

26 Government Policy and the Commercial Value of Academic Information

Robert Cook-Deegan

We all think innovation is important—whether it is biomedical research to discover things that save people’s lives, microcomputers to do the myriad tasks we now delegate to computers, or telecommunication and transportation that make the world smaller and more tightly knit. Innovation can be a good thing—but it is not without cost. This article addresses one of those costs—the cost of secrecy. In our competitive world, innovative ideas and, often, the development of new products and technologies may be cloaked in secrecy. While such secrecy might indeed protect the initiator’s bottom line, does it have a negative impact on society as a whole? In other words, does secrecy become sludge in the scientific pipeline?

Innovation is connected to parts of the economy that grow the fastest. Two factors strongly influence the pace of innovation. One is the level of resource deployment, or how many resources are going into the process of innovation. The other is the speed of information flow in that process. As a backdrop for this discussion of secrecy in science, I would like to illustrate the dramatic expansion of scale in research and development (R&D) that has occurred in the post-World War II period.

My own areas are health policy and health research policy, areas in which there are two general attitudes about secrecy:

- when lives are at stake, secrecy is wrong—because health research is about saving lives, not making money; and
- secrecy may harm innovation.

Robert Cook-Deegan is the director of the National Cancer Policy Board at the National Academy of Sciences. This article is based on remarks he delivered at the AAAS/Massachusetts Institute of Technology Colloquium on Secrecy in Science, held March 29, 1999, Cambridge, Massachusetts.

The first, instinctual, response is that there should be no secrecy, because “if there are lives at stake, of course it is a very bad thing.” This is a widely held attitude, and it will continue to be reflected in public opinion surveys. Federal policy, however, is not premised on that response. In fact, a lot of our policies create incentives to keep secrets for certain periods of time, because we believe that there is some connection between innovation and being able to keep and protect information that is commercially valuable. The policies are driven by a belief that in the long run, we are all better off if some people can protect information for limited periods. I will not further examine the “all secrecy is bad” stance, but will make the case that even if viewed from the narrower framework of the national system of innovation, and on efficiency grounds alone, secrecy can be sludge in the pipeline of innovation.

Anybody who talks about innovation these days has to disavow the pipeline model of innovation. We no longer believe in a simple pipeline model, if we ever did. The pipeline model was developed some time ago and was probably most beautifully laid out in the first few annual reports (in the early 1950s) of the National Science Foundation. The model says that the innovation process begins with basic research, then moves to applied research, then to development, then to application of that development, and then marketing of a product. The notion was that the process entails a flow from step to step—that basic science begets applied research, which then begets development and then, eventually, a product or service. In this model, federal funding for basic research passes the baton to industrial support for development.

There are kernels of truth to the model, if we step back far enough and review a long enough time span. The most apparent one is that in general, the flow is indeed from science toward the market through development. But the flow is far from uniform—there are many eddy currents and reverse flows, and the feedback system is a paragon of cybernetic complexity. Perhaps the key insight of the past few decades is that what is being mapped is not the flow of ideas or widgets, but of people and information—not all people, but those carrying a stock of tacit knowledge about how to do new things.

Another kernel of truth to this model is that in terms of investment in R&D, the public and private sectors are curiously mismatched partners who are intertwined in a delicate dance. While most private R&D is product- or service-oriented, public R&D is typically mission-oriented (e.g., health, defense, and space exploration). And yet, public science induces private investment, and private science induces public R&D.

From the end of World War II until today, life sciences research has grown enormously, most markedly through the National Institutes of Health. The federal government has become deeply involved in the complex ecosystem that supports innovation, and during the 1980s and 1990s, federal R&D funding has become increasingly skewed in the direction of the life sciences, especially biomedical research. (Figure 1 and 2)

Few people realize that the United States is unique in its emphasis on health in terms of its R&D priorities. The U.S. spends more on health research as a fraction of its gross domestic product (GDP) than any other country. Since we have the largest GDP in the world, we have a huge investment in health research compared to any other nation. The relative funding in engineering, chemistry, or physics is quite different, as other countries have a higher fraction of their GDPs devoted to these sciences than does the U.S. (Figure 3)

This is also the case with private investment. To illustrate this trend, I will use R&D spending by members of the Pharmaceutical Research and Manufacturers Association (PhRMA). Sometime in the late 1980s or early 1990s, private annual investment in R&D in the health sector exceeded the federal investment. So the rate of growth in private funding has been even greater than in government, and more private money than public money is now going into health R&D. (Figure 4)

