Robert M.
La Follette School of Public Affairs
at the University of Wisconsin-Madison

Working Paper Series
La Follette School Working Paper No. 2011-014
http://www.lafollette.wisc.edu/publications/workingpapers

Do important inventions benefit from knowledge originating in other technological domains?

Gregory F. Nemet
La Follette School of Public Affairs and Nelson Institute for Environmental Studies
at the University of Wisconsin-Madison

nemet@wisc.edu

Evan Johnson
La Follette School of Public Affairs at the University of Wisconsin-Madison

July 19, 2011

Robert M. La Follette
School of Public Affairs
UNIVERSITY OF WISCONSIN–MADISON

1225 Observatory Drive, Madison, Wisconsin 53706
608-262-3581 / www.lafollette.wisc.edu

Do important inventions benefit from knowledge originating in other technological domains?

Gregory F. Nemet[a,b,c,*], Evan Johnson[a]

[a] *La Follette School of Public Affairs, University of Wisconsin, 1225 Observatory Drive, Madison, WI 53706 USA*

[b] *Nelson Institute Center for Sustainability and the Global Environment, University of Wisconsin, Madison USA*

[c] *Harvard Kennedy School, Cambridge, MA, USA*

Abstract

A frequently made claim in the innovation literature is that important inventions involve the transfer of new knowledge from one technological domain to another. This study uses U.S. patents granted from 1976-2006 to identify the role of knowledge acquired from outside each patent's technological domain. Our results do not seem to support the claim above. Increasing citations to external prior art is a significantly less important predictor of forward citation frequency than citing prior art that is technologically closer. This result is robust across several model specifications and ways of defining whether each flow of knowledge is external. The result is even stronger in the most highly-cited technology categories. We discuss possible explanations for this apparently negative impact of external knowledge—including both measurement issues and challenges associated with assimilating disparate knowledge.

Keywords: invention, patents, citations, spillovers, knowledge flows

*Corresponding author phone: +1(608)265-3469, fax: +1(608)265-3233, email: nemet@wisc.edu.

1. Introduction

Important concepts in the innovation literature build on the notion that important inventions involve the transfer of knowledge from one technological domain to another. This observation derives from the perspective that new technologies are combinations of existing components and principles—and further, that exotic combinations produce the most novel inventions. Of those inventions that become commercialized, payoffs and importance may be especially high. Examples abound of important inventions involving combinations and transfer across technological domains (Mowery and Rosenberg, 1998; Ruttan, 2001; Arthur, 2007). Jet engines for military aircraft provided the fundamental technology for high efficiency natural gas power plants; advances in ball bearings and tires for bicycles enabled development of automobiles; production of long wires for radial tires was instrumental for slicing silicon wafers to produce solar panels. On a larger scale, general purpose technologies such as the steam engine, electric power, chemical engineering, and semi-conductors have had pervasive effects across multiple sectors of the economy. More specific technologies like lasers and synthetic fibers became useful for improving the performance of technologies far afield from their original area of application.

The possibility of substantial societal benefits resulting from novel combinations from disparate sources motivates work within technology policy, for example on the characteristics of knowledge networks, interdisciplinarity, and technology transfer. Firms too have been advised of the benefits of developing broad absorptive capacity, making connections across structural holes, and reaping economies of scope. While few would dispute that the potential for novel combinations is large, government decisions on science and technology inevitably face a trade off between how much to support science and technology conducted within existing domains versus supporting new endeavors that span domains. Corporations too wrestle with tradeoffs associated with improving routines versus initiating new ones; allocating between exploration and exploitation; and to what extent to pursue the benefits of concentration versus those of diversity.

This study aims to add insight on such decisions by using patent data to address the question: *can we observe benefits from assimilating knowledge from other technological domains?* Our data are information on the front pages of all 3 million U.S. patents issued between 1976 and 2006. We run negative binomial regressions on these data to identify the role of knowledge acquired from outside each patent's technological domain. The dependent variable in these regressions is each patent's importance. As a proxy, we use the number of times it is cited by subsequent patents within 10 years. We measure knowledge flows using the citations that each patent makes to previous patents, as well as the technological classifications for each. If a cited patent has a different technological classification from the citing patent, that patent-citation pair is counted as an external knowledge flow. We thus test the hypothesis that patents with more citations to external knowledge are more likely to be important, and thus receive more citations from future patents. We control for time, citation lags, patent breadth, the type of organization patenting, and technological areas with higher propensity for being cited. In order to improve the robustness of the results to patent examiners' decisions, we employ both the U.S. and international patent classification systems. In this paper, we first describe the theoretical issues at stake in understanding the importance of flows of knowledge and the use of patent data to measure them. Section 3 describes our approach to hypothesis testing. Section 4 presents the results and Section 5 discusses interpretation and implications.

2. Knowledge flows, inventions, and patents

The hypothesis we test in this study examines two longstanding concepts in the literature on innovation: first, the notion that novel technologies are assembled through a combinatorial process and second, that important inventions have often been characterized by the transfer of knowledge from one technical area to another. To test this hypothesis, we make use of previous work on the interpretation of information found in patent documents to construct indicators of inventions, invention value, boundaries of technical domains, and knowledge flows.

2.1. Combination, cumulative synthesis, and transfer

The notion of "cumulative synthesis" provides a starting point in characterizing the processes of discovering and developing new technologies. Usher (1954) proposed that technologies are formed through the *combination* of existing components and principles; a major or strategic invention represents the *cumulative synthesis* of many individual inventions (Ruttan, 2001). Schumpeter argued that the crucial role that entrepreneurs play in economic growth arises from their "special function" in "carrying out of new combinations" (Schumpeter, 1934). Nelson and Winter (1982) focus on the role that firms play in assembling combinations of technical, organizational, and market knowledge. Arthur (2007) used historical examples of the invention of 'radical'—that is non-incremental—new technologies, to develop a theory of invention that includes the claim that, "technologies are combinations; they are phenomenon based; and their architecture is recursive." An essential aspect of this argument is that knowledge creation is inherently cumulative; it resembles aspects of an evolutionary process (Gilfillan, 1935). This well-established notion of cumulative synthesis provides some of the basis for more recent work on the benefits of diversification within firms (Pavitt et al., 1989), technological transitions (Geels, 2002), and the theoretical structure of technologies (Arthur, 2009).

The perceived importance of cumulativeness and combination leads many to claim that novel combinations account for much of the historical evidence of particularly novel inventions. Inter-sectoral flows of knowledge are particularly important for these non-incremental inventions: "the most important inventions have had implications across industries" (Mowery and Rosenberg, 1998). More generally, "an important determinant of the rate and direction of scientific progress has been the transfer of concepts from one scientific specialty to another" (Rosenberg, 1994). And, "we know that novel technologies are shaped by social needs; that they issue often from experience gained outside the standard domain" (Arthur, 2007). Using increasingly disaggregated data on R&D, productivity, and patenting, Schmookler (1966) and Scherer (1982a,b) found that R&D investment in one industry had substantial effects on productivity growth in others.

Several studies highlight the role of the transfer of knowledge across boundaries within firms (Rosenkopf and Nerkar, 2001; Suzuki and Kodama, 2004), as well as in universities (Rosell and Agrawal, 2009). Others have linked these findings to the benefits of diversification in firms (Nelson, 1959; Jaffe, 1986; Garcia-Vega, 2006), and of relaxing disciplinary borders in academia (Rosenberg, 2009). Further results include: benefits from combining old and more recent technology (Nerkar, 2003); more variable outcomes when combining unfamiliar technology (Fleming, 2001); and breakthroughs arising from experimentation with combinations of unfamiliar, recent, and novel inventions (Ahuja and Lampert, 2001). Some firms actively exploit the diversity of their client base to combine disparate knowledge (Hargadon and Sutton, 1997). The concepts of technological space, technological domains, and technological distance have been essential for developing empirical tests of the value of combination (Gilsing et al., 2008; Benner and Waldfogel, 2008). As described below, delineation of technological domains is central for this study.

