Research Summaries

The Production of Scientific Ideas

Pierre Azoulay and Joshua Graff-Zivin*

There is considerable evidence that the advancement of science influences productivity in the private sector of the economy. Thus, policymakers typically believe that public investments in science are important for long-run economic growth. But how do new scientific ideas come about? Apocryphal stories of Archimedes' eureka moment, or Newton's otherworldly contemplation interrupted by the fall of an apple, would have us believe that luck is an essential feature. Of course, if luck is all that is necessary to produce breakthroughs, then there is little room for scholarship on the subject. If, on the other hand, scientific knowledge production depends upon individuals, institutions, and incentives, then economic research should play an important role in increasing our understanding in this area.

While the pioneering work of Zvi Griliches, the founding director of the NBER's Productivity Program, set the stage for hundreds of subsequent empirical studies examining the diffusion of various technologies, comparatively little work has focused on the creation of the original technologies in the first place. This dearth of applied research on idea creation has not been for lack of interest but rather principally because of data limitations. As recently as 15 years ago, very little data were available to systematically study the scientific enterprise. Today, economists have at their disposal vast quantities of new data that allow them to link mentors and trainees, collaborators, and intellectual peers to characterize the production team. The data on papers, patents, and citations enable one to trace out the impact of individual bits of knowledge as they are incorporated into the research activities of other research teams, as well as within private sector firms. Together with methodological advances in the analysis of quasi-experimental data, we have begun to credibly characterize this production process, the conditions under which scientists collaborate to create new knowledge, and the benefits that follow.

One important theme that has emerged from the recent literature is the notion that an increased burden of knowledge because of an ever-expanding scientific frontier has led to greater scientific specialization, longer training periods, and to an increased propensity to collaborate.[1] This realization has cast a pall over the potential for ideas-based growth, because it implies that innovation is becoming more difficult over time.[2] Ultimately, whether this pessimism is warranted is an empirical question, which has led us to explore in more detail the impacts of interactions among scientists for the pace of scientific advance, and whether these interactions occur because of geographic proximity, shared intellectual interests, or social connections.

The Impact of Superstar Scientists

While the most important scientific work is much more likely to be produced as part of a collaboration than was the case only 40 years ago,[3] our own work suggests that the central members of these teams — whom we call "superstars" — continue to play an important role in shaping the rate and direction of scientific advance. Over the past ten years, we have gathered biographical information for a sample of 12,000 elite, academic life scientists, and combined these with precise measures of inputs (namely grants from the National Institutes of Health), outputs (publications and patents), and impact (citations to both publications and patents). Furthermore, we have linked these superstars to a much larger population of 200,000 academic life scientists in the United States, corresponding to most of the profession from the immediate post-war era to the end of the previous decade. Thanks to open-source software tools we designed for this purpose, we are able to locate all of these scientists in geographic space, identify their ties through co-authorship and citation networks, and assess the extent to which they work on similar topics.

Our first study in this area focuses on the benefits of exposure to superstar talent derived from formal collaboration.[4] The formation of collaborative teams is the result of a purposeful matching process, making it difficult to uncover the causal effect of collaboration on follow-on individual performance. To overcome the endogeneity of the collaboration decision, we use the quasi-experimental variation in the structure of co-authorship networks induced by the premature and sudden death of active superstar scientists. Our sample comprises 122 of these unfortunate events, and provides a unique opportunity to estimate the impact of the prominent members of scientific teams on their less-heralded collaborators when they work on other projects, as well as to

** Azoulay and Graff Zivin are NBER Research Associates. Their Profiles appear later in this issue.*

probe the mechanisms that undergird this influence. Our results show that upon losing a superstar collaborator in this way, scientists experience a long-lasting productivity decline, with a loss of 5 to 10 percent of their previous quality-adjusted publication output. Given the rich data we have gathered, we are also able to examine several competing explanations for this effect.

One view of the academic reward system provides the backdrop for a broad class of stories with a common thread: that collaborating with superstars deepens social connections, possibly making researchers more productive in ways that have little to do with scientific knowledge, such as connecting coauthors to funding resources, editorial goodwill, or potential coauthors. Yet, we find no differential impact on coauthors of stars well-connected to the NIH funding apparatus, on coauthors of stars more central in the collaboration network, or on former trainees. These findings do not jibe with explanations stressing the gate keeping role of eminent scientists.

Rather, the effects of superstar extinction appear to be driven by the loss of an irreplaceable source of ideas. We find that coauthors close to the star in intellectual space experience a sharper decline in output than coauthors who work on less related topics. Furthermore, the collaborators of stars whose work was heavily cited at the time of their death also undergo steeper decreases than collaborators of superstars who were less well known. Together, these results paint a picture of an invisible college of coauthors bound together by interests in a fairly specific scientific area, which suffers a permanent and reverberating intellectual loss when it loses its star.

