Volume Title: The Rate and Direction of Inventive Activity Revisited

Volume Author/Editor: Josh Lerner and Scott Stern, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-47303-1; 978-0-226-47303-1 (cloth)

Volume URL: http://www.nber.org/books/lern11-1

Conference Date: September 30 - October 2, 2010

Publication Date: March 2012

Chapter Title: Introduction to "The Rate and Direction of Inventive Activity Revisited"

Chapter Authors: Josh Lerner, Scott Stern

Chapter URL: http://www.nber.org/chapters/c12344

Chapter pages in book: (p. 1 - 23)

Introduction

Josh Lerner and Scott Stern

I.1 Introduction

Innovation—whether in the form of new products such as the iPad, new ways of incorporating process technologies such as bar coding, or new management practices—is critical to economic growth. This is particularly true in mature economies such as the United States and Europe, where pressing fiscal and demographic challenges preclude many other avenues to growth. But despite the critical nature of innovation, much remains unclear as to how nations, firms, and academic bodies can encourage this activity. While impressive strides have been made in understanding the economics of innovation over the past few decades, much about this activity remains uncertain or even mysterious.

This volume explores what we do and do not know about this critical area. It is based on the proceedings of the National Bureau of Economic Research (NBER) 50th Anniversary Conference in honor of the influential 1962 volume, *The Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by Richard Nelson. We saw the anniversary of that volume—seen by many as having ushered in the modern era of study of the economics of technological change—as a timely opportunity to not only take stock of the economics of innovation and technological change,

Josh Lerner is the Jacob H. Schiff Professor of Investment Banking at Harvard Business School, with a joint appointment in the Finance and the Entrepreneurial Management Units, and a research associate and codirector of the Productivity, Innovation, and Entrepreneurship Program at the National Bureau of Economic Research. Scott Stern is the School of Management Distinguished Professor of Technological Innovation, Entrepreneurship, and Strategic Management at the Massachusetts Institute of Technology Sloan School of Management and a research associate and director of the Innovation Policy Working Group at the National Bureau of Economic Research.

but also to bring together leading scholars to identify the shape of the field going forward.

As the discussions by Arrow, Nelson, and Rosenberg and Stern later in this volume highlight, the backdrop for the 1960 conference was the growing recognition of the role of technological change in economic growth. This insight—which grew out of the insights of Abramovitz, Kendrick, and Schmookler, and the key work of Solow—highlighted that increased inputs (e.g., more capital expenditures and workers) could only explain a modest fraction of American economic growth over the past century. As Nelson noted in his introduction to the 1962 volume, "The lion's share had to be attributed to something else: to increased productivity or efficiency."

While this insight sparked a desire to understand the nature of technological innovation, there was also a more practical motivation. The Soviet Union's launch of the Sputnik satellite had raised alarms about the United States' competitive positioning, and led to a need to better understand the circumstances through which scientific insights could be translated into new defense and space technologies. Along with the creation of the National Science Foundation and the creation and availability of more data on research and development activities, the time was ripe for a more systematic research program in economics on both the causes and consequences of invention and technological change.

The 1960 conference brought together many leading thinkers of the era, and resulted in a volume whose influence extended well beyond the typical conference volume. A number of papers have resonated through the decades, but none more than the final essay, Arrow's "Economic Welfare and the Allocation of Resources for Invention." In it, Arrow lays out the implications of the conflict between the low social cost of using knowledge and the high cost of producing it, and the subtle ways in which information and knowledge are distinct economic goods. The implications for firms seeking to appropriate returns and for social welfare are substantial, as the author explicates. The five thousand-plus citations that this essay has on Google Scholar are a testimony to the power of these ideas.

But to focus on this one essay misses the richness and breadth of the 1962 volume. The conference brought together a rich array of methodologies, from theory to large-sample empirical analyses to economic history to case studies. The range of topics was broad, from the nature of appropriability to the role of organizational structure in shaping research and development productivity. Not surprisingly, (in light of both the times and the affiliation of a number of the authors with the RAND Corporation), there was a heavy emphasis on the nature of publicly funded innovation, particularly in the defense sector. The conference, as its subtitle suggested, also explicitly sought to draw in perspectives from other social sciences. Thus, the volume simultaneously provided general building blocks for understanding the innovation process and reflected the issues of its day.

The current volume builds on this legacy. Organized under the aegis of the NBER's Innovation Policy Working Group, the conference and this volume seek to honor and assess the original volume, and sponsor new theoretical and empirical contributions on fundamental questions in the economics of innovation and technological change. An explosion of empirical and theoretical research in the economics of technological change, as well as contemporary policy challenges, suggests an opportunity for reevaluation of the traditional innovation policy framework.

Among the questions that we sought to grapple with were:

- How do innovation and diffusion depend on the institutional environment in which new technology is developed and commercialized, and how are the drivers of innovation changing over time? Does the pervasive diffusion of information technology impact the economics of knowledge accumulation itself?
- What is the role of "open" research environments (from scientific communities to the open-source software movement) in innovation? What are the economic and institutional drivers of open-access versus proprietary innovation models, and how does institutional design impact innovation outcomes?
- What determines the allocation of research investment between the public and the private sector (and what should determine that allocation)? What role do universities (and other nonprofit research institutions) play in long-term technical change and economic progress?
- How do innovation and diffusion impact economic growth? Has technical change moderated or exacerbated macroeconomic fluctuations? What is the relationship between innovation and economic inequality, both within and across countries? What is the role of innovation—as a driver or a remedy—in the current economic crisis?

We sought to achieve these goals through two approaches. First, we circulated a call for papers, and encouraged submissions from leading scholars. These essays were refined through discussions in a July 2009 preconference at Laguna Niguel, California, and formal presentations and discussions in the September 2010 conference in Warrenton, Virginia. We recorded and edited the discussants' remarks from this conference to give a fresh perspective on the issues raised by the authors. In addition, we incorporated into the conference three panel discussions, which took a broader view of the issues under consideration. The key presentations from these discussions (though not the lively back and forth that ensued) are also captured in this volume.

