

Gans, Joshua; Murray, Fiona

Article

Scientific teamwork

NBER Reporter

Provided in Cooperation with:

National Bureau of Economic Research (NBER), Cambridge, Mass.

Suggested Citation: Gans, Joshua; Murray, Fiona (2016) : Scientific teamwork, NBER Reporter, National Bureau of Economic Research (NBER), Cambridge, MA, Iss. 1, pp. 18-20

This Version is available at:

<https://hdl.handle.net/10419/178723>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Joshua Gans is a professor of strategic management and holder of the Jeffrey S. Skoll Chair of Technical Innovation and Entrepreneurship at the Rotman School of Management, University of Toronto, with a cross appointment in the Department of Economics. Since 2013, he has been area coordinator of strategic management; he also is chief economist of the University of Toronto's [Creative Destruction Lab](#).

Prior to 2011, Gans was the foundation professor of management (information economics) at the Melbourne Business School, University of Melbourne; prior to that he was at the School of Economics, University of New South Wales. In 2011, he was a visiting researcher at Microsoft Research in New England.

Gans holds a Ph.D. from Stanford University and an honors degree in economics from the University of Queensland. In 2012, he was named a research associate of the NBER, where he is involved in the Productivity, Innovation, and Entrepreneurship Program.

At Rotman, Gans teaches M.B.A. and commerce students network and digital market strategy. He is a coauthor with Stephen King and Robin Stonecash of the Australasian edition of Greg Mankiw's *Principles of Economics*, and lead author of *Core Economics for Managers*, *Finishing the Job*, and *Parentonomics*. Most recently, he has written *Information Wants to be Shared* and *The Disruption Dilemma*.

Gans specializes in the nature of technological competition and innovation, economic growth, publishing economics, industrial organization, and regulatory economics. He has published on these subjects in the *American Economic Review*, *Journal of Political Economy*, *RAND Journal of Economics*, and others.

Scientific Teamwork

Joshua Gans and Fiona Murray

Team performance in many settings has long challenged economic thinking. Even when monetary incentives are present, it is hard to structure those incentives to overcome moral hazard and other issues of free riding, especially when team tasks interact with one another. This is especially true for scientific teams, where the challenges are multiplied: The rewards tend to be non-monetary and thus principals — to the extent they even exist — face additional complexity in structuring those rewards. To add to the challenges, in recent decades science has become more complex and the knowledge frontier is now harder to expand than ever. This manifests itself in many changes, among the most important being a change in the life cycle of scientific careers and an increase in the prevalence and size of research teams.

Along with our coauthor Michael Bikard, we have looked at the choices scientific teams make, both in terms of how they form and in how they signal to the outside world the contributions of individual team members.

Who Gets What?

When entrepreneurs found startups, they agree on a division of equity between themselves and investors. Regardless of the ultimate value of the venture, the division of shares determines what each party owns. When teams form a scientific collaboration, one could imagine the same thing occurring. Two collaborators put their names on a paper and then whatever the paper's scientific value, credit would be divided equally between them.

However, while equity allows for a definitive and legally binding split of future profits, things are not so simple with scientific output. For starters, the total value created by a publication is not necessarily fixed and independent of the number of authors (say, in terms of citations and impact). The total value to the career prospects of authors

from a two-author publication may be more than twice what they would receive had they produced two single-author publications, even of the same quality. Likewise, the value of the publication may be much greater for a team of younger scientists than for an older, more-established group of collaborators. In other words, there is nothing to stop “the market” — a shorthand for the complex process that determines the incremental effect of a new paper on the professional standing of its authors — from assigning shares of the publication's value that sum to more than one for the output of scientific teams.

The composition of teams also matters in the market for scientific attribution, which may look at who is part of the team and be influenced in assigning credit by their prior reputations and skills. Thus while attribution may split evenly among authors, it may also be unevenly distributed by outside observers. The great sociologist of science Robert Merton noted that often a Matthew effect arose in that those scientists who had the better reputation upon entering a collaboration would seem to receive a disproportionate share of the benefits from collaborative output.¹

These issues of attribution introduce a number of complexities. For instance, it is difficult to envisage an economic equilibrium in which a scientist actually contributes less to a project and yet is persistently rewarded more because the market misjudges his contribution on the basis of prior

reputation. The equilibrium should eliminate the misjudging. If disproportionate rewards persist, it is possible that there is an efficiency explanation for this outcome.

How To Organize?

