

Does Winning a Patent Race lead to more follow-on Innovation?

Neil C. Thompson¹ and Jeffrey M. Kuhn^{2,3}

August 2016

Abstract

Competition between firms to invent and patent an idea, or “patent racing,” has been much discussed in theory, but seldom analyzed empirically. This article introduces a way to identify patent races, and provides the first broad-based view of them in the real world. It shows that many of the traditional assumptions behind patent racing models are incorrect. It also looks at the causal effect of a winning a patent race, and thus the benefits of getting patent protection in a competitive technology environment, and shows that firms that win do more significantly more follow-on innovation.

¹ MIT Sloan School of Management. neil_t@mit.edu

² PhD Candidate, UC Berkeley Haas School of Business. The author would like to thank the Hoover Institution Working Group on Intellectual Property, Innovation, and Prosperity for generous financial support during the preparation of this article.

³ The authors would like to thank Douglas Hanley for providing patent-firm matching information. We would also like to thank Pierre Azoulay, Scott Stern, Noam Yuchtman, Rui de Figueiredo, participants at the DRUID Conference (Copenhagen 2016), and the seminar group at the College of Management of Technology at EPFL (Lausanne 2015) for helpful suggestions on drafts of the article.

1 Introduction

On February 14, 1876 two men filed patents for the telephone. One, Alexander Graham Bell, received the patent and started a company that was so successful that even today his name is synonymous with the telephone. The other, Elisha Gray, is usually only known, if at all, as the telephone inventor who wasn't Bell.

Their case is perhaps the most famous example of a "patent race" – where two entities race to invent and then get a legal monopoly for their invention. This situation has been much theorized about by, for example, Loury (1979), Lee and Wilde (1980), and many, many others. Typically this is modeled as racing firms making investments over time to develop an idea, with the one that develops it first getting patent coverage and therefore all of the profits from it.

Individually such work produces clear economic predictions, but collectively it produces strongly divergent predictions – with results that depend heavily on initial conditions, market power, information availability, etc. For example, some models predict that patent races will end almost as soon as they begin, because an initial lead by one entity will result in the laggard exiting the race. (Gilbert and Newbery 1982). Others allow for sustained racing, even if one entity is in the lead. (Fudenberg, Gilbert, Stiglitz and Tirole 1983, Harris and Vickers 1985). More modern models (e.g., Judd, Schmedders and Yeltekin (2012)) are subtler about their predictions, highlighting the range of outcomes that could occur, based on differences in empirical parameters.

Results from the real world have the potential to help focus this debate, but, to the authors' knowledge, this subject has not been tackled empirically, except in very narrow settings, because of the difficulty in observing patent races. For, although it is easy to see the patent that issues as a result of winning a patent race, how could one look for the non-existent patent from the loser? And, without seeing both, how would one know that the winning patent was part of a patent race, and not just a typical instance of invention and patenting?

This article introduces a novel strategy for observing patent races. Our strategy takes advantage of a gap in timing between when a patent is filed and when it is made public. During this window a patent has “priority” – future patents that claim the same invention will be disallowed – but it is not yet public knowledge. Thus, during this window, two patent applications can be filed with the patent office for the same invention with both applicants believing that they are first, but only the first applicant will receive the patent coverage whereas the second will have its claims rejected. Importantly, however, *both* applications will be present in the patent office’s records. Using this administrative data, we are able to observe the presence of these events and construct “patent twins” – sets of patent applications filed at nearly the same time on the same invention. We can then contrast these patent twins with all other patents to observe where and to what extent patent racing is happening within the patent system. We can also contrast the outcomes for the winner of the patent race (the “leader”) with the loser of the patent race (the “follower”) to see what effects winning the patent race has.

Although our method cannot observe all patent races – for example, it misses any ‘over before they start’ races – we nevertheless find that patent racing is a frequent occurrence, with 10-11% of all patents in a race. Moreover, we find strong evidence at the project level that winning a patent race is valuable to the winning firm, leading them to maintain their patent coverage and do more and closer follow-on research. However, despite a large sample we find no evidence of firm-level differences in spending on research and development or in profitability from winning an individual patent race. Moreover, we show that one should *not* expect to see such a difference, because the diverse research (and patenting) portfolios of large firms leads them to both win and lose large numbers of patent races, and thus the net number of ‘wins’ will be small. Indeed, we find that the largest firms win and lose patent races in roughly even proportions.

This article proceed as follows. First we introduce both the theory (Section 2) and reality (Section 3) of patent racing, and show that there are large and consequential differences between them. In

Section 4, we discuss where patent races happen and which firms participate. In Section 5 we discuss how we get a causal estimate of the effect of winning a patent race. Section 6 presents these causal estimates, at the patent, firm, and project level. In Section 7 we summarize the implications of our findings for theory models, and in Section 8 we conclude.

2 Patent Races in Theory

Patent racing occurs when firms compete to produce and patent a new invention first. Because the patent system provides a legal monopoly over the technology to the winner, and (in theory) no protection to the loser, winning a patent race allows a firm to extract above-normal profits by selling the invention in a market where competitors are unable to replicate it. Thus, models of patent racing tend to treat a patent as a winner-takes-all prize that is awarded to one of a continuum of racing firms that discovers an independent, unitary invention. Early models included processes for discovery that were either (i) random, or (ii) deterministic. In early random-discovery models, firms compete vigorously until a patent is obtained. (Dasgupta and Stiglitz 1980, Lee and Wilde 1980, Loury 1979, Reinganum 1981, Reinganum 1982). However, a firm could never pull ahead of its rivals because a firm's past R&D was assumed not to affect its current likelihood of discovering an invention. Instead, firm strategies in equilibrium induce a welfare tradeoff between the benefits of faster innovation and the costs of overinvestment in R&D, yielding the result that firms invest in R&D beyond the socially optimum level when market structure is fixed. (Loury 1979), (Lee and Wilde 1980). In early deterministic-discovery models, each firm made observable progress toward the prize. However, the equilibria of these models was ε -competition, unobserved in the real world, in which a firm with an arbitrarily small head start causes all (identical) rivals to immediately drop out of the race because they have no chance of winning. (Dasgupta and Stiglitz 1980, Gilbert and Newbery 1982).

Subsequent articles identified conditions under which firms could leap-frog each other in making progress. A multi-stage R&D process with random discovery can support leapfrogging under perfect competition because a firm running behind can catch up with enough time and luck, and so they will attempt to do so if the expected return winning the patent race is large enough. (Fudenberg, Gilbert, Stiglitz and Tirole 1983). Leapfrogging can also occur with deterministic discovery if firms differ in terms of their “(1) their valuations of the prize, (ii), their discount rates, (iii) their efficiency at making progress, and (iv) their initial distances from the finishing line”. (Harris and Vickers 1985). Finally, leapfrogging can occur if firms accumulate knowledge over time and reduce R&D investments as knowledge stocks increase. (Doraszelski 2003).

Although these theoretical models invoke widely different assumptions and produce correspondingly disparate predictions, they share several commonalities. First, the models assume that the prize is a unitary patent that is won by a single firm and that slower firms receive nothing for their efforts. Second, the models assume that the prize is immediately and observably awarded upon completion of the race, giving all firms the opportunity to modify their strategies moving forward. Third, the models implicitly assume that races are conducted in serial rather than in parallel. Fourth, the models predict that firms will increase or decrease R&D expenditure as a result of the race. We provide empirical evidence that directly contradicts these assumptions and conclusions.

3 Patent Races in the Wild

The empirical work on patent races is thus far extremely limited. An early study of the disk drive industry finds support for Reinganum’s patent racing model. (Lerner 1997). Later, a laboratory experiment with 36 subjects provided limited support for the Harris and Vickers multi-stage model. (Zizzo 2002). This article extends the empirical literature by being the first study to patent races throughout the patent system, the first to catalogue where they occur, and the first to causally identify

the effects of winning or losing a patent race using observational data. To do this, we start by describing institutional features of how patent racing happens in the wild.

3.1 Continuous Technology

In a canonical patent race, firms compete to patent a discrete idea. In reality, the patenting of technology is more continuous. Two firms may try to claim one idea, but it is typically part of set of claims that also include non-overlapping areas. For example, if firm 1 claims a set of ideas, A , first while racing with firm 2, which claims a set of ideas, B , then we would expect firm 1 to 'win the race' on $A \cap B$, and uncontestedly win the remaining claims in A , whereas firm 1 'loses the race' on $A \cap B$, but still uncontestedly wins the non-intersecting claims in B .

Depending on the size of firm 2's non-intersecting claims we would expect to see them abandon the application (if no claims remain, or there is too little value in those that remain to justify the filing or legal costs) or pursue the patent to issuance but with a narrower scope (if the non-overlapping claims still have value). In practice, abandonment is the rarer outcome. More common is for firm 2 to pursue, and get, a narrower patent. In the races we observe, only 33% of those that lose a patent race abandon their patent (compared to ~12% for those that win). The remaining 67% of firms that *lose* a patent race *still end up with a patent*, but that patent will have narrower scope. That is, they will have protection over a smaller area in 'technology space,' or their protection in an area will be less strong.

In a separate article, we report that ~70% of patents from a random sample have their scope narrowed during prosecution (as judged by patent lawyers and metrics we developed). (Kuhn, Roin and Thompson 2016). We apply this same metric in this article to look at patent scope narrowing in patent races. Because patent races are (by definition) a sign of a contested technology area, and thus of a high amount of potential conflict in patent claims, it is not surprising that we find that above-average patent narrowing occurs for *both the winning and losing* patent, and that those that lose a patent race

are particularly hard hit. Figure 1 shows that this is true. Whereas 72% of all patents get narrowed during the prosecution process, 91% of winning patents ('leaders') and 95% of losing patents ('followers') get narrowed.

