NBER Post-Doctoral Fellowship Proposal

The purpose of the patent system is to foster innovation by providing market exclusivity as a reward for contributing innovation to the sphere of public knowledge, but the extent to which the patent system achieves this goal remains an open question in innovation policy. As explained by Merges and Nelson (1990), the economic consequences of patents largely depend on their scope, and policymakers face a tradeoff in defining that scope. On one hand, broader patents provide greater direct incentives for innovation by covering a wider range of competing products, enabling innovators to capture more of the social returns from their inventions through higher prices. On the other hand, broader patents also cover more subsequent improvements to the inventions, which may reduce the incentives for follow-on innovation. Resolving this tradeoff and setting socially optimal policy would require understanding the effects exerted by the scope of a patent's legal right to exclude on follow-on innovation in related technology, which remains a largely unanswered empirical question.

Patent Racing

During the course of pursuing my PhD, I have made significant progress in empirically investigating the causal impact of patent scope. For instance, competition between firms to invent and patent an idea, or "patent racing," has been much discussed in theory, but seldom analyzed empirically. In one article, I introduce an empirical way to identify patent races, and provide the first broad-based view of them in the real world. The article shows that many of the traditional assumptions behind patent racing models are incorrect. Importantly, the article shows that the loser of a patent race typically still receives a patent, albeit one having much narrower scope than the winner of the race. It also looks at the causal impact 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 produce significantly more follow-on innovation than firms that lose. (Figure 1 and Figure 2).

For this project, I constructed a new data set to address problems of measurement and endogeneity. In genetics research, it is common to disentangle the impact of nature vs. nurture using twin studies. Scholars have imported this concept into innovation research with "Paper Twins"—areas where a single discovery is made by multiple research teams and published simultaneously in the scientific literature (Bikard, 2014). My project takes this concept a step further and considers "Patent Twins"—instances where the same invention is filed with the patent office within a narrow window of time. Perhaps the most famous example of this is the patent filings by Alexander Graham Bell and Elisha Gray on the invention of the telephone within a few hours of each other. I identified patent twins by first employing optical character recognition techniques to convert from image to text more than 50 million pages of correspondence between patent applicants and examiners, and then extracting from that text the patent claim rejections that provide evidence of tight linkages between patent applications. The paper is under review at the RAND Journal of Economics. (Thompson & Kuhn, 2016).

Although the study of patent twins sheds new light on the effects of receiving a broader patent on cumulative innovation, the focus on patent races limits the generality of the conclusions that may be drawn about the effects of patent scope. After all, most patents are not involved in races, and the effect of patent protection outside of races may be very different than within races. A more general investigation of the effects of broader patent rights requires, first, a generalized measure and causal instrument for patent scope and, second, a robust and cross-industry approach for measuring cumulative innovation.

Causal inference on patent scope

In a recent paper, I introduce and validate a simple, practitioner-grounded metric for patent scope: the number of words in a patent's first claim. (Kuhn, Thopmson & Roin, 2016). We examine why this metric works and its possible limitations and then test it against the judgments of legal experts. We find that adding words to a patent's first claim during examination strongly predicts expert views of changes to the patent's scope across most (but not all) technology classes. (Figure 3). In our validation exercise, our measure outperforms other measures commonly used in the literature, some of which have virtually no predictive power. Our validation exercise also allows us to compare, quantitatively, the performance our metric to other measures commonly used in the literature. We find that our measure is significantly better, and that commonly-used techniques are either limited in applicability to a specific area of technology or have virtually no predictive power at all.

We then extend our measure to show how it can be used to draw causal inferences about patent scope based on random assignment of patents to patent examiners, by adopting an examiner-toughness approach similar to Sampat and Williams (2015). Specifically, we measure the number of words added to the claims of all patents examined by each patent examiner, relative to the patent examiner's peers in the same area of technology and leaving out the focal patent. We find that an examiner's toughness score strongly predicts the number of words added to the focal patent. (Figure 4). We implement this causal inference approach to show that an exogenous increase in patent scope significantly increases the likelihood of a patent becoming standards essential.

