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

Permalink
https://escholarship.org/uc/item/7fz268rp

Author
Egan, Edward John

Publication Date
2012

Peer reviewed|Thesis/dissertation

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

by

Edward John Egan

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Michael L. Katz, Co-chair
Professor Toby E. Stuart, Co-chair
Professor David J. Teece
Professor Bronwyn H. Hall

Spring 2012

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

Abstract

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

by

Edward John Egan

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor Michael L. Katz, Co-chair

Professor Toby E. Stuart, Co-chair

This thesis examines the effect of the nature of a startup firm's underlying technological relationships, as measured by patent citations, on its commercialization strategy, likelihood of success, and value when successful. The innovation-lineage view of patent citations that has been articulated and explored in the literature to date has left several empirical puzzles and is insufficient for the task. I therefore propose and validate an economic view of patent citations. Using variation in whether a cited patent is in-term or expired at the time that a citing patent is applied for, I provide evidence that patent citations can and do convey economic information directly pertaining to a patent-holder's right to exclude. I then exploit whether citations are in-sector or not, and whether it is the number of cited owners or cited patents that explains meaningful variation in outcome measures, to uncover a fourth degree of heterogeneity in patent citations – whether citations convey information about complements and substitutes. Using samples that almost comprise the population of successful startups and failed venture-capital-backed firms that experienced their liquidity event between 1986 and 2006, and a reference dataset of all U.S. patents from 1963-2006, I find that the economic information in patent citations is of first-order importance and that startups file for patents on technologies that are at least partly substitutes. Furthermore, my results are consistent with start-ups that achieve initial public offerings patenting more radical inventions and acting as a force for creative destruction. The substitution effect of some patents explains both why receiving more citations does not necessarily make patents more valuable and why positive value effects are more likely to be found in the tail of citation count distributions. Overall, the economic view provides richer information than the innovation-lineage view, has strong construct validity, and provides a foundation for further research using patent-citation-based measures.

Dedication

Many people have been pivotal in changing the course of my life and so leading me to the completion of this dissertation and the conferment of a doctoral degree in business administration. James Baring, the 6$^{\text{th}}$ Lord of Revelstoke, passed away in February of this year. James was my first true mentor; we founded a software firm together, which laid the foundation for my interest in startups, venture capital and patents - the topic of this dissertation. Michael and Mary Kennedy were the first to suggest that I consider pursuing a doctorate degree in business. Michael, a barrister, and Mary, a financier, introduced me to the world of high-finance and guided me through the perils of the dot-com boom. And Livia Mahler, a partner at Greenstone Venture Capital, hired me as an associate, tutored me in the ways of VC, and supported the creation of my second firm, which provided data and analysis to the Canadian venture-capital industry.

However, I dedicate this work to Jim Brander, the Asia-Pacific Professor of International Trade at the Sauder School of Business, University of British Columbia. Jim, I will always be your student. Please accept this small tribute as a token of my appreciation for everything that you have taught me, all of the time that you have devoted to the work that we have done together, and the enormous positive impact that you have had on my life.

Contents

List of Figures

List of Tables

Acknowledgments

I would like to thank Michael Katz for his input on this research project as a whole; Steve Tadelis for his guidance through the earlier stages of this project's development; Jim Brander, Thomas Hellmann and Neil Thompson for their feedback on specific sections and approaches; and various participants at the Berkeley Innovation Seminar, the USPTO 'Patent Statistics for Decision Makers 2011' conference, and the Haas Business and Public Policy Seminar for useful comments. I would also like to thank Toby Stuart for his encouragement and friendship, without which I might not have reached this milestone.

During my time at Berkeley, numerous faculty members have assumed the role of my 'advisors', some formally as a member of either my oral examination committee or my dissertation committee, and others on a more ad-hoc basis. Michael Katz read every paper that I ever gave him and returned them promptly with detailed notes that provided everything from thoughtful comments on structure and approach to reflections on my use of the normative gerund. Toby Stuart stood by me when I needed an advisor, worked tirelessly with me in an attempt to get access to a novel dataset, wrote references for me at short notice, and provided excellent, and much needed, career advice and emotional support. Steve Tadelis took the time to meet with me weekly for a year and guided my exploration of cooperative game theory and models of the theory of the firm. Dave Mowery chaired my oral's committee, graded the paper component of my comprehensive examination, and provided me with a solid foundation in innovation research; I wish him a happy retirement. Christine Parlour acted as my finance advisor, and was always available for frank advice on both my research and academic prospects. Suzanne Scotchmer joined my oral examination committee as my outside member and brought her broad perspective to my work. Bronwyn Hall got me my job at the NBER and was instrumental in my choice of topic; this work stands on Bronwyn's shoulders more than any other giant. And finally, David Teece provided two fundamental contributions to my doctorate degree: His 1986 paper on the appropriability of technological innovation provides a theoretical underpinning for this research, and at a crucial moment in my studies he became a member of dissertation committee and showed me both far I had come and just how little I still had to go before the finish line.

Numerous other faculty members have generously given me their time and imparted some of their knowledge to me, either informally or in the classroom. I am particularly grateful to Bob Anderson, Ben Hermalin, John Morgan, Denis Nekipelov, Dmitry Livdan, Jim Powell, Matthew Rabin, Adam Szeidl, and Alexei Tchistyi. Likewise, no graduate student stands alone; I thank my peers for being there for me when I needed them. Brian Ayash, Harm Schutt, Ron Berman, Neil Thompson, Orie Shelef, Bryan Hong, and Justin Tumlinson, deserve particular mention. I hope I have repaid you all in kind. Finally, I wish to thank two non-academics for their support. Chris Kramer, you've always been a good friend to me; and Izzie Egan, you are all the family I need.

Financial support for this research was provided in part by the Social Science and Humanities Research Council of Canada award #752-2008-1770 and a Ryoichi Sasakawa Young Leaders Fellowship from the Institute for Business Innovation (IBI) at the Haas School of

Chapter 1

Introduction

This thesis asks the research question: What can patent citations tell us about the complementary and substitution-based economic relationships that innovative startups have with other firms?

It is important to understand both what patent citations measure and how they affect innovative firms, in particular successful startup firms. Patent citations map the observable relationships between aspects of the commercial innovation processes of firms. Successful startups are thought to play a vital role in introducing certain types of radical new innovations.

Startup firms are crucial drivers of the modern knowledge-based economy, particularly in the United States. Recent estimates suggest that venture-capital-backed startups alone accounted for 17.6% of gross domestic product, and 9.1% of private sector employment, in the United States in 2006[1]. Startup firms have created and commercialized much of the so-called disruptive innovation that has defined the modern era. Semiconductor microchips (Intel Corp.), computer operating systems (Microsoft Corp.), internet search, auctions, and shopping (Google Inc., eBay Inc., and Amazon.com Inc.), ipods®, ipads®, and iphones®, (Apple Corp.), genetically-engineered medicines (Genzyme Corp.), and green energy (Solyndra Corp.), provide everyday examples of technologies created and commercialized by what at the time were startup firms. Such firms have changed the way that we live and work in recent decades.

For certain types of inventions, patents provide their holders with some intellectual property-rights. In particular, a valid patent conveys the right to exclude usage of the invention by others for a pre-specified period of time. However, the invention must meet certain requirements to qualify for patent protection, and then must be disclosed in a patent filing. As a result, there are many inventions that are never patented. Depending on the nature of the knowledge that codifies an invention, and the competitive and technological environment, firms may opt to keep their inventions as trade-secrets, protect them through an

---

[1]Source: National Venture Capital Association (2007), report on 'Venture Impact' compiled by Global Insight.

assertion of copyright, or even release them as open-source or license-free solutions. However, patents are often given particular focus in the literature on innovation for several reasons.

First, patents must be publicly disclosed and filed with a central authority (the patent office), which makes them particularly tractable for empirical analysis. This may also make them better suited than other types of intellectual property protection for startups, which, by necessity, must often disclose their technologies to outside investors in order to secure commercialization financing. As I will show later, startups in the U.S. account for a disproportional amount of patenting activity relative to other privately-held firms.

Second, patents, through their collection of claims, represent discrete inventions. Although patents are hugely heterogeneous on almost every measureable dimension, this property allows researchers to treat a patent as a unit of *innovation*, which should have both private and social value.[2] As startups are privately-held prior to the liquidity event that allows them access to commercialization capital from public markets, there is often very little information available on their either private value or their contribution to social welfare. The academic literature has, as a result, made extensive use of the number of patents held by a startup firm to proxy for both of these value measures.

And third, a patent should cite the prior-art on which it is based, which includes relevant prior patents.[3] This provides a patent with an innovation-related lineage, which has been used as a broad, and generally poorly specified, constructional foundation for the citation-based measures in the literature. Specifically, citations have been taken to represent knowledge flows, technological positioning, and technological importance.[4] Given that patents are heterogeneous, if citations contain information about a patent's innovation-lineage, then it naturally follows that citation counts may be useful in shedding light on the underlying heterogeneity in the amount of innovation that a patent represents.

The typical argument in the literature relies on the number of citations-received representing a patent's quality, and that this in turn is correlated with private and social value: A patent with more citations-received has a better technological position, is of greater technological importance, or is the source of more knowledge, and so is of higher quality and more valuable than a patent with fewer citations-received. It is, therefore, not surprising that weighting patents by the count of citations-received has become a defacto standard practice when using patents as measures of value, and that this approach is particularly common in the analysis of startup firms, where other value measures are lacking. However, several studies have found that, in some samples of patents or patent-holders, receiving more citations

---

[2]Following Schumpeter (1934), innovation should be defined as the commercialization of invention. In the context of this thesis, I envisage startup firms as creating inventions, and then pursuing commercialization financing to turn them into innovations. However, most of the rest of patent literature treats patents as commercialized inventions.

[3]Patent citations are directional. I use the nomenclature *citations-made* and *citations-received* to indicate the (counts of) citations made from a patent to the patents it cites, and the citations received from the patents that cite it, respectively. Other authors use *backwards citations* and *forward citations*, or *out-degree* and *in-degree*, synonymously with these terms.

[4]E.g., Jaffe et al. (1993), Trajtenberg et al. (1997), and Albert et al. (1991).

is not correlated with greater value. These findings have generally been ignored, perhaps because they run contrary to mainstream intuition and no alternative value argument for a patent citation is easily conceived within the current paradigm.

The innovation-lineage view of citations has enjoyed considerable success by allowing researchers to address various policy questions. By using citations to reflect knowledge spillovers between entities, within industrial clusters, or even across sub-national and national jurisdictions, researchers have assessed the effects of changes to Intellectual Property (I.P.) policies, estimated the diffusion rates of innovations, and examined optimal cluster sizes and compositions.[5] However, this view of citations has too little formal structure to be used to address most strategy or firm-level economics questions. It is not just the assessment of the private and social value of patents falls into this category; a startup's technological dependencies are thought to influence both whether it seeks to finance the commercialization of its intellectual property through an initial public offering or an acquisition and the role that it plays in industry dynamics. To answer strategy and firm-level economics questions we need to understand not just whether a patent provides more knowledge or has better technological position than another patent, but also we need to understand the characteristics of the knowledge that are being transmitted through a citation or other discriminating details of the interaction mechanism between a patent and its peers.

To this end, in this thesis I consider the possibility that patents participate directly in economic processes, so that when a patent cites another patent this represents an economic relationship between two participants in an economic process. Specifically, I consider patents to represent inputs to firms' production functions or outputs that are traded in economic markets. Therefore, when a patent cites another patent, I consider it to represent that both patents are inputs or outputs, that they have a some kind of substitutability or complementarity with one another, and that each patent-holder will have specific rights to exclude the usage of its patented invention.

There are two potential approaches that one could take to establish the validity of this 'economic view' of patent citations. The first is to carefully examine a small sample of patents that cite each other and manually determine whether each and every citation conveys information about complementary or substitution-based relationships. However, there are several problems with such an approach. Most importantly, while I am theorizing that a firm's patent citations convey information about economic relationships, I do not anticipate that every citation will embody this information: many citations may convey information about knowledge flows.[6] And some citations may provide information about both comple-

---

[5]Henderson et al. (1998), Jaffe and Lerner (1999), Jaffe et al. (1993), Jaffe and Trajtenberg (1998), Audretsch and Feldman (2004) and many others. For a discussion of research addressing the knowledge spillovers between universities and firms see Agrawal (2001). Jaffe et al. (2000) provides evidence of the importance of knowledge flows between inventors employed by firms.

[6]I do not require knowledge flows and direct economic information to be mutually exclusive. That is, a citation may convey information that the cited patent is, say, being used as a complementary input along with the citing patent in the citing patent-holder's production function, and, at the same time, the citation may also convey information about, say, the nature of the technology that provides the basis for the citing

mentary and substitution effects: perhaps some claims in the cited patent are being replaced by the citing patent, while other claims detail components that will be needed. So the classification of complement and substitute information is likely highly problematic. Ideally, one would want a set of independent experts to perform the evaluations, and then inter-judge correlations could give assurances of construct reliability. The experts would need to be economists with a deep understanding of patents, the technology that they protect, and the market structure and supply-chains of the patent-holders.

From this basis, one could establish how the presence of complements and substitutes affects the value and commercialization strategy of start-up firms by empirically exploring the relationship between different classifications of patent citations and their patent-holder's performance. Likewise, one could attempt to determine the relative importance of patent citations that convey economic information against those that represent knowledge flows, and so forth. But there are serious limitations to this approach. Aside from concerns over the judgments used in the classification, there are immediate generalizability issues. On average in my sample of successful startups with patents, firms have about 10 patents each (for publicly-traded firms the average conditional on patenting is about 100 patents). Over the last 30 years, a typical patent has made about 10 citations. Even a small sample of firms with a high incidence of 'between firm' citations would need multiple experts to classify thousands of patent citations. And the results of such an analysis would still be left open to potential criticism regarding the representativeness of the classified patent citations and their firms.

I therefore adopt the second potential approach, which is to consider multiple tests that each provide different evidence of the presence and balance of complement and substitute effects and see whether they provide a coherent and consistent story. The advantages to this approach are that: I can conduct my analysis on a dataset that approaches the population of startup firms over a 30-year period; I can explore the implications of different types of patent citations at the same time as I shed light on the meaning and relative importance of these types; and I do not have to rely on subjective judgments. On the other hand, my approach carries some disadvantages. Specifically: my understanding of patent citations will necessarily be coarse, as I sacrifice the finely-detailed and highly-nuanced understanding of patent citations that would come from a manual classification for an aggregate understanding with well-defined generalizability; and I must be careful to use a set of tests that can cleanly identify the economic view. Identification requires that I can reject the economic view altogether if patent citations do not convey economic information, and if patents do convey economic information then I must be able to distinguish between portfolios that are predominantly comprised of complements, or of substitutes, or those that contain a mixture.

My approach uses a novel empirical design that leverages three important levels of heterogeneity within patent citations (i.e., different patent citations convey different meanings that become apparent with an appropriate decomposition), which are each interesting in their own right. The correlation between citations of certain 'types' and the observable outcomes

---

patent-holder's new input.

experienced by their patent-holders then allows me to identify the economic view and to document the existence of a fourth level of heterogeneity in patent citations - that of citations that convey complement information versus those that convey substitute information.

The first level of heterogeneity I consider is between citations to patents that are still valid, and so can exclude usage of their inventions, and those that have expired. I make this distinction at the time that the citing patent is filed, and refer to these two types of citations as in-term and out-of-term citations, respectively. Patents provide the right to exclude usage for a fixed duration, and once the patent has expired the patent-holder has no further rights.[7] If citations reflect how patents participate in economic processes, then whether cited patents are in-term or not should matter. I consider citations to out-of-term patents as representing knowledge flows of some kind, as in the innovation-lineage view, as they cannot represent direct economic claims. I therefore examine whether patent citations contain information about direct economic relationships by considering whether in-term citations have different effects from out-of-term citations, with the difference not explainable by the 'age' of the citation but instead the expiration cut-off.

Second, I explore the importance of heterogeneity in the degree of competition between citing and cited patent-holders. Some citations are made to firms in the same primary sector of operation as the patent-holder, while others are not. I expect this to be important because competitive effects should influence economic outcomes. Whether citations represent complements or substitutes, I expect that citations to firms in the same sector will be much more influential than citations to non-competitors. For complements, if the start-up firm faces a patent thicket (a situation where products are complex, requiring many patented complementary inputs, the requisite inputs are diversely-held, and at least some requisite input-holders have the ability to engage in hold-up), then in-sector citations are those directly into the thicket. Likewise, for substitutes, an in-sector citation represents the replacement of/by a competitor, which will have first-order consequences for rent appropriation. Because citations should cover *relevant prior-art*, I expect that citations out-of-sector are more likely to represent complements than substitutes. Substitution out-of-sector *can* be irrelevant, whereas requirements for underlying technology are relevant irrespective of their source.[8] That is not to say that a potential out-of-sector citation that embodied substitution information would always be ignored or overlooked, but instead that I believe it is far more likely that it could be.

Within in-term citations, for both citations to close competitors and other citations, I consider a third level of heterogeneity and distinguish between citations to different owners

---

[7]Patents filed after December 12th, 1980, are subject to maintenance fees (also called 'renewal' fees). Maintenance fees are due three times during the term of a patent. If a patent-holder does not pay a maintenance fee within a prescribed window, its patent automatically expires. This is discussed in detail later.

[8]A well-developed market for I.P. would undermine this argument. However, such a market does not currently exist, and while a firm could license its I.P. out-of-sector, it would then face information problems. This reduces the likelihood that out-of-sector I.P. is relevant. Anecdotally, many patents are thought to 'sit on the shelf' unused, when at least a portion of them could be repurposed outside of their firms.

and multiple citations to the same owner. This is done by counting the number of owners cited and creating a measure of the number of patents cited that does not contain this ownership information. This is important because it allows me to identify imperfect substitutes. When a patent-holder's right to exclude is still valid, the patent-holder will be able to force bargaining over provision for complements or place economic constraints on rent-sharing in output markets for substitutes. Bargaining theory in economics suggests that surplus should be divided according to the assignment of veto rights, and that having more than one veto right should confer no additional benefit. Ownership should therefore be paramount in citations that convey complement information. In the rare case of perfect substitutes, ownership will also be paramount, as having more than one perfect substitute conveys no further economic advantage to its owner and the rents are shared in the market according to the number of owners. However, for imperfect substitutes, which are much more common than perfect substitutes, ownership is not always important. What matters is the number and strength of other substitutes in the market and the market structure. As I will discuss later, in some cases involving imperfect substitutes, I expect that each additional citation, beyond the first to each separate owner, can convey information. Therefore, I can examine the relative importance of ownership to patent counts with respect to some economic outcome measure, like value or commercialization strategy, to establish the presence of imperfect substitutes.

I then turn to an examination of how different types of citations, particularly in-term citations and in-term, in-sector citations, influence two economic outcomes: the value of a startup firm at its liquidity event, and the type of event that it sought and experienced.[9] If patent citations represent complements then I expect that every additional citation-received will correlate with additional value, as each citation indicates a new need and new potential rents from licensing or refusing to license.[10] However, if patent citations represent substitutes then I expect that citations will decrease value, as each additional citation potentially indicates further replacement and loss of rents (though all rents may be lost with the first replacement).

Regarding the type of liquidity event that a startup achieves, I appeal to arguments of allocative efficiency from the literatures on the 'Theory of the Firm' and the resource-based view of the firm. I will refer to the effect of citations, which will measure the underlying technological relationships between the startup and other firms, on the liquidity event type as the 'strategy effect'; I argue that a commercialization strategy of securing an IPO will be optimal with in the context of some technological relationships, and the commercialization strategy of seeking an acquisition will be optimal in others. With citations representing complements, and if patent thickets pose material hurdles to new entrants, then startups making the most citations have the greatest exposure to the thicket and its associated hold-

---

[9]I also consider if citations influence whether a startup firm achieved success, but this is done as a robustness check to provide further evidence that the innovation-lineage view is inadequate.

[10]Refusing to license can be valuable. It can not only stop the erosion of rents from new entry, but can also alert the patent-holder to new commercial possibilities that directly involve its technology. A patent-holder that owns a cited complement may be well-placed to 'invent around' a new complement and capture the complement's associated rents itself.

up problem. An acquisition by an incumbent with a portfolio that allows the startup to navigate the thicket solves the startup's cospecialized asset problem and is the most efficient outcome. On the other hand, radical innovations are by definition completely unlike the previous generation of technology. With citations representing substitutes, and incumbents incurring *costs* (i.e., not just lost economic rents) in a potential acquisition related to their historic technologies that an entrant securing capital from public markets for a first wave of commercialization does not incur, an initial public offering is the efficient outcome for a firm that makes the most citations.[11]

I can therefore indirectly examine whether patent citations represent complementary or substitution effects by examining the nature of the correlation between citations and both a startup firm's value at its liquidity event and its optimum commercialization strategy, subject to several caveats. One potential problem inherent in the value effect arises with a finding of an 'indeterminate mixture' of complements and substitutes. In this case, the positive value effect of complement citations-received may be offset by the negative value effect of substitutes citations-received. Each citation-received likely contributes differing value: for complements, some licensing fees (or the rents associated with the strategic choice not to license) will be larger than others; and for substitutes, not only might the degree of substitution differ between citing patents but differences in market structures can also lead to wildly differing value consequences. As an example, the value consequences of the first perfect substitute patent arriving to change a monopoly to a Bertrand duopoly will be drastic, but a slightly superior substitute arriving to a incumbent in a broad oligopoly might be negligible. Combining the effects of complements and substitutes together can easily lead to an insignificant value result, as the variance of any line of best fit between two groups with opposing effects, one positive and the other negative, will necessarily be large. Thus, I expect a statistically significant positive effect when the citations are overwhelmingly dominated by complements and a statistically significant negative effect when they are overwhelmingly dominated by substitutes, but most likely insignificant effects for the mixtures in between.

An assessment of the effect of citations on the firm's commercialization strategy suffers from a different problem: a finding that citations-made are associated with an increased likelihood of an acquisition would be the result of a joint test of the hypotheses that citations contain complement information, that there is a patent thicket, and that an acquirer could successful mitigate the thicket. I do not have a hypothesis for the effect of complement citations when there is no patent thicket; there may be synergies between a startup and an acquirer if they share the same technology-base but these are likely of second-order importance to market-entry considerations. Nevertheless, a finding of a correlation between citations-made and an increased likelihood of an initial public offering is consistent with the presence of substitutes and so this test can add information to my overall story through consistency and coherency.

---

[11]Incumbents will lose at least a portion of the economic rents derived from their historic technologies when a radical new technology is brought to market. They incur these losses whether they acquire the startup or the startup achieves an initial public offering. However, I will later argue that they incur (additional)

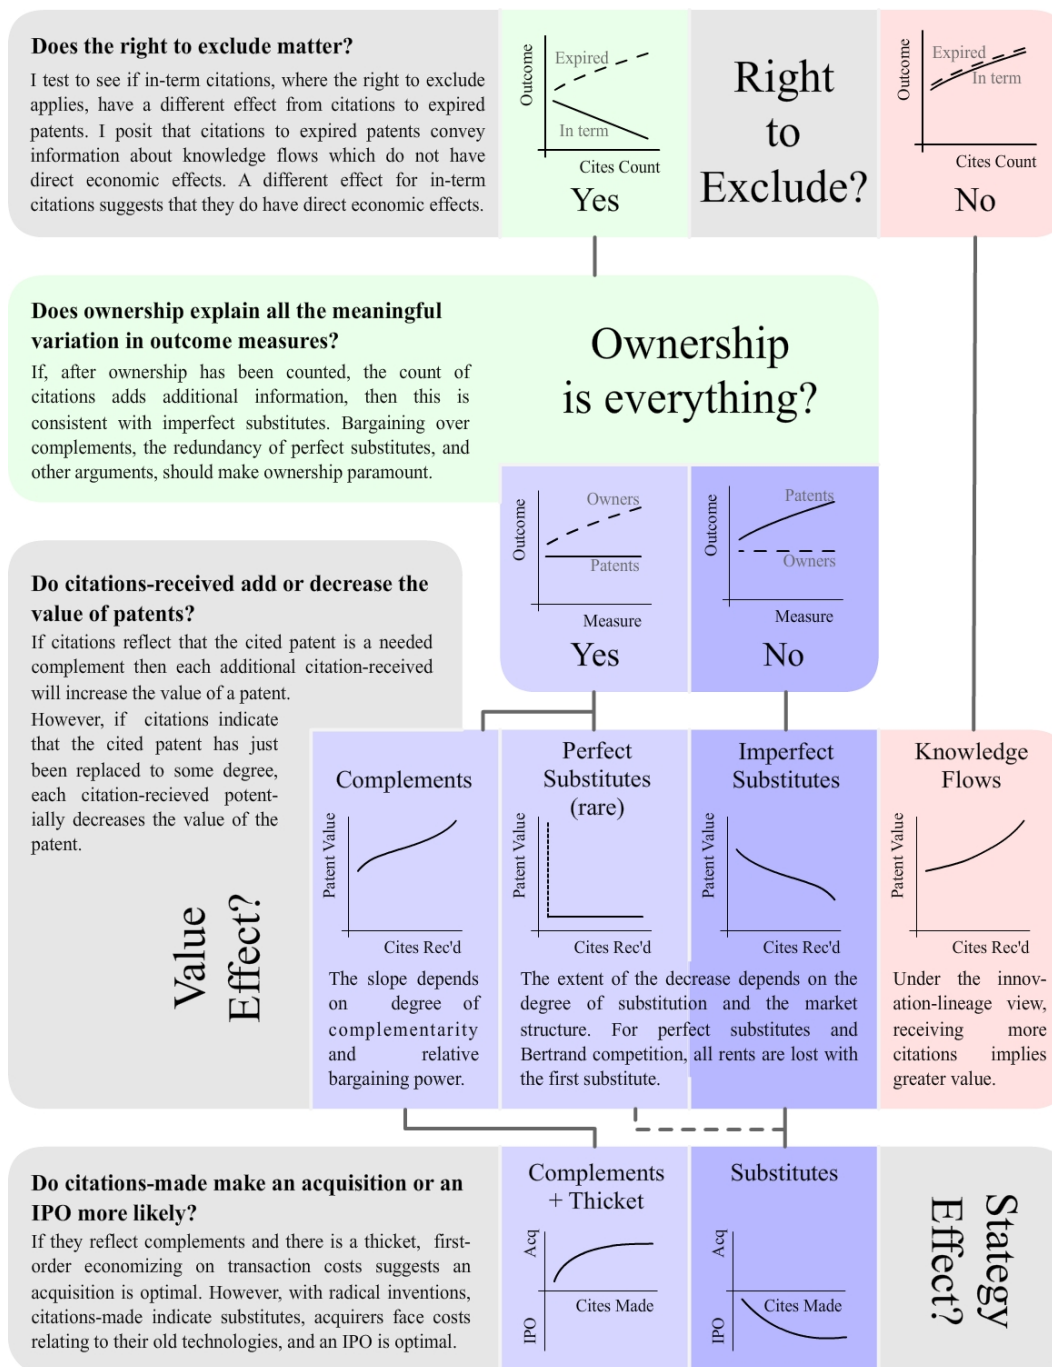Figure 1.1: Structure of the analysis

Figure 1.1 shows the structure of my analysis. I begin by considering whether in-term and

costs relating to their historic technology-base in an acquisition.

out-of-term citations have different effects in order to determine whether the right-to-exclude embodied in a citation has an effect.[12] I then refine my understanding of the mechanism behind the right-to-exclude effect by examining whether ownership is important, and by considering the sign of the value effect and the sign of the commercialization-strategy effect. Each of the four tests provides overlapping information, so that together they can provide a coherent and consistent story.

To assuage concerns regarding the identification of my approach, figure 1.2 shows five 'consistent' result sets, each with a unique interpretation, along with five examples of 'inconsistent' result sets. Each column provides four results together in a set: 1) whether in-term citations, where the right-to-exclude exists, have a different effect from out-of-term citations, with the difference not explained by the age of the citations; 2) whether ownership explains none, some, or all of the meaningful variation in citation counts; 3) whether the correlation between value and citation-received counts is positive, not statistically different from zero, or negative for in-term citations; and 4) whether in-term citations-made correlate with the startup achieving an IPO or an acquisition, or whether the expected effect is unspecified. Below each column is an interpretation.

The first five columns are consistent with some balance of complements and substitutes. For each of them, in-term citations have a different effect from out-of-term citations, and this difference is not explainable by the age of the citations, so that the right to exclude has an empirically demonstrable effect.

In the first two columns, ownership explains all of the meaningful variation in citation counts and the value effect is positive, consistent with complements. In the second column, the strategy effect indicates that citations-made are associated with an increased likelihood of an acquisition, consistent with complements in the presence of a patent thicket. The third column represents the reverse situation. Ownership explains none of the variation in citations, the value effect is negative, and the strategy effect indicates an association between more citations-made and an increased likelihood of an initial public offering, all consistent with (imperfect) substitutes.

The fourth and fifth columns describe the effects for mixtures of complements and substitutes. With mixtures, I expect that ownership will explain some but not all of the variation in citation counts, particularly for in-sector citations. Likewise, I expect a null value effect, or at least a muted positive or negative effect. The strategy effect is also slightly more nuanced for a mixture. Absent a patent-thicket I expect a positive correlation between citations and the likelihood of an IPO, arising from the presence of substitutes, and that this effect will be strongest in-sector where substitutes are more likely to be included as relevant prior-art. However, in the presence of a patent-thicket, the complement effect may overcome the substitute effect, leading to a positive correlation between citations and the likelihood of an acquisition.

In each of the first five columns, the results are consistent. Each result within a result set

---

[12]I also assess the relative importance of citations that convey economic information and those that represent knowledge flows empirically.

Figure 1.2: Identifying complements and/or substitutes

Test	Consistent Results					Inconsistent Results				
Right to Exclude	✓	✓	✓	✓	✓	✗	✓	✓	✓	✓
Ownership Explains	All	All	None	Some	Some	:	All	None	Some	None
Value Effect	+	+	−	O	O	:	−	+	++	−
Strategy Effect	?	Acq	IPO	Acq	IPO	:	:	:	:	Acq
Finding	Complements	Complements + Thicket	Substitutes	Mixture + Thicket	Mixture	Right to exclude meaningless	Ownership & Value Inconsistent	Ownership & Value Inconsistent	Ownership & Value Inconsistent	Strategy Inconsistent

tells the same story. My four tests each reinforce one another to provide a single coherent interpretation. In columns six through ten, the results are inconsistent. In other words, one test provides a finding consistent with, say, complements only, and another test provides a finding consistent with, say, substitutes only. It is crucial that I could generate inconsistent results. I would not want to, but without the possibility of inconsistent results my theory would not be falsifiable. I now discuss the major ways in which this could happen.

The sixth column shown is the first inconsistent column and describes the situation where the right-to-exclude does not convey information. This could happen either when in-term citations do not behave statistically differently from out-of-term citations, or where any difference is explainable by the age of the citation. In such a case, I would be unable to assert that in-term citations behave differently from knowledge flows, and so could not claim that they contain direct economic information pertaining to complementary and substitution effects.

The seventh and eighth columns show value effects that are at odds with the extent to which ownership explains variance. In the seventh specification, ownership considerations imply complements and the value effect implies substitutes, and vice versa for the eighth

specification. In the ninth column, ownership implies a mixture of complements and substitutes, but the value effect is extremely positive, consistent only with pure complements.[13] The tenth column has ownership and value effects that are consistent with pure substitutes, but a strategy effect that is consistent with complements in the presence of a patent thicket. There are other considerations aside from those presented in figure 1.2 that need to be consistent in order for all of my analysis to support a coherent story, and I will address them as I go. However, the figure should provide reassurance that my theory is easily rejectable.

I find results that follow the pattern in the fifth column, which is consistent with a mixture of complements and substitutes.