At the same time, we should acknowledge that in some respects, we were less ambitious than the 1960 volume. Given the explosion of economics research into innovation, not to mention the great growth of work on the topic in the strategy, technology management, and social psychology literatures, we decided to keep a sharp disciplinary focus. (It should be noted,

though, reflecting the broadening of the economics literature during this period, a number of papers cross over into topics that have typically been the purview of economic sociologists and other social scientists.) Nor did we try to explicitly duplicate the many case studies in the 1962 volume, though many papers use field research techniques.

The success of this conference was a function of many people and organizations. The Scientific Committee—Philippe Aghion, Ken Arrow, Richard Nelson, Manuel Trajtenberg, and Hal Varian—helped to shape the vision and agenda for the conference. They also helped considerably to boost this effort through their contributions. Both Marty Feldstein and Jim Poterba, the current and former NBER presidents, were uniformly supportive of this idea. Patrick Gaule, the 2010 to 2011 NBER Innovation Policy fellow, played a key role in organizing the conference and the production of the volume. Hal Varian gave a thoughtful and provocative after-dinner talk at the event. Carl Beck, Brett Maranjian, and Rob Shannon of the NBER conference department provided critical logistical support. The Kauffman Foundation was generous in their support of this initiative, playing a key role in supporting the Innovation Policy effort at the NBER more generally as well as funding this particular conference. We are particularly grateful to Bob Litan, Carl Schramm, and Bob Strom.

I.2 Broad Themes

In developing the agenda for the conference, we sought to focus on forward-looking research that offers direction for the field going forward. As one considers the essays as a whole, it is useful to highlight four thematic clusters: the university-industry interface, the interdependency between market structure and innovation, the sources of innovation, and the social impact of innovation.

I.2.1 The University-Industry Interface

One topic that received relatively little scrutiny in the 1962 volume, but much more attention here, was the university-industry interface. Certainly, many authors espoused a belief that basic research was important for innovation, and called for more work in the area. But to a large extent, much of the focus was on government and corporate research. This volume considers the consequences of academic research to a considerably greater extent.

Of course, this shift in emphasis largely reflects real-world developments over the past five decades. The passage of the Bayh-Dole Act in 1980 and the proliferation of multibillion-dollar companies founded on university research (from Genentech to Google) have vastly increased the economic profile of these activities. Moreover, the revolution in data availability—particularly, the detailed data on citations of papers and patents—has facilitated work in this area.

Gans and Murray take a broad, conceptual look at the issues associated with funding academic research. They begin with a paradox: when agencies funding scientific research emphasize basic research over translational projects, they are criticized for their impracticality, but when they emphasize near-term mission-oriented R&D projects, they are criticized for crowding out what industry would have done otherwise and backing redundant efforts. To help sharpen our thinking about these issues, the authors present a model in which the supply of and demand for public funds plays out in a world where private funding sources also exist.

In their model, public officials can decide not simply which projects to fund, but also what requirements regarding scientific openness to add. They show that the choices regarding funding sources—and the impact of publicly imposed requirements around disclosure—will vary not only with the scientific merit of the research proposal, but also with the immediacy of its applicability to commercial uses. In particular, they highlight that providing unrestricted public funds (i.e., without any disclosure requirements) may lead to many researchers who would otherwise be industry funded accepting public dollars: this can actually lead to fewer projects being funded overall without consequent gains in openness. Though some of the key issues raised here have long been recognized—indeed, both Nelson (1959) and Nelson's careful study of the transistor in the 1962 volume raised related issues—Gans and Murray provide fresh insight into the subtle ways that public and private funding interact, and the role that government policy (e.g., mandating openness) plays in shaping the production and use of knowledge.

Azoulay, Graff Zivin, and Sampat look at the consequences of academic mobility: to what extent does the movement of high-achieving faculty members affect both scientific and commercialization activities at their old and new schools? To examine this, they look at articles published by and patents granted to the mobile scientist before he departed for the new school, comparing these to similar outputs by scientists who did not move. In this way, they hope to limit the heterogeneity that can distort simpler comparisons.

The analysis suggests that the citations to a departing scientist's articles from the university where he or she departs are barely affected by the move. But citations to the departing scientist's patents (whether made in articles or patents) decline sharply at the originating school. This suggests that the physical proximity of the researcher is important to ensuring knowledge flows to industry. Not surprisingly, citations to the scientist's work at his or her new location increase dramatically once the move is complete. The authors offer the intriguing conclusion that barriers to scientific mobility may actually be socially detrimental, as they prevent the kind of knowledge gains from the mixing of ideas.

Kahn and MacGarvie also explore the impact of scientific mobility, focusing on the Fulbright Foreign Student Program, which since 1946 has brought students from many countries to undertake graduate studies in the United

States, with the expectation that they spend at least two years in their home nation before they can return. Like the prior paper (though with a substantially smaller sample and less exact controls), they compare the output of the Fulbright scientists with a set of otherwise similar scientists who studied in the United States without such a return requirement.

Tracing the subsequent career of these researchers, the authors find that the Fulbright scientists (relative to the controls) spent more than twice as many of their postgraduation years outside the United States when compared to controls. While the program does increase collaborations between US scientists and those based in the emerging world, Fulbright scientists from poorer nations or those with a weaker scientific tradition have fewer publications and less of an impact. This effect is not seen among those scholars from wealthier nations or those with a stronger scientific base.

These last two chapters suggest one profound difference between the two volumes: the vast increase in data availability. The richness of citation and personnel data has given us both the ability to test relationships that previously could only be discussed abstractly or else explored only in case studies. It also underlines the importance of phenomena that were previously not fully appreciated, such as the impact of geographic proximity on knowledge spillovers, a topic that received little mention in the 1962 volume.

I.2.2 Market Structure and Innovation

A second cluster of chapters focuses on a question that goes back at least to Schumpeter, but was brought back to prominence within economics with Arrow's 1962 paper: what is the relationship between market structure and innovation? Bresnahan, Greenstein, and Henderson focus squarely on a central puzzle in this line of research: why are incumbents who are able to succeed within a given technological trajectory often so ineffective at being able to take advantage of a new technological trajectory? This question is particularly salient once one considers the many advantages that incumbents are able to leverage in introducing new technology.