2.2. Interpretation of patent data

To operationalize and test these concepts, we make use of patent data and previous studies of what they represent. Patents provide an attractive way to measure inventive activity for several reasons: comprehensive data are publicly available, the technical characteristics are described in detail, the definition of what constitutes a patent in the U.S. has changed little for over 200 years, and every patent is categorized by experts using a standard classification scheme (Griliches, 1990; Watanabe et al., 2001; Jaffe and Trajtenberg, 2002; Hall et al., 2005; Popp, 2005). To be sure, using patents involves many well-documented limitations. For example, all patents are not equally important, not all inventions are patentable, firms use alternative means to protect their intellectual property, and sometimes they patent strategically (Harhoff et al., 1999; Bessen, 2005).

2.2.1. Forward patent citations as a measure of value

Claims about the role of inter-domain knowledge flows emerge from historical case studies of *important* technologies. Our ability to examine those claims against a broader set of inventions therefore requires distinguishing important inventions from others. We make use of previous work that shows that more frequently-cited patents tend to be more valuable. This literature finds a positive relationship using indicators of value based on: sales-based estimates of social value (Trajtenberg, 1990), stock market value of the assignee firm (Hall et al., 2005), interviews with inventors (Harhoff et al., 1999), payment of patent renewal fees (Griliches et al., 1987; Harhoff et al., 1999, 2003), whether a lawsuit was filed (Allison et al., 2004), and filling patents for the same invention in multiple countries (Lanjouw and Schankeman, 2004). Recent work has developed composite indicators using multiple measures (van Zeebroeck, 2011). Forward citations are a far from perfect measure of value, however: they still account for only a small part of the variation in value (Bessen, 2008); a full citation history takes decades to establish, introducing truncation issues (Lin et al., 2007); and they ignore future inventions that are not subsequently patented (Mariani, 2004). We acknowledge that our chosen measure of value is a “noisy” indicator (Jaffe et al., 1998) and make efforts to reduce noise by controlling for truncation issues and variables omitted in other studies.

2.2.2. Patent classes as technological domains

We use patent classifications to delineate technological domains. After a patent application is filed, the patent examiner responsible for that technical area assigns patents to at least one or more patent classifications. The U.S. PTO for example defines over 100,000 detailed classifications based on technical characteristics in the "Examiner Handbook to the U.S. Patent Classification System."[1] These detailed subclasses are structured under about 400 3-digit higher-level "classes." Hall et al. (2001) grouped these 400 classes into 36 technical categories, each of which they then assigned to one of six broad technical categories. As in previous work, we use these groupings to delineate technical domains (Benner and Waldfogel, 2008).[2] Like Jaffe (1986) we view the technical, rather than product or market, orientation of the patent classification system as an asset for our study, since we are interested in the technical flows of knowledge. Similarly, even though many of these citations are added by examiners (Alcacer and Gittelman, 2006), there is evidence that these additions enhance the picture of whence knowledge derives, as it corrects for incentives to ignore potentially overlapping prior art (Criscuolo and Verspagen, 2008; Schoenmakers and Duysters, 2010). We address concerns about the validity, precision, and consistency of patent classification (Dahlin and Behrens, 2005) in four ways: first we collapse the multiple dimensions of technological proximity to a simple dichotomy between inside and outside of a technical domain; second, we avoid use of the most detailed level of clas-

[1] www.uspto.gov

[2] We do not make use of the >100,000 detailed classes because these distinctions likely over-state their precision with the trend toward extensive cross-referencing of patent applications into multiple classes.

sification; third, we use multiple classification systems; and fourth, we use nested sub-categories within classifications.

2.2.3. Backward patent citations as knowledge flows

Many studies have found that citations from one patent to another provide a means for measuring the flows of knowledge (Trajtenberg, 1990; Jaffe et al., 1998, 2000). Detailed technology case studies of citation networks generally find that cited prior does include precursor inventions (Mina et al., 2007; Fontana et al., 2009; Barbera-Tomas et al., 2011). These results support subsequent work employing backward citations to measure knowledge flows across geographical space (Maurseth and Verspagen, 2002; Mariani, 2004) and to a lesser extent social space (Sorenson et al., 2006). We treat backward patent citations as flows across *technological space* (Benner and Waldfogel, 2008), and make only the modest assumption that that citations from one patent class to another represent knowledge flows across technological domains.

3. Data and methodology

We constructed variables indicating how frequently each patent was cited by subsequent patents (*forward citations*), as well as how many and what types of patents each patent cited as prior art (*backward citations*). We used the resulting data set to evaluate several tests of the following null hypothesis:

H_0: the extent to which a patent cites prior art outside its technological area has no significant effect on the number of times it is cited by subsequent patents.

Based on the discussion in section 2, we expected to reject the null hypothesis, and that increasing the amount of external prior art would have a positive effect on forward citation frequency. The practical challenges in testing this hypothesis center on: (1) establishing boundaries of technical domains (2) establishing the value of patents, and (3) controlling for the other determinants of value.

3.1. Patent data and timing

We use patents granted by the U.S. Patent and Trademark Office (PTO) from 1976–2006. Patent citation data come from a revised version of the National Bureau of Economic Research (NBER) Patent Citation Data File (Hall et al., 2001). We impose a ten-year window—on both forward and backward citation pairs—to minimize truncation bias An alternative method to address truncation is to estimate adjusted patent citations using the shape of the overall citation lag distribution (Hall et al., 2005) or the patents available for citation over time (Dahlin et al., 2004). We choose the blunter method for its simplicity and reduced tendency to inflate the importance of recent patents with a single early forward citation. We choose a 10-year window, rather than the shorter 5-year window used in related work (Mariani, 2004; Nemet, 2009) to account for the possibility that novel combinations take longer to be used in subsequent inventions. The price paid is that the 10-year window restricts our evaluation to patents i granted from 1986–1996 (n=1,020,484). Truncation issues are minimal because every patent i has exactly 10 years available to receive forward citations, and 10 years of previous patents to which it can make citations as prior art (Fig. 1). For the forward citations, the 10-year window begins at the grant date of each patent i. For example,

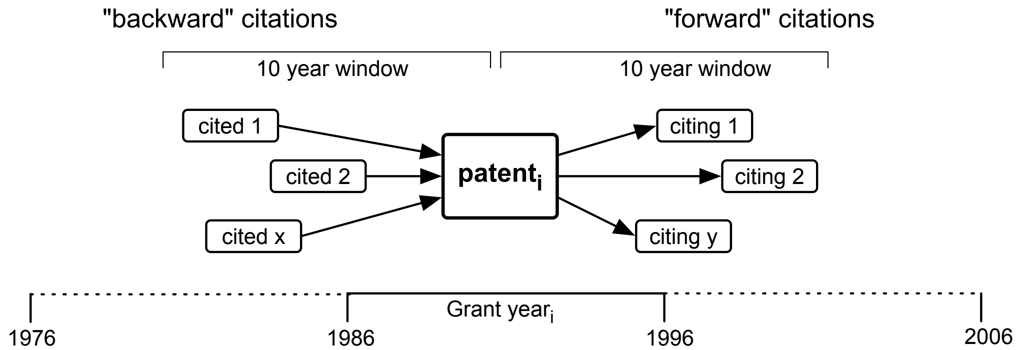

Figure 1: Schema for patent citations showing forward and backward citations to and from patent of interest, i. Arrows indicate flows of knowledge.

for a patent i, granted in 1990, forward citations include those citing patent i that were granted before 2000. Backward citations from patent i include citations to patents that were granted after 1980.