My findings support opposite meanings for in-term and out-of-term citations, and so are consistent with in-term citations providing information related directly to economic processes. This result is a necessary condition for my results to support the economic view; without this result I would reject the economic view altogether. Furthermore, I find that the aggregate effect of all citations is dominated by the effect of in-term citations, which suggests that the economic effects of complements and substitutes are of first-order importance and that the effects of knowledge flows are a second-order consideration.[14]

With respect to ownership considerations, my results are consistent with the presence of a mixture of complements and substitutes. For out-of-sector citations, ownership is indeed paramount. For in-sector citations, ownership explains some of the meaningful variation, but patent counts beyond ownerships still remain important. Citations to firms in the same sector of operation are comparatively rare, so, unsurprisingly, I find that overall ownership is paramount. Crucially, I also find that the meaningful variation in out-of-term citations is not explained by ownership. This provides an additional consistency check: Out-of-term citations represent knowledge flows and each patent from the same former owner should contribute additional knowledge, as these patents are now essentially 'ownerless'.

The value-effect estimation reinforces that startup patent citations describe a mixture of complement and substitute relationships. I find a null result for the correlation between citations-received and value, which is consistent with the positive effect of complements being countered by the negative effect of substitutes. Furthermore I also find a null result for the correlation between *citations-made* and value. My value theory is essentially symmetric - if citations represent that the cited patent is a complement to the citing patent, then citations-received increase value and citations-made decrease it, and vice versa for when citations represent that the citing patent is substitute.

There have been previous studies in the literature that have reported a null value effect for citations. This was a puzzle under the innovation-lineage view of patent citations, but is

---

[13]I could judge whether the value effect is extreme or muted relative to other value effects in the literature that have reported their balance of complements and substitutes, like Harhoff et al. (1999).

[14]Out-of-term citations are rarer than in-term citations, but this would not necessarily make knowledge flows a second order consideration. I check for the presence of knowledge flows in in-term citation by examining the correlation between in-term and out-of-term citations, and take this in account. And I will later show that in-sector citations, which are much rarer than out-of-term citations, are the most important of all citations.

reconciled by a supposition of a suitable balance of citations representing complements and substitutes under the economic view advanced in this thesis.[15] Moreover, some studies that report a positive value effect either directly report the overwhelming presence of complements or have samples that are strongly associated with complements, suggesting both that the economic view can be supported in samples other those comprised solely of startups, and that the startups in my sample are patenting relatively more substitutes that the firms in these studies. Particularly, this study presents evidence consistent with the population of startups patenting relatively more substitute technologies than the population of publicly-traded U.S. manufacturing firms over roughly the same period.

Finally, my commercialization-strategy results provide evidence of a positive correlation between in-term citations-made and IPOs (or, correspondingly, a negative correlation with acquisitions). This is supportive of startups in my sample having a mixture of patents that contain a material proportion of substitutes. Two additional strategy-effect findings provide other insights and reinforce the consistency of my overall results. First, the effect of substitutes should be symmetric. That is to say, if citations-made imply that patent is replacing previous technologies, then citations-received should imply that the patent is being replaced by later technologies. Thus, just as citations-made should make an initial public offering more likely, citations-received should make it less likely. I find that citations-received do act like the diametric of citations-made. Second, I argue that substitutes are more likely to be recorded in in-sector citations, and that in-sector citations should be the most important in determining the strategic outcome. I find evidence to support both of these arguments. In fact, my results suggest that in-sector technology relationships, as measured by in-term, in-sector, patent citations, are the single most important determinant of the liquidity-event strategy that a patent-holding startup pursues. The explanatory power of the measures of these relationships is roughly equivalent to the explanatory power of the interaction between the firm's industry and the timing of its liquidity event, which is typical the biggest factor considered in either the IPO literature or the mergers and acquisitions (M&A) literature.[16]

Overall, this thesis sheds light on how the technological relationships that startup firms have with other patent-holders influence their performance. That startup firms are associated with substitute technologies is an important empirical finding it is own right. The Schumpeterian view that entrepreneurs create value though the commercialization of inventions that achieve creative destruction is central to many policy arguments concerning the importance of entrepreneurs to innovation and economic growth. Previous studies have

---

[15]Other puzzles include why citation value effects are typically noisy, with large standard errors, and why value effects are often driven by the tail of citation distributions. These are addressed later.

[16]The literature on the factors that cause a nascent firm to choose between an acquisition (or coopertation) and an initial public offering (or competition) is essentially limited to Teece (1986) and Gans and Stern (2003). That the nature of technology endowments, rather than I.P. strength and contracting issues, should directly affect this choice is the single most important contributions of this thesis. However, there are literatures on the performance of firms that consider either IPOs or acquisitions, and the venture-capital literature generally treats both as successful outcomes. See Stuart and Sorenson (2003), Hellmann et al. (2005), Cochrane (2005), Cumming (2008), and others.

documented apparently causal relationships between entrepreneurial activities and innovation. For example, Kortum and Lerner (2000) find evidence consistent with venture-capital causing research and development efforts. This work takes the next step and advances our understanding of the nature of the innovative efforts undertaken by startups. I find results consistent with successful startups that replace the most incumbent technologies in their sectors of operation securing initial public offerings, and those that replace fewer technologies, particularly technologies outside of their sector, as well as those who have their technologies quickly replaced, being acquired. As such, my results are suggestive of a story where startups play an important role in industry dynamics, and where the radical innovation efforts of startups that achieve initial public offerings might act to move the technological foundations of their industries from one generation to another. These shifts in technology might provide new products and services, and new business models, as direct benefits. They could also offer indirect benefits, such as clearing away existing patent thickets.

This thesis makes an important contribution to the literature by supplanting the innovation-lineage view of patent citations with a new economic view. Although many of the high-level arguments used in this thesis, and at least one of the high-level findings, have appeared before in the literature, I believe this work is the first to synthesize them into a coherent framework and analysis.[17] For example, as I will discuss in later sections, previous authors have considered patent citations to represent complements or substitutes in certain contexts; this work simply considers that both possibilities might apply broadly. Likewise, the null value effect for citations-received for startups has been reported elsewhere in the literature; this work replicates that finding using a sample that approaches the population and explains how and why it should be true.[18]

Nascent literatures involving patent thickets, patent pools, patents in standards, and patent licensing and litigation, have all emerged in recent years. All of these literatures depend, at their very cores, on patents acting as complements or substitutes. My findings that patent citations convey information about complementary and substitution effects, and that this information is of first-order importance, dominating information about knowledge flows, can provide a fundamental building block for future research.

The remainder of this work is structured as follows: Chapter 2 provides the theoretical foundations for both the meaning of patent citations and for how they will influence a firm's value and a startup firm's commercialization strategy; chapter 3 reviews the literature on patent measures and their applications to startups; chapter 4 discusses my data and measures; chapter 5 details my empirical approach; chapter 6 provides results and analysis; and chapter 7 summarizes my findings, discusses them in the context of the literature, and concludes. An appendix provides details of how I classified firms into NAICS-based sectors.

---

[17]The high-level finding that citations-made are correlated with initial public offerings rather than acquisitions for a sample of startup firms (from certain sectors) was included in an old version of Cockburn and MacGarvie (2006). This finding was removed from later versions of the paper, in part because all discussion of the effects of patents on 'exit' (i.e., liquidity event type) was moved to Cockburn and MacGarvie (2009), though perhaps also because it was unexplainable in their framework.

[18]This finding was reported by Shane and Stuart (2002), discussed later.

Chapter 2

Theoretical Foundations

2.1 The Right to Exclude

A patent gives its assignee the legal right to exclude others from using an invention.[1] The invention is described in a number of claims, each of which should meet requirements for novelty, usefulness and non-obviousness, and must pertain to some patentable subject matter. A patent applicant owes a duty of candor to the patent office to disclose all relevant prior-art, which is generally done in the form of citations. These citations might be to other patents, academic works, or other material, and when a patent is granted there is a presumption of validity against this declared prior-art.[2] That is not to say that granted patents are automatically valid – they may still be challenged in court – but that in the granting of the patent, the patent office attests that a patent examiner has reviewed the submitted prior-art, added additional citations if needed, and believes that claims in the patent meet the necessary criteria for the assignee to be given exclusionary rights. My view of patent citations builds upon the foundation of patents as providing exclusionary rights. If a citation represents an economic relationship, such as the provision of an input or the trading of a competitive product, then the right to exclude will influence this relationship and the economic outcome that results.

Patents confer rights only for limited duration. Prior to June 8[th], 1995, patents had a statutory term of 17 years from the date of granting, and subsequent to this date the statutory term was 20 years from the date of application. After the patent expires, the assignee can no longer lawfully exclude usage of the invention. I refer to a citation as 'in-term' if the cited patent would still have conferred protective rights under the statutory term at the time of the application of the citing patent, and as 'out-of-term' if the cited patent's statutory term would have expired before the application of the citing patent was filed.

However, statutory terms can be prolonged if there were administrative delays or exten-

---

[1]Lemley and Shapiro (2005) point out that patents actually confer the "right to *try* to exclude" others from using an invention (original emphasis).

[2]This provides incentives for applicants to declare as much relevant prior-art as possible.

sions were granted, or shortened if the patent is not renewed, its rights are disclaimed, or it is later determined to be invalid. Some patent-holders opt to not renew their patents to the full term available, as renewal is costly. A revision to the patent act enacted on the 12$^{\text{th}}$ December, 1980, introduced 'maintenance fees' for U.S. patents. Following the act's enactment, patent-holders were required to pay to prevent their patents from expiring. Maintenance fees, which are more commonly referred to as renewal fees, are required $3\frac{1}{2}$, $7\frac{1}{2}$, and $11\frac{1}{2}$ years after a patent is granted, and cost \$900, \$2,300, and \$3,800, respectively. Although these fees are small, they have apparently had a material impact.[3] Scotchmer (1999) reports that, on average across all patent classes and types of patent-holder, "no more than 50% of patents are renewed past 10 years."[4] Although patent-holders in my sample of startups, perhaps particularly the venture-capital-backed startups that comprise 30% of my sample, are probably more likely than the average patent-holder to renew their patents, there is likely still substantial variation in renewals. Furthermore, for citations-made the cited patent-holder is generally not a startup in my sample, but rather a firm, individual, or institution in my reference dataset.

There are also other factors that might affect a patent's effective term length or otherwise affect whether a patent citation is in-term. A material portion of citations are added after the application, either by the patent examiner or the applicant (or the applicant's representatives).[5] The application process can take several years and citations could be added at essentially any point. Likewise, in some cases, patent-holders may manage to extend their terms through various administrative procedures or opt to disclaim their terms. In particular, patent applicants can apply for a patent term adjustment to extend the term of the patent for the time that it took for the patent office to conduct the prosecution of the patent application. On the other hand, a material proportion of patent applications are actually 'continuations', where an applicant adds new claims to a previously disclosed invention. For a continuation application, the start of the term is determined by the filing date of the original invention.[6]

Ideally, I would like an 'in-term' citation to represent that the cited patent's exclusory rights were valid at the time of the application of the citing patent, and an 'out-of-term' citation to represent that the cited patent had actually expired. The factors I have just described make my definitions of in-term and out-of-term imprecise, but they are still useable for my specific empirical analysis. Although my out-of-term citations may contain a small number of in-term citations, and my in-term citations will likely contain some out-of-term citations (as renewals are probably the single biggest factor in changing effective term lengths), I still have two groups that have very different term compositions; my in-term group will be sub-

---

[3]Lemley (2000) estimates a conservative average cost of a U.S. patent application at U.S.\$20,000. Renewal fees are therefore approximately 35% of this cost, before considering additional legal costs in processing each renewal.

[4]Likewise the USPTO claims that "less than forty percent of patentees pay all three maintenance fees". Source: USPTO.gov - Patent Term Adjustment (PTA) Questions and Answers.

[5]Alcácer and Gittelman (2006) and Cotropia et al. (2010).

[6]This is called the 'priority date' of the application.

stantially more in-term than my out-of-term group. The imprecision in my term definitions will reduce the statistical power of my analysis, but does not undermine its basic meaning.

I exploit the variation between in-term and out-of-term citations in my data to determine whether they provide different information with respect to some economic outcome. If citations do not reflect economic relationships directly, I would expect that in-term and out-of-term citations will have the same effects, or that any difference in effects will be attributable to the elapsed time between the patents (i.e., how old the knowledge in the cited patent is when the citing patent's application is filed) and not attributable to the expiration cut-off. In-term citations may embody the right to exclude and out-of-term citations generally do not, therefore I will test the hypothesis that the right to exclude, which is an economic right that affects both the cited and citing patent holders, is associated with a difference in the effect of citations on an outcome variable.

Put another way, I will use the effects of in-term and out-of-term citations to conduct an empirical contrapositive proof. Out-of-term citations are those to patents that are most likely expired. The invention underlying an expired patent has passed into the public domain. Although a citing patent could still use an expired patent as a complementary input or replace it as a substitute output, and so forth, this no longer has any economic bearing on the cited patent-holder. Therefore I treat out-of-term citations as conveying information about knowledge flows, rather than first-order economic consequences pertaining to the relationship between the cited and citing patent-holders. When I test to see whether in-term citations have different effects from out-of-term citations, I will be testing the proposition that in-term citations convey information that is *not* not-first-order-economic in its nature.

2.2   Complements and Substitutes

In economic theory both goods (i.e., outputs) and inputs are often described as complements or substitutes. A complement good is something that is used in conjunction with another good: Printers and ink cartridges for example. A (perfect) complement input is one that is required in the firm's production function in order to get a non-zero output: The Leontief production function, for example, produces the minimum of a function of one of two inputs. A substitute good, on the other hand, is one that can be used by consumers instead of another good to get the same utility. A Bic and a Pentel ballpoint pen may be functionally, and even atheistically, indistinguishable to a consumer and so are substitute goods. Likewise a substitute input is one which can be used instead of another input. With identical output elasticities, two substitute inputs can be used completely interchangeably, so that a firm can literally swap one unit of one input for one unit of another, and still produce the same output.

The patent office should not issue patents that embody the same 'inventive step', so they should not allow patents that cover the same technological foundation of an invention. But this does not mean that multiple patents cannot cover perfect economic substitutes. For example, the 'beverage sleeve', a cardboard insulator that protects the hand from direct

contact with a hot disposable coffee cup, was first patented in 1993.[7] This sleeve now has many subtlety different technological forms (i.e., their cardboard is corrugated in different ways, the two edges of material are joined together differently to form the sleeve, and so forth) and several subsequent patents have been issued to cover them. But they all provide essentially the same utility to consumers, are used in the same way, and to all reasonable extents and purposes they are perfect substitutes.[8]

Although I believe that multiple patents on perfect economic substitutes are possible, this is likely very rare in comparison to cases of imperfect substitutes. Imperfect substitutes might be inferior or superior – that is a substitute might replace an invention with something that is better or with something that is not as good. Inferior substitutes would still allow firms access to markets, albeit most likely at a reduced price. For example, some patents that contribute to standards may be substitutes. Patents that contribute to a standard and are filed before the standard is formalized likely cover complementary components. However, patents are also added to a standard after it has been announced, sometimes as a part of a formal revision to the standard. In this case, many of the new patents may be replacing the older technologies with improvements.[9]

Outside of simple mathematical models, defining something as a complement or substitute can be problematic. The degree of complementarity and substitution covers a continuum of possibilities ranging from inferior to perfect to superior. Even characterizing something as an input or an output might depend on how we characterize output markets. Consider, for example, Apple Corp's 'user interface gestures' patent that pertains to the distinctive interaction technology used in ipods®, ipads®, and iphones®.[10] Is this software a perfect complement in its own right (so that consumers purchase a bundled good from Apple – a device and some software), or is it a complementary input that allows Apple to vertically differentiate their good from other imperfect substitutes?

For my purposes, it suffices to call a patent a complement if needs or is needed by something to realize an economic function, and a substitute if it replaces or is replaced by another patent. I will then judge whether the collection of patents in a startup firm's portfolio exhibits characteristics and effects that are consistent with portfolios comprised predominantly of complements, predominantly of substitutes, or a mixture of the two.

A patent on an invention that is a superior substitute for a complement might make citations to another patent-holder to say, "I am replacing your patent with something better", and receive citations from other patent-holders to say, "I potentially need your patent as an input into mine". If all startups held just such patents, then I would use various different properties of complements and substitutes to empirically identify citations-made as substitutes and citations-received as substitutes. These properties include whether their effects

---

[7]U.S. Patent No. 5,425,297.

[8]To the beverage sleeve connoisseur, some are clearly better than others. However, I find mediocre beverage sleeves to be unrankable, and only describable as "another beverage sleeve".

[9]Patents filed after the standard could have been 'in the pipeline' before the standard was formalized, or could provide extensions to the standard as well.

[10]U.S. Patent No. 7,656,394.

are driven by ownership or not, their value-effect, and their effect on the commercialization-strategy of startup firms. Each of these properties in discussed in detail in the next sections.

First though, it is important to understand that the right to exclude manifests itself differently for relationships based on complements and those involving substitutes. And that for this reason, citations to complements are more likely between firms that do not have a competitive interaction. When something is needed (i.e., a complement) the right to exclude allows the owner of the needed item to decide whether or not to provide it, and on what terms. The competitive environment might determine whether or not the owner will provide it, but not whether or not it is needed. On the other hand, when something is being replaced, the owner of the replaced item has no say in its replacement. All interaction between the replaced owner and the new substitute's owner happen through the competitive environment. As such, with substitutes the right to exclude is better viewed as a right to participate. Each holder of a substitute patent can participate in the market, and non-patent-holders are excluded.

When a patent expires, the owner's rights vanish. A citation to an expired complement versus an in-term complement is different in that the former case the citing firm can acquire the complement as if it were a commodity, and in the latter case must enter into a (possibly fruitless) relationship with the cited patent-holder.[11] Likewise, a substitute to an expired patent is a substitute to a commodity, signaling potential entry into a market that can be served by anyone, whereas a substitute to an in-term patent indicates the citing firm has the right to compete in a market that is constrained by the number of other substitute patents. Therefore the contrast between in-term and out-of-term citations is similar for complements and substitutes, even though the mechanism for the interaction is different: In-term citations describe direct economic relationships involving the cited patent-holder; Out-of-term citations provide knowledge about public-domain technologies that are needed or being replaced.[12]

Because the economic interaction between patent-holders of substitute patents occurs through their competition, citations to substitute patents held by owners that have had and will have no competitive interaction are essentially meaningless. Patent law requires patents applicants to cite *relevant prior-art*. What is relevant to an applicant and to the patent examiner may be subtlety different, as the patent applicant has a greater incentive to use prior-art claims strategically, whereas the patent examiner may be more concerned with whether the technology is correctly defined. And each party may have different search costs; the applicant likely knows their competitive environment better, and patent examiners specialize in examining patents with the same technological characteristics.[13] However, I

---

[11]With imperfect complements, the citing firm doesn't strictly require the cited complement, but would benefit from it.

[12]It is worth noting that a citation to an expired substitute is only meaningful if the new invention is a superior replacement, in some way, for the previous technology. It is hard to envisage a benefit to patenting technologies that are inferior or perfect substitutes to technologies that are already in the public domain.

[13]Unfortunately, in my data I am unable to discern between citations added by the applicants and the examiners. The data have recently become available. Further research could use it to explore these issues,

believe that both parties are less likely to include citations to substitutes held by firms that operate in completely different sectors from the applicant, and the difference is only a question of degree.

This argument is tempered somewhat by the possibility of licensing-market competition, as opposed to product-market competition. A firm operating in a different sector that has does not compete in a startup's product-market, may still compete in the licensing market. However, firms may have limited knowledge of the production activities in other sectors and may be unaware that their technologies can be repurposed. Even a 'perfect' market for intellectual property, which suffered from no market failures, might still be unable to assist firms in providing substitutes technologies outside of their sectors. A patent is the codified knowledge of the inventive-step underlying an invention. Manufacture of the actual invention, as well as its successful commercialization, may require 'show-how' as well as 'know-how'. Thus, a market for intellectual property is not the same as a market for inventions or technologies.

The extent of licensing out-of-sector is an empirical question, and, to the best of my knowledge, there is no research that has considered the proportion of licensed technologies that are complements or substitutes. Nevertheless, it seems reasonable to assume that the competitive interactions of substitute patent-holders are less commonly mediated through licensing markets, and much more likely to occur in product markets.

2.3 The Importance of Ownership

The right to exclude embodied in an in-term patent citation implies 'bargaining over provision' with complements or 'constraints on rent-sharing' in output markets with substitutes. I now turn to a discussion of how ownership considerations enter into these two interaction mechanisms.

Supposing that the cited patent is something that the citing patent requires, or would benefit from including; then the cited patent-holder can use their right to exclude to either withhold the product or to force licensing (or other interactions that depend on bargaining) between the two parties. Bargaining theory advanced by Nash Jr (1950) and Shapley (1953) is based upon ownership and not the counts of inputs.[14] That is, these bargaining theories suggests that if two parties must come together to create some value, where one party holds 99 inputs and the other holds just 1 input, then both parties have effectively a single veto apiece (the party with 99 inputs can provide them all or not), and so, all else equal, the resulting proceeds from the combination should be split equally among both patent-holders. Put another way: ownership is paramount, and the count of citations beyond the first to each different owner is immaterial.

---

but this is beyond the scope of this thesis.

[14]In the infinite version of Rubinstein (1982)'s bargaining theory the outcome is based on the number of players, not their relative number of turns.

These bargaining theories are applicable provided that each citation reflects a need that adds positive joint-surplus if fulfilled. Without product-market competition, fulfilling each of the citing firm's needs will add joint-surplus, as the cited firm suffers no loss that offsets the citing firm's gain. With product-market competition, it is possible that fulfilling some needs provides positive joint-surplus but other needs incur zero or negative joint-surplus. In this case, and in the case where the cited patent would want to refuse all license requests, these bargaining theories are not strictly applicable.[15]

However, on a per patent basis, negotiating for complementary inputs can only have two effective outcomes with respect to an owner: either the owner licenses or refuses. An owner who offers to license some complementary inputs and not others either blocks the production of the technology because the refused complement is perfect, or doesn't because the refused complement is sufficiently imperfect. Therefore, all meaningful variation with respect to the patent's effects on economic outcomes takes place at the ownership level, and patent counts will not add additional information. My measures of ownership are aggregated up from a per patent basis to capture this. That is, I will count owners-cited for each patent, and average this over the startup's portfolio.[16]

There is an additional strategic complexity to negotiations involving complements. In the section detailing the theory on commercialization strategy for startup firms, I will discuss patent thickets. Patent thickets can lead firms to hold patents on complementary inputs into all of their competitors, creating a situation where if one firm were to refuse to provide a complementary patent it could experience a reciprocal refusal from the competing firm. Broad cross-licensing agreements may exist to prevent exactly this kind of 'patent war' between competitors, but regardless I anticipate that a repeated game could support equilibria where a given firm always licenses all of its patents or always refuses to license everything.[17] Either scenario would again make ownership, rather than patent-counts, of primary importance for complements.

With substitute patents, the right to exclude can place constraints on rent-sharing. Substitute patent-holders interact through the competitive effects of their product-markets or licensing-markets. The market structure and relative qualities of the substitutes will dictate the rents that patent-holders earn, and the potential change in rents from a new patent's filing. For product-market interaction, if two substitute patent-holders are sufficiently hori-

---

[15]Each of these bargaining theories explicitly requires super-additivity of components. Negative joint-surplus is possible if, for example, a firm would lose a monopoly to become part of an oligopoly if it licensed an input that facilitated the entry of a competitor. Even trades that would have positive surplus gross of transaction costs might be infeasible once legal fees and other expenses are taken into account. Furthermore, while a trade may have positive joint-surplus in a static welfare analysis, the cumulative nature of innovation might make the long-term joint-surplus negative.

[16]An alternative approach is to count all owners-cited across the startup's entire portfolio. This is more problematic, as it is possible that a cited-owner would be willing to license some patents to the startup to further the development of one technology but refuse to license for the development of another technology, and would potentially leave meaningful variation in the citation counts.

[17]It is also possible that firms could be involved in a 'Mexican standoff' where none of the firms takes a license and none of the firms sues for infringement.

zontally differentiated, so that there is effectively no competitive interaction between them, then a citation indicating substitution conveys no direct economic information and is essentially redundant. Substitute citations are therefore more likely to pertain to either mildly horizontally-differentiated product-markets, or some kind of vertical differentiation.

In horizontally-differentiated product markets, firms sell the same basic products to the same overall market but each firm provides products whose *unorderable characteristics* appeal to specific customer preferences. Quality-based, or vertical, differentiation is objectively orderable: some products are of higher quality than others, and producers of lower quality products must charge lower prices.[18] Products that have different unorderable characteristics are not 'better' than one another, just different. Location choice for bricks-and-mortar stores provides the prototypical example from the literature.[19] In a typical model, firms sell identical goods but choose to serve geographically distinct markets. Transport costs then make a customer in a given location prefer one firm over another. Although horizontal-differentiation can be a strategic choice, it can also be enabled by technologies that give rise to unorderable characteristics and can be patented.

Firms operating in mildly horizontally-differentiated product markets may have patents on their core technologies (i.e., the technologies that allow them to create or improve the production of their basic goods) or patents that protect their niche (i.e., patents on the technologies that allow them to differentiate themselves from their competitors). These patents might cover several complementary inputs or be substitutes for one another. If a competing firm patented a substitute technology for one or more of a firm's many substitute niche patents, ownership might be particularly important; The first time that a competitor cites one of several substitute niche patents, to convey that it has gained access to the underlying technology's economic function, is effectively a declaration that the competitor can now enter the firm's niche. Additional citations might still convey information about improvements in the competitor's capabilities within the niche, but this is almost surely a second-order consideration. Thus, in this case, the best measure of the competitive impact of patents conveying substitute citation information in mildly horizontally-differentiated product markets is the count of cited or citing owners.

However, if a competing firm patented a substitute technology for one or more complementary inputs, or a patent on core technologies, each citation should convey additional information. A substitute that replaces, say, three complements is not the same as a patent that replaces just one. And if some number of complements are all required to create a product in the niche, then the number complements that have been substituted is key. Likewise, if the substituted patents are in the firm's core, then, as the competitor already has access to its own basic core technology, each citation conveying substitution indicates that the competitor has now gained additional technologies that each facilitate incremental improvement. Thus, in these cases, the best measure of the competitive impact of a new substitute patent

---

[18]For seminal models of vertical differentiation see Dixit and Stiglitz (1977) and Shaked and Sutton (1982).

[19]Hotelling (1929) and Salop (1979).

is the count of the patents that it has replaced.

In vertically-differentiated product markets, consumers' willingness to trade-off quality for price, determines the number of products, taking into account their relative differentiation, that can earn positive rents. I refer to this as the 'thickness' of the market. When a patent cites other patents to indicate that it has replaced them to some degree, it is informative to count the number of patents that have been replaced, but uninformative to additionally count the number of owners that hold a patent that has been replaced. The information contained in the number of patents replaced is incomplete – some of these patents may have been partially occluded from using their right to participate and others might be pushed off the quality ladder altogether. But still, a vertically-differentiated product market of a given thickness supports a certain number of products, irrespective of who manufactures them. Thus, if patents cover products directly (i.e., the patents are on outputs), counting patents replaced, and not owners replaced, conveys the pertinent information. If patents cover inputs that enable the creation of outputs, the effect of replacing a patent is more difficult to anticipate. A new substitute patent-holder may be able to replace many products or none. However, again, is it patent replacement and not owner replacement that drives any economic effects. Therefore, each patent cited or each citation-received matters in vertically-differentiated product-markets.

Perfect economic substitutes are an exceptional case, but then a market for perfect substitutes cannot be properly characterized as vertically differentiated. For perfect substitutes, ownership is always paramount; holding more than one perfect substitute per owner conveys no further economic benefit, and if an outside party replaces one of a firm's perfect substitutes then it replaces them all. As such the only information in citation counts comes from the number of owners involved, and the number of times that an owner cites or is cited by a patent is irrelevant. However, as I have said before, although multiple patents covering perfect economic substitutes are possible, I expect them to be rare. Multiple patents on perfect *technological* substitutes are only possible to the extent that the patent office makes mistakes, and accomplishing exactly the same economic outcome (i.e., creating identical inputs or products) using different technology is highly problematic; some variation in the quality of the substitute (i.e., an imperfect substitute) is likely.

Finally, it is important to consider licensing-market competition, rather than product-market competition, as the medium for interaction between firms holding substitute patents. In licensing-market competition, the number of cited or citing owners provides the primary measure of competition in the market. The more owners of substitutes that a firm cites, the more potential licensing-market competition the firm faces. While the relative quality of substitute technologies may matter, this is firmly a second-order consideration.[20]

---

[20]The marginal cost of assigning rights to the production of another unit to the same licensee are zero. The marginal cost of providing the right to produce units to a new licensee are the contracting costs. Competition on price between a number of imperfect, but completely orderable, substitutes should lead to a price of the marginal cost of a new contract plus the difference in value to the user between the best substitute and the next best substitute for the best substitute patent, and the price of the contracting cost for all other substitute patents in the hierarchy. Furthermore, the user should be entitled to manufacture unlimited units.

Overall, therefore, for substitute citations, the proportion of meaningful variation in an economic outcome measure that is explained by ownership will depend on the nature of the substitute citations. For citations that convey information about perfect substitutes, or imperfect substitutes to various substitute technologies that can each support a competitor's niche in a horizontally-differentiated product market, or citations that convey information about substitutes that will compete in licensing markets, ownership should be paramount. However, for citations that convey information about imperfect substitutes in vertically-differentiated product markets, or those that convey information about the substitution of complements in horizontally-differentiated product markets, or those that replace other substitutes that support a competitor's core in horizontally-differentiated product markets, each individual cited or citing patent is important, and ownership conveys essentially no information.

To summarize, I associate cases where the meaningful variation in some economic outcome comes from ownership counts rather than citation counts with complements and with certain types of substitutes. However, only for certain types of imperfect substitutes should the count of patent citations beyond ownership be important. I expect that perfect economic substitutes will be rare; and I expect that relative to citations conveying complement information, citations conveying substitution information will also be uncommon.[21] But I have no priors about the relative incidence of the various types of imperfect substitutes in my sample. These factors will make it difficult for me to identify substitutes by considering whether the explanation of variation in an outcome measure is attributable to patent count or ownership counts, but my best opportunity to do so will be in the consideration of in-sector citations, where substitute citations are more likely, and where product-market interactions are more likely than licensing-market interactions.

Whether cited or citing patent counts, rather than ownership of citations, explains the meaningful variation in some outcome therefore provides my first differential test with regards to complements and substitutes. It builds upon the right to exclude – I only expect ownership to matter for in-term patent citations; evidence that ownership mattered in out-of-term citations would contradict my maintained assumption that expired citations convey 'ownerless' knowledge flows – and can be tested using any economic outcome where patents citations have a statistically significant effect. In the following sections I describe two further differential tests for complements and substitutes.

2.4   The Private Value of Citations

A firm's patent portfolio might comprise predominantly complements, substitutes, or an indeterminate mixture. I earlier defined a patent as a complement it needs or is needed by something, and as a substitute if it replaces something or is replaced by something. For

---

This is not what we observe in practice, most likely because of product-market competition.

[21]I will later describe a survey paper that reported a ratio of 12:1 for complements to substitutes in a sample of patents held by German inventors.

clarity I now introduce four terms: citing complement, cited complement, citing substitute and cited substitute.

From the perspective of citations-received, which have traditionally been used as a private value measure under the innovation-lineage view, a patent that is a complement is a required or useful input to, or a value-adding output with, the citing patent. That is, the citing patent either needs to license the *cited complement* in order to commercialize its patent, or by its very existence it demonstrates a new market need for the cited complement. Cited complements, therefore, become more privately valuable with every citation received.

A *cited substitute*, conversely, is something that has just been replaced to some degree. Patents are filed over time, and a citing patent is necessarily filed after the patents that it cites. As the cited patent came first, it may have previously been a complementary input or a substitute output, and so forth, but now the citing patent has arrived and at least partially replaced it, it has become a cited substitute. The extent to which the citing patent replaces the cited patent depends on the degree of substitution and the market structure. As examples, the citing patent might be an inferior substitute, and so reduce the rents to the cited substitute only slightly if the two patents cover products that are sold in a thick vertically-differentiated product-market; or if the citing patent is a perfect substitute and Bertrand competition applies, the rents to the cited substitute might drop to zero the moment the citing patent becomes commercially available. Citations-received weakly decrease value for cited substitutes. In a sample comprised solely of substitutes, one could attempt to uncover the average value function, which could be informative about both the average degree of substitution and market structure. However, I will simply be using the sign of the average incremental value of a citation-received as a differential test for the presence of complements and substitutes.

