

THE NBER/SLOAN PROJECT ON INDUSTRIAL TECHNOLOGY AND PRODUCTIVITY: INCORPORATING LEARNING FROM PLANT VISITS AND INTERVIEWS INTO ECONOMIC RESEARCH[†]

Knowledge Spillovers and Patent Citations: Evidence from a Survey of Inventors

By ADAM B. JAFFE, MANUEL TRAJTENBERG, AND MICHAEL S. FOGARTY*

It is well understood that the non-rival nature of knowledge as a productive asset creates the possibility of “knowledge spillovers,” whereby investments in knowledge creation by one party produce external benefits by facilitating innovation by other parties. At least since Zvi Griliches’s (1979) seminal paper on measuring the contributions of R&D to economic growth, economists have been attempting to quantify the extent and impact of knowledge spillovers. One line of research of this type has utilized patent citations to identify a “paper trail” that may be associated with knowledge flows between firms.¹

Very little of this research has attempted to determine the modes or mechanisms of communication that actually permit knowledge to flow. Further, most of the work has simply assumed that citations or other proxies are sufficiently correlated with knowledge flows to allow statistical analysis of the proxies to

be informative regarding the underlying phenomenon of interest.

This paper reports on a preliminary attempt to improve this situation. The idea for this survey came from R&D managers whom we were interviewing to test whether the picture of knowledge flows produced by patent citations was consistent with the managers’ impressions. One of these managers commented that he could not let us talk to the scientists who worked for him, but that he could not stop us from contacting them via the postal address that appears on their issued patents.

The survey results suggest that communication between inventors is reasonably important, and that patent citations do provide an indication of communication, albeit one that also carries a fair amount of noise.

I. Survey Design

We surveyed two groups, one in which we asked inventors about citations made in their patents to previous patents (the “citing inventor” survey), and one in which we asked inventors about citations received by their patents from subsequent patents (the “cited inventor” survey).² For the *citing* inventor survey, we asked the inventor a series of questions about the extent, timing, and nature of communication that they had with the inventors of three previous patents. Two of these previous patents were patents that appeared among the citations made

[†] *Discussants:* Robert J. Gordon, Northwestern University and NBER; Samuel S. Kortum, Boston University and NBER.

* Jaffe: Brandeis University, Waltham, MA 02454-9110, and NBER (e-mail: ajaffe@brandeis.edu); Trajtenberg: Eitan Berglas School of Economics, Tel-Aviv University, Tel-Aviv 69978, Israel, and NBER (e-mail: manuel@ccsg.tau.ac.il); Fogarty: Center for Regional Economic Issues, Wickenden 311, Case-Western Reserve University, Cleveland, OH 44106 (e-mail: msf3@post.cwru.edu). We thank Richard Anderson for his help in editing and administering the survey, and Melinda Miller and Margeret Lister Fernando for research assistance.

¹ Patent citations or references appear on the front page of a granted patent. They serve the legal function of identifying “prior art” upon which the current invention builds. For more detail, see Jaffe et al. (1993).

² The survey instruments were “pretested” on about 20 experienced inventors, and the survey questions were refined in response to their questions and comments. Space limitations preclude a detailed description of the survey and survey samples. For these details, see Jaffe et al. (2000).

by the surveyed inventor's patent. The third previous patent was a "placebo" patent that was not cited by the surveyed inventor, but which was granted in the same patent class and year as one of the actually cited patents. In the survey questionnaire this placebo was not identified or distinguished in any way; all three of the earlier patents were referred to as "cited patents."

For the *cited* inventor survey, we picked one of the actually cited patents about which the citing inventor was questioned, and we contacted the primary inventor. This cited inventor was asked about communication with the citing inventor, and was also asked to form a judgment, based on his or her reading of the citing patent, about the likelihood that the citing inventor had utilized knowledge represented by the cited inventor's patent. Finally, both the citing and cited inventors were also asked a series of questions about the economic and technological significance of their inventions.

We received 166 partial or complete responses to the *citing* survey, based on patents granted to these inventors in 1993. We received 214 partial or complete responses to the *cited* inventor survey, based on patents granted to these inventors between 1985 and 1993.

