

NBER WORKING PAPER SERIES

KNOWLEDGE SPILLOVERS AND LEARNING IN THE WORKPLACE:
EVIDENCE FROM THE U.S. PATENT OFFICE

Michael D. Frakes
Melissa F. Wasserman

Working Paper 24159
<http://www.nber.org/papers/w24159>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2017

We are grateful for the very helpful comments received from Rochelle Dreyfuss, Jacob Goldin, Mark Lemley, Matt Notowidigdo, Arti Rai, Kyle Rozema, Dave Schwartz, Neel Sukhatme, Neil Thompson, and seminar participants at the Annual intellectual Property Scholars Conference and the Intellectual Property Statistics for Decision Makers Conference. We are also grateful to Bhaven Sampat for providing patent citations data. The work was funded by University of Illinois at Urbana Champaign Research Board, Award 12088 and the Cornell Institute for the Social Sciences Small Grant Award. We are grateful to Matt Berry at the National Center for Supercomputing Applications at the University of Illinois for collecting data from the Patent Office's PAIR database. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Michael D. Frakes and Melissa F. Wasserman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Knowledge Spillovers and Learning in the Workplace: Evidence from the U.S. Patent Office
Michael D. Frakes and Melissa F. Wasserman
NBER Working Paper No. 24159
December 2017
JEL No. J01,M50,O30

ABSTRACT

Using application-level data from the Patent Office from 2001 to 2012, merged with personnel data on patent examiners, we explore the extent to which the key decision of examiners—whether to allow a patent—is shaped by the granting styles of her surrounding peers. Taking a number of methodological approaches to dealing with the common obstacles facing peer-effects investigations, we document strong evidence of peer influence. For instance, in the face of a one standard-deviation increase in the grant rate of her peer group, an examiner in her first two years at the Patent Office will experience a 0.15 standard-deviation increase in her own grant rate. Moreover, we document a number of markers suggesting that such influences arise, at least in part, through knowledge spillovers among examiners, as distinct from peer-pressure mechanisms. We even find evidence that some amount of these spillovers may reflect knowledge flows regarding specific pieces of prior art that bear on the patentability of the applications in question, as opposed to just knowledge flows regarding general examination styles. Finally, we find evidence suggesting that the magnitude of these peer examiner influences are just as strong, or stronger, than the influence of the examination styles of supervisors.

Michael D. Frakes
Duke University
School of Law
210 Science Drive
Box 90362
Durham, NC 27708
and NBER
Michael.frakes@law.duke.edu

Melissa F. Wasserman
University of Texas
College of Law
727 E Dean Keeton Street
Austin, TX 78705
mwasserman@law.utexas.edu

An online appendix is available at <http://www.nber.org/data-appendix/w24159>

I. INTRODUCTION

The economics literature has become increasingly interested in understanding how the behaviors of individual agents are shaped not just by the various economic incentives that they face but also by their social interactions with others. One setting where peer influence is likely to be of critical import to economic growth is the workplace. To what extent are worker decisions impacted by the corresponding behaviors of their co-workers, even when we focus on non-team-based tasks? A still small, but growing number of studies have begun to tackle this question and have started to demonstrate the critical role of social interactions within the workplace. However, various uncertainties and open questions remain. For instance, how do the magnitudes of these peer influences compare with other key determinants in the workplace—e.g., supervisor influences? Moreover, are co-workers responding to each other due to pressures to conform to social norms, or are knowledge spillovers causing co-workers to learn from one another? And, what are the nature of any such spillovers? Do they reflect flows regarding specific, technical knowledge or do they reflect something more general?

In this paper, we confront these questions and the empirical challenges accompanying them while studying the behavior of patent examiners within the U.S. Patent Office. Although context undoubtedly matters in all questions of this nature, the institutional setting surrounding the Patent Office and the rich data on individual examiner behaviors that the Patent Office makes available offers a number of unique and novel tools by which we may approach these challenging inquiries.

One of the key benefits of exploring workplace behavior in the patent examiner context is the tractability offered by the relatively homogenous nature of examiners' jobs. At the core, examiners are tasked with reviewing patent applications and determining whether a patent should be granted covering the underlying invention, a decision that can readily be codified and a decision

that will be the focus of this study. While this benefit may be more easily obtainable in low-skilled worker settings, it is arguably rare to find high-skilled settings amenable to codification and measurement of this sort. Further helpful is the fact that U.S. patent examination is a predominantly isolated and individual task (supervisory oversight aside), making it easier to separate peer-based knowledge flows from what is simply the product of joint team-based efforts.

An additional benefit of the Patent Office context is that we are able to identify and observe each examiner's peer group. Examiners are organized into operational units within the Patent Office called Art Units, each of which is managed by a Supervisory Patent Examiner (or SPE). Each Art Unit consists of roughly eight to fifteen patent examiners who review applications in similar technological areas. Examiners in Art Units generally work in close proximity to one another in the Patent Office—e.g., same floor, same section of the hallway, etc. In our empirical investigation, we treat examiners within the same Art Unit as the relevant peer group; however, we acknowledge that examiners may indeed socially interact with others from outside of these organizational units. To the extent that examiners from other Art Units likewise impact examiner behavior, our results may be seen as a lower bound for the extent of examiner peer influence.

In order to estimate examiner peer effects, we collected data on individual patent applications filed with, and disposed of by, the Patent Office over a 12-year period, with records reflecting the nature of the disposition of those applications and, importantly, the name of the associated examiner and the Art Unit to which they belong. To these data, we merged additional information that we collected via the filing of various Freedom of Information Act Requests, including information about each examiner's tenure at the Patent Office, the names of the SPEs within the corresponding Art Units, and the dates when examiners begin telecommuting.

The identification of peer effects is a task that faces several well-known econometric problems (Manski 1993). At the outset, we note that applications themselves are effectively randomly assigned to examiners within Art-Units.¹ This key fact alone, however, does not cure all sources of endogeneity. To overcome concerns that examiners of similar dispositions may be allocated to similar peer groups—which might otherwise explain any correlated behaviors—our specifications include examiner fixed effects. Of course, even if the composition of peer groups is randomly determined, one might observe correlated behaviors within groups not as a result of actual peer influences but due to unobservable factors that are common to the group—e.g., due to changes in supervisory policies. We take several approaches in alleviating these concerns, beginning with the inclusion of SPE fixed effects in some specifications. This analysis explores how examiners’ grant rates change as the granting tendencies of the peers within their Art Unit change over time while accounting for turnover in supervisors over that time period. Secondly, we estimate specifications with a rich set of Art-Unit-by-year fixed effects (or, alternatively, Art-Unit-by-bi-year effects). These specifications calculate scores reflective of peers’ grant rates at an Art-Unit-by-month level and thereafter explore how a given application’s likelihood of being allowed changes within a given Art-Unit-by-year cell as the granting proclivities of the examiners within that cell (other than the examiner associated with the given application) likewise change.

Finally, to confront the so-called “reflection” problem—e.g., a concern as to whether group behavior affects individual behavior or merely reflects or aggregates individual behavior—we take an approach inspired by Cornelissen et al. (2017) and create peer scores at any point in time based

¹ If applications were assigned within Art Units based on quality—e.g., all of the highest quality applications would go to a particular examiner—that might tend to produce a negative association between individual examiner behavior and peer behavior. Lemley and Sampat (2012) and Frakes and Wasserman (2017) interviewed a number of examiners to confirm the assumption that sorting of this nature does not occur and that applications are randomly assigned within Art Units. A recent paper, however, by Righi and Simcoe (2017) documents evidence of within-Art-Unit assignments based on sub-technology specializations. However, Righi and Simcoe’s analysis finds no evidence to suggest that applications are sorted across examiners based on the importance or claim breadth of the applications.

on the long-term, lifetime grant rates of the examiners comprising that peer group, as opposed to the peer grant rates at that precise time.² To what extent do the collective inherent grant rates of the peers that an examiner faces at a point in time influence her own grant rate at that time? With this construction, changes in the peer score over time capture temporal changes in the composition of the peer group as opposed to temporal changes in the granting practices of a given, stable set of peers. By abstracting away from any effect that contemporaneous co-worker behavior may have on examiner behavior, this approach may likewise lead to lower-bound estimates of the degree to which examiners influence each other's practices. Moreover, by de-emphasizing contemporaneous effects through the use of peer scores based on time-invariant grant rates, this approach to resolving the reflection problem also alleviates concerns that the peer-to-individual grant rate associations we observe are driven by time-varying common unobservables.

While identifying true peer effects in the first place is a task that confronts various econometric issues, identifying the mechanisms underlying any such effects faces challenges of its own. If any peer influences do exist, do they derive from a story of peer pressure in which an examiner's own views towards granting patents is shaped by some degree of shame in departing from a known social norm or do they derive from a story in which examiner's learn how to conduct examination reviews through their social interactions with peer examiners? To attempt to separate these stories, we take advantage of the temporal breadth of our data and explore the dynamics of any observed peer effects. If peer influences follow from a learning mechanism, we would expect that examiners would be most influenced by their peers soon after the affected examiners start

² In the alternative, we attempt to create even more pre-determined peer scores by calculating each examiner's overall grant rates in the years preceding the year in which the subject application is being disposed of by the relevant examiner. The results are virtually identical across these alternative constructions. We use lifetime rates as the primary specification as the purely pre-determined approach will tend to leave few observations for examiners early in the sample period to characterize granting tendencies. We also consider other alternatives to determining individual examiner effects in our construction of peer effects, including those that shrink individual examiner effects towards the mean using signal-to-noise reliability factors (Kane and Staiger 2008).

their jobs with the Patent Office. Under a learning story, we would then predict that in the ensuing years the practice styles learned during their initial years would persist and that future changes in peer composition would have weaker influence. Moreover, under a learning story, we would predict that new examiners are influenced to a greater degree by their more experienced peers rather than by their similarly inexperienced co-workers.

Investigating dynamics of this nature will not only allow us to shed light on the mechanisms underlying any peer influence, but may also further support the identification of peer effects as a general matter. For instance, to the extent that the relationship between examiner grant rates and peer grant scores is indeed the strongest in the case of new examiners, especially in the case of new examiners surrounded by more experienced peers, it is also likely the case that (a) these associations represent effects originating from the peers themselves rather than the other way around (thereby further appeasing reflection problem concerns) and (b) the correlated behaviors that we observe are not merely the result of shocks common to the entire Art Unit.

Ultimately, our results suggest a striking degree of peer influence within the Patent Office that is likely to arise to some degree—though perhaps not exclusively—through knowledge spillovers among examiners, with findings consistent with each of the predictions of the learning story. In the face of a one standard-deviation increase in the inherent grant rate of her peer group, an examiner in her first two years at the Patent Office will increase her own grant rate by roughly 7.6 percentage points, representing a roughly 0.15 standard-deviation increase in her grant rate. Moreover, subsequent changes over her career in the composition of her peer group are associated with notably weaker influences on her grant rate relative to the peer effect during her early years with the Patent Office. Further, results from lagged specifications suggests that peer influences tend to persist over time, rather than being fleeting in nature. Collectively, these findings suggest

that examiners establish somewhat durable practice “styles” early in their career that generally persist even in the face of subsequent changes in their workplace environment. Finally, we find that these early-career effects are stronger when we construct peer scores based on the inherent grant rate of the more experienced co-workers surrounding her.

To put these magnitudes in perspective, we compare the degree to which examiners appear to learn from their co-workers to the degree to which they learn from the Supervisory Patent Examiner (SPE) overseeing their Art Unit. For these purposes, we draw on information from each SPE’s tenure as an examiner—to characterize that SPE’s own views towards patent examination—and estimate similar specifications that draw on within-Art-Unit changes over time in the granting propensities of assigned SPEs. Through this exercise, we determine that peer influences on new examiners are considerably stronger than supervisory influences.