----- Insert Figure 1 about here -----

The greater narrowing of patent twins suggests that patent races happen disproportionately in areas crowded with competing patent claims. This might be driven by the market structure in these areas, or perhaps by the type of technology. More interesting for theory, however, is that it might indicate that patents in races reflect smaller innovative steps. If so, they would be 'closer' to existing patents and thus would be more likely to be narrowed by them.

3.2 Winning a Patent Race

Patent races also diverge from theory in the mechanics of how they are won. Instead of a universally-acknowledged victory of a known benefit, there is typically no announcement of victory, and trailing firms may continue to race even after the race has been won. Moreover, even once the victor is known, the spoils of that victory (patent protection) are often unclear until years afterwards.

The USPTO makes the existence and content of a patent application public to the world no sooner than 18 months after its earliest priority date, which is typically its filing date. This delay creates a window of time in which competing firms will be unaware that they have lost the race and thus may continue racing. Such firms may then try to claim patent protection for ideas that have already been submitted to the patent office by the earlier-filing firm. In such cases, the patent office will ultimately reject the claims that duplicate coverage.

The determination of whether a patent gets issued at all, and if so what claim scope is approved, typically takes 3 years to complete. (USPTO 2015). As such, it is only at this point that patent racers

know with some certainty the value of the patent protection prize that they have won. (See, e.g., Gans, Hsu and Stern (2008)).

Because patent races differ from theory in having these ambiguous endings, it is common for two patent racers to file for patent protection on the same idea within a short period of time, in a more-modern, firm-version of what Alexander Graham Bell and Elisha Gray did 140 years ago. The administrative records from these pairs of applications are the empirical window through which we observe patent racing.

3.3 How the patent system lets us observe patent races

Our empirical analysis is grounded in the mechanics of patent examination, the process by which the USPTO determines whether to grant a patent application. Patent examination is often modeled as a binary outcome—a patent application is either granted or not. In practice, nothing could be further from the truth. Patent examination is an extremely complex and frequently lengthy process of procedure and negotiation. Both the patent applicant's choices and each of a series of decisions reached at the patent office exert considerable influence on not just whether a patent application is ultimately granted, but also the time required to come to a final decision as well as the scope and strength of the claims of a patent if one is issued.

At a high level, the lifecycle of a patent application is divided into pre-examination, examination, and issuance. During pre-examination, the patent application is classified by technology area and assigned to a patent examiner knowledgeable in the relevant field. The examination phase begins when the patent examiner has time to begin work on a new application. During the examination phase, the patent examiner evaluates the application to ensure that it complies with every requirement of U.S. patent law. The USPTO is required by U.S. law to grant an application for a patent unless it can find a reason to reject the application, but the requirements imposed on patent applicants are numerous and

complex. Reasons for a rejection can range from the administrative (e.g., a patent applicant has made a procedural mistake in the examination process) to the descriptive (e.g., a patent application is improperly formatted) to the substantive (e.g., a patent application is insufficiently inventive).

In this article we focus on substantive rejections. In particular, on the case when an examiner rejects one or more of the claims in a patent application on the grounds that there is prior art – that is, the invention that is being claimed had already been disclosed to the world. Prior art rejections are divided into two types—novelty and non-obviousness.⁴ In a novelty rejection, the patent examiner argues that the claimed invention is completely described in a prior art reference that precedes the filing date of the application. In a non-obviousness rejection, the patent examiner argues that even though the patent claim is novel, it is an obvious step over some combination of prior art references. We focus on novelty rejections, rather than obviousness ones, for reasons that are outlined below.

Under U.S. law, a patent claim is anticipated (i.e. is not novel and therefore not patentable) if a single prior art reference discloses all of the features recited in the claim. Federal courts have interpreted the novelty requirement as meaning that “[a] claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference.”⁵ Further, U.S. law requires that “[t]he identical invention ... be shown in as complete detail as is contained in the ... claim” in order to support a novelty rejection. Thus, a novelty rejection specifically highlights the appearance of an invention that was earlier disclosed, but which a later-filed application has tried to claim. The following example helps to clarify the nature of novelty rejections.

On Date September 1980 Daniel and Vicky Lyer, filed an application for “a wearer adjustable article carrying harness assembly adapted to releasably hold an article, such as a camera, against the

⁴ These are also known as 102 and 103 rejections, based on the sections of the patent law that describe them.

⁵ The Manual of Patent Examination Policy (MPEP) section 2131. The MPEP is a lengthy manual written by the USPTO that provides practical guidance to patent attorneys and patent examiners regarding the patent examination process. It aggregates and interprets the vast collection of U.S. code, judicial decisions, administrative rules, and administrative decisions that collectively govern the U.S. patent application system.

body of the wearer in a carrying position while not in use...” In March 1982 this patent, 4,320,863 was granted.

In October 2005, Nicholas Woodman, the founder and CEO of the action camera manufacturer GoPro, filed an application for “a configurable wrist or arm worn camera harness...that facilitates convenient carry, access, and secure use of a camera during participation in a physical activity...”. In that application, Woodman made 47 claims for patent protection to cover his invention. However, the patent examiner rejected much of Woodman’s claim coverage for lack of novelty, saying “Claims 1-3, 5-12 and 30-45 are rejected under 35 U.S.C. 102(b) as being clearly anticipated by Lyer et al.” As a result, Woodman dropped all these claims. Nevertheless, Woodman still ended up with patent 6,955,484 which covered the 17 claims that remained after the novelty (and other) rejections.

This example highlights that patent applications with rejected claims can still produce an issued patent, a fact which will turn out to be the norm for patent races. This example also shows some of the standard phrases that can be used to identify novelty rejections (“rejected under 35 U.S.C. 102(b)”). Novelty rejections are analytically useful for us for several reasons. Firstly, the examiner indicates the earlier disclosure which described the invention. When that disclosure is a patent, this provides a direct indication that the examiner believes both patents cover the same invention. Under the right conditions, such a rejection can be indicative of a patent race. Moreover, the existence of the potential patent race was determined by a third party expert evaluating the claims contemporaneously as part of their job. This determination saves us, as econometricians, from having to estimate the existence of the race from indirect measures or from after-the-fact retrospection about technological similarity.

Importantly for our identification strategy, a patent applicant who receives a novelty rejection only has a limited set of options: they can abandon the application, argue that a rejection is improper, or amend it. Those that amend *cannot add new matter to the application*—the application’s inventive

contribution is fixed at the time of filing. However, they can narrow their claims, so that they don't overlap with the earlier work, and pursue the issuance of this more-modest patent.

3.4 Data

When a patent examiner evaluates a patent application, they prepare various documents and post them to the USPTO's Patent Application Information Retrieval (PAIR) system. For patent applications that have not yet been published, these documents are accessible only to the patent applicant and the applicant's designated representatives. All documents generated during the examination of the application become publicly accessible via the same system when an application is published as a patent publication or issued patent, typically 18 months after filing. Documents for patent applications that are abandoned before being published remain private forever, which severely limits the analysis that can be done on these patents. Thus, for most of our analysis we focus on instances where both the leader and the follower both issued as patents, which in fact make up the majority of all patent races in our data.

Through an agreement with the USPTO, Google scraped the PAIR system and aggregated publically available patent examination documents in ZIP files that are accessible through Google's website.⁶ The full data includes PAIR data for 5,027,882 patent applications filed between April 23, 1984, and June 10, 2013. Data availability considerations restrict this further to a sample of patents issued between January 1, 1998 and February 7, 2012.

----- **Insert Table 1 about here** -----

We extracted from the PAIR data 4,231,280 Office Action documents. Office Actions are published in a PDF image format that renders their text inaccessible in a digital format. Accordingly, we

⁶ <https://www.google.com/googlebooks/uspto-patents-pair.html>

first applied the optical character recognition library Tesseract to recover the text of Office Actions in a digital format. Because the application of the Tesseract library to so many documents required many years of computation time, we divided the computation among 10,000 CPUs accessed via the Amazon Elastic Compute Cloud (Amazon EC2) service.

We searched the text of each patent application's Office Actions to identify patent rejections. Patent rejections are not listed in an Office Actions in an easily accessible format or location, but rather are embedded throughout the text of the document. When stating claim rejections, however, patent examiners employ formal legal language to clearly indicate the legal status of each claim. For example, a patent examiner issuing a novelty rejection ("35 U.S.C. 102") would include a sentence such as "Claims 1-20 are rejected under 35 U.S.C. 102 as anticipated by U.S. Patent No. 1,234,567 by Smith et al." In the paragraphs that follow, the patent examiner would support the rejection with documentary evidence comparing the claims of the application to the disclosure of Smith. We applied regular expressions to the text of each Office Action to identify rejection sentences, determine the statute used to support the rejection, and extract the patent or patent publication number of the reference used to support the rejection.

We linked each patent in the dataset with bibliographic data collected from other publically available documents issued by the USPTO. Each week the USPTO issues an XML document that include a wealth of tabularized data describing each patent application published or issued during the week.⁷ From this data we extract such fields as the patent filing date, the patent issue date, and all patent citations made by each issued patent.

⁷ <https://www.google.com/googlebooks/uspto-patents-applications-biblio.html>

We limit our sample to those pairs where neither patent application claims priority to a previously-filed application to avoid evaluating continuations or first-foreign-filed patents as potential patent races, both of which could imply disclosure patterns different than our 18-month window.⁸

Each issued patent is also linked with patent maintenance fee payment data parsed from the USPTO's maintenance fee payment data file.⁹ To maintain a patent in force, the patentee must pay a fee to the USPTO by each of 3.5, 7.5, and 11.5 years following the issue date of the patent. The maintenance fee payment data file lists for every patent whether or not the patentee made each of these payments. Maintenance fee payment data can provide some indication as to the private value of a patent, as a patentee would only pay a maintenance fee if the patent was more valuable than the fee amount.