We now face situations that are analogous to situations that confronted computing research in the 1960s and 1970s. For example, consider the policy pursued by the Defense Advanced Research Projects Agency (DARPA) regarding chip design. DARPA is a part of the Department of Defense that funds research for all of the armed services. It was created in the wake of Sputnik. Staff at DARPA identified very large scale integrated circuits—that is, the chips we now take for granted—as being important to the future of microelectronics. They identified chip-making as a critical technology. Chip manufacturing was capital intensive; therefore, private firms were doing most of the work. (An analogous situation exists in the genomics R&D environment: the private sector is generating data and technology; competition limits data sharing, standard setting, and academic collaboration; and lead firms seek to protect themselves from harm resulting from information flow.)

Each firm was trying to develop a chip that nobody else could use in order to lock customers into their line of products. As a result, many products were incompatible. What we now take for granted as stan-

Figure 1
The Lasker Legacy: NIH's Postwar Growth




Figure 2
Research Funding
 (Adjusted to 1992 Dollars)






Figure 3
Health R&D as %GDP



Sources: Michaud and Murray, Harvard School of Public Health.

Figure 4
Private R&D Exceeds Public



Sources: NSF FFRD Survey; PhRMA Annual Survey, 1997; ADAMHA 1992 (corrected)

standard operating systems were not yet developed, and if the trend had continued, they never would have been. It was not clear in that framework—where most of the R&D took place in private sector silos—whether products would be mutually compatible or who would design chips in the future (since few graduate students were being trained). DARPA decided that there was a serious choke point that endangered the future of integrated circuit design. What they pursued as a partial solution was a research program strategically placed in academic institutions. It was an attempt, among other things, to supply academic groups to keep up with cutting-edge research in chip design.

DARPA did two basic things. It funded researchers who were then linked to graduate students and formed an academic network. They also funded a fabrication facility to which graduate students or academic researchers could send a design over the Internet (the ARPANet in those days). It was a way to lower the barriers to innovation in design, allowing a large number of people from different institutions to say, “Please make this chip, send it back to me, and I will see how well it works.” No single university would have been able to support that, but with a chip fabrication facility available to all universities, it became possible for chip design work to be rather open. Another important indirect result of the DARPA initiative was a textbook (written by professors Carver Mead and Lynn Conway) that served as a rule set and guiding principles for chip design. The textbook allowed an exchange of information and informal adoption of standards that was not happening under the proprietary model.

Why did DARPA pick universities? Partly because there is a strong presumption at universities that research will result in information that is openly published. In addition, universities house cheap, highly expert labor. Graduate students are underpaid, and in a fiscal production model that is a good thing. That source of cheap labor is also leveraged into the future because the research program trains people who are going to be company managers and designers of future chips. Universities have two other effects on the innovation ecosystem. One is that if the work is openly published, it is available to all firms in a way that if the work was done at a single firm and not published, it would not be. Much industrial research remains unpublished, even in biomedicine, where publishing is respected in the biotechnology and pharmaceutical firms. The wide availability of information to diverse firms is not a universal rule, but in general, research is more available if it is performed in the academic sector. Moreover, it is usually a lot easier for academic groups

to collaborate with industrial partners of different sizes and levels of sophistication. So the research done in universities is valuable to many firms.

The grander scale of R&D and the direct relevance of R&D to practical application have led to many changes in the university. One is that universities now seek patents to a greater degree. That has led some universities to see technology licensing primarily as a profit center, or business operation, of the university. In the postwar era there was an active debate in the major universities that pitted the federal government against industrial research. For example, the Stanford Research Institute was created to produce direct industrial investment in order to keep the university from having to depend on another unreliable partner—the federal government. In many sectors there was an either/or framework for industrial research funding.

The most successful major research universities figured out that the winning strategy was to use federal funding as the base and use industrial funding to augment (but not supplant) that base. The strategy that has worked seemingly over and over in field after field is to create a cadre of highly conspicuous, well-known scientists who publish frequently and become known not only to their scientist colleagues but also to firms working in that area. These stars are the people who are going to attract industrial funding. They are also going to get patents and publish the most. The origins of biotechnology, it has been well documented, grew from a star system in molecular biology, and every university administrator has come to recognize that. Finally, it has become abundantly obvious that the enormous sums of federal R&D money create a mountain of information that can be mined by private firms for valuable applications. Building mountains of knowledge is a big part of what universities do these days.