One problem with this approach is that the complement and substitute value functions for citations-received are unknown. For example, it is possible that a first substitute citation will reduce rents by more than a first complement citation will increase them. And likewise, given that sufficient substitution citations can drive the value of a patent to zero, and that every complement citation can still add further value, it is possible that a late (i.e., an incremental citation that occurs after many such citations) substitute citation will no-longer reduce rents but a late complement citation will still increase them. Furthermore, at some point a potential citation to a substitute becomes irrelevant. Once a technology has been fully substituted, and superior quality inventions dominate the market, its patent's right to participate has been made obsolete and their is no reason to cite it. Thus a patent with many citations might be more likely to be a complement, or at least be engaged in largely complementary relationships. This would lead to a finding that the positive value of citations is largely driven by the tail of citation distributions, which is consistent with some results discussed in my literature review.

Some mixtures of complements and substitutes may exhibit positive value effects or negative value effects. There is one obvious potential method for calibrating mixtures: I could use the value effects of samples dominated by complements, reported in the literature, that have determined their complement-substitute proportions. One study in particular,

Harhoff et al. (1999), reports a ratio of 12 complement patents to 1 substitute patent and an average, linear, implied positive value effect of about \$250,000 per citation.[22] However, this method relies on the value effect being generalizable across samples. Although I expect the presence of complement and substitute information in patent citations to be generalizable outside of my sample of startups, I do not necessarily expect different samples of complements and substitutes to embody the same underlying value function. For example, complementary inputs produced by universities might embody more basic science than complementary inputs produced by commercial firms, and so have different value implications with respect to citations that reflect their use.

As a result, the best I can say is that a sample that exhibits a positive value effect for citations-received is associated with a predominance of complements, and a sample that exhibits a negative value effect for citations-received is associated with a predominance of substitutes, where predominance is determined by the aggregation of effects rather than counts of complements or substitutes. In between these two extremes is the possibility of an insignificant value effect. The significance of a linear estimation using least-squares is based upon the contrast between the mean effect and its variance, which is calculated using square distances. The variance of an estimate that contains two heterogeneous groups, each with opposing effects, will necessarily be quite large. As such, for many mixtures it is likely that the value effect will be insignificantly different from zero, even though the coefficient might be materially positive or negative. For citations-received, a null value effect is therefore consistent with the presence of both complements and substitutes, even though the relative proportions can not be determined.

Citations-made have the opposite value effect to citations-received. For *citing complements*, every citation the patent must make to a pre-existing complement-holder potentially reduces the rents available to the citing patent-holder. In order to commercialize the patent, the patent-holder must secure the rights to use the complement patents. In a simple bargaining model, with no disproportionate bargaining power (which includes patience or outside options, depending on the model used), the rents will be split equally between the citing patent-holder and the n holders of all complement patents, so that the patent-holder will get $\frac{1}{1+n}$ of the rents.[23] For *citing substitutes* every citation-made increases the rents to the

---

[22]The data in Harhoff et al. (1999) is described in Scherer et al. (1997), who report that in their full sample patent values fit a log-normal distribution with an approximate mean of \$12.1m. Harhoff et al. (1999) report the results of a linear model of citation value, which has a coefficient on log value of approximately 0.1 for their full sample. Thus, $0.1 \times \ln(12.1m) \approx 0.25m$. However, a major point of Scherer et al. (1997) and Harhoff et al. (1999) is that patent values and citation premiums are highly non-linear.

[23]It should be noted that citations-made have different value functions from citations-received. In the case of citing complements, the fraction of the rents that accrue to the citing patent must be sufficient to cover the costs associated with the patent application. One might envisage a data generating process where patent applicants draw an intrinsic value for their invention as well as the patent's technological characteristics (i.e., whether the patent is complement or substitute and how many citations it must make). Only patents with very high intrinsic values could support giving away a large fraction of their rents to cited complements. But, in equilibrium, the value to a patent-holder of a patent that makes many complement citations will be low.

patent-holder. Every additional citation-made indicates another pre-existing invention which has just been replaced to some degree, with the rents now potentially captured by the citing substitutes patent-holder.[24]

A single patent might behave as substitute with respect to some of its citations and a complement with respect to others. Perhaps some claims in the patent replace old inventions and other claims outline the requirements for pre-existing technologies needed to build this new invention. However, in this analysis I aggregate up the notion of complements or substitutes to the patent portfolio, as my unit of analysis is the startup firm. Inherent in this aggregation is a loss of information. One patent could need three medium-cost complements and substitute for one high-value invention, and another patent could substitute for a low-value invention. In aggregate, a portfolio of these two patents might be, on average, an indeterminate mix of citing complements and substitutes.

The innovation-lineage view used only citations-received and not citations-made as a value measure. Under that view, a patent that receives a lot of citations is 'important', whether because of its technological position, the knowledge that flows out of it, or other such arguments. But a patent that makes a lot of citations is at best 'broad'.[25] It is unclear how this should map into value, but it does not seem natural to associate 'breadth' with a negative value premium. In my world of complements and substitutes, the effects of citations-made and received are essentially symmetric; both have clear and opposite private value implications. This is important for two reasons. First, I can use citations-made and citations-received to get a perspective on industry dynamics. This is discussed in the next section. And second, some potential findings are not consistent with the innovation-lineage view: citations-received should not have a null or negative value premium, and citations-made should not have a negative value premium.

2.5 Commercialization Strategy

Teece (1986) claims that there are three fundamental considerations that affect the commercialization strategy surrounding a new invention: the appropriability regime, the dominant design paradigm, and the need for cospecialized-assets. The appropriability regime consists of the nature of the technology, that is whether the knowledge embodied in the technology is tacit or codified, and the efficacy of its legal protection mechanisms. I consider only startup firms that have one or more patents, and as startups must disclose at least some information about their technologies to outside investors who will provide commercialization financing, it is likely that the firms in my sample use patents as their dominant intellectual property

---

[24]Note that each additional citation-made might also indicate that there is now greater competition in the market, so that the total rents are reduced. However, the new substitute will still gain new positive rents from each act of substitution.

[25]Making a lot of citations might also reflect the 'thoroughness' of the applicant. However, all prior-art should be cited, and each citation should be relevant. Furthermore, as previously noted, most citations are added by patent examiners.

protection mechanism.[26] However, I cannot see whether the firms in my sample are also using trade secrets or other protection mechanisms. Furthermore, I am also unable to infer whether protection for patents varies across industries or even across firm types (i.e., startups may enjoy less protection from patents, as the legal process for upholding a patent's validity can be extremely costly and these firms may need to commit all of their capital to pre-commercialization activities). In my analysis I will explore the mean value of a startup patent by industry, but these estimates are made without considering R&D expenditure and other variables measuring the technological opportunity that a firm faces, which might allow me to at least partially disentangle differences in I.P. regime strength from the variation in propensity to patent.[27] Nevertheless, patents represent (at least partly) codified rather than purely tacit knowledge, the legal protection that they offer is generally greater than that afforded other forms of intellectual property, and the patents in my sample appear to be highly valued by outside investors, so I will proceed to use Teece's commercialization-strategy framework under the assumption of a strong appropriability regime.

Teece's dominant design paradigm is drawn from Kuhn's stages of development of a branch of science, and closely resembles Abernathy and Utterback (1978)'s dominant design construct for industries.[28] There are two phases: the preparadigmatic phase, which occurs before a dominant design has emerged, where there is no single generally accepted approach to creating a product or service; and the paradigmatic phase where a standard basis for development has emerged and will continue until overturned by revolutionary science.[29] To the extent that startup firms are creating 'radical innovations', I expect that the filing of their patents will indicate that their industry has now entered the preparadigmatic phase, and conversely I associate 'incremental innovation' with the paradigmatic phase that occurs after a dominant design has stabilized.[30]

Cospecialized assets are those that need relation-specific investments from one or both sides. As Williamson (1971, 1979, 1991) argues, in the context of incomplete contracts, relation-specific investment leads to the possibility of hold-up, which, under the assumption of first-order economizing, should in turn lead to vertical integration as the optimal solution for the firm.[31] Put another way, when a firm considers investing in an asset that is dependent on a relationship with a third-party, it exposes itself to the risk of hold-up by that third-party – the third-party can give one set of terms before the investment and then demand better terms afterwards.[32] This can only happen when firms can't contract over all the states of

---

[26]The disclosure problem inherent in the financing or transfer of intellectual property is known as "Arrow's Information Paradox". One purpose of a patents is to, at least partially, mitigate this problem. See Arrow (1958). However, non-disclosure agreements may also mitigate this problem.

[27]I do not have representative, let alone comprehensive, R&D data for the firms in my sample.

[28]Kuhn (1970).

[29]This is elaborated further in Teece (2006).

[30]Henderson and Clark (1990).

[31]Patent-pools, cross-licensing agreements, and other mechanisms can also provide solutions to this problem, and are discussed shortly.

[32]Grossman and Hart (1986) and Hart and Moore (1990) provide formal models that capture one essential component of transaction-cost economics - that of underinvestment arising from cospecialization.

the world (perhaps because this is prohibitively expensive), as otherwise the contract can forbid the possibility of renegotiation. But with incomplete contracts, the renegotiating firm can claim that the environment has changed beyond that envisaged in the original contract. One solution to this problem advanced by Williamson is to merge the contracting firms into one entity, so that the incentive to renegotiate, except when the world really has changed, disappears.[33]

Complements and patent thickets

Patents that cite complementary input patents held by third-parties can act as cospecialized-assets. A firm that needs outside patents in order to commercialize its invention might want to merge with, or be acquired by, a holder of these outside patents.[34] This argument is not new in the patent literature. Hall and Ziedonis (2001) observed what they called the 'patent paradox' in publicly-traded semiconductor firms. These firms reported in surveys that they patented frequently but did not use patents in order to appropriate rents. Hall and Ziedonis (2001) argued that these firms faced a *patent thicket*. Products were complex, requiring many diversely-held inputs, and firms faced the possibility of hold-up by the input holders. To prevent hold-up, and allow the commercialization of new inventions, firms would patent a large number of inputs that were needed by their rivals.[35] With all producing firms holding patents on inputs into each other's products, they would face mutually assured destruction if one of them attempted hold-up.[36] Broad cross-licensing agreements may provide credible commitments not to do this, and guarantee that each firm can commercialize.[37] However, one side-effect of this behavior is that firms would have to hold patents for defensive purposes. These are patents that the firms weren't using directly, and so might report that they did not use to appropriate rents.

Ziedonis (2004) used a patent citation-based measure to further this argument. Ziedonis (2004) found that, within publicly-traded semiconductor firms, those with a greater fragmentation of ownership, that is those whose citations-made were spread more diversely across outside owners, patented more. Ziedonis interpreted this finding as consistent with the presence of a patent thicket in the semiconductor sector, and argued that those firms with the greatest exposure to the thicket were patenting complementary inputs in order to prevent hold-up. Underlying Ziedonis' argument is the notion that within the thicket, a firm would

---

[33]Williamson's hold-up mechanism can, in its simplest form, be reduced to the following 'equation': incomplete contracts + cospecialized-assets + opportunism = hold-up. The steps are as follows: i) firms A and B enter into an incomplete contract for firm B to provide goods or services to firm A; ii) firm A undertakes investment in an asset that depends on the contract, cospecializing it to work with firm B ; iii) firm B engages in opportunism by claiming that a state of the world not envisioned in the original contract has occurred; iv) firm A can not verify that the state of the world has not actually occurred; v) with the original contract void, firm B renegotiates better terms.

[34]Or an entity with either an implicit or explicit pre-existing arrangement to use these patents.

[35]These patents are often referred to as 'blocking patents'.

[36]Williamson (1983).

[37]Grindley and Teece (1997).

want to hold a patent or patents on inputs that are needed for each product that a rival produces, or at least a key product produced by each rival firm.[38] The number of a rival's products, and the extent to which they are key to a rival's business, dictates the leverage that a firm would have in negotiations.

Within the thicket, there is an assumption that for strategic reasons (i.e., firms are producing competing products and blocking a competitor from the market increases a firm's rents) only firms with sufficient leverage can secure the necessary rights. Although Ziedonis does not articulate it, she, like me, associates complements with bargaining and expects the negotiations to take place at the firm level, rather than the patent level. Whereas my reasoning led me to count the number of owners-cited, Ziedonis instead used the fragmentation of ownership as a measure of this transaction cost.[39]

Not all complements should be associated with patent thickets. Some, perhaps most, patents on complements appear to be licensed on a Fair/Reasonable And Non-Discriminatory (FRAND) basis.[40] For example, FRAND licensing is often a requirement for patent-pools, which should be composed entirely of complements.[41] Outside of a patent thicket, a firm with a citing complement would still incur transaction costs in assembling the required intellectual property to commercialize its patent. And the resulting negotiation should still take place at the owners-cited level, rather than the patents-cited level. But the firm does not face the prospect of hold-up if each patent needed can be licensed on FRAND terms.[42]

The possibility of patent thickets has caused concern among some academics, practitioners, and regulators alike. A patent thicket, based either on the hold-up mechanism I describe above, or arising from the patent office and the courts inadvertently allowing multiple patents covering the same underlying technologies, could have a very negative effect on innovation. As examples, a patent thicket might discourage new innovation, as a new firm might not be able to commercialize its products, cause wasteful duplication of effort, as existing firms would need to produce inventions within their competitor's research lines, and act a as friction on investment in (particularly cumulative) research and development efforts as firms could be wary of hold-up. Furthermore, policy responses that prevent cooperation between direct competitors could exaggerate these adverse effects.[43]

---

[38]Note that this could be achieved with a single patent that is used by all rivals, or many different patents, each of which is used by a separate rival in a separate product. A 'key' patent/product is one that is used to generate material rents and for no substitutes exist.

[39]Fragmentation measures the dispersion of ownership, independent of its count. That is, an unbiased fragmentation measure would record the same information if a patent cited two patents, each held by a different owner, or three patents, each held by a different owner. In both cases there would be maximal dispersion of ownership. It is immaterial to a correctly calculated fragmentation measure that in the former case there was two owners and in the latter three owners.

[40]The fraction of complement patents that are licensed on a FRAND basis is an outstanding empirical question, as is where, how, why, when and even if patent thickets emerge/exist.

[41]Including substitute patents in a patent pool would be an act of collusion, and potentially subject to penalties under anti-trust law.

[42]FRAND terms might still cause a royalty-stacking problem, where the sum of all royalties makes commercialization prohibitive.

[43]Shapiro (2001).

However, at least with respect to the barrier-to-entry concern, a finding that startup firms create substitute patents should provide grounds for some optimism. I have argued that complements, which could lead to an involvement in a thicket, are more likely in an incremental innovation regime that is typical in the paradigmatic phase (i.e., after a dominant design has emerged). And that conversely a preponderance of substitutes is more likely in the preparadigmatic phase of more radical innovation, when firms are competing to produce the dominant design. Substitute technologies can then act to reset the innovation stage, providing a new technological basis that is, at least temporarily, free of the problems of patent thickets.

If startup patents are mostly citing complements and the startup is operating in a thicket, in much the same way as the patents of established publicly-traded firms (particularly in the semiconductor sector) are supposed to be, then the greater the number of owners-cited the greater the startup's exposure to a hold-up problem. A startup could potentially solve this problem through an acquisition, if the acquirer held the necessary cross-licensing agreements and/or blocking patents. To offset the benefit from acquisition, I posit that acquisitions are costly to the entrepreneur, perhaps because imperfect alignment of the two firms leads to only a partial appropriation of value from the acquisition or because the entrepreneur derives some ego-rents from independent operation. Then I predict that the least-exposed firms will seek initial public offerings and the firms with more dramatic cospecialized-asset problems will seek acquisitions, leading to a positive correlation between the number of owners-cited and the likelihood of an acquisition.

However, I would stress that this prediction of the strategy effect for complements is really a test of the joint-hypotheses that the citations represent complements, that there is a thicket, and that an acquirer could mitigate the thicket. Absent a thicket, I have no clear prediction for the strategy effect of complement patents.[44] On the one hand, startups with complements will incur transaction costs related to negotiating licenses that might be reduced in an acquisition, particularly by an acquirer that had broad cross-licensing agreements already in place. And likewise, citing complements might be associated with more complex products whose commercialization might benefit from other cospecialized assets, such as specialized plant or an established brand, that an acquirer could provide. But on the other hand, especially in some sectors (such as Biotechnology), there has been a recent trend of vertical-suppliers seeking IPOs. With strong intellectual property rights, a market for technologies and technological cooperation is possible.[45] A firm with complementary technologies, that are not cospecialized (perhaps because rights to use the complements are available on a FRAND basis, or because the complements are effectively commodities), and a distinct niche, can license or buy the needed inputs, use its I.P. rights to protect its products from appropriation, and participate in a vertical as an independent entity.

---

[44]In the sub-sample of acquisitions, I examined the incidence of a target citing or being cited by its future acquirer. Both of these occurrences are extremely rare (i.e., occur in less than 1% of cases). This effectively eliminates the direct argument that a startup's complement patents will be cospecialized with the acquirer's patents.

[45]See Gans and Stern (2003).

Aside from the volume of citations, I also expect that to whom these complement citations are made will matter. I refer to citations made to or received from other firms in the same primary sector of operation of the startup as in-sector citations. Collectively, I refer to in-sector citations, along with citations to publicly-traded firms, and citations in-category as *competitive citations*. With citations representing complements, competitive citations are those to and from firms that are already operating in any thicket faced by the startup.

Substitutes and technology-specific assets

If startup patents are mostly citing substitutes, then there is an opposing 'allocative efficiency' argument to the first-order economizing of integration associated with cospecialized-assets. I will frame this argument using the notion of *technology-specific assets*, but it is applicable to other non-technology resources and draws upon Wernerfelt (1984)'s 'resource-based view of the firm'. The essential premise is that an incumbent can incur costs, above and beyond the erosion of its economic rents from the substitution, relating to its historic resources when it purchases a new resource, which an entrant, securing capital from public markets for a first wave of commercialization activity, does not incur.

The Cournot merger problem provides a simple example.[46] Holding industry profits constant, with each firm holding perfect substitutes, and $(N-1)$ firms in the market before the startup enters, the startup and a potential acquirer will each earn $\frac{1}{N}$ of the total profits from the market, and combined rents of $\frac{2}{N}$, if the startup enters the market. However if an incumbent acquires the startup, its post-acquisition rents will be $\frac{1}{N-1}$, which is the same as before the startup considered entry. This leads to a cost to an acquisition in the form of lost rents whenever $\frac{2}{N} > \frac{1}{N-1}$, or whenever $N > 2$. This argument is tempered with imperfect substitutes but the value of the 'new' technology to the acquirer can still remain below the value of the same technology to outside investors. In this example, this is because the new technology renders the acquirer's old technology obsolete.

Allocative efficiency underpins our understanding of the "Theory of the Firm".[47] There are three partially overlapping perspectives that provide the foundation our current understanding of the theory of the firm: Agency theory, Transaction-cost Economics (TCE), and the 'capabilities' perspective.[48] The agency perspective views the firm as a set of contractual relationships and focuses on the conflict of interest. The TCE perspective emphasizes asset-specificity - the extent to which an asset is specialized to a transaction, and such that some of its value would be lost if the asset were redeployed. And the capabilities perspective suggests that a firm derives its competitive advantage from the possession and coordination of resources that are difficult to trade, transfer or imitate, and so define it as an entity. Each

---

[46]The seminal reference is Salant et al. (1983). For a more recent summary see, for example, Head and Ries (1997).

[47]This branch of the economics literature has its roots in Coase (1937)'s seminal paper. A review of (formalizable) theories of the firm is provided by Gibbons (2005).

[48]These are summarized and contrasted in Teece (2012).

theory ultimately suggests that firm boundaries will be dictated by allocative efficiency, and I will draw arguments from each of them.

I begin by supposing that a potential entrant, a technology-creating startup, has discovered and patented a radical invention. The nature of this invention is such that its commercialization will permanently and irrevocably change the innovation landscape, and render much of the previous technologies held by an incumbent obsolete. The question at hand is whether, for reasons of allocative efficiency, this potential entrant should be acquired or undergo an initial public offering and remain a standalone entity.

A potential acquirer faces three choices: i) keep its old technology and either shelve or resell the new technology; ii) adopt the radical technology and sell its old technology; or iii) pursue both technologies. The first option is clearly inefficient. Buying the technology to resell it incurs transaction costs, and shelving the technology prevents progress and the gains to social welfare that would accrue as a result. Furthermore, shelving the technology exposes the firm to the risk that someone will invent-around the intellectual property to achieve the same radical change.

The second option exposes the potential acquirer to a technology-specific asset problem. Technology-specific assets might be the old technologies themselves, or assets or other resources whose value is tied to them. Human capital, tangible assets, and in-tangible assets might all be technology-specific.

Human capital can be intricately tied to the technology-base of the firm. Particularly R&D staff, whose knowledge and skills may be tied to the path-dependent evolution of the previous technology, but also production workers, sales staff, and other employees, may have knowledge and skills that are coupled to firm's the old technologies.[49] For both legal reasons and because workers may embody tacit knowledge, cultural norms, and other resources, a firm can not simply terminate the employment of its workers and replace them. Likewise, a firm will generally not be able to capture all of the benefits of retraining workers to work with a new technology.[50]

Innovation is rarely a one-shot game. Instead innovation is generally viewed as cumulative process.[51] A single invention may spawn several patents that are highly complementary, but with cumulative innovation, the likelihood of dependencies between different technologies that originate from the same stream increases dramatically. Physical plant related to the development and commercialization of a stream of innovations will tend to become more and more specialized over time. Even the business model of the firm, the firm's entire "organizational and financial architecture", may be a technology-specific asset – Teece (2010) argued that business models should be 'coupled' with the characteristics of a firm's innovation process and its outputs.

Transferring a specialized tangible asset to a third-party leads to a cospecialized-asset problem. That is, a firm can not sell a specialized asset to a third-party without exposing

---

[49]Cohen and Levinthal (1990).

[50]The seminal reference for a firm's inability to capture the returns to training is Becker (1964). Acemoglu (1997) provides a model where innovation exacerbates labor market inefficiencies.

[51]Levin et al. (1987).

the third-party to a hold-up problem, and as such would have to underprice the asset. With intangible assets, this specificity problem is amplified by information problems. Intangible assets often require tacit knowledge, which is difficult, if not impossible, to transfer. Patents may document the 'know-how' of an invention, but still require 'show-how' to implement their underlying technologies.[52] Likewise, trade secrets suffer from Arrow's information paradox, and although it is possible to license a firm's brand, it is extremely difficult to transfer it outright. Thus for information reasons ranging from adverse selection to appropriation-risk, one should expect even unspecialized intangible assets to be underpriced.

Aside from literal assets, firms have many resources, ranging from their routines, organizational structures and governance mechanisms, to their relationships with customers and suppliers, and their understanding of their product market and commercialization strategy. Any or all of these could be technology-specific, and at the extreme of a specialized business model it would seem almost impossible for a firm to buy and use a radical new innovation efficiently without fully reinventing itself.

Thus under the second option of adopting the radical technology and selling its old technology, we should not expect the potential acquirer to get the fair market value for its old technology; the potential acquirer faces a cost of disposing of its existing technology that the new entrant does not face. Capital markets can act to enforce allocative efficiency.[53] Rational investors would provide financing to the entrant to achieve an initial public offering, knowing that a potential acquirer will not provide the same return on investment. The future discounted cash-flows of the potential acquirer would reflect the value of the new technology less the losses incurred in disposing of the old technology.[54] However, it is important to note that capital markets consider only producer surplus and not consumer surplus. Thus, market-power considerations can dominate allocative-efficiency arguments. For example, in the Cournot merger problem at the start of this section, there was only a cost to an acquisition with $N > 2$. A monopolist earns greater rents than the sum of two duopolists. As such, a firm with extreme market power might be willing to incur the losses associated with an acquisition in order to maintain the rents that accrue from its dominant position.

The third option of pursuing both technologies is plagued by similar problems. Firms can and do pursue multiple independent lines of business. By effectively becoming a conglomerate a potential acquirer could maintain its existing technology alongside a disparate new technology. However, this is not without costs.

From an agency perspective, a conglomerate faces a multitasking problem, as described in Holmstrom and Milgrom (1991). Managers should allocate resources (including effort) to both lines of business in an optimal fashion, but may face different incentives to allocate resources to each line. In the case at hand, the new line of business, embodying the new technology, would be at the forefront of a radical shift in technologies, whereas the old line

---

[52]Hegde (2011).

[53]The primary economic rationale for having capital markets is to efficiently allocate capital.

[54]For publicly-traded acquirers, the firm's stock price, particularly its stock-price reaction to an acquisition announcement, provides the disciplining mechanism directly. For privately-held acquirers, the terms on which the firm can secure or renegotiate its debt offers some of this disciplining mechanism.

of business would be in decay. Ego-rents alone might provide sufficiently different incentives to cause sub-optimal allocations of resources.

Conflict between the two lines of business is also problematic. Scientists and research personnel with sufficient skills to move to the new technology may want to move, leading to training costs for their replacements. Sales staff on commission may want to sell the new product and not the old, especially if it offers new sales opportunities. Additional resources brought into the firm for the commercialization of the new technology may be coveted by the managers of the old technology, and time and effort might be wasted fighting over firm resources that are inherently constrained. New staff brought on for the new technology must be trained in the culture, routines and organizational structures of the firm, as may the existing staff if these need to undergo transformations. And the firm may end up trying to support two fundamentally incompatible business models.

All of these costs are born by a potential acquirer but not by an entrant that obtains commercialization financing from an IPO and remains a standalone entity. For these and other reasons, conglomerates face a diversification discount in financial markets.[55]

Now suppose that instead of a radical invention that permanently and irrevocably changes the landscape, the landscape itself is in flux. That is to say, suppose that a potential acquirer finds itself in the preparadigmatic phase of an industry, where firms are experimenting with new ways of achieving functionality and a dominant design is yet to emerge. This might be characterized by successive generations of substitute patents.

In this case the potential acquirer is less likely to have specialized its assets to a specific technology, as the optimal structure for an incumbent is one that supports learning and experimentation. Firms that face routine upheaval are more likely to have invested in dynamic capabilities.[56] As innovation is path-dependent and cumulative, the firm has a incentive to bring in new outside inventions; the firm needs to learn from them and develop its own platforms for the commercialization in the paradigmatic phase. Chesbrough (2003) argued that, particularly in the last 30 years, firms that pursue closed-innovation strategies, and attempt to conduct all research and development in-house, unduly limit their efforts and face being locked-out from breakthroughs originated outside of the firm. Consistent with Cohen and Levinthal (1990)'s notion that a firm should develop 'absorptive capacity', Chesbrough (2003) claims an open-innovation strategy should be optimal for a technology firm. Under an open-innovation strategy, a firm still conducts in-house R&D as this is crucial to understanding outside technologies, but it also sources a material proportion of its research and development from outside of the firm. In this fashion, a firm can leverage break-through research and development conducted world-wide, rather than just in its own lab.[57]

The combination of less specialization and incentives to bring in outside substitute

---

[55]See Lang and Stulz (1994) for early evidence. That the decision to diversify is endogenous, and how this affects the discount, is addressed in Campa and Kedia (2002).

[56]Teece and Pisano (1994), Teece and Shuen (1997), Teece et al. (2005), and others.

[57]Chesbrough (2003) provides a much richer strategic doctrine than is presented here. In particular, an open innovation strategy also requires firms to license, sell, or otherwise transfer away technologies that are invented in-house but are either underutilized or 'on the shelf'.

technologies leads to a different calculus for potential acquirers. Now the allocative efficiency costs to an acquisition are small and the benefits are considerable. Furthermore, the startup may be unsuitable for an initial public offering. With the dominant design still un-established, the final technology-base that the startup will use to generate cash-flows is still unknown. In addition to regular business risk (i.e., the startup may be more likely to fail altogether), investors who supported an IPO of a startup in this phase would face an allocative efficiency risk – the startup's final technology may involve cospecialized assets.

If citations represent substitutes, I can therefore make two predictions. The more citations that a startup's patents make, the more likely it will be to achieve an IPO. Each citation replaces a pre-existing technology, and makes it more likely that the startup holds a patent on a radical invention that will irrevocably change the landscape. But on the other hand, the more citations that a startup's patents receive, the more likely it is that these patents are covering technologies in the pre-paradigmatic phase, making the startup firm a suitable acquisition target. And with its technology already replaced to some degree before its liquidity event, a startup would not have a sound commercialization proposition if it sought an initial public offering and tried to remain an independent entity.

Again, competitive citations should be key. Citations to firms that are operating in the same primary sector as the startup, should represent that a technology will have a radical impact on the startup's competitors, or that the startup's technology is a suitable source invention for learning and experimentation. As competitive citations will have the greatest allocative efficiency impact both when citations represent complements (and there is a patent thicket) and substitutes, I therefore hypothesize that competitive citations will be particularly important in predicting the liquidity event type of a startup.

Allocative efficiency and value

I now have two competing sets of hypotheses regarding the effects of citations on the commercialization strategy of startups, one for citations-made and another for citations-received. Much like with the value-effect, the 'strategy effect' of citations can be used as a differential test to distinguish between complements (in a patent thicket) and substitutes. However, the allocative efficiency arguments that underpin my hypothesized strategy effect are related to a value argument; the most efficient strategic outcome is that one where the entrant realizes the most value. With citations representing substitutes, or a mixture of complements and substitutes without a patent thicket, I would expect startups to achieve higher valuations in initial public offerings than in acquisitions. And, conversely, if citations represent complements and there is a patent thicket, then I would expect an acquisition that solves a startup's cospecialized asset problem to yield a higher valuation than a hypothetical IPO. Therefore, I have a potential endogeneity issue: firm values and strategies are simultaneously determined, and both may reflect information contained in citations.

This problem is amplified if the most valuable firms, allocative efficiency aside, tend to seek either acquisitions or initial public offerings. One might suppose that initial public

offerings provide, on average, more funds for less equity than acquisitions.[58] Furthermore, it seems likely that the number of possible IPOs in a given year is limited by investor demand.[59] Much of the monies for IPOs comes from institutional investors (i.e., pension funds, insurance companies, hedge funds, and other large investors) who have limited capital, and further limit their allocations of capital to investment in new issues. A constraint on the number of IPOs, combined with a lower value for acquisitions, would make acquisitions a 2nd place prize for startup firms.[60]

Therefore citations could reflect value, which could predict the organizational form, and, concurrently, citations could reflect the firm's technology characteristics, which could predict its commercialization strategy. Fortunately, in my results I will find that startup citations are an indeterminate mixture of complements and substitutes, and so citation counts do not reflect value. However, firm value is still potentially an endogenous regressor in any estimation of the firm's liquidity event type. I will discuss this again later.

Like both the value effect and ownership analyses, an analysis of the strategy effect can be made using citations-made or citations-received. Again the predictions are symmetric. For cited complements, more citations-received provides a greater ability to over any thicket; and for cited substitutes, more citations-received indicates greater replacement by other firms, making the startup less suitable for an IPO but still able to provide a learning opportunity for an incumbent that faces a changing technology market. This symmetry of effects allows me to treat citations-made and citations-received separately, and provides an interesting window on to the industry dynamics associated with the commercialization of technologies by startup firms.

---

[58]Whether or not this is true is an open empirical question. Although the reverse scenario gives rise to the same issue, in my final sample, the mean firm valuations for IPOs and acquisitions are $246m and $108m, respectively, described in table 4.2.

[59]The number of IPOs per year in my sample ranges from 8 to 54 with a mean of 38, whereas the number of acquisitions ranges from 14 to 123 with a mean of 51. For successful venture-capital-backed startups, Hellmann et al. (2005) report that acquisitions are around three times more common than IPOs.

[60]There is also likely a limit for M&A demand. This is particularly relevant if M&As are actually more valuable than IPOs. The 3rd place prize for a startup might raise commercialization financing from another source or to fail.