Bresnahan, Greenstein, and Henderson undertake detailed case studies of two historically important transitions—the introduction of the personal computer (PC) and the browser—to evaluate this question. Their analysis allows them to both assess the adequacy of existing theories (e.g., anticannibalization concerns, or the potential for organizational barriers within incumbents) and to identify key patterns that seem to characterize the process of creative destruction. Their analysis points to a novel driver of creative destruction—*diseconomies of scope* induced by the presence of necessarily shared assets within the firm. When the strategic commitments made by an incumbent are necessarily reflected in business activities for *both* the old and the new technological trajectory, the incumbent may not simply be able to create a "firm-within-a-firm" to preempt competitive entry. The fact that the incumbent must simultaneously sell both the new and the old technologies

may put them at a disadvantage in both technologies relative to an entrant; these disadvantages can be observed through the significant organizational conflicts that accompany technological transitions.

The case study approach (though out of favor at traditional economics journals) allows Bresnahan, Greenstein, and Henderson to undertake a close reading of the evidence. This leads to a novel hypothesis about the underlying forces that may be at the heart of many cases of incumbent failure in the face of the gale of creative destruction.

Spulber offers a complementary perspective on this question by considering how the strategic interaction between incumbents and innovators in the market for ideas shapes (and is shaped by) the potential for product market competition. On the one hand, if the market for ideas is efficient (e.g., there can be perfect, low-cost transfer of both new designs and process innovations), then incumbents and entrants will have an incentive to cooperate (rather than compete) in the commercialization process. However, when technology transfer (of either product designs or processes) is imperfect, then innovators will have an incentive to enter the product market (and so start-up innovation will be associated with increased competition).

The question then arises: when is entry more likely? While Spulber considers a range of cases (and some of the results are subtle), an overarching lesson of the analysis is that the incentives for entry are higher when the underlying technologies are more (horizontally) differentiated from each other. Since the gains from cooperation are higher when the degree of differentiation is lower, the likelihood of entrepreneurial entry is higher under conditions of high product differentiation and imperfect technology transfer. Spulber highlights the idea that the impact of start-up innovation on market structure depends crucially on the nature of strategic interaction between start-ups and established firms, and that such strategic interaction is itself going to depend on the specific nature of the innovations impacting an industry at any point in time.

Daron Acemoglu also considers how strategic interaction impacts the relationship between innovation and incumbency, but places emphasis on a dynamic setting that incorporates not simply the rate but also the *direction* of innovative activity. Specifically, Acemoglu considers an environment where there are multiple potential "research lines," but only one research line is commercially active at any point in time. Acemoglu is particularly interested in cases where there is a chance that the commercially active research line will at some point be made obsolete (e.g., as the result of exhausting a natural resource), and focuses attention on the underlying incentives to invest in the alternative (but not yet commercially viable) technology line. Since the returns from innovation are only realized for those generations where the research line is commercially active, the private returns to innovation in the alternative line will be low unless there is a high likelihood that the currently active line is about to made obsolete.

This analysis highlights an important externality: while the social planner would prefer investments on the alternative research line (so that the level of this technology is at a high when the other technology is made obsolete), the private incentive to invest in the alternative research line is too low. In considering the impact of alternative policy approaches, Acemoglu surfaces a novel argument for public funding of a diverse set of research approaches: researchers with different incentives, capabilities, or perspectives may contribute to a more diverse research portfolio, and so contribute to economic growth.

Finally, Carl Shapiro turns our attention to how these types of analyses can inform innovation policy. In an essay that clearly captured the prize for the most clever chapter title, Shapiro offers a synthetic assessment of how the lessons of the economics of innovation inform merger analysis. Shapiro contrasts two dominant perspectives that inform merger analysis: Arrow versus Schumpeter. Where the Arrow approach suggests the positive impact of product market competition on innovation, the Schumpeter perspective focuses instead on the innovation inducements due to scale and the prospects of market power.

Shapiro emphasizes that these two perspectives—often taken to be contradictory—are not at all incompatible with one another, at least as they apply to policy analysis. While recognizing that the relationship between innovation and market structure is quite complex, Shapiro focuses on three key principles that build on the insights of both the Arrow and Schumpeter perspectives. By so doing, they can help us understand the impact of mergers on innovation incentives. Specifically, Shapiro highlights the idea that innovation is enhanced when (a) firms have the prospect of either gaining or protecting sales by providing additional value to consumers (the Contestability Principle), (b) the level of intellectual property protection is higher (the Appropriability Principle), and (c) complementary assets can be combined to enhance innovative capabilities (the Synergy Principle). Illustrating the role of these principles in clarifying the innovation impact of mergers in particular cases and circumstances, Shapiro's essay highlights the role of careful economic analysis in helping to clarify policy analysis, and how long-standing conceptual frameworks can be enriched by careful, formal reconsideration.

Taken together, this second group of essays provides a very useful delineation of our understanding of the relationship between innovation and market structure. Fundamentally, the economic analysis of market-based innovation incentives relies on a dynamic understanding of how innovation shapes (and is shaped) by industrial organization. These dynamics are themselves dependent on both the nature of competitive interaction between different technologies, and the organizational consequences of innovation. Interestingly, as a number of the authors and discussants remark, there are too few systematic studies of this process, and it has been difficult to bridge

the gap between the types of qualitative and theoretical insights emerging from these chapters and the type of empirical research that tends to dominate scholarly discussion. That gap surely represents an important direction for future research.

I.2.3 The Sources and Motivations of Innovators

A third cluster of chapters focuses more directly on the incentives and motives of inventors and innovators, and highlights the role of institutions in shaping the behavior of individuals and firms in producing new technologies.

Moser and Rhode consider the impact of formal intellectual property rights—specifically, the Plant Patent Act of 1930—on innovation. While standard economic theory suggests that the introduction of formal intellectual property protection should enhance appropriability and the incentives to innovate, there are only a very small number of cases where economists are able to observe whether a *change* in intellectual property law results in a change in the degree or nature of innovation in a particular area. Moser and Rhode focus on the impact of the Plant Patent Act on patenting and innovation in roses, which were the plant variety most impacted by the Act (nearly half of all plant patents between 1930 and 1970 were for rose varieties). An important element of their analysis is that they are able to distinguish between the impact of the Act on patenting (which of course increased) versus the impact on innovation (which they measure in terms of new rose registrations).

