KAUFFMAN
The Foundation of Entrepreneurship

KAUFFMAN DISSERTATION EXECUTIVE SUMMARY

Part of the Ewing Marion Kauffman Foundation's Emerging Scholars initiative, the Kauffman Dissertation Fellowship Program recognizes exceptional doctoral students and their universities. The annual program awards up to fifteen Dissertation Fellowship grants of $20,000 each to Ph.D., D.B.A., or other doctoral students at accredited U.S. universities to support dissertations in the area of entrepreneurship.

Since its establishment in 2002, this program has helped to launch world-class scholars into the exciting and emerging field of entrepreneurship research, thus laying a foundation for future scientific advancement. The findings generated by this effort will be translated into knowledge with immediate application for policymakers, educators, service providers, and entrepreneurs as well as high-quality academic research.

Kauffman Dissertation Fellowship Program KDFP Kauffman Dissertation Fellowship Program KDFP

www.kauffman.org/kdfp

The Economics of Patent Citations: Startup Commercialization Strategy, Value, and Success

Edward J. Egan

Abstract

This thesis proposes and validates an economic view of patent citations for start-up firms. I exploit variation in whether a cited patent is active or expired, whether citations are between firms in the same sector or not, and the number of citations accruing to each owner. I then use the effects of patent citations on three important economic outcomes for start-up firms to provide direct evidence that patent citations convey information about a patent-holder's right to exclude. I also develop new measures of the technological relationships between firms: I show that certain citations convey information about complements and others substitutes.

A Summary of *The Economics of Patent Citations: Startup Commercialization Strategy, Value, and Success*

Edward J. Egan

Overview

> "No matter the aspect of innovation under consideration, someone, somewhere, will claim that it can be measured with a patent citation."

When an inventor applies for a patent they have a duty of candor to disclose all prior art "material to patentability" to the patent office. Patent applicants disclose prior art by recording citations in their patent filing. As there is a paucity of other innovation measures, researchers use these 'patent citations' as measures of almost every aspect of innovation.

Patent citations have been taken to represent knowledge flows (e.g., Jaffe et al. 1993), technological positioning (e.g., Trajtenberg et al. 1997), technological importance (e.g., Trajtenberg et al. 1997), complementary relationships (e.g., Ziedonis 2004), substitution relationships (e.g., Lerner 1994), and much else besides. Each of these interpretations of the meaning of patent citations is reasonable. For example, when inventors in one geographical location record citations to inventions created elsewhere this likely signifies that they were aware of these other inventions and that knowledge has flowed. Or if a patent receives many citations from other patents then the original patent is probably an important antecedent to future invention. It although it is less clear why a patent that uses patented complementary inputs might cite these inputs, making citations to indicate substitution would help a patent application to demonstrate novelty. When a new invention replaces or improves upon some older invention, the old invention is very likely to be material to patentability.

But which is interpretation is correct? Do all citations convey information about knowledge flows as well as about complementary components? Or do some citations mostly convey substitution information and others mostly technological importance? In *The Economics of Patent Citations* I consider the effects of different types of citations – citations to active patents versus those to expired patents, citations within the same patent class or between firms in the same sector, and so forth – through three economic lenses to determine the information content of patent citations.

Each of the three economic lenses takes place in the context of start-up firms. Start-up firms can succeed or fail, undergo initial public offerings or acquisitions to raise capital, and have different values at each of these outcomes. I use the effect of patent citations on these three outcomes to reveal the semantics of citations made and received to patents held by start-up firms.

Patents expire at the end of their statutory term and if patent-holders elect not to pay their renewal fees. The invention detailed in an expired patent is in the public domain – the former patent-holder no longer has any exclusionary rights. If patent citations convey information about knowledge flows it does not matter whether a cited patent has expired

or not. Property rights do matter, however, if patents are complementary inputs into other patents or if patents substitute for other patents. I use a start-up firm's strategic choice to IPO or be acquired to consider whether or not property rights matter. Through this economic lens it is clear that citations to expired patents have different effects from citations to active patents. Moreover, the difference in these effects is not driven by the age of the cited patent. I therefore infer that patent citations contain information about property rights and not just knowledge flows.

Firm value provides another lens to differentiate between the meaning of patent citations. If patent citations largely convey substitution information, then the more citations a patent receives the less valuable it will become – each citation would represent a new, potentially superior, alternative invention that can compete to serve the same demand. On the other hand, if patent citations largely convey information about complementary inputs then each citation received would indicate a new potential usage for the patent, making it more valuable. In contrast to the literature studying patent citations in the context of established firms, I find evidence that when a start-up firm receives more patent citations its value decreases. This is consistent with citations to active patents held by start-up firms containing information about substitution effects between property rights.