We retrieved information on publicly traded firms via the Compustat database. Linkages between patents and firms were provided by Douglas Hanley, who parsed USPTO patent applicant data to associate each patent filed by an institution with an institution identifier, and who linked patents issued to publicly traded companies with a Compustat identifier. (Hanley 2015). Firms that are not present in Compustat are dropped from any sample that considers firm-level characteristics.

3.5 Defining Patent Races

3.5.1 *Patent Twins*

To support our analysis into the effects of winning or losing a patent race, we introduce the notion of a “patent twin”. By analogy to Bikard’s ‘paper twins’, a patent twin refers to two patents filed at roughly the same time, covering the same discovery. (Bikard 2015). We first present an example, and then codify this definition.

⁸ For example, a continuation would indicate that the invention was first described in an earlier patent, and thus that the correct timing of revelation of the described information would be the publication of that earlier patent.

⁹ <https://www.google.com/googlebooks/uspto-patents-maintenance-fees.html>

In April 2005, Cardiac Pacemakers Inc. filed U.S. Patent No. 7,640,057, with the title: “Methods of providing neural markers for sensed autonomic nervous system activity”. Eight months later, Pacesetter Inc. filed U.S. Patent No. 7,672,722, with title: “Hardware-based state machine for use in discriminating near field signals from far field signals for use in an implantable cardiac stimulation device”. Both applications sought protection for techniques for identifying cardiac events, as well as other aspects of their discoveries. Ultimately both were granted patents, but some of the scope that Pacesetter Inc.’s wanted was refused by the patent examiner on the basis of the earlier patent submission by Cardiac Pacemakers.

We define patent twins as pairs of patent applications that exhibit two characteristics. First, the two applications include sufficient technological overlap to claim patent protection for the same invention. In practice we will operationalize this by looking for cases when an examiner has rejected claims by the later-filed application (which we term the “follower”) on the basis of those claims already being present in an earlier-filed one (which we call the “leader”) – that is, through novelty rejections. Second, the difference in time between the filing of the leader and the follower is sufficiently small that it would be unlikely for the applicants to know the content (and likely even the existence) of the others’ patent application. In practice we will utilize two time windows for this analysis:

- (1) “All Twins 18”: Applications within 18 months of each other
- (2) “All Twins 6”: Applications within 6 months of each other

The 18-month window reflects time between when the USPTO would receive and make known to the world the contents of the leader application. We find this to be the most-compelling high-level definition for a Twin.

In Section 6, however, we will infer a causal interpretation to our estimates for winning a patent race by arguing (and testing) that with a small enough time window, the difference between the ‘leader’ and ‘follower’ in a patent twin is as good as random. That is, we will make a regression discontinuity

argument. For this, we will want as narrow a bandwidth as possible. The 6-month window was chosen because it shows excellent covariate balance between leaders and followers, and still provides a sufficient sample size for our analysis.

In order to look at firm-level outcomes, there will also be instances where we narrow to samples where both twin patents were filed by public firms. To parallel the definitions above these will be referred to as “Compustat 18” and “Compustat 6.”

3.5.2 Taking our definition to the data

The set of patent applications linked by novelty rejections is large and includes many instances where the filings are years apart. Such a situation might occur, for instance, if the follower is simply at a much earlier stage of technological development than the leader and is unaware of the state of the prior art, or could not afford to do a prior art search. We do not regard such situations as patent races because there is no active competition for the idea – one has already completed the race, and that information is public knowledge prior to the other filing a patent application. To identify actual competition, we narrow to application pairs, linked by novelty rejections, with “sufficiently small” time between their filing dates. Figure 2 shows the distribution of these differences in dates. Figure 3 is a sub-figure that focuses on the 0-2 year window after filing.

----- Insert Figure 2 and Figure 3 about here -----

For our description of the phenomenon of patent racing, we define ‘sufficiently small’ as the window of secrecy provided to patent applicants under U.S. law. Each patent application is secret upon filing. However, by default, every application filed on or after January 1, 2001 is made public 18 months after filing, so an applicant for a later-filed application cannot search for that earlier-filed application if it has been filed less than 18 months prior. Our definition for a patent race is therefore: two firms filing patents covering the same invention within a window where the application of the other would not

have been publically known.¹⁰ Later, to do a regression discontinuity estimation, we will narrow the window even further to 6 months. However, for the remainder of this section we will focus on the 18-month window.

Our methodology purports to find races for a common technology. Under such circumstances we should expect a great deal of similarity between the patents (although not perfect similarity, because each will also contain non-overlapping technology). We can test our success at this by comparing the textual similarity of technology descriptions for patent twins as compared with a reference group (say, patents in the same class). If our method is working, the technology descriptions of patent twins should be much closer. To test this, we use the textual similarity measure developed by Younge and Kuhn (2015), which treats each patent specification (description of the technology) as a bag-of-words, and looks for similarity in the usage of non-common words. Figure 4 shows the distribution of these similarity scores.

----- Insert Figure 4 about here -----

We find that, as predicted, patent twins are much more similar than are patents in the same main class. In addition to being relatively more similar, they are also very similar in absolute terms. For reference, less than 3% of patent pairs have a similarity score greater than 0%, and only .54% have a similarity score of 25% or greater.¹¹

Thus we conclude that our method for identifying patent races, in addition to being based on criteria that mean it *should* correspond to patent races, can be shown to identify instances of patenting that are both close in time and close in the technologies where patent protection is sought. Thus, we conclude that this definition of patent twins allows us to observe patent racing.

¹⁰ Notice, this is a *de facto* definition. Presumably there are firms within this group that were racing for patent coverage, but were unaware of it.

¹¹ Note: because the background rate of similarity is 10% or lower, this graph renormalizes, taking 10% similarity as the minimum measure (score 0).

4 Where do patent races happen? Who races?

Overall we find that 10-11% of patents are in patent races. This section describes where these races occur, and which firms do the racing.

4.1 Technology areas

Figure 5, Figure 6, and Figure 7 show the technology areas where patent races occur. Here, patents are divided by their classification into USPTO technology centers.¹² Figure 5 shows, in absolute numbers across all the years of our data, that the most patent races happen in the semiconductor and communications divisions, and fewest happen in the biotechnology division. However, Figure 6 shows that, as a share of all patents in those areas, Computers has the largest share (15.7%), with Communications next (13%), and then Semiconductors (10.1%). Biotechnology also has the smallest share, with 4.9%.¹³ Finally, Figure 7 recapitulates Figure 6, but instead of reporting the share of *all* patents in that technology center, it considers those in patent races as a share of those that had novelty rejections (at any point in time). This confirms that it is not just a prevalence of novelty rejections in these areas that are driving these results, but instead a higher rate of contemporaneous patenting of the same technologies – that is, more patent racing.

----- Insert Figure 5, and Figure 6, and Figure 7 about here -----

The reason behind these differences is unclear, and calls for more research. Plausible explanations could include (i) technology lifecycles, (ii) patent effectiveness, or (iii) market structure.

¹² These eight divisions serve as administrative units into which the USPTO organizes its patent examiners.

¹³ All the share calculations are done on a single year, 2005, to ensure the least possible impact from truncation of our data (discussed later).

These patterns could reflect technology lifecycles if areas like computers and communications require smaller inventive steps to get patent coverage. If so, these could lead to a greater density of patents at the frontier of knowledge and more competition (i.e. patent races) for the open spaces.

Alternatively, differences in patent effectiveness could cause this distribution of patent races. Biotechnology and Mechanical are technology areas thought to have a stronger ability to enforce patents, whereas Computers and Communications are thought of as having a weaker ability to do so. (See, e.g., Cohen, Nelson and Walsh (2000)). If firms deem themselves to have a significant chance of losing the race, they might choose not to race in the first place if they think that they would have little ability to continue the work in the face of a strong competitor patent.

Lastly, the distribution of patent races might reflect market structure. For instance, there could be fewer patent races in some areas simply because there are fewer (or less direct) competitors. This could even be an endogenous result of the cost of patent racing if patent racing in areas like pharmaceuticals is so costly that it has led to consolidation in industry structure.

4.2 Firms

How are firms that patent race different from other firms that patent? To answer these questions, we narrow our data to publically-traded companies filing applications within 18-months of each other (that is, to our “Compustat 18” dataset). Linking patents with publicly traded firms is possible by using a correspondence generously provided by Douglas Hanley that allows us to match patents to their Compustat firm.

The restriction to public companies also implies narrowing our sample to instances where both the applications of the leader *and the follower* issue as patents because the USPTO does not require applicants to provide the relevant information needed to match patents to firms until patent issuance. As a result of this data constraint, this analysis misses patent races where the follower abandons their

application. However, high value patents are unlikely to be abandoned (Kuhn, Roin and Thompson 2016), so any selection introduced is likely to affect mostly low-value patents.

Figure 8 illustrates the share of our data in these technology centers for each of these samples, again using 2005 as the representative year. It shows that narrowing to public-firm patent races leads to a slight over-representation of computers and communications and a slight under-representation of biotechnology, e-commerce, and mechanical, as compared with the full sample.

----- Insert Figure 8 about here -----

Table 2 and Table 3 compare firm characteristics for those associated with twin patents to those associated with patents in general.¹⁴ Not surprisingly, firms in patent races are larger and do more R&D, a difference that exists even on a per employee basis, where racing firms have slightly higher sales and R&D per employee. As Table 3 shows, they are also much more prolific patenters, both in terms of filed and issued patents.