It was with such challenges in mind that we examined choice of potential collaborators regarding team production. Do two collaborators team up or go their own way?² Our first approach imagined an asymmetry between collaborators akin to that which arises in lab settings in the natural sciences. A project was, initially, controlled by a pioneer scientist who could improve the project by eliciting the contribution of a junior scientist (or postdoc or graduate student). Like any good outsourcing arrangement, the pioneer would happily pay for value. Thus, if the junior scientist contributes enough to outweigh any lost share in value accruing to the pioneer, then the pioneer would enter into the arrangement.

Of course, the collaboration could also take another form. The pioneer might publish interim results while the junior scientist might publish separately his or her own follow-on results. The entire corpus would add the same increment to the knowledge frontier as an integrated collaboration. The difference lies in how the contributions of each party would be valued in the market for scientific attribution.

The most significant thing we found in our analysis of the organizational choices made by scientists was that if “the market” designated who gets what share in a co-authored work, it would favor an attribution rule that did not sum to more than one. Why? Because any other attribution rule would lead to scientists choosing to co-author rather than publish separately when it was otherwise less efficient to do so. In other words, when a full range of organizational

choices is considered, the market for attribution may not freely reward all contributors, but rather must allocate attribution sparingly so as not to overly distort the decision to collaborate rather than to work separately on a scientific project.

Are Teams Optimal?

If economists had the luxury of designing attribution shares, they might ask what type of attribution shares would be optimal. In reality, there is no central designer and who gets what is resolved by norms — and evolving norms at that. So what norms have evolved and how might we measure them?

That was the question we explored with Bikard.³ We analyzed a unique dataset of the annual research activity of 661 MIT faculty scientists over three decades and examined their choices of whether to collaborate or not. The idea was that by observing their publication outcomes, we could infer, in any given year, a particular scientist's portfolio of collaboration choices. If, in turn, we assumed that the scientist was maximizing the total volume of attributed citations less the costs, if any, associated with collaborating, we might be able to understand whether their choices were optimal.

The figure at left illustrates our findings. It shows that if scientists were (i) maximizing the total attributed number of citations their output generated per year and (ii) attributed a share $1/n$ of the credit for papers with n authors, then any collaboration with more than three authors would be, on average, sub-optimal for them. This suggests that the scientists were facing large costs in terms of time wasted and drawn from other projects when they were part of large teams.

Our data show that scientists made continual “mistakes” in engaging in large team collaborations. We therefore had to ask if their revealed preference in this regard might suggest a different attribution rule than the simple $1/n$ rule. Using this insight, we fit our data to a number of alternatives of the form $(1/n)^b$. We found that the best fit for b that would explain the behavior as optimal was $b = 1/2$. In other words, scientists in our MIT sample appeared to behave as if the attribution rule allocated $1/\sqrt{n}$ share of the total value of a publication to each coauthor. Importantly, with this rule, the sum of the attribution shares would exceed one. This suggests



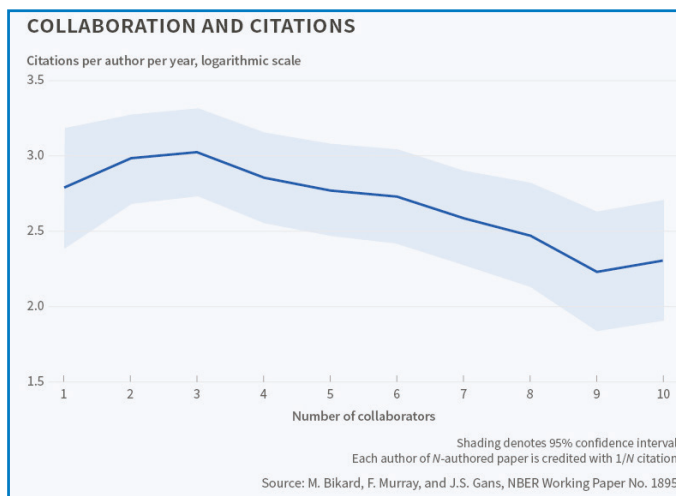
Fiona Murray is the William Porter (1967) Professor of Entrepreneurship, the associate dean for innovation at MIT's Sloan School of Management, and co-director of the MIT Innovation Initiative. She is also a member of the British Prime Minister's Council on Science and Technology and an NBER research associate.

Murray holds a B.A. and M.A. in chemistry from Merton College, University of Oxford, and an M.S. in engineering sciences and a Ph.D. in applied sciences from Harvard University.