In my job market paper, I apply the causal instrument to study the market for ideas. (Kuhn, 2016). Patent scope is central to the sale of ideas, which can spur economic growth and provide significant gains from trade. Awarding an inventor a patent on a new idea partially solves a commitment problem that would otherwise prevent the inventor from selling the idea. (Arrow, 1962). In the absence of a patent, a prospective buyer cannot credibly promise not to steal the idea should the inventor reveal it, while the inventor cannot credibly promise to reveal the idea should the prospective buyer pay for it. A firm's ability to use a particular patent to overcome this transactional hurdle derives from two factors: (1) the scope of the patent's legal right to exclude and (2) the effectiveness of that legal right in providing market exclusivity. I first show that an exogenously broader patent is more likely to be sold, holding fixed the underlying idea. I then show that the effect of an increase in patent scope on the likelihood of a sale increases dramatically in certain industries (Figure 5) and in thicker (Figure 6) and less congested markets (Figure 7). These results shed new light on how the practical impact of an exogenous increase in the legal right to exclude varies with endogenous industry and market characteristics.

Measuring cumulative innovation

Measuring the impact of a patent on cumulative innovation presents a substantial challenge because it requires identifying subsequent inventions technologically similar to the focal patent. Current measures of patent similarity rely on the manual classification of patents into taxonomies, but such classification schemes are too coarse to identify the few patents most closely related to a focal patent. In a project with Kenneth Younge (EPFL), we leverage information retrieval theory and Big Data methods to develop a machine-automated measure of patent-to-patent similarity based on a cosine distance applied to the full text of pairs of patents. (Younge & Kuhn, 2015). We show that the measure can be used to illustrate firm technological focus. (Figure 8). We validate the measure and demonstrate that it significantly improves upon existing patent classification systems. Finally, we construct a vector space

having more than 700,000 dimensions and calculate the measure for every pair of more than 5 million patents, producing more than 14 trillion dyadic observations. The resulting measure allows the direct observation of the patents most similar to the focal patent without relying on manual classification of patents into taxonomies. (Figure 9).

In a subsequent paper, we apply the patent-to-patent textual similarity measure to examine the data generating process underlying patent citations, another commonly used measure of patent impact and technological relatedness. (Kuhn, Younge & Marco 2016). We show that a very small but growing number of patents contribute an extremely large number of backward patent citations. (Figure 10). This effect, combined with the overall increase in the number of patents issued, has led to an exponential increase in the total number of patent citations generated over time, which are disproportionately generated by these high-citing patents. (Figure 11). Applying the patent-to-patent textual similarity measure, we show that the patents cited by these high-citing patents are particularly dissimilar to the patents that cite them, which has caused the patent-to-patent similarity of citing-cited pairs of patents to decline substantially over time. (Figure 12). These changes to the data generating process call into question the validity of applying to new data citation-based measurements originally validated on older data.

Although this analysis reveals problems with existing measures, it also lights a path toward new and more powerful techniques to measure cumulative innovation. For example, I employ patent-to-patent similarity in the patent twins project to show that the winners of patent races tend to produce follow-on patents that are more similar to the winning patent than the losers produce for the losing patents. (Figure 2). As another example, I use patent-to-patent similarity in my paper on the market for patents to measure patent rights congestion—the extent to which property rights in a technological market overlap. (Figure 7). I plan to extend these initial efforts to more robustly measure follow-on innovation by analyzing the technological neighborhood of a focal patent in both the past and the future. This text-based measure can be supported with the explicit linkages revealed by corrected and reinterpreted patent citations.

Next steps

I have thus constructed many of the components of the empirical analysis necessary to directly investigate the effect of awarding broader patent rights on cumulative innovation. A post-doctoral fellowship at the National Bureau of Economic Research would be the perfect position for me to combine these elements into a common whole. Most importantly, the post-doctoral fellowship would provide access to many scholars central to innovation policy, who no doubt could provide invaluable guidance. Indeed, the feedback I received from participants at the NBER's Productivity Seminar has already proved invaluable for developing the paper that introduces and validates the causal instrument for patent scope.

In addition, a post-doctoral fellowship at the NBER would assist me in my efforts to make data available for other scholars. The NBER has long played a central role in facilitating research in innovation through data accessibility. Most notably, the NBER patent citation data set has proved a crucial resource for many thousands of scholars. Although I have made publicly available limited samples of instruments and measures I have constructed, a position as an NBER post-doctoral fellow would allow me to advance this agenda by constructing and publishing comprehensive, researcher-friendly versions of my data for general purpose use.