----- Insert Table 2a, Table 2b, and Table 3 about here -----

5 Methodology for getting a causal estimate of the effect of winning a patent race

5.1 Regression Discontinuity Approach

As previously outlined, the firm-level sample considered so far has been “Compustat 18,” which represents patents linked by a novelty rejection that were filed within 18 months of each other. Within this 18-month window firms may race for patent coverage without having the earliest-filed patent

¹⁴ Because a firm can have both leader and follower patents in our sample, they may be present in both sets of summary statistics. The values presented should thus be thought of as patent-weighted averages across firms.

application revealed. That does not mean, however, that within this window firms are the same. One might imagine that patenting leaders might be systematically different than patenting followers, for example they might do more R&D or work with better lawyers. And, indeed, tests of covariate balance at the 18-month window show a lack of balance at this point.

This creates a challenge for one of our empirical goals, estimating the causal effect of winning a patent race, because it indicates selection. To address this concern, we adopt a regression discontinuity approach. Imagine that two patents covering the same invention get filed at the patent office on the same day. Absent intrusive (and possibly illegal) corporate espionage, the exact timing for each application on that day is likely to be as-good-as-random. Under these conditions, small changes in filing time (the running variable) will introduce large changes in patent protection (the treatment) – and thus differences in firm outcomes should be interpreted as arising causally from winning the patent race. Figure 9 shows this schematically:

----- Insert Figure 9 about here -----

Our estimation focuses on the effect of this loss of patent coverage, due to losing the patent race.

Unfortunately for the empiricist, there is a dearth of same-day, same-invention filings at the patent office, so we must consider a larger bandwidth for the comparison. The larger the bandwidth we choose, the more data we get to consider, but the greater the chance that other factors confound our estimates.

We choose a 6-month window as our bandwidth. We feel that this is a plausibly short window wherein small random events (e.g. illness of a patent attorney) could lead to changes in the order of filing. If this assertion is correct, that is, if filing order within these six months is as-good-as-random, we should see good covariate balance on pre-treatment variables between those filing patents that will

ultimately win the patent race and those that will ultimately lose it. We test this in several ways, as it is central to our causal interpretation of the effect of winning a patent race.

5.2 Testing whether those that win and lose patent races look similar, pre-treatment

Testing for covariate balance on the *full set* of data is only possible in a limited way because some leader and follower patent applications are abandoned, and the USPTO publishes very little data on abandoned applications. As a result, the balance check on the full set of data focuses on the small number of covariates that are both available and of some interest.

----- Insert Table 4, Figure 10, and Figure 11 about here -----

Although all the differences in Table 4 are statistically significant, the significance is largely a product of our large sample size. In terms of economic importance, we observe limited differences. Slightly more patent *leaders* are small entities, 20.9% vs 18.5% for followers. The length of their patent titles are almost identical—61.8 vs 60.9 characters. We are also able to observe some characteristics about the entity that files the patent (which may be the firm, but is commonly their attorney) by looking at the USPTO customer numbers associated with leader and follower patents. We find that slightly fewer (87.5%) of leaders have USPTO customer numbers, vs 88.3% of followers, although their customer numbers are associated with slightly higher numbers of patents. Because you must have a USPTO number to file electronically, these differences probably reflect small changes in the timing of different attorneys going digital. To emphasize how similar leaders and followers are at the time of filing, Figure 10 and Figure 11 show the full distribution of title length and the number of patents associated with a customer number. Thus, we conclude that although there is evidence for small differences between the leaders and followers at the time of filing, the potential for these to bias our estimates is limited.

Because balance on the full set of data can only be looked at with a limited number (and relevance) of covariates, we also look for balance across a much larger set of variables, *post-abandonment*.

Because, not surprisingly, followers are almost twice as likely to abandon their patent application (32.8%) as leaders (17.9%), the following tests are conditional on issuance, and the effect of differential abandonment rates must be kept in mind in interpreting the results.

Table 5 compares Student's T-test values for firm-level financial characteristics of leaders and followers in the sample. Because, by construction, follower patents are filed later than leader patents in our sample, we analyze firm variables as of the filing year *of the leader* for both.

The firms associated with leaders and followers look very similar, both in terms of the magnitudes of their characteristics and the lack of statistical significance between them. They exhibit no significant differences on revenue, assets, number of employees, or cost of goods sold as of the year in which the leader application was filed. R&D levels are statistically different at a threshold of 10%, but the magnitude of this effect is incredibly small.

----- **Insert Table 5 about here** -----

Beyond firm financial performance, we also look for difference in previous patenting behavior. Table 6 performs t-tests for differences in means for these firm-level patenting characteristics. We find small differences in the number of patents for each, but in opposite direction, which leaders having more issued patents and followers having more applications.

----- **Insert Table 6 about here** -----

Finally, Table 7 compares the characteristics of the leader and follower *patents* in the sample – that is, the applications that result in winning or losing the patent race. Leader and follower patents exhibit no significant difference in the number of claims at filing. However, leaders do have slightly longer first claims on average. Because longer claims are associated with *narrower* (i.e. less good) patent scope (Kuhn, Roin and Thompson 2016), this difference suggests that the follower patents are broader. We suspect that this difference reflects the higher abandonment rate for follower patents, where those

follower patents with the smallest scope would be more likely to be abandoned, leaving a pool with slightly *broader* scope remaining in the sample.

Table 7 also shows that leader patents have slightly more inventors listed on their patents. The (unlogged) economic magnitude of this effect is 0.17 inventors. This difference could also be a result of the differential abandonment rates, or could indicate that firms that devote more inventive resources to a project are (slightly) more likely to win patent races.¹⁵ Figure 12 shows the distribution of inventors, and Figure 13 shows the dose-response curve for this effect.

----- Insert Table 7 about here -----

----- Insert Figure 12 and Figure 13 about here -----

Taken together, Table 4, Table 5, Table 6, and Table 7 suggest that firms that win patent races in our sample look very similar, both on overall firm observables and patenting behavior, to those that lose patent races. Those small differences that remain in our 6-month sample, we will control for explicitly in our regressions. Inasmuch as there remain any differences in samples induced by the differential abandonment rates, they will work against our finding results for our key findings (e.g. follow-on research).

The finding that firms that win patent races are similar to those that lose them may seem counter-intuitive. In the broader world, we expect that firms with more resources, technical expertise, etc. would be more likely to win patent races. Many of such races, however, should be sufficiently lop-sided that the less able firm fails to file a comparable invention within 6 months of their abler competitor. In such

¹⁵ Interestingly, there is suggestive evidence of firms being in a mixed strategy equilibrium where they allocate inventor resources in a way that mimics an auction strategy. That is, where each allocates additional resources to raise their chances of winning the patent race, but where they also a desire to win by as little as possible because hiring inventors is expensive. (Bell et al. 2015). Such a mixed strategy model would show victories and loses in patent races over many different numbers of inventors. Figure 14 shows that this is exactly the case for IBM.

instances, these races would not appear in our data. By construction, our regression discontinuity approach focuses on patent races that are closer, and thus where we would expect this balance.

Because of our regression discontinuity approach, our estimates should be interpreted as local average treatment effects, centered on patent races where the competition between firms is credible enough that neither firm drops out before filing a patent application. Within this group, we observe good balance on observables, which we interpret as support for our as-good-as-random assumption regarding the timing of patent filings.

5.3 Firm Portfolio approaches to Patent Racing

A key implication of having as-good-as-random allocation of patent races winners in the 6-month window is that we should not observe firms systematically winning or losing patent races. That is, if IBM were to win all its patent races in this window, that would argue strongly against our approach for getting causality. We test this prediction explicitly on the top 100 firms in our patent racing sample. If the as-good-as-random assumption is correct, the distribution of firm victories of patent races should reflect repeatedly drawing [win, lose] with 50:50 odds, a number of times equal to the number of races they participate in. That is, the modal firm should win 50% of their patent races and there should be a binomial distribution around it. Figure 15 shows the empirical distribution of the victories of the top 100 firms in the solid line, and the corresponding binomial distribution, as simulated with 10,000 replications, in the dotted line.

----- Insert Figure 15 about here -----

The simulated curve replicates remarkably well the variation in firms' success at winning patent races. The only portion that it fails to replicate are the extreme tails of the distribution, suggesting that in the real world firms win or lose patent races *less* systematically than a binomial model would imply. We think that this would likely reflect competitive pressures in the market where firms that are

systematically losing invest more (or drop out), whereas those that are systematically winning have competitors that invest more (or drop out). These dynamics would produce the lesser dispersion seen in Figure 15.

This test suggests that, conditional on a patent race being close enough that both firms file within 6 months of each other, most of the remaining difference is chance, providing support for our identifying assumption.

6 Causal Effect of Winning a Patent Race

We examine the causal effect of winning a patent race on firm behaviour at three levels:

- (i) the extent of patent protection actually obtained and maintained
- (ii) the changes in firm-level outcomes (e.g. R&D), *conditional on (i)*
- (iii) the follow-on research done by firms (as measured by patenting), *conditional on (i)*

Throughout this section, key results will be shown in two ways: (i) unconditional means, and (2) estimates adjusted for any remaining covariate imbalance via the following regression form (where the level of the analysis will depend on the estimand):

$$Outcome = \alpha + \beta Leader + \gamma_i Control_i + \epsilon$$

Here the coefficient of interest is β , the effect of winning a patent race (i.e. being the “Leader”). Control variables include filing year fixed effects, Technology Center fixed effects, and controls for variables with imperfect balance (notably: the number of inventors). Standard errors are clustered at the firm x year level.

6.1 Patent protection obtained and maintained

By virtue of losing a patent race, we expect that a follower firm will be more likely to (1) abandon their patent (because it will have less value absent the contested claims), (2) get less scope in their final patent protection for ones that do issue,¹⁶ and (3) take longer in the prosecution of their patent. We also expect that, conditional on issuance, follower firms would be less likely to keep their patents in force by paying the maintenance fees required in the third and seventh year after issuance (regressions 4 and 5). Figure 16 presents the unconditional means on these variables and Table 8 presents the regression results.