Start-up firm success provides a third lens through which we can discern the meaning of patent citations. Success and value are, of course, related. One might expect that start-up firms who build up a suitable portfolio of highly-cited patents become valuable and succeed and that other start-up firms who do not will fail. But this is not the case. Having more patents is predictive of success but receiving more citations is predicative of failure, at least in important sectors like the semiconductor industry. Much like the results of the analysis of start-up value, this result is consistent with patent citations conveying information about substitution. The data suggests securing exclusionary rights that are not weakened or replaced by patents on later inventions is the recipe for success.

Using these three lenses, *The Economics of Patent Citations* focuses on answering a single question: What are the semantics of start-up firm patent citations? However, in the process of answering this question it provides two other types of contribution to the literature on innovation and entrepreneurship. First it establishes some empirical regularities which are relevant to a number of other open questions in the literatures. And second it provides a number of important insights into the economics of patents that can serve as the foundation for new understanding. In the remainder of this overview I provide some examples of these two other types of contribution.

Patent thickets have become an important topic to practitioners, policy makers and academics in the fields of innovation and entrepreneurship in recent years (see, for example Hall et al. 2012). A patent thicket is when patents form relationships that have the potential to cause market failure (see Egan and Teece 2014). These relationships have four main forms; three of these forms are based on substitution between patents and one relies on patents being complementary inputs into products. Naturally, statistics based on patent citations have been used extensively, and essentially exclusively, as measures of patent thickets. The results in this thesis do not answer the question of whether patent thickets exist. They do, however, suggest start-up firm patent citations do not have appropriate semantics to be used as measures of patent thickets.

Some empirical regularities uncovered in the analyses provide evidence to support or

reject certain technology commercialization strategy (TCS) theories, or at least put them into context. The result that firms that do more patenting are more likely to seek initial public offerings rather than acquisitions provides an example. Cockburn and MacGarvie (2011)'s patent thicket based TCS theory suggests that firms must build up defensive portfolios of patents in order to compete in the market. This is supported – though not identified – in the data. Gans and Stern (2003)'s 'market for ideas', however, argues that intellectual property rights like patents are key to cooperation, and that without patents firms may be forced to compete. This is not supported in the data, although the data is not ideal to test Gans and Stern (2003)'s TCS theory – the market for ideas matters when the appropriability regime varies across firms and the data contains only start-ups that have protected their intellectual property with patents.

Other results in *The Economics of Patent Citations*, particularly taken together, suggest a new TCS theory might be appropriate. High-technology products are often characterized as modular and complex (see, for example, Ethiraj and Levinthal 2004), in the sense that they are made up of systems of complementary components. As a consequence, start-up firms may face a choice between inventing a single high-quality component and cooperating with incumbents to commercialize it, or attempting to create a full system of components in order to compete as a standalone entity in the marketplace. Egan (2014) builds on this foundation, elaborates on the relationship between technology strategy and commercialization strategy, and provides some empirical identification by exploring the effects of a shock to start-up firms' commercialization strategy. Specifically, it suggests that patent citations to active patents held by firms in the same sector as the start-up firm convey information concerning component-wise substitution. It then uses these patent citations to provide measures of a start-up firm's choice to pursue a component or system based technology strategy.

Data

The primary data source is the NBER patent dataset, which contains data on every utility patent application at the United States Patent and Trademark Office (USPTO) from 1963-2006. The NBER patent data records all citations made subsequent to 1975, when citations were first stored in an electronically readable format, and has almost all assignment records for this period too. In total the data contains records on a little over 3.2m patents assigned to 4.86m entities.

Data on acquisitions are taken from the SDC Mergers and Acquisitions database from 1980 to 2010. The dataset includes only completed acquisitions for 100% of shares of the target firm, where the target was a U.S. private company that was acquired by either a public or private U.S. firm. For acquisitions by public firms for amounts above the mandatory disclosure limit, the data very closely approximate the population of acquisitions. The data on initial public offerings are taken from the Global New Issues (GNI) database from 1986 to 2010. As these data are extracted from offering prospectus and other mandatory security filings, these data represent the entire population of IPO events. Each firm must be privately-held and never have been publicly-listed or acquired before its initial (i.e., first) public offering to be included.