Their empirical evidence poses an important challenge for the standard theory: after 1930, the number of registrations by American nurseries actually fell, and European nurseries accounted for an increasing share of new rose registrations. Instead of increasing the rate of innovation, it seems that the Plant Patent Act may have had the consequence of increasing the relative importance of commercial nurseries relative to hobbyists in the American industry, and spurred the use of patents as a defensive and strategic tool in the context of litigation. Importantly, the findings of Moser and Rhode are made more plausible by the fact that there are important nonpecuniary motivations on the part of (at least an important group of) innovators in this area; prior to the Plant Patent Act, both hobbyists and public sector breeders played an important role in establishing distinctive American rose varieties, but their role was diminished thereafter.

Mokyr and Meisenzahl offer a complementary perspective, offering an economic history approach of the peculiar nature of innovators and their motivations and interests during the British Industrial Revolution. Their analysis focuses in particular on the body of individuals who advanced technology and innovation during this period. Moving beyond the celebration of specific individuals responsible for macroinventions such as the steam engine, Mokyr and Miesenzahl focus in particular on "tweakers"—indi-

viduals involved in the process of incremental improvement and refinement central to cumulative technical progress. Their analysis builds on a novel database of such individuals, and offers a portrait of their careers.

Most notably, Mokyr and Miesenzahl provide suggestive evidence that formal intellectual property rights such as patents likely played (at best) a limited role in the incentives and compensation of tweakers. Instead, their primary incentives seem to be associated with the reputation-based and first-mover advantages associated with innovation, as well as the rewards to be gained through prize mechanisms or nonpecuniary rewards such as membership in societies and the like. Similar to Moser and Rhode, this analysis suggests that, in the presence of multiple innovation incentive instruments, the traditional arguments for patents may be weakened. Perhaps more importantly, the chapter opens up a critical window into both the motives and training underlying incremental innovation. As such, the chapter addresses an important concern: one of the enduring challenges among students of technology has been the difficulty of moving beyond the study of formalized, often patent-oriented innovation to the many more informal processes through which technologies are improved.

Finally, Boudreau and Lakhani directly confront the impact of innovator preferences on innovation and research productivity. Their chapter reports on an actual field experiment that tests for the influence of "sorting" on innovator effort. They focus in particular on the potential heterogeneity among innovators in whether they prefer a more cooperative versus competitive research environment. The focus of the field experiment is a real-world multiday software coding exercise in which participants are able to express a preference for being sorted into a cooperative or competitive environment (i.e., incentives in the cooperative environment are team-based, while incentives in the competitive environment are individualized and depend on relative performance). Half of the participants are indeed sorted on the basis of their preferences, while half of the participants are assigned to the two modes on a random basis.

Boudreau and Lakhani find strong evidence that sorting matters: those who prefer a competitive regime exert twice as much effort when they are assigned to that regime, and those who prefer a cooperative regime also increase their effort by 50 percent when they are assigned to their preferred regime. In addition to the sheer novelty of their experimental approach for the economics of innovation, their substantive results once again highlight the important role that motivation and preferences play in understanding innovative activity. Not simply a matter of providing appropriate incentives for effort, innovators exhibit strong preferences over the organization and incentives in their work environment, and the ability to match workers with their preferences has significant effects on overall research productivity.

Similar to the findings of the earlier volume, these detailed empirical stud-

ies of the motives of innovators pose a significant challenge to traditional economic models of incentives for innovators. For example, in all three studies there seems to be a significant role and interaction with the broader innovation "community." The historical evidence from the tweakers of Mokyr and Miesenzahl and the rose growers in Moser and Rhode suggests that the patent system, in particular, either played a limited role or (in the case of Moser and Rhode) may actually have undermined innovation incentives on the part of individual growers. As we discuss further later on, a great deal of the panel discussions and commentary at the conference focused on the drivers of volunteer contributors, which we may refer to as "wiki-motives." What is the impact of traditional innovation policy instruments such as patents or prizes in environments when innovators are motivated by recognition and community concerns rather than monetary payoffs? How important are such motives in understanding aggregate innovative effort, and how has this varied across time and context?

I.2.4 The Social Impact of Innovation

A final grouping of chapters grapples with what is undoubtedly the most challenging issue in the economics of technological change: assessing the social consequences of innovation. As Paul David points out in his discussion, an implicit assumption of policymakers today is that more innovation is undoubtedly a good thing. Economic theory takes a more cautious view, suggesting that the private sector can engage in too much as well as too little innovation.

Part of the reason for the presence of misconceptions, of course, is that the assessment of innovations' social impact is a daunting task. While industrial organization economists have made great strides in developing structural models that allow social welfare calculations over the past few decades, the types of dynamic changes that characterize important innovations defy ready characterization. The three chapters in this section take differing approaches to this challenging issue.

Lerner and Tufano explore the broader impacts of financial innovation. This class of breakthroughs—which attracted no real discussion in the 1962 volume—has broad impacts: not only do financial services represent a significant economic share (estimates in the United States run as high as over 30 percent[1]), but in an ideal world, they enable households to have new choices for investment and consumption, and firms to raise capital in larger amounts and at a lower cost than they could otherwise. At the same time, financial innovation has been criticized by Paul Krugman and others as a key driver of the recent global financial crisis.

In this chapter, the authors review the literature on financial innovation

1. Available at: http://www.ggdc.net/databases/10_sector.htm.

and highlight the similarities and differences between financial innovation and other forms of innovation. The chapter proposes a research agenda to systematically address the social welfare implications of financial innovation. To complement existing empirical and theoretical methods, the authors propose (and take some initial steps toward) the examination of case studies of systemic (widely adopted) innovations, explicitly considering counterfactual histories had the innovations never been invented or adopted.