----- Insert Figure 16 and Table 8 about here -----

Column 1 confirms that 13.3%*** fewer leader patents are abandoned than follower patents. For those patents do issue, column 2 shows that follower patents get their scope narrowed more during patent protection. Follower patents have, on average, 15*** more words added to the first independent claim. As discussed by Kuhn, Roin and Thompson (2016), the addition of words to a claim typically narrows the scope of patent protection. Thus, not surprisingly, follower patents are narrowed more during examination.

Column 3 shows that the follower patents take 0.7*** years longer to issue. This delay is presumably due to additional back and forth between the examiner and the follower firm's lawyers to try to revise the claims in light of the conflict, and likely implies both higher legal costs for the follower firm as well as less ability to enforce the patent, as the claims are not issued until later.

Finally, columns 4 and 5 show the renewal rates of patents at the first and second maintenance payments.¹⁷ There is no significant difference in renewal rates at the first maintenance period, with both leader and followers renewing about 92% of the time – perhaps because this is the lowest fee, at

¹⁶ This result is partially a mechanical result for how the patent twins are identified.

¹⁷ For the second maintenance period, this effect is *not* conditioned on the first payment. That is, it is a cumulative difference.

\$1,600. By the second maintenance fee, this difference has grown, and leader patents are 3.8pp* more likely to renew – although again against a backdrop of very high renewal rates. Indeed, it is telling that the renewal rate for all patents is 85.3% for the first maintenance fee and 64.4% for the second, so the rates are significantly higher in patent races.

Overall, we see strong and consistent effects that losing a patent race results in a less valuable potential patent and thus that applicants are more likely to abandon that patent, either during the prosecution process or later.

6.2 Effect on Firm Level outcomes

As already shown in Figure 15 and discussed in section 5.3 larger firms win and lose patent races with roughly equal frequency. As such we should expect little-to-no-difference in firm-level outcomes from winning a patent race. Table 9 confirms that this is indeed the case:

----- **Insert Table 9 about here** -----

This precise null result is a large departure from theory, which predicts large effects on R&D from winning a patent race (see, e.g., Fudenberg, Gilbert, Stiglitz and Tirole (1983)), because those models typically assume a single project per firm. In reality, our analysis confirms that they take a portfolio approach. Thus, to investigate the effects predicted by theory models we must narrow our analysis to project-level outcomes.

6.3 Follow-on research, conditional on the issuance of both patents

Like much of the literature, we will investigate project-level research using citations by later patents as a way to infer related future research. However, we will also validate our citation results using a separate, textual similarity-based measure. We also analyse citations differently than other articles to mitigate some of the weaknesses inherent in citation-based measures.

The correct interpretation of patent citations is unclear, with many authors offering very different versions (see Kuhn and Younge (2016), for a more detailed discussion). For our part, we believe that it is credible that follow-on citations can sometimes reflect follow-on research in an area, but that they are equally likely to reflect differences in the citation behaviour in patents. That is, differences can reflect either a change in the propensity to do further work, or changes in the propensity to cite. In some cases, we will be unable to distinguish these. In other cases, we will make arguments based on the institutions around patenting that argue for one interpretation or the other.

We consider two different measures of citations, and three different sources of citations in our analysis. The measures we consider are: the total number of citations accruing to a patent and the percentage of patents that accrue at least one citation.¹⁸ We include the latter measure as a robustness check, to ensure that highly-cited patents do not mask differential changes to lightly-cited ones. We consider three sources for these citations: which we call “self cite”, “twin cite”, and “other cite”). All three sources reflect citations from later patents, but we segment them into those by the same firm (“self cite”), by the other firm in the patent race (“twin cite”), and by other firms not in the patent race (“other cite”). Table 10 shows our estimates for these effects.

----- Insert Figure 17 and Table 10 about here -----

From columns (1) and (5) it is clear that leader patents get cited more, to a level that is both economically and statistically significant. Most of this effect derives from others firms citing the leader patent more than the follower patent, columns (2) and (6). Although it is possible that this effect reflects the larger patent scope first enumerated by the leader patent, we find it at least as credible that it reflects examiners becoming more reliant on the leader patent for demarcating future claims. That is, examiners become familiar with the patent, and thus use it for rejections more frequently, leading to it

¹⁸ Methodological note: because a novelty rejection to the follower patent mechanically adds a citation to the leader patent, we exclude it from this and all analyses of citations.

being inserted into the citation record more often. In theory, these interpretations could be disambiguated by looking at whether the examiner or the applicant added these citations. But, because our data is not broken apart this way, we do not attempt separate these effects.

Columns (3) and (7) show that the company that wins the patent race is more likely to cite its own leader patent in the future than is the other firm to cite its follower patent. Because both firms are already clearly aware of their own patents, we feel that it is more straight-forward to interpret this difference as indicating that the leader firm pursues more follow-on innovation by virtue of winning the patent race. Under this interpretation, the effect of winning the patent race is large, leading firms to be 11% more likely to do follow-on research on their patent, and to do 14% more follow-on research in that area.

Columns (4) and (8) focus on the citations by the competing firm in the patent race. They show that follower firms cite the leader patent 70% more than the leader firm cites the follower patent. This is unsurprising as the legal process of rejection makes the follower firm very familiar with the leader patent, but the converse is not true. Moreover, patent applicants are subject to a statutory duty of disclosure that requires them to acknowledge prior art they are familiar with. This duty frequently leads patent applicants to copy patent citations, such as references used to support a rejection, across groups of related patents. (Kuhn and Younge 2016). Thus, we believe the increase in citations by the competing firm are most likely because of a change in citation behavior, rather than follow-on research behavior.

Having shown that firms that win a patent race do more follow-on research, we also test whether that follow-on research is more similar. We test this proposition by directly applying a measure of the similarity between the patent twins and the follow-on work done by each firm. We expect that firms winning patent races will have follow-on patents 'close' to the leader patent, whereas firms that lose the patent race will need to invent around, and thus will be 'further' from their follower patent (or the leader patent). We measure similarity based on a cosine distance metric, calculated on the words in

the description (aka specification) of each invention (see Younge and Kuhn (2015) for a detailed discussion of this metric). High similarity indicates that similar content words are used to describe both inventions, whereas low similarity indicates the usage of very different content words.

----- Insert Figure 18 about here -----

Figure 18 shows that, indeed, the content of the follow-on patenting by the leader is 'closer' to the patent that won the pace race than the follower's work is to their patent that lost the patent race. This relative decrease in proximity by followers suggests that many are abandoning that direct area of research and instead moving to adjacent areas. The right-most part of this graph, showing very-high similarity reflects patent divisionals or continuations. Not surprisingly, the leader firm files more divisionals / continuations, reflecting continuing interest in the specific area of the patent race win.

Thus we conclude that, following a patent race victory, firms do more follow-on research. And that the follow-on research that they do is likely to be more similar.

7 Ten Implications for Theory Models

Collecting our findings from across the article on the implications for patent racing theory models, we find that:

1. Patent races frequently persist through to a mature point of technology development where multiple racers can file patents.
2. Patent racing is rarely 'winner-take-all', in that usually both winners and losers get patent coverage.
3. Knowledge of who wins a patent race is revealed much after-the-fact.
4. The 'prize' of a patent race is unknown until even later, when the patent scope is determined.

5. There is evidence suggestive of firms allocating inventive resources as a mixed strategy auction.
6. Large firms participate in many patent races, and thus their efforts should be thought of as a portfolio-based approach.
7. Because large firms have a portfolio of research, the effects of winning a patent race typically accrue at a project level rather than at a firm level.
8. Patent races happen in disproportionately 'crowded' technological areas.
9. Patent races happen disproportionately in areas characterized by weaker patent effectiveness.
10. Winning a patent race causes firms to do more follow-on research, and that research is 'closer' to the technological area of their patent race win.

8 Conclusion

This article provides a first large-scale look at how patent races happen in the real world. Using detailed administrative data from patent examinations and a novel approach for identifying patent races, we (i) establish some initial key facts about who races and where, and (ii) test the assumptions and conclusions that underpin theoretical models of patent racing.

It addresses theoretical research in patent racing that highlights the complex trade-offs inherent in setting patent policy to reach socially optimal outcomes for patent races. (Judd, Schmedders and Yeltekin 2012, Wright 1983). Crucially, these models have depended on empirical parameters that have been unobservable. This article provides a way for such theory models to be in dialogue with empirical work to help delineate this space.

Contrary to the conclusions of some models, we find that patent racing is common, with 10-11% of all patents in races, with large firms being more likely to race. In areas of technology that are usually thought of as having weaker patents (computers and communications) patent races are more common, whereas in areas with stronger patents (biotechnology and mechanical) they are less frequent. We also observe that patents in patent races are much more likely than the average patent to have their scope narrowed during prosecution. From this, we infer that patent races are more likely to occur in areas with a high-density of intellectual property protection per unit of “technology space”. This could, for example, occur if firms in patent races were taking smaller innovative steps in order to be the first to file, and thus remained closer to the status quo, which one would expect to be more crowded with patents.

We also consider the causal effects of a patent race on those who win (the “leaders”) and those who lose (the “followers”). To estimate these effects we use a regression discontinuity approach, focusing on cases where both the leader and the follower file patents on the same invention in a 6-month window.

We find strong evidence of firms that lose the patent race being more likely to abandon their application and more likely to let the protection lapse by not making maintenance payments. Despite this evidence of poorer outcomes for followers, the more-stark interpretation of patent races in the theory literature as winner-take-all is clearly incorrect, with 67% of followers still getting patenting coverage from being in the race.