Only firms that have one or more patents prior to their liquidity event are considered. As

firm patenting behavior undergoes systematic and dramatic changes depending on whether the firm experienced an IPO or an acquisition, only patents, and citations to and from patents, that were filed prior to the startup's liquidity event were considered. When start-up firms are acquired, patenting drops precipitously because the start-up firm generally ceases to legally exist. The patents that belong to acquired firms are cited less after an acquisition, though this fact has no clear theoretical basis.

The main sample in *The Economics of Patent Citations* is restricted to firms that had either IPOs or acquisitions, which took place between 1986 and 2006. The main sample comprises 4,731 successful, patent-holding startup firms, made up of 1,230 that secured IPOs and 3,501 that were ultimately acquired. An additional sample of 1,311 'failed' venture-capital-backed startups is used for supplementary analysis. A startup is defined as failed if it did not secure either an acquisition, an IPO, or a subsequent round of financing for a period of four years after its last recorded round of financing. Data on venture-capital-backed firms are taken from Thomson VentureXpert from between 1980 and 2010.

The analysis also uses a subsample which requires that any successful startup firm was valued between U.S. $10m and U.S. $2.5b inclusive at its IPO or acquisition. This mitigates concerns that many acquisitions might be essentially 'fire-sales' or other non-successful financing events. The subsample comprises 801 IPOs and 1,070 acquisitions that held at least one patent prior to their IPO or acquisition. It is this sample that I will use in this summary to report some example results.

Example Results

In this summary I will review just three results. Each is important in their own right. *The Economics of Patent Citations* relies upon the synthesis of a number of results in order to tell its story. As such, I will only be able to foreshadow its overall understanding of patent citation semantics.

I discussed some of the main categories of potential patent citations semantics in the overview: Patent citations might convey information about knowledge flows, technological positioning and technological importance. They might also convey information about complements and substitutes. One important distinction between these two categories of potential semantics is that the latter is influenced by exclusionary rights and the former is not.

I define a patent citation as 'in-term' if is made to a patent that confers exclusionary rights to its patent-holder. Patents can lapse and lose their exclusionary rights at the end of their statutory term and if the patent holder does not pay renewal fees at certain intervals. 'out-of-term' citations are those made to lapsed patents. In Table 1 below, I report the results of logit regressions where the dependent variable takes the value one if a successful, patent-holding start-up firm secured an acquisition and zero if it secured an initial public offering (IPO). I use log counts of in-term and 'out-of-term' citation as dependent to explain whether successful start-ups firms were more likely to secure an acquisition as compared with an IPO.

In this analysis, I control for the start-up firm's portfolio time distance (the mean time between its patent applications and its liquidity event), as well which sector-year combination

its liquidity event took place in, its modal patent class (where patent classes are aggravated in categories using the categorization in Hall et al. 2001), and a polynomial function of its value at its liquidity event. I also use a non-linear function of citation age as an explanatory variable. The age of the citation is defined as the amount of time elapsed between the cited patent's application and the citing patent's application. Knowledge flows may differ in both quantity and quality over time; older citations – that is citations to older patents – might convey more basic knowledge or less relevant knowledge. The point of transition from in-term to out-of-term for exclusionary rights is also related to citation age. If a cited patent is renewed to its fully statutory term then citing patents that are filed within 20 years will make in-term citations to the cited patent.

Specification 1 of Table 1 shows that citations made to active patents have an opposite effect on a start-up firm's commercialization strategy (i.e., the choice to IPO or be acquired) than citations made to expired patents. Specification 2 shows that this result is almost entirely robust to the inclusion of measures of citation age. Specifications 3 and 4 show the same result in a different way. They use the fraction of citations made to active patents as an explanatory variable. There is a highly statistically significant correlation between an increase in this fraction and a decrease in the likelihood of an acquisition as compared with an IPO. And again, this effect is largely unaffected by the inclusion of measures of citation age.

