

Experiments in Open Innovation at Harvard Medical School

What happens when an elite academic institution starts to rethink how research gets done?

BY EVA GUINAN, KEVIN J. BOUDREAU AND KARIM R. LAKHANI

HARVARD MEDICAL SCHOOL SEEMS an unlikely organization to open up its innovation process. By most measures, the more than 20,000 faculty, research staff and graduate students affiliated with HMS are already world class and at the top of the medical research game, with 12 Nobel laureates in medicine or physiology and approximately \$1.4 billion in annual funding from the U.S. National Institutes of Health (NIH).

But in February 2010, Drew Faust, president of Harvard University, sent an email invitation to all faculty, staff and students at the university (more than 40,000 individuals) encouraging them to participate in an “ideas challenge” to generate research topics in Type 1 diabetes. Eventually, the challenge was shared with more than 250,000 invitees, resulting in 163 research ideas and hypotheses. These were narrowed down to 12 winners, and multidisciplinary research teams were formed to submit proposals on them. The goal of opening up idea generation and disaggregating the different stages of the research process was to expand the number and range of people who might participate. Today, seven teams of multidisciplinary researchers are working on the resulting potential breakthrough ideas.

In this article, we describe how leaders of Harvard Catalyst — an organization whose mission is to drive therapies from the lab to patients’ bedsides faster and to do so by working across the many silos of

Harvard Catalyst sought to examine how every stage of the Type 1 diabetes research process could benefit from openness — from generating ideas, to evaluating and selecting them, to encouraging experts from diverse related fields to consider joining research teams.



THE LEADING QUESTION

How do open-innovation concepts apply to a world-class R&D organization?

FINDINGS

- ▶ The open-innovation process brought forth proposals that might have been lost in a standard grant review.
- ▶ The initiative revealed demand among researchers for cross-disciplinary work.
- ▶ Add open aspects to the current innovation process rather than supplanting it.



Harvard Medical School — chose to implement principles of open and distributed innovation. As architects and designers of the experiments, we share firsthand knowledge about what it takes for a large elite research organization to “innovate the innovation process.”

Harvard Catalyst’s Experiment

Harvard Catalyst, the pan-university clinical translational science center situated at Harvard Medical School, wanted to see if “open innovation” — now gaining adoption within private and government sectors — could be applied within a traditional academic science community. Many insiders were highly skeptical. The experiment risked alienating top researchers in the field, who presumably know the most important questions to address. And there was no guarantee that an open call for ideas would generate breakthrough research questions.

But leaders at Harvard Catalyst suspected open innovation might be an effective vehicle. Harvard Catalyst also has a mission to share infrastructure, expediting resources to collaborators across Harvard Medical School. That is no small task, for Harvard Medical School is a huge and complex institution, with 17 separate health centers operating independently, where research projects are typically housed within one discrete center.

Harvard Catalyst’s experiment in open innovation not only was about finding novel solutions to difficult problems but also was intended to explore how all aspects of the scientific research process could be systematically opened up: from formulating research questions to evaluating research proposals and ultimately to encouraging scientific experiments that brought in new ideas and approaches to long-standing problems. The trick was to do this while integrating with the traditional research process.

Guiding Principles of Open Innovation

Research in open innovation in recent years has clarified how a wide range of economic principles and innovation practices can be used to harness the collective energy and creativity of large numbers of contributors. This may take many forms — open-source software, open apps platforms, prize-based contests, crowdsourcing and crowdfunding, dis-

tributed contributions to shared goods and repositories, and so on. But core to open-innovation systems is the intent to grant access to large numbers of actors.¹

In the private and government sectors,² open innovation has been particularly useful in addressing nonroutine or non-paradigmatic problems that may require novel approaches.³ (The approach has also been used to find low-cost solutions to more routine problems.) Organizations as diverse as Eli Lilly, Procter & Gamble, Disney and NASA have all initiated successful programs to tap into the knowledge and expertise of individuals and teams outside their own organizational boundaries.

At Harvard Medical School, interest in open innovation began to surface as Harvard prepared its application to become one of 60 NIH-funded Clinical and Translational Science Centers. While Harvard’s primary motivation was to preserve its historical NIH funding for clinical research, Harvard Catalyst’s principal investigator, Dr. Lee Nadler, and his leadership team strongly embraced the NIH’s overarching goal of bringing multidisciplinary investigators from different institutions and diverse fields together to solve high-risk, high-impact problems in human health-related research.