We support these findings through a range of robustness and falsification checks. For instance, we find that peer influences are weaker when constructing peer scores based on the set of examiners that telecommute for at least 4 days a week—i.e., peers that are less present at the office. Moreover, we find stronger signs of peer-based learning and influence in the case of rejections based on obviousness grounds relative to the case of rejections based on lack-of-novelty grounds. This is intuitive insofar as one might predict a stronger scope for learning in the application of the obviousness standard given that it is arguably more nebulous and challenging to apply in comparison with lack-of-novelty rejections. Finally, we move beyond viewing the job of examiners as simply rejecting or allowing patent applications and consider a more nuanced behavior of examiners: affirmatively working with applicants to narrow initially invalid claims to the point where they become allowable. Consistent with the granting/rejecting results, we continue to document strong peer influences in the case of these claim-narrowing behaviors.

Though the workplace peer effects literature has, to our knowledge, yet to dig deeper into peer effects mechanisms than coarsely distinguishing between standard peer pressures and knowledge spillovers, employers and policymakers may indeed wish to know the more precise nature of any such mechanisms. For instance, is the information flow among patent examiners one that respects general examination practice styles and strategies? Or, something more specific and technical? For instance, are examiners learning of specific pieces of prior art—e.g., particular prior patents—from their peers that may bear on the patentability of the applications they are presently reviewing? In an additional empirical exercise, we attempt to uncover specific knowledge flows of this sort taking advantage of another rich dimension to the data available in the patent setting: micro-level patents citations data. We find that when reviewing applications, examiners are significantly more likely to cite to a prior art reference that is among the set of “pet” or favorite prior art references of their peer examiners when those peer examiners are not telecommuting—and are thus socially accessible—relative to when those peer examiners are telecommuting. This finding lends support to a claim that at least some degree of the knowledge flows among examiners capture a rich degree of specificity.

This analysis holds a number of potentially important policy implications given, in part, the significant social welfare consequences of examiners’ granting decisions. Should examiners be overly permissive in their practices and routinely grant patents on inventions that are already known or represent only a trivial advancement over current scientific understanding, they may burden society with the deadweight losses associated with monopoly protection without reaping the benefits of spurred innovation (Nordhaus 1969). In addition, invalidly issued patents can inhibit follow-on discoveries in markets characterized by cumulative innovation (Scotchmer 1991, Sampat and Williams 2014, Galasso and Schankerman 2014). Scholars and commentators have

argued that the Patent Office may indeed be issuing too many patents; others have emphasized the equitable implications and deadweight losses associated with the substantial heterogeneity in grant rates that have been observed across examiners (Frakes and Wasserman, 2017). To begin to address any problems associated with elevated and/or inconsistent granting practices, it is critical to first understand the determinants of such practices. This paper demonstrates the key role that peer learning has to play in the process, a finding that may hold various implications for the ways in which the Patent Office may seek to train and allocate new hires.

The rest of the paper proceeds as follows. In Section II, we provide a background on the related literature and on the patent examination process. In Section III, we discuss our data and methodology. In Section IV, we present our results. Finally, in Section V, we conclude.

II. Background

II.A. Literature Review

This paper contributes to several literatures, beginning with the growing literature exploring peer effects within the workplace. In a recent, path-breaking analysis, Cornelissen et al. (2017) estimates peer effects within firm-occupation groups in an entire local labor market in Germany, focusing on peer effects on wages as opposed to concrete behaviors of workers. The breadth of the sectors included in their analysis allows them to separately test for peer influences in low-skilled and high-skilled settings. They document weaker peer influences in the latter setting, which leave the authors to suggest a potentially weak role for knowledge spillovers in the workplace given that spillovers are more likely to be relevant in such high-skilled sectors.

To further separate knowledge spillover effects from peer-pressure effects, Cornelissen et al. also estimate distributed lag specifications and document lagged peer effects among high skilled sectors and near zero coefficients of the lagged peer score in the low-skilled sample. They suggest

that these findings support some role for knowledge spillovers in the skilled settings and a lack of spillovers in the low-skilled settings. Though the authors do not fully spell out their interpretation of this lagged coefficient, a near zero value for this coefficient would tend to suggest that a *temporary* change in the peer score of interest would have only a contemporaneous effect and no lingering effect (a zero lag would have a different interpretation in the face of a hypothesized *permanent* change in the peer score). Presumably, if one thought that workers would learn from their peers, then one might indeed expect a long-term effect of even a temporary shock. In other words, if knowledge spillovers exist, a given peer shock may alter worker behavior now and in the time ahead even in the face of a subsequent alteration of that peer group in future periods.

While these lagged findings do suggest some degree of learning in high skilled settings, Cornelissen et al.'s analysis does not fully explore this learning story and does not identify the extent to which peer influences help shape initial practice styles early in workers' careers that may persist throughout their careers. Instead, their analysis effectively groups together new and seasoned workers to test for more generic markers of learning, whether initial or ongoing learning. By attenuating an estimation of early-career effects, this analysis may arguably be seen as underselling the role of peers in determining worker behavior and the heterogeneous pathways that workers set out upon. Moreover, should one indeed find stronger evidence of peer effects early in a worker's career, such a finding would further support a learning interpretation and further cut against the possibility of standard peer-pressure effects, including peer-pressure effects that manifest with a delay and that might lead to lagged coefficients of the sort documented in Cornelissen et al. (2017).

Jackson and Bruegmann (2009) come closer to documenting a story of this nature. Using longitudinal data on student achievement and teacher characteristics, they find that teachers

experience greater test score gains among their students when they are surrounded by more effective peer teachers. Moreover, in the associated web appendix, they find that new teachers are even more sensitive to peer quality, supportive of a learning mechanism. They note, however, that they are unable to disentangle a story in which new teachers truly learn from their peer teachers (reflective of knowledge spillovers) from a story in which being surrounded by more effective peers merely gives new teachers more time to engage in self-learning (a learning-by-doing mechanism unrelated to knowledge spillovers). The idea behind this latter story is that teachers do share some common tasks in overseeing the teaching of a particular grade; being surrounded by better peers may reduce the amount of time a given teacher needs to spend on these common tasks, opening up more time for individual learning. Importantly, this school teacher setting differs from the patent examiner context in that new patent examiners do not share analogous common tasks with her their peers, in which event the patent context may face fewer concerns over separating a learning-from-co-worker story from a co-worker-induced self-learning story.³

In addition to building on the above co-worker-related studies, our analysis contributes to a larger literature on learning within the workplace. The management science and organizational theory literatures, among others, have long recognized that early moments within careers at particular organizations are especially critical in determining how workplace practice styles are developed, often theorizing that initial hiring conditions may become “imprinted” on employees.⁴ Behind this imprinting theory is the contention that new hires are more malleable than experienced workers within an organization and, in light of the uncertainty surrounding their new jobs, are

³ Other workplace peer effects studies include Guryan et al. 2009 (professional golfers), Gould and Winter 2009 (professional baseball players), Mas and Moretti 2009 (supermarket workers), Waldinger (2012) (academic scientists), and Azoulay et al. (2010) (academic superstars). Our study is also related to Ho’s (2017) experimental work on peer review within government agencies. Guryan et al. (2009) likewise look for heterogeneous workplace peer effects by the experience level of professional golfers and find more peer sensitivity at higher experience levels, perhaps inconsistent with their expectations of greater sensitivity of new golfers.

⁴ For a recent survey paper regarding imprinting theories, see Marquis and Tilcsik (2013). Among many others, example analyses of imprinting are found in Allen and Meyer 1990 and Baron et al. 1999. These studies are also related to a literature that explores the importance of initial conditions in developing individual styles of behavior more broadly. See, for example, Malmendier & Nagel (2011).

more likely to be operating from a blank state (DiRenzo 1977; Ashforth and Saks 1996).⁵ Once imprinting has occurred and these styles have been established, they may persist even in the face of subsequent environmental changes.⁶

Finally, our paper builds upon a growing literature empirically analyzing patent examiner behavior. Early studies in this still nascent literature demonstrated a substantial degree of heterogeneity in patent office outcomes across patent examiners at the U.S. Patent Office (Cockburn, Kortum, & Stern, 2003; Lichtman, 2004; Mann, 2014). Subsequent studies have begun to unpack the origins of this heterogeneity.⁷ In this vein, Lemley and Sampat (2012) estimate a monotonically increasing relationship between years of examiner experience and examiner grant rates. Frakes and Wasserman (2017) decompose the experience correlation found in Lemley and Sampat (2012) into its various parts. Frakes and Wasserman's analysis suggests that much of that positive correlation between experience and grant rates may be due not to experience itself but to increases in examiner pay-grade levels, which are themselves associated with substantial reductions in the amount of time given to examiners to review applications (which, in turn, may crowd out time to find and articulate bases of rejections). Another key factor driving the experience-grant-rate correlation is the examiner's hiring-year cohort, a factor identified in Frakes and Wasserman (2017) but more fully explored in Frakes and Wasserman (2016). Examiners hired prior to 2002—at a time in which the philosophy of the Patent Office was one characterized by a very permissive granting style—tended to exhibit higher grant rates throughout their careers. In particular, they would tend to maintain higher rates even in the face of a changing philosophy of

⁵ As stated by Dokko et al. 2009, during this the initial period of employment, it is believed that “cognitive models that . . . [individuals] hold can be challenged and replaced with scripts and schema that are more congruent with the new environment.”

⁶ In this light, our analysis and the imprinting literature in general is likewise related to a related literature about path dependence and historical happenstance (see, for example, David 1985).

⁷ Other studies have not necessarily attempted to explain this heterogeneity, but have attempted to embrace it for identification purposes. That is, a number of recent studies have taken advantage of heterogeneity in the granting tendencies of examiners, along with the random assignment of applications to examiners, to estimate the effect of receiving a patent on various outcomes, including the effect of receiving patents on follow-on innovation (Williams and Sampat 2014).

the Patent Office in the mid-2000s that called for a more restrictive stance. New examiners hired in the mid-2000s, on the other hand, did reflect that changed philosophy. In other words, initial career conditions appear to matter significantly in explaining patent examiner behavior. Missing from Frakes and Wasserman (2016, 2017) is an appreciation of the role that co-workers may play in those initial environments.

II.B. Background on Patent Examination Process

Every patent application filed with the Patent Office contains a specification describing the invention, and a set of claims that defines the metes and bounds of the rights the applicant is seeking. Incoming applications are first routed to an Art Unit, an organizational unit consisting of eight to fifteen patent examiners who review applications in the same technological field. Upon arrival, the Supervisory Patent Examiner (SPE) of that Art Unit randomly assigns the application to a specific examiner (Lemley & Sampat, 2012). That examiner will typically begin her examination by conducting a prior art search, that is a search of previous patents, patent applications, or other publications, that are material to the patentability of the claimed invention.

After completing this search, the examiner will assess the patentability of the invention in light of the criteria outlined in the Patent Act. Two of these key criteria are the novelty and nonobviousness requirements. Examiners may reject an application for lack of novelty if they determine that the claimed invention is covered, in its entirety, by a single prior publication or patent. An obviousness rejection is a little more complicated. Such a determination requires an examiner to start with a prior art reference that covers only a portion of the invention and then piece together additional references or rely upon what is known to one of ordinary skill in the art in order to determine whether it would be “obvious” to modify any one of the cited prior art references to achieve the claimed invention.

It is critical to emphasize that there are two types of examiners working within each Art Unit. Examiners at pay grades GS-13 and below on the General Schedule pay scale constitute the first type and are known as “secondary” or “assistant” examiners. After completing a special evaluation program, examiners may be promoted to become “primary” examiners (generally reaching GS-14 at this time). Primary examiners have full authority to sign off on all aspects of their reviews without the need for supervisory review. Though primary examiners continue to be assigned their own applications to review, they also help serve as quasi-supervisors for assistant examiners. That is, while assistant examiners independently review, and complete the bulk of the work associated with, the applications assigned to them, they must have their reviews and decisions approved by a supervisor—either by a primary examiner or by their SPE.