This first paper focused on the effects of exposure to superstar talent through collaboration, but our second effort highlights geographic co-location as the channel of influence. We use as a source of variation the job transitions between distant institutions in our sample of elite scientists.[5] There again, the challenges involved in establishing causality loom large, since scientists might choose to switch jobs at least in part based on the prospects of deeper interactions with colleagues or firms. We use a novel identification strategy that exploits labor mobility in a sample of 9,483 elite academic life scientists to examine impacts on the citation trajectories associated with individual articles (resp. patents) published (resp. granted) *before* the scientist moved to a new institution. This longitudinal contrast purges our estimates of most sources of omitted variable bias that can plague cross-sectional comparisons. However, the timing of mobility itself could be endogenous. To address this concern, we pair each moving scientist/article dyad (resp. scientist/patent dyad) with a carefully chosen control article or patent associated with a scientist who does not move to a new position. In addition to providing a very close match based on time-invariant characteristics, these controls share very similar citation trends prior to the mobility event. By analyzing the data at the matched-pair level of analysis, this simple difference-in-difference framework provides a flexible and non-parametric methodology for evaluating the effects of labor mobility on knowledge flows. Indeed, conditional on the assumption that the matching algorithm we employ successfully pairs articles and patents of comparable quality, we are able to present the findings in a straightforward, graphical form.

The results reveal a multifaceted story. We find that article-to-article citations from the scientists' origin location are barely affected by their departure. In contrast, article-to-patent citations, and especially patent-to-patent citations, decline at the origin location following a superstar's departure, suggesting that spillovers from academia to industry are not completely disembodied. We also find that article-to-article citations from scientists' destination locations markedly increase after they move. To the extent that academic scientists do not internalize the effect of their location decisions on the circulation of ideas, our results raise the intriguing possibility that barriers to labor mobility in academic science limit the recombination of individual bits of knowledge, resulting in a diminished rate of scientific exploration.

We are currently extending our research in this area along two dimensions. The "superstar extinction" study left open the question of whether we would observe the same negative impact on non-coauthors working in the same field. Our next study aims to ascertain whether coauthorship is required in order to be a full-fledged member of the invisible college of scientists formed around the star while alive. This question is important because it provides a lens through which we can examine whether scientific ideas are accessible to any trained scientist — exemplars of the pure public goods described in our economic textbooks — or whether they should be thought of as at least partially excludable — as would be the case if career success all but required direct connections with the scientific elite.

Second, we revisit our earlier work on the effect of elite scientist mobility, but with a focus on the creation of new knowledge as opposed to the diffusion of preexisting knowledge. Using a novel software tool, we can identify the peers of individual scientists based solely on their shared intellectual interests as indicated by keywords that tag publications — without any reference to linkages through coauthorship or citation. This opens the door to a deeper understanding of the process through which scientists position themselves in "intellectual space," and to the evolution of scientific fields over time.

Incentives for Scientific Exploration

A distinct but related part of our innovation research agenda is how scientists choose projects, and the extent to which funding systems shape these choices. In collaboration with Gustavo Manso from the University of California at Berkeley,[6] we document that the features of incentive schemes embodied in the design of research contracts exert a profound influence on the subsequent development of breakthrough ideas.

In particular, we study the careers of researchers who can be funded through two very distinct mechanisms: investigator-initiated R01 grants from the NIH, or support from the Howard Hughes Medical Institute (HHMI). HHMI, a non-profit medical research organization, plays a powerful role in advancing biomedical research and science education in the United States. It has also adopted practices that should provide strong incentives for breakthrough scientific discoveries: the award cycles are long (five years, and typically renewed at least once); the review process provides detailed, high-quality feedback to the researcher; and the program selects "people, not projects," which allows for the quick reallocation of resources to new approaches when the initial ones are not fruitful. This stands in sharp contrast with the incentives offered to life scientists funded by the NIH. The typical R01 grant cycle lasts only three-to-five years, and renewal is not very forgiving of failure. Feedback on performance is limited in its depth and projects are funded based on clearly defined deliverables.

The contrast between the HHMI and NIH grant mechanisms naturally leads to the question of which incentives result in a higher rate of production of particularly valuable ideas. In the absence of a plausible source of exogenous variation for HHMI appointment, we estimate the treatment effect of the program by contrasting HHMI-funded scientists' output with that of a carefully matched group of NIH-funded scientists who focus their research on the same subfields of the life sciences as HHMI investigators, and who received prestigious early career prizes.