3.2. Dependent variable: forward citation frequency

We use "forward citations" as a proxy for the importance of a patent. Counts of forward citations received within 10 years provides the dependent variable, *citations received*$_i$, used in all models. Table 2 provides a summary of the variables used, which are described in this section.

3.3. Knowledge flows

The key variables of interest for this study are those that measure knowledge flowing from one technical domain to another. We use backward citations to indicate flows of knowledge from each cited patent to each patent i.[3] We refer to each citation from patent i to a previous patent as a citation

[3]Citations *to* prior art use language that implies the reverse direction of knowledge flows; knowledge is assumed to flow *from* the cited patent *to* the citing patent.

pair and we construct several measures characterizing the relationship for each citation pair. For each patent i, we construct counts of each type of pair.

Our basic approach is, for each pair, to compare the technical classification assigned to each cited patent with the classification for the citing patent. To avoid strategic citations to prior art, we exclude self-citation pairs (9.8%) from the data set by removing instances in which the assignee for patent i and the assignee for the patent that i cites are the same (Hall et al., 2005). We count the total number of patents that patent i cites to define $\textit{citations made}_i$. We consider prior art *external knowledge* if it does not reside in the same classification as patent i. We assume that the classification to which a patent is assigned represents that patent's *technical domain*, and that any patent with a different classification is *external* to that technical domain. We make use of the patent classification system, which assigns each patent primarily to one of over a 100,000 detailed sub-classes. The U.S. patent office groups these detailed classes to higher level categories called "classes."[4] Hall et al. (2001) (HJT) aggregate these higher-level categories further. They use the 428 3-digit classes defined by the U.S. PTO and group them into 36 "sub-categories" and 6 top level "categories." We use these grouping to define *sub-classes*, *classes*, and *super-classes* below.

We use the hierarchy defined by these three highest levels of aggregation to define the technological relationship for each citation pair. Fig. 2 shows

[4]These classification systems change over time, which is especially problematic with a data set that spans 31 years. To address this issue, we use the reclassification that occurred in 2008 so that early patents have assignments within the same hierarchy of classes as later patents.

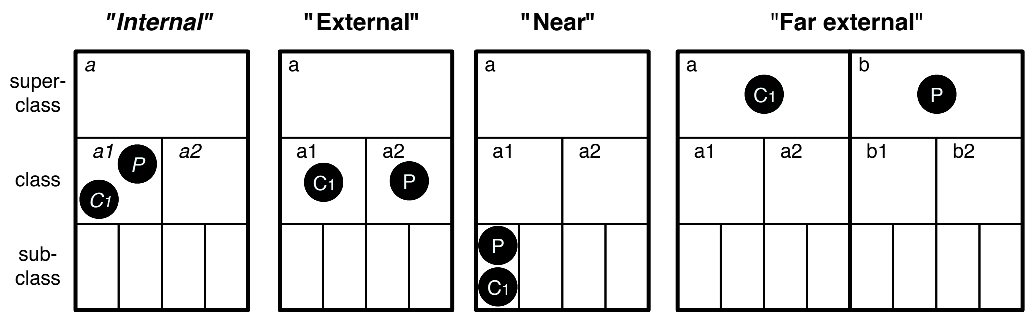

Figure 2: Definitions of external knowledge flows. P is patent of interest, patent i. C_1 is prior art cited by patent i.

three levels of aggregation, which we refer to generically as *super-class*, *class*, and *sub-class*. We code each citation pair as *external*, *far external*, or *near* by comparing their classifications at each level in the hierarchy. For example, a pair is considered *internal* if both patent i and the cited patent are in the same *class*. If they are in different classes, they are considered an *external* pair. The pair is coded as *near* if the pair share the same *sub-class*. The pair is coded as *far external* if the two are in different *super-classes*. We use this latter measure as the primary indicator of external knowledge flow; we use *near* to represent the opposite. This coding scheme is used to develop 3 variables using counts of backward citation pairs for each patent i: $far\ external_i$, $external_i$, and $near_i$. Given the observation that patent counts in general increase forward citations (Bessen, 2008), we add counts of *other* citations made, defined as: $other = citations\ made - (far + near)$. Because we later drop it, *external citations* are not included in the definition of *other*. We find that using *other* citations, rather than *total* citations made avoids collinearity in the regressors. Using USPTO classifications *other* has a correlation with *far* of 0.13 and with *near* of 0.14. *Near* and *far* have a correlation

Table 1: Patent classification structure. Values in parentheses show number of categories.

	U.S. Patent and Trademark Office	International Patent Classification
super classes	HJT "category" 1 Computers and Communications 2 Drugs and Medical 3 Electrical and Electronics 4 Chemical 5 Mechanical 6 Others	IPC "section" 1 Human Necessities 2 Operations; Transport 3 Chemistry; Metallurgy 4 Textiles; Paper 5 Fixed Constructions 6 Mechanical; Light; Heat 7 Physics 8 Electricity
classes	HJT "sub-category" (36)	IPC "class" (124)
sub-classes	3-digit PTO "class" (428)	IPC "sub-class" (1053)

coefficient of 0.11. Results under the IPC definitions are similar: 0.09–0.15, as well as under strict definitions: 0.06–0.20. Results on correlations are included in the appendix and are discussed in the next section.

In order to improve the robustness of the results to inconsistencies in patent examiners' classification decisions, we employ a parallel but distinct taxonomy, the International Patent Classification system (IPC). We arrive at a comparable hierarchy of super-classes (8), classes (124) and sub-classes (1053) and code each citation pair according to these categories as well. Finally, we develop a third coding for each citation pair using the intersection of the coding for the U.S. system and the international system. This intersection creates the strictest definition of external. Table 1 shows the parallel classification structures used in this analysis.

3.4. *Other determinants of patent value*

We control for time trend, citation lags, patent breadth, the type of organization patenting, and technological areas with higher propensity for being cited.

3.4.1. *Time of patenting*

The 10-year forward citation window helps control for patent vintage effects. Still some periods may be more technologically productive and thus raise the potential for citation among recently granted patents. Further, the general increase in patenting over time increases citations probabilities for later patents. We add a variable for the year in which a patent is granted ($grant\ year_i$). We also include a variable for the year in which the application was made ($application\ year_i$) but observe no difference in results when using that indicator.

3.4.2. *Citation lag*

We define citation lag as the time between the year in which a cited patent was granted and the year in which a citing patent was granted. Previous work has found that: spillovers from one sector to another peak rather quickly in the life of a patent, for example within two years (Verspagen and De Loo, 1999); and also that citation lag is a significant, although small, determinant of forward citations (Criscuolo and Verspagen, 2008). We create $citation\ lag_i$, the mean of the citation lags for all of the patents cited by a patent i within 10 years. Previous work suggests that we should expect patent value to decrease with increases in mean citation lag.

3.4.3. Technological field

Because some technological areas are inherently more dynamic, especially during periods when technological opportunity is high (Nelson and Wolff, 1997; Nelson, 2003), we measured mean citations received for each of the 6 broad technological categories defined in the most aggregated classification scheme: HJT category. We found that three of these—computers and communications; drugs and medical; and electricity and electronics—received significantly and substantially more forward citations than the others. We add dummies for these three technological categories to control for high technological opportunity in these areas ($computers_i, medical_i$, and $electrical_i$).

3.4.4. Patent breadth

Since some patents stake out a larger swath of intellectual property than others, we use the number of *claims* each patent i makes as an indicator of its breadth. The expectation is that broader patents, those with more claims, are more likely to be cited frequently.

3.4.5. Type of inventor organization

We use information about the type of organization whence each patent originated using the U.S. PTO categorization of each assignee. Possible categories for assignee type are: (1) unassigned (typically individual inventors); (2) U.S. non-government organizations (typically corporations); (3) Non-U.S., non-government organizations (typically corporations); (4) U.S. individuals; (5) Non-U.S. individuals; (6) U.S. Federal Government; and (7) Non-U.S. governments. We add binary variables for U.S. corporate assignees ($U.S.corp = 1$ if type 2) and government assignees ($govt = 1$ if type 6 or 7).