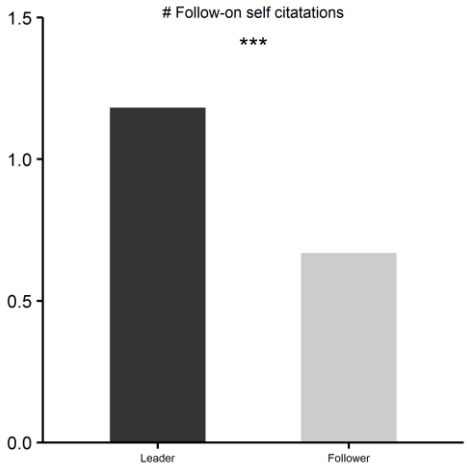


Figure 1. Patent race leaders cite the winning patents more frequently than race followers cite the losing patents.

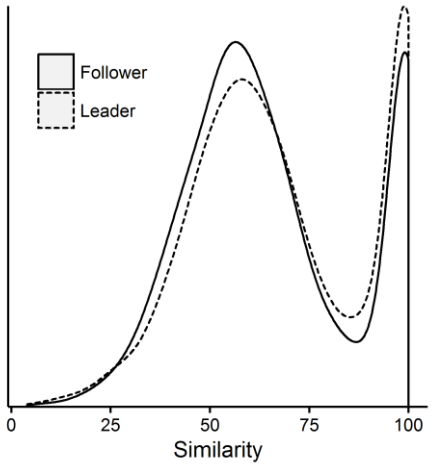


Figure 2. Patent race leaders do follow-on research that is more similar to their winning patent, than do followers of their losing patents (Right-tail reflects continuations).

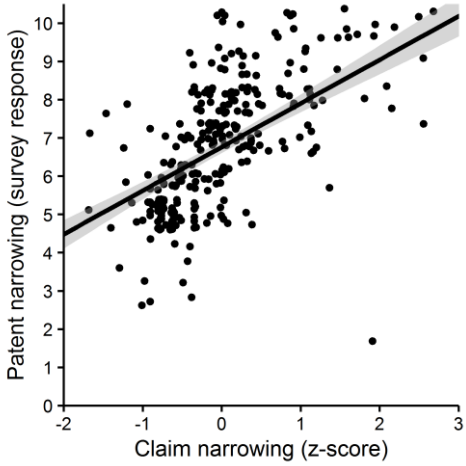


Figure 3. Normalized words to the first independent claim strongly predicts patent attorney evaluations of overall patent narrowing.

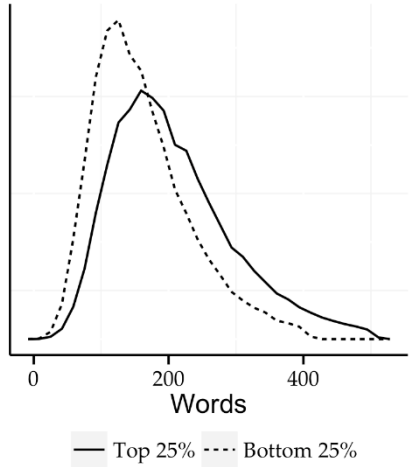


Figure 4. Patent examiners in the upper quintile of toughness cause applicants to add substantially more words to the claims than examiners in the lower quintile.

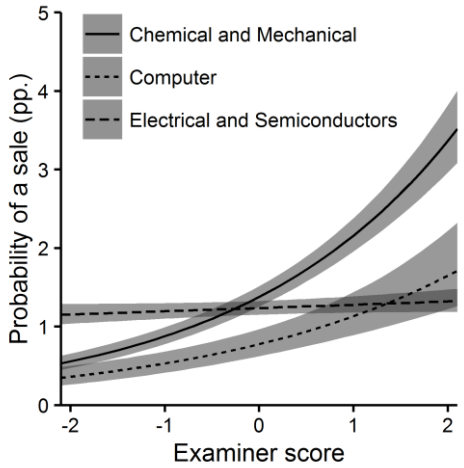


Figure 5. The effect of patent scope as proxied for by examiner score on the likelihood that a patent is sold varies considerably across industries.

