market structure of the intermediaries of the market for patent rights. Moreover, we found that transaction costs affect the workings of the market for patent rights by reducing the proportion of patents sold and generating at the same time a selection into trade. We see these results as a first step towards understanding the broader policy question of whether patent transactions increase total welfare.

At the same time, this paper may have some limitations. We likely examine a domain where the sale of patents may be considerably more useful than for the larger firms. In this respect our estimates may overestimate the overall gains from patent markets in the economy. Nevertheless, our model and empirical methodology is with ease portable to other related domains such as large firms. At the same time our sample does not contain the licensing of patents, and thus the benefits of the market for patent rights only account for the outright sale of all patent rights. A limitation of our setting is that it does not allow for the possibility that firms may choose to increase patenting as well as to specialize in what they excel. Our data also exhibited some limitations such as not linking the type and the name of firms acquiring patents, not reporting the sale prices, and not collecting the contractual terms in the patent transactions. These limitations provide opportunities for further research.

References

Acs, Z. J., and D. B. Audretsch (1988): "Innovation in Large and Small Firms: An Empirical Analysis," *American Economic Review*, 78 (4), 678–690.

Arora, A. (1997): ""Patents, Licensing, And Market Structure In Chemicals," *Research Policy*, 26, 391–403.

Arora, A., and M. Ceccagnoli (2006): "Patent Protection, Complementary Assets and Firms'Incentives for Technology Licensing," *Management Science*, 52, 293–308.

Arora, A., A. Fosfuri, and A. Gambardella (2001a): *Market for Technology: The Economics of Innovation and Corporate Strategy*. The MIT Press.

——— (2001b): "Specialized technology suppliers, international spillovers and investment: evidence from the chemical industry," *Journal of Development Economics*, 65, 31–51.

Arora, A., and A. Gambardella (2010): "Ideas for Rent: An Overview of Markets for Technology," *Industrial and Corporate Change*, 19 (3), 775–803.

Arrow, K. (1962): "Economic Welfare and the Allocation of Resources for Invention," in *Rate and Direction of Inventive Activity*. NBER and Princeton Univ. Press, Princeton, NJ.

——— (1983): "Innovation in Large and Small Firms," in *Entrepreneurship*, ed. by J. Ronen. Lexington Books, Lexington, MA.

Astebro, T. (2002): "Noncapital Investment Costs and the Adoption of CAD and CNC in U.S. metalworking industries," *RAND Journal of Economics*, 33 (4), 672–688.

Bental, B., and D. Peled (2002): "Quantitative growth effects of subsidies in a search theoretic R&D model," *Journal of Evolutionary Economics*, 12 (4), 397–423.

Berman, B. (2002): "From Tech Transfer to Joint Ventures, Part 1," Cafezine.

Bessen, J. (2008): "The value of U.S. patents by owner and patent characteristics," *Research Policy*, 37 (5), 932–945.

Degnan, S. A. (1998): "The Licensing Payoff from U.S. R and D," *Journal of the Licensing Executives International Society*, 33 (4).

Deng, Y. (2007): "Private value of European patents," *European Economic Review*, 51 (7), 1785–1812.

Dykeman, D., and D. Kopko (2004): "Recording Patent License Agreements in the USPTO," *Intellectual Property Today*, August, 18–19.

Figueroa, N., and C. J. Serrano (2011): "Patent Trading Flows of Small and Large Firms," University of Toronto Working Paper.

FTC (2011): "The Evolving IP Marketplace: Aligning Patent Notice and Remedies with Competition," Government Printing Office, Washington DC.

——— (2012): "Patent Assertion Entity Activities Workshop," Washington DC.

Galasso, A., M. Schankerman, and C. J. Serrano (2011): "Trading and Enforcing Patent Rights," *RAND Journal of Economics*, NBER Working Paper No. 17367, Forthcoming.

Gans, J., and S. Stern (2003): "The Product Market and the Market for "Ideas": Commercialization Strategies for Technology Entrepreneurs," *Research Policy*, 32, 333–350.

Gans, J. S., D. Hsu, and S. Stern (2008): "The Impact of Uncertain Intellectual Property Rights on the Market for Ideas: Evidence from Patent Grant Delays," *Management Science*, 54, 982–997.

Goffe, W. L., G. D. Ferrier, and J. Rogers (1994): "Global Optimization of Statistical Functions with Simulated Annealing," *Journal of Econometrics*, 60, 65–100.

Griliches, Z. (1992): "The Search of R&D Spillovers," *Scandinavian Journal of Economics*, 94, 29–47.

Grindley, P., and D. J. Teece (1997): "Managing Intellectual Capital: Licensing and Cross-licensing in Semiconductors and Electronics," *California Management Review*, 39.

Grossman, G. M., and E. Helpman (1991): "Quality Ladders in the Theory of Growth," *The Review of Economic Studies*, 58 (1), 43–61.

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

Hall, B. H., and R. H. Ziedonis (2001): "The patent paradox revisited: an empirical study of patenting in the US semiconductor industry, 1979-1995," *RAND Journal of Economics*, 32 (1), 101–128.

Harhoff, D., F. Narin, F. M. Scherer, and K. Vopel (1999): "Citation Frequency and the Value of Patented Inventions," *Review of Economics and Statistics*, 81, 511–515.

Jovanovic, B., and P. L. Rousseau (2008): “Mergers as Reallocation,” *The Review of Economics and Statistics*, 90 (4), 765–776.

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

Lemley, M. A., and C. Shapiro (2005): “Probabilistic Patents,” *Journal of Economic Perspectives*, 19 (2), 75–98.

McFadden, D. (1989): “A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration,” *Econometrica*, 57, 995–1026.

Merrill, S. A., R. C. Levin, and M. B. Myers (2004): *A Patent System for the 21st Century.* The National Academies Press, Washington D.C.