Table 2: Variables used to characterize each patent, i.

Name	Description
Dependent variable:	
Citations received	Count forward cites within 10 years
Knowledge flows (counts):	
Far external	Citations made, coded as *far*
External	Citations made, coded as *external*
Near	Citations made, coded as *near*
Other citations	Citations made, not *near* or *far*
Technological fields (HJT):	
Computers	Computers and communications (1/0)
Medical	Drugs and medical (1/0)
Electrical	Electricity and electronics (1/0)
Other characteristics:	
Claims	Breadth: number of claims made
U.S. Corp.	Assignee: U.S. corporation (1/0)
Govt.	Assignee: government (1/0)
Grant year	Year patent was issued
Citation lag	Mean backward citation lag for all cites made

Table 3: Descriptive statistics (n=1,020,484 observations).

Variable	mean	median	std. dev.	min.	max.
Citations received	7.13	4	11.10	0	754
Citations made	3.79	3	4.39	0	213
Claims	12.64	10	10.38	1	868
Grant year	1991.3	1991	3.115	1986	1996
Citation lag	5.13	5	1.96	0	10
Computers	0.127		0.333	0	1
Medical	0.089		0.284	0	1
Electrical	0.183		0.387	0	1
U.S. Corp.	0.428		0.495	0	1
Government	0.011		0.105	0	1

The hypothesis is that corporations have the best information about both technology and market opportunities and are thus more adept at generating and commercializing valuable inventions. Empirical work support this hypothesis (Bessen, 2008).

4. Results

We regress the knowledge flow variables and controls on counts of forward citations using negative binomial models. Our focus is on the influence of external citations.

4.1. Descriptive statistics

Table 3 shows descriptive statistics for all variables other than the knowledge flow variables. Note that the mean for *citations received* is higher than that of *citations made* because the number of patents granted each year has grown substantially, by over 50% from 1986–96.

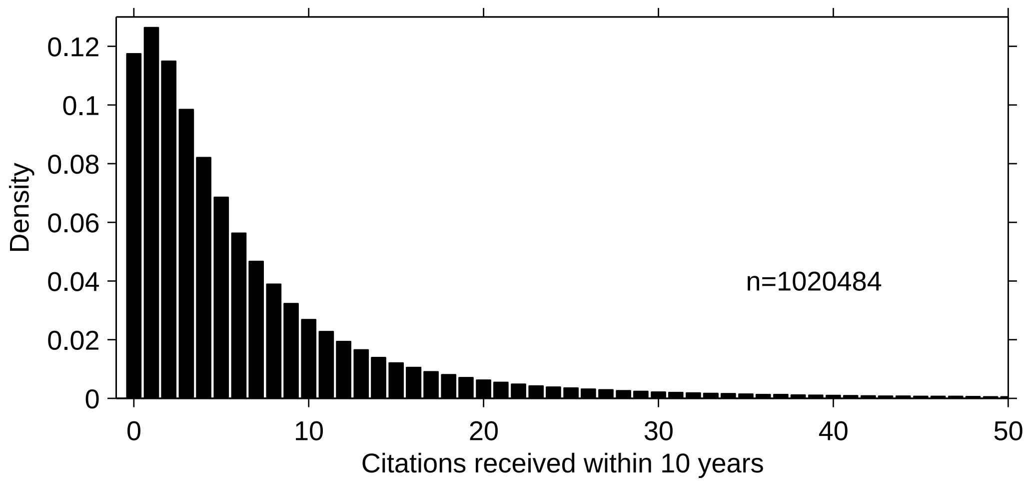

Figure 3: Distribution of values for the dependent variable, citations received within 10 years. One percent of observations had values above 50 and are not displayed on this figure.

Fig. 3 shows the probability distribution for the dependent variable, citations received within 10 years. Only 12% of patents had no forward citations and less than 1% had more than 50 citations. Note that this shape of the distribution is sensitive to the choice of citation window. A 5-year citation window would produce many more patents with zero citations, and a longer window would shift the central tendency to the right.

Table 4 shows the mean share of backward citations for each type of citation pair relationship. On average, 21% of a patent's backward citations are to prior art categorized as far external using the USPTO system. Only 11.5% are coded as far external in both the USPTO and IPC. Over 50% of patents' backward citations are to prior art coded as technologically *near*, meaning it resides in the same 3 digit US PTO classification as the citing patent. Given our interest in making distinctions among types of backward citation relationships, we are particularly concerned about collinearity in

Table 4: Portion of total citations that are coded as far external, other, and near. Values are means across n=878,226 observations with > 0 citations made within 10 years. The three types of citation pairs are shown under three classification systems.

	USPTO	IPC	US+IPC
Far external	0.211	0.229	0.115
Other	0.202	0.220	0.436
Near	0.588	0.550	0.449

these variables. The tables in the appendix show the correlation coefficients (Table 9) and indices of collinearity (Table 10). We find no correlations above 0.20 and a maximum variance inflation factor (VIF) of 1.09, well below the level of concern.

4.2. Estimation approach

Our estimation approach is as follows. First, because the dependent variable consists of count data, we consider a Poisson regression. However, the variance, 123 is much larger than the mean 7.1, indicating the data are overdispersed, and suggesting that a negative binomial regression would be more appropriate. This choice is confirmed by our results, which show α values consistently significantly different from zero. That the mode for citations received is 1, not zero, implies that the variable does not have an excessive number of zero values, and thus supports the use of the negative binomial regression model. We use robust estimators to avoid heteroskedasticity; all z-statistics reported are based on robust standard errors. Second, as a robustness check, we also run the negative binomial regressions using *grant year* fixed effects. We include grant year as a control, but it is possible that events may have occurred in individual years that are not accounted for by the general increase in patenting activity captured by the inclusion of the grant year trend variable. This allows identifying effects within each of the 11 years in which patents could be issued.

To assess the robustness of the results to bias in patent examiners' classification decisions for each patent, we employ various ways of defining *externality* and specify 7 models. The first 3 models (1–3) use counts of *far external* citations and the next 3 (4–6) add counts of *near* citations. Models 1 and 4 define near and far using USPTO classifications, while 2 and 4 use those from the IPC. Models 3 and 6 use the strict definitions of near and far; each pair has to be coded the same in USPTO and IPC to be considered near or far. Model 7 is the same as 6 but adds the variable *citation lag*. We do not use this variable in other models because it requires dropping all observations with no citations to prior art within 10 years (about 12% of the patents). We drop the use of *external* knowledge flows from our models as it represents an intermediate measure of technological proximity and is a weak indicator relative to the variable for *far external*, on which we focus our analysis.

4.3. Regression results

A first general observation is that the results are stable across the 7 specifications using the negative binomial regression model (Table 5). The results are robust to the classification scheme used to define external and near. They are also robust to whether counts of near are included. Coefficient values are, without exception, significant. Controls are in the expected directions. We include the results for negative binomial regressions with grant year fixed effects in the Appendix (Table 8).