While SPEs do help sign off on the reviews of assistant examiners, they no longer review applications of their own (unlike primary examiners). However, their supervisory functions go beyond checking the work of assistant examiners, as SPEs are also tasked with overseeing the training (initial and ongoing) of the examiners within their Art-Unit and with making sure that their Art Units implement general Patent Office policies and directives.

Our analysis below endeavors to account for the different roles of assistant examiners, primary examiners and SPEs. In particular, given the dual roles of primary examiners, we make sure in our analysis below to separate GS-14 examiners from sub-GS-14 examiners when constructing peer groups in order to identify true “peer” effects. In other words, we aim to understand how the behavior of assistant patent examiners are influenced by their peer assistant examiners. Nonetheless, we also take advantage of this institutional feature and estimate how the granting behavior of assistant examiners are affected by the granting philosophies of the group of primary examiners working within their Art Units. This exercise effectively allows us to compare

the magnitudes of pure peer effects with quasi-supervisor effects. We then take that comparison one step further and compare each of these effects with pure supervisor effects. To capture these latter effects, we use the granting practices of SPEs when they previously acted as examiners to characterize the granting philosophies of SPEs and thereafter observe how granting behaviors of examiners change in connection with within-Art-Unit changes in such philosophies—driven by within-Art-Unit changes in SPEs over time.

III. Data and Methodology

III.A Data

We collected data on individual patent applications from the Patent Office’s Patent Application Information Retrieval (PAIR) database, covering roughly 1.4 million utility patent applications that were filed on or after March, 2001 and that reached a final disposition—i.e., excluding ongoing applications—by July 2012. These data include, among other things, information on whether or not the application was granted, the name of the examiner charged with reviewing the application, the Art Unit to which the application was assigned, and information on the bases of rejections associated with the application—e.g., whether it was subject at any point to an obviousness rejection or a lack-of-novelty rejection.⁸

Through a series of Freedom of Information Act (FOIA) requests, we also collected a range of information about examiners, including the year in which they joined the Patent Office (left censored at 1992), and their GS-level over each year in our sample. Moreover, for those examiners participating in the Patents Hoteling Program (PHP), which allows examiners to work from home at least 4 days a week, we collected information on the precise day in which they started to

⁸ The rejection-criteria data was collected based on a textual analysis of office actions following the execution of optical character recognition programs to office actions uploaded to the PAIR database. Further details on this data collection can be found in the Online Appendix to Frakes and Wasserman (2017).

telecommute.⁹ We then merged these examiner-specific fields with the application-level data (using a fuzzy-name-matching application).

Through additional FOIA requests, we collected information about the identity of the Supervisory Patent Examiner (SPEs) for the Art Unit associated with that application. For those SPEs that were promoted to that position during our sample, we observed information about the applications that those individuals reviewed while they were patent examiners prior to such promotions, allowing us to calculate their pre-SPE grant rates, a metric that we use to proxy for their general granting dispositions. For those SPE's promoted to that rank prior to the beginning of our sample, we were unable to determine their examination style. Overall, we are able to assign pre-SPE grant rates for the SPEs associated with given applications for roughly 38% of the applications in our sample.

Table 1 provides summary statistics for the key variables in the analysis. Across all applications and all examiners, applications are granted roughly 70 percent of time throughout our sample.¹⁰ Table 1 also presents the mean peer / supervisory grant rate across the three relevant peer / supervisory groups in our analysis: (1) assistant examiners, (2) primary examiners and (3) SPEs. As found in Frakes and Wasserman (2017), examiners' grant rates tend to increase as they rise within the ranks in the Patent Office. With this in mind, the mean peer / supervisory grant rates unsurprisingly increase as we move across these three groups, with the pure peer score—the grant rate of the assistant examiners—coming in at roughly 65%.

III.B. Methodology

⁹ To be eligible for the PHP, patent examiners must have achieved a GS-12 level, have positive performance ratings, and have worked at the Agency for at least two years. The PHP began in 2006. Over 86% of those eligible to participate in the Patent Office's teleworking programs in fact elect participation.

¹⁰ In calculating grant rates, it is important to clarify that Requests for Continued Examinations do not count as rejections of one application and filings of a new applications; rather, they count as an intermediate step within the same application and thus do not contribute to an increase in the grant-rate denominator. Continuation applications, however, are counted as new applications.

To begin to explore how patent examiners learn from their peer examiners, we estimate the following specification, restricted to the applications completed during an examiner's first six years at the Patent Office (this restriction being discussed further below and also lifted below):

$$\begin{aligned}
GRANT_{aikt} = & \alpha + \boldsymbol{\gamma}_i + \boldsymbol{\delta}_k + \boldsymbol{\delta}_t + [\boldsymbol{\theta}_s] + \beta_1(PEER_{ikt}) \\
& + \beta_2(\mathbb{1}(EXPER_{ikt} = \{3,4\})) \\
& + \beta_3(\mathbb{1}(EXPER_{ikt} = \{5,6\})) \\
& + \beta_4(\mathbb{1}(EXPER_{ikt} = \{3,4\}) \times (PEER_{ikt})) \\
& + \beta_5(\mathbb{1}(EXPER_{ikt} = \{5,6\}) \times (PEER_{ikt})) + \boldsymbol{\beta}_6 \mathbf{X}_{aikt} \\
& + \varepsilon_{aikt}
\end{aligned} \tag{1}$$

where a indexes the individual application, i indexes the individual examiner, k indexes the Art Unit to which the application is assigned, and t indexes the year in which the application is disposed of by the examiner. $GRANT_{aikt}$ indicates whether or not the given application was allowed by the examiner. $EXPER_{ikt}$ captures the experience level in years of the examiner (based on the year of application disposition matched with annual Patent Office rosters). Art-Unit and year effects are captured by $\boldsymbol{\delta}_k$ and $\boldsymbol{\delta}_t$, respectively. Examiner fixed effects are captured by $\boldsymbol{\gamma}_i$, allowing us to account, among other things, for endogenous allocations of examiners with particular characteristics to certain peer groups. \mathbf{X}_{aikt} includes other application characteristics including whether the applicant is a large or small entity (as the Patent Office uses such terms for fee-setting purposes), whether the application has foreign priority, the duration of the examination period (and its square), and the GS-level level of the examiner at the time of disposition.¹¹

$PEER_{ikt}$ represents the peer score of interest for our analysis. To calculate this score, we begin by determining the lifetime grant rate for all examiners in the sample (the percentage of

¹¹ The entity size of the applicant is not included in the machine-ready PAIR data publicly disseminated by the Patent Office; however, we were able to collect this information through our own Optical Character Recognition analysis of the bulk PAIR files.

applications that they allow over their tenure at the Patent Office). Given random assignment of applications to examiners, this rate should be indicative of the examiner’s general disposition towards allowing patents—i.e., we should not be concerned about high grant rate examiners systematically being assigned the highest quality applications (Sampat and Williams 2014). For each application, we then calculate $PEER_{ikt}$ by taking the average of the lifetime grant rates for all examiners in the same Art-Unit-by-year cell as the application in question, leaving out the contribution of the examiner reviewing the application at issue. Rather than drawing on contemporaneous changes in the granting practices of a stable set of peers, this approach identifies the influences of peers’ granting practices on an examiner’s grant rates by drawing on changes in the composition of the peer group over time within an Art Unit.

In certain specification checks, we include SPE fixed effects, θ_s , to account for changes in supervisors within Art Units over time, which might otherwise represent shocks common to all examiners within the affected Art Unit (which, in turn, might otherwise explain peer associations). The SPE rosters that we obtained from the Patent Office allow us to assign SPEs to over 90% of the observations in our sample. In yet another specification check meant to address concerns over common unobservables, we calculate peer grant scores at the Art-Unit-year-month level and impose a full set of Art-Unit-by-year fixed effects to allow for yearly changes in Art-Unit-specific supervisory policies. We leave this latter specification as a robustness check in that it relies heavily on within-year changes in peer composition. Considering that applications are actually completed over a process of time that often spans greater than a year, such a fine-grained temporal analysis requires relatively strong assumptions about the critical significance of the time of application

disposition itself.¹² With this in mind, we also take a more in-between approach and estimate specifications with Art-Unit-by-bi-year fixed effects, allowing more time within Art-Unit-by-time cells to observe fluctuations in peer composition.

As further robustness exercises, our analysis below takes several alternative approaches in constructing the relevant peer score. For instance, rather than using overall lifetime grant rates to characterize a peer examiner’s granting proclivity, we consider a peer examiner’s overall grant rates up to year $t-1$ to characterize that her granting proclivity at year t . Moreover, instead of using overall grant rates, we estimate specifications that use lifetime grant rates that are adjusted for certain characteristics of examiners over that lifetime, mainly their experience levels, paygrade levels, and the years and Art Units in which they practiced. We perform these adjustments by regressing an examiner’s grant rates on these various characteristics and a set of examiner fixed effects and using the estimated fixed effects to characterize an examiner’s lifetime granting proclivities. In Figure 1, we depict a kernel plot of the distribution of estimated examiner fixed effects across the various examiners in our sample, demonstrating the substantial degree of examiner grant-rate heterogeneity underlying this empirical exercise.

In our primary specifications, we limit the above analysis to applications reviewed by assistant examiners—that is, by examiners at GS-levels 13 and below—while also making sure to construct the peer score at time t based on the composition of other assistant examiners in their Art-Unit at time t . This ensures that we are picking up pure peer effects since none of the examiners in this group would be serving any supervisory function over the others. In the alternative, we estimate specifications similar to that above but replace $PEER_{ikt}$ with a similar

¹² Recent research by Frakes and Wasserman (2016) arguably support this assumption to some degree, however, in demonstrating the often insufficient and cursory nature of the decisions that examiners make in their early rounds of review, perhaps as a result of procrastination behaviors.

measure meant to reflect: (1) the mean inherent grant rates of the primary patent examiners (GS-level 14) practicing in the same Art Unit at year t and (2) the inherent grant rates of the Supervisory Patent Examiners overseeing that Art Unit in year t . These latter specifications allow us to explore quasi-supervisory and supervisory influences, respectively, with which we can compare to the degree of peer influences.

Key to the above specification is its ability to explore how peer (and/or supervisory) influences evolve with an examiner's tenure at the Patent Office. As such, the key coefficients of interest in the above specification are those capturing how peer effects vary by the experience group of the examiner reviewing the application in question, where we use two-year experience bins and where we focus, at least in our primary specifications, on examiners within their first six years at the Patent Office.¹³ With the reference group being examiners in their first and second year at the Patent Office, the estimated coefficient of $PEER_{ikt}$, β_1 , indicates the degree to which an examiner's grant rate is associated with her peer examiner's granting tendencies during her first two years at the Patent Office, drawing on within-Art-Unit changes in peer composition over that time to derive that estimate. β_4 , the coefficient of the interaction between the peer grant score and the indicator for being in the third and fourth year at the Patent Office then captures the degree to which subsequent changes in peer composition during this next stage of her career are associated with the grant rate that she applies during that third and fourth year, where the magnitude of this effect is interpreted with reference to first-and-second-year effect. We group experience bins into two-year groups in order to allow for temporal variation within those groups in peer composition.

Given the critical nature of this dynamic analysis to our empirical exercise and to our aim to look for markers of examiner learning, we attempt to construct this dynamic specification in the

¹³ The results are robust to alternative groupings of experience bins.

most balanced way possible. As such, we limit the analysis in our primary specifications to those examiners that we can observe throughout our sample period at each of their first, second, third, fourth, fifth and sixth years of work experience and limit our analysis to those applications that they review over that time. Accordingly, in our primary specifications, we drop those examiners that either stay at the Patent Office for a shorter period of time during our sample or that started at the Patent Office prior to our sample beginning (which we can flag with the backdated rosters we received from the Patent Office). Nonetheless, as robustness exercises, we also estimate the same specification indicated above but in an unbalanced way—e.g., including examiners that start their fourth year at the Patent Office at the beginning of our sample and use their observed behavior to identify peer influences at the fourth year and beyond. In yet other specification checks, we estimate unbalanced specifications that follow examiners over their entire careers.