Chapter 3

Literature Review

In his seminal paper, Griliches (1998) foresaw the problem that I am trying to address concerning the meaning of patent citations. He specifically asked, "What aspects of economic activity do patent statistics actually capture?"[1] Patents and their citations have variously been estimated as measures of social value, private value, knowledge flows, closeness of an invention to basic science, and the breadth of applicability of the invention and so technological importance, to name just the most common approaches.[2] The literature on patent and citation-based measures is very large.[3]

The vast majority of past 'citation-based measures' papers have considered only the heterogeneity between patents in terms of their citation counts, and not the heterogeneity in different kinds of patent citations.[4] This led their authors to make statements along the lines of 'the volume of citations-made or -received correlates with some outcome y' but left them unable to say why or how this should be so. A small minority of papers, notably Trajtenberg et al. (1997), Lerner (1994), Ziedonis (2004), and Hall et al. (2005), each considered a single dimension of the underlying heterogeneity within patent citations.[5] These papers made statements along the lines of 'based upon an assumption that variation in citations on some dimension (patent classes, ownership, etc.) means x, then x being correlated with some outcome y means such and such.' Sometimes these papers are able to add that 'x is correlated with y in a way that is consistent with the original assumption we made about the meaning of x.' For example, Trajtenberg et al. (1997) argues that when a patent makes citations that are more broadly spread across patent-classes, then the patent

---

[1]And points out that this is not the same as asking, "What would we like them to measure?"

[2]Seminal work includes Jaffe et al. (1993), Trajtenberg et al. (1997) and Albert et al. (1991).

[3]A special issue of 'Economics of Innovation and New Technology' was devoted to the variation in uses of patent citations. Gay and Le Bas (2005) provides an overview.

[4]An excellent discussion of patent citations, who adds them (65% are added by patent examiners), and why it matters, is provided by Alcácer and Gittelman (2006) and Cotropia et al. (2010).

[5]For Trajtenberg et al. (1997) this was variation in citations across USPTO patent classes; for Lerner (1994) it was between International Patent Classification patent classes; for Ziedonis (2004) it was in the dispersion of cited ownership, other than the patent-holder; and for Hall et al. (2005) it was variation in the degree of self-citation.

represents an invention that is in some sense closer to basic science than another patent that cites more narrowly. Trajtenberg et al. (1997) find a result consistent with this meaning for the spread of cited patent classes in a sample of patents from both universities and corporations.

I use three different dimensions of heterogeneity in patent citations, in conjunction with two outcome measures, to uncover evidence of a fourth dimension of heterogeneity – whether patent citations convey information about complements or substitutes. Specifically, I use in-term and out-of-term patent citations to test whether patent citations can convey economic information. Then I use variation in assignment, that is whether patent citations are assigned to the same owner or different owners, as well as the variation in the degree of competition between cited-owners and patent-holders, in conjunction with value-effect and strategy-effect tests, to uncover whether the economic information I find pertains to complements and/or substitutes. Ownership variation has been used before in the literature, but this work ties it to bargaining theory and complements, as well as certain types of substitute relationships, making its reasoning explicit for the first time. Term and sector-based variation have not been used before in the literature, and their theoretical impacts have not been articulated.

Furthermore, this thesis is fundamentally different from work in the literature in that it attempts to overturn the existing paradigm of meaning for patent citations, and provide a new constructional foundation. I will present evidence that the economic-information component of patents dominates the knowledge-flow component, and that the effects of the economic component are consistent with information about complements and substitutes. My analysis takes place in the context of startup firms, but I expect economic information to be ubiquitous in patent citations. As such I believe that the economic view put forward in this thesis informs many of the previous findings in the literature.

Under the economic view, patent citations reflect enhancements to social value, consistent with Trajtenberg (1990). Complements indicate new combinations, and substitutes indicate welfare enhancing progress.[6] The private value of patent citations, discussed above, is also now better understood: Null findings on value, such as those in Shane and Stuart (2002) and Sampat (2004) are meaningful; the wide variance associated with value estimations, reported in Hall et al. (2007) and others, is explained by the opposing effects of complements and substitutes; and the positive value effect from the tail of citation distributions, documented in Scherer et al. (1997), Hall et al. (2005), and elsewhere, is given a theoretical justification. Likewise, the notion of an invention's 'closeness to basic science', considered in Trajtenberg et al. (1997), can be revisited in terms of whether or not basic science inventions replace existing technologies and then provide building blocks for the next generation of inventions. And self-citations, explored in Hall et al. (2005) and elsewhere, could now provide insight into whether firms are creating successive generations of technology or assembling portfolios of complementary technologies, and this could be used to explore the performance consequences

---

[6]Schumpeter's notions that entrepreneurs generate value by the "carrying out of new combinations" (Schumpeter 1934, page 66) and through "Creative Destruction" (Schumpeter 1942, page 82) aptly reflect the two sides of how patent citations reflect enhancements to social value.

of these innovation strategies.[7]

Both Lerner (1994) and Ziedonis (2004) implicitly subscribed to my economic view of patent citations. In each case though, they neglected the opposing possibility of complements or substitutes, respectively. My contribution here is to provide a theoretical justification for the measures appropriate to their contexts, articulate it carefully, and empirically demonstrate that economic information is present in patent citations.

In addition, the innovation and entrepreneurship literatures that apply patent citation-based measures would benefit from an understanding of patent citations that is grounded in economic theory and that has well-articulated construct validity. This is particularly true in the context of the literature on the performance of startup firms, where patent and citation-based measures are frequently the only measure of firm value available. The innovation-lineage paradigm has become accepted by researchers, with the consequence that the construct validity of patent citations as (statistically significantly positively correlated) measures of value, in particular, is taken as given.[8]

The following quotations highlight the acceptance of patent citations as a defacto standard measure of patent quality and so value:

> *"The standard measure of patent quality … is the number of citations the patent receives"* – Gompers et al. (2003)

> *"…patent citations [are] a good indicator of the value of the invention"*
> – Dushnitsky and Lenox (2005)

As I will show, there have been findings in the measures literature that suggest that the correlation between citations-received and value may be spurious for certain types of inventions or firms, but within the innovation-lineage view there is simply no way to reconcile these findings with theory. As a result, the findings that do support a correlation are the ones that are touted.

The history of the exploration of the relationship between patent citations and value begins with Trajtenberg (1990). Trajtenberg (1990) proposed citation-weighed patent counts as measures of patent value after finding a strong relationship between citations and social value in Computed Tomography scanner inventions. This was followed by results that support the notion that citations-received are positively correlated with private value by Harhoff et al. (1999) and Hall et al. (2005).[9]

Harhoff et al. (1999) surveyed inventors on the value of their 964 inventions in the U.S. and Germany and found that inventions reported as more valuable are associated with higher

---

[7]A self-citation is a citation made from and received to patents held by the same patent-holder.

[8]The literature that takes this construct validity as given is probably several orders of magnitude larger than the literature on the measures themselves.

[9] Hall et al. (2005) is a reduced but published version of Hall et al. (2000), which was released as an NBER working paper. I will cite Hall et al. (2005) when discussing this work, but readers should understand that this paper was released five years before publication, and that the working paper contains considerably richer institutional detail.

counts of citations-received. Harhoff et al. (2003a) re-confirmed the basic results of this work, as well as adding the effects of other patent measures, such as a measure of patent scope, and a measure of the importance of non-patent citations. Hall et al. (2005) examined the Tobin's Q of 4,800 U.S. manufacturing firms over 30 years and found that citation-weighted patent measures are more highly correlated with the market value of the firm than simple patent stocks. However, a careful reading of this paper reveals some intricacies to the main result. It is the long-run citations, received after the value measure is taken, that are most informative; and the citations-received before the valuation have only a comparatively small effect on value.[10] Likewise, self-citations are more important than outside citations, particularly for firms with small patent portfolios. Furthermore, citation effects are driven by the tail of the distribution; firms with below the median number of citations-received experience no value effect from citations, whereas firms with more than three times the median number of citations-received experienced a large value effect from a citation.

Hall et al. (2005), in its various forms, is the most cited reference to justify the use of citation counts as proxies for invention or firm value.[11] This is likely due to a combination of factors: The paper has very large sample, effectively covering the population of publicly-traded manufacturing firms over 30 years; the empirical approach is very sophisticated – essentially reflecting the state-of-the-art for econometric estimation using patent data; and the authors are highly respected econometricians and economic theorists, who were jointly responsible for the various efforts at early computerization of patent data.[12]

However, other (mostly later) work was much more cautious regarding a relationship between citations-received and private value. Hall et al. (2007) considered patents held by 7,168 publicly-traded European firms and found that citations-received provide just a small explanation of value and then only in certain fields, and that for business method and software-patents citations are essentially meaningless. Overall, they report that "although quality adjusting these patents is significant [using forward citations] ... it adds only about 0.1 per cent to the explanatory power of the regression." Hall et al. (2007) report that other studies of European firms, including Blundell et al. (1999), Bloom and Van Reenen (2002), Toivanen et al. (2002), and Greenhalgh and Rogers (2006), had failed to find any relationship between patent citations and the market value of their inventions.[13]

Sampat (2004) examined the patent characteristics and value of patents produced by university laboratories. They found that citations are significantly related to the likelihood of licensing but not to licensing revenues. This work implicitly draws attention to the possibility of a selection effect. Griliches (1998) repeated and popularized the stylized fact that only approximately 50% of patents are used. It is therefore possible that citations might select between useful and useless patents, but have little to say about the value of patents conditional on use. Particularly for certain types of non-practicing entities (NPEs), like

---

[10]These citations are observable to the econometrician but not to those determining the value.

[11]Google Scholar estimated citation counts in November 2011 as 1,290 for Hall et al. (2000) and 582 for Harhoff et al. (1999).

[12]Hall et al. (2001).

[13]See Hall et al. (2007), table 1, page 12.

universities, there is a natural explanation for Sampat (2004)'s findings under the economic view of patent citations. As an NPE does not participate in competitive product-markets, an invention may be less likely to receive citations representing substitution until it is licensed.[14] An invention with commercial applications might accrue more citations indicating a need prior to licensing, and then accrue a mixture of further complement and new substitute citations leading to a null value effect.

Aside from citations, the two most common measures of patent use are *renewals* and *litigation*.[15] Lanjouw et al. (1998) provided some evidence that patent counts are very imperfect measures of innovation output, and hence both private and social value, but that they could be improved substantially by considering renewals. Lanjouw and Schankerman (2004) considered both renewal and litigation, and found a positive statistical relationship between citations-received, as well as other measures such as the number of claims-made and the number of jurisdictions in which the patent is filed, on both renewal and litigation. And Bessen (2008), among others, used renewal data to assess the value of patents, and found that citations-received, as well as some derivative citation measures, are positively and significantly associated with renewal. Renewal should be associated with complement citations, but litigation could be associated with either complement citations or substitute citations: Litigation might occur because a firm is using a complement without licensing it, or represent an attempt to claim that a replacement technology's patent isn't valid or that its use has still infringed the substituted technology.[16]

Lanjouw et al. (1998), Lanjouw and Schankerman (2004), and Bessen (2008) do not measure value conditional on use, and I will consider a sample of startup firms that are far more likely than the average patent-holder to have useful patents.[17] I could find only a single

---

[14]This argument is less sound for 'patent trolls' that buy up patents and then license them. A producing entity might want to explicitly cite a patent held by a troll to indicate substitution and so provide it with a shield from litigation. University Technology Transfer Offices (TTOs) realize their rents almost exclusively from licensing, rather than litigation. As a result, a patented university technology that has not yet been used may be less likely to accrue substitute citations, as the citing entity may be unaware of potential applications for the technology in their sector and are not likely to be threatened by the prospect of litigation if they create a replacement technology.

[15]Many patent offices require patent-holders to pay renewal fees to prevent their patents from expiring. As mentioned earlier, the USPTO requires renewal fees three times during the statutory term of patent.

[16]Failure to renew would be rational if the technology has been replaced. This provides an intriguing possibility for future research: If one makes the, admittedly heroic, assumption that citations received between renewals mostly reflect complements, and that citations received after the last renewal before a failure to renew a patent are mostly substitutes, one could code the population of citations-made and -received after December 1980 as either complements or substitutes. One could then experiment with value measures, much as in this thesis, to see whether this coding added information, and whether complements do indeed have positive (negative) value effects for citations-received (-made) and substitutes do indeed have negative (positive) value effects for citations-received (-made).

[17]Mann and Sager (2007) provide evidence of a correlation between the progression through a startup firm's (venture-capital-based) financing process and the accumulation of patents. Although the causation underlying this correlation is inherently difficult to assess, one should still expect to see successful firms having more patents. However, within my framework where citations represent the extent of complement or substitute-based relationships, I have no clear hypothesis for how the volume of citations will be related to

paper that reported estimates of the effects of citations on firm value for startup firms.[18] This paper reported a value effect that was not statistically significantly different from zero for "a variety of citation-based indicators of the importance and radicalness of the patents licensed." Their results therefore suggest that for startup firms citations, on average and in aggregate, do not reflect value.

In the context of citations-made, Jaffe (1986) found that there are both positive and negative effects on profits and market value from having cited other firms, noting that on balance the effects are positive. This would be consistent with cited patents acting as a mix of complements and substitutes, but this argument was not advanced in their paper.

This research is not the first to advance the hypothesis that citations might represent complementary or substitution-based economic relationships, but I believe that it is the first to consider both possibilities together, the first to attempt to determine the balance of complement and substitute effects, and the first to test the hypothesis that citations could conceivably contain this kind of information.

The earliest claim that citations represent the substitution effect of a new invention appears to have been made by Lerner (1994). Lerner (1994) suggested that patents that act as substitutes to broad range of classes are much more valuable than those that are substitutes to only a narrow range of classes, and provided evidence to support this notion using Herfindahl type measures constructed using citations-made and -received.[19]

The origin of the claim that citations represent complements, particularly requirements for complementary inputs, appears to have coincided with the development of the patent thicket literature. Hall and Ziedonis (2001), Shapiro (2001), and others, set the stage by suggesting that patents on complementary inputs provide a mechanism for the patent thicket, and Ziedonis (2004) operationalized a measure of the concentration of ownership among a patent's citations-made to explore this theory. Ziedonis (2004) calculated a Herfindahl-type measure of the fragmentation of ownership in citations-made for the patent portfolios of 67 publicly-traded semiconductor firms and showed that firms facing a more fragmented technology-base ultimately patent more; a result that is consistent with defensive patenting in the presence of a patent thicket.

Noel et al. (2006) used this ownership fragmentation measure to represent the costs of negotiating in a patent thicket for publicly-traded firms. Cockburn and MacGarvie (2009) used this measure as a proxy for the density of the patent thicket to study the effect of thickets on the extent and timing of investment into startup firms in 27 carefully defined software markets, as well as the effects on their hazard rates for securing an initial public

---

success; more citations-made could represent that the startup has a greater degree of both cited complements and cited substitutes.

[18]Shane and Stuart (2002) provided these findings, reported in considerable detail, in footnote 13 on pages 161-162. I hereby repay them.

[19]The calculation of such Herfindahl-type measures requires an adjustment for citation-count bias, detailed in Hall (2005). Without this adjustment the measures capture both the dispersion across classes/ownership/etc. and the number of citations-made or received. Ziedonis (2004), discussed below, uses the unbiased measures, but Lerner (1994) and many other papers (even after 2005) do not.

offering or an acquisition.[20] Cockburn and MacGarvie (2006) used the number of owners-cited, rather than a Herfindahl-based measure, to represent the extent of the patent thicket, to study the relationship between entry and patenting in the software industry. Cockburn and MacGarvie (2006) reported that firms are more likely to enter a market (where entry is defined as launching a product into a particular market segment, like 'dental practice management software') when they have patents in that market, and when those patents receive more citations. They interpret these findings as consistent with patent-holding firms gaining suitable leverage (i.e., complementary input patents) to overcome the barriers posed by a thicket.

The closest application of the notion of patents as complements and substitutes to that in my thesis appeared in Harhoff et al. (2003a). A companion paper, Harhoff et al. (2003b), reported that, in their survey of German inventors, about 35% had patents with some blocking power (i.e., that acted as complementary inputs), compared with just 3% that reported a case of patenting substitutes.[21] Harhoff et al. (2003a) advanced a sparse, symmetric-oligopoly framework to explore how the asset value of a patent might differ from its renewal value in three cases: simple innovation where one firm gets a patent that other firms don't; cumulative innovation with one patent acting as a complementary input; and the possibility that two patents (perhaps held by the same firm) cover perfect economic substitute outputs. In their paper, the asset value of a patent was defined as the rents that accrue to it through usage (including blocking), and the renewal value was defined as the value of the patent being kept out of the public domain. Thus, they argue that in the 'simple innovation' and 'cumulative innovation with a blocking-patent' cases, the asset value will exceed the renewal value, but in the case of the substitutes the renewal value will exceed the asset value as "it is better to have only one competitor matching one's own technology than to have many equally skilled adversities".

Harhoff et al. (2003a) include citation counts in their exploration of the effects of patent characteristics on reported values. However, citations are not integrated in any way into their discussion of complements and substitutes, and they do not attempt to determine whether citations could contain this kind of economic information or infer what the balance of complements and substitutes that is reportedly present in their sample implies about the economic outcomes that they examine.[22] Implicit in their economic reasoning is the argument that, when a new substitute arrives, a pre-existing patent becomes less valuable; and when a new need for blocking-patent arises, the blocking-patent must become more valuable. But

---

[20]Cockburn and MacGarvie (2009) report that holding a greater number of patents is positively associated with an initial public offerings; a result that is supported in my analysis. Their paper does not report a finding for the effect of the count of citations-made or -received.

[21]Which begs the question, "What about the other 62%?"

[22]In this thesis, I am questioning whether patent citations convey information about economic complements and substitutes, but I do not assert that every patent citation conveys this information or that citations convey it equally in any fashion. I fully expect that some in-term patent citations might reflect knowledge flows or other constructs. If this proportion is material then this would provide another reason for ownership to explain less than all of the meaningful variation in citations.

there is no suggestion that citations might indicate when these types of relationships are taking place.

The nascent literatures on patent pools, patents in standards, and patent licensing and litigation, all build on the notions of patents acting as complements or substitutes for one another. This thesis provides a foundation from which these literatures could begin to operationalize citation-based measures to represent the nature of the relationships between patents. Concerns about *royalty-stacking*, a practice where patents on multiple complementary inputs are licensed separately so that the cumulative license royalties are extortionate or prohibitive was voiced in Lemley and Shapiro (2006) and Geradin et al. (2008). One solution to this problem, particularly when the complementary patents are widely-held, is the creation of a patent pool. Patent pools are run my independent organizations; firms agree to allow the pool to negotiate license agreements on their behalf, and the pool disburses royalty payments back to its contributors according to a schedule of previously agreed rules.[23] Patent pools have re-emerged in recent years, as anti-trust authorities have come to understand that the costs of collusion can be more than offset by the benefits from allowing I.P.-holders to cooperate, particularly in environments characterized by cumulative innovation and complex products.[24] However, anti-trust authorities should, and do, object to the formation of pools that contain substitutes.

Patents in standards might be complements or substitutes. Even when a standard is formed, it is possible that substitute technologies that achieve the same functionality could have been included. However, after a standard is formed, the likelihood of a substitute technology being included in a standard increases dramatically. Standards undergo improvements and revisions, and many patents included in a revision to a standard may be replacing previous technologies. This is problematic from an anti-trust perspective because standards may be associated with patent pools; a pool for a standard may notionally allow manufactures to buy the right to implement the standard from the pool.[25]

Finally, I note that there is a nascent literature on the *market for technology*, where startup firms either compete in output markets or cooperate by licensing their inventions or sell themselves to incumbents, in order to appropriate the value from their innovative efforts. This literature, epitomized by Gans et al. (2000), argues that startups face an existential choice between competition and cooperation with incumbents: Cooperation is enabled by strong intellectual property rights, and its benefits increase when the startup faces a cospecialized-asset problem; competition, however, is preferred when the investment costs are low and the technology is not protected by patents, so that the "severe disclosure threat

---

[23]Lerner and Tirole (2004) provides rich institutional detail on patent pools.

[24]Lerner and Tirole (2008).

[25]It is important to note that not all patents required to implement a standard will necessarily be available from the pool. For example, this is appears to be with case MPEG-2 pool, where numerous developers that made declarations that they have patents relevant to the implementation of the standard, do not have patents in the pool. See Layne-Farrar and Lerner (2011) for a discussion of when firms contributing to standards decide to join pools.

tends to foreclose the ideas market".[26] As such, Gans et al. (2000) embodies the same Teecean arguments as this thesis. The main difference between our work lies in the economic factor that drives the relative returns to cooperation versus competition. For Gans et al. (2000) the strength of intellectual property protection drives this trade-off, with "[i]mperfections in the market for ideas... spur[ing] a competitive strategy by start-up innovators". Whereas in this thesis I argue that the primary economic factor is the nature of the technology itself: whether the invention is a complement participating in incremental innovation, or a substitute that is more likely associated with radical innovation and so creative destruction.

---

[26]Note that for Gans et al. (2000), the cospecialized-assets are distribution channels, brand names, or manufacturing expertise, and not complementary patents.

Chapter 4

Data and Measures

4.1 Data

My primary data source is the NBER patent dataset, which contains data on every utility patent application in the United States at the USPTO from 1963-2006.[1] This dataset is an extension of the dataset described in Hall et al. (2001). 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. I consider only firms that have one or more patents throughout all of my analyses.

The acquisitions are taken from the SDC Mergers and Acquisitions database from 1980 to 2010. I consider 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, this will very closely approximate the population of acquisitions. Privately-held firms are sometimes required to disclose material acquisitions under applicable Security Exchange Commission regulations, but not always. Therefore, it is less certain that I have something approaching the population for these firms. However, SDC (now owned by Thomson Reuters Corporation) collects data from surveys and press releases as well as from securities filings, and acquisitions above a certain threshold are very likely to be included. Low-value acquisitions are much more problematic, and I will later exclude them from my sample, even though my results are essentially unchanged by their inclusion.

The data on initial public offerings come from the Global New Issues (GNI) database (now also owned by Thomson Reuters). As these data are extracted from offering prospects and other mandatory security filings, these data represent the entire population. I consider data from 1986 to 2010; prior to 1986 GNIs data collection was not automated. To be included

---

[1]In this thesis, I will use data pertaining to patent applications rather than granted patents. The granting of a patent does convey additional information for startups; Greenberg (2010) documents a value premium to startup patent grants. Likewise, the granting of a patent will likely influence its citations-received. However, I have no reason to believe that the failure to achieve a grant introduces any systematic bias into my analyses.

in my dataset, IPOs must have completed listings on any U.S. exchange, but must not be a Leveraged Buyout (LBO) or a spin-off.[2] The firm must be privately-held and never have been publicly-listed or acquired before its initial (i.e., first) public offering. I also restrict the acquisitions in a similar way. I do not include acquisitions of firms that were private then sought a public listing and then were acquired. As a general rule I take the first of multiple liquidity events.

When I constructed the patent portfolios for the startups in my sample, I considered only patents that were filed prior to the startup's liquidity event.[3] This was done as I found that patenting behavior undergoes systematic and dramatic changes after a liquidity event depending on the type of event the firm experienced. For firms that are acquired, patenting drops precipitously because the startup generally ceases to legally exist. The volume of citations-received after a liquidity event also depends on the event type in much the same way: the patents that belong to acquired firms are cited less after an acquisition. It is unclear why this would be the case under the innovation-lineage view of citations, but under an economic view one might suppose that the technologies covered by patents held by acquired startups are used less by outsiders after the acquisition.[4] Regardless of the explanation, there is a clear, and apparently causal, correlation between patenting and citing behaviors after a liquidity event and the event type, which would introduce econometric problems when I estimate the event type using patents and citations.[5]

I therefore required that liquidity events took place between 1986 and 2006, and that firms had one or more patents prior to their liquidity event.[6] Prior to the constraint on the timing of liquidity events, I had 5,446 successful startups with one or more patents prior to their IPO or acquisition. This constraint removed 161 acquisitions prior to 1986, and 519 after 2006, as well as 35 initial public offerings.[7] Thus, after applying the constraint on timing, I had 4,731 successful startup firms, made up of 1,230 IPOs and 3,501 acquisitions.

---

[2]I identified, and excluded, 51 spin-offs in my data.

[3]My datasets were joined by name-based matching. Custom matching software provided normalization-based and algorithm-based matches, which were validated using state of incorporation and event-date information. In the U.S. companies are incorporated in a U.S. state, rather than federally. As a result, it is possible (though unlikely, as trademarks are federal), that two firms can have the same name but be incorporated in separate locations. Likewise, in very rare cases, I found firms with identical names in the same state, but operating in different time periods. Every effort was made to ensure correct matches, but it is possible that a small number of errors persist in my data. I do not believe that these have any material effect on my analysis.

[4]An alternative explanation is that acquirers do not, for some reason, apply for continuances for their acquired patents at the same rate as firms that achieved an IPO. Citations may be made to specific claims. With fewer additional claims added, acquired patents may receive fewer citations.

[5]I address this issue by including an indicator variable for IPOs or acquistions when I use citations-received after the liquidity event in a robustness check on my main results.

[6]By restricting attention to liquidity events that took place inside of the range of the patent data (i.e., prior to 2006, inclusive), I ensure that my analysis does not suffer from an automatic truncation-to-zero issue for citations-received.

[7]This period was characterized by a poor IPO market, and only a fraction of all IPOs have patents prior to their offering.

I refer to this as my 'full sample', which is described in the first set of rows of table 4.1.

Data on venture-capital-backed firms were taken from Thomson VentureXpert from between 1980 and 2010; the coverage of this dataset is regarded as approaching the population of venture-capital-backed firms. I checked that my start-ups included all successful VC-backed firms, and constructed a VC-backed indicator variable accordingly. In addition, I constructed a sample of 'failed' venture-capital-backed startups, where I defined a startup 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. I was therefore able to construct a set of 1,311 failed VC-backed firms to cover the entire 1986-2006 period.

Disclosure of firm values is particularly problematic in low-value acquisitions. I was concerned that many of these acquisitions might be essentially 'fire-sales' or other non-successful liquidity events. And I need values to be disclosed for my value analysis. Thus I constructed my 'value sample' as the subset of my full sample where value was disclosed. My value sample is described in the second set of rows of table 4.1, and will be used only to check the robustness of my results. A total of 2,019 (almost 60%) of my acquisitions were discarded because they did not have disclosed firm values.

I then further refined my sample by requiring that any successful startup firm was valued between U.S. \$10m and U.S. \$2.5b inclusive at its liquidity event. This involved throwing out a further 412 acquisitions and 76 initial public offerings. The lower bound was imposed primarily to remove low-value and unrepresentative acquisitions, and the upper bound was imposed as there was only a single acquisition but 36 initial public offerings above this level: The largest acquisition was a little under U.S. \$2.7b, and the largest IPO a little over U.S. \$77b. I did this in an attempt to ensure common support (in terms of firm value) for IPOs and acquisitions. A more sophisticated approach would involve matching of IPOs to acquisitions. However, with sufficient firm-value controls, a common support should suffice to ensure that I am comparing apples to apples. That said, I stress that my main results are unaffected by these choices; I simply choose to report the most cautious results.

After all of the constraints, I was left with a 'final sample' of 801 IPOs (excluding spin-offs) and 1,070 acquisitions (of privately-held firms that had never sought a public listing for their stock), that occurred between 1986 and 2006, with firm values between \$10m and \$2.5b, for startup firms that held at least one patent prior to their liquidity event. This gave me a final sample of 1,871 successful startups, 579 (31%) of which were venture-capital-backed, and a companion sample of 1,311 failed (venture-capital-backed) startups. My final sample is detailed in the third set of rows of table 4.1.

Below the description of my final sample in table 4.1, is the description of three sub-samples. Single patent-holders, that is startups with just one patent prior to their liquidity event, account for approximately 30% of my final sample. I will use this sub-sample for robustness checks, as it affords a straight-forward mapping between firms and patents. Multiple-patent holders, which account for the other 70% of my final sample will be useful when I consider self-citations. A self-citation is a citation between two patents in the same patent-holder's portfolio; thus it is only meaningful to test the impact of self-citations in a sub-sample of firms with two or more patents. The final set of rows of table 4.1 describe

Table 4.1: Description of samples

	N	Firm Value μ	No. Patents μ	Avg. C. Made μ	Avg. C. Rec'd μ
Full Sample					
IPOs	1230	708	15.2	11.9	5.4
Acquisitions	3501	84.1	13.4	9.5	6.9
Total	4731	316.4	13.9	10.1	6.5
Value Sample					
IPOs	879	708	13.2	12	5.2
Acquisitions	1482	84.1	8.3	10.8	7.2
Total	3672	203.4	8.5	11.4	7.6
Final Sample					
IPOs	801	246.1	10.9	12.3	5.3
Acquisitions	1070	108.4	9.8	11	7.4
Total	1871	167.4	10.3	11.6	6.5
Single Patent-holders					
IPOs	193	199.9	1	9.9	5
Acquisitions	352	93.4	1	9.1	8.1
Total	545	131.1	1	9.4	7
Multiple Patent-holders					
IPOs	608	260.8	14.1	13.1	5.5
Acquisitions	718	115.8	14.1	11.9	7
Total	1326	182.3	14.1	12.5	6.3
VC-backed sub-sample					
IPOs	310	250.3	10.2	15.4	6.3
Acquisitions	269	111	6.2	15.5	9.4
Total	579	185.6	8.4	15.5	7.8
Failed Firms	1311	-	5.7	11.8	9.7

the sub-sample of venture-capital-backed startups, which also account for about 30% of my final sample. The venture-capital-backed sub-sample is the only sub-sample with a higher incidence of IPOs than acquisitions. As reported in Hellmann et al. (2005), acquisitions are generally about three times more common than IPOs for venture-capital-backed firms. Even taking out low-value acquisitions, acquisitions remain more than twice as common as IPOs. My sub-sample of successful venture-capital-backed startups closely approximates the population of successful venture-capital-backed-startups that hold patents. That is to say, holding patents, and not other constraints, is what has made IPOs so much more common for VC-backed firms. The final line of table 4.1 describes the companion sample of failed venture-capital-backed firms that I will use in my analysis of the influence of citation-based measures on success.

Figure 4.1: The distribution of citations received
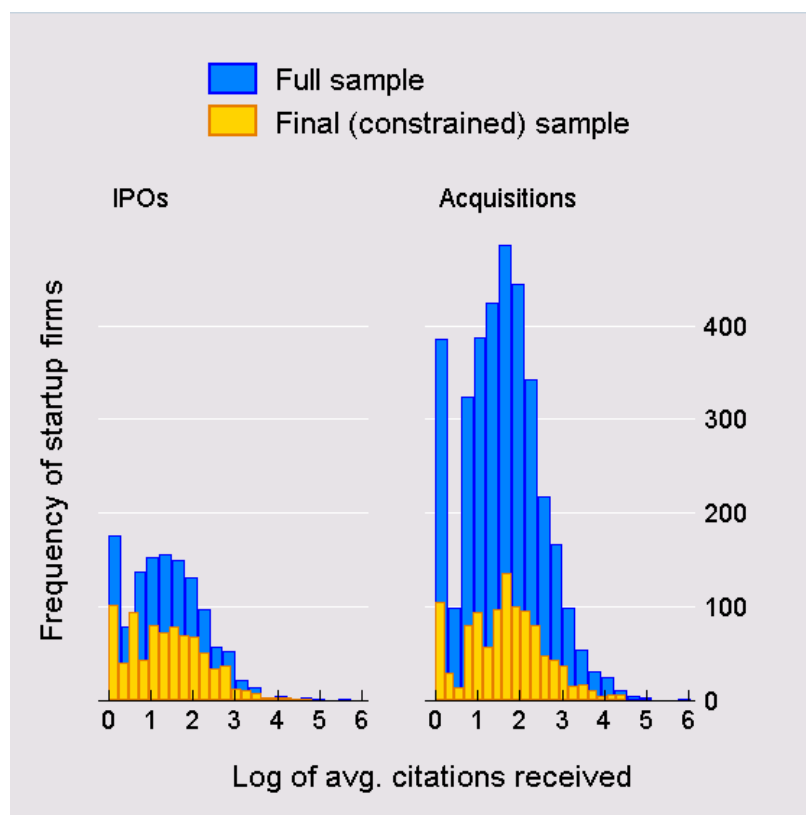
The two biggest differences between my full sample and final sample lie in the mean values of firms and the average number of patents held by firms. The mean value of IPOs is reduced dramatically from \$708m to \$246m, reflecting the removal of the high-value tail. And the mean value of acquisitions increases by around 30% from \$84m to \$108m, reflecting the removal of the low-value tail. The number of patents held by a typical startup drops for both IPOs and acquisitions from my full sample to my final sample. However, the mean number

of citations-made and -received look remarkably similar between the two samples, which suggests that I have not selected a sample based on the characteristics of my explanatory variables. Further evidence that the distribution of citations-received was not materially altered by sample selection is presented in figure 4.1, where I provide histograms for the log of the average count of citations-received to patents held by startups that achieved initial public offerings and acquisitions for both my full sample and my final sample.