Her research is focused on the economics and sociology of innovation and entrepreneurship. Much of her work evaluates how policies and programs shape the commercialization of science, and how they can positively impact the role of women in entrepreneurship. She is also interested in the organizational economics of science and in how changes in science funding shape the ways in which laboratories and inter-lab collaborations are structured.

Murray has done extensive work with entrepreneurs, governments, large corporations, and philanthropists designing and evaluating policies and programs that shape vibrant innovation ecosystems, such as prizes, competitions, accelerators, patent licensing rules, and proof of concept funding.

She has published in diverse journals: *Science*, *Nature*, *Proceedings of the National Academy of Sciences*, *Management Science*, *New England Journal of Medicine*, *American Journal of Sociology*, *Organization Science*, and *Journal of Economic Behavior & Organization*.



that the prevailing norms were encouraging collaboration disproportionately to individual publication.

What Drives Attribution?

We know from our own experience in evaluating our peers that the process of dividing credit for joint work is not formulaic. In particular, when we are presented with the work of a team, we try to parse the contributions of individual members.

In another collaborative paper, we explore this process by considering again a pioneer and a follow-on scientist.⁴ Both can contribute to a project. However, it is the pioneer who determines the prevailing sharing arrangements. When both actually contribute, this increases the likelihood that the project is of high quality. Indeed, we assume that to get very high quality you need both scientists to make a substantive contribution. In this event, the market knows what is going on and so divides attribution between the authors.

Things get tricky if the project is good but not of the highest quality. In that situation, by looking at the output alone, the market for scientific attribution cannot work out the underlying process. The pioneer alone surely could have generated that work. If the pioneer had been a sole author, the market would have given him all of the attribution. But what if there are two names on the paper?

If one scientist has contributed considerably more than the other, “the market” would like to find out who contributed more and attribute more credit to that author. Interestingly, this gives rise to two potential equilibria. In each one, all credit is given to one author or the other. In one of these, the follower scientist only puts in effort if the pioneer has already achieved a promising result, as the follower will share in the reward by also making a significant contribution. However, if the pioneer has

not achieved such a result, the follower puts in no effort and guarantees a low quality result precisely because the market would not attribute any share to either of them. Of course, that assessment is self-fulfilling precisely because the follower does not deserve any credit. A mirror equilibrium holds where the follower receives all of the credit. In each case, the market assessments turn out to be correct because they shape the incentives of scientists to conform to those assessments.

Our principal purpose in this paper is not to consider whether to invite another researcher to become a coauthor but, rather, when to do so. One degree of flexibility pioneer scientists have — if they lead their own labs with some autonomy — is that they can employ junior scientists but can potentially separate that working relationship from the credit or formal attribution that junior scientists receive. Senior scientists might wait until they see their own contribution and that of the junior scientist before inviting the junior scientist to be a coauthor. The senior scientist may never choose to do this, but suppose, perhaps to send a signal to others in their lab, that they commit to putting a junior scientist on the paper only if the junior’s contribution is significant.

While this arrangement might seem precarious for the junior scientist, it facilitates attribution in “the market.” If the market for scientific attribution understands that the junior scientist is only a coauthor on the paper if the junior made a significant contribution, then in the ambiguous range where it would otherwise be hard to tell who was the main contributor, “the market” can now tell. What is more, this all adds up to maximal incentives for the junior scientist to put effort into generating a significant contribution. The junior scientists are better off for this arrangement. We show that, of all of the organizational arrangements that could have been chosen, leaving the decision of whether to credit

the junior scientist until the end is Pareto optimal.

Conclusion

The research presented here is an initial foray into understanding how the choices of scientific teams are shaped by market assessments of individual performance. It is part of a broader agenda that we think of as the organizational economics of science. By demonstrating that such market assessments are likely to be important, it presents initial insights but also conjectures about what “the market” is. That remains an open theoretical and empirical question. Our work yields some insights but in many respects only highlights the reality that understanding scientific work — in academia and in industry — will require much more research, both theoretical and empirical.

¹ R. Merton, “The Matthew Effect in Science,” *Science*, 159(3810), 1968, pp. 56–63.

[Return to Text](#)

² J. S. Gans and F. Murray, “Credit History: The Changing Nature of Scientific Credit,” in A. Jaffe and B. Jones eds., *The Changing Frontier: Rethinking Science and Innovation Policy*, Chicago, University of Chicago, 2014, pp. 107–131.