Field takes a close look at the boom-bust pattern that characterizes many industries. During the boom period, there is a dramatic accumulation of physical capital—think of the huge efforts to lay broadband during the Internet boom of the late 1990s—followed by a contraction. In the short run, it is easy to see how such a contraction leads to a decline in productivity, as excess capacity lies unused.

But this chapter is interested in a more challenging question: what are the long-run consequences of these boom-bust cycles? To what extent are the resources accumulated during booms the appropriate ones, or do they represent wrong-headed investment decisions brought about by a frenzied market? To examine these questions, Field examines the experiences of railroads during the Great Depression. This was a difficult period for the industry: the economic downturn, along with increased competition from automobiles and trucks, led to a sharp contraction in demand for railroads. Moreover, access to capital was largely cut off after a period of heavy expenditures. He shows that the industry undertook a major restructuring to utilize labor and capital resources more effectively. Both capital and labor inputs declined substantially. Yet logistical innovation enabled railroads to record slightly more revenue ton miles of freight and book almost as many passenger miles in 1941 as they had in 1929. Adversity seems to have triggered a wave of innovation in this industry.

In the final chapter in this cluster, Bresnahan focuses on the recombination and reuse of key general purpose technologies (GPTs), which he defines as widely used discoveries capable of ongoing improvement that enable complementary innovations. He argues that a critical factor behind the creation of these key technologies is the extent to which the broad prospects for reuse can be anticipated.

Bresnahan distinguishes between two kinds of knowledge. He argues that technical knowledge—the understanding of how a firm can transform a technology into a product—is relatively commonplace. But an understanding of market demand and how an invention might be used in other sectors—which he refers to as entrepreneurial knowledge—is a rarer and more valuable asset. Because of the scarcity of entrepreneurial knowledge, the returns from developing a GPT may be much lower than they would be otherwise. But over time, through a process of innovations and product introductions, this scarce entrepreneurial knowledge may become much more widely known. He illustrates his theory with a number of cases from

the information technology industry, where important GPTs were only developed after numerous false starts.

These three chapters take very different approaches to understanding the broader impact of innovation on social welfare. Despite the challenging nature of these questions, and the absence of well-accepted answers, the importance of this topic remains a major challenge to the economics of technological change.

I.3 Panel Discussions

In addition to the formal papers (and discussions), the conference included three panel discussions. By design, the panels were intended to be provocative, and to identify key research challenges going forward. Though each of the three panels were different in both style and substance, each significantly expanded the scope of discussion within the conference, and highlighted some of the central limitations of current models or empirical methodologies. The volume includes short contributions by nine of the panel chairs and participants, based on transcripts of their remarks.

The first panel—"The Impact of the 1962 *Rate and Direction* Volume: A Retrospective"—explicitly linked the 1960 and 2010 conferences, and included commentaries by two of the central participants in that earlier effort. Rosenberg and Stern began the discussion with a critical assessment of the 1962 volume, with a focus on identifying why that earlier volume turned out to be so influential on subsequent scholarship. The central contention of their remarks is that the Rate and Direction Conference can be interpreted as a reaction to the work by Solow and others highlighting the *aggregate* implications of technological change. More than simply a debate about the nature of the "residual," the 1960 conference focused attention on the central *economic* questions raised by inventive activity, innovation, and technological change. Specifically, the original conference highlighted (a) the nature of innovation as an economic good, (b) the economics of the organization of research and development, and (c) the industrial organization of innovation-intensive industries and sectors, with a particular focus on dynamics and evolution. As a marker in the history of economic thought in this area, a central contribution of the earlier conference was to crystallize the questions and issues that would come to dominate the *microeconomics* of innovation and technological change for the foreseeable future.

Nelson expanded on these themes. He focused on some broad lessons that have emerged since the earlier conference and also on important methodological issues that have been raised. Nelson noted that an important divide exists between the type of theory and empirical research emphasized within the United States (and within the NBER) and the interdisciplinary, evolutionary approach that has been emphasized by researchers such as those at the Science Policy Research Unit (SPRU) in the United Kingdom. Nelson

argued that some of the underlying tensions between these two camps were foreshadowed in the earlier volume: the largely empirical tradition pioneered by Kuznets and Schmookler (and reflected in the NBER Productivity Program and its growth under Zvi Griliches) sat alongside (sometimes uncomfortably) the detailed case studies or innovation systems studies emphasized within the evolutionary tradition.

Arrow took a broad view of the issues that both conferences grappled with. His comments crystallized why economists have had such difficulty in clarifying the nature of innovation as an economic good: "How can you have a theory of the unexpected?" Arrow highlighted the idea that the economics of innovation must confront and incorporate some of the unusual properties of innovation, both in terms of its production (e.g., the significant level of uncompensated effort toward inventive effort, in areas ranging from medicine to Wikipedia) and use.

These panelist remarks (and subsequent discussion) highlighted how the peculiar nature of innovation poses an ongoing challenge to theory and measurement. They illustrated why the wide-ranging and exploratory nature of the 1962 volume has had such a significant and long-lived impact on subsequent work.

In many ways, the second panel—"Innovation Incentives, Institutions, and Economic Growth"—built on the first panel, reconsidering the implications of innovation and technological change. Paul David opened that panel with a deliberate mission—"mass provocation." He focused his remarks on the underlying (though often implicit) assumption among economists that a higher rate of innovation is almost always preferred. David pointed out that the social impact of technological change depends not only on innovation but on diffusion. The ultimate impact of research investments depends on how those research investments are organized, and the complex process by which technologies are improved and adapted over time and context. Without considering the dynamic process by which social systems adapt and incorporate technological change, it is difficult to consider the net impact of new technologies on human welfare.

Philippe Aghion looked at a related question, the implications of advances in endogenous growth for both macroeconomics and microeconomics. Aghion argued that a major contribution of theories of economic growth that explicitly endogenize the production and diffusion of technology is to identify the potential policy impacts of different types of intervention. Aghion stated that contemporary policy matters insofar as it facilitates a higher level of innovative investment and shifts the long-run growth rate. A range of recent evidence highlights the role of ensuring the ability to protect ideas (e.g., a stronger patent system) in economic growth, and the potential benefits of "industrial policy" measures.