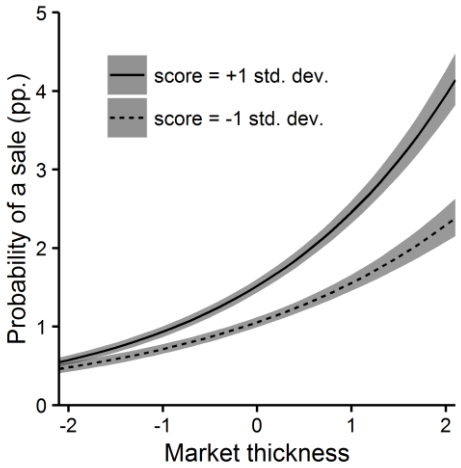


Figure 6. The effect of patent scope as proxied for by examiner score on the likelihood that a patent is sold is substantially larger in thicker patent markets.

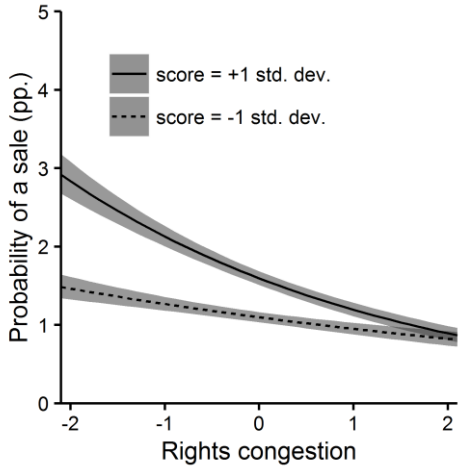


Figure 7. The effect of patent scope as proxied for by examiner score on the likelihood that a patent is sold is substantially smaller in more congested patent markets.

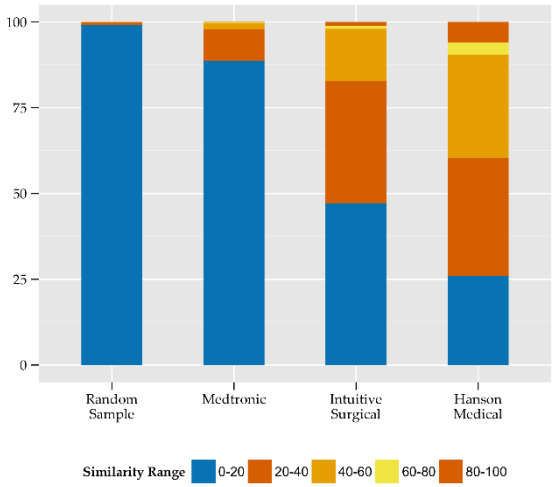


Figure 8. Patent-to-patent similarity of random pairs of patents as compared to pairs drawn from within increasingly focused medical device firms.

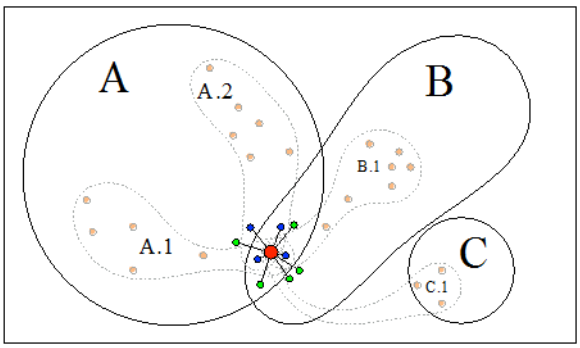


Figure 9. Conceptual illustration of patent-to-patent similarity measure (direct lines) as opposed to patent class (lettered areas).

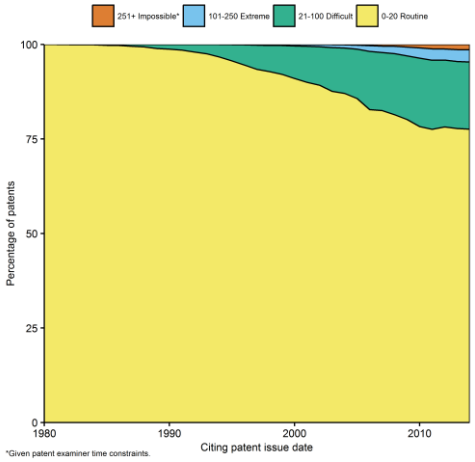


Figure 10. Proportion of patents issued over time by backward patent citation count.