Turning now to secrecy, in her book *Regional Advantage*, AnnaLee Saxenian compares the MIT Route 128 region with Silicon Valley and looks specifically at the early 1980s, when there was a depression in the computing industry. She found that recovery in Silicon Valley was much faster than it was along Route 128. She ascribed it to the way the firms behaved and the way individuals behaved within those firms. She observed that in Boston people tended to be loyal to a single firm. Those firms tended to be relatively insular, and workers focused on job stability. In contrast, within Silicon Valley the loyalty was to the technology more than the firm. Companies like Hewlett-Packard and Fairchild thought spin-offs were cool, not something to be fought. In Massachusetts,


when Data General split from DEC, there was a bitter battle with lingering memories. Across the country, Hewlett-Packard and Fairchild were comfortable with their employees starting new companies, and Hewlett-Packard invested heavily in the local community well beyond its immediate business interests.

Those are the principal reasons Saxenian used to explain why Silicon Valley was quicker to recover than Route 128. Another part of the equation, however, is secrecy. Saxenian did not focus on this as much, but I believe it can explain some of the difference as well. The research ethos in the Boston area was heavily influenced by work in the wartime era. Wartime research was big at Stanford, too, but it not nearly as big or dominant. The prototypical institution at MIT was the Lincoln Labs, and in Stanford it was the linear accelerator. The ethos of those two institutions was very different. A significant difference was in the degree of information sharing. Most of the work at Lincoln Labs was secret, and that secrecy was strongly enforced. The linear accelerator facility was mainly a tool for open research. Communication was much more open on the West Coast.

That said, why should we care about information sharing? Two points about biotech illustrate that the rules about secrecy and patenting intellectual property may play out differently in biotech than they have for other technologies. First, patents matter more in pharmaceuticals and biotech than they do in any other major sector of the economy. Pharmaceutical firms routinely invest ten years in developing a lead compound and bringing it to market. The reason they can do that is because they can rely on strong patent protection that will survive that whole period of investment. In contrast, telecommunications and software firms do not rely on patents to nearly the same degree because their product cycle and time horizon for investment are considerably shorter.

The other thing that is distinctive about biotech is that it grew out of and is much more highly dependent on academic research than are other sectors. The most straightforward evidence in support of this comes from Edwin Mansfield's surveys. He asked CEOs, "How important is academic research to your business?" He asked people in telecommunications, pharmaceuticals, and other sectors. People in pharmaceuticals routinely reported that about one-fourth of their products depended completely at some point on academic research and another fourth would have been delayed significantly without a base of NIH-supported academic research. In most other sectors, CEOs reported figures that were half that or less—a twofold higher dependency on academic research in pharmaceuticals compared to other sectors.

Figure 5
Drug/Medical Patents



Source: "Universities as a source of commercial technology: a detailed analysis of university patenting, 1965-1988," Rebecca Henderson, Adam Jaffe and Manuel Trajtenberg, Review of Economics and Statistics.

Econometric studies corroborate those surveys, correlating the number of patents, citations, ownership of patents, and other measures. A much higher fraction of patents based on DNA technology, for example, is owned by academic institutions than are other kinds of patents, such as chemicals or instruments. The DNA Patent Database (at the Foundation for Genetic Medicine) indicates that among the more than 1,000 DNA patents issued between 1980 and 1993, roughly 30 percent are owned by academic institutions. The academic patent ownership rate is five percent or lower for most patent classes.

Figure 5 shows the academic and industrial patent trends for academic and industrial firms in the pharmaceutical patent classes. The line on the bottom, which represents the number of patents filed and obtained by academic institutions from 1965 to 1992, indicates a very rapid rise. Two things are worth noting here. First, the rise began before the Bayh-Dole Act was passed in 1980—the ethos had already shifted. The other is that it is not just academic patenting per se, but also a rise in the citations to academic patents. In the same publication from which these data were extracted, this high level of citation of academic patents is not evident in chemistry, telecommunications, or other areas. In fact, in the

Bayh-Dole era there have been more and more patents obtained by academic institutions, but fewer and fewer citations to them. That is the main point of the paper by Henderson and colleagues from which these data were extracted. Biotech, however, is maintaining this importance of the academic research base.