Pakes, A. (1986): “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks,” *Econometrica*, 54 (4), 755–784.

Pakes, A., and D. Pollard (1989): “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 57, 1027–1058.

Pakes, A., and M. Schankerman (1984): “The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources,” in *R&D, Patents and Productivity*, ed. by Z. Griliches, pp. 98–112. University of Chicago Press, Chicago.

Pakes, A., and M. Simpson (1989): “Patent Renewal Data,” *Brookings Papers on Economic Activity. Microeconomics*, 1989, 331–410.

Putnam, J. (1996): “The Value of International Patent Rights,” Ph.D. thesis, Yale University.

Rosenberg, N. (1996): “Uncertainty and Technological Change,” in *The Mosaic of Economic Growth*, ed. by R. Landau, T. Taylor, and G. Wright. Stanford University Press.

Saha, A., H. A. Love, and R. Schwart (1994): “Adoption of Emerging Technologies Under Output Uncertainty,” *American Journal of Agricultural Economics*, 76, 836–846.

Schankerman, M. (1998): “How Valuable is Patent Protection? Estimates by Technology Field,” *RAND Journal of Economics*, 29 (1), 77–107.

Schankerman, M., and A. Pakes (1986): “Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period,” *Economic Journal*, 96, 1052–1076.

Serrano, C. J. (2006): “The Market for Intellectual Property: Evidence from the Transfer of Patents,” Ph.D. thesis, University of Minnesota.

——— (2010): "The Dynamics of the Transfer and Renewal of Patents," *RAND Journal of Economics*, 41, 686–708.

Shapiro, C. (2001): "Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting," in *Innovation Policy and the Economy*, vol. 1, pp. 119–150. National Bureau of Economic Research, Inc.

Silveira, R., and R. Wright (2010): "Search and the Market for Ideas," *Journal of Economic Theory*, 145 (4), 1550–1573.

Teece, D. J. (1977): "Technology Transfer by Multinational Firms: The Resource Cost of Transferring Technological Know-how," *Economic Journal*, 87, 242–261.

——— (1986): "Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy," *Research Policy*, 15, 285–305.

Ughetto, E., and C. Odasso (2010): "Patent Value: Seller and Bidder Perspectives," Working Paper, Politecnico di Torino.

Appendix: Proofs, Tables and Figures

Lemma 1. The value function $V_a(x,y)$ is continuous and increasing in the current return of the holder of the patent, x, and the return factor of the potential buyer, y. The option value $EV_{a+1}(x',y'|x,y,d)$ is decreasing in a.

Proof. The option value $EV_{a+1}(x',y'|x,y,d)$ depends on whether or not the patent has been sold, kept, or renewed at age a. If the patent is sold $(d=1)$, $EV_{a+1}(x',y'|x,y,d) = EV_{a+1}(x',y'|y,S)$; if the patent is kept by the current owner, $EV_{a+1}(x',y'|x,y,d) = EV_{a+1}(x',y'|x,K)$; the option value $EV_{a+1}(x',y'|x,y,d) = 0$ whenever the patent is allowed to expire. For convenience, let us consider the option value be $EV_{a+1}(x',y'|z,d)$ where $z \in \{y,x,0\}$ depending on whether the patent was sold (y), kept (x), or allowed to expire (0). The proof is separated into two parts. Parts (i) shows that the option value is continuous and increasing with z; part (ii) proves that the option value is decreasing with a. Both proofs are by induction on a. Let L be the maximum legal length of patent protection.

Part (i). Since $E[V_{L+1}(x',y'|z)] = 0$ for all z, the initial condition of the inductive argument is satisfied. Now, suppose that the result holds for $a+1$, it suffices to show that the result also holds for a. The hypothesis of the inductive argument is that $E[V_{a+2}(x',y'|z)]$ is continuous and increasing in z. Recall that $V_{a+1}(x,y) = \max\{V^S_{a+1}(x,y), V^K_{a+1}(x,y), 0\}$ and that $V^S(x,y) = y - c_a - \tau + \beta E[V_{a+2}(x',y'|y], V^K(x,y) = x - c_a + \beta E[V_{a+2}(x',y'|x]$. So, $V_{a+1}(x,y)$ is also continuous and weakly increasing in $z \in \{y,x,0\}$. Finally, we want to show that if $V_{a+1}(x,y)$ is also continuous and increasing in $z \in \{y,x,0\}$, then $E_{g^e}[V_{a+1}(x',y'|z)]$ is also continuos and increasing in z. Since the joint distribution of the internal and external returns $F(.)$ is independent of z,we can show that if $z_1 \geq z_2$, then $E[V_{a+1}(x',y'|z_1)] \geq E[V_{a+1}(x',y'|z_2)]$.

$$\begin{aligned}
E[V_{a+1}(x',y'|z_1)] &= \int_{u^e}\int_{u^i} V_{a+1}(u^i z, u^i z u^e_a|z_1) F_{g^i}(\mathrm{d}u^i) F_{g^i}(\mathrm{d}u^e) \\
&\geq \int_{u^e} V_{a+1}(u^i z, u^i z u^e_a|z_2) F_{g^i}(\mathrm{d}u^i) F_{g^i}(\mathrm{d}u^e) \\
&= E[V_{a+1}(x',y'|z_2)]
\end{aligned}$$

To prove continuity, take any $z \in R^+$, $E_{g^e}[V_{a+1}(x',y'|z)]$ will be continuous at z if for every sequence $\{z_n\}$ such that $\lim(z_n) = z$, we can show that $\lim\limits_{z_n\to z}(E[V_{a+1}(x',y'|z_n)]) = E[V_{a+1}(x',y'|z)]$. Since $F_{g^e}(.)$ and $F_{g^i}(.)$ are independent of z.

