KNOWLEDGE SPILLOVERS AND LEARNING IN THE WORKPLACE: EVIDENCE FROM THE U.S. PATENT OFFICE[1]

BY: MICHAEL D. FRAKES AND MELISSA F. WASSERMAN

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

[1] Frakes (corresponding author): 210 Science Drive, Box 90362, Durham, NC 27708 (michael.frakes@law.duke.edu). 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.

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 offer 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 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.[2] 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. Secondarily, we estimate specifications with a rich set of Art-Unit-by-year fixed effects (or, alternatively, Art-

[2] 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.

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 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.[3] 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

[3] 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. 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).

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 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.