I want to use genomics as a case example to illustrate disparate approaches to promoting commercial application. There appear to be at least five business models at work in genetics research that are working simultaneously in a real-time experiment about what works best. The Human Genome Project was originally conceived as a public infrastructure project—the federal government was going to create maps and databases that would serve as an interstate highway for human genetics that would be useful to everybody. It was a government enterprise, and the idea was “Build it and they will use it.” That origin carries forth into current government genome projects. The Department of Energy (DoE) and the NIH have growing programs. Those programs include large-scale sequencing centers, the big groups that get tens of millions of dollars per year to do rapid sequencing. This approach is shared by the Wellcome Trust (which, while a private nonprofit funder, behaves similarly—or better—than government funders) and several foreign collaborators. The rule is open access and rapid disclosure of data.

In Palo Alto in 1992–1993, a firm called Incyte, which began as a Genentech contract firm, turned to sequencing expressed DNA sequences from the human genome, using sequencing as a way to grab lots of genes very quickly. They evolved a business model and became a publicly traded firm in 1993 based on plans centered on creating a database of sequences and tools to analyze it. Incyte’s basic strategy was to license access to their database to individual firms, mainly large pharmaceutical companies. The business plan has evolved, but that remains an important underlying strategy.

Another company, Human Genome Sciences (HGS), was formed in 1993. It took a similar scientific approach, but with a different business strategy. HGS initially had just one big investor, SmithKline-Beecham. HGS thus had one main licensee and also sought to forge collaborations with academic groups. This effort was somewhat circumscribed because not many academic research groups were willing to work with HGS, in part because of the strictures on disclosure that were associated with it. That brings us back to the secrecy issue.

Finally, along came the giant of the pharmaceutical business, Merck, who decided to fund the same sort of work that was going on at Incyte

and Human Genome Sciences. But Merck decided to pay Washington University in St. Louis—an academic research center—to generate the sequence and put it into the public domain immediately. Private R&D money was being used to create public domain information. A similar effort has just been announced by a consortium of pharmaceutical and biotech firms to create a public domain database of information about human genetic variation.

Celera Genomics has entered the field in the last few months, and it is planning to do massive amounts of sequencing of the entire human genome and sell access to these and other genomic data on a subscription basis. Their model is more or less the Lexis-Nexis database model. Celera currently states that it will not attempt to claim rights to future discoveries, or reach-through, when subscribers use their data to discover something valuable. Nor will they impose nondisclosure agreements beyond protections for the data in their database.

The final model, which is really not a model but a hodgepodge of disparate approaches, is based on mining the research created by thousands of academic groups at universities. Most research in molecular biology is done by people who are not employed by a company. They are working in the lab doing sequencing and mapping of genes, working at institutions whose policies regarding intellectual property are all over the map—such as licensing, speed of disclosure, and willingness to restrict publication.

Why is this happening in genomics? A few features defining genomics seem likely to generalize to many other fields. First of all, it costs a lot of money. The main thing these companies can do that academic researchers cannot is spend lots of money up front on two things—computers and sequencing machines, which are very expensive and require serious management to make them work. Companies that purchase hundreds of new sequencing machines will have the capacity to generate sequencing information much faster than any academic group can hope to match. For genomics companies and firms in other rapidly moving fields, the value of the information is when it is fresh—its commercial value decays rapidly. Its scientific value may or may not drop rapidly, but its commercial value drops very rapidly with time, because that value derives in large part from exclusivity—and a big part of exclusivity is who finds the information first.

Two features are shared by chip design and genomics. One is that the research is not basic or applied—it is both. Research creates extremely interesting information that is also extremely useful. The other sim-

ilarity is that you need lots of equipment to interpret the information—to make it meaningful. The information is so plenteous and complex that it is useful only if you can use computers to help analyze it.

Information flow problems, too, are beginning to crop up in genomics and other fields. There are several reasons that individuals may want to keep information quiet. Early in the genome debate, one of the reasons that the National Research Council did a report on mapping the human genome was a concern among molecular biologists that human geneticists were sitting on data, and hoarding pedigree and related resources. That was, in fact, true. Secrecy was not because of money and it was not because of patents or commercial interests. People who spend a lot of time creating a pedigree and putting work into constructing a clinical profile of every member of that pedigree are reluctant to share that information because it is closely linked to their ability to publish and advance their careers.

Damming information flow has been and still is a serious problem in human genetics and other areas. (In fact, one of the reasons for the Human Genome Project was to break the lock that human geneticists had on the information flow.) Information logjams typically resulted from such common practices as pursuing worldwide patents, seeking a royalty stream, signing nondisclosure agreements, and establishing licensing agreements for access to materials or information.