Table 1: Term considerations and the right-to-exclude

The dependent variable is a binary variable taking the value one if the firm experienced an acquisition, and zero if it experienced an IPO. Coefficients are estimated using a logit regression. Standard errors, calculated using a Huber-White sandwich adjustment to correct for heteroskedasticity, are reported in parenthesis. All estimates include controls for the log of the value of the firm, the log of the value of the firm squared, the mean time elapsed between the portfolio's application date and the liquidity event date, as well as sector-cross-year and modal patent category fixed effects. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

	Specification 1	Specification 2	Specification 3	Specification 4
Log No Patents	-0.032 (0.086)	-0.053 (0.087)	-0.068 (0.084)	-0.088 (0.086)
Log Avg. Cites Made				
In Term	-0.358*** (0.113)	-0.319*** (0.114)		
Out of Term	0.403** (0.195)	0.369* (0.211)		
Fraction Cites In Term			-5.718*** (1.656)	-7.640*** (1.782)
Avg. Citation Age		0.034 (0.031)	0.040 (0.029)	-0.001 (0.031)
(Avg. Citation Age)^2		-0.000 (0.000)		-0.001*** (0.000)
Log Firm Value	0.055*** (0.021)	0.085*** (0.025)	0.096*** (0.022)	0.113*** (0.024)
Log Firm Value Squared	-1.819*** (0.465)	-1.757*** (0.471)	-1.746*** (0.473)	-1.712*** (0.476)
Portfolio Time Dist.	0.074 (0.049)	0.069 (0.050)	0.068 (0.050)	0.064 (0.051)
Sector x Year F.E.	yes	yes	yes	yes
Modal Pat. Cat. F.E.	yes	yes	yes	yes
Constant	3.970*** (1.214)	3.553*** (1.258)	8.632*** (2.077)	10.765*** (2.215)
R-Squared	0.3469	0.3489	0.34935922	0.3521
N	1509	1509	1509	1509

The results in Table 1 make no sense in a paradigm where citations only convey information about knowledge flows, technological position, or technological importance. In order for the status of cited property rights to have an effect on some economic outcome, like

the commercialization strategy of start-up firms, there must be some property rights based interaction between the cited patents and the citing patents.

Substitution and complementarities both depend on property rights, albeit in different ways. When one patent is covered by an active property right and another substitutes for it, the property right holder has no say and his rights have been diminished. And when a second patent needs an earlier patent as a complementary input, the original property right holder can exert their property rights by denying usage or setting licensing terms. In both cases, there is a property right interaction when the original patent is active but the original patent holder feels no effect once rights have expired.

A theory would be needed to distinguish between complements and substitutes based on the effects that patents have on choice between an acquisition and an IPO. Several such theories have been advanced in the literature, but discerning between these theories is outside the scope of this summary. Nevertheless, prior the analysis we could imagine that both in-term and out-of-term citations might convey information about knowledge flows, and that for either or both types of citation this information could be dominant. The results presented in Table 1 are not consistent with knowledge flow information being dominant in in-term citations. Instead they very strongly suggest that property-right status matters in patent citations. This has been surprisingly underexplored in the literature and has important ramifications.

Table 2 explores a different dimension of patent citation semantics. It considers whether the identity of the owner of cited patents matters. Specifically, it looks at the effects of citations made to established firms in the same sector as the start-up firm or to publicly-traded incumbents irrespective of their sector.

Table 2: Exploring the importance of competitive citations

The dependent variable is a binary variable taking the value one if the firm experienced an acquisition, and zero if it experienced an IPO. Coefficients are estimated using a logit regression. Standard errors, calculated using a Huber-White sandwich adjustment to correct for heteroskedasticity, are reported in parenthesis. The explanatory variables are citations-made in-sector, to publicly-held firms, and within patent-category. For in-sector and to-public citations, the measures are refined to consider only in-term citations, and then ownership of in-term citations and the residual impact of citation counts net of ownership for in-term citations. All estimates include controls for value of the log of the value of the firm, the log of the value of the firm squared, the mean time elapsed between the portfolio's application date and the liquidity event date, as well as sector-cross-year and modal patent category fixed effects. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