Table 5: Estimates of negative binomial regressions. Dependent variable is counts of citations received within 10 years. Second row indicates classification system used.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	US	IPC	US+IPC	US	IPC	US+IPC	US+IPC
Far citations	0.0523***	0.0496***	0.0493***	0.0522***	0.0498***	0.0485***	0.0428***
	(69.61)	(69.62)	(48.11)	(69.46)	(69.82)	(47.22)	(42.93)
Near citations				0.0532***	0.0579***	0.0517***	0.0394***
				(107.16)	(114.63)	(93.90)	(72.67)
Other citations	0.0553***	0.0563***	0.0554***	0.0596***	0.0528***	0.0585***	0.0510***
	(129.93)	(130.67)	(137.02)	(72.63)	(67.60)	(104.53)	(92.78)
Citation lag							-0.0531***
							(-76.26)
Claims	0.0122***	0.0122***	0.0122***	0.0122***	0.0123***	0.0122***	0.0119***
	(91.68)	(91.75)	(91.72)	(91.75)	(91.78)	(91.74)	(87.10)
Grant year	0.0392***	0.0392***	0.0392***	0.0392***	0.0392***	0.0392***	0.0393***
	(94.96)	(94.95)	(94.99)	(95.02)	(94.88)	(95.12)	(91.65)
Computers	0.656***	0.655***	0.655***	0.654***	0.654***	0.655***	0.612***
	(173.74)	(172.84)	(173.21)	(172.98)	(172.75)	(173.12)	(158.39)
Medical	0.516***	0.514***	0.516***	0.516***	0.513***	0.519***	0.575***
	(98.17)	(97.70)	(98.18)	(98.37)	(97.40)	(98.45)	(103.78)
Electrical	0.263***	0.262***	0.263***	0.263***	0.261***	0.264***	0.234***
	(79.05)	(78.75)	(79.03)	(79.30)	(78.39)	(79.53)	(67.70)
U.S. Corp.	0.166***	0.166***	0.166***	0.165***	0.166***	0.165***	0.177***
	(61.97)	(62.08)	(62.03)	(61.57)	(62.16)	(61.69)	(63.50)
Constant	-76.80***	-76.79***	-76.81***	-76.81***	-76.75***	-76.89***	-76.77***
observations	1,020,482	1,020,482	1,020,482	1,020,482	1,020,482	1,020,482	878,226
Log-likelihood	-3.0E+06	-3.0E+06	-3.0E+06	-3.0E+06	-3.0E+06	-3.0E+06	-2.6E+06
Pseudo R^2	0.0310	0.0310	0.0310	0.0310	0.0310	0.0310	0.0310
$\ln(\alpha)$	-0.0355***	-0.0356***	-0.0355***	-0.0356***	-0.0356***	-0.0357***	-0.0878***
$p : \beta_{far} < \beta_{other}$	0.0004	0.0000	0.0000	0.0000	0.0033	0.0000	0.0000

[a]

[a] Robust z-statistics in parentheses. *** p<0.01, ** p<0.05, * p<0.1

4.3.1. Positive coefficients on knowledge flow variables

The coefficients for all three knowledge flow variables—*far*, *near*, and *other* citations—are positive and significant in every specification tested, including fixed effects models. The sizes of the coefficients are also reasonably consistent across models. As a robustness check, Table 8 in the Appendix shows results with grant year fixed effects—to account for any environmental changes in particular years not accounted for in the general increase in patenting over time. All three knowledge flow variables remain positive and significant across the 7 model specifications.

4.3.2. Controls

Across all 7 models, controls are significant and in the expected directions. Signs and sizes of each coefficient are stable across models. Computers, medical, and electronics patents raise citation frequency, with the effects of computers and medical about twice as large as electronics. Patents filed by U.S. corporations have a positive coefficient, although this effect is smaller than filing a patent in one of the highly cited technology areas. Grant year is positive and stable across models. We expected this result because the growth in patenting over time allows for higher forward citation probabilities over time; the total number of patents in the ten-year forward citations window rises over time. Claims is positive, significant and stable. Broader patents cover a wider swath of technological territory and thus are more likely to be cited. Citation lag is included in model 7, which is a variation on model 6. It is negative, which is consistent with the notion that knowledge depreciates, so older prior art is less valuable than more recent art, as also found in Schoenmakers and Duysters (2010). In the fixed effects results, signs and significance of control variables are all similar to those in the ungrouped results.

4.3.3. Comparing effects of external to other citations

Counts of external citations are positive predictors of forward citations, but are they more important than other types of citations? On a first look, the positive and significant result on far external citations appears to support the hypothesis that assimilating external knowledge stimulates more important innovations. However, a comparison across the three types of knowledge flow variables provides evidence of the opposite. We compared coefficients of far external citations and other citations. Fig. 4 compares these coefficients across all 7 model specifications for the negative binomial regressions and negative binomial fixed effects regressions. We also include ratios for technology specific subsets of models 4 and 6, described below. The vertical axis in the figure represents the ratio of the coefficients for far external and other citations ($\beta_{far}/\beta_{other}$). The dashed line, where the ratio equals 1 is where there is no difference between citations to external prior art and citations to other prior art. In the region above the line, external citations have a larger effect on forward citations than do other citations. In the region below, external citations have a weaker effect. One can see that in almost every case the coefficients for external are nearly the same, but smaller, than those for other citations. The ratios are generally smaller in the fixed effects models than in the original negative binomial. These differences are significant. Wald tests were run to test the hypothesis that $\beta_{far} < \beta_{other}$. In almost every specification, differences were significant with $p < 0.01$. The difference in model 5 using fixed effects is barely above the 0.01 threshold. These test

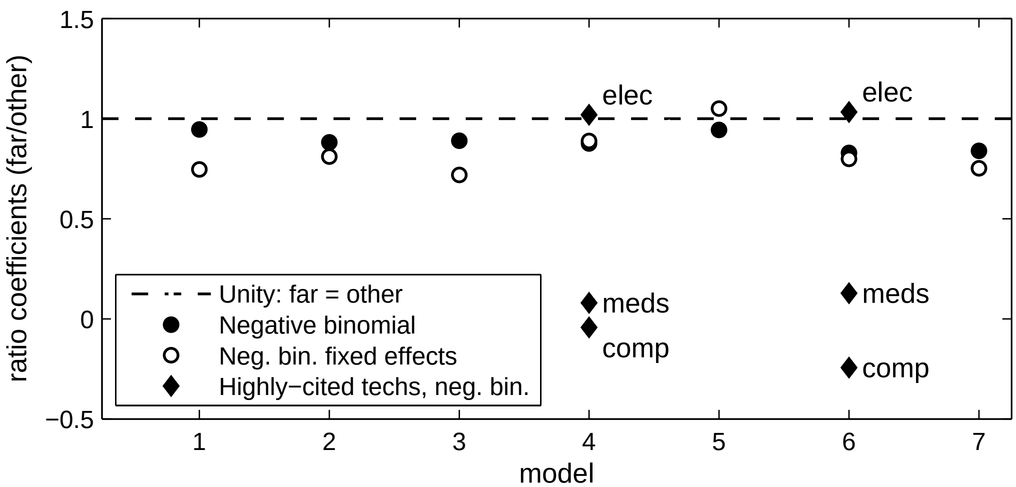

Figure 4: Comparison of coefficients for far external to those for other citations. Model numbers refer to specifications in Tables 5 and 8. Vertical axis = $\beta_{far}/\beta_{other}$.

results are shown in the bottom row of the regression results tables.

Finally, we find that, with only one exception, the effect of near citations is always larger than that of far citations. This results holds for all model specifications, including fixed effects and technology subsets. In only model 7 is the coefficient for near smaller than that of far. Moreover, in the fixed effects models, near citations are more important than other citations. The relative importance of near citations is mixed in the pooled negative binomial regressions as well as in the technology subsets.

4.4. Results within technological categories

Because of the strong results found on the three controls for technological categories—computers, medical, and electronics—we look at the effect of external knowledge within each of these highly-cited technology categories. Table 6 shows the results using the strict definition of external (model 6 from

Table 5). The results point to important differences from the pooled results. All coefficients remain significant, however, the knowledge flow variables have changed. The category for computers and communications shows the most dramatic change. External citations have a significant negative effect on forward citations. Near and other are positive. Fig. 4 shows this negative effect, in both model 4 and model 6. Unsurprisingly a Wald test (Table 6) shows a significant difference between far and other citations. Results for medical and drugs are similar. The coefficient for far external is not actually negative, but the ratio comparing it to other citations is very small. For the two most highly-cited technology categories, external citations are much less effective than other types of citations. There is even evidence that, for computer technologies, an external citation lowers forward citations frequency. For the third technology category, electrical and electronics, the results look quite similar to the pooled estimates, although the difference between external and other citations is not significant.