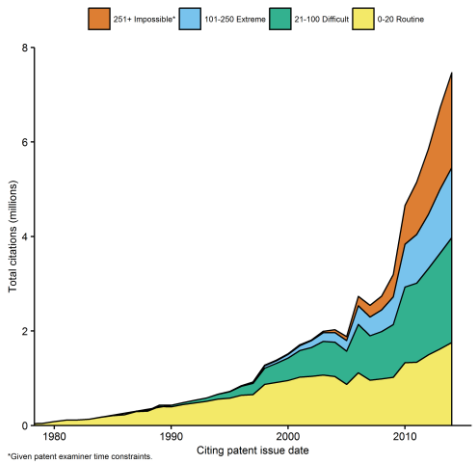


Figure 11. Number of patent citations generated over time by backward patent citation count of citing patent.

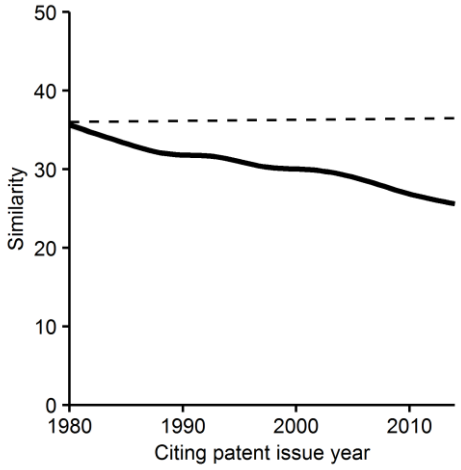


Figure 12. Patent-to-patent textual similarity between citing and cited patents. Dotted line estimates OLS model that includes backward citation count fixed effects.

JEFFREY M. KUHN

jeffrey_kuhn@haas.berkeley.edu • http://jeffreymkuhn.com
1285 64th Street, Apt. C • Emeryville, CA 94608 • 510.457.6858

EDUCATION

Ph.D.	Business and Public Policy Haas School of Business, University of California Berkeley	Expected 2017
J.D.	Law School of Law, University of California Berkeley	2008
B.S.	Computer Science & Mathematics Purdue University	2005
B.A.	Psychology Purdue University	2005

LEGAL BAR MEMBERSHIP

UNITED STATES PATENT BAR	2010 – present
CALIFORNIA STATE BAR	2008 – present

RESEARCH INTERESTS

My research lies at the intersection of innovation and firm strategy. In particular, I study the causal effects of the patent system on firm decisions and outcomes, which provides important insight into how innovative firms establish and sustain a competitive advantage.

My job market paper investigates how the delineation of property rights affects patent sales. The sale of patents can spur economic growth and provide significant gains from trade, but frictions in the market for ideas may lead to the misallocation of resources by obstructing trades that would be efficient from an institutional perspective. A firm's ability to use a particular patent to overcome this transactional hurdle derives from two factors: (1) the scope of the patent's legal right to exclude and (2) the effectiveness of that legal right in providing market exclusivity. I first show that a broader patent is more likely to be sold by employing a causal instrument that provides a plausibly exogenous shock to the scope of a patent's legal right to exclude, holding fixed the underlying idea. I then examine variation in the effectiveness of the right by interacting the instrument with endogenous firm, industry, and market characteristics. These results shed light on how firms profit from innovation and also connect the important but understudied market for patents, widely believed to be illiquid and inefficient, with fundamental research about how markets function in other contexts.

THESIS COMMITTEE

Rui de Figueiredo	Noam Yuchtman	Brian Wright	Robert Merges	Ned Augenblick

WORKING PAPERS

Kuhn, Jeffrey. "The Right to Exclude: Property Rights and Frictions in the Sale of Patents."
Job Market Paper.
Selected for the AOM Strategy Division Dissertation Consortium (2016).

Thompson, Neil, and Jeffrey Kuhn. "Does Winning a Patent Race Lead to More Follow-on Innovation?" ***Under review.***
Nominated for the Best Paper Award, Druid Conference, 2016. Presented at the EPFL MTEI seminar (2016), the Conference on Empirical Legal Studies (CELS) (2016), and the Roundtable for Engineering Entrepreneurship Research (REER) (2016).