	Specification 1	Specification 2	Specification 3	Specification 4	Specification 5
Log No Patents	0.194* (0.106)	0.190* (0.106)	0.172 (0.108)	-0.031 (0.086)	0.143 (0.106)
Log Avg. C.-Made In Sector	-0.405* (0.216)				-0.336* (0.198)
In Term		-0.429* (0.221)			
Owners			-0.472* (0.263)		
Resid. Cites			-4.984*** (1.526)		
Log Avg. C.-Made To Public	-1.691*** (0.149)				-1.772*** (0.144)
In Term		-1.708*** (0.151)			
Owners			-2.057*** (0.184)		
Resid. Cites			-0.713 (0.845)		
Log Avg. C.-Made In Category				-0.17* (0.103)	0.575*** (0.180)
Log Firm Value	0.046** (0.018)	0.047** (0.018)	0.047** (0.019)	0.057*** (0.021)	0.068*** (0.022)
Log Firm Value Squared	-1.626*** (0.550)	-1.621*** (0.552)	-1.756*** (0.572)	-1.809*** (0.466)	-1.739*** (0.552)
Portfolio Time Dist.	0.068 (0.056)	0.067 (0.056)	0.079 (0.058)	0.075 (0.049)	0.082 (0.056)
Sector x Year F.E.	yes	yes	yes	yes	yes
Modal Pat. Cat. F.E.	yes	yes	yes	yes	yes
Constant	4.462*** (1.343)	4.477*** (1.349)	4.848*** (1.394)	3.569*** (1.207)	4.014*** (1.387)
R-Squared	0.602	0.604	0.611	0.343	0.608
N	1509	1509	1509	1509	1509

Specification 1 of Table 2 reports the effects of citations made 'in sector' and 'to public firms'. Both types of citations reduce the odds of an acquisition relative to an IPO. This finding is new and again could be interpreted through the lens of many different transaction commercialization strategy theories or beget new ones. However, arguably the most important aspect of this finding might not be immediately apparent: Citations in-sector, as well as citations to publicly-traded firms, have extraordinary explanatory power. Pseudo-R^2s from logit regressions do not measure goodness of fit in the same was as R^2s from OLS regressions. However, they can still be informative. Using a variety of Pseudo-R^2 measures, I find that using citations in-sector explain around 20% of the variation in commercialization strategies. Put another way, using Wald and Likelihood Ratio tests, citations in-sector add approximately the same amount of information to the regression analyses as sector $\times$ year fixed effects. As a general rule, patent citations are weak measures. The analysis in *The Economics of Patent Citations* suggests that this is simply because authors have failed to identify which patent citations are meaningful and which are noise.

Table 2 shows that the identity of ownership matters in patent citations. In Table 2 I explore whether there is information in the count of patent citations to certain owners beyond the count of these owners themselves. Specifications 2 and 3 add a measure of citation counts net of ownership to the count of owners cited in sector and the count of publicly-held firms cited, which were used before. The count of owners cited should matter that the count of patent cited when patent citations convey information about complements. This because complementary information is about needs – the citing patent-holder needs the cited patent-holder's patents – which can be fulfilled by licensing a portfolio of patents from each incumbent. Substitution information, on the other hand, is about replacement. When a citing patent replaces a cited patent it (generally) will not replace all of its patent-holder's other patents as well. The results in Table 2 are therefore consistent with in-term in-sector citations conveying substitution information but in-term out-of-sector (though still to a publicly-traded incumbent) citation conveying complementary information.

In a final example analysis, I consider the effects of patent citations on firm value. There is a large literature that has considered this question using total citation received (i.e., without decomposing citations into various types or considering citations made). Two papers, Hall et al. (2005) and Harhoff et al. (1999), have popularized the stylized fact that citations received are positively correlated with firm value. This is consistent with patent citations conveying information about technological importance or complementary relationships (as the more citations a firm receives the greater the number of products it is a complementary input it to, and so the more valuable it becomes). This stylized fact has is still repeated despite the large number of papers have found no relationship between citations received and firm value, and at least one paper (Shane and Stuart 2002) has found that citations received decrease firm value.

In Table 3, I use the log counts of patent citations received, and the log count of the owners of these citations, to explain log firm value in Ordinary Least Squares regressions. Each analysis sector $\times$ year $\times$ organizational form fixed effects, where the organizational form is an IPO or an acquisition. It also uses state fixed effects to limit the analysis to firms within each U.S. state.

Table 3 shows that patents are valuable but that, with one important exception, citations received (or made) are not. The result that provides the exception concerns citations received

to active patents from incumbents in the same sector as the start-up. *The Economics of Patent Citations* argues that citations to active patents in-sector largely convey information about substitution. This is consistent with the result in Specification 2 – the more citations that a start-up receives to its active patents from direct rivals, the greater the degree of replacement of its technology, and the less it is worth. The null results for total citations received (in Specification 1) can also understood as a consequence of opposing semantics. Some citations convey information about complementarities and technological importance, both of which can lead to increases in firm value, and this offsets the effects of citations that convey information about substitution.