4.5. Whence did knowledge flow?

Finally, we briefly look at the sources of knowledge flows for the most highly cited patents. We restrict this analysis to the set to patents in the top 25th percentile of citations received (≥ 9). Echoing approaches by Schmookler (1966) and Scherer (1982a,b), Table 7 represents an input-output table for knowledge flows across technology areas, We only include citation pairs which were defined as far external using the USPTO definition, which is why the diagonal is empty. One can think of knowledge flowing from the categories for each row to the categories for each column. Column totals show the share of each technology category among highly-cited patents receiving

Table 6: Technology areas with high citation propensity: estimates of negative binomial regressions (model 6). Dependent variable is counts of citations received within 10 years. Far external and near use the strict definition of external (U.S.+I.P.C.).

	Computers & communications	Medical & drugs	Electrical & electronics
Far ext. citations	-0.0137***	0.0104***	0.0467***
	(-5.27)	(3.62)	(20.22)
Near citations	0.0300***	0.0678***	0.0544***
	(24.45)	(41.39)	(43.22)
Other citations	0.0534***	0.0814***	0.0452***
	(48.76)	(39.21)	(36.49)
Claims	0.0131***	0.0147***	0.0121***
	(43.97)	(31.23)	(39.94)
Grant year	0.0471***	0.0220***	0.0535***
	(45.12)	(13.35)	(58.41)
U.S. Corp.	0.345***	0.232***	0.154***
	(50.76)	(22.81)	(26.20)
Constant	-91.85***	-42.20***	-105.1***
observations	129,642	90,638	187,018
Log-likelihood	-451386	-298014	-564575
Pseudo R^2	0.0229	0.0227	0.0175
$\ln(\alpha)$	-0.110***	0.264***	-0.105***
$p : \beta_{far} < \beta_{other}$	0.0000	0.0000	0.2910

[a]

[a]Robust z-statistics in parentheses. *** p< 0.01, ** p< 0.05, * p< 0.1

Table 7: Inter-domain knowledge flows for highly cited patents (top quartile for citations received). Values indicate percentages of total far external citation pairs (n=666,630 pairs).

	Destination (technology category for patent i)						
	Chem	Comp	Med	Elec	Mech	Other	Total
Source:							
Chem	-	0.8	4.1	3.0	3.8	6.2	18
Comp	0.8	-	0.5	8.4	3.8	1.6	15
Med	3.9	0.4	-	0.8	1.1	1.6	8
Elec	3.0	9.1	1.2	-	4.6	2.8	21
Mech	3.5	4.2	1.1	4.4	-	6.4	20
Other	5.9	1.6	2.0	2.8	6.7	-	19
Total	17.1	16.2	8.9	19.3	20.0	18.5	100

external knowledge; electrical and mechanical are highest. Row totals indicate each sector's contribution of knowledge flows to highly-cited patents in other sectors. The same two categories were most important. Setting aside the "other" categorization, the 3 largest flows are: from electronics to computers, from computers to electronics, and from electronics to mechanical.

5. Discussion

Our main finding is that citations to external prior art are significantly less important to subsequent forward citations than are other types of citations. These results are robust across patent classification schemes, which allowed for three different definitions of externality. They are also robust to the use of a model using grant year fixed effects. The results are even stronger when applied to the two most highly-cited individual technology categories: computers and medicine. In these most highly-cited areas, external citations had an especially adverse effect. Conversely, citing prior art that is technologically near has a stronger effect than far citations, and in most cases a stronger effect than all other types of citations. These results do not appear to support the hypothesis that external knowledge leads to more important inventions. Rather, we see significant effects in the opposite direction.

5.1. Why do citations in the technological proximity increase value?

These results apparently contradict the substantial set of claims in the literature about the importance of inter-sectoral knowledge flows. We suggest a few interpretations. A first possibility is that the results are biased due to one or more measurement problems. This problem could stem from our fundamental assumption that patents provide useful measures of inventions and knowledge flows. More specifically, we make three important assumptions in our use of patent data. We assume that forward citations received is a useful proxy for the value of an invention. Counts of forward patent citations may be ill-suited to characterizing the importance of those inventions that feature in the case studies that emphasize external flows of knowledge. Note that this problem requires more than an acknowledgement of the 'noisiness' of patent citations. Rather, it implies that these important historical inventions are disproportionately less likely to be patented and cited. There is however a large body of work now supporting the assumption that citations proxy for value. Distinguishing between radical and incremental inventions, as in Hurmelinna-Laukkanen et al. (2008), while not feasible with a large data set, may add information that avoids bias. We also assume that backward citations represent knowledge flows. One could argue that inventors are less likely to cite prior art that is external, perhaps because there is less concern among inventors about overlapping intellectual property claims for distant technological domains. It is also possible that patent examiners fail to add prior art from external domains because those areas lie beyond the examiners' own expertise. Still, the strong results on the importance of the total number of citations suggests that inventors benefit from being diligent in their search for and descriptions of prior art. We also assume that patent classification systems provide a useful delineation of boundaries between technological domains. Perhaps examiners assign patents inconsistently and that classification-based technological domains overlap. Our efforts to use only high-level classes, to compare results across multiple classification systems, and our construction of a strict definition that combines systems alleviate these concerns. That the results are consistent across the classification schemes make it unlikely that the use of patent class assignments is a source of bias. We provided justifications based on the literature for each of these assumptions in section 2.

A second explanation is that incorporating external knowledge is hard to do well and is risky. The literature on absorptive capacity shows that efforts to assimilate external knowledge take investment and considerable time to develop. Coordination costs are higher due to delays and higher likelihood of project termination (Cuijpers et al., 2011). They often require expertise and relationships that are outside a firm's typical activities. In addition, even if resources, patience, and expertise are available, the incorporation of external knowledge for new inventions likely involves more risk. Inventors and firms lack the deep familiarity they have with more proximate knowledge. Unforeseen bottlenecks may emerge; incompatibilities may arise; prolonged revision and iteration may be required, perhaps extending beyond what a firm or investor is willing to tolerate. Normatively, studies that code projects by their technical uncertainty show the need for quite different management strategies (Shenhar, 2001). Even so, the literature suggests that the combinations that do work out are so valuable that they justify the uncertain returns on investment (Scherer and Harhoff, 2000). Perhaps firms are risk averse or the benefits of resulting inventions are difficult to sufficiently appropriate. So one explanation for the weak effect of external knowledge is that successful inventions are few and failures prevalent. The distribution of payoffs for inventions may be so extremely skewed that even the highly-skewed distribution of forward citation counts (Fig. 3) does not adequately characterize the magnitude of the payoffs for the most successful inventions.

A third explanation is that inventions that emerge from the exploitation of inter-industry knowledge flows may be particularly difficult to assimilate into subsequent inventions, even if they become useful and commercially successful themselves. These inventions from hybrid fields may not provide easy to use building blocks that enable further combinations. They may display unfamiliar, perhaps even awkward forms, with difficult to predict properties. They likely lack the adherence to standards that facilitates progress within technical areas. Their suitability for subsequent innovations may require a prolonged period of testing, understanding, and incremental improvement. This latter issue may be particularly relevant in this study due to our imposition of a ten-year forward citation window. The exceptional novelty of hybrid inventions may inhibit, or at least delay, their ability to serve as prior art for subsequent inventions. One possibility would be to assess value based on the next generation of forward patent citations (Popp, 2006). A related possibility is that these challenges make it especially difficult to incorporate more than a small amount of external knowledge in a new invention. Perhaps we see rather weak effects from external citations because many inventions encounter diminishing returns to their external citations.