Most importantly, we show that winning a patent race affects follow-on innovation, both in direction and magnitude. Contrary to many models that assume that firms participate in a single race, we do not observe this effect at the firm level – and we show why one should not expect to observe such an effect because firms participate in a portfolio of races. Focusing down at the patent level, however,

we observe that leaders do significantly more follow-on research (14%) and that the work that they do is closer to the content of the original patent than is the follow-on research by followers.

Collectively, our results suggest that patent races are frequent and that they have important effects on the rate and direction of innovation by the firms. Moreover, our results provide empirical feedback on the assumptions and conclusions of patent racing models that can be used to model these better in the theory literature.

References

- Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen, "The lifecycle of inventors," (mimeo Harvard, 2015).
- Bikard, Michaël, "Peer-Based Knowledge Validation: A Hurdle to the Flow of Academic Science to Inventors," Available at SSRN 2333413, (2015).
- Cohen, Wesley M, Richard R Nelson, and John P Walsh, "Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not)," (National Bureau of Economic Research, 2000).
- Dasgupta, Partha, and Joseph Stiglitz, "Industrial structure and the nature of innovative activity," *The Economic Journal*, (1980), 266-293.
- Dasgupta, Partha, and Joseph Stiglitz, "Uncertainty, industrial structure, and the speed of R&D," *The Bell Journal of Economics*, (1980), 1-28.
- Doraszelski, Ulrich, "An R&D race with knowledge accumulation," *RAND Journal of economics*, (2003), 20-42.
- Fudenberg, Drew, Richard Gilbert, Joseph Stiglitz, and Jean Tirole, "Preemption, leapfrogging and competition in patent races," *European Economic Review*, 22 (1983), 3-31.
- Gambardella, Alfonso, Dietmar Harhoff, and Bart Verspagen, "The value of European patents," *European Management Review*, 5 (2008), 69-84.
- Gans, Joshua S, David H Hsu, and Scott Stern, "The impact of uncertain intellectual property rights on the market for ideas: Evidence from patent grant delays," *Management Science*, 54 (2008), 982-997.
- Gilbert, Richard J, and David MG Newbery, "Preemptive patenting and the persistence of monopoly," *The American Economic Review*, (1982), 514-526.
- Harris, Christopher, and John Vickers, "Perfect Equilibrium in a Model of a Race," *The Review of Economic Studies*, 52 (1985), 193-209.
- Judd, By Kenneth L, Karl Schmedders, and Şevin Yeltekin, "Optimal Rules for Patent Racing," *International Economic Review*, 53 (2012), 23-52.
- Kuhn, Jeffrey M, Benjamin N Roin, and Neil C Thompson, "Does Patent Scope Affect Cumulative Innovation? Evidence from Claim Narrowing During USPTO Examinations," 2016.
- Kuhn, Jeffrey M, and Kenneth A Younge, "Patent Citations: An Examination of the Data Generating Process," Available at SSRN 2714954, (2016).

Lee, Tom, and Louis L Wilde, "Market structure and innovation: a reformulation," *The Quarterly Journal of Economics*, (1980), 429-436.

Lerner, Josh, "An empirical exploration of a technology race," *The Rand Journal of Economics*, (1997), 228-247.

Loury, Glenn C, "Market structure and innovation," *The Quarterly Journal of Economics*, (1979), 395-410.

Reinganum, Jennifer F, "Dynamic games of innovation," *Journal of Economic Theory*, 25 (1981), 21-41.

Reinganum, Jennifer F, "A dynamic game of R and D: Patent protection and competitive behavior," *Econometrica: Journal of the Econometric Society*, (1982), 671-688.

United States Patent and Trademark Office (USPTO), *Traditional Total Pendency, Including RCEs, Oct. 2014 - Nov. 2016* (<http://www.uspto.gov/corda/dashboards/patents/kpis/kpiWithRCE.kpixml>,

Wright, Brian D, "The economics of invention incentives: Patents, prizes, and research contracts," *The American Economic Review*, 73 (1983), 691-707.

Younge, Kenneth A, and Jeffrey M Kuhn, "Patent-to-Patent Similarity: A Vector Space Model," Available at SSRN, (2015).

Zizzo, Daniel John, "Racing with uncertainty: a patent race experiment," *International Journal of Industrial Organization*, 20 (2002), 877-902.

Tables and Figures

Table 1. Summary stats.

Patents	5,027,882
Earliest filing date	April 23, 1984
Latest filing date	December 11, 2014
Claim rejections	18,335,938
Novelty rejections	1,265,038
Mean specification length (words)	9,274
Mean number of claims	17.11
Mean number of inventors	2.592

Table 2a. Comparison of the firm characteristics of twin patents to firm characteristics of all patents. Full sample 2005.

Variable	All publicly traded firms	Twin firms (18 month)	t	p	df	lower	upper	difference
Revenue	8.918	9.228	-6.783	0	2,543	-0.400	-0.221	0.311
RD expenditure	6.521	6.983	-13.070	0	2,556	-0.531	-0.393	0.462
Assets	9.276	9.551	-6.819	0	2,576	-0.355	-0.196	0.276
Employees	3.588	3.806	-6.281	0	2,482	-0.287	-0.150	0.218
Cost of goods sold	8.306	8.545	-4.995	0.00000	2,522	-0.333	-0.145	0.239

Table 2b. Comparison of the firm characteristics of twin patents to firm characteristics of all patents, on a per-employee basis.¹⁹

Variable	All publicly traded firms	Twin firms (18 month)	t	p	df	lower	upper	difference
Revenue / employee	298.0	310.9	-3.581	0.0003	2,605	-19.867	-5.809	12.838
RD / employee	37.7	41.0	-3.208	0.001	2,597	-5.377	-1.298	3.337
Assets / employee	474.0	462.7	1.324	0.185	2,544	-5.410	27.926	11.258
Cost of goods sold / employee	170.6	165.5	2.109	0.035	2,722	0.358	9.852	5.105
Employees (logged)	3.6	3.8	-6.281	0	2,482	-0.287	-0.150	0.218

Table 3. Comparison of the (logged) firm characteristics for twin patents to firm characteristics for all patents. Full sample 2005.²⁰

Variable	All firms	Twin firms (18 month)	t	p	df	lower	upper	difference
Prior apps. filed (2 yrs)	4.434	6.639	-55.194	0	2,425	-2.283	-2.126	2.205
Prior apps. filed (5 yrs)	4.991	7.280	-53.638	0	2,426	-2.373	-2.205	2.289
Prior pats. issued (2 yrs)	3.701	5.534	-31.013	0	2,336	-1.950	-1.718	1.834
Prior pats. issued (5 yrs)	4.173	6.111	-30.350	0	2,337	-2.063	-1.812	1.938

¹⁹ The first four variables are measured in thousands of dollars per employee. The final variable is measured in thousands of employees, logged. Full sample 2005.

²⁰ Note: because this is a firm level variable weighted by patent count, firms with large patent portfolios play a disproportionate role in these calculations.

Table 4. Application-level balance

Variable	Leaders	Followers	t	p	df	lower	upper	difference
Is small entity	0.209	0.185	9.727	0	95,845	0.020	0.030	0.025
Title Length	61.804	60.935	4.301	0.00002	96,217	0.473	1.264	0.868
Has Cust. No.	0.875	0.883	-3.791	0.0002	96,497	-0.012	-0.004	0.008
Cust. No. Patents (log)	7.979	7.860	9.040	0	85,250	0.093	0.144	0.118

Table 5. Financial characteristics (logged).

Variable	Leaders	Followers	t	p	df	lower	upper	difference
Revenue	9.280	9.357	-1.310	0.190	4,490	-0.192	0.038	0.077
RD expenditure	7.010	7.097	-1.849	0.065	4,496	-0.180	0.005	0.087
Assets	9.613	9.659	-0.876	0.381	4,529	-0.150	0.057	0.046
Employees	3.805	3.763	0.871	0.384	4,450	-0.052	0.134	0.041
Cost of goods sold	8.560	8.589	-0.454	0.650	4,507	-0.153	0.095	0.029

Table 6. Firm-level patent characteristics. Covariates logged.

Variable	Leaders	Followers	t	p	df	lower	upper	difference
Prior apps. filed (2 yrs)	6.678	6.752	-1.357	0.175	4,555	-0.181	0.033	0.074
Prior apps. filed (5 yrs)	7.385	7.430	-0.775	0.438	4,559	-0.159	0.069	0.045
Prior pats. issued (2 yrs)	5.733	5.588	1.796	0.073	4,565	-0.013	0.303	0.145
Prior pats. issued (5 yrs)	6.358	6.156	2.318	0.020	4,564	0.031	0.373	0.202

Table 7. Patent characteristics. Covariates logged.

Variable	Leaders	Followers	t	p	df	lower	upper	difference
Number of claims at filing	3.147	3.159	-0.623	0.533	3,490	-0.051	0.027	0.012
Length of first claim at filing (words)	4.671	4.464	12.021	0	3,607	0.174	0.242	0.208
Number of inventors	0.874	0.816	2.828	0.005	3,401	0.018	0.098	0.058

Table 8. Patent protection and maintenance results.²¹

	<i>Dependent variable:</i>				
	Abandonment (pp.)	Claim Narrowing (words)	Pendency (years)	Renewal Fee 1 (pp.)	Renewal Fee 2 (pp.)
	(1)	(2)	(3)	(4)	(5)
Is Leader	-0.133*** (0.023)	-14.866*** (3.556)	-0.693*** (0.049)	0.895 (0.823)	3.759* (1.943)
Sales at filing (logged)		0.653 (5.833)	0.178*** (0.052)	2.139* (1.220)	-0.227 (2.733)
R&D at filing (logged)		0.858 (4.607)	0.015 (0.044)	1.268 (0.961)	4.999** (2.428)
Employees at filing (logged)		2.570 (6.134)	-0.161*** (0.054)	-5.787*** (1.375)	-8.983*** (2.799)
Earlier-filed patents (2 years, logged)		-3.171 (2.560)	-0.099*** (0.029)	-1.196* (0.633)	-1.226 (1.394)
Number of inventors (logged)		-0.147 (3.089)	0.147*** (0.031)	-1.166 (0.790)	-0.122 (1.898)
Observations	215,235	2,775	3,681	3,517	1,537
R ²	0.290	0.435	0.929	0.917	0.787
Adjusted R ²	0.290	0.429	0.928	0.917	0.783
F Statistic	3,386.712***	81.349***	1,642.735***	1,385.623***	232.610***

Note:

*p<0.1; **p<0.05; ***p<0.01
 Model 1 run for all twins and includes filing year fixed effects.
 Models 2-6 run for Compustat 6 and include technology center and filing year fixed effects.
 Standard errors are clustered at the level of firm x filing year.