Younge, Kenneth, and Jeffrey Kuhn. "Patent-to-Patent Similarity: A Vector Space Model." ***Under review.***
Presented at the USPTO Chief Economist visiting speaker series (2015), the SKEMA KTO research seminar (2015), the DRUID Conference (2016), and the AOM Advances in Patent Data PDW (2015).

Kuhn, Jeffrey, Neil Thompson, and Benjamin Roin. "Causal Inference on Patent Protection."
Presented at the NBER Productivity Seminar (2016), the Searle Center Conference on Innovation Economics (2016), the EPFL MTEI seminar (2015), and the MIT LABi lunch (2015).

Kuhn, Jeffrey, Kenneth Younge, and Alan Marco. "Patent Citations: An Examination of the Data Generating Process."
Presented at the Searle Center Conference on Innovation Economics (2016), the AOM IP Strategy: Patents and Citations session (2016), and the Munich Summer Institute (2016).

LAW JOURNAL ARTICLES

Kuhn, Jeffrey (2011) "Information Overload at the U.S. Patent and Trademark Office: Reframing the Duty of Disclosure in Patent Law as a Search and Filter Problem."
Y*ale Journal of Law and Technology*, Vol. 13.
Selected for republication in the *Intellectual Property Law Review* (2012) by Thomson Reuters, as one of the best intellectual property law review articles published in 2011.

Merges, Robert, and Jeffrey Kuhn (2009) "An Estoppel Doctrine for Patented Standards."
California Law Review, Vol. 97.

Kuhn, Jeffrey (2007) "Patentable Subject Matter Matters: New Uses for an Old Doctrine."
Berkeley Technology Law Journal, Vol. 22.

INVITED PRESENTATIONS

Conference on Empirical Legal Studies, 2016.
NBER Productivity Seminar, 2016.
Searle Center Conference on Innovation Economics, Northwestern Pritzker, 2016.
DRUID Conference, 2016.
Advances in Patent Data PDW, AOM Annual Meeting, 2016.
Patents and Citations session, AOM Annual Meeting, 2016.
Visiting Speaker Series, Office of Chief Economist of USPTO, 2015.
Laboratory for Innovation Science and Policy Lunch, MIT, 2015.
Knowledge, Technology and Organization Research Seminar, SKEMA, 2015.
Management, Technology & Entrepreneurship Seminar, EPFL MTEI, 2015.
Advances in Patent Data PDW, AOM Annual Meeting, 2015.

FELLOWSHIPS AND AWARDS

IP^2 Fellow, Hoover Institution, Stanford University — 2015 – 2016
Working Group on Intellectual Property, Innovation, and Prosperity

Google Research Award — 2014

Robert N. Noyce Memorial Fellow — 2012 – 2014

Ryoichi Sasakawa Young Leaders Fellow — 2011 – 2012

PROFESSIONAL ACTIVITIES

Teaching:	Competitive and Corporate Strategy, Executive MBA, graduate student instructor (2016), Technology & Strategy MBA Seminar, EPFL, guest lecturer (2015), Technology & Strategy MBA Seminar, Purdue University, guest lecturer (2015).
Referee service:	*Strategic Management Journal* (2016), *Strategic Management Society* (2016), *Yale University Press* (2015).
Advising:	M.S. thesis committee member, Christian Zommerfelds, EPFL (2016).
Editing:	Senior Article Editor, *Berkeley Technology Law Journal* (2005-2008), Advisor, Law and Technology Writing Workshop (2007-2008).
Seminar management:	Berkeley Innovation Seminar (2012-2014).
Workshops:	NBER Digitization Tutorial (2015, 2016), NBER Entrepreneurship Boot Camp (2014).

INDUSTRY AND PROFESSIONAL EXPERIENCE

Weaver Austin Villeneuve & Sampson LLP, Oakland, CA — 2007 – 2011
Patent Attorney

University of California, Berkeley, School of Law — 2006 – 2008
Research Assistant for Robert P. Merges, Professor of Law and Technology

OTHER

Computing Theory: Distributed computing, algorithms, complexity theory, graph theory.

Programming: Python, R, Stata, SQL, LaTeX, Matlab, C, C++, Java, JavaScript.

Cloud Computing: Google (Compute, BigQuery, Storage), Amazon (EC2, RDS, S3).

Interests: Chess, Go, literature, SCUBA.

Married since 2011.