In May 2008, Harvard Catalyst received a five-year NIH grant of \$117.5 million, plus \$75 million from the university, the medical school and its affiliated academic health-care centers. These funds were designated to educate and train investigators, create necessary infrastructure and provide novel funding mechanisms for relevant scientific proposals. However, the funds did not provide a way to engage the diversity and depth of the whole Harvard community to participate in accelerating and “translating” findings from the scientist’s bench to the patient’s bedside, or vice versa. Could open-innovation concepts be applied within a large and elite science-based research organization to help meet that goal?

Applying Open Innovation to Open Science

In contrast to R&D activities in the private sector, a significant hallmark of academic science is its openness.⁴ Scientists build their reputations by publishing in top-ranked peer-reviewed journals. The scrutiny of peer review by experts in the field is

designed to ensure that valid and significant scientific theories and evidence are advanced, while less convincing work is weeded out. Funding decisions are based on prior work, the credibility of the investigator in the field and the potential impact and feasibility of the proposed research. Competition is tough, with fewer than 10% of papers accepted in the top scientific journals and fewer than 10% to 20% of research proposals funded.⁵

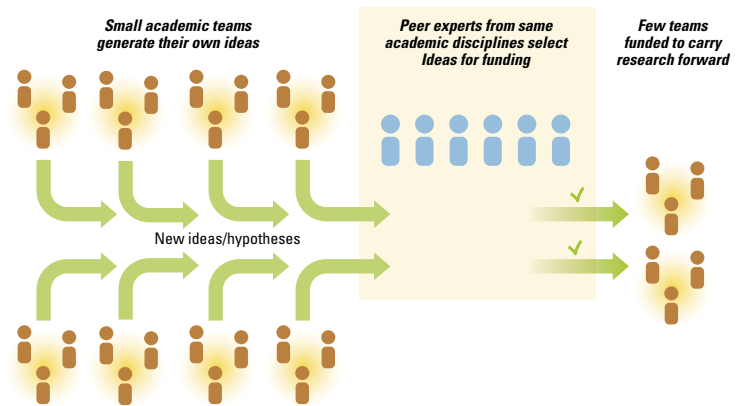
While the academic scientific process is more open than most commercial endeavors, it still follows the logic of a closed innovation system in which a few individuals determine the direction and execution of innovative efforts. A scientific team, typically led and assembled by a principal investigator, first determines hypotheses to explore. The research questions and hypotheses are based on the training, education and expertise of the principal investigator and his or her collaborators. The direction of research is also conditioned by the priorities set by the scientists' peers through specialist meetings and calls for funding proposals. Peer-review panels then evaluate whether a proposal is responsive to the funder's request and suitable for funding. The research team conducts the experimental work and, depending on the quality and significance of the results, sends out the results for another peer review in hopes of publishing the results. Hence, while the outputs of academic science are open, the *inputs* and the process are closed and primarily limited to experts within a field. This results in a fully integrated innovation system in which all steps in the research process are accomplished by the research team. (See "The Integrated Innovation Process in Academic Medicine.")

Achieving wider exposure is not simply a matter of removing barriers (legal, informational, etc.) but hinges on investments to create interest (such as marketing and messaging) in a problem. Another key issue is reducing the costs incurred when outsiders work on newly exposed problems, by architecting smaller, bite-sized problems and supporting infrastructure and platforms to facilitate interactions with large numbers of actors.

Harvard Catalyst sought to examine how every stage of the Type 1 diabetes research process could benefit from openness — from generating ideas, to evaluating and selecting ideas for funding, to