$$\begin{aligned}
\lim_{z_n\to z} E[V_{a+1}(x',y'|z_n)] &= \lim_{z_n\to z}\int_{u^e}\int_{u^i} V_{a+1}(u^i z_n, u^i z_n u^e_a|z_n) F_{g^i}(\mathrm{d}u^i) F_{g^i}(\mathrm{d}u^e) \\
&= \int_{u^e}\int_{u^i} \lim_{z_n\to z}(V_{a+1}(u^i z_n, u^i z_n u^e_a|z_n)) F_{g^i}(\mathrm{d}u^i) F_{g^i}(\mathrm{d}u^e) \\
&= \int_{u^e}\int_{u^i} V_{a+1}(u^i z, u^i z u^e_a|z)) F_{g^i}(\mathrm{d}u^i) F_{g^i}(\mathrm{d}u^e)
\end{aligned}$$

where second step follows, in particular, because $F(.)$ is independent of z, and the last step follows because $V_{a+1}(u^i z, u^i z u^e_a|z)$ is continuos in z (i.e., $\lim\limits_{z_n\to z}(V_{a+1}(u^i z_n, u^i z_n u^e_a|z_n)) = V_{a+1}(u^i z, u^i z u^e_a|z)))$.

Part (ii). We want to show that $E[V_{a+1}(x',y'|z)]$ is weakly decreasing in a. The initial condition of the inductive argument requires to show that $E[V_{L+1}(x',y'|z)] \leq E[V_L(x',y'|z)]$. Since patents are active for L periods, then $E[V_{L+1}(x',y'|z)] = 0$. Moreover, by definition of $V_L(x',y'|z) \geq 0$, so it must be the case that $E[V_L(x',y'|z)]$. The induction hypothesis is that $E[V_{a+2}(x',y'|z)] \leq E[V_{a+1}(x',y'|z)]$. It suffices to show that $E[V_{a+1}(x',y'|z)] \leq E[V_a(x',y'|z)]$. Recall that $V_{a+1}(x',y'|z) = \max\{y' - c_{a+1} - \tau + \beta E[V_{a+2}(x'',y''|y')], x' - c_{a+1} + \beta E[V_{a+2}(x'',y''|x')], 0\}$

$\leq \max\{y' - c_a - \tau + \beta E[V_{a+1}(x'',y''|y)], x' - c_a + \beta E[V_{a+1}(x'',y''|x')]\} = V_a(x',y'|z)$, where the inequality holds because $c_a \leq c_{a+1}$ and because of the induction hypothesis. Finally, we must show that

$$\begin{aligned} E[V_{a+1}(x',y'|z)] &= \int_{u^e}\int_{u^i} V_{a+1}(u^i z, u^i z u_a^e|z) F_{g^i}(\mathrm{d}u^i) F_{g^e}(\mathrm{d}u^e) \\ &\leq \int_{u^e}\int_{u^i} V_a(u^i z, u^i z u_a^e|z) F_{g^i}(\mathrm{d}u^i) F_{g^e}(\mathrm{d}u^e) \\ &= E[V_a(x',y'|z)] \end{aligned}$$

where the inequality follows because $V_{a+1}(x',y'|z) \leq V_a(x',y'|z)$, $F_{g^e}(.)$ is independent of a, and F_{g^i} is such that given z the $\Pr(x' > x|z,a)$ decreases with a. ■

Lemma 2 For a fixed τ, g^e, g^i and a , we can show that: (a) If $V_a^S(x,y) - V_a^K(x,y) \geq 0$, then $V_a^S(x,y) - V_a^K(x,y)$ is increasing with x, (b) the $V_a^K(x,y) - V_a^E(x,y)$ is increasing with x.

Proof. This Lemma will be useful in the proofs of Proposition 1 and 2. For convenience, let me define $\widetilde{V}_a(x)$ as $V_a^K(x,y)$, that is

$$\widetilde{V}_a(x) = x - c_a + \beta E[V_{a+1}(g_a^i x, g_a^i x g_a^e)]$$

So, since $V_a^S(x,y) = V_a^K(x,y) - \tau$, then $V_a^S(x,y) = \widetilde{V}_a(y) - \tau$.

Let us start proving part (a) of the Lemma. That is the case when the owner of the patent is indifferent between selling the patent or keep it (i.e., any per period revenue such $x > \widehat{x}_a$).

The proof is by induction on a.

1) Suppose $a = L$ (i.e., the last period).

$$V_L^S(x,y) - V_L^K(x,y) = xg^e - \tau - x = x(g^e - 1) - \tau$$

Recall that $V_L^S(x,y) - V_L^K(x,y) \geq 0$, so it must the case that $g^e > 0$. Then, $V_L^S(x,y) - V_L^K(x,y)$ is decreasing with x.

2) Suppose that the result holds for $a+1$. So, the induction hypothesis is that for a fixed x, the $[V_{a+1}^S(x,y) - V_{a+1}^K(x,y)]$ is increasing with x. Then, it suffices to show that for a fixed x, the $[V_a^S(x,y) - V_a^K(x,y)]$ is increasing with x. Without loss of generality we can consider that $y > x$.

We can rewrite $[V_a^S(x,y) - V_a^K(x,y)]$ as

$$\begin{aligned} V_a^S(x,y) - V_a^K(x,y) &= y - c_a - \tau + \beta E[V_{a+1}(g_a^i y, g_a^i y g_a^e)] \\ &\quad -x + c_a + \tau - \beta E[V_{a+1}(g_a^i x, g_a^i x g_a^e)] \\ &= (y - x - \tau) + \beta E[V_{a+1}(g_a^i y, g_a^i y g_a^e) - V_{a+1}(g_a^i x, g_a^i x g_a^e)] \end{aligned}$$

Since $y > x$ and $y = xg_{a-1}^e$ then It is obvious to show that $(y - x - \tau)$ is increasing with x. Let us look at the second term $E_{g^e}[V_{a+1}(g_a^i y, g_a^i y g_a^e) - V_{a+1}(g_a^i x, g_a^i x g_a^e)]$. Since F_{g^e} and F_{g^i} are independent of x, it just remain to be shown that $[V_{a+1}(g_a^i y, g_a^i y g_a^e) - V_{a+1}(g_a^i x, g_a^i x g_a^e)]$ is increasing with x.