Table 3: Derivative citation measures and value

The dependent variable is the log of firm value, measured at the liquidity event (i.e., at IPO or acquisition). Coefficients are estimated using Ordinary Least Squares. Standard errors, calculated using a Huber-White sandwich adjustment to correct for heteroskedasticity, are reported in parenthesis. All estimates include controls for the mean time elapsed between the portfolio's application date and the liquidity event date applied at the sector level where applicable, as well as sector-cross-year-cross-acquisition-indicator, state of incorporation, and modal patent category fixed effects. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

	Specification 1	Specification 2	Specification 3	Specification 4	Specification 5
Log No. Patents	0.235*** (0.040)	0.238*** (0.040)	0.237*** (0.041)	0.236*** (0.041)	0.232*** (0.040)
Log Avg. O. Rec'd	-0.063 (0.054)	-	-	-	-
In-term In-Sector	-	-0.172* (0.098)	-	-	-
Log Avg. C. Made	-	-	-0.027 (0.047)	-	-
Log Avg. O. Made	-	-	-	-0.022 (0.051)	-
In-term In-Sector	-	-	-	-	0.046 (0.046)
Portfolio Time Dist.	yes	yes	yes	yes	yes
Sector x Year x Acq F.E.	yes	yes	yes	yes	yes
State F.E.	yes	yes	yes	yes	yes
Modal Pat. Cat. F.E.	yes	yes	yes	yes	yes
Constant	3.366*** (0.315)	3.353*** (0.315)	3.392*** (0.325)	3.376*** (0.322)	3.319*** (0.314)
R-Squared	0.5062	0.5069	0.5057	0.5056	0.506
N	1677	1677	1677	1677	1677

The most important point made in Table 3 does not depend on any interpretation of patent citation semantics. Patents are wildly heterogeneous – some protect minor incremental inventions with little commercial value while others protect fundamental inventions that will change the world. As a consequence hundreds of papers in the fields of innovation and entrepreneurship have used patent citations in the hope that they might shed light on this heterogeneity. Using data on the near population of successful patent holding start-up firms, *The Economics of Patent Citations* provides evidence that only one type of patent citation sheds light on patent value, and shows that for this type of patent citation there is a negative relationship between citations received and firm value.

References

Cockburn, I.M. and MacGarvie, M.J. (2011). Entry and patenting in the software industry. Management science, vol. 57 (5), pp. 915–933.

Egan, Edward J (2014). How start-up firms innovate: Technology strategy, commercialization strategy, and their relationship. SSRN Working paper no. 2364096. Available at SSRN: http://ssrn.com/abstract=2364096 or http://dx.doi.org/10.2139/ssrn.2364096.

Egan, Edward J. and Teece, David J. (2014). A taxonomy of patent thickets. U.C. Berkeley Working Paper.

Ethiraj, Sendil K and Levinthal, Daniel (2004). Modularity and innovation in complex systems. Management Science, vol. 50 (2), pp. 159–173.

Gans, J.S. and Stern, S. (2003). The product market and the market for ideas: Commercialization strategies for technology entrepreneurs. Research Policy, vol. 32 (2), pp. 333–350.

Hall, B.H., Helmers, C., Von Graevenitz, G., and Rosazza-Bondibene, C. (2012). A study of patent thickets. Draft Report to the UK IPO, pp. 1–66.

Hall, B.H., Jaffe, A.B., and Trajtenberg, M. (2001). The nber patent citation data file: Lessons, insights and methodological tools.

Hall, Bronwyn H., Jaffe, Adam, and Trajtenberg, Manuel (2005). Market value and patent citations. The RAND Journal of Economics, vol. 36 (1), pp. pp. 16–38.

Harhoff, D., Narin, F., Scherer, F.M., and Vopel, K. (1999). Citation frequency and the value of patented inventions. Review of Economics and Statistics, vol. 81 (3), pp. 511–515.

Jaffe, A.B., Trajtenberg, M., and Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. Quarterly Journal of Economics, vol. 108 (3), p. 577.

Lerner, J. (1994). The importance of patent scope: an empirical analysis. The RAND Journal of Economics, pp. 319–333.

Shane, S. and Stuart, T. (2002). Organizational endowments and the performance of university start-ups. Management Science, pp. 154–170.

Trajtenberg, M., Henderson, R., and Jaffe, A. (1997). University versus corporate patents: A window on the basicness of invention. Economics of Innovation and New Technology, vol. 5 (1), pp. 19–50.

Ziedonis, R.H. (2004). Don't fence me in: Fragmented markets for technology and the patent acquisition strategies of firms. Management Science, pp. 804–820.