Paul Romer also contributed remarks to the panel (not included in this volume), focusing on the dynamic interplay between different types of the-

ory (e.g., verbal versus formal). Consistent with Aghion's discussion, Romer emphasized the specific contribution that models of endogenous growth have played; in one example, Romer highlighted the central role that appropriability conditions play in determining the rate of aggregate long-term technological change. The panel put a spotlight on the central role of policy and institutions in shaping the long-term rate and direction of technological change, and the value of bridging more narrow studies of the innovation process with more aggregate treatments in order to clarify the long-term drivers of economic growth.

These themes then were reinforced in the final panel discussion—"The Art and Science of Innovation Policy." After brief remarks by Bronwyn Hall, Glenn Hubbard focused on some of the challenges of developing and implementing well-designed innovation policy initiatives. Hubbard pointed out the disjunction between arguments for particular policies—for example, a particular tax rate or regulatory change—and the broader evidence that the rate and impact of innovation reflect broader measures of the overall innovation environment. Hubbard also emphasized the disjunction between academic and policy approaches. He also highlighted the role of certain types of institutions—for example, long-term interagency working groups—in facilitating a more sophisticated innovation policy-making process.

Dominique Foray reinforced these ideas, focusing in particular on the limited influence of economic science on policy making. Reflecting on his experiences within Europe, Foray argued that policy debates are often characterized by a low level of empirical sophistication, and that conditional statements or caveats often result in a diminished impact of rigorous economic analysis. Foray also highlighted that the bulk of innovation policy initiatives have been focused on enhancing the overall rate of innovation, but that an increasing share of innovation policy challenges are now about the direction of innovation (e.g., addressing climate change).

Finally, Trajtenberg considered the broader legacy of the Nelson-Arrow paradigm (with its focus on appropriability and the role of government support for early-stage research) on innovation policy. Trajtenberg highlighted that many of the central challenges facing innovation policymakers cannot be addressed directly through the Nelson-Arrow framework. For example, while the Nelson-Arrow framework assumes a single potential public funder, the question facing policymakers today is how much should an individual country fund, given the global nature of research and the potential to benefit from research conducted in other jurisdictions. Similar to the other panelists, Trajtenberg also remarked on the limited influence of rigorous economic analysis on actual policy, and suggests a focus on more policy-oriented research.

These panel discussions raised a rich array of issues. While there were more questions than answers, they suggested a variety of topics that should reward scrutiny by researchers in the years to come.

I.4 Crosscutting Insights and Themes

Taken as a whole, the chapters and discussions highlight some crucial and novel insights into the economics of innovation and technological change, and the role of policy and institutions in shaping innovation, diffusion, and ultimately, the social returns to technological change. While this volume cannot capture the full range of these more subtle implications, it is worthwhile to highlight a few central and novel ideas that were surfaced during the conference.

I.4.1 Innovation Externalities

The conference raised the hypothesis that the underinvestment problem is more pervasive, more subtle, and perhaps more pernicious that is usually understood. Building on the classic treatments of Nelson (1959) and Arrow (1962) emphasizing appropriability, a great deal of economics research has focused on how to provide sufficient market-based incentives for innovation (without inducing dissipative racing or rent-seeking).

A number of papers in the conference suggested, however, that our traditional understanding of the appropriability problem does not go far enough. Bresnahan, for example, emphasizes the idea that the history of general purpose technologies suggests that the conditions giving rise to their initial development usually arise in the context of a narrow application. This analysis suggests that innovation incentives are shaped by the prospective returns associated with that narrow application, rather than the returns associated with the diffusion of the general purpose technology. Externality problems can arise when the information about the potential impact of a new technology is widely diffused, so that commercialization of a general purpose technology depends on the coordination of multiple economic actors. In that case, no single actor can understand or appreciate the potential social impact of that innovation from an ex ante perspective. As Ben Jones emphasized in his discussion, "the fact that you can't identify the recombinant possibilities ex-ante means that you can't easily solve the bargaining problem in practice." Accordingly, the level of investment focused on general purpose innovations will be low.

Though different in its specifics, a similar theme runs through the analysis of Acemoglu. In that chapter, potential innovators will have little incentive to invest in an immature technology that cannot earn immediate commercial returns, even though the improvement of that technology over time will yield significant social benefits once an older technology becomes obsolete. Of course, it is possible that property rights could be specified in a way so that early innovators in alternative technologies retained some claim on the returns that ultimately arise from research lines that they are associated with; however, such rights would themselves pose a disincentive for later-stage innovators.

Gans and Murray broaden the scope further in their analysis of disclosure and knowledge accumulation. They highlight the idea that, even if the incentives for research investment are appropriate, the incentives for disclosing the knowledge resulting from that research are shaped by the strategic and institutional environment. Not only is there a significant gap between the private and social incentives for disclosure, but what seems to be a straightforward policy solution—such as mandating the disclosure of publicly funded research—can actually reduce the net level of disclosure (by pushing researchers to accept privately funded research contracts that mandate secrecy).

Together, these insights (and others dispersed throughout the volume) suggest that a central insight of the 1962 volume—the gap between the private and social returns to inventive investment—remains not only relevant today but is likely to stand as a central concern in the economics of innovation for the foreseeable future.

I.4.2 Agency Costs and Innovation

Another theme had to do with the impact of agency costs on the success of innovation projects. This theme—which was only dealt with implicitly in the 1962 volume—cut across a variety of the papers. Innovative projects are a natural place to see agency problems at work. Typically, there is substantial uncertainty as to whether a project will work or what the output will be, making the monitoring of effort or contracting on outputs difficult. The researcher is likely to have far more information about the intricacies of the project than his or her supervisors or financiers. In many cases, there are few tangible assets associated with the project making contracting particularly difficult. Market booms and bust may lead to dramatic shifts in the assessment of projects and availability of financing.

Against this backdrop, it is not surprising that economists have highlighted two important agency problems. The first has to do with the way in which innovators are rewarded. The second has to do with the way in which the firm itself is structured, and in particular the trade-offs associated with firm scope. While a number of the papers in the 1962 volume explored the role of individual researchers and the organization of firms, few (the Rubenstein paper is a notable exception) grappled with agency issues. Of course, this reflects the fact that agency theory was not formalized until the 1970s, and that the extent of agency problems in innovation was not thoroughly delineated until works such as Holmstrom (1989) and Aghion and Tirole (1994).