We can rewrite the equivalent of the induction hypothesis as

$$[V_{a+1}^S(x,y) - V_{a+1}^K(x,y)] = \widetilde{V}_{a+1}(y) - \tau - \widetilde{V}_{a+1}(x)$$

which is weakly increasing with x.

There are three cases to study.

Case (1): fix x and consider a realization of g_a^i and g_a^e such $V_{a+1}(g_a^i y, g_a^i y g_a^e) = V_{a+1}^S(g_a^i y, g_a^i y g_a^e)$, and $V_{a+1}(g_a^i x, g_a^i x g_a^e) = V_{a+1}^K(g_a^i x, g_a^i x g_a^e)$. That is in K region with $(g_a^i x, g_a^i x g_a^e)$ and in S region with $(g_a^i y, g_a^i y g_a^e)$. Let $\lambda = g_{a-1}^e g_a^e$. Then,

$$\begin{aligned}
V_{a+1}(g_a^i y, g_a^i y g_a^e) - V_{a+1}(g_a^i x, g_a^i x g_a^e) &= V_{a+1}^S(g_a^i y, g_a^i y g_a^e) - V_{a+1}^K(g_a^i x, g_a^i x g_a^e) \\
&= V_{a+1}^S(g_a^i x g_{a-1}^e, g_a^i x g_{a-1}^e g_a^e) - V_{a+1}^K(g_a^i x, g_a^i x g_a^e) \\
&= V_{a+1}^S(\lambda g_a^i x, \lambda g_a^i x g_a^e) - V_{a+1}^K(g_a^i x, g_a^i x g_a^e) \\
&= \widetilde{V}_{a+1}(g_a^i x g_{a-1}^e g_a^e) - \tau - \widetilde{V}_{a+1}(g_a^i x) \\
&= \widetilde{V}_{a+1}(\lambda g_a^i x) - \tau - \widetilde{V}_{a+1}(g_a^i x)
\end{aligned}$$

where the last expression is an increasing transformation of the equivalent of the induction hypothesis (i.e., $\widetilde{V}_{a+1}(y) - \tau - \widetilde{V}_{a+1}(x)$). Therefore, $V_{a+1}(g_a^i y, g_a^i y g_a^e) - V_{a+1}(g_a^i x, g_a^i x g_a^e)$ is increasing with x.

Case (2): fix x and consider a realization of g_a^i and g_a^e such that, in K region with $(g_a^i y, g_a^i y g_a^e)$, in K region with $(g_a^i x, g_a^i x g_a^e)$. Let $\lambda = g_{a-1}^e$. Then,

$$\begin{aligned}
V_{a+1}(g_a^i y, g_a^i y g_a^e) - V_{a+1}(g_a^i x, g_a^i x g_a^e) &= V_{a+1}^K(g_a^i y, g_a^i y g_a^e) - V_{a+1}^K(g_a^i x, g_a^i x g_a^e) \\
&= V_{a+1}^K(g_a^i x g_{a-1}^e, g_a^i x g_{a-1}^e g_a^e) - V_{a+1}^K(g_a^i x, g_a^i x g_a^e) \\
&= \widetilde{V}_{a+1}(g_a^i x g_{a-1}^e) - \widetilde{V}_{a+1}(g_a^i x) \\
&= \widetilde{V}_{a+1}(\lambda g_a^i x) - \widetilde{V}_{a+1}(g_a^i x)
\end{aligned}$$

where the last expression is an increasing transformation of the equivalent of the induction hypothesis (i.e., $\widetilde{V}_{a+1}(y) - \tau - \widetilde{V}_{a+1}(x)$). Therefore, $V_{a+1}(g_a^i y, g_a^i y g_a^e) - V_{a+1}(g_a^i x, g_a^i x g_a^e)$ is weakly increasing with x.

Case (3): fix x and consider a realization of g_a^e and g_a^i such that, in S region with $(g_a^i y, g_a^i y g_a^e)$ and in S region with $(g_a^i x, g_a^i x g_a^e)$. Let $\lambda = g_{a-1}^e$. Then,

$$\begin{aligned}
V_{a+1}(g_a^i y, g_a^i y g_a^e) - V_{a+1}(g_a^i x, g_a^i x g_a^e) &= V_{a+1}^S(g_a^i y, g_a^i y g_a^e) - V_{a+1}^S(g_a^i x, g_a^i x g_a^e) \\
&= V_{a+1}^S(g_a^i x g_{a-1}^e, g_a^i x g_{a-1}^e g_a^e) - V_{a+1}^S(g_a^i x, g_a^i x g_a^e) \\
&= \widetilde{V}_{a+1}(g_a^i x g_{a-1}^e g_a^e) - \widetilde{V}_{a+1}(g_a^i x g_a^e) \\
&= \widetilde{V}_{a+1}(\lambda g_a^i x g_a^e) - \widetilde{V}_{a+1}(g_a^i x g_a^e)
\end{aligned}$$

where the last expression is an increasing transformation of the equivalent of the induction hypothesis (i.e., $\widetilde{V}_{a+1}(y) - \tau - \widetilde{V}_{a+1}(x)$). Therefore, $V_{a+1}(g_a^i y, g_a^i y g_a^e) - V_{a+1}(g_a^i x, g_a^i x g_a^e)$ is increasing with x.

