KAUFFMAN DISSERTATION EXECUTIVE SUMMARY

KAUFFMAN
ENTREPRENEURSHIP SCHOLARS

KAUFFMAN DISSERTATION FELLOWSHIP

Part of the Ewing Marion Kauffman Foundation's Entrepreneurship Scholars initiative, the Kauffman Dissertation Fellowship recognizes exceptional doctoral students and their universities. The annual program awards Dissertation Fellowship grants to Ph.D., D.B.A., or other doctoral students at accredited U.S. universities to support dissertations in the area of entrepreneurship.

Since its establishment in 2003, this program has helped to launch world-class scholars into the exciting and emerging field of entrepreneurship research, thus laying a foundation for future scientific advancement. The findings generated by this effort will be translated into knowledge with immediate application for policymakers, educators, service providers, and entrepreneurs as well as high-quality academic research.

Ewing Marion
KAUFFMAN
Foundation

www.kauffman.org/kdf

THREE ESSAYS ON THE COSTS OF INDUCING INNOVATION

Kyle R. Myers

A DISSERTATION

in

Health Care Management & Economics

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2017

Abstract

This dissertation explores the costs and benefits of inducing innovation. First, I test whether the NIH can direct resources to specific areas of science. I find that it can, and that science funded through these directed efforts is more productive than average. Second, I identify how willing scientists are to adjust their research in exchange for resources: the elasticity of direction. Estimates suggest the directional adjustment costs of science are large enough to warrant policy attention. Finally, in joint work with Mark Pauly, we explore the growing costs of pharmaceutical R&D and find support for the notion that demand growth can lead to R&D productivity declines.

Category

Categories: Technology & Innovation; Productivity Growth
Keywords: economics; science policy; R&D management; productivity

Takeaways

1. The government can successfully stimulate new research, despite the large amount of control afforded to academic scientists.

2. Scientists must undertake large costs when adjusting the direction of their research; costs large enough that it may be worthwhile to subsidize these adjustments.

3. As evidenced by the pharmaceutical industry, when demand for new products grows, firms will invest in more difficult R&D and productivity will likely decline – not as a problem, but in hopes of larger rewards.

Longform Summary

Can the government "direct" science?

A classic problem pointed out by economists Richard Nelson and Kenneth Arrow around 1960 was that private firms will never invest the socially optimal amount of resources in "knowledge generation", or in other words, science. This is because it is very difficult to appropriate the full value of information since it is a "public good" - it is very difficult to prevent non-purchasers from obtaining the information and, if they do, one person learning the information doesn't prevent other from also learning it. Thus, like other public goods such as defense and sanitation, there is good reason for the government to intervene. So it is not surprising that the U.S. federal government is the single largest funder of science in the world. However, to solve Nelson and Arrow's problem, the government must not simply fund "science," but it must fund the particular types of science that private firms will underinvest in. That is to say it must fund *and direct* science.

Typically, public organizations rely on peer-review panels to assess the value of potential funding opportunities. And there is growing evidence from the National Institutes of Health (NIH), the focus of the first part of this dissertation, that such panels are able to identify promising lines of research and draw on the reviewers' expertise without them succumbing (too much) to strategic incentives that might otherwise result in inefficient allocations. Thus, by

and large, scientific funding agencies have relied on scientists themselves to identify and pursue topics they think most valuable.

However, it is not clear a priori that the allocation of funds ultimately decided by scientists themselves will be optimal from society's perspective. Their priorities might include what is in society's best interest, but will also likely include what is in their own best interest. Given there is almost certainly a mismatch – and more work needs to be done to determine just how "bad" this mismatch truly is – this project examines the ability of the NIH, the single largest public funder of science, to choose where in the spectrum of science these effects occur.

While the NIH traditionally solicits any proposal for research that may "enhance health, lengthen life, and reduce illness and disability," roughly 25% of the annual budget (~$8.2 billion in 2016) is allocated to requests for applications (RFAs) that focus on specific scientific subjects. What might limit the effectiveness of these RFAs? First, since science is funded through a multitude of public and private sources, other funders may anticipate growth in particular funding streams at the NIH and reduce their own allocations. Second, given the relatively fixed supply of scientists at any given time, each whom is specialized in very particular aspects of biomedicine, it is unclear that the extra dollars set aside are enough to increase the supply of a particular type of science.

To investigate the aggregate effect of RFAs, I utilize the NIH's administrative database to construct a new dataset of roughly 900 RFAs, 8 years of grant applications from more than

80,000 scientists, along with all of the scientists' publications, regardless of funding source. This dataset facilitates two complementary analyses: (1) when scientific subjects (i.e. "neoplasms", "hormones") are targeted, do we actually see new applications, awards and eventually publications on these same subjects; and (2) how productive are these new projects compared to when scientists submit ideas of their own?

Overall, my results indicate that with an average size of $11 million and 6-7 grants, the targeted contests induce a 8% (12%) increase in the rate of science proposed in all (successful) applications. The evidence also suggests that the newly proposed science is of roughly comparable quality (per success rates) to counterfactual science, while being nearly twice as productive in terms of publication mentions per application mention. These productivity gains at the level of science appear to be almost entirely due to differences across projects - compared to the traditional "investigator-initiated" mechanism, projects awarded via RFAs generate 1.8 times as many publications per dollar awarded.