In this volume, the impacts of agency problems in innovation are far more widely recognized. Mokyr and Miesenzahl highlight the many incentives that were offered to British inventers during the Industrial Revolution, ranging from prizes to patents to consulting opportunities. The results from large-sample studies of the effects of incentives on individual researchers

are more mixed. Boudreau and Lakhani show how incentives impact the output of software programmers, but that this relationship is mediated by the coders' preferences regarding incentive schemes. Moser and Rhode show that increased incentives—in the form of stronger intellectual property rights for plant varieties in the United States—seem to have not led to more innovation by American nurseries relative to their European counterparts.

At the firm level, Bresnahan, Greenstein, and Henderson explore why two very successful software firms became increasingly unable to respond to competitive threats from new rivals. They argue that neither fears of cannibalization nor the inability to recognize competitive threats were critical: rather, the need to share key assets across old and new businesses created severe organizational conflict.

This volume, then, reflects the increased appreciation of agency problems as a barrier to innovation, and the organizational response that can address them. In a theme we will return to in the final section of this essay, the proliferation of new organizational firms and incentive schemes in research-intensive industries suggests that opportunities for research into these issues are far from being exhausted.

I.4.3 The Analysis of Innovation Policy

A third commonality in the volume is the focus on policy analysis, and the role of innovation within economic policy more generally. While a whole collection of chapters in the current volume focus on the university-industry interface, these interactions were (essentially) in the background of the 1962 volume. Whereas universities were once seen as ivory towers, policymakers have increasingly come to regard innovation resulting from university research (or collaborative projects) as central drivers of regional economic growth.

Though the university-industry interface is seen as ever more important, there has been less attention to how the rules and policies governing these interactions matter. For one example, though the policy origins of the Fulbright program are remote, the program is a primary driver of how foreign graduate students in the United States are trained. The Fulbright program rules have an important impact on the ultimate research productivity of those involved in the program, particularly those from less developed environments. Similarly, Gans and Murray emphasize how the rules governing the disclosure of publicly funded research not only affect that research directly but also the governance of research that is funded by the private sector. As Scotchmer emphasized in her discussion of Gans and Murray, "disclosure rules and other details of public funding should be chosen with an eye to how they affect the funding choices of innovators." More generally, the conference highlights the central role of economic governance and policy for understanding the university-industry interface, and points toward the value of examining specific policies and institutions.

A second domain for policy analysis is the intersection between innovation and antitrust. As highlighted by Shapiro, the impact of antitrust policy on innovation is increasingly salient, and an emerging set of principles may allow economists to offer more concrete policy guidance for policymakers in this area. Indeed, Shapiro builds on a number of prior analyses published in the NBER's "Innovation Policy and the Economy" series in developing these principles. However, there is still a significant gap between the type of principles emphasized by Shapiro and the ability to apply those principles in real time to cases that pose potentially significant antitrust and innovation incentive concerns.

Finally, it is useful to note what might be seen as a nonfinding: in the one chapter in this volume that directly examines the impact of intellectual property policy, Moser and Rhode find little evidence that enhanced patent protection enhances the rate of innovation (and indeed one can interpret their findings as suggesting the opposite). Moreover, as emphasized in the panel discussion of Paul David, it is not clear that the primary goal of innovation policy should simply be to maximize the rate of innovation itself. Ultimately, the policy debate over intellectual property should be guided by the goal of maximizing social welfare, not simply innovations.

I.4.4 New Approaches for Studying Innovation

Over the past few decades, there has been a much greater emphasis placed in the economics literature—led by fields such as labor and development—on ensuring the careful identification of the causal effects behind the phenomena under study. A particular emphasis has been on the development of research methodologies that can address concerns about causality, such as experiments and regression discontinuity approaches.

This movement has posed real challenges for students of the economics of technological change. In the overwhelming majority of cases, given the substantial economic stakes at work and the magnitude of the investments, it is impossible to get a corporation or government to agree to run an experiment in lieu of its usual project management approach. It is a very different thing to randomize the teaching of a few third grade classes than the funding of a potentially multibillion-dollar drug! Moreover, the complexity of the innovation process does not lend itself well to the classic hour-long laboratory experiment.

As a result, the approach to addressing causality concerns has been twofold. First, there has been an emphasis on the development of careful matching approaches, which enables the undertaking of difference-in-difference analyses with a minimum of potential biases (as illustrated by the Azoulay, Graff Zivin, and Sampat chapter in this volume). The second has been to find circumstances where some exogenous shifts have allowed the use of an instrumental variable, such as the consequences of the rise of the Nazis and the consequent expulsion of Jewish scientists (Waldinger 2009) and the US

pension reforms that greatly increased the flow of funds into venture capital in the early 1990s (Kortum and Lerner 2000).

This volume has two empirical chapters that represent substantial methodological departures in the economics of technological change, and thus deserve some special comment. First, Boudreau and Lakhani adopt a field experiment approach, exploiting the flexibility of web-based software development schemes to offer different incentive schemes to programmers. This approach seems to be an extremely promising one. While it can be argued that the incentive issues are different in software than other arenas—with the relatively finite project scale and the ability of skilled programmers to address a relatively broad array of challenges—this chapter should trigger many other innovation experiments in the years to come. Second, Lerner and Tufano adopt a counterfactual approach to explore the social consequences of a number of financial innovations. While this methodology remains controversial in economic history, it seems desirable to further explore its applicability to addressing some of the broad challenges in assessing the social impact of innovation.

Finally, the transformation of another methodology well represented in the 1962 volume deserves comment. As Nelson observes in his remarks, many of the participants in the original conference felt that some of the most valuable insights came from the case studies that represented a substantial share of the program.