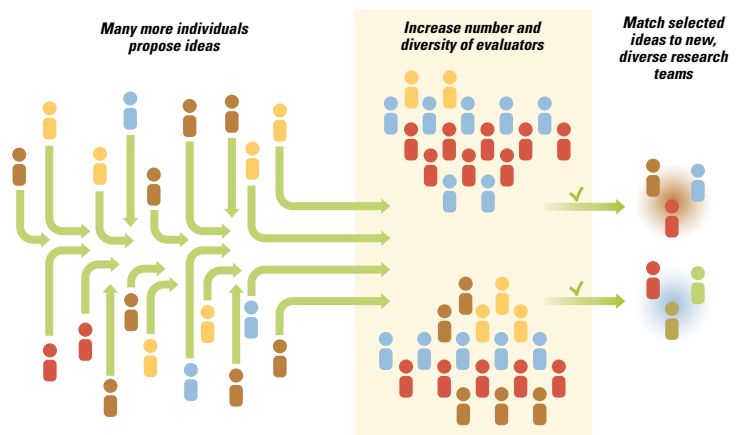
THE INTEGRATED INNOVATION PROCESS IN ACADEMIC MEDICINE

While the outputs of academic science are open, the inputs and the process are primarily limited to experts within a field.



OPENING UP THE INNOVATION PROCESS IN ACADEMIC MEDICINE

The core insight driving Harvard Catalyst's experiments was that all stages of the previously narrow and fully integrated innovation system — from hypothesis generation to idea selection to execution — can be disaggregated, separated and opened to outside input.



encouraging experts from diverse related fields to consider joining research teams.

Generating Research Questions

With support from Dr. William Chin, the executive dean for research at Harvard Medical School and a former vice president of research at Eli Lilly (an early adopter of open innovation), Harvard Catalyst started with the front end of the innovation system by opening up the process of generating research questions. Albert Einstein captured the importance of this aspect of research:

“The formulation of a problem is far more often essential than its solution, which may be merely a matter of mathematical or experimental skill. To raise new questions, new possibilities, to regard old problems from a new angle, requires creative imagination and marks real advances in science.”⁶

Thus, instead of focusing on identifying individuals who might tackle a tough research problem, Harvard Catalyst wanted to allow an open call for ideas (beginning with Faust’s email to the entire Harvard community) in the form of a prize-based contest to determine the direction of the academic research. This might lead to potentially relevant questions not currently under investigation or largely ignored by the Type 1 diabetes research community.

Harvard Catalyst partnered with the InnoCentive online contest platform to initiate the idea generation process. Using American Recovery and Reinvestment Act (ARRA) funding from NIH, Harvard Catalyst and InnoCentive broadcast invitations to participate in a challenge titled “What do we not know to cure Type 1 diabetes?” The challenge, which was open for six weeks in 2010, was advertised throughout the Harvard and InnoCentive communities, and in the journal *Nature* as well. Participants had to formulate well-defined problems and/or hypotheses to advance knowledge about Type 1 diabetes research in new and promising directions.

Participants needed neither the resources nor the detailed information to implement their ideas, and they were not expected to provide solutions. Rather, they had to define problems or areas that called for further exploration and research. Any area of Type 1 diabetes research (molecular causes, detection and diagnostics, novel treatments and optimized treatment regimens, patient maintenance and care, etc.) was fair game for a submission. Contestants could participate in teams or individually, using InnoCentive’s virtual solving rooms.

Harvard Catalyst offered \$30,000 in awards. Contestants were not required to transfer exclusive intellectual property rights to Harvard Catalyst. Rather, by making a submission, the contestant granted Harvard Catalyst a royalty-free, perpetual, non-exclusive license to the idea and the right to create research-

funding proposals to foster experimentation.

In total, 779 people opened the link at InnoCentive’s Web site, and 163 individuals submitted 195 solutions. After duplicates and incomplete submissions were weeded out, a total of 150 submissions were deemed ready for evaluation. The submissions encompassed a broad range of therapeutic areas including immunology, nutrition, stem cell/tissue engineering, biological mechanisms, prevention, and patient self-management. Submitters represented 17 countries and every continent except Antarctica. About two-thirds came from the United States. Forty-one percent of all submissions came from Harvard faculty, students or staff, and 52% of those came from Harvard Medical School. Responders’ ages ranged from 18 to 69 years, with a mean age of 41.

As hoped, the open call for ideas attracted many people with limited knowledge of or expertise in diabetes. A mere 9% of responders defined themselves as having intimate knowledge of Type 1 diabetes issues. About 47% of responders had some prior knowledge, and for about 42%, the field and the problem were new. Although only .1% of the U.S. population has Type 1 diabetes, about 46% of respondents to the call reported that they either had Type 1 diabetes or knew a family member or friend with the disease. On average, responders spent 27 hours working out a question for the challenge.