[Return to Text](#)

³ M. Bikard, F. Murray, and J. S. Gans, “Exploring Tradeoffs in the Organization of Scientific Work: Collaboration and Scientific Rewards,” *NBER Working Paper No. 18958*, April 2013, and *Management Science*, 61(7), 2015, pp. 1473–1495.

[Return to Text](#)

⁴ J. S. Gans and F. Murray, “Markets for Scientific Attribution,” *NBER Working Paper No. 20677*, November 2014.

[Return to Text](#)

Productivity and Misallocation

Chang-Tai Hsieh and Pete Klenow

The starting point of a large body of recent research on economic growth is the notion that differences in aggregate total factor productivity (TFP) may not be driven solely by technology but rather in part by allocative efficiency. The key building block of this literature is the idea that firms differ, and we do not necessarily want all the resources to be allocated to one firm. For example, suppose that there are a number of firms in a country and the output Y_i of each firm is given by a standard production function, $Y_i = A_i F(K_i, L_i)$, where K_i is the firm’s capital stock (equipment and structures), L_i is the firm’s labor input (skill-weighted hours worked by its employees), F is the production function which combines capital and labor, and A_i is residual firm productivity.

If each firm produces different products, we do not want all the inputs allocated to the firm with the highest A_i , as we value having access to a variety of differentiated products. Instead, what we want is for resources to be allocated across firms to equalize the revenue productivity of the firm, or $P_i A_i$. Resources are misallocated when revenue productivity differs between firms. Reallocation increases aggregate TFP and generates growth when resources flow to firms with high revenue productivity.

Micro-data from manufacturing censuses suggest substantial gaps in revenue productivity across firms within India and China.¹ The gaps are also present in U.S. data, but are much smaller. Figure 1 plots the dispersion of revenue productivity in the three countries. In India and China, revenue productivity of firms in the 90th percentile

exceeds that of firms in the 10th percentile by a factor of five. In the U.S., the equivalent gap in revenue productivity is a factor of three. These gaps in revenue productivity between firms may contribute to substantial gaps in aggregate TFP. In a standard model, aggregate TFP would increase by 43 percent in the U.S. in 1997, by 115 percent in China in 1998, and by 127 percent in India in 1994 if resources were to be reallocated to equalize revenue productivity across firms.

We now have a large body of evidence on gaps in revenue productivity at the microeconomic level, largely thanks to the detailed, firm-level data available for a growing number of countries. A project spearheaded by Santiago Levy at the Inter-American Development Bank provides detailed evidence on these gaps for a large number of countries in Latin America.² There is similar evidence from microeconomic data for a number of countries in Europe.³ These studies find wide gaps in revenue productivity, consistent with substantial misallocation.

The literature has largely focused on measuring the static effects of firm-level gaps in revenue productivity, but the firm-level gaps are likely to also have important dynamic effects. If more-efficient establishments face larger distortions, it undermines firms’ incentives to invest in better technology. Put differently, there are two effects of resource misallocation — the static effect and the dynamic effect of resource misallocation on growth in firm productivity. This has been highlighted in several case studies.⁴ Evidence from firm-level censuses in India and Mexico is also consistent with the presence of



Chang-Tai Hsieh is the Phyllis and Irwin Winkelried Professor of Economics at the University of Chicago’s Booth School of Business. He is the co-organizer, with Nick Bloom and Chris Woodruff, of NBER’s Working Group in Development, Productivity, and Innovation, and a member of the Steering Group of the International Growth Centre at the London School of Economics.

Before joining the faculty at Chicago, he was an associate professor and professor at the University of California, Berkeley, and an assistant professor at Princeton University. He recently served as the coeditor of the *Journal of Economic Perspectives*. He is a fellow of Academia Sinica, a recipient of the Sun Yefang Award from the Chinese Academy of Social Sciences, and a Sloan Foundation Research Fellow. His research focuses on development and economic growth.

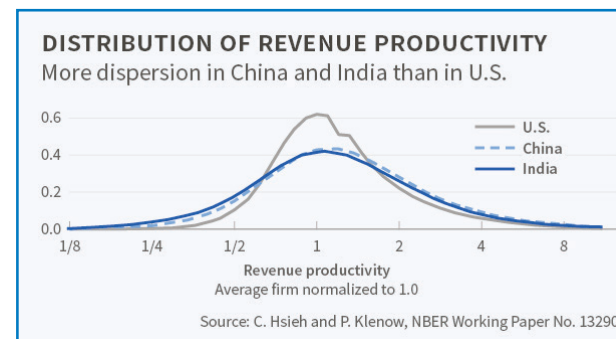


Figure 1