Table 9. Firm-level financial results.

	<i>Dependent variable:</i>					
	Sales	COGS (D.V.s 1-5 as an increase in logged variable)	Pre-tax Income	R&D	Employees	Profitability
	(1)	(2)	(3)	(4)	(5)	(6)
Is Leader	-0.006 (0.012)	-0.004 (0.013)	-0.036 (0.033)	-0.004 (0.012)	-0.007 (0.011)	-0.002 (0.002)
Sales at filing (logged)	0.097*** (0.031)	0.121*** (0.033)	-0.047 (0.100)	0.129*** (0.031)	0.160*** (0.028)	-0.015*** (0.006)
R&D at filing (logged)	0.007 (0.020)	0.030 (0.025)	0.031 (0.059)	-0.033 (0.022)	-0.059*** (0.020)	-0.007 (0.004)
Employees at filing (logged)	-0.158*** (0.028)	-0.206*** (0.034)	0.005 (0.076)	-0.154*** (0.028)	-0.152*** (0.024)	0.018*** (0.006)
Earlier-filed patents (2 years, logged)	-0.005 (0.012)	-0.016 (0.014)	-0.015 (0.025)	-0.018* (0.010)	0.013 (0.009)	0.005** (0.002)
Number of inventors (logged)	0.008 (0.008)	0.006 (0.011)	-0.001 (0.022)	0.004 (0.010)	0.006 (0.008)	0.003* (0.002)
Observations	3,143	3,144	2,166	3,135	3,093	3,140
R ²	0.594	0.530	0.437	0.460	0.391	0.093
Adjusted R ²	0.590	0.526	0.430	0.455	0.385	0.085
F Statistic	156.887***	121.325***	57.252***	91.091***	67.723***	11.062***

Note:

*p<0.1; **p<0.05; ***p<0.01
 Outcomes measure within-patent changes between the filing date of the leader and the later of the issue date of the leader and follower.
 All models run for Compustat 6 and include technology center and filing year fixed effects.
 Standard errors are clustered at the level of firm x filing year.

²¹ Constant terms suppressed where they have no interpretability due to other fixed effects. Also true for Table 9.

Table 10. Patent-level citation results.

	<i>Dependent variable:</i>							
	Total (Logged)	Other (Logged)	Self (Logged)	Twin (Logged)	Has Any	Has Other	Has Self	Has Twin
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Is Leader	0.813*** (0.029)	0.490*** (0.029)	0.138*** (0.021)	0.703*** (0.016)	0.323*** (0.014)	0.234*** (0.016)	0.109*** (0.015)	0.760*** (0.013)
Sales at filing (logged)	0.043 (0.034)	0.055* (0.029)	0.039 (0.032)	0.016 (0.019)	-0.021 (0.015)	-0.008 (0.017)	0.029 (0.018)	0.014 (0.013)
R&D at filing (logged)	-0.017 (0.027)	-0.006 (0.026)	-0.063*** (0.022)	0.010 (0.012)	0.021* (0.013)	0.014 (0.014)	-0.032** (0.015)	-0.005 (0.009)
Employees at filing (logged)	-0.075** (0.034)	-0.069** (0.031)	-0.061* (0.032)	-0.026 (0.017)	0.011 (0.013)	0.006 (0.017)	-0.036* (0.019)	-0.008 (0.012)
Earlier-filed patents (2 years, logged)	0.036** (0.017)	-0.004 (0.015)	0.079*** (0.016)	-0.003 (0.007)	0.008 (0.008)	-0.004 (0.008)	0.047*** (0.008)	0.0001 (0.005)
Number of inventors (logged)	0.064*** (0.024)	0.047** (0.023)	0.065*** (0.018)	0.003 (0.012)	0.006 (0.011)	0.013 (0.013)	0.040*** (0.013)	-0.012 (0.009)
Observations	3,670	3,670	3,670	3,670	3,670	3,670	3,670	3,670
R ²	0.732	0.587	0.291	0.650	0.819	0.674	0.351	0.774
Adjusted R ²	0.730	0.584	0.286	0.647	0.818	0.672	0.347	0.772
F Statistic	368.785***	191.924***	55.359***	250.649***	611.039***	278.961***	73.119***	462.371***

Note:

*p<0.1; **p<0.05; ***p<0.01
 Citations counted 6 years after filing.
 All regressions run for Compustat 6.
 All regressions are conditional upon issuance.
 All regressions include technology center and filing year fixed effects.
 Standard errors are clustered at the level of firm x filing year.

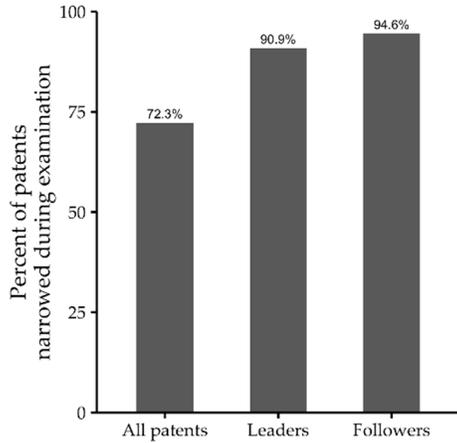


Figure 1. Percentage of patents narrowed during examination at the USPTO.

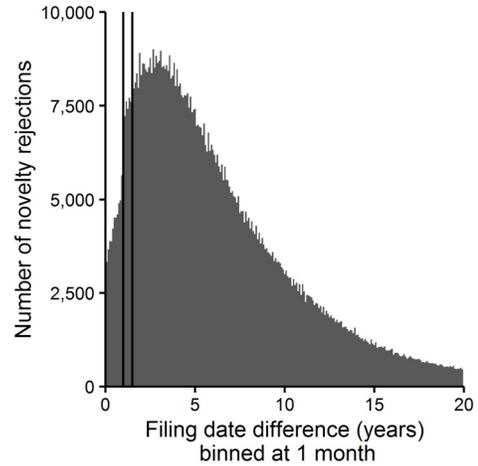


Figure 2. Histogram of rejections by difference in filing date.

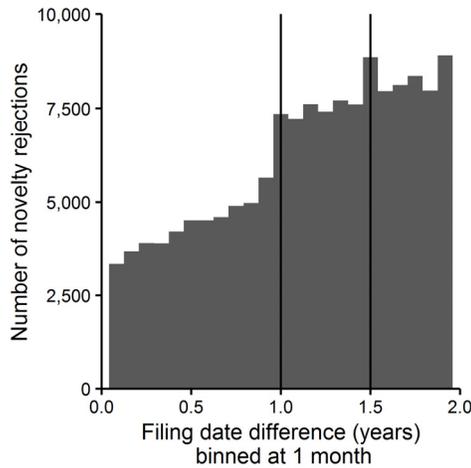


Figure 3. A continuation of Figure 2, with a detailed view of years 0-2.

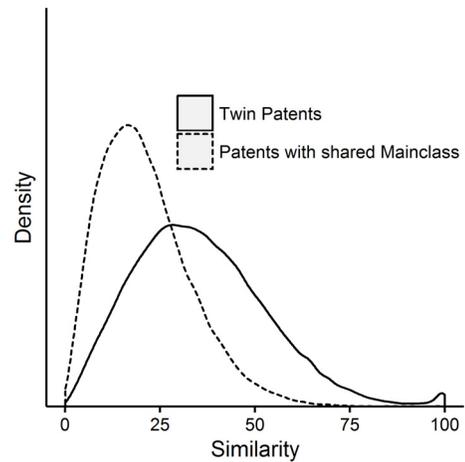


Figure 4. Similarity (as measured by cosine distance) between patents in the same (main) patent class or patent twins. Scores only calculated for cases with similarity greater than a (low) minimum threshold for computational reasons.

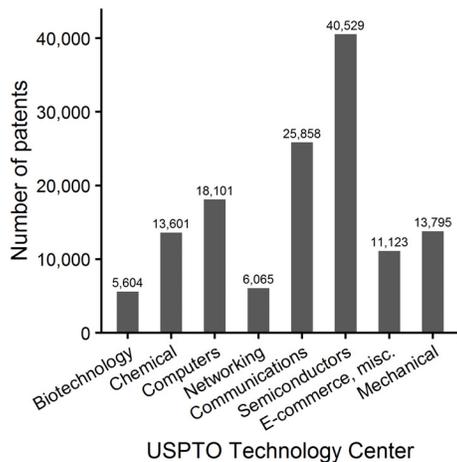


Figure 5. Patents in patent races.

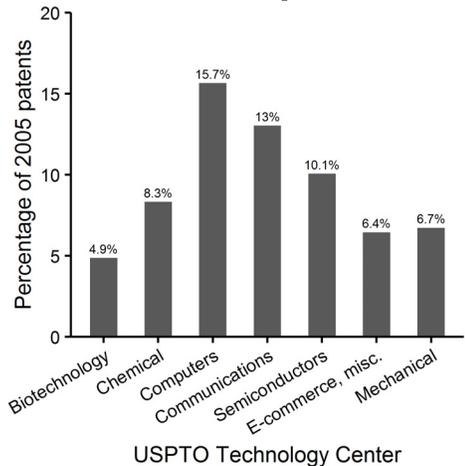


Figure 6. The percentage of patents issued in 2005 that are in at least one twin, binned by USPTO technology center. The networking technology center is omitted because it was started in 2005 and not yet issuing patents.