5.2. Conclusion

This discussion proposed several explanations for the result we observed—that additional flows of knowledge from distant technological domains are associated with fewer forward citations received relative to adding citations to more proximate knowledge. Some of these explanations arise from the possibility that the construction of our variables, and their interpretation, produced biased results. We took care to address these concerns as much as possible, by referencing prior literature and extensively assessing robustness across definitions. Still, understanding of what patents and patent citations actually represent, particularly in their characterization of inventions, value, and knowledge flows, is incomplete. The reliability of results in a study like this depends on improved understanding of how to interpret patent data. Evaluating international patent data might also enhance reliability. Incorporating other measures of the value of inventions, such as expert elicitation or case studies of commercialization, would obviate the need to make important assumptions on the meaning of forward citations. Scalability of such labor-intensive techniques is a limitation.

These potential sources of measurement bias may be small or offsetting. More mechanistic explanations exist as well. Compared with more familiar areas, working with external knowledge is costly, time-consuming, and risky. Hence this study may accurately characterize the role of the many technical failures, conversely it may under-estimate the magnitude of the important, but rare, successful inventions arising from cross-technology knowledge flows. Still, the possibility exists that external knowledge flows really do have an adverse effect on invention. They may be distracting, too difficult, and too prone to failure to justify investment. Resolving this open interpretation has important policy implications. If inter-industry knowledge flows are valuable then firms should develop capabilities to develop them. If their payoffs are too uncertain for risk-averse firms, then it may be necessary for governments to reduce uncertainty—for example by supporting cross-cutting knowledge exchange in intellectual communities, conferences, and longer term collaboratives. If inventions resulting from inter-industry knowledge flows are important but more difficult to appropriate than other inventions, then governments may need to increase incentives for the development of cross-sectoral science and technology. The results from this study suggest that the payoffs from inter-sectoral knowledge flows, if they do exist, are obscured in patent data and may be offset by substantial risk and appropriability concerns.

Acknowledgments

We thank the Wisconsin Alumni Research Foundation for support. This paper benefitted from comments by participants at the Association for Public Policy Analysis and Management Fall Conference, as well as during seminars at Syracuse University and at Harvard University where Nemet conducted much of this work as a Visiting Scholar. We especially appreciate suggestions from Stuart Bretschneider, Jonathon Mote, David Popp, and Michael Scherer.

References

Ahuja, G., Lampert, C. M., 2001. Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions. Strategic Management Journal 22 (6-7), 521–543.

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

Allison, J. R., Lemley, M. A., Moore, K. A., Trunkey, R. D., 2004. Valuable patents. Georgetown Law Journal 92 (3), 435–479.

Arthur, W. B., 2007. The structure of invention. Research Policy 36 (2), 274–287.

Arthur, W. B., 2009. Combination and Structure. Free Press, New York, Ch. 2, pp. 27–44.

Barbera-Tomas, D., Jimenez-Saez, F., Castello-Molina, I., 2011. Mapping the importance of the real world: The validity of connectivity analysis of patent citations networks. Research Policy 40 (3), 473–486.

Benner, M., Waldfogel, J., 2008. Close to you? bias and precision in patent-based measures of technological proximity. Research Policy 37 (9), 1556–1567.

Bessen, J., 2005. Patents and the diffusion of technical information. Economics Letters 86 (1), 121–128.

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

Criscuolo, P., Verspagen, B., 2008. Does it matter where patent citations come from? inventor vs. examiner citations in european patents. Research Policy 37 (10), 1892–1908.

Cuijpers, M., Guenter, H., Hussinger, K., 2011. Costs and benefits of inter-departmental innovation collaboration. Research Policy 40 (4), 565–575.

Dahlin, K., Taylor, M., Fichman, M., 2004. Today's Edisons or weekend hobbyists: Technical merit and success of inventions by independent inventors. Research Policy 33, 1167–1183.

Dahlin, K. B., Behrens, D. M., 2005. When is an invention really radical?: Defining and measuring technological radicalness. Research Policy 34 (5), 717–737.

Fleming, L., 2001. Recombinant uncertainty in technological search. Management Science 47 (1), 117–132.

Fontana, R., Nuvolari, A., Verspagen, B., 2009. Mapping technological trajectories as patent citation networks. an application to data communication standards. Economics of Innovation and New Technology 18, 311–336.

Garcia-Vega, M., 2006. Does technological diversification promote innovation? an empirical analysis for european firms. Research Policy 35 (2), 230–246.

Geels, F. W., 2002. Technological transitions as evolutionary reconfiguration processes: a multi-level perspective and a case-study. Research Policy 31 (8-9), 1257–1274.

Gilfillan, S., 1935. The Sociology of Invention. Follett, Chicago.

Gilsing, V., Nooteboom, B., Vanhaverbeke, W., Duysters, G., van den Oord, A., 2008. Network embeddedness and the exploration of novel technologies: Technological distance, betweenness centrality and density. Research Policy 37 (10), 1717–1731.

Griliches, Z., 1990. Patent statistics as economic indicators: A survey. Journal of Economics Literature 28 (4), 1661–1707.

Griliches, Z., Pakes, A., Hall, B., 1987. The value of patents as indicators of inventive activity. Cambridge University Press, Cambridge MA, Ch. 4, pp. 97–124.

Hall, B. H., Jaffe, A. B., Trajtenberg, M., 2001. The NBER patent citation data file: Lessons, insights and methodological tools. Working paper, NBER.

Hall, B. H., Jaffe, A. B., Trajtenberg, M., 2005. Market value and patent citations. Rand Journal of Economics 36 (1), 16–38.

Hargadon, A., Sutton, R. I., 1997. Technology brokering and innovation in a product development firm. Administrative Science Quarterly 42 (4), 716–749.

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

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

Hurmelinna-Laukkanen, P., Sainio, L. M., Jauhiainen, T., 2008. Appropriability regime for radical and incremental innovations. R&D Management 38 (3), 278–289.

Jaffe, A. B., 1986. Technological opportunity and spillovers of research-and-development - evidence from firms patents, profits, and market value. American Economic Review 76 (5), 984–1001.

Jaffe, A. B., Fogarty, M. S., Banks, B. A., 1998. Evidence from patents and patent citations on the impact of nasa and other federal labs on commercial innovation. The Journal of Industrial Economics 46 (2), 183–205.

Jaffe, A. B., Trajtenberg, M., 2002. Patents, citations, and innovations: a window on the knowledge economy. MIT Press, Cambridge, Mass.

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

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

Lin, B. W., Chen, C. J., Wu, H. L., 2007. Predicting citations to biotechnology patents based on the information from the patent documents. International Journal of Technology Management 40 (1-3), 87–100.

Mariani, M., 2004. What determines technological hits? geography versus firm competencies. Research Policy 33 (10), 1565–1582.

Maurseth, P. B., Verspagen, B., 2002. Knowledge spillovers in europe: A patent citations analysis. Scandinavian Journal of Economics 104 (4), 531–545.

Mina, A., Ramlogan, R., Tampubolon, G., Metcalfe, J. S., 2007. Mapping evolutionary trajectories: Applications to the growth and transformation of medical knowledge. Research Policy 36 (5), 789–806.

Mowery, D. C., Rosenberg, N., 1998. Paths of Innovation: Technological Change in 20th-Century America. Cambridge University Press, Cambridge.

Nelson, R. R., 1959. The simple economics of basic scientific research. Journal of Political Economy 67 (3), 297–306.

Nelson, R. R., 2003. On the uneven evolution of human know-how. Research Policy 32 (6), 909–922.

Nelson, R. R., Winter, S. G., 1982. An Evolutionary Theory of Economic Change. The Belknap Press of Harvard University Press, Cambridge, MA.