Our results support the hypothesis that appropriately designed incentives stimulate exploration. In particular, we find that the effect of the HHMI program increases as we examine higher quantiles of the vintage-adjusted distribution of citations. Our preferred econometric estimates imply that the program increases overall publication output by 39 percent relative to early career prize winners; the magnitude jumps to 96 percent when we hone in on the number of publications in the top percentile of the citation distribution. Symmetrically, we also uncover robust evidence that HHMI-supported scientists "flop" more often than the control group: they publish 35 percent more articles that fail to clear the (vintage-adjusted) citation bar of their least well cited pre-appointment work. This suggests that the HHMI investigators are not simply stars on the rise who are anointed by the program, but rather they appear to place more risky scientific bets after their appointment.

These findings are important for at least two reasons. First, they demonstrate the impact of nuanced features of research contracts for the rate and direction of scientific progress. Given the prominent role that scientific change is presumed to play in the process of economic growth, this has important implications for the organization of public and private research institutions. Second, they offer empirical support for the theoretical model developed by Manso,[7] and as such may provide insights relevant to a wider set of industries that rely on creative professionals, ranging from advertising and computer programming to leadership roles at the upper echelons of the corporate world. Many questions remain, and will constitute part of our research agenda going forward.

---

1 *B. Jones, "The Burden of Knowledge and the Death of the Renaissance Man: Is Innovation Getting Harder?" NBER Working Paper No. 11360, May 2005, and* Review of Economic Studies, *76 (2009), pp. 283–317.*

2 *P. Azoulay, J. Graff Zivin, and G. Manso, "NIH Peer Review: Challenges and Avenues for Reform," NBER Working Paper No. 18116, June 2012, and* Innovation Policy and the Economy, *Volume 13, J. Lerner and S. Stern, eds., forthcoming from the University of Chicago Press.*

3 *S. Wuchty, B. Jones, and B. Uzzi, "The Increasing Dominance of Teams in the Production of Knowledge,"* Science, *316 (May 2007), pp.1036–9.*

4 *P Azoulay, J. Graff Zivin, and J. Wang "Superstar Extinction," NBER Working Paper No. 14577, December 2008, and* Quarterly Journal of Economics, *25 (2010), pp. 549–89.*

5 *P. Azoulay, J. Graff Zivin, and B. Sampat, "The Diffusion of Scientific Knowledge across Time and Space: Evidence from Professional Transitions for the Superstars of Medicine," NBER Working Paper No. 16683, January 2011, and* The Rate and Direction of Inventive Activity: A New Agenda, *J. Lerner and S. Stern, eds., University of Chicago Press, April 2012.*

6 *P. Azoulay, J. Graff Zivin, and G. Manso, "Incentives and Creativity: Evidence from the Academic Life Sciences" NBER Working Paper No. 15466, October 2009, and* The RAND Journal of Economics, *42 (2011), pp. 527–54.*

7 *G. Manso, "Motivating Innovation",* Journal of Finance, *66 (2011), pp. 1823–60.*

NBER Profile: *Pierre Azoulay*

Pierre Azoulay is a Research Associate in the NBER's Program on Productivity, Innovation, and Entrepreneurship. He is also an Associate Professor at MIT's Sloan School of Management.

Azoulay received his Ph.D. in Management from MIT in 2001. He taught at Columbia University's Graduate School of Business from 2001–6 before joining the MIT Sloan School faculty in July 2006. His research focuses on the economics of science, technological innovation, and entrepreneurship.

Azoulay grew up in Paris, and now lives in Newton, MA with his wife Andrea, an architect, and two daughters, Sivan (6) and Orli (4). The young age of his children may explain why he barely remembers an earlier period of his life when he enjoyed reading books unrelated to work.

NBER Profile: *Joshua Graff Zivin*

Joshua Graff Zivin is a Research Associate in the NBER's Programs on Environmental and Energy Economics, Health Economics, and Health Care. He is also a Professor in the School of International Relations and Pacific Studies and in the Department of Economics at the University of California, San Diego (UCSD).

Graff Zivin holds an undergraduate degree in economics and psychology from Rutgers University and a Ph.D. in Environmental and Resource Economics from the University of California, Berkeley. Prior to joining the UCSD faculty, he taught at Columbia University from 1998–2008. In the 2004–5 academic year, he served as Senior Economist for Health and the Environment on the Council of Economic Advisers.

Graff Zivin currently serves on the Board of Editors of the *Journal of Economic Literature*. His research interests are broad and include the environment, health, economic development, and innovation.

Graff Zivin lives in Solana Beach, California. He enjoys running, cooking, and spending time with his wife and three sons.