PAPER ABSTRACTS

The Right to Exclude: Property Rights and Frictions in the Sale of Patents (job market paper)

The exchange of ideas, facilitated in part by the sale of patents, can spur economic growth and provide significant gains from trade by facilitating the exploitation of complementary assets. However, frictions in the market for ideas may lead to the misallocation of resources by obstructing trades that would be efficient from an institutional perspective. This paper investigates how the delineation of property rights affects patent sales. Using a causal instrument that allows me to identify random variation in the protection afforded by a patent relative to the underlying idea, I show that a broader patent is more likely to be sold. However, the effect is moderated by such factors as market congestion, market thickness, and patent effectiveness.

Does Winning a Patent Race Lead to More Follow-on Innovation? (under review)

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 an empirical way to identify patent races and is the first to provide a broad-based view of patent racing 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 victorious firms produce significantly more follow-on innovation.

Patent-to-Patent Similarity: A Vector Space Model (under review)

Current measures of patent similarity rely on the manual classification of patents into taxonomies. In this project, we leverage information retrieval theory and Big Data methods to develop a machine-automated measure of patent-to-patent similarity. We validate the measure and demonstrate that it significantly improves upon existing patent classification systems. We compute the measure for every pair of patents at the USPTO, generating more than 14 trillion values. Moreover, we illustrate how a pairwise similarity comparison of any and every two patents in the USPTO patent space can open new avenues of research in economics, management, and public policy. We make the data available for future scholarship through the Patent Research Foundation.

Patent Citations: An Examination of the Data Generating Process

Existing measures of innovation often rely on patent citations to indicate intellectual lineage and forward impact. We show that the data generating process for patent citations has changed substantially since citation-based measures were originally validated. Today, far more citations are created per patent, and the mean technological similarity between citing and cited patents has fallen significantly. These changes suggest that the use of patent citations for scholarship needs to be re-validated. We develop a novel vector space model to examine the information content of patent citations, and show that methods for weighting informative citations can substantially improve the information content.

Causal Inference on Patent Protection

The economic consequences of patents—both good and bad—largely depend on their scope, and some patents are much broader than others. Broader patents cover a wider range of competing products, enabling innovators to capture more of the social returns from their inventions through higher prices. Broader patents also cover more subsequent improvements to the inventions, which affects the incentives for follow-on innovation. This paper presents and validates a simple, practitioner-grounded metric for patent scope: the number of words in a patent's first claim. In our validation exercise, our measure strongly predicts expert views of changes to the patent's scope across most (but not all) technology classes and also outperforms other measures commonly used in the literature. We extend our measure to show how it can be used to draw causal inferences about patent scope based on random assignment of patents to patent examiners. Lastly, we implement this causal inference approach to show that an exogenous increase in patent scope significantly increases the likelihood of a patent becoming standards essential.

UNIVERSITY OF CALIFORNIA, BERKELEY

BERKELEY • DAVIS • IRVINE • LOS ANGELES • MERCED • RIVERSIDE • SAN DIEGO • SAN FRANCISCO • SANTA CRUZ

SANTA BARBARA

WALTER A. HAAS SCHOOL OF BUSINESS
545 STUDENT SERVICES BUILDING
BERKELEY, CALIFORNIA 94720-1900

Noam Yuchtman
Associate Professor of Business and Public Policy
Haas School of Business
yuchtman@haas.berkeley.edu

November 21, 2016

Dear Colleague,

I write to most strongly recommend Jeffrey (Jeff) Kuhn for a Post-Doctoral Fellowship on Innovation Policy. Jeff is a unique scholar working on the economics of innovation, combining extraordinarily strong skills in engineering, experience and deep institutional knowledge acquired as a patent attorney, and the ability of an applied economist to pose and test interesting hypotheses about the institutional environment in which innovation takes place. Jeff has applied this toolkit to produce a portfolio of high impact projects, generating important new descriptive facts about the process of patenting, as well as estimating causal effects of the outcomes of the patent-granting process. In addition to producing impressive scholarship, Jeff will be a wonderful colleague: he is humble, thoughtful, an enthusiastic contributor in seminars, and always willing to provide public goods. In short, I believe that he is an ideal candidate for the Fellowship.