These gains combined with a lack of any apparent substitution away from non-treated subject areas suggests that, so long as the additional value of knowledge generated on the RFA-treated subjects is at least half as large as in the non-treated subjects, these RFAs are a cost-effective use of public resources committed to basic science. However, RFAs intended for small businesses do not exhibit the clear signs of inducement observed in the basic science grants. Thus, while these early-stage commercial grants have been seen to alleviate frictions unique to

emerging firms, these results highlight the difficulties that policymakers may face when redirecting scientists at this stage of development.

What is the elasticity of the direction of science?

A fundamental duty of any manager is to determine the allocation of resources across people and projects: the direction of work. But when the best path forward is unclear, managers often relinquish some control to their workers. This is especially true in the case of managing scientists, whereby managers (that don't know who's best for the job) target resources to what they think is important, and let the scientists choose what they think they're best at. Here, resources are quite often both inputs – dollars to buy equipment – and incentives – rewards "paid" to the scientists. But importantly, these incentives that must be valued enough that a scientist is willing to change what they are working on – their direction – which is likely no small task given how specialized each scientists' knowledge and skillsets are. This paper provides evidence to the cumulative magnitude of the costs a scientists must undertake when changing their direction by asking: to what extent do scientists require more resources in order to change their research a certain amount? Or, what is the elasticity of the direction of science?

Understanding this tradeoff is essential, because as scientists become more difficult to move, the likelihood that managerial intervention is worthwhile will decrease. In order to estimate the elasticity of direction, I examine the choices of academic biomedical scientists in response to a series of targeted grant contests at the National Institutes of Health (NIH), a setting both

empirically useful and policy relevant. The NIH has increasingly become an active manager of the biomedical workforce, and like many public and private organizations that utilize similar mechanisms to manage high-skilled workers, their optimal research policy will depend on how elastic their workforce is.

Administrative data allows me to observe a large number of individuals, consider a large number of new opportunities with well-defined payoffs. And importantly, a text algorithm allows me to quantify a scientists' change in direction.

Conditional on applying to the NIH for funding, I find that scientists are indifferent between $1.7 million in research grant funding and a 1 standard deviation decrease in the scientific similarity between the new project and their prior work. Importantly, I find that this differential is driven entirely by scientists "experienced" with the NIH; evidence that older, more successful scientists are less willing to adjust the trajectories of their work.

Estimates of these kind are necessary for any manager intent on redirecting their workforce through indirect means. And while such indirect means, like these targeted contests implemented by the NIH, have been proposed as useful "low-powered incentives," understanding this elasticity is necessary to ensure these incentives are not too low-powered.

Traditionally research policies, such as those designed to address the "bottleneck" in translational biomedical research have not considered that allocating funds to particular areas

is not only necessary to conduct the research but it must also incentivize the choice of scientists' to redirect their research to that area. My finding of sizeable adjustment costs suggests this needs to be taken into consideration.

What has happened to pharmaceutical R&D productivity?

Considerable criticism has been expressed about the growing research and development costs for new drugs. These critiques often cite estimates of the average out-of-pocket costs to bring a new drug to market, which have increased in real terms from $230M in the 1980s to $540M in the 1990s to $1,300M in the 2000s. More broadly, the declining pace of innovation and growing costs of R&D per output have been documented across the U.S. economy. Still, it remains unclear to what extent these trends are driven by forces such as mismanagement, regulatory burdens, or are instead the expected outcomes of rational firms making investments in more "difficult", but still highly demanded ideas.

Instead of pinpointing specifics of these declines, the goal of this part of my dissertation is to reframe the discussion to one of demand-induced investment amidst scarce ideas. We do this by connecting the classic notion that the rate of innovation is directly related to demand growth, with David Ricardo's famous point - demand and productivity will be inversely related when inputs (here, profitable new drug ideas) are rare.

This connection guides our investigation of a what we term the "R&D production function," essentially, how many new drugs are approved for every dollar the industry spends. But a classic challenge we face when evaluating production functions is that firms do not randomly fund new projects – they fund what they think will be the easiest (low costs) and have the largest returns (high value). Thus, it is not clear if year-to-year changes in R&D investments are truly causing the year-to-year changes in new drug production we observe.

To overcome this, we utilize the fact that a firms' optimal investment level depends primarily on the future size of the market – when people are willing to spend more to buy a product, firm's should in turn be willing to spend more to create that product. Based on prior work, we construct a measure of the size of the market and use it to identify only the subset of R&D investments made in response to these changes in demand. Thus, by looking only at this particular subset of investments, we can clearly understand why this particular R&D was undertaken.

Our first set of results shows a complicated picture. In contrast to the many broad critiques of this industry, we find some nuanced results – over the past thirty years we find that the "elasticity of new drugs with respect to market size" (e.g. how many more drugs does a larger market get) is very stable. And by markets, we are referring to drug classes such as anti-infectives or growth hormones. That is to say, whether it was 1980 or 2010, larger markets received the same larger amount of new drugs. However, during this same timeframe, all markets received roughly one quarter as many new drugs. This nuance is important because it

indicates that the drug industry continued to allocate its resources in the same manner and therefore the “returns to scale” across different projects was the same. That the decline in new drugs was common to all markets over time suggests that the growth in R&D costs was driven by forces related to the fixed costs of operating in this market and other industry-wide strategic changes in drug development.