The history of case studies—or to use the preferred modern parlance, clinical studies—in economics over the past century has been a bumpy one. The representation of such studies in major journals dropped precipitously after the 1950s, reflecting both the strides made by theoretical and empirical researchers and the uneven quality of many of the published cases. But some in the profession still feel that such studies can yield valuable insights into the richness of real-world phenomena, and suggest future directions for theoretical and empirical explorations. Such sentiments led to the initiation of the Sloan "Pin Factory" Project at the NBER and the launch of the clinical section of the *Journal of Financial Economics*.

Here, the field-based methodology is still present, though with a twist. There are a number of case-based chapters in the volume, including Bresnahan, Greenstein, and Henderson, Gans and Moser, Lerner and Tufano, and Moser and Rhode. These chapters can be seen as continuing the field-based tradition in the 1962 volume, but with a more developed theoretical and/or empirical structure than many of those earlier works.

I.5 The Agenda Going Forward

While the range of topics covered in this volume is substantial, there are also substantial lacunae. It is useful to highlight three critical issues that deserve more attention going forward.

The first of these has to do with the globalization of innovation. During the twentieth century, innovation was dominated by a handful of nations, such as the United States, Germany, and Japan. The twenty-first century—as witnessed, for instance, by the changing distribution of patent filings—has already seen a substantial disruption to this established order.

Lying behind this shift is a variety of factors. Governments such as those of China and Singapore have accepted the importance of academic science to economic development, and sought to lure faculty to their national universities, often with substantial investments. Corporations have increasingly sought to exploit the substantial cost savings associated with engineering talent in emerging economies, and what has often started with the overseas transfer of routine technical tasks has expanded in scope and magnitude. Venture capitalists, whether based in Silicon Valley or in emerging economies, have also been an important engine to the diffusion of research and innovative activities.

This rapid globalization of innovation poses many challenges to economists. How does the globalization of innovation affect our understanding of the economics of innovation? For example, the innovation system in many Western nations is characterized by a central role for university technology transfer offices in commercializing academic research, the prevalence of younger firms as strategic partners to and competitors with established players, and the challenges that many incumbents have faced in maintaining their initial innovative thrust. The extent to which these patterns will continue to hold in emerging economies is open to debate, and would reward close scrutiny.

A second area that deserves more scrutiny is the changing nature of incentives for innovators. Over much of the twentieth century, the structure of corporate research efforts, with their academic-type laboratories and weakly powered incentives for researchers, were largely static. However, the organization of research has seen a sharp transformation in recent years. Companies are increasingly relying on strategic alliances and other types of collaborations, and are increasingly proactive in aligning their internal research activities with the innovation system in which they reside. More strikingly, both private and public sector efforts have started to focus on relatively unfamiliar approaches, such as the widespread use of prizes, the proliferation of corporate venture schemes to facilitate spin-outs (and spin-ins), and attempts to harness the creativity and ideas of users and consumers. As was noted at various points in the volume, there seems to have been a significant increase (or at least an increasing awareness) of the roles that volunteers, freelancers, and users play in the innovation process. While a number of chapters in the volume touch on the changing nature of innovation incentives, only Boudreau and Lakhani directly address these issues (and probably not accidentally, do so using a quite novel methodological approach). The study of the subtle ways in which incen-

tives matter for innovation will provide grist for research for the foreseeable future.

A third area is an old—but ongoing—one: the appropriate measurement of the consequences of innovation. As discussed earlier, the measurement of the social welfare consequences of innovation poses some daunting challenges, which defy easy solutions. But even more modest goals, such as accounting for the impact of innovative products in national accounts and price indexes, remain problematic.

It might be surprising that these issues remain problematic, given that economists have been thinking about them since the work of Kuznets, Schmookler, Abramowitz, Griliches, and Solow in the 1950s. Given the central role of innovation and technological change in long-term economic growth, it is perhaps surprising that so few of the chapters in this volume directly examine the welfare implications of innovation, and none of the empirical chapters undertake a detailed welfare analysis. At one level, this absence underscores the intellectual history of the conference and the participants, including the microeconomic and phenomenological orientation of the 1962 volume. At a deeper level, however, it highlights a challenge that was raised by Paul David, Ken Arrow, and Dick Nelson in their panel commentaries. The presumed benefits arising from innovation are indeed not only hard to measure, but are in many cases difficult to conceptualize. For example, while there has undoubtedly been progress in the ability to measure the rate of commercialization of particular types of technologies (e.g., university disclosed inventions), does an increase in this rate imply an increase in social welfare? As Acemoglu pointed out in the conference discussion, there are general equilibrium effects that can often be as important as the main effects when undertaking such welfare calculations, and therefore a great deal of caution should be applied.

Moreover, the nature of technological change—most dramatically, the growing importance of the Internet, particularly the set of applications often referred to as "Web 2.0"—has highlighted the limitations of earlier approaches. Perhaps the most dramatic limitation has been the inability of economic frameworks to account for activities that are free: people around the world are spending more time on blogs, Facebook, and YouTube, and consuming less of many traditional media. And while economists can account for the loss of revenue that newspapers have experienced or declining prime-time television advertising rates, the benefits of these alternative activities resist ready quantification. Building better tools for assessing innovations that are systemic in nature is an ongoing challenge.

The chapters collected in this volume are of necessity limited in scope, as is this survey of the broader territory. One conclusion, though, is inescapable: the study of the rate and direction of inventive activity remains highly vibrant, and is likely to reward scholars from multiple perspectives in the years to come.

References

Aghion, P., and J. Tirole. 1994. "On the Management of Innovation." *Quarterly Journal of Economics* 109: 1185–207.

Arrow, K. J. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity,* edited by R. Nelson, 609–26. Princeton, NJ: Princeton University Press.

Holmstrom, Bengt. 1989. "Agency Costs and Innovation." *Journal of Economic Behavior and Organization* 12 (3): 305–27.

Kortum, S., and J. Lerner. 2000. "Assessing the Contribution of Venture Capital to Innovation." *RAND Journal of Economics* 31: 674–92.

Nelson, R. R. 1959. "The Simple Economics of Basic Scientific Research." *Journal of Political Economy* 49: 279–306.

Waldinger, Fabian. 2009. "Peer Effects in Science—Evidence From the Dismissal of Scientists in Nazi Germany." Centre for Economic Performance. Discussion Paper no. 910. London School of Economics.