In the results section below, we also highlight several additional modifications to the above design, allowing us to more fully explore an examiner learning story. All empirical specifications are clustered at the Art Unit level to account for auto-correlation within Art Units over time.

IV. Results

IV.A. Peer effects

We present the results from our primary specification in Column 1 of Table 2. As reflected by the estimated coefficient of the peer score variable, we find that a change from 0 to 100 percentage points in the mean inherent grant rate of an examiner’s peer group is associated with a roughly 43 percentage point increase in her own grant rate during her first and second year at the Patent Office. Interpreted differently, these results suggest that a one standard deviation increase in the peer grant rate is associated with a 7.6 percentage-point—or a roughly 0.15 standard-deviation—increase in an examiner’s own grant rate at the beginning of her career. During the

third and fourth's year of an examiner's tenure, we find that a subsequent one-standard deviation change in the grant rate of her peers is associated with a 3.1 percentage-point lower increase in her own grant rate than it was during her first and second year. Furthermore, a one-standard deviation increase in the peer grant score during her fifth and sixth years at the Patent Office is similarly associated with about a 3.2 percentage-point lower increase in her grant rate relative to the first and second year effect. In Column 2, we show that these results are nearly identical when using an unbalanced sample of examiners over their first six years of their career.

In Table 3, we demonstrate that this pattern of results—i.e., a strong early career effect that soon weakens—is robust (subject to some fluctuations in magnitudes) to a number of alternative specifications, including those that (1) include SPE fixed effects (for those Art Units-year cells for which we were able to collect data on the relevant SPE), (2) impose Art-Unit-by-year fixed effects (while calculating peer grant scores at the month-year level) and Art-Unit-by-bi-year fixed effects, (3) calculate peer scores (inherent peer granting tendencies) by using peer examiner grant rates in all of the years prior to the year of the application in question, (4) calculate peer scores by using risk-adjusted lifetime peer grant rates, adjusting for certain characteristics of the examiners (mainly, experience levels and GS-levels), (5) calculate peer scores by using an empirical Bayesian estimation approach that effectively modifies the above risk-adjusted approach to further adjust predicted examiner grant rates to reflect their reliability by shrinking noisier estimates towards the mean,¹⁴ (6) likewise use an unbalanced sample of examiners but that also follow examiners over their entire careers as opposed to just their first six years. In Figure II, we graphically depict the results from this latter full-career specification.

¹⁴ This approach is inspired by various studies in the teacher value-added literature, for instance Kane and Staiger (2008) and Chetty et al. 2014). Further details on this estimation strategy are provided in the Online Appendix.

This strongly declining influence of peer composition beyond an examiner's early career moments is consistent with a model of examiner learning. As is well supported empirically by the examiner-heterogeneity literature (Cockburn et al. 2002), the Patent Office extends considerable discretion to examiners in conducting their reviews. But, how do examiners learn to operate within these bounds of discretion? Under various models of learning, one would predict that examiners would draw on certain sources of information—whether from self-experience or from communication with others about their experiences—to develop a practice style within this range of discretion. Moreover, under such models, they may be especially likely to do so early during their careers when they are most impressionable.¹⁵ Once developing that style, one would predict that further stimuli in subsequent years would be less influential in further shaping granting practices (Marquis and Tilcsik 2013). However, considering the nature of their jobs—i.e., dealing with evolving technologies—we would not necessarily predict that learning processes (and the influence of peers on learning) would come to a complete halt later in an examiner's career; we may simply predict a diminished role for learning. Not only do the dynamic pattern of results present in Columns 1 and 2 of Table 2 and in Table 3 provide support for the claim that examiners begin to develop their practice style soon after starting with the Patent Office, but they also support the hypothesis that much of the source of information behind this learning comes from an examiner's peers. The fact that peer influences do not completely diminish later in an examiner's career may be consistent with the above supposition that examiners may continue to learn to some degree throughout their careers. However, that residual peer effect may also be consistent with

¹⁵ Consider for instance models of imprinting which broadly suggest that individuals—e.g., employees—may during particularly sensitive moments begin to take on characteristics reflective of the environment during those moments and persist with those characteristics moving forward even in the face of subsequent environmental changes. Marquis and Tilcsik (2013). When applying such models to individual employees, this literature often focuses on early career moments as capturing these sensitive times during which the subjects are especially malleable. Other models might lead to similar implications through predictions that employees are more likely to invest in building human capital—invest in learning—when they have longer time horizons in front of them—i.e., when they have less experience (Jackson and Brueggemann 2009).

some role for standard peer pressure behaviors among examiners. We return to this ongoing learning point below when discussing lagged effects and when discussing knowledge flows regarding specific pieces of prior art.

Under a story of learning, in addition to predicting a diminishment in the influence of stimuli later in an examiner's career, we would also predict a persistent effect of any stimuli that they experience in the past. To explore the durability of initial influences more directly, we regress the incidence of an application being granted on the peer grant score along with a 2-year lag of that peer score, in addition to the various controls included in equation (1).¹⁶ A positive coefficient of the lagged term would either suggest a persistent effect of a past temporary shock or an effect from a permanent shock in peer composition that grows over time. Either interpretation would document persistence in peer effects that is arguably more consistent with a model in which examiners learn from one another than with a model in which examiners simply wish to conform in order to avoid social stigma (in the latter case, one would perhaps expect more fleeting responses to temporary shocks in peer composition). We note at the outset of this exercise that we cannot properly evaluate this question while looking at examiner decisions in their first and second year at the Patent Office—after all, any relationship between an examiner's grant rate in the period of time prior to her second work anniversary and the peer score for that Art Unit a full two years prior would not be informative on any learning behavior of that particular examiner (since she was not there at that prior time). However, refining our balanced sample a little further, we can ask how examiner behavior between their third and sixth year at the Patent Office is shaped not only by their contemporaneous peer scores but also by their lagged peer scores. We simplify this endeavor

¹⁶ We choose a 2-year lag time period as opposed to a 1-year lag (as used in Cornelissen et al. 2017), again considering that patent applications are themselves processed over a period of time often spanning a year, in which we event, we hesitate to model temporal dynamics on too fine-grained of a basis. That being said the results from the 1-year lagged specifications are nearly identical to the 2-year lag specifications we choose as our primary specification. Moreover, these results also generalize to a consideration of longer lagged periods.

by not interacting these contemporaneous and lagged effects by separate experience groups over this range (a specification that would otherwise entail two dimensions of interactions). Moreover, in light of this choice to test for lagged responses in general—without separately testing for lagged responses for each experience interaction term—we also estimate specifications that test for lagged responses throughout an examiner’s entire career (doing so in an unbalanced approach), as opposed to limiting our focus solely to examiners between their third and sixth years.

We present the results from these lagged specifications in Columns 4-7 of Table 4. In the balanced sample specification pooling examiners over their third-through-sixth years at the Patent Office, we estimate a coefficient of the contemporaneous peer grant score of roughly 0.25 and a coefficient of the lagged peer score of roughly 0.12, consistent with the expectations of a learning mechanism behind the documented peer influences. The extent of this lagged effect is robust to the estimation of an unbalanced sample that consider examiners in all years beyond their second year and to the inclusion of SPE fixed effects, as demonstrated by Column 5 and 6 of Table 4. The fact that we observe lagged effects on average even when tracking examiners throughout their entire careers may also suggest that the residual role for peer influences we observe later in an examiner’s career (discussed above) may indeed be reflective of ongoing learning as opposed to simple peer pressure.

These patterns of persistence complement prior research in Frakes and Wasserman (2016) and Frakes and Wasserman (2017), which found that a key determinant of an examiner’s grant rate is the year in which she is hired, combined with the general philosophy of the Agency’s central administrators at such times. Examiners starting in the mid-2000s—at a time when the Agency Director Jon Dudas espoused a more restrictive stance towards granting—consistently exhibited a roughly 7-10 percentage point lower grant rate throughout their careers relative to examiners

starting prior to 2003—an era commonly perceive to be characterized by a more permissive granting culture (perhaps evidenced by the Patent Office itself when stating in its 2001 Corporate Plan that its primary mission was to help “customers” get patents). The magnitude of these hiring-year cohort effects are of a comparable size to the findings in the present analysis, with the difference in grant rates between examiners starting in the restrictive era versus the permissive era being as large as the impact of a one-standard-deviation increase in the inherent grant rate of the peers that surround new hires at the Patent Office. This suggests that peer effects may be as influential in shaping examiners’ granting practices as large high-level fluctuations in the stated missions of Agency heads.

To further put the magnitude of our findings into perspective, we also compare our findings to those set forth in Frakes and Wasserman’s (2017) analysis of examiner time allocations, another key determinant of patent examiner behavior identified in the literature to date. Conceptually, time allocations are of critical import to observed granting practices given that examiners are legally expected to allow applications if they are not able to find and articulate a basis of rejection in the allotted time. Frakes and Wasserman (2017) find that as examiners ascend from GS-7 to GS-14—a path that essentially cuts in half the amount of time examiners are given to review applications—they experience a roughly 10-19 percentage-point rise in their grant rates. As such, the effect of a 1-2 standard deviation fluctuation in the peer grant rate facing an examiner at the beginning of her career is nearly as influential in shaping her granting practices as is a roughly 100% change in the amount of time given to examiners to review applications.

IV.B. Supervisory Effects

Placing the magnitude of the above findings into even further context, we also estimate similar specifications exploring the relationship between examiner grant rates and the inherent

grant rates of the GS-14 “primary” examiners practicing in the affected examiner’s Art Unit, a particular group of peer examiners that also serve a quasi-supervisory role in helping to sign off on the reviews of assistant examiners. As demonstrated by Column 3 and 4 of Table 2 (focusing on balanced and unbalanced samples, respectively), we estimate a similar dynamic pattern in the case of this peer/supervisor group as we do in the case of the pure peer group consisting of peer assistant examiners—that is, a strong relationship between an assistant examiner’s grant rate and the inherent grant rates of the group of primary examiners in her Art Unit at the time, though one that is strongest when that assistant examiner is new to the Patent Office. The magnitude of this relationship is roughly equal to the pure peer effects documented in Columns 1 and 2 of Table 2, capturing how assistant examiners respond to their peer assistant examiners.

We extend this supervisory analysis to the consideration of SPE effects in Column 5 and 6 of Table 2. We continue to document an influence that dissipates with examiner experience. That is, we do find that new examiners’ granting practices may be shaped by the inherent grant rates of the SPEs overseeing their Art Units (as captured by that SPEs granting history before they were SPEs), and we find that changes in SPE compositions within Art Units later in an examiner’s career are associated with a weaker and weaker influence on an examiner’s grant rates at those later moments. In comparing the magnitude of these SPE effects with those of the assistant examiner peer effects or the primary examiner peer effects, it is important to note that this SPE analysis is performed on a subset of Art-Unit-by-year groups for which we have data on previous grant rates of the relevant SPEs. To best form the comparison group for these SPE effects, in Column 6 of Table 3, we replicate the pure peer effects analysis on this subset of Art-Unit-by-year cells. Doing so, we find that the magnitude of the pure peer effects is nearly three times as large as the magnitude of these supervisory effects.

All told, the evidence suggests that the composition and the granting backgrounds of an examiner's peers appear to be just as or more influential on an examiner's early-career granting decisions than the granting background of an examiner's supervisors.

IV.C. Other Specification Checks

We next consider several falsification exercises, beginning with the estimation of a 2-year lead coefficient of the peer grant score in Columns 1-3 and 7 of Table 4. If changes in peer composition would cause changes in assistant examiner grant rates, one would not expect to observe this assistant examiner response prior to the point in time in which the peer composition changed.¹⁷ The results confirm that there are indeed no observed lead effects of concern.