Figure 4.2: The number of IPOs and acquisitions over time
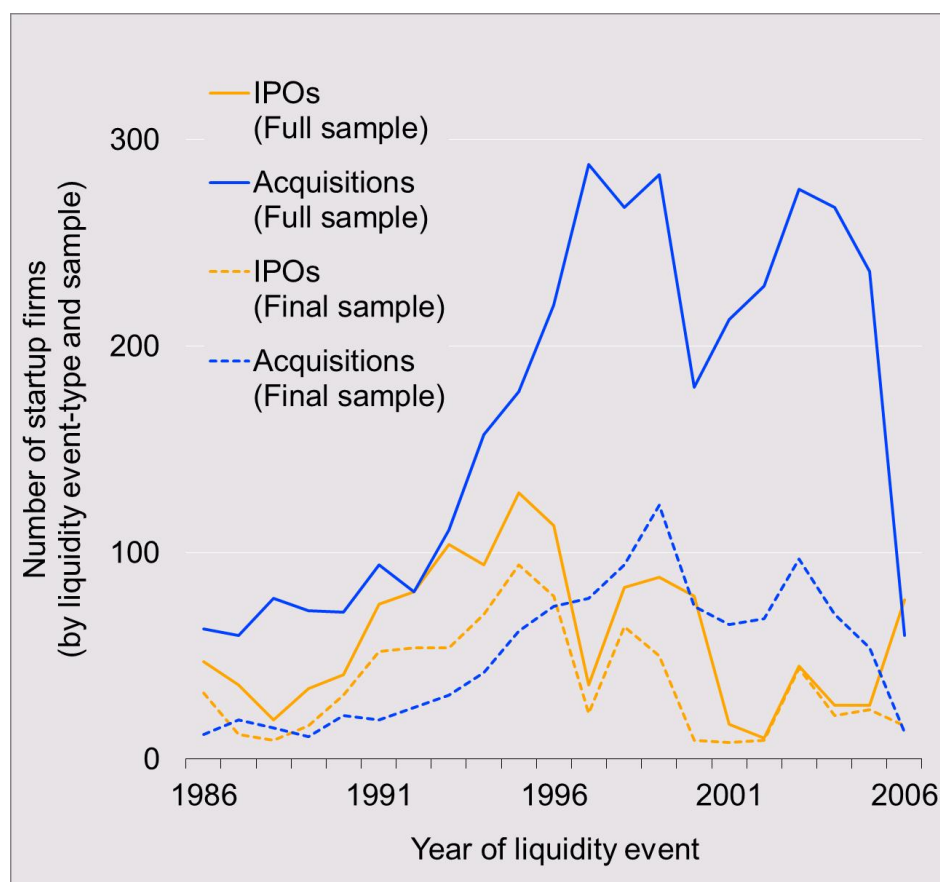
In figure 4.2, I provide a visual description of how the composition of my full and final samples vary over time. The solid gold line plots the number of startup firms achieving an initial public offering each year in my full sample. There are two peaks, one in the early '90s and a second, corresponding to the Internet boom, in the late '90's. The IPO market also appears to make a recovery in 2006. The dotted gold line, representing IPOs in my final sample, follows a very similar pattern at a slightly lower level. The solid blue line depicts acquisitions in my full sample. It also has two peaks, occurring around three to six years after the peaks in the IPO market. The dotted blue line, depicting acquisitions in my final sample, is dramatically below the solid blue line, as over $\frac{2}{3}$rds of the acquisition in my full

sample were discarded (primarily because their value wasn't disclosed). However, the dotted blue line does appear to follow the same overall trends as it's solid blue counterpart.

4.2   Measures

Citations might come from the firm itself, another firm in my sample of successful startups, a firm in my control sample of failed venture-capital-backed firms, a publicly-traded firm, or another firm, organization, government, or individual from the population of patent applicants who filed for a patent between 1963 and 2006, that I will use as a reference set.

For every citation I have the application and grant date of both the citing and the cited patent. Although citations were electronically recorded only beginning in 1975, a cited patent might have been filed many years previously. The determination of whether a citation is classified as in-term or out-of-term depends on the whether the citing patent was applied for after June 8$^{\text{th}}$ 1995. If it was and the cited patent was applied for less than 20 years previously, or if it wasn't and the cited patent was granted less than 17 years previously, then I declare the citation to be 'in-term'. Otherwise citations are declared to be 'out-of-term'. I count the number of in-term and out-of-term citations for each patent in a startup's portfolio, and then average these counts across the portfolio to create my in-term and out-of-term citation count measures.

I determine the ownership of patents through USPTO assignee codes.[8] On a per patent basis, for each patent in a startup's portfolio, I count the number of citing and cited owners. These counts are then averaged to create a measure for each startup. A simple example may be instructive: Suppose a startup has two patents, the first of which cites two patents owned a single patent-holder and the second of which cites one patent held (naturally) by a single (possibly different) patent-holder. Then this startup would make an average of $\frac{2+1}{2} = 1\frac{1}{2}$ citations to an average of $\frac{1+1}{2} = 1$ owners.

When a citation involves a successful startup, a failed startup, or a publicly-traded firm, I know the industry classification of the citing or cited firm, and this allows me to calculate whether the startup and the citation-owner operate in the same sector. I make use of the industry classification of all of the startups in my initial sample (i.e., my full sample before liquidity-event year and other constraints), not just the 1,871 that remain in my final sample used in the analysis, for this purpose.[9] The assignment of NAICS codes to operating sectors

---

[8]The recording of USPTO assignee codes begins in 1975, and is not complete or perfect in the data. A total of 9.93% of patent records from 1976-2006 are missing assignee codes, and a further 14.84% of records have an 'unassigned' assignee code. Furthermore, I anticipate that a small percentage assignee codes are 'incorrect' and, although I attempt to rectify this, sometimes multiple assignee codes are given to the same entity. I have no reason to believe that any of these problems introduce any systematic bias into my results.

[9]I have a close match to the population of publicly-traded firms from joining the patent data to COMPUSTAT. Coverage of NAICS codes in COMPUSTAT varies from 63% in 1976 to 100% in most years, with an overall average of 92% of firms coded. My full sample of startups and publicly-traded firms together account for a little under 70% of all U.S. corporate patent holders, with the missing 30% being privately-held companies.

is detailed in the appendix.

I calculate an average *portfolio time-distance* variable to measure the mean time between patent applications and the liquidity event for all of the startup firms in my sample. The portfolio time-distance is important for two reasons: firms with patents further from their liquidity event have a longer window in which to accumulate citations; and, as I will discuss later, citation counts suffer from inflation over time, so that a patent in the 1990's will generally make and receive (per year) more citations than a patent in the 1980's.

I have considerable detail on both the characteristics of the startups in my final sample, and their patent portfolios, that I use to construct control variables for the regression analysis. Specifically, one might be concerned about geographic effects, industry effects, time effects, and technology-class effects, as well as the impact of firm values on the liquidity event type, and vice versa. I discuss how I address these concerns in the section on my empirical approach.

4.3 Descriptive Statistics

Between 1986 and 2006 there were a little over 2.43m patent applications filed with the USPTO.[10] These applications were assigned to approximately 3.67m assignees. The NBER patent data provides information on the type of the assignee for approximately 2.4 million patents from 1976-2006. These types include permutations of 'foreign' and 'U.S.' with 'corporation', 'government' (state and federal), 'university', 'institute' and 'hospital'. U.S. corporations account for 1.2m (or 50%) of these patents and foreign corporations for 1.1m (or 44%), while all other categories account for less than 2% of patents each, and less than 6% combined together.

My reference dataset of publicly-traded firms comprises largely (92%) U.S. firms, and a little under $\frac{3}{4}$ of its 1.05m patents are assigned to these U.S. firms. Therefore my reference set of publicly-traded U.S. firms account for approximately 60% of all the patents assigned to U.S. corporations.[11]

---

[10]The 'grant rate' for patents at the USPTO varies considerably over this time period. Using simple counts of applications filed and patents granted, Jaffe and Lerner (2006) and Hall and Harhoff (2004) provide evidence consistent with a decline from about 70% to about 50%. However, Lemley and Sampat (2008) estimates that about 75% of applications result in at least one patent. Grant rates are difficult to calculate for several reasons including: the increasingly long processing time at the USPTO; the possibility that claims may be amended, added or retracted; that some applications are actually continuances; and that in some cases a single application may result in multiple patents. Taking all of these factors into account. Quillen Jr et al. (2002) estimate an effective grant rate of approximately 85%. Furthermore it is likely that grant rates vary according to the type of the applicant. Individuals with constrained resources may face a lower grant rate than corporations. The former Chief Patent Council to a large computer manufacturer once boasted to the author that he had "never lost a patent application".

[11]This is directly comparable to the 50-65% (depending on year) figure for the 1965-1995 period reported by Hall et al. (2005).

My final sample of 1,871 successful start-ups filed a total of 26,709 patent applications in this period.[12] Therefore my final sample of successful startups accounts for a material 1% of U.S. patenting and about 2% of all U.S. patenting by U.S. corporations, despite accounting for less than 0.05% of the assignees.[13] However, my original sample of 1,525 IPOs, 4,181 acquisitions and 1,311 failed VC-backed firms (i.e., the sample I had before excluding spin-offs, placing constraints on when firms experienced their liquidity events, or restrictions on firm values, totaling 7,017 firms, had 93,000 patents, or about 8% of the total assigned to U.S. corporations. This leaves at most 32% of U.S. corporate patents for privately-held firms that would not become public or get acquired from 1976-2006.

Recently, in a given year, there have been approximately 6.5m active privately-held corporations and about 4,000 active publicly-traded firms incorporated in the United States.[14] Startup firms therefore appear much closer to publicly-held firms than an 'average' privately-held firm in terms of the volume of their patenting activities. In the appendix, in table 9.2, I provide counts of the numbers of patent-holders (i.e., owners) and patents in each sector for my final sample of IPOs, acquisitions and failed firms, as well as for publicly-traded firms listed on the Nasdaq, NYSE and AMEX exchanges.

Startups in my final sample of 1,871 firms have roughly comparable quantities of patents to publicly-traded firms in the Computer Media, Internet, Life Sciences, Biotech and Chemicals sectors. In all other sectors, the volume of filings-per-firm is roughly an order of magnitude higher for public firms. The overall mean value of a successful startup at its liquidity event was U.S.$167.4m, and the overall mean number of patents filed before the liquidity event was 10. This compares with a mean market value for publicly-traded manufacturing firms of U.S.$916.3m, with an average stock of 76 patents, reported by Hall et al. (2005).

A naive view that judged patents, or citation-weighted patents, as comparable units of innovation might therefore consider startups to account for a little under 10% of the U.S. corporate innovative output, second only to public firms. However, in this thesis I hope to get across a more sophisticated message: It is not just the volume of patents that matters but the nature of the innovation that they represent. I begin my exploration of the nature of startup innovation in table 4.2, where I report the means and medians (and number of observations for which a variable was equal to zero) for my main citation-based measures separately for both the acquisitions and initial public offerings in my final sample of successful startups.

There are five important observations regarding table 4.2. First, every one of my variables is skewed. I therefore follow the convention of using log variables in my analyses. Second,

---

[12]A total of 995 (53%) of the successful startups in my sample filed only a single patent, the remainder had accumulated a portfolio of patents prior to their initial public offering or acquisition.

[13]It should be noted that not all 'successful' innovative startups need to secure commercialization financing from an initial public offering or an acquisition. Aside from other sources of investment, such as private equity, Angels, debt and Friends/Family/Founders, some successful startups may be self-financing. As such, one might expect that there is a material proportion of privately-held firms that are both highly successful and innovative.

[14]Asker et al. (2010).

Table 4.2: Descriptive statistics

	IPO ($N = 801$)			Acquisition ($N = 1070$)		
	μ	p50	# $= 0$	μ	p50	# $= 0$
Firm Value (U.S.$m)	246.1	117.7	0	108.4	40.5	0
No Patents	10.9	4.0	0	9.8	3.0	0
Avg. Cites Recd	5.3	3.0	93	7.4	4.3	101
Avg. Cites Made	12.3	7.0	17	11.0	7.0	45
Avg. Portfolio Time Dist.	4.2	2.8	0	6.5	4.7	0
Avg. C. Made In Term	12.0	7.0	17	10.6	7.0	45
Avg. C. Made Out of Term	0.3	0.0	621	0.4	0.0	713
Avg. Owners Made	7.4	5.0	29	6.7	5.0	65
Avg. O. Made In Term	7.2	5.0	29	6.4	4.7	66
Avg. O. Made Out of Term	0.2	0.0	636	0.3	0.0	757
Avg. C. Made In Cat.	9.4	5.5	23	8.5	5.5	60
Avg. C. Made to Public	10.4	5.6	160	0.8	0.0	985
Avg. C. Made in Sector	6.6	1.2	379	0.4	0.0	1031
Avg. C. Rec'd. In Term	5.3	3.0	93	7.3	4.2	103
Avg. C. Rec'd. Out of Term	0.0	0.0	781	0.1	0.0	979
Avg. C. Rec'd. In Cat.	4.5	2.2	114	6.0	3.1	127
Avg. C. Rec'd. to Public	3.0	1.2	179	3.5	1.2	241
Avg. C. Rec'd. in Sector	0.9	0.2	366	0.7	0.0	595

firm values differ considerably between initial public offerings and acquisitions despite the requirements regarding values I placed on my sample in an attempt to make them comparable. I take measures to control for these differences that are detailed in the analysis section, but even so this difference does not drive any of my results. Third, the average count of citations-received differs between IPOs and acquisitions. This difference is explained almost entirely by a systematic difference in the length of the window, from the mean application time to the liquidity event, that firms had in which to receive citations to their patents. Again, I refer to the length of this window as the average portfolio time-distance, and it is almost $1\frac{1}{2}$ times as long for acquisitions as it is for IPOs. Fourth, there is a systematic difference between the two organizational-form choices with respect to the citations-made to publicly-held firms and citations-made in-sector. This difference is not apparent in the citations-received from publicly-held firms or citations-received in-sector, and is explored later. And fifth, for many of my measures a large fraction of firms report zero. That is, for example, 379 (47%) out of 810 IPOs and 1031 (97%) out of 1,070 acquisitions do not make even a single citation to a firm in the same sector of operation as them.

Chapter 5

Empirical Approach

The first step of my empirical approach is to show that some citations do contain economic information by contrasting the effect on an economic outcome variable of citations that cannot contain direct economic information with those that plausibly could contain such information. The second step is then determining whether patent counts, beyond ownership of cited or citing patents, explain at least some of the variation in an outcome variable, as I claim they should if patent citations represent certain imperfect-substitution relationships. These two steps can be achieved by direct observation: for the first I look for a statistically significant difference in the coefficients for in-term and out-of-term citations, and for the second, I look to see whether patent counts beyond ownership explain some of the meaningful variation in an outcome measure. Neither step requires, a priori, a hypothesis regarding a specific magnitude or sign of a coefficient. That is to say, the first step does not require a prediction of the sign or magnitude of the effect of knowledge flows or the effect of (direct) economic information, merely that these two have statistically different effects, and in the second step the impact of ownership will be judged relative to the effect of the economic information in citations. However, I will need a statistically significant effect for citation-based measures on at least one outcome measure to conduct this analysis.

I consider three economic outcome measures, each of which addresses a fundamental question concerning the relationship between the innovation activities and the performance of startup firms. These are whether or not a startup succeeds in reaching a liquidity event; the value of a successful startup at its liquidity event; and the organizational form choice of the startup at this event (i.e., whether the startup secured an initial public offering or underwent an acquisition). I will find statistically significant effects for citations in only one of these measures – the organizational form of startups at their liquidity event. The first two steps of my analysis, in which I determine the presence of economic information and the relative importance of ownership, are therefore undertaken in this context.

The third and fourth steps of my empirical approach rely on indirect observation. In these steps I will extract information about the balance on complements and substitutes in my sample of successful startups from the effects that patent citations have on two outcome variables: the value of a startup at its liquidity event, and the commercialization strategy

that the startup chose to pursue (i.e., the type of liquidity event that it sought and achieved). In these steps I have clear hypotheses regarding the signs (and significance) of coefficients; each sign is associated with a specific interpretation about the balance of complements and substitutes. For the value effect and citations-received, a positive sign is associated with complements, a negative sign is associated with substitutes, and a null finding is associated with an indeterminate mixture. For citations-made, the value-effect interpretations are reversed: a positive sign is associated with substitutes and a negative sign with complements. A null finding remains associated with an indeterminate mixture. Likewise, for the strategy effect and citations-made, a positive correlation with the likelihood of an acquisition is associated with complements in the presence of a patent thicket and a negative correlation (or equivalently a positive correlation with the likelihood of an IPO) is associated with substitutes. Again, these are reversed for citations-received.

In broad terms, the first step acts as a test of whether there is economic information in patent citations, and the second, third and fourth steps act as differential tests of whether this information pertains to complements or substitutes, or a mixture of both. The analysis is 'successful' if all four steps provide information that taken together gives a consistent and coherent story. However, there are several other tests that will be conducted along the way that should also support the story. These include a requirement that ownership should not explain out-of-term variation, which should pertain to ownerless knowledge flows and not economic information to be consistent. And likewise, I have argued that substitutes citations are more likely with in-sector patent-holders than out-of-sector patent-holders. Thus, results otherwise consistent with the presence of substitutes would be further supported by findings that in-sector patent citation effects are less explained by ownership than out-of-sector patent citation effects, and undermined otherwise.

The fifth step of my analysis takes a different approach. The economic view does not provide hypotheses regarding the influence of citations on the likelihood of success distinct from that of a value effect - it is possible that firms can be successful by creating new combinations (i.e., patenting inventions that cite many complements) or by engaging in creative destruction and replacing the last generation of technology (i.e., patenting substitute inventions), putting aside whether a firm's patents are more valuable or not for reasons relating to complements or substitutes. Therefore, from the perspective of the economic view this analysis does not add new information (i.e., it is a repeat of the value-effect analysis with a binary outcome measure indicating high-value and so success, or low value and so failure). Absent distinct hypotheses regarding success, the economic view is supported when the findings for patent citations' 'success effect' are the same as those for the value effect. The natural assumption is that success is related to achieving some value threshold. If the value effect for citations-received, for example, is positive, reflecting complements, then success should be more likely with a high count of citations-received. Conversely, if the value effect for citations-received is negative, reflecting substitutes, then success should be more likely with a low count of citations-received. A finding of a null value effect would be consistent with a finding of a null success effect.

However, this fifth step of my analysis is informative about the innovation-lineage view.

The innovation-lineage view claims that citations-received proxy for the quality and so value of a patent. One defense of a null finding for the relationship between the count of citations-received and value is that there is insufficient heterogeneity in patent values due to a sample selection problem. Successful startups are likely to have high-valuable patents. It is possible that all startup patents are so inherently valuable that receiving additional citations has no meaningful value impact. Therefore, a null finding for a value effect could be still be consistent with the innovation-lineage view if there were a material success effect.

5.1 Economic Meaning in Patent Citations

My outcome measures are at that the firm level, making the firm my unit of observation. Each firm in my sample has a portfolio of one or more patents prior to its liquidity event, and each of these patents may cite or be cited by between zero and several hundred other patents.[1] I do not have any firms with solely out-of-term citations, making it more difficult to cleanly identify how an outcome varies with term considerations.

Regressing the counts of in-term and out-of-term citations on an outcome variable and examining the difference in coefficients for statistical significance is problematic, as there is an obvious omitted covariate – the 'age' of the citation. For a patent filed before June 8th 1995, the statutory patent term provides rights that persist for 17 years after the patent is granted.[2] Suppose that a cited patent was filed 16 years previously or 18 years previously; I would like to attribute the difference in the coefficients to the cut-off, and not to the difference in ages, which is masked by the in-term and out-of-term assignment.

Fortunately, I find that in-term and out-of-term citation counts are both statistically significant with opposite signs. A linear age effect would not reconcile these findings.[3] I therefore consider a quadratic average citation-age effect. As a second approach, I also consider the fraction of citations that are in-term versus out-of-term in combination with an average age effect. This allows me to trade-off the relative importance of an age effect against whether the citations' ages fall on one side of the boundary or the other.

I have argued that for reasons relating to bargaining for complements and rent-sharing constraints for substitutes, ownership should matter. Specifically, if patent citations convey economic information then the counts of owners-cited should be an informative component of the count of patents-cited. The counts of owners-cited and patents-cited will be correlated

---

[1]U.S. Patent No. 5,774,660, belonging to Resonate Inc. (a startup that sells 'load balancing' network technology, which was VC-backed and achieved an IPO in 2000) is the most cited patent in my sample, having received 354 citations between its application in 1996 and the end of my patent data in 2006. My record for Resonate Inc. shows that it applied for two other patents prior to its liquidity event. U.S. Patent No. 6,587,438 received 19 citations and U.S. Patent No. 6,182,139 received a single citation. Resonate's average of $124\frac{1}{3}$ citations-received is the highest achieved by a successful startup in my sample.

[2]A patent may not be renewed to full term, and its term may be lengthened or shortened in other ways, which introduces noise into this argument.

[3]As the group of 'young' citations has a negative slope and the group of 'old' citations has a positive slope, the age effect must at least trace a quadratic U-shaped curve to fit the data.

for purely mechanical reasons. If a patent cites just one patent, it must cite just one owner: The number of owners cited must be less than or equal to the number of patents cited.[4] To address this I adopt a two-stage regression approach. In the first stage I regress the count of owners on the counts of citations and store the residual, and in the second stage I regress both the count of owners and the residual on the outcome variable.[5]

The residuals from the first stage give me a measure of the count of citations that is orthogonal to the count of ownership. I can then consider the relative explanatory power of ownership separate from the additional effect of patent counts by comparing the z-scores of the coefficients of these two variables in the second stage regression.[6]

A similar approach is taken when I consider citations-made and citations-received together. Citation counts suffer from class-based inflation, as demonstrated in the literature on patent networks.[7] Particularly in the last 30 years, the expected number of citations that a given patent will make has risen substantially over time, perhaps because of the increasing technological complexity of innovation over this period.[8] A typical old patent might make, say, 3 citations and receive 4, and a typical new patent might make, say, 11 citations and receive 9. I therefore anticipate that there will be a mechanical correlation between citations-made and -received, as when the counts of citations-made are low, it is likely that the counts of citations-received are also low, and likewise for high counts. I anticipate that this will be partly addressed by year fixed effects, as well as patent class-based fixed effects, but nevertheless some unwanted correlation may remain and therefore I will also use a two-stage regression approach to isolate the component of citations-received that is uncorrelated with citations-made in my analysis.

5.2 Influences on Commercialization Strategy

I cannot observe when citations represent complements or substitutes directly, but I can do it indirectly by using three differential tests which together can provide a consistent and

---

[4]This excludes the possibility that a patent has multiple assignees. Approximately 2% of the patents in the NBER patent data have multiple assignees. Less than 1% of the patents assigned to startups in my sample were also assigned to another party. I experimented with several different methods to deal with this phenomena, including the construction of 'multiple assignee' variables that were included in unreported regressions. No significant effects relating to multiple assignment were found, and in this thesis I use (and count) only the first-listed assignee.

[5]The first stage is performed using Ordinary Least Squares on log counts, and the second stage uses a logit regression, as the outcome variable is a binary indicator taking the value 1 for an acquisition and 0 for an IPO.

[6]The comparison is based on z-scores, rather than t-statistics, as the coefficients in question will be determined using maximum likelihood estimation in a logit regression.

[7]This arises from a property called 'homophily' in network formation models, where nodes (patents) tend to form edges (to cite) to similar nodes (patents). An excellent discussion of the empirical properties of patent networks, and which formation models most closely resemble reality, is provided in Leskovec et al. (2005).

[8]Marco (2007).

coherent picture. Two of these three differential tests are conducted in the context of the effect in-term citations have on the liquidity-event type of the startup.

I consider this event type to be an organizational form choice for the startups in my sample – a choice that is surely endogenous to a large number of factors. I take the view that a startup's 'choice' of an initial public offering or an acquisition reflects its equilibrium outcome given a commercialization strategy that is rationally optimal based upon exogenously given technological endowments. Put simply, I assume that firms cannot choose which citations to make because they cannot choose which technologies they are going to need or replace when they come up with a new invention. I envisage a process where firms invent something and then choose how to try to commercialize it. Implementing the commercialization process that best matches the technology then results in the firm securing commercialization financing from either an IPO or an acquisition.

I argued that if a firm is producing citing substitutes then it should seek an initial public offering, and if it is producing cited complements in the presence of a patent thicket then it should seek an acquisition. I claimed that this should be particularly true with competitive citations, which represent relationships with direct competitors operating in the same sector as the startup. And I argued that ownership should explain all of the meaningful variation in outcome measures except in certain cases involving imperfect substitutes, where citations to individual patents will matter. Furthermore, I noted that competitive citations should be more likely to convey substitution information if there is any, as a citation to a non-competitor conveying substitution can be irrelevant.

I use a logit regression framework to examine how a startup's technological characteristics influence whether it seeks an initial public offering or an acquisition, and so conduct these two differential tests.[9] My samples comprise successful startups and the outcome variable is a binary variable taking the value '1' if the startup was acquired and '0' if it achieved an initial public offering. Key explanatory variables include the count of in-term citations and the count of in-term owners-cited, as well as equivalent measures in-sector.

There are many characteristics of the firm, which may be correlated with the nature of its technology, that likely affect the type of liquidity event. I will therefore include a variety of control variables in my reported specifications. If a control variable is positively correlated with the nature of the firm's technology, then its inclusion will lead the coefficients on my citation measures to understate the technology effect. This is because the control variable will explain some of the meaningful variation in the explanatory variable. And, conversely, a negative correlation will lead to an overstatement of the technology effect.

The firm's value at the liquidity event, the industry that the firm operates in, the timing of the event (particularly with regard to the industry's life-cycle) and the technological classes of the firm's patents, might all predict the liquidity event type and be correlated with whether the firm's patents represent complements and/or substitutes. Firm value, in particular, might be highly problematic. A firm's value at its liquidity event is the end product of everything that the firm has done until then. At the end of chaper 2, section

---

[9]The logit regression is detailed in McFadden (1973).

5, I discussed the relationship between allocative efficiency and value, and noted a direct endogeneity issue – citations may predict both value and organization form, and value may predict organizational form. I will later report a null value effect for patent citation-based measures of any kind, which helps alleviate this concern. Omitting firm value from the regression analysis is also problematic. The 'experiment' that I would like to conduct is to take two firms that are identical in all respects and then endow them with different technological natures (without making them have different values or changing them in any other way). Comparing a U.S.$100,000 acquisition with a U.S.$70b IPO is a comparison of apples to oranges. Clearly many things other than just their technologies are vastly different between two such firms.

I therefore adopt two approaches. The first approach is to use a sample of startups that has a common value support, and then to include controls for firm value. As I can observe distributional differences in firm values for IPOs and acquisitions between the supports, I will experiment with quadratic and higher-order polynomial specifications, as well as quantile-based fixed-effects for capturing the effects of firm value. As it happens, none of these controls will materially influence the coefficients on my patent citation-based explanatory measures. The second approach is use my full sample of successful start-up firms, without imposing a common value support, and not to include value as independent variable. As I will later show, all of the statistical significance and the signs of the coefficients on my citation-based explanatory measures are the same using this approach. The coefficients are even fairly similar in magnitude; these coefficients are robust to the inclusion of value measures as explanatory variables.

Both initial public offerings and mergers come in sector-based waves. The late 1990's, for example, was an era characterized by easier initial public offerings for Internet firms, whereas the telecommunications sector had two merger waves in beginning in 1987 and 1997 as a result of deregulation.[10] I therefore use 'sector x year' fixed effects, and report the average effect of citations within sector-years. Likewise, I use the categorization of patent classes put forward by Hall et al. (2001) to construct modal patent-category fixed-effect control variables.[11] However, none of these variables will drive my results and the coefficients on my explanatory variables are essentially unchanged by their inclusion.

When I consider competitive citations, I include a measure of 'in-category' citations to measure the effect of making citations into the same category of technology classes. I expect that technology classes will be associated with sectors, and therefore that in-category citations will act at least partly as a proxy for competitive citations.[12] I judge whether competitive citations are important by how much information they add to the model.

Within a logit framework, assessing the 'goodness of fit' of a model is somewhat difficult. In Ordinary Least Squares (OLS), the R^2 of the model provides a statistic on the $[0, 1]$ interval that describes the amount of the explained variability. This statistic is comparable

---

[10]Harford (2005).

[11]I experimented with both course- and fine-grained categorizations.

[12]I will find that they behave like somewhat like a weak version of the count of competitive citations but that they also contains some other information.

across models and even datasets. Logit regressions use maximum likelihood estimation, and there is no corresponding measure of explained variability. Instead, a number of authors, including the logit's creator, have put forward various pseudo-R^2 measures, which are only comparable across models with identical outcome data. In my tables, I follow convention and report McFadden's pseudo-R^2. However, in the text I will also report McKelvey & Zavoina's pseudo-R^2, which Veall and Zimmermann (1996) demonstrate has the closest relationship to an OLS R^2, as well as the *Count* and *Adjusted Count* pseudo-R^2. The Count pseudo-R^2 treats any observation with a predicted outcome-probability greater than 0.5 as having an outcome of 1, and otherwise as having an outcome of 0, and then computes: $R^2 = \frac{Correct}{Total}$. The Adjusted Count pseudo-R^2 takes into account the null model, by including a factor, n, that measures the most frequent outcome: $R^2 = \frac{Correct-n}{Total-n}$. As such, the Adjusted Count pseudo-R^2 measures the proportion of correct predictions beyond that which would be achieved by random assignment.

5.3 Estimating the Effect of Citations on Value

In my estimations of a value effect for citations, I use an approach that differs from that employed by Hall et al. (2005) in three important ways. First, I estimate the importance of the contribution of patent citations to value directly, rather than through Tobin's Q.[13] My measure of value is the long-term value of the firm, calculated using the market capitalization at the initial public offering or the sale price of the firm at acquisition. I explored other measures of value, notably the proceeds raised at IPO and the proceeds to the vendor net of fees for acquisitions, but using these measures did not materially alter my results.

All of the measures of value I tried represent, with varying degrees of noise, the return to the original shareholders at a 'changing of the guard'. That is, at IPO or acquisition, the firm is taking on new shareholders and offering liquidity, perhaps after a lock-up period, to old shareholders for the net present value of the future cash-flows that the firm is expected to earn. I have a large dispersion in the average counts of citations made by and received to patents that a startup has filed before its liquidity event and a large dispersion in firm values at the liquidity event; I simply seek to determine whether the variation in one is associated with the variation in the other.

However, I must be careful not to attribute any association in the variations of these measures with correlated covariates. The most important potential controls are the organizational form choice, as I know that IPOs have higher average values than acquisitions (even though I winsorized my sample to force a common distributional support), and time

---

[13]Tobin (1969) defined a measure, Q, as the ratio between market and replacement value of assets. When taken as a left-hand side variable there is often an implicit assumption that the value of the firm is a function of the sum of the physical and intangible assets of the firm: $V = f(A+K)$. This allows empiricists to estimate $\log Q = \log(\frac{V}{A})$, as the contribution of investment in intangible assets to the value of the firm. Griliches (1981) describes this econometric estimation in detail.

and industry effects. Accordingly, I use 'sector x year x acquisition-indicator' fixed effects.[14] Thus I estimate the average effect of citations on firm value within each combination of organizational form choice, liquidity event year, and primary sector of operation, which should alleviate concerns about the effects these variables could have either separately or together.