Let us now prove part (b) of the Lemma. We want to show that the $V_a^K(x,y) - V_a^E(x,y)$ is increasing with x. We know that $V_a^E(x,y) = 0$. So, $V_a^K(x,y) - V_a^E(x,y) = V_a^K(x,y)$. For a fixed g^e and g_a^i, Lemma 1 shows that $V_a^K(x,y)$ is increasing with x. Therefore, $V_a^K(x,y) - V_a^E(x,y)$ is also increasing with x. ■

Proposition 1. For each age, there exist cutoff rules $\widehat{x}_a(\theta)$ and $\widehat{g}_a^e(x,\theta) > 1$ such that patents with an improvement factor above $\widehat{g}_a^e(x,\theta)$ will be sold. Among the patent holders that met a potential buyer with improvement factor below $\widehat{g}_a^e(x,\theta)$, the patents with per period returns x above $\widehat{x}_a(\theta)$ will be renewed, and those with returns x below $\widehat{x}_a(\theta)$ will be retained by their current owners.

Proof. We first show that there exist a cutoff rule $\widehat{g}_a^e(x,\tau)$ defined as the improvement factor that makes a potential seller indifferent between whether or not to sell the patent. Two cases must be analyzed: (a) The seller is indifferent between selling and letting the patent expire (i.e., $V_a^S(x, x\widehat{g}_a^e(x,\tau)) = 0$); and (b) The seller is indifferent between selling and keeping the patent (i.e., $V_a^S(x, x\widehat{g}_a^e(x,\tau)) = V_a^K(x, x\widehat{g}_a^e(x,\tau))$).

Let us start with case (a). Lemma 1 implies that, for a fixed a, $\widetilde{V}_a(x)$ is an increasing function of x. We can show that for any $\tau > 0$, if $g^e = 0$, $\widetilde{V}_a(xg^e) - \tau < 0$. We can also show that for any $\tau > 0$ there exist a sufficiently high g^e, for example $\overline{g^e}$, that $\widetilde{V}_a(x\overline{g^e}) - \tau > 0$. Therefore, by Bolzano's theorem, there exist an improvement factor $\widehat{g}_a^e(x,\tau)$ such that

$$\begin{aligned}
\widetilde{V}_a(x\widehat{g}_a^e(x,\tau)) - \tau &= 0 \\
V_a^S(x, x\widehat{g}_a^e(x,\tau)) &= 0
\end{aligned}$$

In case (b) the potential seller is indifferent between selling and keeping the patent. That is $V_a^S(x,y) = V_a^K(x,y)$, which implies that

$$\widetilde{V}_a(x\widehat{g}_a^e(x,\tau)) - \tau = \widetilde{V}_a(x)$$

Now, let us consider a $\tau > 0$ and a $g^e = 1$. We can show that $V_a^S(x,y) - V_a^K(x,y) < 0$ because $V_a^S(x,y) = V_a^S(x,x) - \tau = V_a^K(x,y) - \tau$. On the other hand, for fixed τ and fixed x, Lemma 1 shows that $V_a^S(x,y)$ is increasing in y. So, for sufficiently high g^e we can show that $V_a^S(x,y) - V_a^K(x,y) > 0$. Finally, by Lemma 2 we know that for a fixed τ, the difference $(V_a^S(x,y) - V_a^K(x,y))$ is increasing in x. By Bolzano's theorem, there exist a $\widehat{g}_a^e(x,\tau)$ such that $V_a^S(x, x\widehat{g}_a^e(x,\tau)) = V_a^K(x, x\widehat{g}_a^e(x,\tau))$, which is that

$$\widetilde{V}_a(x\widehat{g}_a^e(x,\tau)) - \tau = \widetilde{V}_a(x)$$

The second part of the proof shows that there exist a per period return $\widehat{x}_a(\theta)$ such that patents with per period return $x < \widehat{x}_a(\theta)$ will be allowed to expire while the rest of the patents will be renewed. The cutoff rule $\widehat{x}_a(\theta)$ is defined as the patent revenue that makes a patent owner indifferent between keeping the patent and letting it expire (i.e., $V_a^K(\widehat{x}_a, y) = 0$). Recall that $V_a^K(x,y) = x + \beta E[V_{a+1}(x', y')|x, y, K] - c_a$ and the function $V_a^K(x,y)$ is increasing with x by Lemma 1. Moreover, if $x = 0$, $V_a^K(x,y) < 0$; and if x is positive and sufficiently large $V_a^K(x,y) > 0$. Therefore, Bolzano's theorem implies that there is a $\widehat{x}_a$ such that $V_a^K(\widehat{x}_a, y) = 0$). ■

Proposition 2. The following properties characterize the functions $\widehat{x}_a(\theta)$ and $\widehat{g}_a^e(x,\theta)$:

- **(a)** The function $\widehat{x}_a(\theta)$ is increasing with patent age a at the renewal dates.
- **(b)** If the transaction cost τ is positive, then
  - **(b.1)** For a fixed patent age a, the function $\widehat{g}_a^e(x,\theta)$ is monotonically decreasing with x
  - **(b.2)** For a fixed patent return x, the function $\widehat{g}_a^e(x,\theta)$ is monotonically increasing with patent age a

Proof. Part (a). We want to show that $\widehat{x}_a(\theta)$ is increasing with patent age a at the renewal dates. From proposition 1 we know that that there exist a renewal date a and patent revenue $\widehat{x}_a$ such that $V_a^K(\widehat{x}_a, \widehat{x}_a g^e) = 0$. Similarly we can obtain that $V_{a+1}^K(\widehat{x}_{a+1}, \widehat{x}_{a+1}g^e) = 0$ because the schedule of renewal fees is increasing with a. Note that $V_{a+1}^K(\widehat{x}_{a+1}, \widehat{x}_{a+1}g^e) = V_a^K(\widehat{x}_a, \widehat{x}_a g^e)$. Because the function $V_a^K(.,.)$ is increasing with a and x (see Lemma 1). then it must be the case that $\widehat{x}_{a+1} \geq \widehat{x}_a$ for the above equality to hold.