We next consider a falsification exercise based on an evaluation of peer influences on the use of lack-of-novelty rejections versus obviousness rejections. Both of these standards are similar in ensuring that patents not be granted to inventions that are effectively already present in society. However, obviousness determinations are commonly perceived as being more indeterminate and subjective in nature than lack-of-novelty rejections. With this greater scope for discretion, one would arguably expect to observe more markers of learning and peer influence in the case of obviousness rejections. We test this in Columns 1 and 2 of Table 5, estimating specifications similar to that estimated in Column 1 of Table 2 but where the dependent variable equals the incidence of any obviousness (or lack of novelty) rejection during the course of the examination process (even if the application is ultimately allowed in later stages of review) and where the peer score represents the inherent obviousness rate (or inherent lack of novelty rate) of the peer

¹⁷ Again, we elect not to take too fine-grained of a temporal approach here and choose to track behavior over a two-year period considering that the duration of examination reviews often span a period of time in excess of a year. Moreover, considering that peer grant scores are calculated based on the time of disposition of applications and that applications indeed take some time to process, it may be possible to observe some amount of anticipation effects (e.g., new examiners in an Art Unit exerting some influence on current examiners before those new examiners begin completing their first reviews). However, one might not expect any such anticipation effects to be substantial and, in fact, we do not observe strong anticipation effects anyway.

examiners. In the case of obviousness rejections, we find a pattern of early-career peer effects that diminish with examiner experience, very similar to the grant-rate results. In the case of lack-of-novelty rejections, we find little evidence of peer effects at any level of experience.

As an additional specification check, we further break down the pure peer score—the mean inherent grant rate of the peer assistant examiners—into more specific peer scores based on the experience of those peers. If the above findings are reflective of knowledge spillovers and learning, one might predict that the channel of influence would be weaker in the case of new peer examiners and stronger in the case of seasoned peer examiners. We explore these predictions in Columns 3 and 4 of Table 5, estimating separate specifications where we calculate a peer score based on peer assistant examiners who are in their first and second years at the Patent Office (Column 3) and who are in their third year and beyond at the Patent Office (Column 4). As predicted, we estimate a weak relationship between a new examiner’s grant rate and the inherent granting practices of her similarly situated junior peers, with a point estimate of the coefficient of this junior peer score being one-fourth of the size of the new hire peer-effect documented in Table 2 and one-fourth of the size of the senior peer effect estimated in Column 4 of Table 5.

Next, we consider another separation of the assistant examiner peer group based on the peer examiners’ participation in the Patent Office’s telecommuting program, whereby eligible examiners may work from home for all but 1 or 2 days during a bi-week period. Naturally, one may expect that the group of telecommuting examiners would have a weaker peer influence on new examiners (who are themselves not eligible to work from home) relative to the group of non-telecommuting examiners. Non-telecommuting peers are actually in the office day-to-day with new examiners and thus in a position to exert social influence in the first place. Knowledge spillovers, after all, have been argued to be more likely to occur through high-frequency repeated

social contact (Von Hippel 1994). Of course, telecommuting examiners may still have some influence, even if weaker, considering that they will typically spend some time in the Patent Office over a bi-week period and considering that these examiners are among the most senior in the Art Unit (i.e., a fact which bias against finding a differential).

We present the results of this final exercise in Columns 5 and 6 of Table 5. In the case of the non-telecommuting peers, we estimate a strong peer effect on new hires that, as above, diminishes with examiner experience. In the case of telecommuting peers, the point estimate for the coefficient of the peer score—capturing the effect of the telecommuting peers on new examiner grant rates—is around $\frac{1}{2}$ of the magnitude of the corresponding effect for the non-telecommuting examiners. These results lend further support to an interpretation of the above results as indeed arising from social interactions with peers.

Finally, we reiterate that the above approach has characterized peer granting tendencies by the mean inherent grant rate of surrounding peers. It is possible, however, that individuals may be influenced by different aspects of the distribution of inherent grant rates across peers—e.g., they may be especially influenced by the highest grant-rate examiners around them. We consider distributional effects of this nature in the Online Appendix, separately constructing peer scores based on the examiners at the different percentiles of the peer grant-rate distribution within the Art-Unit-by-year cell. Our results suggest that movements in peer composition at both the top and bottom of the distribution of peer granting tendencies are associated with peer effects on individual examiner behavior that are similar in magnitude and pattern to the mean effects estimated above. That is, for instance, an increase in the peer grant rate at the 25th percentile of the within-Art-Unit peer granting distribution is associated with a large (mean) increase in the grant rate of the affected examiner during the first two years of her career, followed by a weakening of this peer influence

as her career proceeds. As such, the peer influences we document do not simply arise from the highest grant-rate peers. However, we do find that the effect of changes in peer composition at the higher percentiles are associated with larger effects on individual examiners relative to corresponding changes in peer composition at the bottom of this distribution. For instance, we find that the point estimate of the effect on new examiners' grant rates of an increase in the peer score is 0.13 percentage-points—or roughly 37%—larger in absolute terms in the case of specifications that base peer scores on the 75th percentile of inherent peer grant rates relative to those that base peer scores on the 25th percentile. In other words, examiners may indeed be less influenced by the granting tendencies of their lower grant-rate peers.

Aside from these distributional considerations regarding which peers constitute the peer score, the primary specification estimated above is also arguably limited in its parametric and linear treatment of the peer grant score variable in estimating its influence on individual grant rates. Accordingly, in the Online Appendix, we also take a more non-parametric approach, whereby, instead of using the value of the inherent peer grant score as the key regressor of interest, we include a series of dummy variables capturing the incidence of that peer score falling into the various quartiles of the distribution of peer grant scores *across* Art Units.¹⁸ To simplify this exercise, we do not interact each of these dummies with the series of experience bins and instead estimate an experience-invariant specification. Our results demonstrate the robustness of our findings to this non-parametric alternative, with examiner grant rates rising monotonically across each of the peer grant score quartiles. In particular, we find that grant rates rise by 2.3, 7.7 and 9.6 percentage points as we ascend into the second, third, and fourth quartiles of the distribution of mean peer grant rates across the full sample. With a two-quartile jump being roughly on par with

¹⁸ For these purposes, we are still looking at the average peer to construct the peer score. But, we are then looking at the distribution of average peer scores across the different Art Units.

a one-standard deviation change in the peer grant rate, and in light of the standard-deviation interpretation of the primary results discussed above, the magnitude of these findings is consistent with the more parametric approach taken in the baseline specifications.

IV.D. General versus Specific Knowledge Flows

To the extent the above findings reflect some degree of knowledge flow among examiners, one might wonder the nature of such flows. Are examiners learning from other examiners general styles and strategies towards examination—e.g., general search strategies or general views towards the technological advancements necessary to surpass novelty and nonobviousness requirements? Or, are examiners imparting more specific information to each other that be driving some part of the above findings? For instance, are examiners learning about specific pieces of prior art from their peers? To shed some light on these questions—in particular, to test for the presence of this latter more specific channel—we collect data from 2000-2010 on the patents that are cited by each patent issued over that period, a dataset that allows us to explore citation patterns *by* a given patent and citations patterns *to* a given patent, subject to the limitations of the given time period.¹⁹

To explore whether examiners appear to be learning about specific patents from their peers, we conduct a simple exercise. First, for each patent issued at time t , we determine whether the examiner reviewing said application cited a patent among the set of “pet” patents most frequently cited throughout their careers by the other examiners also in the same Art Unit at time t , where we define an examiner’s set of “pet” prior art by the 10 patents that they most frequently cite throughout their career.²⁰ Motivating this approach is the observation by others in the literature that examiners frequently turn to the same set of patents (Abrams and Sampat 2017) as pieces of

¹⁹ Given this construction, naturally, the pool of patents that are targets of citations span a longer period of time considering that patents issued beyond 2000 will nonetheless cite older patents. We are grateful to Bhaven Sampat for providing us with these citation data.

²⁰ The findings documented below are not sensitive to this precise cut-off.

prior art when conducting their reviews, a set of personalized information and preferences that examiners may impart to their peers.

Of course, just observing that examiners cite, with some probability, patents that are among the favorite pieces of prior art of those peers is not, in and of itself, telling regarding the transmission of knowledge regarding the existence of those pieces of prior art. After all, peers work within the same area of technology and thus there is likely to be some degree of correlation in the set of patents that examiners within an Art Unit cite. Nonetheless, if examiners were indeed learning about specific pieces of prior art from their peers, we might predict that they would be more likely to learn from those peers around them regularly. That is, we might predict that the above citation likelihood would be stronger when confining the relevant peer group to those examiners in the same Art Unit at time t that are not telecommuting at such time relative to a situation when we confine the peer group to those telecommuting at time t .

Telecommuting examiners tend to be more experienced and of higher GS levels than non-telecommuting examiners, creating a concern that examiners may differentially turn to pet patents of their senior peers for reasons beyond just accessibility to those peers. Of course, this concern would tend to bias *against* finding an accessibility effect. Nonetheless, to appease this concern, we perform this comparison between the likelihood of citing to pet patents of one's telecommuting peers versus non-telecommuting peers while only looking at peer assistant examiners at either GS-level 12 or 13. We also confine this comparison to only those Art Units and years where there are at least one telecommuting examiner in the Art Unit-by-year cell.

With these restrictions, we find that examiners cite to the pet patents of their current non-telecommuting peers roughly 2.1 percent of the time, while only citing to the pet patents of their telecommuting peers roughly 0.25 percent of the time. That is, examiners are roughly 8.2 times

more likely to cite to the favorite patents of their local, accessible peers than those peers working from home. Of course, some difference here is to be expected given that at any point in time an examiner may be surrounded by more non-telecommuting peers than telecommuting peers. In fact, again when confining ourselves to Art-Unit-by-year cells in which there are at least one non-telecommuting examiner, there are nearly 6 time more non-telecommuting examiners than telecommuting examiners. Nonetheless, the differential in citation likelihoods exceeds this differential in peer examiner counts, suggesting that accessibility to peers may indeed impact the likelihood that an examiner may cite to the favorite patents of those peers, which in turn may suggest that peers do impart some degree of specific information to each other regarding pieces of prior art that may be relevant to each other's current applications.

One concern with this simple comparison of means of course is that Art Units may vary in the degree to which their examiners telecommute and the degree to which they cite from their peers' favorite patents—e.g., the above differential could just be a reflection of a situation in which a certain Art Unit happens to have a low propensity of telecommuting examiners but a generally high degree of citing to peers' favorite patents, though to the same degree across telecommuting and non-telecommuting peers within that Art Unit. As such, in Table 6, we formalize this comparison so as to better isolate the difference in citation likelihoods based on accessibility of peers. For these purposes, we stack two separate samples of individual patents issued by assistant examiners. The dependent variable across both such sub-samples indicates the incidence of the examiner associated with the issued patent citing a patent that is among the set of pet patents for the peer group at the relevant time. For the first sub-sample, this measure focuses on the incidence of citing a pet patent of their telecommuting peers. For the second sub-sample (stacked on top of the first), that dependent variable captures the likelihood of citing a pet patent of the non-

telecommuting peers. We also include a measure of the number of examiners associated with the relevant Art Unit-by-year cell. This measure for the first sub-sample reflects the number telecommuting examiners in the Art-Unit-by-year cell; for the second sub-sample, it reflects the number of non-telecommuting examiners. We take two approaches to parameterizing these examiners counts: first, a semi-parametric approach in which we include the examiner count and its square and, second, a more non-parametric approach in which we include a series of dummy variables indicating the various quartiles of the examiner count distribution.