Likewise, I use 'state of incorporation' fixed effects because a material number of the start-ups in my sample are venture-capital financed (31%), and these firms are concentrated in California and Massachusetts. VC-backed firms in other states may experience a price discount or premium.[15]

I conduct my estimation, with these controls, in an OLS regression framework with my value measure as the dependent variable. The more sophisticated approach of estimating a production function brings associated costs. Hayashi (1982), and more recently other authors such as Miao (2005), emphasize the problem: Econometricians would like to estimate the contribution of production inputs on marginal Q, but do so using measures of average Q. Unfortunately, without specifying an underlying model of the structure of the firm, it is difficult to say when or whether average Q and marginal Q will coincide. If these measures don't coincide then the meaning of any estimates is lost. Using Hayashi's model of the structure of the firm leads to two eponymous 'Hayashi conditions' for when marginal Q will equal average Q: 1) When markets are perfectly competitive (i.e., the firm is price-taking); and 2) When production functions are homogenous of degree one. Both of these conditions are unlikely to be met in reality, and particularly in my sample of start-ups. Miao (2005) demonstrates that these conditions aside, agency costs will also lead to a difference between marginal and average Q.

Second, I rely solely on time-based fixed effects to address citation truncation issues. Hall et al. (2005) use a more sophisticated approach, where citations counts are adjusted to compensate for truncation. The consequence of my choice is that I cannot estimate a premium to a citation, instead I can only say whether on average within each year (measured at the time of the liquidity event) more citations are significantly associated with more value or not. As it turns out, my results will be consistent with the systematical rejection of the existence of a premium, so I incur no cost from this choice. The benefit from the choice is that I can be sure that I have not accidentally influenced my results through the use of a citation-weighting scheme.

Third, I do not include R&D expenditure, in part because I do not have sufficient data on R&D, particularly for acquisition targets.[16] The effect of R&D expenditure can be decomposed into three components: 1) a component associated with patent stock filed prior to the liquidity event; 2) a component associated with potential patent stock that may be filed subsequent to the liquidity event; and 3) a component that is orthogonal to patent

---

[14]In unreported regressions I control for the organizational form choice, liquidity event year, and sector, directly and separately. These controls do not affect my results.

[15]More generally, a startup may experience value effects from operating within a geographic cluster of related firms.

[16]Hall et al. (2005) note that there are problems in the reporting of R&D expenditure even for publicly-traded firms. Approximately one third of the firms in their sample do not report any R&D at all.

stock. When I estimate the effect of patents on value I will be measuring both the first and second components of R&D expenditure listed. This will lead to an overestimation of the contribution of patents to value, and my estimates are therefore biased upwards and should not be used to assign an average value for a patent. Fortunately this has no bearing on my research question.[17] I am interested in the effects of citations once I control for the size of a firm's patent portfolio. It is important for my analysis that patents are valuable (at least on their own without citations), but a statistically significant positive coefficient on the number of patents held by startups suffices to provide supporting evidence.

Aside from its simplicity, my estimation approach has one important benefit: it is how it is used in the literature.[18] The vast majority of applications take the log of the number of patents and the log of citations-received together, either as two variables or in the form of a single log of citation-weighted patents variable, and use it as a proxy for patent portfolio value in an OLS. Likewise the vast majority of applications do not adjust citations for truncation (instead they typical rely, as I do, on time-based controls) and do not include R&D expenditure effects. I estimate the patent portfolio value in exactly this way so that my results will be immediately applicable.

---

[17]Because my sample consists solely of firms that have one or more patents, the estimate would be biased upwards anyway.

[18]Prominent examples include Stuart (2000), Gompers et al. (2003), Dushnitsky and Lenox (2005) Hsu and Ziedonis (2007), and Haeussler et al. (2009).

Chapter 6

Results and Analysis

6.1 Commercialization Strategy

I begin my analysis with an exploration of the effects of citations-made on the commercialization strategy of startup firms. In table 6.2, I report the results of logit regressions where the dependent variable is binary and takes the value 1 if the firm experienced an acquisition and the value 0 if it experienced an initial public offering. As such, a negative coefficient on an explanatory variable (e.g. the log of the average of the count of citations-made by patents in the firm's portfolio) indicates that the measure is correlated with an initial public offering. In causal language, which is warranted if the technology of the firm is genuinely exogenously-given (it certainly precedes the liquidity event), a negative coefficient indicates that the variable in question makes it more likely that the firm will achieve an initial public offering.

In table 6.1, I provide results for the effect of the log of the average count of citations-made on the commercialization strategy of startup firms for three different samples and with various control variables. In specification 1, I report the results for my full sample of startups, which secured a liquidity event between 1986 and 2006, without including any control variables. For some observations in this sample, the value of the firm is not disclosed, and for others the value falls outside out the bounds that I required for my final (constrained) sample. The log of the average count of citations-made has a coefficient of -0.299, which is statistically significant with a p-value of less than 0.01.

In specification 2, I report the coefficient for citations-made for my full sample when portfolio time-distance, sector-year fixed-effects, and modal patent-category fixed-effects are included. The coefficient decreases by about 10% but is essentially robust to the inclusion of these controls. In specification 3, I include log firm-value and log firm-value-squared controls. In order to include these controls, I have to estimate the coefficients using my 'value sample', where the firm value for each observation is disclosed. The coefficient increases by on the order of 20%, but again is very similar to before. In the section that discusses the analysis of the value of startup firms, I will find that none of my citation-based measures predict value,

Table 6.1: The relationship between organizational form and citations-made

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. Samples used include my full sample, my sample where value is specified for every observation, and my final (constrained) sample with common value supports for acquisitions and IPOs. Potential control variables include 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 (Full Sample)	Specification 2 (Full Sample)	Specification 3 (Value Sample)	Specification 4 (Constrained Sample)
Log No Patents	-0.344*** (0.036)	-0.561*** (0.054)	-0.034 (0.082)	-0.022 (0.086)
Log Avg. Cites Made	-0.229*** (0.038)	-0.202*** (0.072)	-0.241** (0.103)	-0.243** (0.101)
Log Firm Value	-	-	-1.592*** (0.236)	-1.808*** (0.466)
Log Firm Value Squared	-	-	0.052** (0.026)	0.074 (0.049)
Portfolio Time Dist.	no	yes	yes	yes
Sector x Year F.E.	no	yes	yes	yes
Modal Pat. Cat. F.E.	no	yes	yes	yes
Constant	2.066*** (0.101)	0.032 (0.419)	2.753*** (1.023)	3.777*** (1.207)
R-Squared	0.0283	0.272	0.4171	0.3446
N	4731	4030	1980	1509

and so I feel assured that these controls are truly exogenous to my explanatory variables.

In specification 4, I report the coefficient using the same controls but in my final sample, where the values of startups are constrained to lie between $10m and $2.5b. Despite the winsorization of my final sample to ensure a common support on firm value across initial public offerings and acquisitions, there is a material and significant difference in the mean value for these two commercialization-strategy outcomes: an unequal, unpaired t-test between the two samples gives initial public offerings a greater mean by U.S.$137m, which is significant with a p-value of less than 0.0001. Between the supports there are distributional differences in the value of the firms. To mitigate concerns, I report results based on my final sample with both a linear and a higher-order value control. I experimented with other controls, including quantile-based firm value fixed effects, but none of these controls materially altered my results. The change in the coefficient between specifications 3 and 4 that arises as a result of imposing this common value-support is so small that the coefficients are essentially identically.

Therefore, through out the rest of this section, I use linear and quadratic firm value controls, 'sector x year' and modal patent-category fixed-effects, and estimate all coefficients using my final sample. In table 6.1, specification 4, I find a negative correlation between citations-made and acquisitions. The number of patents is apparently unimportant in determining the commercialization strategy of a patent-holding startup, but the nature of its technology matters. I have argued that a positive coefficient would be associated with the startup having a balance of patent-based relationships such that the effect of citing-complements in the context of a patent thicket is of greatest importance, and a negative

coefficient should be associated with a dominant effect for citing-substitutes. My findings therefore support a dominant effect for citing-substitutes for startup patents. In table 6.2, I take the first step towards showing that patent citations can convey this type of economic information.

Table 6.2: 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

Citations that are out-of-term cannot have direct economic consequences; the right to exclude that was granted to the cited patent has expired at the time that the citing patent is filed. Out-of-term citations can, of course, convey knowledge flows, or otherwise have indirect consequences. In specification 1, I show that in-term citations have the opposite effect of out-of-term citations with respect to an outcome variable (the type of liquidity event the startup experiences). I treat this an empirical proof by contrapositive: no right-to-exclude implies a negative correlation, the right-to-exclude implies a positive correlation, and both correlations are highly statistically significant.[1] My evidence is consistent with

---

[1]I did not advance a hypothesis for the effect of knowledge flows on the organizational form choice of startups. Why out-of-term citations are correlated with an increased likelihood of an acquisition is beyond the scope of this thesis. To address the question of whether economic information pertaining to the right-to-exclude matters in patent citations, I need in-term and out-of-term citations to have statistically different effects. I would have sufficient evidence of this if out-of-term citations had no effect or were correlated with the likelihood of an initial public offering, providing that the difference in effects was still statistically significant (and could not be explained by the age of the citations). The finding that out-of-term citations

in-term citations representing when the right to exclude matters, and conveying a direct economic meaning.

Furthermore, in-term citations are correlated with my outcome in the same way as the overall citations-made count (reported in table 6.1, specification 4, as -0.243^{**} (0.101).) . Out-of-term citations act to reduce the size of the overall citation count effect, but the in-term citation effect dominates. This suggests that the information in in-term citations is of greater consequence than that in out-of-term citations, and so that the direct economic information in patent citations is of primary importance.[2] A logit regression estimates the log-odds ratio as the regression coefficient.[3] A coefficient of -0.25 is economically meaningful, representing an odds ratio of 0.78, which in turn implies that an additional log-citation (or 2.718 citations) would be associated with an expected change in the odds of an acquisition by about $\frac{3}{4}$, relative to that of an initial public offering. Regardless of their relative importance, both the effects of in-term and out-of-term citations-made are economically meaningful and merit attention.

One possible omitted covariate that could undermine this finding is the age of the citation. Out-of-term citations are necessarily older, being cited after the 17 or 20 year expiration date has elapsed. With a change of sign between the effects of in-term (younger) and out-of-term (older) citations, any age effect must be at least quadratic. In specification 2, I show that the in-term and out-of-term citation effects are largely robust to the inclusion of a quadratic age-effect measure, maintaining about 90% of their magnitude, though the out-of-citation effect loses some statistical significance. In specifications 3 and 4, I check that age effects are not driving my results in another way; I try using a measure of the fraction of citations that are in-term along with a quadratic age term. The result supports a greater fraction of in-term citations being negatively correlated with the outcome, and only a very small effect attributable to the square-age of the citations.

If in-term and out-of-term citations are not considered, and an average citation-age variable is used as the sole explanatory variable (i.e., aside from control variables), I find that linear citation-age has a coefficient of -0.039 (0.032) and is not a statistically significant predictor of the commercialization strategy. When in-term and out-of-term citation measures are also included to take into account term-effects, the coefficient on citation-age becomes positive and weakly statistically significant at 0.047^{*} (0.025). Thus, while this evidence is extremely weak, it would appear that if there is any relationship between the age of citations, term-effects aside, and the commercialization strategy of the firm, then it is that firms that make citation to older technologies are more likely to be acquired.

In table 6.3, I explore the importance of ownership in citations. In specification 1, I report the effects of all citations-made decomposed into the effect attributable to the count of owners-cited and the effect attributable to the component of the count of citations-made

---

do matter is important. I leave it to further research to explain why.

[2]I will address the correlation between in-term and out-of-term citations, which may partly capture the extent to which in-term citations also convey knowledge flows, shortly.

[3]The odds ratio is given by $\frac{p(Acq)}{p(IPO)} = e^{\beta_0 + \beta_1(cites) + \cdots}$.

that is orthogonal to the count of owners-cited. The ownership effect is strongly statistically significant and approximately equivalent in magnitude to the citations-made effect reported in table 6.2, specification 1 (i.e., the effect before the decomposition). The component of the overall citations-made count that does not convey ownership information is not statistically significant. In specification 2, table 6.3, I perform the same decomposition for in-term citations, with the same basic result. All of the meaningful variation in the count of in-term citations-made is explained by ownership. In specification 3, I repeat this analysis for out-of-term citations. At a first glance, the result appears to support the importance of ownership in out-of-term citations-made as well. However, this specification neglects an important consideration: in-term citations and out-of-term citations are correlated.

I find a correlation of 0.47, significant with a p-value less than 0.001, between in-term and out-of-term citations-made in my sample of successful startups. Therefore when a firm's patent portfolio makes a lot of in-term citations, which I have found are explained by in-term ownership, it will also tend to make a lot of out-of-term citations, which for mechanical reasons will also be explained by in-term ownership. This gives rise to the highly statistically significant negative correlation for the out-of-term ownership measure in specification 3. In specification 4, I take this correlation into account and show that ownership is actually unimportant for out-of-term citations. To do this I use a two-factor first-stage regression: I regress the in-term count of owners-cited and the out-of-term citation-made count component that is orthogonal to ownership (i.e., the component used in specification 3) on the count of all citations-made (i.e., both in-term and out-of-term). I find that the meaningful variation in citations-made is explained by a combination of in-term ownership and out-of-term citation counts (that are independent of out-of-term ownership). The remainder of the citations-made measure, reported as '(Cites-Owners), 2 factor' in table 6.2, is noise. Furthermore, the coefficient on in-term owners-cited has increased by approximately 5% from specification 2 to specification 4, as a result of the removal the component of in-term owners-cited that is correlated with out-of-term citations-made (and so knowledge flows). This suggests that knowledge-flows play only a very minor role in in-term citations, at least with respect to the effects of in-term citations in predicting the organization form of startups.

In unreported regressions I also explored the use of the (unbiased) fragmentation of ownership measures, used in Noel et al. (2006) and Cockburn and MacGarvie (2009), and put forward by Ziedonis (2004). This variable measures only the dispersion of ownership, as distinct from the count of ownership. I found that this measure had no predictive power: it did not explain any of the meaningful variation in my ownership variables, or contribute any new information to my models of startup firms' commercialization strategies. I therefore infer that it is the count of ownership that matters and not its dispersion, which is consistent with my theoretical arguments regarding ownership.

I have argued that, when firms exercise their right to exclude to force bargaining or to participate in markets characterized by certain substitute relationships, ownership should be paramount. I take the results in table 6.2 as strong evidence that ownership is important in patent citations and that in-term citations do indeed convey economic information concerning the usage of patent-holders rights to exclude. I now refine my understanding

Table 6.3: Exploring the importance of ownership in citations-made

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 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. The residual measures are calculated using OLS. The measure (Cites-Owners) is the residuals from regressing the log of the number of owners-made on the log of the number of citations-made. The 'in-term' and 'out-of-term' measures are the residuals from the log of the number of owners-made in-term on the log of the numbers of citations-made in-term, and likewise for out-of-term, respectively. The '2 Factor' measure is the residuals from a regression of the log of the number of owners-made in-term and the log of the number of citations-made out-of-term on the log of the number of citations-made irrespective of whether these citations are in or out of the term window. ***, **, 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	-0.023 (0.086)	-0.021 (0.086)	-0.039 (0.085)	-0.035 (0.086)
Log Cites Made Measures:				
Owners	-0.263** (0.114)			
Owners in-term		-0.292** (0.114)		-0.315*** (0.190)
Owners out-of-term			-0.304*** (0.116)	
Residual Measures:				
(Cites-Owners)	-0.201 (0.300)			
(Cites-Owners), in-term		-0.206 (0.305)		
(Cites-Owners), out-of-term			0.364* (0.190)	0.354* (0.190)
(Cites-Owners), 2 Factor				-0.204 (0.319)
Log Firm Value	-1.808*** (0.466)	-1.806*** (0.466)	-1.824*** (0.464)	-1.822*** (0.465)
Log Firm Value Squared	0.074 (0.049)	0.074 (0.049)	0.075 (0.049)	0.075 (0.049)
Portfolio Time Dist.	yes	yes	yes	yes
Sector x Year F.E.	yes	yes	yes	yes
Modal Pat. Cat. F.E.	yes	yes	yes	yes
Constant	3.729*** (1.208)	3.765*** (1.208)	3.846*** (1.205)	3.894*** (1.211)
R-Squared	0.345	0.345	0.346	0.347
N	1509	1509	1509	1509

of the composition of a startup's portfolio, in terms of citing complement and citing substitute patents, by examining how the importance of ownership varies with the degree of competition between cited patent-holders.

If a firm has had and will have no competitive relationship, whether through product-markets or licensing-markets, with another firm then, because substitution effects take place through the medium of competition, citing the other firm to convey substitution can be irrelevant. This is not the case with complements. If a firm needs a technology from another firm, whether or not the other firm is a competitor may be important in whether the other firm will provide the technology, but has no bearing on existence of the underlying need. Licensing-markets may facilitate out-of-sector substitution, but markets for technologies are still in their infancy and out-of-sector firms face information problems, particularly in identifying when substitution could take, or has taken, place. Therefore I expect out-of-sector

citations are more likely to convey complement information, and should consider in-sector citations if I want to maximize my chances of finding evidence of substitutes.

The descriptive statistics provided in table 4.2 showed that in-sector citation are extremely rare for startups that get acquired but reasonably common for firms that achieve an IPO: 96% of startups that get acquired do not make a single citation into their sector, but 53% of those that achieve a public offering make at least one. Overall, only about 25% of all firms make one or more citations in-sector. I argued in my commercialization-strategy theory section that IPOs should be associated with competence destroying substitution. This is most likely to occur in-sector. On the other hand, I also argued that non-competence destroying technologies might be suitable for acquisitions, as they would not compromise technology-specific assets but they would provide valuable learning opportunities. As such, an acquirer can purchase substitute technologies too, but is more likely to do so when the technologies replace inventions belonging to firms that are out-of-sector.

In table 6.4, I employ three measure of citations-made at different degrees of competition. The most 'competitive' citations are those to firms in the same primary sector of operation as the startup. Beyond these firms is the broader competitive environment of publicly-traded firms that have access to global capital markets.[4] I also consider citations-made to patents in the same patent-class-based category as the citing patent, as these may represent potential competitors in the licensing market.[5]

In specification 1, table 6.4, I report the results for citations-made to firms in the same sector of operation as the startup and to publicly-traded firms. Both measures have statistically significant negative coefficients which are materially larger than the coefficient for the count of all citations-made which was reported in table 6.2. Citations to firms in the same sector, or even to publicly-held firms, are comparatively rare but of greater importance. In unreported regressions I tried these measures separately. Without including citations to publicly-traded firms, the coefficient on citations-made to firms in the same sector was $-1.722^{***}(0.157)$. The reduction in this coefficient when citations to publicly-traded firms are included is attributable to the correlation between the two measures.[6] Patents with citations to firms within the same sector and to publicly-traded firms behave like citing-substitutes in my sample. Aside from the sign and significance of the coefficients, supporting evidence is provided in specification 2, where I show that all of this effect is driven by in-term citations.

In specification 3, I consider how ownership varies with the degree of competition of the citations. I provide results consistent with ownership explaining more of the meaningful variation in citation-based measures as the degree of competition lessens. For citations-made to firms in the same sector of operation, owners-cited explains some variation, but the component of the count of citations that is orthogonal to ownership is far more important, as

---

[4]I previously reported that publicly-traded firms account for almost two-thirds of the patents filed by all corporations in the United States at the USPTO, making them a reasonable next degree of competition in startup's innovation environment after firms that are in the same sector of operation as the startup.

[5]Again, I use the categorization of patent classes put forward by Hall et al. (2001).

[6]This correlation is at least partly mechanical, because I use the NAICS codes of publicly-traded firms, as well as both failed and successful startups, to determine whether a citation is in-sector.

72

Table 6.4: 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

judged both from the size of the coefficient and from its z-score. The z-score for owners-cited was −1.79, whereas for the citation-count component it was −11.18.[7] For citations-made to publicly-traded firms, ownership explains all of the meaningful variation. The signs of these coefficients are negative, indicating that the more in-term citations-made to firms in the same sector of operation as the startup, as well as to the firms in the broader competitive environment of other publicly-traded incumbents, the more likely it is that the startup will achieve an initial public offering. This is consistent with my theory that substitutes that embody racial innovation and destroy incumbent competencies are those that are best suited to initial public offerings. For in-sector citations, I find that ownership is comparatively unimportant; the count of in-sector citations-made is the strongest predictor of an IPO. I previously argued that counts of citations-made, beyond the count of owners-cited, should only be important with certain types of imperfect substitutes. For out-of-sector citations, I find that owners-cited explains all of the meaningful variation. I argued that for certain types of substitute relationships, particularly those that take place within licensing-markets, this should be true. Although ownership is also associated with bargaining and complements, the sign of the coefficients suggests that citing substitutes are overcoming any effect that

[7]Note that to make a comparison of the importance of coefficients using z-scores, the measures underlying the coefficients must be statistically independent. This is guaranteed by construction for my owners-cited and orthogonal-citation measures.

might arise from citing-complements in the context of patent-thickets.

In specifications 4 and 5, I consider citations to patents in the same patent-class-based category. Without the inclusion of other competitive citation measures, the coefficient on citations-made within the same patent category is negative and weakly significant, however, when my other competitive citation measures are included, the sign of the coefficient flips to positive and increases in significance. As such, I find: that in-category patent citations do contain useful information; that one component of this information is correlated with measures of sector-based competition; and that, absent other measures of sector-based competition, the sector-based competition component of in-category patent citations is dominant, at least with respect to a startup firms commercialization strategy.[8]

My theory suggests that the result for the component of citations-made in-category that is orthogonal to sector-based considerations has two possible interpretations. The first interpretation is that these citations represent complements, and patent-thickets act at the technology level, rather than the sector level. However, I found in my data that less than 1% of acquired firms make even a single citation to their acquirers. With very little technological overlap between acquirers and their targets, it seems unlikely that acquirers will be able to mitigate the problems associated with a thicket that operates at the technology level. The second interpretation is that these citations represent substitutes to technologies from outside of the firm's sector that have been repurposed to be useful inside of the immediate competitive environment. As these are not technologies that are destroying competencies within the sector, they are not radical innovations that will change the landscape. However, they are suitable technologies for an incumbent to learn from. In section 6.2, I will turn to an analysis of the impact of citations on firm value. I will find that citations do not predict firm value for either acquisitions or IPOs. I do not know the necessary proportion of substitutes to complements to find a null value effect, and I expect that a wide range of proportions may give this result. Nevertheless, a null value effect for acquisitions suggests that startups that are acquired hold technologies that are at least partly substitutes. As citations-made in-category are the only citations that I have found that appear to increase the likelihood of an acquisition, I favor the second interpretation.

Perhaps the most important result in table 6.4 comes from the 'goodness of fit' of my logit regressions when I include competitive citations. In the tables, I report the standard McFadden pseudo-R^2 of each model. The McKelvey & Zavoina's pseudo-R^2, which most closely approximates an OLS R^2, the Count, and the Adjusted Count pseudo-R^2s for specification 1 in table 6.4 were 0.805, 0.884, and 0.747, respectively. For my first-cut citation-made estimation from table 6.2, specification 1, these values were 0.558, 0.773, and 0.506, respectively. These values are almost the same as for a base model without citation information, which produced McKelvey & Zavoina, Count and Adjusted Count statistics of 0.541, 0.787 and 0.531, respectively.[9] Adding competitive citations has added substantial explanatory

---

[8]In unreported regressions I decomposed the effects of citations-made in-category, and found that the non-sector-based competitive effect was driven entirely by in-term ownership with a coefficient of 0.628***(0.232).

[9]Though I stress that adding citation information did add information, as evidenced by Wald test that

power to my model: The McFadden pseudo-R^2 almost doubles and most other measures of the goodness-of-fit have increased by around 50%.

A practitioner who is trying to predict the commercialization-strategy outcome for a startup would be aware of nearly all of my right-hand side variables (i.e., they know the industry classification, geography, and other characteristics of the startup, including its innovation history). If the practitioner can form accurate estimates of the relative number of IPOs and acquisitions in the next period then the Count pseudo-R^2 is the most appropriate measure. Otherwise, the Adjusted Count pseudo-R^2 might most closely represents a practitioner's best guess. Adding competitive citation considerations improves the Count pseudo-R^2 from around 80% to almost 90%, and improves the Adjusted Count pseudo-R^2 by almost 25% from just 50%. Thus adding competitive citation information might allow a practitioner to accurately predict whether a startup will achieve an IPO or an acquisition between $\frac{3}{4}$'s and $\frac{9}{10}$'s of the time, depending on his or her knowledge of the upcoming market for these events.

Wald tests provide an alternative approach to establish the importance of these measures. Wald tests are parametric, with an underlying assumption that the difference between the true parameter(s) and the proposed parameter(s) is normally distributed, and estimate whether a parameter or parameters are (simultaneously) different from a proposed value, which in my case will be zero. A highly-significant Wald test, using a zero proposed-effect value, conducted over several parameters at the same time suggests that the parameters in question are jointly adding information to the model.[10] A Wald test for both the citations-made in-sector and to publicly-traded firms measures in specification 1, table 6.4, gives a $\chi^2_{(2)}$ statistic of 207.97 with an associated p-value of $6.91e^{-46}$, and including the third measures of citations-made in-category, in specification 4, gives a $\chi^2_{(3)}$ statistic of 238.23 with an associated p-value of $2.30e^{-51}$. These p-values are zero to any reasonable approximation, re-affirming the notion that these measures have impressive explanatory power.[11] Considering each of these three measures separately, I find that citations-made in-category, my patent-class-based competition measure, is least important with a $\chi^2_{(1)}$ of 2.71, which has a p-value of just 0.100. The $\chi^2_{(1)}$ values for citations-made in-sector and to publicly-traded firms are 120.29 and 218.09, with p-values of $5.47e^{-28}$ and $2.36e^{-49}$, respectively.

Finally in this section, I examine the effect of citations-received. As I have already discussed, the count of citations-received is likely to be correlated with the count of citations-made for two reasons: both measures suffer from inflation over time, though year fixed-effects, particularly in conjunction with industry and patent-class-based fixed-effects, should at least

produced a $\chi^2_{(1)}$ statistic of 5.75, which is significant at the 0.05 level.

[10] A likelihood ratio test can also be used to test whether an effect is present in the data or not. However, likelihood ratio tests cannot be estimated using heteroskedastically robust standard errors. I anticipate considerable heteroskedasticity in my analysis because my measures are on very different scales and I employ many categorical variables. I therefore use Wald tests through-out.

[11] To put these *p*-values in context, it is a billion-billion times more likely that these measures convey information about a non-zero effect (assuming the model and test are correctly specified) than it is that a 50lb dog will successfully quantum-tunnel through a half-inch of aluminum (Orzel 2009).

partly address this concern; and in network formation models nodes exhibit homophily, so nodes with matching innate characteristics are drawn together and the frequency of the characteristic in the population will affect both the likelihood that a node is cited as well the its propensity to make citations. The nature of the technology embodied in a patent most likely provides the innate characteristics for the matching process.

A simple interpretation under the innovation-lineage view would be that some areas of technological innovation are dense and others are sparse, but this puts into question the construct validity of citations-received as representing value and citations-made as representing breadth (as distinct from value). In my framework of economic interactions the interpretation of this correlation is more natural. When a patent replaces a large number of inventions the area of technology is in flux and a dominant design is yet to be established, so the patent may be more likely to be replaced by a similarly large number of inventions. And likewise, when a patent has a large number of dependencies, the invention that it covers is inherently complex, the dominant design is stable, and further incremental innovation is likely to build upon it and use it.

I have argued that when an industry is in flux, incumbent firms have an incentive to acquire technologies that have already been replaced in order to learn from them. Likewise, a startup with a replaced technology is still experimenting. The firm's final technology may have suffer from a cospecialized-asset problem making independent operation (and an IPO) inefficient. On the other hand, if the industry is stable, firms are competing to sell complex products made up of a large number of complements, and a patent thicket has emerged, a firm whose technology is widely needed has a strong position and can remain independent. Therefore, I hypothesized that a startup receiving a large number of substitute citations should seek an acquisition, and one receiving a large number of complement citations in the presence of a patent thicket should seek an IPO.

Before I explore the effects of citations-received in table 6.5, it is important to note that citations-received are truncated in my sample at 2006. As I am using time-based fixed effects and estimating an average effect within years from 1986-2006, and as citations take time to accumulate, the average year will have less variation in citations-received than in citations-made. This may affect both the magnitude of coefficients and decrease the likelihood of finding a significant result.

In specification 1 of table 6.5, I find that citations-received are positively correlated with acquisitions, though statistically insignificant. The coefficient, at 0.131, is smaller in magnitude than that for citations-made, which I reported in table 6.2 as -0.243.

Specification 2 reports the results for citations-received when the measure of citations-made is also included. The measure for the count of citations-received is now statistically significant. This is consistent with the substitution effect of citations being dominant, and internally consistent with my result on citations-made (i.e., I argued that the best candidate for an IPO is a firm that replaces a lot of, particularly in-sector, technology, and then does not get its technology replaced). The coefficient on citations-made is approximately 10% larger than it was for the estimation of that measure alone (as in table 6.2). In specification 3, I show that this inflation in the effect on citations-made is attributable to the correlation

Table 6.5: The effect of citations-received

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 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. The residual measure is calculated using OLS: 'Resid. C.-Recd' is the residual from regressing the log of the number of citations-made on the log of the number of citations-received. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

	Specification 1	Specification 2	Specification 3
Log No Patents	-0.107 (0.085)	-0.070 (0.086)	-0.070 (0.086)
Log Avg. C.-Made		-0.272*** (0.104)	-0.233** (0.101)
Log Avg. C.-Recd	0.131 (0.101)	0.177* (0.093)	
Resid. C.-Recd			0.177* (0.093)
Log Firm Value	0.063** (0.029)	0.043** (0.022)	0.043** (0.022)
Log Firm Value Squared	-1.909*** (0.463)	-1.924*** (0.456)	-1.924*** (0.456)
Portfolio Time Dist.	0.087* (0.049)	0.087* (0.049)	0.087* (0.049)
Sector x Year F.E.	yes	yes	yes
Modal Pat. Cat. F.E.	yes	yes	yes
Constant	3.750*** (1.160)	4.192*** (1.163)	4.404*** (1.179)
R-Squared	0.344	0.348	0.348
N	1567	1567	1567

between citations-made and citations-received. The coefficient for the citations-received measure remains unchanged between specifications 2 and 3, and retains its significance, but the coefficient for citations-made is now reduced to -0.233, though with little change to its standard error.

Overall, my results in table 6.5 support my hypothesis that citations contain information about economic complementarities and substitution effects between inventions, and provide a consistent picture of startup patent-holders filing for some substitute patents. I suggest that citations-made and citations-received convey the same information – that of a mixture of complements and substitutes, with the substitute effects dominating – but that each measure is separately important.

The correlation between these measures is material. Using log measures the correlation is 0.18, significant at the 0.001 level, and this does materially affect coefficients in the estimation of the effect of citations on the commercialization strategy of the firm. I expect that this correlation will also affect other outcome measures and so caution authors accordingly.

6.2   Firm Value

Thus far I have provided evidence to support the notions that patent citations convey economic information and that startup firms in my sample are patenting a mixture of complements and substitutes, such that the effect of their substitutes on their commercialization-strategy dominates any effect from complementary effects arising in the context of a patent thicket.

I now turn to my exploration of the impact of patent citations on firm value. The innovation-lineage paradigm claims that the more citations a patent receives the more valuable it becomes. Yet the empirical evidence has not always born out this claim, creating a puzzle in the literature. I can resolve this puzzle as in my framework patent citations might represent substitution effects, which would potentially reduce the value of a patent with each citation-received, as well as complementary effects, which would lead to an increase in the value of the patent as more citations accrue.