II. Results Regarding Extent, Timing, and Nature of Communication

We discuss here the responses to four questions asked of the citing inventors about their cited and placebo patents. One of the questions sought responses on a Likert scale to a question regarding the overall degree of familiarity of the citing inventor with the cited invention. For the patents that were in fact cited, 28 percent of the responses indicated a 4 or 5 on the Likert scale, indicating a significantly high familiarity; just under half of the respondents rated their familiarity at the low end of the scale. In contrast, over 80 percent of the respondents rated their familiarity with the placebo patent at the lowest possible level.

Another question related to when citing inventors learned about the cited invention. For the "true" citations, about 38 percent of respondents indicated that they had learned about the cited invention either before or during the development of their own invention. About one-

third indicated that they had learned about it after essentially completing their invention, including cases in which they learned about it during the preparation of their own patent application. A little less than one-third of respondents indicated that, despite the presence of the patent citation, they had not learned about the cited invention before receiving our survey. This is not surprising, because citations to inventions unknown to the inventor can be generated by the inventor's patent attorney or the patent-office examiner.

A third question relates to the mode of knowledge spillover. Even for the true citations, only about 18 percent of citing inventors indicated that they had had either direct communication or had been exposed to some kind of presentation or demonstration of the cited invention. Another 18 percent indicated that they learned through "word of mouth" or had read the patent document itself. Consistent with the answers regarding timing, almost 40 percent indicated that it was the process of their own patent application that had caused them to learn of the previous invention.

The last question in this section, perhaps ambitiously, tried to get at the issue of the nature of assistance that the citing inventor may have received from the cited invention. Respondents were given a set of choices that we thought possible and were also invited to "write in" their own responses. Again, about 60 percent of the respondents indicated some specific way in which they had benefited from the cited invention; the single most common response was that the cited invention represented a concept that could be improved upon.

The other responses provided by the inventors also provide some insight into the nature of possible interactions. Examples include the following:

"The technology from patent 1 was incorporated in the product which used my invention."

"[N]ew market for our new technology!"

"The other patents gave credibility to our idea—they showed our ideas were 'feasible' to the people not intimately involved in our idea."

Other explanations confirmed that many citations derive from the patent process and probably are not related to any spillover:

“[D]id not learn of patents before filing—therefore these patents were not a factor in our work.”

“[A] patent cited by the patent examiner with no direct ties to my patent.”

Assuming that these responses can be taken at face value, they suggest that a significant, but not overwhelming, fraction of the “links” indicated by a patent citation correspond to some kind of spillover. Across the different aspects captured by each of these questions, typically one-quarter of the responses correspond to a fairly clear spillover; perhaps one-half of the answers indicate no spillover, and the remaining quarter indicate some possibility of a spillover. It is also clear that addition of citations by the inventor’s patent lawyer or the patent examiner is the primary reason for citations to patents unknown to the inventor.

In order to test formally whether the respondents’ answers relating to citations are different from those for the placebos, we constructed a composite spillover index for each cited patent, using the answers to all four questions. This index was constructed by consolidating the possible answers to each question to produce a score of 0, 1 or 2, and then adding these scores across the four questions. We then undertook an ordered probit analysis of this score, using as regressors variables that would seem likely to foster communication between the cited and citing inventors, variables that might foster the inventor’s remembering that communication occurred, other controls, and a dummy variable for whether the score pertains to a true citation as opposed to a placebo. Results of these analyses are presented in Jaffe et al. (2000).

Overall, the results confirm that citations can be interpreted as providing a (noisy) signal of spillovers. The difference in spillover score between the citations and placebos is quantitatively and statistically significant. The other variables generally have plausible and often significant effects. Overall, the spillover score is higher if the cited patent is more recent. Interestingly, analyses that separate the true citations

from the placebos show that this combined effect mixes a significantly positive effect for the citations with a significantly negative effect for the placebos. For the citations, this is consistent with more recent patents being more useful, and older citations being more likely to be non-spillovers included by the lawyer or examiner. It could also reflect the possibility that the inventor’s memory of actual communication is better with respect to more recent technology. Conversely, for the placebos, the spillover index is lower the more recent the cited invention. Since these represent patents that were not, in fact, cited, there should not have been communication; thus the negative coefficient for the placebos is consistent with the inventors’ giving more accurate answers with respect to more recent patents, and more often “mistakenly” indicating communication with respect to older patents.