With this structure, we then regress the likelihood that the examiner cites one of her peers' favorite patents on an indicator variable for the non-telecommuting status of that peer group, while including a control for the associated number of examiners (either telecommuting or non-telecommuting depending on the relevant sub-sample) and while including issued patent fixed effects. This effectively allows us to compare—within a given issued patent—the likelihood that the associated examiner cited to one of her non-telecommuting peers relative to one of her telecommuting peers while accounting for the stronger likelihood to do so based on the mix of non-telecommuting and telecommuting peers in the Art Unit at the time of issuance. With these layers of control, we find that examiners are roughly 0.3 to 0.4 percentage points—or roughly 19 to 25%—more likely to cite to their non-telecommuting peers' favorite patents.²¹

On a final note, we do not find a strong experience gradient in the degree to which examiners cite their peers' pet prior art. This may reinforce the point addressed above that examiners may engage in some degree of ongoing learning throughout their careers. This continued learning is perhaps more likely to consist of specific, technical information of the sort

²¹ These findings are robust to alternative approaches that control instead for Art Unit and year fixed effects or Art-Unit-by-year fixed effects (in lieu of patent-specific fixed effects).

explored in this sub-section, at least relative to the early career learning which may also involve a greater degree of learning over general practice styles.

IV.E. Peer Effects in Claim Narrowing

Our analysis thus far has primarily viewed the job of patent examiners as either allowing or rejecting patent claims. However, the iterative nature of the patent examination process— involving a back and forth between patent examiners and applicants over multiple rounds of review—creates the opportunity for a more nuanced dimension to the job of patent examiners: claim narrowing. The exclusionary power of a patent depends not just on the presence of that patent in the first place but also on the breadth of the claims underlying that patent. For instance, a patent on “skis” would tend to exclude a broader range of competition than a patent on “composite downhill skis.” As such, not only may examiners develop a practice style regarding their proclivities to allow or reject patent claims, but they may also develop a practice style regarding how they work with applicants to narrow their claims to the point that they comply with the legal patentability requirements (Kuhn and Thompson 2017).

Accordingly, in a final empirical exercise, we extend the above empirical framework to explore the relationship between the degree to which a given patent was narrowed throughout the examination process and the inherent narrowing proclivities of the associated peer group at the relevant time and in the relevant Art Unit. For these purposes, we collected data from Jeffrey Kuhn and Neil Thompson on the number of words added throughout the examination process to the first claims in the patents issued during our sample.²² Following, Kuhn and Thompson (2017),

²² The data received from Kuhn and Thompson, however, is subject to certain exclusions in that that they focused on patents issued subsequent to January 1, 2005, while also excluding continuation applications and applications in the biotechnology area. Given the resulting implications for our sample size, we elect with this exercise to estimate experience-interaction specifications that do not impose balance restrictions—that is, we focus on examiners over the first six years of their careers without imposing requirements to follow each include examiner throughout the full extent of those 6 years.

we use this word-added measure as a reflection of the degree of claim narrowing, considering that longer claims generally impose a greater number of conditions that must be met before patent infringement is found. In Panel A of Table 7, we present the results from this approach, whereas in Panel B, we present results from specifications that form the claim-narrowing variable in percentage terms—i.e., normalizing words added by the number of words in the first claim of issued patents. We also show results with and without the inclusion of SPE fixed effects.

Consistently across these specifications, we find little association between the degree of claim narrowing for a given patent and the inherent claim narrowing tendencies of the relevant peer examiners during the first 2 years of the affected examiner's career. However, as examiners move into the later experience bins—in their third year and beyond—the degree to which they add words to claims begins to be more strongly associated with the claim narrowing tendencies of their peers. For instance, from Panel B, during the third and fourth years of an examiner's career, we find an increase of roughly 0.15-0.17 in the number of words that an examiner adds to the first claim of an issued patent as that examiner experiences a change in the composition of her peer group that represents an increase of 1 word in the average inherent words-added of that peer group. The magnitude of that peer effect does not appear to increase further as the examiner moves into her fifth and sixth year at the Patent Office.

To summarize, we likewise find strong peer effects in claim narrowing practices of examiners; however, these peer influences do not appear to emerge until an examiner has garnered some degree of experience at the Patent Office. This delayed influence is not necessarily inconsistent with the granting-focused learning story set forth above, where peer influences were strongest early on in an examiner's career. Deciding to allow or reject patents is something that examiners are necessarily asked to do from the very beginning of their careers at the Patent Office.

Accordingly, the learning process over granting styles—and the potential scope for peers to help shape that learning—is something that will naturally commence immediately upon the onset of an examiner’s career. The task of affirmatively working with applicants to narrow their claims before allowing them is not necessarily a task that examiners must perform and thus is not a skill that they will necessarily begin to develop from day one at the Patent Office. It may take examiners some time to gather enough experience with reviewing applications and evaluating claims in any sense before they are even in a position to start developing a deliberate claim-narrowing practice style of this nature.

Consider the well-known idiom—one must learn to walk before they can run.²³ The same may simply be true for the sequence of learning facing patent examiners—that is, an examiner must learn to reject claims generally before learning to work with applicants to narrow claims to the point of legal permissibility. For this reason, it may not be unreasonable to think that there will be a delay in the onset of peer influences over claim-narrowing practices. If anything, in light of this theorized sequencing of learning behavior, any such observed delay in peer influences in the case of claim narrowing behavior may only reinforce an interpretation of the above documented peer effects on grant rates as arising from a learning / knowledge spillover mechanism.

To support this interpretation of a delayed peer effect as arising from a delay in developing this claim narrowing skill in the first place, we also estimate a simple regression of the number of words added to an issued patent on a series of examiner experience bins in addition to various controls, including examiner fixed effects, year fixed effects, Art Unit fixed effects, examiner GS-level fixed effects, applicant entity size and application foreign priority status. In Figure A2 of the

²³ That is, someone learning to walk may be influenced by the walking styles of those around them during those critical moments, but they may be less influenced by the running styles of those around when they are simply learning to walk. At some point though, they reach a level of comfort with walking that they can begin to start working on the more nuanced challenge of running, and perhaps it is at this point that they may begin to be influenced by the running styles of their peers at those times.

Online Appendix, we plot the estimated coefficients of the experience group dummies. We find that as examiners ascend from the first experience group (0-2 years) to the second experience group, there is an increase of roughly 9 in the average number of words that are added throughout the prosecution process of the patents they issue, representing a roughly 18% increase over the mean. As such, consistent with the notion that examiners may delay in looking to their peers for guidance on claim narrowing strategies because they may be holding off developing such skills in the first place, we find that examiners indeed narrow claims to a much smaller degree at the beginning of their career.

V. Conclusion

Knowledge spillovers have been central to many models of economic growth and technological change (Krugman 1991, Romer 1986, Lucas 1988). While much of the theoretical and empirical discussions surrounding knowledge spillovers have focused on knowledge transmission *across* firms or across geographical units (Audretsch and Feldman 2002), knowledge transmissions *within* firms may also contribute significantly to these same macroeconomic outcomes, in addition to the productivity outcomes of individual firms (Jackson and Bruegmann 2009). Our analysis has attempted to overcome some of the key empirical challenges involved with testing for the presence and degree of spillovers within firms. In the process, we have found strong evidence that a worker's practice style may be shaped early in her career and that one of the key factors shaping her behavior is the corresponding practice styles of her peers during those critical early moments of impressionability. Though our analysis focuses on just one employment setting—the U.S. Patent Office—its analysis demonstrates just how strong of a role that peers can play in high skilled work settings, even when focusing on work tasks that are somewhat isolated and non-team-based in nature.

Regardless of the generalizability of these findings beyond patent examiners, the findings hold various implications for U.S. patent policy. Much attention has been paid to the levels and variability of grant rates produced by examiners. Policies designed to remedy any harms resulting from these patterns of behavior must start with an understanding of the sources of such behaviors. The fact that examiners may learn so substantially from their peers—perhaps more so than from their supervisors—may be important for such purposes. Among other things, that knowledge may help the Patent Office in how it allocates examiners to Art Units—for example, the mix of junior and senior examiners it wants to maintain within Art Units, or the placement of particularly generous or harsh examiners. It may also bear on how the Agency wishes to structure and allocate training efforts among the different classes of examiners within the Agency.

Our findings also hold implications for a particular personnel policy within the Patent Office—i.e., its telecommuting program. To understand this connection, first consider one of the key parallels between our analysis of within-firm knowledge spillovers and the literature on spillovers across geographical units: the concept of proximity. A number of studies (for instance, Jaffe, Trajtenberg and Henderson 1993, Audretsch and Feldman 1996) have demonstrated the critical importance of geographical closeness in the transmission of knowledge and have generally demonstrated the degree of regional clustering and concentration in innovative activity. The findings from these geography-focused studies are arguably consistent with our findings of stronger peer effects in the case of peers that do not telecommute and that are present day-to-day in the Patent Office—i.e., more proximate peers. The significant effects from proximate workers within firms may signal a strong role for peers to play in overall workplace efficiency. If peer effects are properly overseen and directed (so as to produce positive and not negative spillovers), the ability of new workers to learn from their peers can lead to potentially substantial productivity

gains. Such gains may be dampened to the extent that workers no longer interact in person. As such, our analysis sheds light on some of the consequences that may befall the general movement we have observed across a number of economic sectors to allow employees to work from home. While telecommuting may reduce a number of transaction costs for firms and for employees, it may impose transaction costs in the transmission of tacit knowledge (Von Hippel 1994). Whether the former gains outweigh the latter losses is a subject for the future and ongoing research of the consequences of telecommuting (Bloom et al., 2015, Frakes and Wasserman 2016, Mas and Pallais 2016).

REFERENCES

- Allen, Natalie J., and John P. Meyer. 1990. "Organizational Socialization Tactics: A Longitudinal Analysis of Links to Newcomers' Commitment and Role Orientation," *Academy of Management Journal* 33: 847-858.
- Ashforth, Blake K., and Alan M. Saks. 1996. "Socialization Tactics: Longitudinal Effects on Newcomer Adjustment," *Academy of Management Journal* 39(1): 149-178.
- Audretsch, D. B. and M. P. Feldman. 1996. "R&D spillovers and the geography of innovation and production," *American Economic Review* 86(4): 253-273.
- Audretsch, D.B. and M.P. Feldman. 2004. "Knowledge spillovers and the geography of innovation," Chapter 61 in Handbook of Regional and Urban Economics: 2713-2739
- Azoulay, Pierre, Joshua S. Graff Zivin, Jialan Wang. 2010. "Superstar Extension," *Quarterly Journal of Economics* 125 (2): 549-589.
- Baron, J. N., Burton, M. D., & Hannan, M. T. 1999. "Engineering bureaucracy: the genesis of formal policies, positions, and structures in high-technology firms." *Journal of Law, Economics, and Organization*: 15(1): 1-41.

- Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying. 2015. "Does Working from Home Work? Evidence from a Chinese Experiment," *Quarterly Journal of Economics* 130(1): 165-218.
- Chetty, Raj, John N. Friedman, and Jonah R. Rockoff. 2014. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates," *American Economic Review* 104(5): 406-11.
- Cockburn, I., S. Kortum, and S. Stern. 2003. "Are All Patent Examiners Equal? Examiners, Patent Characteristics, and Litigation Outcomes," in W. M. Cohen and S.A. Merrill (Eds.), *Patents in Knowledge-Based Economy* (Washington, DC: National Academies Press, 2003).
- Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg. 2017. "Peer Effects in the Workplace," *American Economic Review* 107(2): 425-56.
- David, Paul. A. 1985. "Clio and the Economics of QWERTY", *American Economic Review* 72: 332-7.
- Dokko, Gina, Steffanie L. Wilk, and Nancy P. Rothbard. 2009. "Unpacking Prior Experience: How Career History Affects Job Performance," *Organization Science* 20: 51-68.
- Frakes, Michael D. and Melissa F. Wasserman. 2016. "Patent Office Cohorts," *Duke Law Journal* 65: 1601-1655.
- Frakes, Michael, and Melissa Wasserman. 2016. "Procrastination in the Work Place: Evidence from the U.S. Patent Office," NBER Working Paper 22987.
- Frakes, Michael, and Melissa Wasserman. 2017. "Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Microlevel Application Data," *The Review of Economics and Statistics* 99(3): 550-63.