In this section I present the results of regressing the count of citations-received on the value, for the successful startup firms in my sample, measured at the time of their liquidity events. My results are consistent with the only other paper in the literature that considered the effects of citations on firm value for startups: I find that patent citations are (on average) not valuable for startup firms.[12] I attribute this result to the presence of a sufficient number of cited substitutes in my sample to offset the effect of the cited complements. Thus, overall, I find evidence consistent with the startups in my sample patenting an indeterminate mix of substitute and complementary technologies. Taking the positive result for citations-received from publicly-traded firms from the literature as true, I therefore infer that the startups in my sample, which approaches the population of startups with patents, are producing more substitute technologies than publicly-traded firms. This is consistent with my previous results on commercialization strategy, and provides a coherent picture of startup firms that achieve IPOs as engaging in Schumpeterian creative destruction; putting an industry into flux by overturning the previous dominant design paradigm and starting a period of radical innovation from which a new dominant design will later emerge.

In table 6.6, I report four specifications. In the first specification, I report the effect of the log of the count of citations-*received* on the log of the value of the firm using my 'value sample' (i.e., my full sample where all observations have a disclosed firm value) and no controls. I begin by estimating the simplest possible model using OLS:

$$\log(Firm\ Value) = \alpha + \beta[\log(No.\ Patents)] + \gamma[\log(No.\ Cites\ Rec'd)] + \epsilon$$

The coefficient on the *number of patents* is significant and positive, suggesting that the more patents a startup has, the more likely it is to be acquired. However, the coefficient on *citations-received* is not statistically significantly different from zero. In specification 2, I repeat this analysis using my final sample (i.e., the sample where values are constrained to lie

---

[12]Again, Shane and Stuart (2002) previously reported this finding, but for a much smaller sample of startup firms.

between \$10m and \$2.5b, and so low-value acquisitions, which may represent fire-sales, and unrepresentative high-value IPOs, are not included). Again the coefficient is not statistically significantly different from zero.[13]

In specification 3, I include time-based controls. Time-based controls, including the portfolio time-distance, which measures how long, on average, has elapsed between the filing of a firm's patents and its liquidity event, and liquidity-event year fixed-effects, are likely to be crucial because I am not using normalized citations-received. A firm that files patents on average in, say, 1980, and achieves its liquidity event in, say 1990, is likely to have received a different number of citations than a firm that files patents on average in, say, 2004, and achieves its liquidity event in, say, 2005. This is because of the differences in the duration in which firms can receive citations before their liquidity events, and because my data on patent citations is truncated at 2006. Therefore, of all of my controls, time-based controls are likely to be the most important when estimating the effects of citations-received. However, again, I find that the average count of citations-received has no statistically significant effect on value.

In specification 4, I attempt to find an effect for the average count of citations-received on value in my final sample by including a full set of controls. Many characteristics of both the firm and its liquidity event may predict firm value. My hope is that by using a very aggressive set of controls, I can remove all of the variation associated with other observable characteristics and so find an effect for my measures of the nature of the technological relationships of the firm. However, yet again, I find that the average count of citations-received has no statistically significant effect on value.

I therefore stress that my results (or lack thereof) are not driven by my controls. Throughout this section I report the results using the most aggressive control set I have available, partly as this provided the only statistically significant result in the section (mentioned below), and partly so that my value analysis is consistent with my strategy analysis.[14]

In table 6.7, I consider the effect of the log of the count of citations-received on firm value for various sub-samples of my final sample. In specification 1, I consider whether citations-received predict value for IPOs, and, in specification 2, I consider whether citations-received predict value for acquisitions. Neither coefficient is statistically significantly different from zero, and the difference is coefficients is not statistically significant either. Thus, based on the reasoning in my framework, successful startups create technologies that are at least partly substitutes, whether they ultimately experience either an acquisition or an IPO. There may

---

[13]Using parametric tests other than OLS, and non-parametric tests, did occasionally yield significant results even with full controls. In particular, using a quantile regression gave a strongly statistically significant negative result for citations-received in my sample - a result that is consistent with substitution value effects being more important on average.

[14]In total, I have 430 indicators from 'sector x year x acquisition' indicators, 51 state indicators, including one for protectorates, and six patent category indicators to make a total of 489 fixed effect indicators. Thus inclusion of my control variables reduces my degrees of freedom and statistical power. However, many of my controls are ultimately omitted because of collinearity, and by removing 'unwanted' variation I can potentially increase my chances of finding a result.

Table 6.6: Using citation counts as a proxy for 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. Samples used include my value sample (i.e., all observations in my full sample where value is declared) and my final (constrained) sample which places a common value support across acquisitions and IPOs by requiring firm values to fall between \$10m and \$2.5b. Estimates potentially 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 (Value Sample)	Specification 2 (Constrained Sample)	Specification 3 (Constrained Sample)	Specification 4 (Constrained Sample)
Log No. Patents	0.430*** (0.038)	0.258*** (0.033)	0.269*** (0.033)	0.235*** (0.040)
Log Avg. Cites Recd	-0.028 (0.037)	0.015 (0.030)	-0.041 (0.032)	-0.017 (0.044)
Portfolio Time Dist.	no	no	yes	yes
Year Fixed Effects	no	no	yes	(implicit)
Sector x Year x Acq F.E.	no	no	no	yes
State F.E.	no	no	no	yes
Modal Pat. Cat. F.E.	no	no	no	yes
Constant	3.230*** (0.084)	3.879*** (0.070)	3.626*** (0.141)	3.353*** (0.314)
R-Squared	0.0565	0.0421	0.0842	0.5056
N	2361	1871	1871	1677

be a difference of degree between IPOs and acquisitions in terms of their substitute effects, but this is too fine a distinction for me to detect. However, the difference may be one of the type of substitution; I earlier argued that firms that IPO displace technologies in their own sector, whereas firms that are acquired may displace technologies in the same patent-class but that have historically been used outside of their own sector.

Following the findings of Hall et al. (2005), I paid careful attention to self-citations, the possibility of non-linear effects in citation-counts, and 'unexpected' citations that arrived after the liquidity event. When I considered only self-citations, that is citations to patents within the same portfolio, the coefficient was -0.178 (0.11) for my full sample and -0.176 (0.133) for my sub-sample with two or more patents, reported in specification 3. My sub-sample with two or more patents is arguably the correct one for this estimation, as self-citation is only possible for firms with more than one patent.

When I employed binary indicators which took the value one if the firm made or received citations, and zero otherwise, as well as quantile-based measures, that I hoped would pick up the effects for the tails of the distribution, I likewise found that all of these measures produced no significant effect. In specification 4, I report the results of using indicator variables for whether the count of citations-received fell into the 1st quartile, 2nd quartile, 3rd quartile, or the 90th percentile of the citation-received distribution. None of these indicators were statistically significant. Likewise, I tried limiting my analysis to only venture-capital-backed firms, or firms with a single patent in their portfolio, again with no result. Within the VC-backed subsample the coefficient on *patents* was somewhat smaller at 0.166* (0.087), but the coefficient on citations-received was robustly zero at 0.025 (0.090). For the subsample of firms

Table 6.7: Predicting value in important sub-samples

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. Samples are based on my final sample and include IPOs only, acquisitions only, and the sub-sample of patent-holders with 2 or more patents 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
	IPOs	Acquisitions	>= 2 Patents Sample	Final Sample
Log No. Patents	0.240*** (0.054)	0.228*** (0.058)	0.265*** (0.061)	0.236*** (0.040)
Log Avg. Cites Recd	-0.094 (0.068)	0.029 (0.059)	-	-
Avg. Self Citations	-	-	-0.176 (0.133)	-
1st Quartile	-	-	-	0.037 (0.138)
2nd Quartile	-	-	-	0.123 (0.121)
3rd Quartile	-	-	-	-0.004 (0.118)
90th Percentile	-	-	-	0.018 (0.122)
Portfolio Time Dist.	yes	yes	yes	yes
Sector x Year x Acq F.E.	yes	yes	yes	yes
State F.E.	yes	yes	yes	yes
Modal Pat. Cat. F.E.	yes	yes	yes	yes
Constant	3.353*** (0.347)	0.851* (0.460)	3.452*** (0.380)	3.276*** (0.333)
R-Squared	0.52685394	0.34579602	0.55900198	0.50670975
N	708	969	1273	1677

with only one patent the coefficient on citations-received was robustly zero at -0.034 (0.026).

Furthermore, I repeated the full analysis in this section using citations-received after the liquidity event. For this analysis it was crucial to control for the organization form (i.e., IPO or acquisition) as citations to acquired firms dropped precipitously after the liquidity event and this fed back into the value. However, with this control in place, I again found no significant correlation with firm value for citations-received, or any derivative citation-based measures, for either my full sample, any sub-sample, or for any industry (discussed below). These citations are observable to the econometrician but not the investors valuing the firms, and I was therefore optimistic that if any patents would convey useful information about private values, then these would. I was therefore particularly surprised that they didn't.

In table 6.8, I report results using some of my derivative citation-based measures. Measures used include the log of the average number of citing-owners, of citing-owners in-term and in-sector, of citations-made, of owners-cited, and of owners-cited in-term and in-sector, reported in specifications 1 through 5, respectively. Only one of these measures is statistically significantly different from zero. The log of the number of citing-owners in-term and in-sector is significantly negatively correlated with firm value, at the $p \leq 0.1$ level.

Aside from my other derivative citations measures, which include citations in-sector, to

Table 6.8: 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

publicly-traded firms, in-category, and so forth, as well as their ownership-count variants, I also experimented with my full variety of derivative citation measures used in the literature.[15] And other than citation-based measures, I considered the number of claims made by the patents held by startups in my sample, which is often taken as a measure of the *scope* or *breadth* of the patent. None of these measures produced a effect that was statistically significantly different from zero.

As such, I do not believe that my 'result' for the log of the number of citing-owners in-term and in-sector should be taken seriously.[16] Nevertheless, if I were to put an interpretation on it, I would have to say that the more in-sector owners that make an in-term citation to a patent held by a startup in my final sample, the less that patent will be worth. This is consistent with a substitution effect and directly contradicts positive value findings reported elsewhere in the literature.

In one last attempt to find a statistically significant value effect, I now consider the

---

[15]The most common derivative citations measures used in the literature are called the 'generality' and 'originality' measures. These are Herfindahl-type measure on the dispersion of citations across patent classes and were created by Trajtenberg et al. (1997). Lerner (1994) employed a variation on these measures, using the dispersion across International Patent Classifications instead. Ziedonis (2004) popularized a Herfindahl-type measure for the dispersion across ownership. As I discussed before, Hall (2005) provides a bias correction factor for these measures. I tried all of these measures in their bias-corrected forms, though without an effect for citations-made or citations-received I generally anticipated that they would have no effect.

[16]The significance of results should be adjusted for the number of statistical tests undertaken. This is discussed in the sector-based analysis below. The citing-owners in-term in-sector result is significant at the 1-in-10 level. I undertook what can only be described as a major deep-sea trawling expedition, testing literally hundreds of variable combinations, to find this one result.

possibility that citations-received have different meanings in different sectors. In table 6.9, specification 1, I repeat the result from table 6.6, specification 4. Readers should note that, in order to keep the results together and fit the table on a single page, the coefficients for the number of patents and for the number of citations-received are presented side-by-side in table 6.9. In the second and third specifications, I report the results for citations-received for startups in different sectors. The second specification estimates the effects for startups in four high-level sectors: Information technology, life sciences, chemicals, industrial, and 'other'. The third specification estimates the effects for startups classified in to 'fine-grained' sectors. The definition of the sectors is provided in the appendix.

The sector-based breakdown is important in that it shows that citations-received universally confer no value effect, rather than adding value in some sectors and reducing it in others. Furthermore, this breakdown facilitates a comparison between my results and other work with an industry-specific focus; software, semiconductors, and biotechnology sectors are all popular fields for the application of patent-based measures.

The results in table 6.9 are fairly striking. Patents are positively and statistically significantly related to value but citations-received are not.[17] Between high-level sectors the coefficient on the *number of patents* shows surprisingly little variation, ranging from 0.193** (0.085) for the 'other' sector to 0.326*** (0.090) for industrials, but the differences are strictly insignificant. For the fine-grained sector definitions the results are slightly more interesting. There is considerably more variation in the coefficients which ranges from 0.055 (0.155) for computer media to 0.686*** (0.119) for printing & paper. This variation does manifest itself as some weak significant differences: within high-level sectors I find that telecoms are different from media, and both 'general industrials' and instruments are different from printing & paper. And treating the 'other' sector as a benchmark for a typical patent value, I find that only the printing & paper sector is weakly significantly different from this benchmark.

---

[17]As I am using a log-log regression, the coefficient of 0.235 equates to a value of U.S.\$3.8±1.3m for a patent held by a start-up in my sample, but I again caution this that estimate will be biased upwards. Hall et al. (2005) puts the value of patent before considering (valuable) citations at approximately U.S.\$1.8m for publicly-traded firms. Griliches (1998) reports an older finding for 1950s patents of approximately U.S.\$0.5m of private value per patent, in 1988 dollars. Google Inc. appears to have valued Motorola's 24,500 patents at about U.S.\$0.5m each (Economist 2011). However, the estimates of the mean patent value from other studies appear much larger. The sample used in Harhoff et al. (1999) reported a mean value (calculated from the mid-point of self-reported value ranges) of DM11.1m-19.1m, which translates to about U.S.\$7.4m-12m. And from Sampat (2004)'s regression coefficient of 1.541, I infer a mean patent value of approximately U.S.\$18m for their sample of university-based inventions.

Table 6.9: Exploring differences between sectors

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.

	Log No. Patents	Log Avg. Cites Rec'd	Log No. Patents	Log Avg. Cites Rec'd	Log No. Patents	Log Avg. Cites Rec'd
All	0.235*** (0.040)	-0.017 (0.044)				
Info. Technology			0.207*** (0.061)	0.026 (0.061)		
Computer Hardware					0.176 (0.197)	0.196 (0.187)
Semiconductor					0.170* (0.092)	0.071 (0.157)
Telecoms					0.055 (0.155)	-0.027 (0.168)
Computer Media					0.453*** (0.168)	0.177* (0.107)
Internet					0.500* (0.282)	0.009 (0.200)
Software					0.219* (0.118)	-0.171 (0.130)
Life Sciences			0.289*** (0.078)	0.090 (0.096)		
Gen. Life Sciences					0.356** (0.166)	-0.035 (0.181)
Biotechnology					0.265*** (0.086)	0.135 (0.112)
Chemicals			0.300** (0.146)	-0.523 (0.475)		
Chemicals					0.295** (0.147)	-0.519 (0.482)
Industrial			0.326*** (0.090)	-0.005 (0.137)		
Industrial					0.239 (0.188)	-0.096 (0.240)
Metal					0.224 (0.175)	0.100 (0.213)
Instruments					0.329 (0.236)	-0.211 (0.221)
Paper & Printing					0.686*** (0.119)	0.603* (0.361)
Other			0.193** (0.085)	-0.174* (0.105)		
Other					0.193*** (0.086)	-0.175 (0.107)
Portfolio Time Dist.	yes		yes		yes	
Sector x Year x Acq. F.E.	yes		yes		yes	
State F.E.	yes		yes		yes	
Modal Pat. Cat. F.E.	yes		yes		yes	
Constant	3.353*** (0.314)		3.469*** (0.345)		3.464*** (0.349)	
R-Squared	0.50562563		0.51418699		0.52737951	
N	1677		1677		1677	

However, I have good reasons to believe that the biases inherent in the patent coefficient will vary by sector. It is likely that R&D has different propensities to generate patents, and that there may be different selection for successful patents (discussed below) across sectors. Furthermore, the standard errors reported in the tables are not adjusted with a Bonferroni-type adjustment to compensate for the simultaneous testing of multiple hypotheses, so these results should be viewed with extreme caution. I am therefore unable to support either Levin et al. (1987), who found insignificant differences, or Hall et al. (2005), who found significant and material differences, on whether patents values vary across industries. However, I note that my sector-based breakdown of patenting industries closely resembles an updated version of that used by Levin et al. (1987), suggesting that certain industries do have a natural predisposition to use patent-based intellectual property protection.

The focus in this thesis, though, is on the contribution of citations after the patent effect is accounted for, and not of the effects of the patents themselves. For the high-level sectors, only the citations-received coefficient for the 'other' sector is even weakly significant, and the coefficient is negative. For the fine-grained sectors, only computer media and paper & printing are weakly statistically significant, and then with positive signs. Within high-level sectors there are significant differences only between computer media and software where the difference in coefficients is 0.348* (0.168), and instruments and paper & printing where the difference in coefficients is 0.814* (0.423). Again treating the 'other' sector as a benchmark, only the coefficient on printing & paper is statistically significantly different from this benchmark. Given the very weak significance of the results and the lack of a Bonferroni-type adjustment on the standard errors, I treat this as compelling evidence of a null result. There are 14 fine-grained sectors, so I would expect at least one of these sectors to show a statistically significant result with a p-value of 0.1 by chance alone.

Another way of examining whether citations matter in predicting value is by conducting a Wald test to consider whether the variables are collectively different from zero, and so add explanatory power to the analysis. As before I find that patents do matter (the R-squared with the fixed effects alone is 47% and patents materially and statistically significantly improve this to over 50%). However, again citations-received do not matter: for all sectors together citations-received had an $F_{(1)}$-statistic of 0.149 with a p-value of 0.699; for the course-grained sector-based citations-received measures the $F_{(5)}$-statistic was 1.015 with a p-value of 0.408; and for the fine-grained sectors, the 14 sector-based citations-received measures collectively produced an $F_{(14)}$-statistic of 1.09 with a p-value of 0.360. Clearly in none of my specifications are the citations-received measures even coming close to statistically significantly adding information to my models.[18] Likewise, citations-made, and all of my derivative citation-based measures, also provided no statistically significant results and

---

[18]As mentioned before, Likelihood Ratio (LR) tests are not strictly appropriate to test whether a model is being improved for my specifications, as my measures are heteroskedastic and a LR test will then tend to overstate the significance of the differences between a model with and without certain measures. However, an LR test also firmly rejected the citations-received measures as adding information in specifications 1 and 2. For specification 3, a LR test produced a $\chi^2_{(14)}$ of 21.38, which is very weakly significant with a p-value of 0.092.

added no information to a base model that included only the number of patents.

My results in this section are consistent with patent citation-based measures of any type being completely uninformative about firm value for any sample successful startup firms. But that is not to say that I believe that individual patent citations convey no information about the value of startup patents, just that in aggregate, no matter which dimension one looks down, the average startup patent citation produces no statistically discernible value effect. Given my results in the previous section, I believe that this result is fully consistent with aggregation masking underlying heterogeneity in patent citations. The null value effect of startup patent citations provides a third piece of evidence that startups patent a mixture of complement and substitute technologies, and a fourth piece of evidence consistent with citations conveying direct economic information. Compared with the positive citation-value result for publicly-traded manufacturing firms, such as those studied by Hall et al. (2005), my analysis suggests that startups produce relatively more substitute technologies than incumbents.

The citation-value literature has long struggled with two findings: the value effect of patent citations is extremely noisy, that is there is always a large standard-error in value estimations which makes it difficult to achieve statistical significance; and the positive value effects that have been found appear to be driven by the tail of the citation distribution. My sample is very large, approaching the population of startup firms that achieved an initial public offering or acquisition over a 30-year period, and as a result I have superb statistical power available to find a result even with noisy measures. Yet I don't. But the underlying heterogeneity that I have attempted to shed light on – that some citations represent information about complements and others about substitutes – explains why patent citations are noisy, and why, even with my available statistical power, I haven't found a result. This underlying heterogeneity can also explain why citation value effects are driven by the tail of citation count distributions. After a certain number of citations representing substitutes, a patent has been fully replaced and is no longer a relevant part of the technology landscape (i.e., there is no further incentive to cite it), but complement citations can accrue for ever. Thus when we observe a patent with a very large number of citations, it is more likely that it is a complement and so that each additional citation-received reflects positive value.

6.3   Determining Success

In this final section of my analysis I consider the possibility of a selection effect. Under the innovation-lineage view, citations-received reflect the importance and so value of a patent. I did not find a value effect for citations-received in my sample of successful startups. However, there is still the possibility that the patents held by the firms in my sample of successful startups are all extremely valuable, and so that there is insufficient heterogeneity in their values to find evidence of value in their citation patterns. To address this concern I consider whether citations-received can determine success – that is I suppose that the innovation-

lineage view could still be supported if success is correlated with value. Perhaps a startup needs sufficiently valuable technologies to achieve success, and the technologies of failed startups were not sufficiently valuable. Such a finding would be consistent with the results of Sampat (2004), who found that citations predicted whether a university-owned patent would be licensed but not the licensing revenues conditional on a license.

I do not advance a hypothesis for the effects of patent citations on success beyond this value hypothesis. As such, under the economic view of patent citations, this section of the analysis should be a repeat of the previous section, and can add new information only by undermining my previous findings.

Data on failed startups are generally very difficult to find, and currently essentially impossible to systematically assemble. However, I do have data on close to the population of failed venture-capital-backed startups. In table 6.10, I therefore examine whether citations-received are (positively) correlated with success in the context of venture-capital-backed firms that achieved either success through an acquisition or an IPO, or failed entirely. However, this sample may have generalizability issues. The presence of venture-capital may fundamentally alter both the technology and value of startup firms. For instance, venture capitalists may abandon firms that would enjoy minor success in order to refocus their resources on those that could achieve a major success.[19] Likewise, it is possible that some of the value of patents comes from mitigating information asymmetries with outside investors and VCs may use their reputations to certify their firms.[20] And VCs may actively select firms with specific technologies already under patent or in their R&D pipeline.

In table 6.10, specification 1, I show that the number of patents held by a startup at the time of success or failure is positively and statistically significantly correlated with success, but the count of citations-received is not. The patent selection effect appears to be driven by patenting in just the semiconductors and medical/heath (i.e., life sciences aside from biotechnology) sectors. For the other 12 sectors I find no statistical significance.[21] These two sectors have been the focus of much research in the literature, and it would be somewhat surprising if I found that the patents did not make a firm more valuable in them.

---

[19]Capital providers to non-VC-backed startups may do this as well, though perhaps not as aggressively. A naive ordering of the occurrence of refocusing by capital providers might range from venture-capitalists to angels to debt financing to friends/family/founders.

[20]The seminal reference for VCs using their reputations to certify firms at IPO is Megginson and Weiss (1991). Brander and Egan (2008) showed that VCs also certify their firms to acquirers, and Hsu and Ziedonis (2007) provided evidence consistent with patents providing a certification effect.

[21]Though my sample sizes are small, and the power of tests correspondingly low, particularly once fixed-effects have been taken into account.

Table 6.10: Using citation counts as a proxy for success

The dependent variable is a binary variable taking the value one if the venture-capital-backed firm succeeded (i.e., experienced an IPO or acquisition), and zero if it failed. 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 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, state of incorporation, and modal patent category fixed effects. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels, respectively.

	Log No. Patents	Log Avg. Cites Rec'd	Log No. Patents	Log Avg. Cites Rec'd	Log No. Patents	Log Avg. Cites Rec'd
All	0.264*** (0.102)	-0.128 (0.100)				
Info. Technology			0.226* (0.125)	-0.140 (0.116)		
Computer Hardware					0.406 (0.282)	-0.249 (0.289)
Semiconductors					0.641*** (0.240)	-0.744** (0.327)
Comms. and Media					0.023 (0.234)	-0.002 (0.226)
Internet Specific					-0.063 (0.380)	0.135 (0.232)
Software and Services					0.412 (0.267)	-0.258 (0.170)
Life Sciences			0.424* (0.228)	-0.200 (0.229)		
Medical/Health					0.627** (0.286)	-0.446 (0.303)
Biotechnology					0.140 (0.356)	-0.059 (0.422)
Other High Tech.			0.055 (0.399)	0.500 (0.399)		
Industrial/Energy					0.055 (0.539)	0.669 (0.517)
Consumer Related					0.484 (0.743)	-0.118 (0.768)
Other Products					-0.235 (0.560)	0.979 (0.888)
Portfolio Time Dist.	yes		yes		yes	
Sector x Year F.E.	yes		yes		yes	
State F.E.	yes		yes		yes	
Modal Pat. Cat. F.E.	yes		yes		yes	
Constant	-2.054 (1.422)		-1.607 (1.439)		-1.009 (1.392)	
R-Squared	0.16674991		0.17511453		0.19692188	
N	1273		1273		1273	

However, this thesis is concerned with patent citations, not patent counts. The results in table 6.10 do not support a selection hypothesis regarding patent citations: the number of citations a patent-holder's patents make is not correlated with the likelihood of the patent-holder's success. Even at the fine-grained sector level, where I consider the effect of citations in 14 sectors simultaneously, I find only a single statistically significant result. For the semi-conductor sector I report a statistically significant negative correlation between citations and success – exactly the opposite of what would be needed to support the selection hypothesis, at least under the innovation-lineage view of citations where more is better.[22] Again, though, the standard errors (and so significance) reported in table 6.10 do not correct for the testing of multiple simultaneous hypotheses. Doing so destroys this result: I find significance at the 1-in-20 level with 14 simultaneous tests.

Overall, the results in table 6.10 act to strengthen my claim that the innovation-lineage view of citations-received reflecting value is at best naive in general, and is not supported for startup firms. Specifically, more citations do not add more value to the patents of startup firms contingent on success, as shown in tables 6.6, 6.7, 6.8, and 6.9, and do not increase the likelihood of a startup firm achieving a valuable (successful) liquidity event, as shown in table 6.10. These results do not undermine the economic view of citations, providing success is truly a binary reflection of a continuous value measure.

---

[22]This result is all the more surprising, as a negative value coefficient on citations-received is consistent with substitution, and the semiconductor sector is most closely associated with complements and a patent thicket.

Chapter 7

Discussion and Conclusion

I have shown in my analysis of the commercialization strategy of startup firms that patent citations do convey information pertaining to the economic relationships between patent-holders. I did this first by considering the difference in the effect of in-term and out-of-term citations, where in-term citations were defined as citations-made to patents whose statutory term had not expired at the time that the citing patent was filed. Out-of-term citations are to patents whose statutory terms have expired, so that their patent-holders cannot use them to exert influence in economic relationships based upon their right-to-exclude. Out-of-term patents can still represent knowledge flows and other constructs that do not have a first-order economic impact. I showed that in-term citations have very different effects to out-of-term citations, and that the difference in their effects was not attributable to the age of the citation but rather to the expiration of the right-to-exclude. This acted as a contrapositive empirical proof that in-term citations do convey economic information.

I then considered whether ownership was important in in-term citations but not in out-of-term citations. A patent might cite or be cited by patents held by the same owners. Ownership, and not patent counts, should be important in some types of direct economic relationships, particularly those related to complements, which are based upon bargaining, or those relating to the interaction of certain economic substitutes. On the other hand, without a right-to-exclude, an expired patent has passed into the public domain and is essentially 'ownerless'. I found that on average the meaningful variation of in-term citations was completely explained by ownership but that ownership was unimportant for out-of-term citations. I took this as further supporting evidence that in-term citations convey direct economic information and that out-of-term citations do indeed represent knowledge flows.

I then posited that there was further underlying heterogeneity in in-term citations. That is, I posited that some in-term citations conveyed information concerning complementary economic relationships and others conveyed information concerning substitution-based economic effects. I found that my sample of startup firms, with one or more patents, that achieved a successful liquidity event (i.e., an initial public offering or an acquisition) between 1986 and 2006, had patents that exhibited some substitution effects. I found this in three ways: through the implications of the extent to which ownership explains meaningful variation in an

outcome measure, through a null value effect, and through the insights gained in an analysis of the effects of in-sector and in-category citations-made on the commercialization-strategy of the firms.

First, I argued that in only certain types imperfect substitute relationships should patent counts, and not the count of owners of patents involved in patent citations, be meaningful. When citations convey information about complementary relationships between patents, I argued that the ownership of cited or citing patents should always be paramount. This argument has its roots in bargaining theory. When a patent cites another patent-holder to indicate that it requires several of its patents, the bargaining for these patents takes place at the ownership level. A cited-owner of many complements effective has a single veto regarding the use of the new technology, and can either provide all of the necessary inputs or not. When citations convey substitution information, the argument is more nuanced.

All substitution relationships are mediated either through product-markets or licensing-markets. When a patent-holder cites another patent-holder's patents, whether those patents are themselves complements or substitutes, to convey replacement, the interaction between the patent-holders takes place in a competitive market. I argued that ownership information is paramount with: licensing-market competition, where the number of competitors in the market is key; horizontally-differentiated product markets when substitution, particular of other substitutes, can convey that a firm has gained access to a competitor's niche; and in the rare cases of perfect economic substitutes, where holding more than one perfect substitute confers no additional advantage, and replacing one of a competitor's perfect substitutes replaces them all. On the other hand, I noted that with imperfect substitutes in vertically-differentiated product markets, or in horizontally-differentiated product-markets where the substituted patents are either complements or do not protect the competitor's niche, each and every citation conveys information.

Furthermore, because the economic relationship between holders of two substitutes is mediated by competition, without competition or the threat of future competition, whether in product-markets or licensing-markets, there is no need to cite a firm holding a potential replacement technology. Thus I expected that citations conveying substitution information were more likely to be found in 'competitive' citations. One class of competitive citations is those to firms in the same primary sector of operation as the startup. I found that effect of in-term, in-sector patent citations was not fully explained by ownership, and instead that a measure of the count of patent-citations, that was orthogonal to the count of owners-cited, had a material and highly statistically significant effect. I took this as evidence of the presence of some substitutes. For in-term citations-made to publicly-traded firms, in-term citations to other patent-holders that were not likely to be engaged in product-market competition with a startup, and in-term citations to patents in the same technological class and so perhaps in licensing-market competition with a startup, ownership was paramount and the citation counts did not matter. The citations-made to firms in-sector are comparatively rare, so overall ownership explained all of the meaningful variation in the effect of in-term citations. This overall result could be consistent with either substitutes or complements. However, because there is no need to make a substitute citation to a firm that isn't in the same

product-market or licensing-market (and won't be in the future), I take the result that the effect of out-of-sector and out-of-category citations is driven exhaustively by ownership as evidence of some complements too.

Second, I found that for startup firms the count of citations-*received* is not statistically significantly correlated with value. This result in itself was not new to the literature, but I was able to explain it for the first time. I argued that if a citation-received indicates that the patent has been used as a complementary input into a later invention then each additional citation will increase the value of the cited patent, whereas if a citation-received indicates that the patent has now been substituted by a later invention then each additional citation will (potentially) decrease the value of the cited patent. Therefore if a patent portfolio consists of an indeterminate mix of cited complements and cited substitutes, I should anticipate a null result. This is what I found for every sample and sub-sample of my data, including samples of only startups that achieved IPOs and only startups that were acquired, as well as every citation-based measure (including those based on *citations-made*) used elsewhere in this thesis or tried in the literature.

Previous work that found a positive value effect for citations-received includes Hall et al. (2005) who studied publicly-traded manufacturing firms and Harhoff et al. (1999) who surveyed German inventors. Hall and Ziedonis (2001), Noel et al. (2006) and others associate at some sectors of publicly-traded firms with the presence of patent thickets, the mechanism for which relies upon a large proportion of the firm's patent portfolio being cited-complements. Harhoff et al. (1999) asked survey respondents whether they could identify their patents as complements or substitutes; positive responses came back at around 12:1 in favor of complements. My null result for value is therefore consistent with startups producing relatively more substitutes than firms in these other samples.

It is possible to find a null result for a citation value effect simply because citations are noisy measures of value. My sample of startups is very large: my final sample is the second largest for which this analysis has been conducted in the literature, and my full sample, which exhibits the same null result, has approximately twice the number of patent-holding firms as largest sample in the literature. I should, therefore, have sufficient statistical power to find a result even if measures are noisy. More importantly, my results suggest an explanation for why these measures are noisy. The countervailing effects of two heterogeneous types of citations, those pertaining to complements and those pertaining to substitutes, would make citations a noisy measure of value. Furthermore, substitute citations are only relevant until the cited patent is rendered obsolete. Beyond some count of substitute citations-received, there is no incentive for an outsider to make further citations. Therefore, when we observe a patent that has an extraordinarily high count of citations-received, it is likely that this patent is at least predominantly engaged in complementary relationships. This would explain why, when positive value findings have been found, the value effect of citations-received has been driven by the tail of the citation distribution.