We included the (log of) total citations received by the cited patent to control for the overall “importance” of that patent. It exhibits a positive effect, meaning that more important patents are perceived to have generated greater spillovers, whether because the spillovers are truly greater or simply more likely to be remembered by the respondent. Similarly, cited patents whose inventors reside in the same state as the citing inventor are perceived to have generated greater spillovers. These two variables are not significant when looking only at the placebos, further confirming that citations are meaningful, in the sense that the perceived extent of spillovers is correlated with things that ought to be correlated with spillovers for the citations but is uncorrelated with these things for the placebos.

III. Citations and Perceived Importance

In addition to the use of individual citation links as possible evidence of knowledge flow, a number of authors have utilized the total number of citations received by a patent as an indicator of the relative significance of patents. Both our *citing* and *cited* inventor surveys asked the inventors to rate the “technological significance” and the “economic importance” of the inventions, and also asked whether the patent had been licensed and whether it had been commercialized. In Jaffe et al. (2000), we present results of regressions of the log of total citations

received on these indicators of importance. For this purpose, the *citing* and *cited* responses were combined into one data set. In order to control for variations in citation practice by field and for changes in propensity to cite and extent of truncation over time, all regressions include technology field and grant-year dummy variables. In addition, based on the findings of Jean Lanjouw and Mark Schankerman (1999), we also included the log of the number of claims made by each patent, to allow for the possibility that patents that consist of more claims are more highly cited.

The results provide some evidence that citations are correlated with significance or importance as perceived by the inventors themselves. Each of the indicators is positively correlated with log citations, with the coefficients achieving *t* statistics that vary from just below to just above 2, depending on the question. There is no particular indication as to whether citations are more associated with technological versus economic significance. The claims variable is strongly significant, though its elasticity of about 0.25 suggests strong diminishing returns to increasing the number of claims, as distinct from the constant-returns relationship suggested by Lanjouw and Schankerman (1999). If the claims variable is excluded from the regression, the effect of the perceived-importance variable increases, suggesting that importance, as perceived by the inventor, reflects both the "size" of the patents as indicated by the number of claims, and the importance or significance of each of the claims.

IV. Conclusion

Many of the important concepts in the economics of technological change are fundamentally unobservable. Therefore, we routinely rely on proxies or indicators for the concepts of interest. Often, our only test of the validity of these measures is the extent to which the proxies are correlated in the way that our theory says their underlying concepts should be. In this paper, we provide an additional kind of evidence

about the unobservable process of knowledge transfer, and the relationship of patent citations to that process.

The results suggest a "half-full cup" with respect to the validity of patent citations as indicators of knowledge spillovers. Taking the responses at face value, the likelihood of knowledge spillover conditional on the observation of a patent citation is significantly greater (in both the statistical and quantitative senses) than the unconditional likelihood. Nonetheless, a large fraction of citations, perhaps something like one-half, do not correspond to any apparent spillover. We believe that these results are consistent with the notion of citations as a noisy signal of the presence of spillovers. This implies that aggregate citation flows can be used as proxies for knowledge-spillover intensity, for example, between categories of organizations or between countries. Further work is needed, however, to refine our understanding of the mechanisms by which these flows move and the relationship of those mechanisms to the citation process.

REFERENCES

- Griliches, Zvi. "Issues in Assessing the Contribution of R&D to Productivity Growth." *Bell Journal of Economics*, Spring 1979, 10(1), pp. 92-116.
- Jaffe, Adam B.; Trajtenberg, Manuel and Fogarty, Michael S. "Report on a Survey of Inventors." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 7631, April 2000.
- Jaffe, Adam B.; Trajtenberg, Manuel and Henderson, Rebecca. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Quarterly Journal of Economics*, August 1993, 108(3), pp. 577-98.
- Lanjouw, Jean O. and Schankerman, Mark. "The Quality of Ideas: Measuring Innovation with Multiple Indicators." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 7345, September 1999.