- Galasso, Alberto and Mark Schankerman. 2015. "Patents and Cumulative Innovation: Causal Evidence from the Courts," *Quarterly Journal of Economics* 130 (1): 317-369.
- Gordon J. DiRenzo, 1977. "Socialization, Personality, and Social Systems," *Annual Review of Sociology* 3: 261-295.
- Gould, Eric, and Eyal Winter. 2009. "Interactions between Workers and the Technology of Production: Evidence from Professional Baseball," *Review of Economics and Statistics* 91(1): 88-100.
- Guryan, Jonathan, Kory Kroft, and Matthew J. Notowidigdo. 2009. "Peer effects in the workplace: evidence from random groupings in professional golf tournaments," *American Economic Journal: Applied Economics* 1 (4), 34–68.
- Ho, Daniel. 2017. "Does Peer Review Work? An Experiment of Experimentalism," *Stanford Law Review* 69(1): 1-119.
- Jackson, C. Kirabo, and Elias Bruegmann. . 2009. "Teaching Students and Teaching Each Other: The Importance of Peer Learning for Teachers," *American Economic Journal: Applied Economics* 1(4): 85-108.
- Kane, Thomas J., and Douglas O. Staiger. 2008. "Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation," NBER Working Paper No. 14607.
- Krugman, Paul. 1991. *Geography and Trade* (MIT Press: Cambridge, MA).
- Kuhn, Jeffrey, and Neil Thompson. 2017. "How to Measure and Draw Causal Inferences with Patent Scope," *International Journal of the Economics of Business*, forthcoming.
- Lemley, Mark A. and Bhaven Sampat. 2012. "Examiner Characteristics and Patent Office Outcomes," *The Review of Economics and Statistics* 94(3): 817–827.

- Lichtman, Douglas. 2004. "Rethinking Prosecution History Estoppel," *University of Chicago Law Review*: 151–82.
- Lucas, Robert. 1988. "On the mechanics of economic development," *Journal of Monetary Economics* 22: 3-39.
- Malmendier, Ulrike, and Stefan Nagel, 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?," *The Quarterly Journal of Economics* 126(1): 373-416
- Mann, Ronald. 2014. "The Idiosyncrasy of Patent Examiners: Effects of Experience and Attrition," *Texas Law Review*: 2149–76.
- Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem," *The Review of Economic Studies* 60(3): 531-542.
- Marquis, Christopher, and András Tilcsik. 2013. "Imprinting: Toward A Multilevel Theory," Working Paper, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2198954.
- Mas, Alexandre, and Enrico Moretti. 2009. "Peers at Work," *American Economic Review* 99(1): 112-45.
- Nordhaus, William, *Invention, Growth, and Welfare* (Cambridge, MA: MIT Press, 1969).
- Righi, Cesare, and Timothy Simcoe, "Patent Examiner Specialization," mimeo (2017).
- Romer, Paul. 1986. "Increasing returns and long-run growth," *Journal of Political Economy* 94(5): 1002-37.
- Sampat, Bhaven and Heidi L. Williams. 2014. "How do Patents Affect Follow-on Innovation? Evidence from the Human Genome," available at <http://economics.mit.edu/files/9778>.
- Scotchmer, Suzanne. 1991. "Standing on the Shoulder of Giants: Cumulative Research and the Patent Law," *Journal of Economic Perspectives* 5: 29–41.

Von Hippel, Eric. 1994. Sticky information and the locus of problem solving: implications for innovation, *Management Science* 40: 429-439.

Waldinger, Fabian 2012. "Peer Effects in Science: Evidence from the Dismissal of Scientists in Nazi Germany," *Review of Economic Studies* 79(2): 838-61.

TABLE 1. SUMMARY STATISTICS

	(1)	(2)
	Mean	SD
Panel A. Grant Rate Sample		
Grant	0.695	0.461
Any Obviousness Rejection	0.855	0.352
Any Lack-of-Novelty Rejection	0.662	0.473
Examiner Experience: 1-2 Years	0.102	0.302
Examiner Experience: 3-4 Years	0.155	0.362
Examiner Experience: 5-6 Years	0.158	0.365
Examiner Experience: 6+ Years	0.585	0.493
Assistant Examiner	0.447	0.497
Assistant Examiner Peer Score (Grant Rate)	0.651	0.178
Primary Examiner Peer / Supervisor Score (Grant Rate)	0.772	0.303
SPE Score (Grant Rate)	0.782	0.192
Panel B. Claim Narrowing Sample		
Number of Words Added to First Claim throughout Prosecution of Issued Patents	47.854	62.490
Assistant Examiner Peer Score (Words Added to First Claim)	56.366	26.368

Each observation in Panel A is a given application from the PAIR database that reached a final disposition and that was published in the PAIR records between March, 2001 and July, 2012. Each observation in Panel B is a given issued patent from the Kuhn and Thompson (2017) database matched with the PAIR records from Panel A.

TABLE 2. EFFECTS OF PEER AND SUPERVISOR GRANTING TENDENCIES ON ASSISTANT EXAMINER GRANT RATES, BY YEARS OF ASSISTANT EXAMINER EXPERIENCE

	(1)	(2)	(3)	(4)	(5)	(6)
	Pure Peer Effects (Assistant Examiner Effects)		Quasi-Supervisory Effects (Primary Examiner Effects)		Supervisory Effects (SPE Effects)	
Peer Score	0.426*** (0.075)	0.401*** (0.057)	0.482*** (0.104)	0.341*** (0.056)	0.314*** (0.078)	0.196*** (0.061)
(Omitted: Peer Score X 0-2 Years Experience)						
Peer Score X 2-4 Years Experience	-0.173*** (0.031)	-0.161*** (0.022)	-0.219*** (0.041)	-0.210*** (0.029)	-0.135*** (0.049)	-0.088*** (0.029)
Peer Score X 4-6 Years Experience	-0.182*** (0.049)	-0.191*** (0.037)	-0.298*** (0.068)	-0.312*** (0.049)	-0.190** (0.077)	-0.169*** (0.051)
N	153906	415575	153584	413499	68063	183,268
Balanced Sample?	YES	NO	YES	NO	YES	NO

* significant at 10%; ** significant at 5%; *** significant at 1%. Standard errors are reported in parentheses and are clustered to correct for autocorrelation within given Art Units. Coefficients of the experience group dummies are omitted for purposes of brevity. Each observation is a given application from the PAIR database that reached a final disposition and that was published in the PAIR records between March, 2001 and July, 2012. Each specifications tracks the granting decisions of assistant examiners (GS-level 13 and below) over the first six years of their careers at the Patent Office. Columns 1, 3, and 5 focus on a balanced set of examiners that we can observe practicing at the Patent Office over the entirety of their first six years at the Patent Office. Columns 2, 4, and 6 present results from an unbalanced sample that imposes no such restrictions (only that we restrict the sample to observations within the first six years of experience). All specifications include examiner fixed effects, Art Unit fixed effects, year fixed effects and controls for various application-level characteristics

TABLE 3. EFFECTS OF PEER GRANTING TENDENCIES ON ASSISTANT EXAMINER GRANT RATES: VARIOUS ROBUSTNESS CHECKS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Peer Score	0.334*** (0.088)	0.327*** (0.067)	0.443*** (0.062)	0.329*** (0.048)	0.486*** (0.074)	0.517*** (0.089)	0.906*** (0.213)	0.456*** (0.056)
(Omitted: Peer Score X 0-2 Years Experience)								
Peer Score X 2-4 Years Experience	-0.163*** (0.034)	-0.148*** (0.048)	-0.122*** (0.036)	-0.177*** (0.027)	-0.122*** (0.030)	-0.140*** (0.032)	-0.183*** (0.052)	-0.190*** (0.020)
Peer Score X 4-6 Years Experience	-0.165*** (0.055)	-0.214*** (0.069)	-0.216*** (0.049)	-0.199*** (0.043)	-0.124*** (0.046)	-0.136*** (0.050)	-0.354*** (0.082)	-0.236*** (0.032)
Peer Score X 7+ Years Experience	-	-	-	-	-	-	-	-0.251*** (0.048)
N	145804	152745	152841	150504	153905	153905	68063	521275
Treatment of Art Unit and Time Effects	Art Unit and Year Effects	Art-Unit- by-Year Fixed Effects	Art-Unit- by-Bi-Year Fixed Effects	Art Unit and Year Effects	Art Unit and Year Effects	Art Unit and Year Effects	Art Unit and Year Effects	Art Unit and Year Effects
SPE Dummies?	YES	NO	NO	NO	NO	NO	NO	NO
Balanced or Unbalanced?	Balanced	Balanced	Balanced	Balanced	Balanced	Balanced	Balanced	Unbalanced
Construction of Peer Grant Score at Year t	Lifetime Grant Rates	Lifetime Grant Rates	Lifetime Grant Rates	Grant Rate for Years Prior to t	Estimated Examiner Fixed Effects	Empirical Bayesian Estimator	Lifetime Grant Rates	Lifetime Grant Rates
Limit to Art-Unit- Year Cells With Data on SPE Grant Rate?	NO	NO	NO	NO	NO	NO	YES	NO

* significant at 10%; ** significant at 5%; *** significant at 1%. Standard errors are reported in parentheses and are clustered to correct for autocorrelation within given Art Units. Coefficients of the experience group dummies are omitted for purposes of brevity. Each observation is a given application from the PAIR database that reached a final disposition and that was published in the PAIR records between March, 2001 and July, 2012. Each specifications tracks the granting decisions of assistant examiners (GS-level 13 and below) over the first six years of their careers at the Patent Office (except for Column 7 which tracks them over their whole careers). In addition to the indicated features of the estimated specifications, all specifications include examiner fixed effects and controls for various application-level characteristics.

TABLE 4. DISTRIBUTED LEADS AND LAGS SPECIFICATIONS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
2-Year Lead Score	0.057 (0.061)	0.056 (0.041)	0.009 (0.042)	-	-	-	0.043 (0.044)
Contemporaneous Peer Score	0.302*** (0.082)	0.332*** (0.064)	0.151*** (0.069)	0.253*** (0.079)	0.234*** (0.048)	0.112** (0.052)	0.191*** (0.069)
2-year Lagged Peer Score	-	-	-	0.118* (0.060)	0.141*** (0.043)	0.138*** (0.053)	0.139** (0.068)
N	131575	409752	388813	116812	374417	360708	286041
Balanced Sample (Over first 6 Years of Career)?	YES	NO	NO	YES	NO	NO	NO
Other Restrictions?	NO	NO	NO	Limit to Examiners in their 3 rd -6 th Years	Limit to Examiners Beyond their Second Years	Limit to Examiners Beyond their Second Years	Limit to Examiners Beyond their Second Years
SPE Effects?	NO	NO	YES	NO	NO	YES	YES

* significant at 10%; ** significant at 5%; *** significant at 1%. Standard errors are reported in parentheses and are clustered to correct for autocorrelation within given Art Units. Each observation is a given application from the PAIR database that reached a final disposition and that was published in the PAIR records between March, 2001 and July, 2012. Each specifications tracks the granting decisions of assistant examiners (GS-level 13 and below) over the indicated years of their careers at the Patent Office. All specifications include examiner fixed effects, Art Unit fixed effects, year fixed effects and controls for various application-level characteristics