Part (b.1) [Selection Effect]. We want to show that the cutoff rule $\widehat{g}_a^e(x,\tau)$ is decreasing with x. For convenience, let me define$V_a^K(x,y)$ as $\widetilde{V}_a(x)$, that is

$$\widetilde{V}_a(x) = x - c_a + \beta E[V_{a+1}(g_a^i x, g_a^i x g_a^e)]$$

So, since $V_a^S(x,y) = V_a^K(x,y) - \tau$, then $V_a^S(x,y) = \widetilde{V}_a(y) - \tau$.

Let us first consider when the owner of the patent is indifferent between selling the patent or let it allow to expire. That is any per period revenue such $x \leq \widehat{x}_a$. Proposition 1 implies that there is a cutoff rule $\widehat{g}^e_a(x,\tau)$ and that is defined as follows.

$$\begin{aligned} V^S_a(x, x\widehat{g}^e_a(x,\tau)) &= 0 \\ \widetilde{V}_a(x\widehat{g}^e_a(x,\tau)) - \tau &= 0 \end{aligned}$$

From Lemma 1 we know that for a fixed g^e the function $V^S_a(x, xg^e) = \widetilde{V}_a(xg^e) - \tau$ is increasing with x. So, for the above equality to hold, it must be the case that if x increases then $\widehat{g}^e_a(x,\tau)$ decreases.

Let us now consider the case when the owner of the patent is indifferent between selling the patent or keep it. That is any per period revenue such $x > \widehat{x}_a$.

The proof is by induction on a.

1) Suppose $a = L$ (i.e., the maximum legal length of patent protection). A seller is indifferent between selling and keeping a patent if

$$\begin{aligned} V^S_L(x,y) &= V^K_L(x,y) \\ x\widehat{g}^e_a(x,\tau) - \tau &= x \end{aligned}$$

And we can show that $\widehat{g}^e_a(x,\tau)$ is decreasing with x

$$\widehat{g}^e_a(x,\tau) = 1 + \frac{\tau}{x}$$

2) Suppose now that $\widehat{g}^e_a(x,\tau)$ is decreasing with x for $a+1$. We want to show that the result holds for a. Notice that the induction hypothesis (i.e., $\widehat{g}^e_a(x,\tau)$ is decreasing with x for $a+1$) is equivalent to showing that for a fixed x, the $[V^S_{a+1}(x,y) - V^K_{a+1}(x,y)]$ is weakly increasing in x. Therefore, in order to prove that $\widehat{g}^e_a(x,\tau)$ in decreasing in x for age a, it suffices to show that for a fixed x, the difference $[V^S_a(x,y) - V^K_a(x,y)]$ is increasing in x.

Note that $\widehat{g}^e_a(x,\tau)$ for any a is only defined for x such that $\widehat{g}^e_a(x,\tau) \geq 1 \ \forall a \in [1, L]$. So, without loss of generality we can consider $y > x$. Finally, by Lemma 2, see especially part (2) of the proof, we show that for a fixed τ, g^e, and g^i the $[V^S_a(x,y) - V^K_a(x,y)]$ is increasing with x. As argued above, this results is equivalent to showing that the function $\widehat{g}^e_a(x,\tau)$ is decreasing with x.

Part (b.2) [Horizon Effect]. We want to show that for any $x > 0$ and $\tau > 0$, the cutoff rule $\widehat{g}^e_a(x,\tau)$ is increasing with a. For convenience, let me define $\widetilde{V}_a(x)$ as $V^K_a(x,y)$ as

$$\widetilde{V}_a(x) = x - c_a + \beta E_{g^e}[V_{a+1}(\delta x, \delta x g^e_a)]$$

So, since $V^S_a(x,y) = V^K_a(x,y) - \tau$, then $V^S_a(x,y) = \widetilde{V}_a(y) - \tau$.

Showing that for a fixed x and τ the cutoff rule $\widehat{g}^e_a(x,\tau)$ is weakly increasing with a is equivalent to proving that for a fixed x and g^e the difference $[V^S_a(x,y) - V^K_a(x,y)]$ is decreasing with a. That is,

$$[V^S_a(x,y) - V^K_a(x,y)] > [V^S_{a+1}(x,y) - V^K_{a+1}(x,y)]$$

And rearranging the induction hypothesis we obtain,

$$\begin{aligned}
\widetilde{V}_a(y) - \tau - \widetilde{V}_a(x) &> \widetilde{V}_{a+1}(y) - \tau - \widetilde{V}_{a+1}(x) \\
\widetilde{V}_a(y) - \widetilde{V}_a(x) &> \widetilde{V}_{a+1}(y) - \widetilde{V}_{a+1}(x) \\
y - c_a + \beta E[V_{a+1}(g_a^i y, g_a^i y g_a^e)] - x + c_a - \beta E[V_{a+1}(g_a^i x, g_a^i x g_a^e)] &> y - c_{a+1} + \beta E[V_{a+2}(g_{a+1}^i y, g_{a+1}^i y g_{a+1}^e)] \\
&\quad -x + c_{a+1} - \beta E[V_{a+2}(g_{a+1}^i x, g_{a+1}^i x g_{a+1}^e)] \\
y - x + \beta E[V_{a+1}(g_a^i y, g_a^i y g_a^e)] - \beta E[V_{a+1}(g_a^i x, g_a^i x g_a^e)] &> y - x + \beta E[V_{a+2}(g_{a+1}^i y, g_{a+1}^i y g_{a+1}^e)] \\
&\quad -\beta E[V_{a+2}(g_{a+1}^i x, g_{a+1}^i x g_{a+1}^e)] \\
E[V_{a+1}(g_a^i y, g_a^i y g_a^e)] - E[V_{a+1}(g_a^i x, g_a^i x g_a^e) &> E[V_{a+2}(g_{a+1}^i y, g_{a+1}^i y g_{a+1}^e)] \\
&\quad -E[V_{a+2}(g_{a+1}^i x, g_{a+1}^i x g_{a+1}^e)]
\end{aligned}$$

Without loss of generality, consider per period revenues x and y, $y > x$, such that $[V_a^S(x,y) - V_a^K(x,y)] \geq 0$.