We analyzed the ideas to see how they related to the existing body of medical knowledge and to examine if the challenge had generated novel research questions. Virtually all medical and life-sciences literature in the U.S. is archived by the National Library of Medicine and goes through an independent process of knowledge category assignment by specialized librarians who assign Medical Subject Headings (MeSH) for each article. This is an essential component of enabling researchers to find accurate and timely knowledge in a database of more than 8 million articles. Currently, there are more than 25,000 MeSH terms, and each article in the National Library of Medicine has from 10 to 15 MeSH terms assigned to it. We had a professional medical librarian create knowledge categories for the content in the 150 proposals using the standardized MeSH categories and found that the proposals covered 732 knowledge categories. We then examined the overlap of MeSH categories in our 150 submissions and a

random set of 150 articles on Type 1 diabetes in the National Library of Medicine database. We did this comparison 1,000 times, each time drawing a random sample of 150 articles from the archive, and we discovered that on average the sets had 208 MeSH headings in common — out of a total of 1,452 headings. In parallel, we examined the overlap in MeSH terms in independent samples of 150 Type 1 diabetes articles from the National Library of Medicine archives in pairwise fashion, again repeating the exercise 1,000 times, and we found that the knowledge overlap was significantly higher, at 306 categories out of 1,440 categories. This implies that the focus of our proposals was quite different from what existed in the literature and from the existing body of ideas under investigation within the Type 1 diabetes research community.

Evaluating Novel Proposals to Select the Best

The process of selecting ideas is typically dominated by a relatively small set of senior experts.⁷ Studies of such grant peer reviews show that experts bring significant biases to evaluation.⁸ In addition, an expert's ability to forecast novel and significant outcomes significantly underperforms larger aggregation of opinions.⁹ In keeping with open-innovation principles, Harvard Catalyst opened up the process of reviewing ideas by inviting experts with widely disparate knowledge bases to help select noteworthy research proposals.

In open innovation, the selection of ideas is a major task because the volume of submissions typically exceeds standard practice. To evaluate 150 submissions, Harvard Catalyst substantially increased both the total number of evaluators and the number of evaluations per submission. In the interests of better understanding how someone's academic background affected his or her review, Harvard Catalyst deliberately sought evaluators with a range of expertise and experience in Type 1 diabetes.

Using institutional or departmental affiliation (for example, "Joslin Diabetes Institute" or "endocrinology vs. not") and publication record (for example, "having type 1 diabetes as a MeSH category vs. not"), Harvard Catalyst recruited six cohorts of faculty members to help in the evaluation. The six cohorts consisted of senior Harvard-affiliated

faculty in the Type 1 diabetes field; junior Harvard-affiliated faculty in the Type 1 diabetes field; senior Harvard-affiliated coauthors of Type 1 diabetes experts who were not themselves Type 1 diabetes experts; junior Harvard-affiliated coauthors of Type 1 diabetes experts who were not themselves Type 1 diabetes experts; senior Harvard-affiliated faculty in medical/biological fields removed from Type 1 diabetes as assessed by specialty and authorship/coauthorship criteria; and junior Harvard-affiliated faculty in medical/biological fields removed from Type 1 diabetes as assessed by specialty and authorship/coauthorship criteria. Overall, 142 Harvard Medical School faculty members evaluated 150 submissions for impact on the field and feasibility — resulting in 2,130 evaluations under a completely double-blind review process, in which neither the contestant nor his or her affiliation or credentials were known to the reviewers.

Analysis of the evaluation scores from the reviewers showed quite a bit of disagreement among the six cohorts as to the best proposals. Using the overall average impact score from all reviewers as the base reference case, we observed that many of the proposals that were top ranked by one cohort were ranked less highly by others. In fact, our analysis showed that only four submissions co-occurred in the top 10% for any two of the six cohorts.

The data also demonstrated that one form of traditional review — in which three reviewers are chosen by a funding agency to evaluate submissions — would have yielded extremely noisy outcomes. A random sample of any three Harvard faculty reviewers would have produced outcomes in which the overall best rank ordering of proposals (based on the average of all 142 reviewers) had rapidly decreasing probabilities of being actualized. For example, in a funding regime with 10% of submissions being funded, a random set of three reviewers would fund the top overall submission 81% of the time, rapidly declining to a 56% probability of funding for the overall fifth-ranked proposal. For proposals ranked sixth through 15th overall, the probability of being funded under a traditional three-reviewer regime varied from 44% to about 26%. The numbers get progressively worse as the funding cutoffs are lowered.