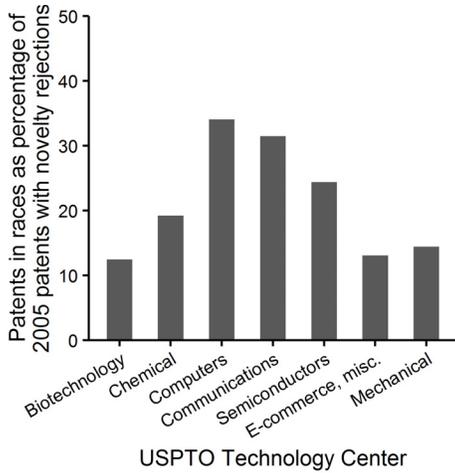


Figure 7. Patents in races, as a percentage of patents with novelty rejections (2005).

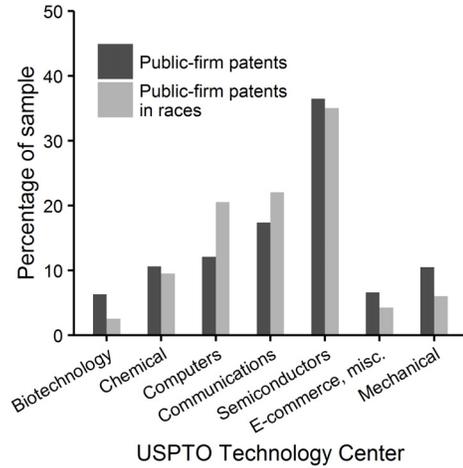


Figure 8. Sample Composition - Public-firm patents in races compared to all public-firm patents (2005).

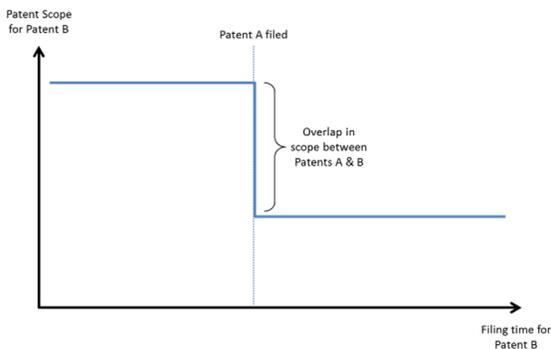


Figure 9. Regression discontinuity setup.

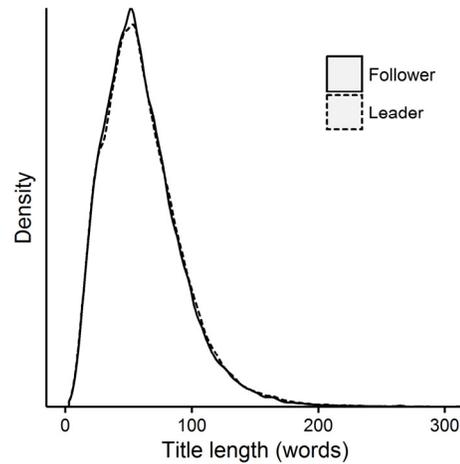


Figure 10. Distribution of title length for leader and follower patents at the time of filing.

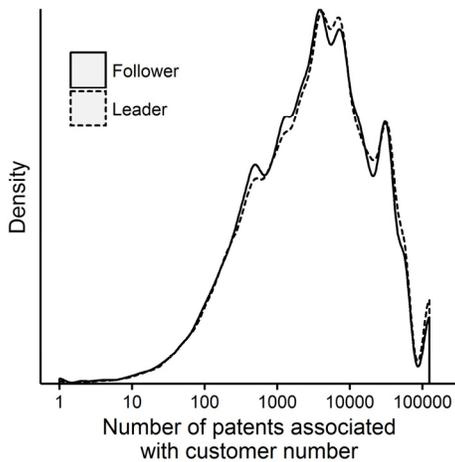


Figure 11. Number of patents associated with the customer number (often the attorney) for leader and follower patents at the time of filing.

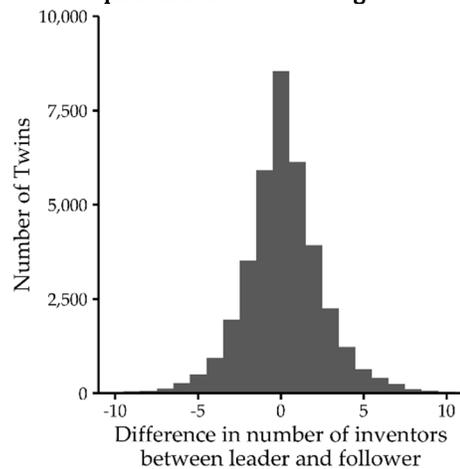


Figure 12: Histogram showing the incidence of differences in the number of inventors between the leader and follower. The full sample of patented twins.

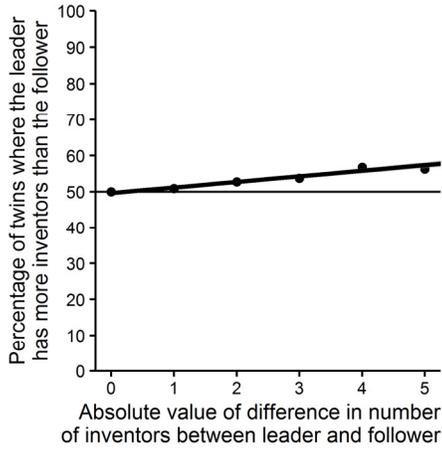


Figure 13. Line plot showing the percentage of twins by the difference in inventors between the leader and follower. The full sample of patented twins.

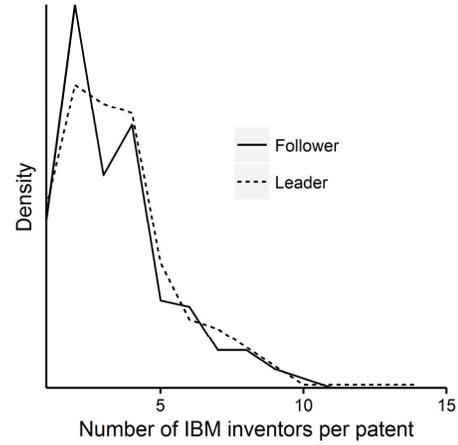


Figure 14: Number of Inventors listed on IBM patents in patent races.

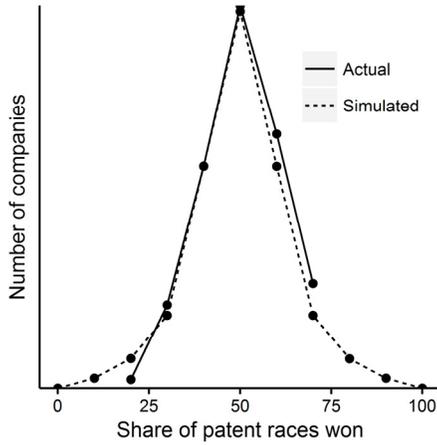


Figure 15. Firm patent race winning percentage for the top 100 firms. The simulation curve plots the result of replicating 10,000 times a random draw [win, loss] with 50:50 odds of winning a race for each patent held by the top 100 firms.

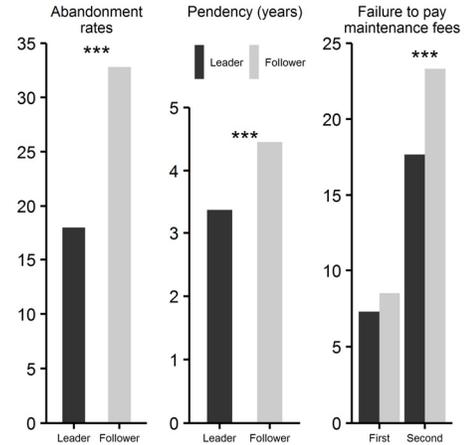


Figure 16. Increase in abandonment rate, pendency, and likelihood of failure to pay maintenance fees.

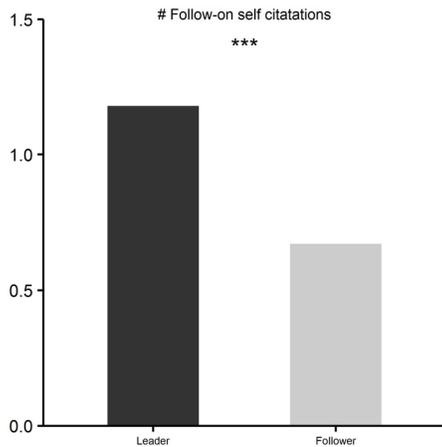


Figure 17. Follow-on self-citations.

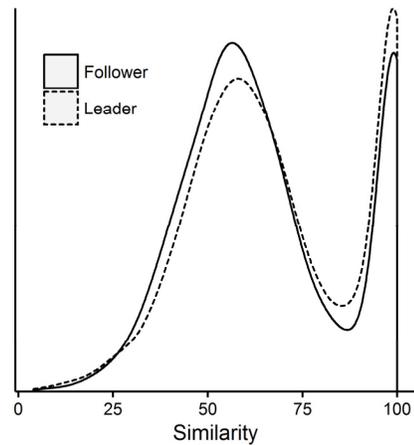


Figure 18. Patent race leaders do follow-on research that is more similar to their winning patent, than do followers of their losing patents (Right-tail reflects divisionals / continuations).