Nelson, R. R., Wolff, E. N., 1997. Factors behind cross-industry differences in technical progress. Structural Change and Economic Dynamics 8 (2), 205–220.

Nemet, G. F., 2009. Demand-pull, technology-push, and government-led incentives for non-incremental technical change. Research Policy 38 (5), 700–709.

Nerkar, A., 2003. Old is gold? the value of temporal exploration in the creation of new knowledge. Management Science 49 (2), 211–229.

Pavitt, K., Robson, M., Townsend, J., 1989. Technological accumulation, diversification and organization in UK companies, 1945-1983. Management Science 35 (1), 81–99.

Popp, D., 2005. Lessons from patents: Using patents to measure technological change in environmental models. Ecological Economics 54 (2-3), 209–226.

Popp, D., 2006. They don't invent them like they used to: An examination of energy patent citations over time. Economics of Innovation and New Technology 15, 753–776.

Rosell, C., Agrawal, A., 2009. Have university knowledge flows narrowed?: Evidence from patent data. Research Policy 38 (1), 1–13.

Rosenberg, N., 1994. Exploring the Black Box: Technology, Economics, and History. Cambridge University Press, Cambridge.

Rosenberg, N., 2009. Some critical episodes in the progress of medical innovation: An Anglo-American perspective. Research Policy 38 (2), 234–242.

Rosenkopf, L., Nerkar, A., 2001. Beyond local search: Boundary-spanning, exploration, and impact in the optical disk industry. Strategic Management Journal 22 (4), 287–306.

Ruttan, V. W., 2001. Technology, Growth, and Development: An Induced Innovation Perspective. Oxford University Press, New York.

Scherer, F. M., 1982a. Inter-industry technology flows and productivity growth. Review of Economics and Statistics 64 (4), 627–634.

Scherer, F. M., 1982b. Inter-industry technology flows in the United States. Research Policy 11 (4), 227–245.

Scherer, F. M., Harhoff, D., 2000. Technology policy for a world of skew-distributed outcomes. Research Policy 29 (4-5), 559–566.

Schmookler, J., 1966. Invention and economic growth. Harvard University Press, Cambridge, MA.

Schoenmakers, W., Duysters, G., 2010. The technological origins of radical inventions. Research Policy 39 (8), 1051–1059.

Schumpeter, J. A., 1934. The Fundamental Phenomenon of Economic Development. Harvard University Press, Cambridge, MA, pp. 57–94.

Shenhar, A. J., 2001. One size does not fit all projects: Exploring classical contingency domains. Management Science 47 (3), 394–414.

Sorenson, O., Rivkin, J. W., Fleming, L., 2006. Complexity, networks and knowledge flow. Research Policy 35 (7), 994–1017.

Suzuki, J., Kodama, F., 2004. Technological diversity of persistent innovators in japan - two case studies of large japanese firms. Research Policy 33 (3), 531–549.

Trajtenberg, M., 1990. A penny for your quotes - patent citations and the value of innovations. Rand Journal of Economics 21 (1), 172–187.

Usher, A. P., 1954. A History of Mechanical Inventions, revised edition Edition. Harvard University Press, Cambridge, MA.

van Zeebroeck, N., 2011. The puzzle of patent value indicators. Economics of Innovation and New Technology 20 (1), 33 – 62.

Verspagen, B., De Loo, I., 1999. Technology spillovers between sectors and over time. Technological Forecasting and Social Change 60 (3), 215–235.

Watanabe, C., Tsuji, Y. S., Griffy-Brown, C., 2001. Patent statistics: deciphering a 'real' versus a 'pseudo' proxy of innovation. Technovation 21 (12), 783–790.

Appendix

This appendix includes results of negative binomial regressions with grant year fixed effects (Table 8), as well as a matrix of correlation coefficients (Table 9) and indices of collinearity in the regressors (Table 10).

Table 8: Estimates of negative binomial regressions showing grant year fixed effects. Dependent variable is counts of citations received within 10 years.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	US	IPC	US+IPC	US	IPC	US+IPC	US+IPC
Far citations	0.0185***	0.0197***	0.0174***	0.0191***	0.0209***	0.0182***	0.0164***
	(66.18)	(67.80)	(39.66)	(68.58)	(70.75)	(41.43)	(35.44)
Near citations				0.0274***	0.0266***	0.0262***	0.0228***
				(163.88)	(168.42)	(140.07)	(105.83)
Citation lag							-0.0388***
							(-79.53)
Other citations	0.0248***	0.0243***	0.0242***	0.0215***	0.0199***	0.0228***	0.0218***
	(204.48)	(199.45)	(212.61)	(99.07)	(74.16)	(144.76)	(128.95)
Claims	0.00627***	0.00627***	0.00627***	0.00625***	0.00625***	0.00626***	0.00619***
	(146.63)	(146.05)	(146.06)	(144.97)	(144.67)	(145.23)	(135.76)
Computers	0.509***	0.508***	0.508***	0.508***	0.507***	0.508***	0.338***
	(204.39)	(204.07)	(204.06)	(204.25)	(203.66)	(204.05)	(101.93)
Medical	0.275***	0.275***	0.276***	0.271***	0.272***	0.274***	0.174***
	(87.93)	(88.03)	(88.57)	(86.54)	(86.78)	(87.47)	(70.18)
Electrical	0.201***	0.201***	0.201***	0.199***	0.199***	0.200***	0.135***
	(85.25)	(85.25)	(85.39)	(84.57)	(84.55)	(84.85)	(71.19)
U.S. Corp.	0.125***	0.124***	0.124***	0.125***	0.125***	0.125***	-0.162***
	(69.02)	(68.96)	(68.83)	(69.29)	(69.11)	(69.11)	(-49.68)
Constant	-0.402***	-0.402***	-0.402***	-0.405***	-0.403***	-0.403***	-0.162***
Observations	1,020,482	1,020,482	1,020,482	1,020,482	1,020,482	1,020,482	878,226
Log-likelihood	-3.0E+06	-3.0E+06	-3.0E+06	-3.0E+06	-3.0E+06	-3.0E+06	-2.7E+06
Groups	11	11	11	11	11	11	11
$p : \beta_{far} < \beta_{other}$	0.0000	0.0000	0.0000	0.0000	0.0103	0.0000	0.0000

[a]

[a]z-statistics in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 9: Correlation matrix for regressors used in model 6 (n=1,020,482 observations).

Variable	1	2	3	4	5	6	7	8
1 Far citations	1							
2 Other citations	0.199	1						
3 Near citations	0.057	0.099	1					
4 Claims	0.088	0.131	0.113	1				
5 Grant year	0.065	0.104	0.107	0.084	1			
6 Computers	0.003	0.120	0.078	0.047	0.064	1		
7 Medical	0.024	−0.004	0.057	0.031	0.038	−0.119	1	
8 Electric	0.005	−0.015	0.040	0.008	0.005	−0.181	−0.148	1
9 U.S. Corp.	0.088	0.129	0.073	0.161	0.022	0.030	0.043	0.031

Table 10: Indices of collinearity for independent variables used in model 6.

	Variance Inflation Factor	Tolerance (1/VIF)	R^2	Eigen value	Condition index
Far citations	1.05	0.950	0.050	1.561	1.000
Near citations	1.04	0.959	0.041	1.181	1.150
Other citations	1.09	0.917	0.083	1.126	1.177
Claims	1.06	0.946	0.054	1.011	1.243
Grant year	1.03	0.971	0.029	0.979	1.263
Computers	1.09	0.921	0.079	0.881	1.331
Medical	1.06	0.946	0.054	0.830	1.372
Electrical	1.07	0.932	0.068	0.771	1.423
U.S. Corp.	1.05	0.954	0.046	0.660	1.537
Aggregate	1.06				1.537