Rather than relying on just one cohort's opinion, Harvard Catalyst chose to aggregate the responses



from all reviewers and provide awards to the 12 best submissions based on the average score. This open process of collecting impressions and ratings from large numbers of well-qualified Harvard faculty members across a range of disciplines brought forth proposals that might well have been lost in a traditional grant review. Winners included a human resources professional with Type 1 diabetes, a college senior, an associate professor of biostatistics, an academic geologist with diabetes, a retired dentist with a family member with diabetes, faculty biomedical researchers and an endocrinologist, among others.

Fostering Interdisciplinary Teams

After selecting the ideas, Harvard Catalyst set out to form multidisciplinary teams. While researchers tend to stay within their domains, Harvard Catalyst wanted to learn if scientists from other life-science disciplines and disease specialties could potentially convert their research hypotheses into responsive experimental proposals in the Type 1 diabetes arena.¹⁰ Harvard Catalyst reached out to Harvard researchers from other disciplines with associated knowledge and invited them to submit a proposal to address one of the selected questions.

Drs. Arlene Sharpe and Laurence Turka of the Harvard Institute in Translational Immunology presented the process and outcomes of the prior exercise to the scientific advisory team of the Leona Helmsley Trust, a charitable organization with an interest in funding novel and unexplored approaches to diabetes. The Leona Helmsley Trust put up \$1 million in grant funding at Harvard to encourage scientists to create experiments based on these newly generated research questions.

Drs. Turka and Sharpe also led Harvard Type 1 diabetes experts in creating a unique request for proposals in five thematic areas, based on the selected research topics. The decision to combine and consolidate the submissions was an attempt to create larger groupings that might further the goal of attracting a more diverse set of researchers.

In addition to normal advertising of the grant opportunity, Harvard Catalyst reached out to researchers whose record indicated that they might be particularly well suited to submit proposals. The Profiles system takes the PubMed-listed publications for all Harvard Medical School faculties and creates a

database of expertise (keywords) based on the MeSH classification of their published papers. Dr. Griffin Weber, then the chief technology officer of Harvard Medical School and the creator of the Profiles system, assisted Harvard Catalyst in taking the coded MeSH categories for the winning proposals — now imbedded in the thematic areas — and matching them through a sophisticated algorithm to the keyword profiles of the faculty. The intention was to move beyond the established diabetes research community and discover researchers who had done work related to specific themes present in the new research hypotheses but not necessarily in diabetes.

The matching algorithm revealed the names of more than 1,000 scientists who potentially had the knowledge needed to create research proposals for these new hypotheses. As expected, some were accomplished diabetes researchers, but many others had expertise that diverged markedly from diabetes investigation. Harvard Catalyst then emailed these faculty members and announced the new funding opportunity. A random half of the researchers were also given the names of other individuals (forming a potential team of three or four) that Dr. Weber's algorithm identified as having complementary skills and knowledge. Emails to this subset not only alerted each researcher to the funding opportunity but also suggested that it might be useful for the listed individuals to work together in developing a submission. Neither step is typical in academic science research.

The outreach yielded 31 Harvard faculty-led teams vying for Helmsley Trust grants of \$150,000, with the hope that sufficient progress in creating preliminary data would spark follow-on grants. These research proposals were evaluated by a conventional panel of Harvard faculty, with expertise weighted toward Type 1 diabetes and unaffiliated with Harvard Catalyst administration. Seven grant winners were announced. Core to the mission of the openness program was that the algorithm for potentially contributory investigators had identified 23 of the 31 principal investigators making a submission and 14 of these 23 had no significant prior involvement in Type 1 diabetes research — a core element of the open-innovation experiment. Seven proposals were funded, five of which were led by principal investigators or co-principal investigators without a history of significant engagement in Type 1 diabetes research.