Jeff has a rich pipeline of projects, ranging from descriptive papers to sharply-identified papers about the causal effects of the outcomes of the patent-granting process. The first paper of Jeff's with which I become acquainted ("Patent Citations: An Examination of the Data Generating Process," with Kenneth Younge) analyzes the process through which patent citations are generated. For social scientists, patent citations have long been the standard measure of the impact of an innovation, so understanding how a particular patent's set of citations is generated is important as a basic input into innovation research. If the citation generating process were fairly stable across time and across types of innovations, this would be quite comforting; on the other hand, if the citation generating process changed over time – particularly if it changed differentially across types of patents – the naïve use of citations as a measure of innovative impact across time would be problematic.

Jeff brings his unique lawyer-economist's lens to the question of how citations have been generated over time, noting that legal changes altered some patent applicants' incentives to cite extremely broadly. He identifies massive amounts of "copy and paste" citing among a small set of patent applications with extremely large numbers of citations. The "copy and paste" citations are far poorer signals of the cited patents' quality and relevance, pointing to a clear need to address the noisiness of this signal. Overall, I find this to be an extremely important paper on the standard measure of an innovation's impact used in the economics of innovation literature.

A second paper of Jeff's ("Patent-to-Patent Similarity: A Vector Space Model", also with Kenneth Younge), highlights Jeff's skills as an engineer. Jeff and Kenneth collect massive amounts of text from patents, and process this text to generate a novel measure of the relatedness of technologies. This measure outperforms the standard administrative "patent class" categorization, literally redrawing the map of the innovation space. Jeff and Kenneth's data and relatedness metric should be extremely useful tools for innovation scholars in the years ahead.

A third paper of Jeff's ("Does Winning a Patent Race Lead to More Follow-on Innovation?" with Neil Thompson) examines the causal effect of winning a patent race – a competition between firms to receive property rights to a particular technology. Of course, firms that innovate ahead of their rivals typically will be different in a variety of observable and unobservable ways compared to the firms they outcompete. This makes causal identification difficult. Jeff and Neil thus identify pairs of "patent twins": nearly identical patent applications, filed at nearly the identical moment in time. Because the filing dates are nearly identical, one would not expect one firm to be systematically different from the other. Taking advantage of the plausible randomness of one firm's earlier patent, Jeff and Neil examine the effects of winning the patent race within the set of these patent twins. They find that winning the race confers more valuable intellectual property and greater follow-on innovation. This is a very clever empirical project that exemplifies Jeff's ability to write papers taking research design and causal inference very seriously.

Jeff's job market paper ("Frictions and Property Rights in the Sale of Ideas") combines many of the strengths highlighted above. In the paper, Jeff argues that the ability to transfer property rights over ideas and technologies across firms is crucial to making firms more productive and to generating returns on innovation. Transfers can take the form of patent licensing (which has received some attention from scholars) or sales (which have received much less scholarly attention). Jeff argues that the ability to buy and sell intellectual property can be extremely valuable to firms, but that the market for ideas and technology suffers from severe imperfections, among them imperfectly defined property rights over ideas.

Jeff asks, how does strengthening the property rights established by a patent affect trade of that patent? This an extremely interesting question: it treats patents as productive assets that may be bought and sold, but – more than many other types of assets – that may come with poorly-delineated property rights. One might wish to know whether expanded (intellectual) property rights would allow this market to function better. Jeff examines one source of expanded property rights: enlarging the scope of patents.

Again Jeff faces a challenging setting in which to identify causal effects: typically, patents with broader scope will differ in a variety ways from patents with narrower scope. To identify the causal effect of more clearly-defined intellectual property rights, Jeff exploits the quasi-random assignment of a patent examiner to a particular application. This generates exogenous variation in patent scope because some examiners grant systematically broader patents than others. Jeff is thus able to convincingly argue that broader patent scope increases subsequent patent sales. This is an important result: it suggests that, indeed, stronger (broader) intellectual property rights may induce exchange in the market for ideas. Broader patents can encourage more trade in patents,

which might in turn generate more efficient use of extremely valuable intellectual property and also encourage more innovation.

In sum, I reiterate my strongest recommendation that you consider Jeff for the Fellowship: his unique combination of an economics toolkit, legal training and experience, and engineering knowledge makes him a truly unique applied economist working on innovation.

Please let me know if I can provide any additional information. Thank you for your time and consideration of this letter.

Sincerely yours,



Noam Yuchtman