Third, drawing from Teece (1986), I argued that if patents represent complements and patent-thickets pose material hurdles to startup firms, then the optimal strategy for firms that make a large number of citations, and so have a large exposure to the thicket, is to

seek an acquisition by an incumbent that could solve their hold-up problem. The patent thicket literature claims that the complex interdependence of technologies, owned by multiple independent and potentially competing patent-holders, has led to a situation where innovation is stifled.[1] The theoretical underpinnings for patent thickets could involve either patents acting as complementary inputs into other patents or the patent office (and courts) inadvertently allowing multiple patents on the same underlying technology. Either way, with complex products and exclusionary rights held by diverse parties, a hold-up problem can act as a barrier to entry for startups and discourage (particularly cumulative) innovation by incumbents.

Hold-up could take place at two levels: the product-level or the technology level. In the former case, I expected that firms that made a large number of in-term, in-sector citations would be the ones that would have the greatest preference for an acquisition by an in-sector incumbent. In the latter case, I expected that firms who made a large number of in-term, in-category citations would face the greatest hold-up problem. I also argued that if citations represent complements then citations-received could mitigate hold-up problems, as firms that receive a large number of citations conveying information that they have complementary inputs into their potential rivals, have a strong position in the thicket.

I also extended Teece (1986)'s framework to consider the effects of substitutes. I argued that incumbents face costs, beyond the erosion of their economic rents, from the acquisition of a radical technology that destroys their technology-base, that a new firm, seeking commercialization financing for the first time with an initial public offering, does not incur. Specifically, I claimed that an incumbent will lose rents from the substitution of its old technologies whether or not it acquires the startup, but if it does acquire the startup it will face additional costs in disposing of its old technology or attempting to commercialize the new technology alongside the old technology. Of course, the technology landscape is continuously changing, and incumbents may have incentives to purchase outside technologies; I argued that they will prefer to acquire technologies that do not disrupt their existing technology-base. I therefore posited that, if citations convey substitution information, startups that make a large number of citations in-term and in-sector would be more likely to hold radical inventions and be more suited to an initial public offering. Conversely, startups that cite technologies out-of-sector, perhaps particular technologies in-category but out-of-sector, would offer incumbents new technologies and knowledge, without disrupting their existing technology-base.

Furthermore, drawing from Teece (1986)'s notion of a pre-paradigmatic phase, where a dominant design is yet to emerge and the industry is in flux, and of a paradigmatic phase, characterized by a stable technology-base and incremental innovation, I argued that a startup whose technology has already been replaced by the time of its liquidity event is more likely to be acquired. With citations representing substitutes, a startup that receives a large number of citations is more likely to be attempting entry in the pre-paradigmatic phase. Without a stable dominant design, incumbents are not so concerned about replacement of

---

[1]Shapiro (2001), Ziedonis (2004), and others.

their technology-base, as it is more important that a technology supports learning than an immediate commercialization potential. Likewise, if a startup does not receive many citations, this would indicate that its technology has not been replaced and that instead it likely forms the basis for the dominant design.

I reported that in-term, in-sector citations appear to be very strongly correlated with initial public offerings. Less than 3% of firms that undergo acquisitions make even a single citation to firms in their sector, whereas 53% of firms that achieve an IPO make at least one citation into their sector before their liquidity event. My results therefore suggest that any effect from complement patents in conjunction with patent thickets operating at the product-market level is dwarfed by that from substitute patents and their creative destruction. Furthermore, I reported that in-term, out-of-sector, in-category citations appear strongly correlated with acquisitions. As I found that less than 1% of targets cite their future acquirers, bringing into question the presumption that the acquirers could solve a patent-thicket problem, I suggest that this finding is consistent with these citations also conveying substitution information.

Putting my findings together creates a coherent and consistent picture of startups as patenting some substitute technologies, and of startups that secure initial public offerings acting as a force for creative destruction. This has important bearings on two branches of the innovation literature: that on patent thickets, and that on the market for ideas.

I cannot say that patent thickets do not exist, just that they don't appear particularly important to startup firms, and don't appear to be acting as a barrier to entry. Furthermore, I can offer the hope that startups may be disentangling any thickets that are in place. The startups in my sample appear to be producing disruptive technologies that displace the patents that came before them. As such they may be acting to renew the technological environment, tearing down the established dominant design and conceivably any thicket that was associated with it, and creating a new basis for the next generation of development.

The 'market for ideas' literature claims that firms that cannot use patents, or other strong intellectual property rights, to protect their technologies will employ competitive strategies, rather than cooperative development strategies.[2] The core argument is that cooperation requires disclosure, which in turn requires strong protection, and a competitive strategy is a response to a market failure arising from the inability to disclose without forfeiting the idea. I find this argument compelling, but suggest that this might be a second-order concern.

My results are consistent with startup firms patenting some substitute technologies, and I find that patent citations indicating substitution effects between firms in the same sector of operation as the startup, where competition is fiercest, are very powerful determinants of whether or not a firm secures an initial public offering. I therefore provide evidence that a competitive strategy is enabled by strong intellectual property protection, not a response to its absence. My story is one of start-up innovation spurring the gale of creative destruction directly: a successful startup achieves an initial public offering if it displaces the greatest

---

[2]Gans and Stern (2003).

amount of pre-existing technology, and is not then displaced itself. Otherwise, it must settle for a cooperative strategy and seek an acquisition.

The empirical study of both innovation and entrepreneurship, including work on both patent thickets and the market for ideas, has relied heavily on patent-based measures, which are among the most visible and tractable signs of both invention and its commercialization. Patent citations, in particular, have been used prevalently; an entire literature, which I described in my literature review chapter, has developed and discussed them. Yet prior to this research, the meaning of a patent citation was still poorly understood.

I name the pre-existing paradigm the 'innovation-lineage' view of patent citations. Patents were taken to measure a unit of innovation, and patent citations provided a mapping of the lineage of innovation. In some papers this innovation-lineage was referred to synonymously as knowledge-flows. In other papers an innovation-lineage-based construct was extrapolated and the count of citations-received measured the technological importance of the patented invention. This led naturally to the view that more citations-received should correlate with greater value, as a patent with a large number of citations-received was either the source of a large body of knowledge or held a commanding position of technological importance. The findings that in many contexts citations-received did not correlate with private value therefore became a puzzle. This research has resolved this puzzle and is able to explain how, why and when a patent citation should indicate the accrual of more value or the subsequent loss of value. As such, I believe that this thesis makes a fundamental contribution to the patent-measures literature, and puts forward a new paradigm as a basis for further research.

My economic paradigm for patent citations suggests that researchers should not blindly count the number of citations-made and -received, but instead should be aware of the different types of underlying heterogeneity within patent citations. In-term citations convey economic meaning whereas out-of-term citations provide information about knowledge flows and the two can produce opposite effects, as they do in my sample with respect to the commercialization strategy of startup firms.

For in-term citations, overall and on average, ownership is paramount; I suggest that, particularly in samples dominated by complements, or if researchers are not able to make a distinction between in-sector and out-of sector citations, researchers should count owners rather than citations as the later only adds noise, and not information, to the former. And it is the count of owners that matters, not its dispersion. I found no support for the Herfindahl-type measure of the 'fragmentation of ownership'. It is unclear in my framework what 'fragmentation of ownership' should measure, whereas the count of ownership has a clear economic meaning as the number of relationships that are relevant in licensing agreements or other bargaining arrangements. As such, I advocate that research applying ownership measures, such as that in the patent-thicket literature purporting to measure the extent of hold-up and/or the magnitude of transaction costs, should follow a variant of the approach taken by Cockburn and MacGarvie (2006) and use the count of *in-term* owners.

However, future researchers should also be wary that not all in-term citations convey the same information. For in-term citations to direct competitors, each citation conveys information beyond its ownership. And some in-term citations appear to represent complementary

relationships, while others appear to represent substitution effects. This thesis provides indirect evidence of this distinction; further research should attempt to address it directly, perhaps by trying to identify when, where and how each citation conveys complement or substitute information in a small sample of carefully analyzed patents and their citations.

The potential power of my new paradigm is in its ability to understand the intricacies in the relationships that are inherent in the development, commercialization, and exchange of patent-based inventions. If we can know when a patent citation indicates that a patent is behaving as a complement or a substitute to another patent, we can develop a working understanding of the development of industries, the waves of creative destruction that befall them, the roles that various types of inventors and their firms play, how private value is captured, and how innovation changes the commercial world. With this research, I have taken a first step towards this understanding, by both carefully deconstructing the meaning of a patent citation and showing how this meaning can be used to understand the commercialization strategy, value and likelihood of success of startup firms. However, much future research is needed to confirm the foundations of my new paradigm, refine it further, and apply it broadly.

Finally, patent-measure paradigms aside, this thesis has brought the role of substitute technologies to the forefront in an analysis of commercialization strategies. A startup's choice to pursue an initial public offering or an acquisition is a prototypical example of a compete versus cooperate decision. Previous work on how the technology of a firm influences this decision, advanced by Teece (1986) and Gans and Stern (2003), considered what was needed for commercialization. I suggest, given the incredible explanatory-power of in-term, in-sector citations, that what has just been replaced is actually more important.

References

Abernathy, W.J. and Utterback, J.M. (1978). Patterns of industrial innovation. Technology Review, vol. 80 (7), pp. 40–47.

Acemoglu, D. (1997). Training and innovation in an imperfect labour market. The Review of Economic Studies, vol. 64 (3), pp. 445–464.

Agrawal, A.K. (2001). University-to-industry knowledge transfer: literature review and unanswered questions. International Journal of Management Reviews, vol. 3 (4), pp. 285–302.

Albert, M.B., Avery, D., Narin, F., and McAllister, P. (1991). Direct validation of citation counts as indicators of industrially important patents. Research Policy, vol. 20 (3), pp. 251–259.

Alcácer, J. and Gittelman, M. (2006). Patent citations as a measure of knowledge flows: The influence of examiner citations. The Review of Economics and Statistics, vol. 88 (4), pp. 774–779.

Arrow, K. (1958). Economic welfare and the allocation of resources for invention.

Asker, J., Farre-Mensa, J., and Ljungqvist, A. (2010). What do private firms look like. Unpublished working paper, New York University.

Audretsch, D.B. and Feldman, M.P. (2004). Knowledge spillovers and the geography of innovation. Handbook Of Regional And Urban Economics, vol. 4, pp. 2713–2739.

Becker, G.S. (1964). Human capital.

Bessen, J. (2008). The value of us patents by owner and patent characteristics. Research Policy, vol. 37 (5), pp. 932–945.

Bloom, N. and Van Reenen, J. (2002). Patents, real options and firm performance. The Economic Journal, vol. 112 (478), pp. C97–C116.

Blundell, R., Griffiths, R., and Van Reenen, J. (1999). Market share, market value and innovation in a panel of british manufacturing firms. Review of Economic studies, vol. 66 (3), pp. 529–554.

Brander, James A. and Egan, Edward J. (2008). The role of venture capitalists in acquisitions.

Campa, J.M. and Kedia, S. (2002). Explaining the diversification discount. The Journal of Finance, vol. 57 (4), pp. 1731–1762.

Chesbrough, H.W. (2003). Open innovation: The new imperative for creating and profiting from technology. Harvard Business Press.

Coase, R.H. (1937). The nature of the firm. Economica, vol. 4 (16), pp. 386–405.

Cochrane, J.H. (2005). The risk and return of venture capital. Journal of financial economics, vol. 75 (1), pp. 3–52.

Cockburn, I.M. and MacGarvie, M.J. (2006). Entry and patenting in the software industry. Tech. rep., National Bureau of Economic Research.

Cockburn, I.M. and MacGarvie, M.J. (2009). Patents, thickets and the financing of early-stage firms: Evidence from the software industry. Journal of Economics & Management Strategy, vol. 18 (3), pp. 729–773.

Cohen, W.M. and Levinthal, D.A. (1990). Absorptive capacity: a new perspective on learning and innovation. Administrative Science Quarterly, pp. 128–152.

Cotropia, Christopher A., Lemley, Mark, and Sampat, B. (2010). Do applicant patent citations matter? implications for the presumption of validity. Ph.D. thesis, School of Public Health, Columbia University.

Cumming, D. (2008). Contracts and exits in venture capital finance. Review of Financial Studies, vol. 21 (5), pp. 1947–1982.

Dixit, A.K. and Stiglitz, J.E. (1977). Monopolistic competition and optimum product diversity. The American Economic Review, vol. 67 (3), pp. 297–308.

Dushnitsky, G. and Lenox, M.J. (2005). When do incumbents learn from entrepreneurial ventures? corporate venture capital and investing firm innovation rates. Research Policy, vol. 34 (5), pp. 615–639.

Economist, The (2011). Valuing patents: Doing the maths.

Gans, J.S., Hsu, D.H., and Stern, S. (2000). When does start-up innovation spur the gale of creative destruction?

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.

Gay, C. and Le Bas, C. (2005). Uses without too many abuses of patent citations or the simple economics of patent citations as a measure of value and flows of knowledge. Economics of Innovation and New Technology, vol. 14 (5), pp. 333–338.

Geradin, D., Layne-Farrar, A., and Padilla Blanco, A. (2008). The complements problem within standard setting: assessing the evidence on royalty stacking. Boston University Journal of Science and Technology Law, Vol. 14, No. 2, 2008.

Gibbons, R. (2005). Four formal (izable) theories of the firm? Journal of Economic Behavior & Organization, vol. 58 (2), pp. 200–245.

Gompers, P., Lerner, J., and Scharfstein, D.S. (2003). Entrepreneurial spawning: Public corporations and the genesis of new ventures, 1986-1999.

Greenberg, G. (2010). Small firms, big patents? estimating patent value using data on israeli start-up firms financing rounds.

Greenhalgh, C. and Rogers, M. (2006). The value of innovation: The interaction of competition, r&d and ip. Research Policy, vol. 35 (4), pp. 562–580.

Griliches, Z. (1981). Market value, r&d, and patents. Economics Letters, vol. 7 (2), pp. 183–187.

Griliches, Z. (1998). Patent statistics as economic indicators: a survey.

Grindley, P.C. and Teece, D.J. (1997). Managing intellectual capital: Licensing and cross-: Licensing in semiconductors and electronics. California Management Review, vol. 39 (2), pp. 8–41.

Grossman, S.J. and Hart, O.D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. The Journal of Political Economy, pp. 691–719.

Haeussler, C., Harhoff, D., and Muller, E. (2009). To be financed or not - the role of patents for venture capital financing. ZEW-Centre for European Economic Research Discussion Paper No. 09-003.

Hall, B.H. (2005). A note on the bias in herfindahl-type measures based on count data. Revue D' Economie Industrielle, Paris Editions, Techniques Et Economiques, vol. 110, p. 149.

Hall, B.H. and Harhoff, D. (2004). Post-grant reviews in the us patent system-design choices and expected impact. Berkeley Tech., vol. 19, p. 989.

Hall, B.H., Jaffe, A.B., and Trajtenberg, M. (2000). Market value and patent citations: A first look.

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

Hall, B.H., Thoma, G., and Torrisi, S. (2007). The market value of patents and r&d: Evidence from european firms.

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

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.
URL: *http://www.jstor.org/stable/1593752*

Harford, J. (2005). What drives merger waves? Journal of Financial Economics, vol. 77 (3), pp. 529–560.

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.

Harhoff, D., Scherer, F.M., and Vopel, K. (2003a). Citations, family size, opposition and the value of patent rights. Research Policy, vol. 32 (8), pp. 1343–1363.

Harhoff, D., Scherer, F.M., and Vopel, K. (2003b). Exploring the tail of patented invention value distributions. Economics, Law and Intellectual Property, Netherlands: Kluwer Academic Publishers, pp. 279–309.

Hart, O. and Moore, J. (1990). Property rights and the nature of the firm. Journal Of Political Economy, pp. 1119–1158.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. Econometrica: Journal of the Econometric Society, pp. 213–224.

Head, K. and Ries, J. (1997). International mergers and welfare under decentralized competition policy. Canadian Journal of Economics, pp. 1104–1123.

Hegde, D. (2011). Asymmetric information and contracts in the market for ideas: evidence from the licensing of biomedical inventions. Unpublished working paper, New York University.

Hellmann, Thomas F., Egan, Edward J., and Brander, James E. (2005). Value creation in venture capital: A comparison of exit values across canadian provinces and us states.

Henderson, R., Jaffe, A.B., and Trajtenberg, M. (1998). Universities as a source of commercial technology: A detailed analysis of university patenting, 1965-1988. Review of Economics and Statistics, vol. 80 (1), pp. 119–127.

Henderson, R.M. and Clark, K.B. (1990). Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. Administrative Science Quarterly, pp. 9–30.

Holmstrom, B. and Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. JL Econ. & Org., vol. 7, p. 24.

Hotelling, H. (1929). Stability in competition. The Economic Journal, vol. 39 (153), pp. 41–57.

Hsu, D.H. and Ziedonis, R.H. (2007). Patents as quality signals for entrepreneurial ventures. Unpublished working paper, University of Michigan.

Jaffe, A.B. (1986). Technological opportunity and spillovers of r&d: Evidence from firms' patents, profits and market value.

Jaffe, A.B. and Lerner, J. (1999). Privatizing r&d: patent policy and the commercialization of national laboratory technologies. Tech. rep., National Bureau of Economic Research.

Jaffe, A.B. and Lerner, J. (2006). Innovation and its discontents.

Jaffe, A.B. and Trajtenberg, M. (1998). International knowledge flows: evidence from patent citations. Tech. rep., National Bureau of Economic Research.

Jaffe, A.B., Trajtenberg, M., and Fogarty, M.S. (2000). Knowledge spillovers and patent citations: Evidence from a survey of inventors. American Economic Review, vol. 90 (2), pp. 215–218.

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.

Kortum, S. and Lerner, J. (2000). Assessing the contribution of venture capital to innovation. RAND Journal of Economics, pp. 674–692.

Kuhn, T.S. (1970). The structure of scientific revolutions, 2nd edition. Chicago: Univ. of Chicago Pr.

Lang, L.H.P. and Stulz, R.M. (1994). Tobin's q, corporate diversification, and firm performance. Journal of Political Economy, pp. 1248–1280.

Lanjouw, J.O., Pakes, A., and Putnam, J. (1998). How to count patents and value intellectual property: The uses of patent renewal and application data. The Journal of Industrial Economics, vol. 46 (4), pp. 405–432.

Lanjouw, J.O. and Schankerman, M. (2004). Patent quality and research productivity: Measuring innovation with multiple indicators. The Economic Journal, vol. 114 (495), pp. 441–465.

Layne-Farrar, A. and Lerner, J. (2011). To join or not to join: examining patent pool participation and rent sharing rules. International Journal of Industrial Organization, vol. 29 (2), pp. 294–303.

Lemley, M.A. (2000). Rational ignorance at the patent office. Northwestern University Law Review, vol. 95, p. 1495.

Lemley, M.A. and Sampat, B. (2008). Is the patent office a rubber stamp. Emory Law Journal, vol. 58, p. 181.

Lemley, M.A. and Shapiro, C. (2005). Probabilistic patents. The Journal of Economic Perspectives, vol. 19 (2), pp. 75–98.

Lemley, M.A. and Shapiro, C. (2006). Patent holdup and royalty stacking. Texas Law Review, vol. 85, p. 1991.

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

Lerner, J. and Tirole, J. (2004). Efficient patent pools. The American Economic Review, vol. 94 (3).

Lerner, J. and Tirole, J. (2008). Public policy toward patent pools.

Leskovec, J., Kleinberg, J., and Faloutsos, C. (2005). Graphs over time: densification laws, shrinking diameters and possible explanations. In Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining. ACM, pp. 177–187.

Levin, R.C., Klevorick, A.K., Nelson, R.R., Winter, S.G., Gilbert, R., and Griliches, Z. (1987). Appropriating the returns from industrial research and development. Brookings papers on economic activity, vol. 1987 (3), pp. 783–831.

Mann, R.J. and Sager, T.W. (2007). Patents, venture capital, and software start-ups. Research Policy, vol. 36 (2), pp. 193–208.

Marco, A.C. (2007). The dynamics of patent citations. Economics Letters, vol. 94 (2), pp. 290–296.

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. Frontiers in Econometrics, pp. 105–142.

Megginson, W.L. and Weiss, K.A. (1991). Venture capitalist certification in initial public offerings. Journal of Finance, pp. 879–903.

Miao, J. (2005). Optimal capital structure and industry dynamics. The Journal of Finance, vol. 60 (6), pp. 2621–2659.

Nash Jr, J.F. (1950). The bargaining problem. Econometrica: Journal of the Econometric Society, pp. 155–162.

Noel, M., Schankerman, M., and for Economic Policy Research (Great Britain), Centre (2006). Strategic patenting and software innovation. Centre for Economic Policy Research.

Orzel, C. (2009). How to teach physics to your dog. Scribner Book Company.

Quillen Jr, C.D., Webster, O.H., and Eichmann, R. (2002). Continuing patent applications and performance of the us patent and trademark office. Federal Circuit Bar Journal, vol. 12, p. 35.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. Econometrica: Journal of the Econometric Society, pp. 97–109.

Salant, S.W., Switzer, S., and Reynolds, R.J. (1983). Losses from horizontal merger: the effects of an exogenous change in industry structure on cournot-nash equilibrium. Quarterly Journal of Economics, vol. 98 (2), p. 185.

Salop, S.C. (1979). Monopolistic competition with outside goods. The Bell Journal of Economics, pp. 141–156.

Sampat, B.N. (2004). Examining patent examination: an analysis of examiner and applicant generated prior art. Ph.D. thesis, University of Michigan.

Scherer, F.M., Harhoff, D., and Vopel, K. (1997). Exploring the tail of patented invention value distributions. ZEW Discussion Papers.

Schumpeter, J.A. (1934). The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle.

Schumpeter, J.A. (1942). Socialism, capitalism and democracy. Harper and Bros.

Scotchmer, S. (1999). On the optimality of the patent renewal system. The RAND Journal of Economics, pp. 181–196.

Shaked, A. and Sutton, J. (1982). Relaxing price competition through product differentiation. The Review of Economic Studies, pp. 3–13.

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

Shapiro, C. (2001). Navigating the patent thicket: Cross licenses, patent pools, and standard setting.

Shapley, L.S. (1953). A value for n-person games. Contributions to the Theory of Games, vol. II (Annals of Mathematical Studies), vol. 28, pp. 307–317.

Stuart, T.E. (2000). Interorganizational alliances and the performance of firms: a study of growth and innovation rates in a high-technology industry. Strategic Management Journal, vol. 21 (8), pp. 791–811.

Stuart, T.E. and Sorenson, O. (2003). Liquidity events and the geographic distribution of entrepreneurial activity. Administrative Science Quarterly, vol. 48 (2), pp. 175–201.

Teece, A., D.J. Pisano and Shuen, A. (1997). Dynamic capabilities and strategy management. Strategy Management Journal, vol. 18 (7).

Teece, D. and Pisano, G. (1994). The dynamic capabilities of firms: an introduction. Industrial And Corporate Change, vol. 3 (3), pp. 537–556.

Teece, D.J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. Research Policy, vol. 15 (6), pp. 285–305.

Teece, D.J. (2006). Reflections on 'profiting from innovation'. Research Policy, vol. 35 (8), pp. 1131–1146.

Teece, D.J. (2010). Business models, business strategy and innovation. Long Range Planning, vol. 34, pp. 172–194.

Teece, D.J. (2012). Management and governance of the business enterprise: Agency, contracting, and capabilities perspectives. In Mueller, Dennis C., ed., The Oxford Handbook of Capitalism, chap. 8. Oxford University Press.

Teece, D.J., Pisano, G., and Shuen, A. (2005). Dynamic capabilities and strategic management. Knowledge Management: Critical Perspectives On Business And Management, vol. 2 (7), p. 234.

Tobin, J. (1969). A general equilibrium approach to monetary theory. Journal Of Money, Credit And Banking, vol. 1 (1), pp. 15–29.

Toivanen, O., Stoneman, P., and Bosworth, D. (2002). Innovation and the market value of uk firms, 1989–1995. Oxford Bulletin of Economics and Statistics, vol. 64 (1), pp. 39–61.

Trajtenberg, M. (1990). A penny for your quotes: patent citations and the value of innovations. The RAND Journal of Economics, pp. 172–187.

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.

Veall, M.R. and Zimmermann, K.F. (1996). Pseudo-r^2 measures for some common limited dependent variable models. Journal of Economic Surveys, vol. 10 (3), pp. 241–259.

Wernerfelt, B. (1984). A resource-based view of the firm. Strategic Management Journal, vol. 5 (2), pp. 171–180.

Williamson, O.E. (1971). The vertical integration of production: market failure considerations. The American Economic Review, vol. 61 (2), pp. 112–123.

Williamson, O.E. (1979). Transaction-cost economics: the governance of contractual relations. Journal of Law and Economics, vol. 22 (2), pp. 233–261.

Williamson, O.E. (1983). Credible commitments: Using hostages to support exchange. The American Economic Review, vol. 73 (4), pp. 519–540.

Williamson, O.E. (1991). Comparative economic organization: The analysis of discrete structural alternatives. Administrative Science Quarterly, pp. 269–296.

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.

Appendix

I classify the NAICS codes of my startups into 14 industries, which I group into 5 high-level sectors: Information technology, biotechnology, chemicals, industrial, and 'other'. This classification is provided in table 9.1.

My industry classification is based upon that of Brander and Egan (2008), who created an assignment of 2002 NAICS codes to the information technology and life sciences sectors. I first updated their classification to include codes adding the 2007 NAICS listing, and then decomposed information technology into computer hardware, telecoms, computer media, internet, and software, and life sciences and biotechnology, respectively. This decomposition was guided by VentureXpert's almost identical decomposition and a matching of NAICS codes to VentureXpert's classification for my sample of venture-capital-backed successful startups, and by the counts of the observations in each sector in my data.

I sorted the remaining unclassified firms by their NAICS codes and attempted to identify additional coherent industries. The identification followed the following rule set: 1) If an entire 5-digit NAIC codes had over 40 observations, none of which were previously classified, and such that together it formed a coherent industry of operation, then extract it; 2) Repeat this for 4-digit, then 3-digit NAIC codes; 3) Aggregate any 3, 4 or 5-digit NAICS which together form a coherent industry of operation; 4) Add any 6-digit industry codes that could be unambiguously assigned to a pre-existing coherent industry of operation; and 5) check a randomly-drawn sample of firms' classifications against a description of their business taken from Thomson.[3]

The exception to this classification system was the instruments sector. I would have aggregated the entire of NAICS 33451 - "Navigational, Measuring, Electromedical, and Control Instruments Manufacturing" into the instrument sector, but code 334515 was already assigned to semiconductors, and codes 334510, and 334516, 33451617, 33451619, were already assigned to biotechnology. A careful review of the firms holding these codes indicated that these assignments were correct, and so the remainder was used to create the instruments sector, which was then supplemented with codes 333314 - "Optical Instrument and Lens Manufacturing" and 335314 - "Relay and Industrial Control Manufacturing".

In my analysis of whether citations influence success, I use VentureXpert's industry classification instead of my NAICS-based classification. This is done because I do not have

---

[3] 40 observations is a reasonable estimate of the threshold at which a t-distribution assumption becomes valid without strong symmetry, unimodality and with outliers.

NAICS codes for failed startups. To make the analysis comparable I assign VentureXpert's communications & media sector to my telecommunications sector, and VentureXpert's industrial/energy sector to my industrial sector. I carefully reviewed these assignments and while they are not perfect, more than three-quarters of firms that did have NAICS codes as well as a VentureXpert industry classification are correctly assigned, and I do not believe that the incorrect assignments have any bearing on my results.

Table 9.1: Classification of sectors by NAICS codes

The official NAICS 2007 code definition is included verbatim for all two, three and four digit codes. NAICS code 516 is taken from the 2002 NAICS definition, this code was removed from use in the NAICS 2007 definition but remains in the data. Thompson Financial marks some Internet firms with a proprietary code of "BBBBBB". The instruments sector is defined solely by six digit codes, but most codes from NAICS 33451 'Navigational, Measuring, Electromedical, and Control Instruments Manufacturing' are included. The information technology and life sciences codes are updated from Brander and Egan (2008). Holding company codes are included in the table but none of the startups in my sample had NAICS codes that would have assigned them into this group.

Sector	North American Industry Classification System (NAICS) Codes
Information Technology	
Computer Hardware	3341 (Computer and Peripheral Equipment Manufacturing)
Semiconductor	3344 (Semiconductor and Other Electronic Component Manufacturing); 42369; 333295; 333994; 334515; 335999
Telecoms	517 (Telecommunications); 3342 (Communications Equipment Manufacturing); 33592; and 561499
Computer Media	3346 (Manufacturing and Reproducing Magnetic and Optical Media)
Internet	BBBBBB (Internet - Thomson Specific); 516 (Internet Publishing and Broadcasting); 518 (Data Processing, Hosting, and Related Services); 4541 (Electronic Shopping and Mail-Order Houses); 51913; 51919; and 61142
Software	5112 (Software Publishers) and 5415 (Computer Systems Design and Related Services);
Life Sciences	
Life Sciences	3391 (Medical Equipment and Supplies Manufacturing); 42345; 42349; 42421; 62199; 446199; 541611 and 541922
Biotech	3254 (Pharmaceutical and Medicine Manufacturing); 5417 (Scientific Research and Development Services); 6215 (Medical and Diagnostic Laboratories); 54138; 236210; 325132; 334510; 334516; 334517; 334519 and 541690;
Chemical	
Chemical	326 (Plastics and Rubber Products Manufacturing); 3253 (Pesticide, Fertilizer, and Other Agricultural Chemical Manufacturing); 3255 (Paint, Coating, and Adhesive Manufacturing); 3256 (Soap, Cleaning Compound, and Toilet Preparation Manufacturing); 3259 (Other Chemical Product and Preparation Manufacturing); 32518; 32521; 32522; 33322; 42383; 42461; 42469; 325131; 325191 and 325199
Industrial	
Industrial	3334 (Ventilation, Heating, Air-Conditioning, and Commercial Refrigeration Equipment Manufacturing); 33312; 333298; 333313; 333319; 333911; 333912; 333913; 333921; 333922; 333923; 333924; 333991; 333992; 333993; 333995; 333996; 333997 and 333999
Metal	331 (Primary Metal Manufacturing); 332 (Fabricated Metal Product Manufacturing) and 3335 (Metalworking Machinery Manufacturing)
Instruments	333314; 334511; 334512; 334513; 334514; 334518; and 335314.
Paper & Printing	322 (Paper Manufacturing); 323 (Printing and Related Support Activities); 333291; 333293; and 333315
Holding Companies	
Holding Companies	55 (Management of Companies and Enterprises); 533 (Lessors of Nonfinancial Intangible Assets (except Copyrighted Works)); 5239 (Other Financial Investment Activities); 5259 (Other Investment Pools and Funds); and 5411 (Legal Services)

Table 9.2: Counts of firms and patents by sector

Columns provide counts of firms and patents by sector for the primary sample of successful startups, comprised of IPOs and acquisitions (broken out), a selection control group of failed VC backed startups, and the reference dataset of publicly traded firms from COMPUSTAT.

Sector	Acquisitions		IPOs		Failed Startups		Public Firms		Total	
	# Firms	# Patents	# Firms	# Patents	# Firms	# Patents	# Firms	# Patents	# Firms	# Patents
Computer Hardware	37	162	41	304	129	839	443	124,419	650	125,724
Semiconductor	87	638	87	1,356	194	1,272	726	122,353	1,094	125,619
Telecoms	74	355	64	420	*155*	*901*	548	64,920	841	66,596
Computer Media	108	389	57	197			14	586	179	1,172
Internet	32	62	24	67	105	313	78	688	239	1,130
Software	89	286	57	422	200	515	611	108,166	957	109,389
Life Sciences	53	523	71	613	167	1,078	432	14,273	723	16,487
Biotechnology	112	1,067	179	1,728	95	645	1429	79,067	1,815	82,507
Chemicals	53	529	24	935			553	74,049	789	76,417
Industrial	51	296	25	339	*159*	*904*	395	27,616	471	28,251
Metal	57	386	13	395			530	14,104	600	14,885
Instruments	33	246	11	99			250	20,699	294	21,044
Paper & Printing	30	453	9	122			241	37,525	280	38,100
Other	254	5,119	139	1,745	107	989	4039	364,664	4,539	372,517
Total	1,070	10,511	801	8,742	1,311	7,456	10,289	1,053,129	13,471	1,079,838