Open Innovation Lessons Learned

Harvard Medical School, with some of the best researchers and practitioners of medicine in the world, would seem an unlikely candidate to embrace open innovation. There is certainly no shortage of talent and motivation at Harvard to drive for breakthroughs in medicine. Yet Harvard Catalyst's experience in opening the innovation system shows that there is significant scope and advantage in adopting open-innovation approaches among even the most elite R&D organizations. The lessons from these experiments, however, are not limited in application to academic medical centers. All organizations that have a mandate to innovate, whether creating the next great cereal product for the consumer market or solving an extremely difficult big-data analytics problem, can benefit from applying a dose of open-innovation principles to their existing innovation processes.

The core insight driving Harvard Catalyst's experiments was that all stages of the previously narrow and fully integrated innovation system — from hypothesis generation to idea selection to execution — can be disaggregated, separated and opened to outside input. (See "Opening Up the Innovation Process in Academic Medicine.") By opening up participation to nontraditional actors, Harvard Catalyst achieved its objectives of bringing in truly novel perspectives, ideas and people into an established area of research.

Each stage of the now disaggregated innovation system had to be actively managed by Harvard Catalyst leadership so that the natural tendency of the organization to rely on pre-existing field-based experts did not lead to premature closure. At each stage of the process, Harvard Catalyst leaders had to first make sure that the process was open and that the output of one open stage led to another open stage.

For example, while generating research hypotheses from diverse actors may have been interesting in its own right, making sure that the most interesting and novel ideas got selected for further pursuit required the leadership to open up the evaluation process to more diverse reviewers. Then, selecting innovative ideas for further research had to be married with an outreach program that would encourage scholars who were not necessarily diabetes experts to consider applying their own expertise

to the newly formed research proposal areas. Along the way, the temptation to resort to the traditional narrow and well-understood closed processes was quite high, as being open required more work and explanation to a large set of stakeholders.

Somewhat unexpectedly, Harvard Catalyst discovered that while academic researchers tend to be very specialized and focused on extremely narrow fields of interest, explicit outreach to individuals with peripheral links to a knowledge domain can engage their intellectual passions. Harvard Catalyst uncovered a dormant demand for cross-disciplinary work, which many leaders within Harvard Catalyst doubted existed. However, as soon as bridges were built, individuals and teams started to cross over. The lesson for managers outside academic medicine is that there may be sufficient talent, knowledge and passion for high-impact breakthrough work currently inside their organizations — but trapped in functional or product silos. By creating the incentives and infrastructure that enable and encourage bridge crossing, managers can unleash this talent.

A key design consideration in launching any sort of open-innovation initiative is the degree of integration between the already existing (closed) and new (open) approaches. Managers risk alienating existing staff if they oversell open innovation. Current employees, particularly at an elite institution like Harvard, may feel threatened and skeptical about the potential of such a radically different way of organizing innovation.

The Harvard Catalyst approach to introducing open innovation was to layer it directly on top of existing research and evaluation processes. Harvard Catalyst executives simply added an open dimension to all stages of the current innovation process. Thus, individuals already in the field did not feel they were being systematically excluded. The entire effort could be viewed as a traditional grant solicitation and evaluation process with the exception that all stages were designed so that more diverse actors could participate. This strategic layering of open dimensions on traditional processes positions open innovations as a tweak to currently accepted practice instead of a radical break with the past.

Harvard Catalyst is now systematically pursuing initiatives to integrate open-innovation principles within the existing research process across any Har-

vard school that is engaged in biomedical clinical lab research that reaches patients' bedsides. Efforts include enabling Harvard researchers to access outside knowledge and talent in solving large-scale data analytics and computational biology problems through the use of innovation contests; regularly opening up the internal grant evaluation process to more diverse individuals so that the best and most novel ideas have a chance of being selected; and consciously searching and soliciting participation in research projects that are accomplished through unconventional teams.

The experiments at Harvard Catalyst have also encouraged other academic organizations to investigate their own potential for using open-innovation initiatives, including the U.S. National Institutes of Health, the Cleveland Clinic and the Juvenile Diabetes Research Foundation. While the road from research question to patient therapy is long and arduous, the open-innovation exercises have demonstrated that new ideas and new people can be systematically harnessed to develop innovations in an important disease area. The Harvard Catalyst experience has shown that the adoption of open-innovation principles is not just for technologists and entrepreneurs. Well-established and experienced innovation-driven organizations can also gain significant benefits from an open approach.

Eva Guinan is an associate professor of radiation oncology at Harvard Medical School and the Dana-Farber Cancer Institute as well as codirector of the innovation and implementation program at Harvard Catalyst.