TABLE 5. EFFECTS OF PEER GRANTING TENDENCIES ON ASSISTANT EXAMINER GRANT RATES: VARIOUS FALSIFICATION EXERCISES

	(1)	(2)	(3)	(4)	(5)	(6)
	Incidence of Any Obviousness Rejection	Incidence of Any Lack-of- Novelty Rejection	Peer Group: Assistant Examiners With Less than 2 Years of Experience	Peer Group: Assistant Examiners With 2 or More Years of Experience	Peer Group: Non- Teleworking Assistant Examiners (2006+)	Peer Group: Teleworking Assistant Examiners (2006+)
Peer Score	0.192*** (0.073)	-0.035 (0.079)	0.098** (0.047)	0.400*** (0.072)	0.462*** (0.075)	0.244*** (0.082)
(Omitted: Peer Score X 0-2 Years Experience)						
Peer Score X 2-4 Years Experience	-0.119*** (0.035)	0.021 (0.050)	-0.082*** (0.034)	-0.185*** (0.031)	-0.166*** (0.038)	-0.116*** (0.057)
Peer Score X 4-6 Years Experience	-0.133*** (0.059)	0.032 (0.080)	-0.065* (0.048)	-0.201*** (0.049)	-0.297*** (0.055)	-0.232*** (0.069)
N	136654	136701	135314	152659	131629	85473

* significant at 10%; ** significant at 5%; *** significant at 1%. Standard errors are reported in parentheses and are clustered to correct for autocorrelation within given Art Units. Coefficients of the experience group dummies are omitted for purposes of brevity. Each observation is a given application from the PAIR database that reached a final disposition and that was published in the PAIR records between March, 2001 and July, 2012. Each specifications tracks the granting decisions of assistant examiners (GS-level 13 and below) over the first six years of their careers at the Patent Office. Each specification focuses on a balanced set of examiners that we can observe practicing at the Patent Office over the entirety of their first six years at the Patent Office. All specifications include examiner fixed effects, Art Unit fixed effects, year fixed effects and controls for various application-level characteristics

TABLE 6: RELATIONSHIP BETWEEN LIKELIHOOD THAT ASSISTANT EXAMINER WILL CITE TO SET OF “PET” / FAVORITE PATENTS OF HER PEER GROUP AND AN INDICATOR VARIABLE FOR THE NON-TELECOMMUTING STATUS OF THAT PEER GROUP (RELATIVE TO THE TELECOMMUTING STATUS OF THAT PEER GROUP)

	(1)	(2)
Non-Tele-commuting Peer Group	0.003*** (0.001)	0.004*** (0.001)
N	326460	326460
Coefficient of Non-Tele-commuting Peer Group as a Fraction of Mean of Dependent Variable	0.19	0.25
Sample	Sample of Issued Patents with Information on Telecommuting Peer Group Stacked on Sample of Issued Patents with Information on Non- Telecommuting Peer Group	
Parameterization of Controls for Count of Telecommuting and Non-Telecommuting Examiners	Relevant Examiner Count and its Square	Dummies for Different Quartiles of Relevant Examiner Count
Issued Patent Fixed Effects?	YES	YES

* significant at 10%; ** significant at 5%; *** significant at 1%. Standard errors are reported in parentheses and are clustered to correct for autocorrelation within given Art Units. There are two observations for each issued patent in our sample (based on applications from the PAIR database that reached a final disposition and that was published in the PAIR records between March 2001 and July 2012). The dependent variable captures the likelihood that the examiner associated with the given patent cited to the set of “pet” / favorite patents frequently cited by the peers in the relevant Art Unit-by-year cell, where the relevant peer group is the set of non-telecommuting examiners (at GS level 12 or 13) for the first observation within each issued patent and the set of telecommuting examiners (at GS-level 12 or 13) for the second observation within each issued patent. We then regress the likelihood that the examiner cited a “pet” patent of the relevant peer group on an indicator for whether the relevant peer group represents the non-telecommuting peers, along with a set of issued patent fixed effects. Each regression controls for the number of examiner in the relevant Art Unit-by-year cell, where this measure reflects the number of non-telecommuting examiners (at GS-level 12 or 13) for the first observation within each issued patent and the number of telecommuting examiners (at GS_level 12 or 13) for the second observation within each issued patent. Columns 1 and 2 reflect the indicated treatment of these examiner count controls. The set of issued patents considered are confined to those issued by assistant examiners.

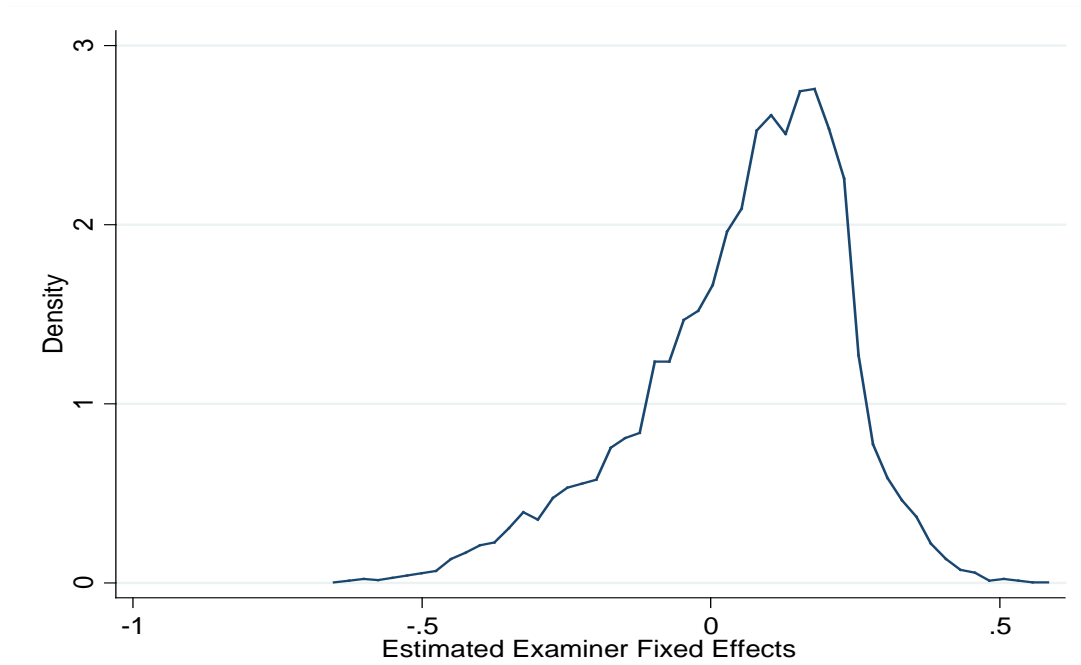
TABLE 7: EFFECT OF CHANGES IN INHERENT PEER NARROWING SCORES ON THE DEGREE OF INDIVIDUAL EXAMINER NARROWING

	(1)	(2)
Panel A. Claim Narrowing Measure: Percentage Increase in Number of Words in First Claim throughout Prosecution of Issued Patents		
Peer Score	0.002 (0.164)	-0.188 (0.281)
(Omitted: Peer Score X 0-2 Years Experience)		
Peer Score X 2-4 Years Experience	0.215* (0.124)	0.375** (0.176)
Peer Score X 4-6 Years Experience	0.221** (0.106)	0.422*** (0.159)
Panel B. Claim Narrowing Measure: Number of Words Added to First Claim throughout Prosecution of Issued Patents		
Peer Score	0.093 (0.089)	0.001 (0.094)
(Omitted: Peer Score X 0-2 Years Experience)		
Peer Score X 2-4 Years Experience	0.149*** (0.040)	0.163*** (0.043)
Peer Score X 4-6 Years Experience	0.150*** (0.050)	0.173*** (0.055)
N	142912	136916
SPE Effects?	NO	YES

* significant at 10%; ** significant at 5%; *** significant at 1%. Standard errors are reported in parentheses and are clustered to correct for autocorrelation within given Art Units. Each observation is a given issued patent from the Kuhn and Thompson (2017) dataset. Specifications are limited to applications reviewed by assistant examiners during the first six years of their career, though, for sample size purpose, we do not impose strong balance conditions that we only track examiners who we can observe over those full six years. All specifications include examiner fixed effects, Art Unit fixed effects, year fixed effects and controls for various application-level characteristics.

FIGURE 1

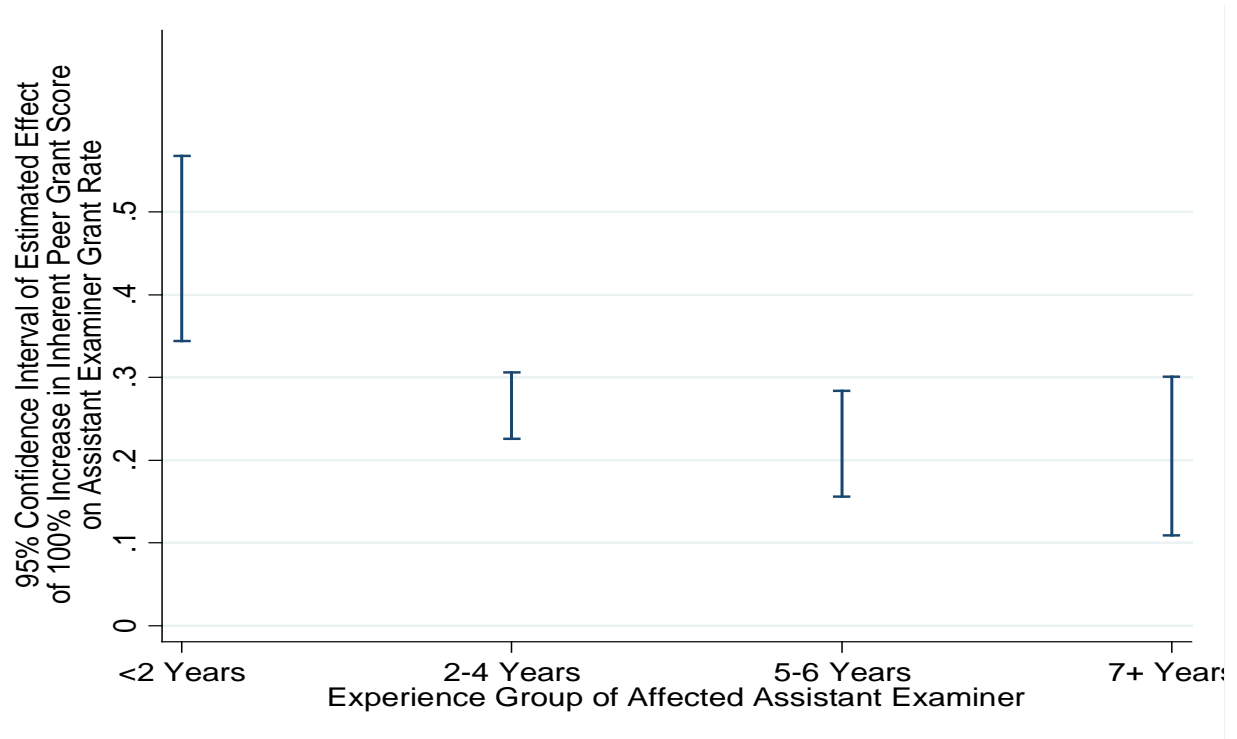
DISTRIBUTION OF ESTIMATED EXAMINER FIXED EFFECTS



Note: this figure presents a kernel density plot (Epanechnikov kernel with “optional” bandwidth) of estimated examiner fixed effects across all examiners in the sample. Examiner fixed effects are derived from the predicted values from a regression of the incidence of the application being granted on a series of an examiner fixed effects, along with year effects, examiner GS levels, examiner experience levels and various application-level characteristics (large entity status of applicant, foreign priority status of applicant, and duration of examination and its square).

FIGURE 2

EFFECT OF INHERENT PEER GRANTING TENDENCIES ON ASSISTANT EXAMINER GRANT RATE, BY YEARS OF EXPERIENCE OF THE AFFECTED ASSISTANT EXAMINER



Notes: this figure presents the results of the coefficients estimated in Column 7 of Table 3.