In the example of pursuing patent rights worldwide, individuals or institutions seeking a foreign patent cannot publicly disclose data before filing the patent application. You can get a U.S. patent, because the United States has a one-year grace period after public disclosure, but you cannot get a foreign patent if you publish first. Sometimes people will sit on data just because the data are valuable, not because they can get a patent or copyright or anything else—they just know that someone will be willing to pay for access to it.

One of the most serious information dams is the nondisclosure agreement, which is usually associated with licensing agreements for data or research materials. In this case, the investigator says “I’m not going to talk about the data pertinent to this project under the following conditions,” and those conditions are specified in the license. That is one of the social norms that changed in many fields of science over the past two decades, and in particular in biomedical research. Whereas nondisclosure used to be anathema, now it is a norm in many academic institutions.

What we have got now is a real continuum of openness, and in fact, most universities would have a continuum within them. The same variation is found in the private sector. Some firms behave a lot like universities—they are very open and encourage scientists to publish in the open literature. In fact, there is evidence that firms who behave most like academic groups and encourage more open communication do better on Wall Street and secure more patents. That is, the more they look like academic groups, the better they do in the marketplace. On the other hand, there are firms that require very rigid protections on their data and there are groups within academe that collaborate with firms that require nondisclosure of all sorts of information.

Turning to public policy, one of the most troublesome aspects of trade secrecy is when it is applied to federally funded research. Nondisclosure can cloak publicly funded science when it is based on agreements with private firms or even academic groups to secure access to data, materials, or industrial funding. Terms of the agreement can hide not only the information generated out of that specific agreement, but also related information generated through federal dollars.

I was once presented with a nondisclosure agreement that took me aback. The main criterion governing public disclosure was whether it was going to harm the business interests of a particular company. Moreover, the final arbiters to interpret those business interests were officers of the corporation. That is, it did not matter who paid for the research or why it was generated, but rather what impact it would have on the company to whose data I might gain access. That makes sense from the perspective of a company, but as public policy, it is terrible—particularly so where federally funded research is encumbered.

So what are we going to do about all this? To simplify, the faster information flows, the faster innovation happens. Thus, we are all better off. That is true for firms, too, who are collectively better off to the degree that information flows freely, but each firm has a marginal incentive to lock up information for itself. We have got a prisoner's dilemma.

What can we do at the collective level? The first thing is to think about what the federal government can do. Policy options that will likely influence the academic information flow revolve around the following:

- Federal funding mechanisms
- Model agreements and guidelines

- Collective norms and standards among universities (licensing practices)
- Bayh-Dole statute modifications
- Patent law

First, giving out dollars for research can be tied to rules about open access to the resulting information. If you get a grant from DoE or NIH to do large-scale gene sequencing, for example, you agree to disclose the information rapidly. NIH and DoE have told their major sequencing groups “If you take our money, you’re going to play by our rules.” And the rules put a premium on quick disclosure.

Government can also develop policies enabling academic scientists to use discoveries made by the private sector. One example is the recent agreement between NIH and DuPont regarding sharing of Cre-lox mice. That is an example of NIH creating a model agreement to set a norm. That model should not be used just for this particular Cre-lox technology, but it also should serve as a model for other technologies that are generated by the private sector, but are useful for the academic community.

Licensing practices are where the most productive policy changes can take place. Many of the relevant policies may well be pursued outside government—by private institutions, both in academia and industry. Policies that allow academic research with low or no fees, carefully consider reach-through and exclusive licenses, and promote research and education as well as income could go a long way toward opening up the innovation process.

Should we go back to the Bayh-Dole statute? There are existing doorways in the statute that have not been used. The government can exercise “march in” rights if a patent owner fails to commercialize a useful invention. That may or may not pertain to secrecy. There is also an “exceptional circumstances” clause that says that if the public health is at stake or other government interests are better served, then government can override the usual default that gives rights in an invention to the grantee or contractor institution.

Some of the controversies surrounding access to federally funded sequencing of pathogenic microbes might be test cases. Government could at least extract commitments to license access to the information. In the defense sector, the government occasionally reclaims intellectual property to further defense needs. The government may be driven to

consider doing the same for health. If someone makes a discovery based on government funding that would be widely useful to those holding other government grants, the Bayh-Dole framework might be modified to enable ready access. There is some talk in Washington that government Bayh-Dole rights might be extended beyond government laboratories to those who are doing federally funded research in university or even small business settings. Some lawyers believe the existing Bayh-Dole statute could be interpreted that way, while others feel that an amendment would be needed.