The proof is by induction.

1) Suppose $a + 1 = L$, so $a = L - 1$. We want to show that

$$[V_L^S(x,y) - V_L^K(x,y)] > [V_{L+1}^S(x,y) - V_{L+1}^K(x,y)]$$

Let us start defining

$$V_L^S(x,y) - V_L^K(x,y) = \widetilde{V}_a(y) - \tau - \widetilde{V}_a(x) = y - \tau - x$$

$$\begin{aligned}
V_{L-1}^S(x,y) - V_{L-1}^K(x,y) &= y + \beta E[V_L(g_{L-1}^i y, g_{L-1}^i y g_{L-1}^e)] - \tau - x + \beta E[V_L(g_{L-1}^i x, g_{L-1}^i x g_{L-1}^e)] \\
&= y - \tau - x + \beta E[V_L(g_{L-1}^i y, g_{L-1}^i y g_{L-1}^e) - V_L(g_{L-1}^i x, g_{L-1}^i x g_{L-1}^e)]
\end{aligned}$$

In Lemma 1 we showed that for a fixed a, $V_a(x,y)$ was weakly increasing in x. So, it must be the case that $V_L(g_{L-1}^i y, g_{L-1}^i y g_{L-1}^e) \geq V_L(g_{L-1}^i x, g_{L-1}^i x g_{L-1}^e)$. The expectation of a random variable that is larger or equal than zero is also larger or equal than zero. So, $E[V_L(g_{L-1}^i y, g_{L-1}^i y g_{L-1}^e) - V_L(g_{L-1}^i x, g_{L-1}^i x g_{L-1}^e)] \geq 0$. Then

$$\begin{aligned}
V_{L-1}^S(x,y) - V_{L-1}^K(x,y) &= y - \tau - x + \beta E[V_L(g_{L-1}^i y, g_{L-1}^i y g_{L-1}^e) - V_L(g_{L-1}^i x, g_{L-1}^i x g_{L-1}^e)] \\
&\geq y - \tau - x \\
&= V_L^S(x,y) - V_L^K(x,y)
\end{aligned}$$

Therefore, $[V_L^S(x,y) - V_L^K(x,y)] > [V_{L+1}^S(x,y) - V_{L+1}^K(x,y)]$.

2) Now, suppose that the relationship holds for $a + 1$. So, the induction hypothesis is

$$\begin{aligned}
[V_{a+1}^S(x,y) - V_{a+1}^K(x,y)] &> [V_{a+2}^S(x,y) - V_{a+2}^K(x,y)] \\
\widetilde{V}_{a+1}(y) - \tau - \widetilde{V}_{a+1}(x) &> \widetilde{V}_{a+2}(y) - \tau - \widetilde{V}_{a+2}(x) \\
\widetilde{V}_{a+1}(y) - \widetilde{V}_{a+1}(x) &> \widetilde{V}_{a+2}(y) - \widetilde{V}_{a+2}(x)
\end{aligned}$$

We will show that it holds for a, that is

$$[V_a^S(x,y) - V_a^K(x,y)] > [V_{a+1}^S(x,y) - V_{a+1}^K(x,y)]$$

We know that

$$\begin{aligned}
V_a^S(x,y) - V_a^K(x,y) &= \widetilde{V}_a(y) - \tau - \widetilde{V}_a(x) \\
&= y - \tau + \beta E[V_{a+1}(g_a^i y, g_a^i y g_a^e)] - x - \beta E[V_{a+1}(g_a^i x, g_a^i x g_a^e)] \\
&= y - \tau - x + \beta(E[V_{a+1}(g_a^i y, g_a^i y g_a^e)] - E[V_{a+1}(g_a^i x, g_a^i x g_a^e)]) \\
&> y - \tau - x + \beta(E[V_{a+2}(g_{a+1}^i y, \delta y g_{a+1}^e)] - E[V_{a+2}(g_{a+1}^i x, \delta x g_{a+1}^e)]) \\
&= y - \tau - c_a + \beta E[V_{a+2}(g_{a+1}^i y, g_{a+1}^i y g_{a+1}^e)] - x + c_a - \beta E[V_{a+2}(g_{a+1}^i x, g_{a+1}^i x g_{a+1}^e)] \\
&= V_{a+1}^S(x,y) - V_{a+1}^K(x,y)
\end{aligned}$$

■

Table A-1: Distribution of the Patent Returns at Age One, Four, and Seven

Patent Returns	Age 1 (1)	Age 4 (2)	Age 7 (3)
50	0.6	9.7	22.1
250	5.2	15.7	25.6
500	11.1	22.6	33.2
1,000	20.6	32.9	44.0
5,000	54.7	64.4	72.3
10,000	69.8	77.1	82.8
25,000	85.2	89.3	92.5
100,000	96.8	97.8	98.6
500,000	99.7	99.9	99.9
Mean	17,999	12,716	9,019