Kevin H. Boudreau is an assistant professor of strategy at London Business School and chief economist at the Harvard-NASA Tournament Lab at the Institute for Quantitative Social Science at Harvard University.

Karim R. Lakhani is the Lumry Family Associate Professor of Business Administration at Harvard Business School and the principal investigator of the Harvard-NASA Tournament Lab. The authors have been architects and designers of Harvard Catalyst's early experiments in open innovation. Comment on this article at <http://sloanreview.mit.edu/x/54317>, or contact the authors at smrfeedback@mit.edu.

REFERENCES

1. K.J. Boudreau and K.R. Lakhani, "How to Manage Outside Innovation," MIT Sloan Management Review 50, no. 4 (summer 2009): 69-76.
2. See the following canonical references on open innovation: H.W. Chesbrough, "Open Innovation: The New Imperative for Creating and Profiting from Technology" (Boston: Harvard Business School Press, 2003) and E. von Hippel, "Democratizing Innovation" (Cambridge, Massachusetts: MIT Press, 2005).

3. In the case of science problems, see the analysis by L.B. Jeppesen and K.R. Lakhani, "Marginality and Problem-Solving Effectiveness in Broadcast Search," Organization Science 21, no. 5 (September/October 2010): 1016-1033. Software problems are analyzed in K.J. Boudreau, N. Lacetera and K.R. Lakhani, "Incentives and Problem Uncertainty in Innovation Contests: An Empirical Analysis," Management Science 57, no. 5 (May 2011): 843-863.

4. Economists studying innovation have investigated the institutional structure of academic scientific research and open science. See, for example, P. Dasgupta and P. David, "Toward a New Economics of Science," Research Policy 23, no. 5 (September 1994): 487-521; P. David, "Common Agency Contracting and the Emergence of 'Open Science' Institutions," American Economic Review 88, no. 2 (May 1998); and S. Stern, "Do Scientists Pay to Be Scientists?" Management Science 50, no. 6 (June 2004): 835-853.

5. See P. Stephan, "How Economics Shapes Science" (Cambridge, Massachusetts: Harvard University Press, 2012).

6. A. Einstein and L. Infeld, "The Evolution of Physics: From Early Concepts to Relativity and Quanta" (New York: Simon & Schuster, 1938).

7. G.M. Carter, "What We Know and Do Not Know About the NIH Peer Review System," RAND Corporation (1982); D.E. Chubin and E. J. Hackett, "Peerless Science: Peer Review and U.S. Science Policy" (Albany, New York: SUNY Press, 1990); S. Cole, J.R. Cole and G.A. Simon, "Chance and Consensus in Peer Review," Science 214, no. 4523 (November 20, 1981): 881-886; and L. Langfeldt, "The Decision-making Constraints and Processes of Grant Peer Review, and Their Effects on the Review Outcome," Social Studies of Science 31, no. 6 (December 2001): 820-841.

8. D.F. Horrobin, "The Philosophical Basis of Peer Review and the Suppression of Innovation," Journal of the American Medical Association 263, no. 10 (March 9, 1990): 1438-41; U. W. Jayasinghe, H.W. Marsh and N. Bond, "A Multilevel Cross-classified Modeling Approach to Peer Review of Grant Proposals: The Effects of Assessor and Researcher Attributes on Assessor Ratings," Journal of the Royal Statistical Society: Series A (Statistics in Society) 166, no. 3 (2003): 279-300.

9. A good example of novices beating experts is the use of prediction markets. See S. Luckner et al., "Prediction Markets: Fundamentals, Designs, and Applications" (Munich, Germany: Gabler Verlag, 2012). For information aggregation through polling versus the use of experts in political predictions, see also N. Silver, "The Signal and the Noise: Why So Many Predictions Fail — But Some Don't" (New York: Penguin Press, 2012).

10. A.M. Weinberg, "Scientific Teams and Scientific Laboratories," Daedalus 99, no. 4 (fall 1970): 1056-1075; and S. Wuchty, B.F. Jones and B. Uzzi, "The Increasing Dominance of Teams in Production of Knowledge," Science 316, no. 5827 (May 18, 2007): 1036-1038.

Reprint 54317. For ordering information, see page 6.

Copyright © Massachusetts Institute of Technology, 2013.

All rights reserved.