However, revisiting or reinterpreting the Bayh-Dole statute might not be the best solution. Perhaps the most productive way to engage these questions, at least for now, is to look at the collective practices—especially at universities, because universities are the point at which industrial and government interests collide and where rapid flow of information is most important.

A final substantive point concerns the high transaction costs of protecting information. People who handle intellectual property, especially technology licensing offices, are well aware that the costs are high. It costs about \$10,000 to get a U.S. patent—double or triple that for significant protection abroad—and it costs money to maintain those patents. In the frenzy for patenting, very few of the patents obtained are ever used. Paying for an unused patent is a poor use of resources. One option is to be more selective in deciding what to patent. At a more common and mundane level, does it really make sense for all investigators to be signing material transfer agreements for everything?

Discretion in what is subject to formal written agreements, all of which cost substantial time and money to craft, would be most welcome. Most reagents and data shared among laboratories are not worth the time, trouble, or added bureaucracy. The counterpoint is that it is difficult to know the value of a discovery soon after it is made—but are we better off with a default path that presumes high value and encumbers information transfer, or one that engages the apparatus of information protection less often? It seems likely that this is another thousand-player prisoners dilemma, where the system as a whole is better off with open disclosure—even if a few individual institutions fail to get as rich as they might.

It seems likely that the effort and cost to negotiate transfer agreements, nondisclosure agreements, and the secrecy that precedes worldwide patent applications are higher than the value of rapid information flow those processes restrict. That is one of the main messages of the June 1998 NIH report on research tools. That report is quite sensible, but

we have not come to our senses yet in the academic community. The Cohen-Boyer patent set a precedent that academic institutions in particular would do well to emulate. Admittedly, it was a patent that many lawyers thought would be invalidated if it were ever challenged, but it was never challenged—in part because Stanford and the University of California kept the price low.

Those universities did two important things. First, they did not charge academic researchers for use of the technology. Second, they kept fees low for the industrial users. The licensing agreement for that patent did set a precedent that has surfaced as a reach-through problem. Most of the quarter of a billion dollars that came to the University of California and to Stanford because of that patent came from its reach-through provision. Licensees who sold a product that used recombinant DNA, even if it was a product outside the scope of the Cohen-Boyer patent claims, agreed to pay a royalty to UC-Stanford. It was those royalty streams on products made using recombinant DNA, rather than the up-front licensing fees for use of the process and vectors, that created most of the money. So Stanford and UC created a precedent for exempting academic researchers from needing a license—a precedent that may be fading into oblivion—but the same licensing agreement also set the stage for reach-through that we are grappling with now. From the perspective of a policy analyst, it seems we have neglected the useful precedents set by Cohen-Boyer and have adopted the most troublesome one.

Standards are unstable and need careful inspection in academic patenting, especially regarding two licensing practices that should raise red flags. One is exclusive licensing and the other is reach-through. I have already addressed reach-through; the other trouble spot is exclusive licensing, which ties use of an invention to one particular firm. If that firm turns out not to be the best choice, innovation can suffer. There are times when only an exclusive license will work—for example, when a discovery is quite close to final market, such as a therapeutic pharmaceutical that needs to be proven in clinical trials. But it should be the exception and last resort rather than the default path. Neither of these practices is wrong in all circumstances, but they seriously encumber future discoveries and should be pursued only with a good deal more care than some universities have been giving them.

Finally, and perhaps most important, the notion of technology licensing offices as profit centers needs to be balanced against the traditional research and educational missions of universities. We will know that is

happening when we hear university presidents making pronouncements about technology licensing policies, and when they assert that academic norms of openness should trump short-term licensing royalties. Articulation of academic norms is probably the most important policy step that needs to be taken. There are good models in place. If you look, for example, at the mission statement for MIT's Office of Technology Licensing, it makes clear that its money-making goals are subordinate to MIT's educational and research missions. The fact is, however, that the technology licensing offices at many other institutions do not operate under this or a similar mission statement.

If there were one thing I could do to reduce the level of secrecy that can impede innovation, it would be to elevate norms of open communication of knowledge in academe, and clearly to subordinate the business interests in intellectual property management to education and research at academic institutions. It is unrealistic, and would be counterproductive, to expect private firms (especially small ones) to follow the same norms of openness, but if I have convinced you of anything, I hope it is that there is enormous value to all players—including those in industry—in having academic information flow freely.