Table A-2: Distribution of the Value of Patents at Age 1, 4, and 7

	Patent value at age 1 (1)		Patent value at age 4 (2)		Patent value at age 7 (3)	
Percentile	Value (US$ 2003)	Cum % of Total	Value (US$ 2003)	Cum % of Total	Value (US$ 2003)	Cum % of Total
5	787	0.0	0	0.0	0	0.0
10	1,472	0.1	134	0.0	0	0.0
25	4,584	0.7	1,426	0.2	431	0.0
50	16,184	3.7	8,664	2.3	4,180	1.2
75	54,950	13.9	35,616	11.2	22,231	8.8
80	74,198	18.1	48,881	15.2	31,379	12.5
90	160,829	32.3	110,991	29.1	75,121	26.1
98.5	763,483	67.5	550,287	65.4	390,948	63.3
99.8	2,432,580	88.0	1,777,078	86.9	1,303,508	86.2
Mean	76,598		52,669		35,835	

Table A-3: Distribution of the Value of Traded and Untraded Patents

	Untraded patents (1)		Traded patents (2)	
A. Distribution of the value of patents				
Percentile	Value (US$ 2003)	Cum % of Total	Value (US$ 2003)	Cum % of Total
5	655	0.0	4,881	0.0
10	1,184	0.1	9,574	0.3
25	3,354	0.8	23,499	1.8
50	10,605	4.0	58,321	7.7
65	20,546	8.45	99,373	14,7
75	33,958	13.7	149,628	22.1
80	44,983	17.6	188,998	27.2
90	98,253	30.8	355,916	42.7
95	189,792	44.2	609,837	56.7
98.5	503,271	64.9	1,293,629	74.7
99.8	1,711,007	87.1	3,758,570	90.8
Mean value	50,162		164,670	
B. Summary statistics				
Percentage (of all patents)	0.77		0.23	
Value of patents as a percentage of the total value	0.51		0.49	

Table A-4: Probability that a Patent is Traded over the Life Cycle

Percentile Initial return at age one	Probability of Trade $[\tau = 5,495]$ (1)	Probability of Trade $[\tau = 2,747]$ (2)
<5	0.020	0.021
5-10	0.034	0.035
10-25	0.048	0.064
25-50	0.122	0.215
50-65	0.252	0.36
65-75	0.343	0.44
75-80	0.397	0.477
80-90	0.455	0.514
90-95	0.511	0.555
95-97.5	0.544	0.565
97.5-99.8	0.551	0.564
All sample	0.231	0.296

Table A-5: Distribution of the Gains from Trade of Traded Patents

Percentile	Gains from trade Value	Cum % of Total
5	17	0.0
10	134	0.1
25	972	0.7
50	3,416	3.7
65	6,600	8.4
75	10,829	13.9
80	14,357	18.1
90	30,970	32.3
95	57,620	44.6
97.5	102,736	56.6
99.8	590,432	88.0
Mean value	15,008	

Table A-6: Value of Traded Patents and Gains from Trade in the Market for Patents

Value of Traded Patents With and Without Option to Sell the Patent and Gains from Trade [$\tau = \$5,495$]

Percentile Initial returns at age one of traded patents	Market for patents active Average Patent value (US$ 2003) (1)	Market for patents non-active Average Patent value (US$ 2003) (2)	Gains from Trade (%) (3)
<5	466	466	0.000
5-10	1,190	1,190	0.000
10-25	3,066	3,038	0.010
25-50	10,839	10,283	0.054
50-65	25,029	23,138	0.082
65-75	44,681	40,872	0.093
75-80	66,073	59,933	0.102
80-90	112,510	102,167	0.101
90-95	221,857	201,346	0.102
95-97.5	403,067	363,322	0.109
97.5-99.8	970,818	875,411	0.109
Mean (all)	164,670	149,662	0.100

Table A-7: Percentage Changed in the Gains from Trade of Traded Patents and Cummulative Gains from Trade After Decreasing the Costs of Technology Transfer by Fifty Percent

Percentile Initial return at age one	Change in the Gains from Trade (1)	Additional Average Gains from Trade $[\tau = 2,747]$ (Cum %) (2)	Average Gains from Trade $[\tau = 5,495]$ (cum %) (3)
<25	6.662	0.004	0.000
25-50	1.397	0.068	0.005
50-65	0.589	0.182	0.026
65-75	0.348	0.306	0.063
75-80	0.239	0.385	0.098
80-90	0.167	0.598	0.234
90-95	0.101	0.743	0.386
95-97.5	0.051	0.818	0.542
97.5-99.8	0.036	0.935	0.891
All patents	0.106		

Figure 2: Probability that an active patent is being allowed to expire for previously traded and untraded patents

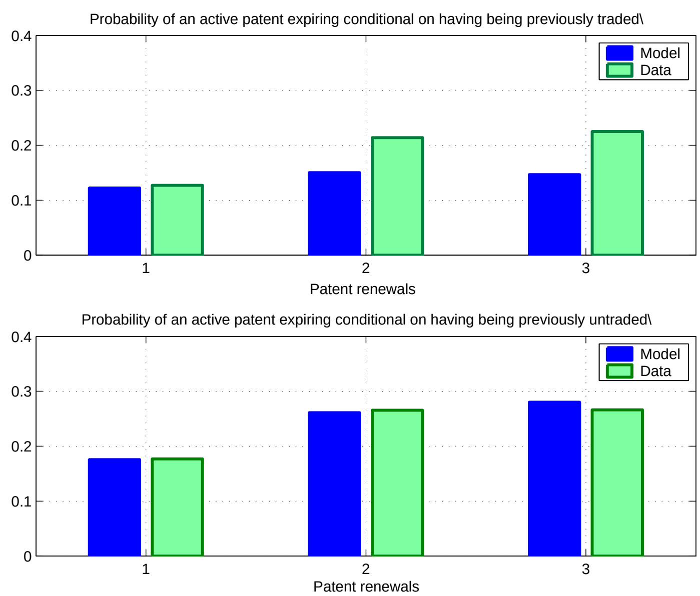

Figure 3: Probability that an active patent is traded for previously traded and untraded patents