14

Technological Innovation and Organizations

Pierre Azoulay and Josh Lerner

1. Introduction

In a handbook that takes a broad view of organizational structure, it might be wondered whether a chapter devoted to the relationship between technological innovation and organization is appropriate. After all, is not innovation just one relatively modest facet of economic activity?

There are two responses to this question. First, since the 1950s, economists have understood that technological innovation is critical to economic growth. Our lives are more comfortable and longer than those of our great-grandparents on many dimensions. At the heart of these changes has been the progress of technology. Innumerable studies have documented the strong connection between technological progress and economic prosperity both across nations and over time. Since the work of Abramowitz (1956) and Solow (1957), the importance of technological change has been generally understood, an understanding that has been deepened by studies in recent years documenting numerous positive effects of technological progress in specific areas, such as information technology (Bresnahan et al. 2002).

Second, innovation represents a particularly extreme ground for understanding organizational economics. This setting is one where information and incentive problems, which exist in the backdrop of many models of organizational structure and effectiveness, are front and center. These problems are at the heart of the innovation process. As Holmström (1989: 307) observes in his classic essay:

> One would expect modern finance theory to give good general advice on how to manage investments into research and development. But a quick look at finance textbooks reveals answers that are based on a very stylized conception of the problem and rather less illuminating than one would hope.

We thank Harvard Business School's Division of Research for financial support. Bob Gibbons, John Roberts, and participants in a meeting of the NBER Working Group in Organizational Economics devoted to reviewing early drafts of chapters from this handbook provided helpful comments. All errors and omissions are our own.

Thus, if we can understand organizational issues in the innovation setting—where performance often cannot be accurately measured, where knowledge often flows between entities, and where contracts may not be enforceable—we are likely to have a better understanding of these questions more generally. In this sense, the study by an economist of the relationship between innovation and organizational structure is akin to that of a physicist studying reactions in extremely young stars to gain insights that will be useful here on earth.

In this chapter, we begin by examining the relationship among internal organization, innovation inputs, and research productivity. We then turn to the relationship between contracting and innovation, and finally to the consequences of innovation for organizations.

Technological innovation is a large topic with many empirical and theoretical challenges. Nonetheless, we can highlight several conclusions that emerge from our discussion:

- Our understanding of contracting and innovation is probably the most developed of any of the major topics delineated here. At least in part, this reflects the close connections between theoretical and empirical research on this topic.
- In far too many other topics, however, there has been a disconnect between the insights of theoreticians and the work by empirical researchers.
- Most of what economists know about technological innovation and organizations does not stem from the mining of traditional datasets. On the contrary, some of the more lasting insights come from the collection of original fine-grained data at the project level. Case studies also inform the state of economists' knowledge on these topics (see Baker and Gil, this volume).
- An essential difficulty facing large-sample empirical research has been an inability to distinguish between association and causation, and, in some cases, a failure even to think carefully about this distinction. Empirical research is likely to advance considerably if some of the experimental methodologies employed in labor economics and related fields can be adopted here.
- Many valuable insights into these questions can be gleaned from fields outside economics, including sociology and more traditional managerial studies.
- The greatest gains in future years will be achieved by looking inside the "black box" of the firm and understanding the internal workings of the innovation process.

Reflecting the substantial size of this topic, we have limited our discussion in several important ways. First, we have focused on technological, not organizational, innovations. Readers interested in the adoption of new management techniques and their consequences for firm performance could turn to the work of Nick Bloom, John Van Reenen, and their collaborators (Bloom and Van Reenen [2010] provide a concise introduction to this fascinating topic). We have also focused only on research concerning the development of new ideas and approaches, and not on the diffusion of those ideas. Finally, we have ignored industry-level perspectives, such as the study of shake-outs and the industry life cycle, which have been the subject of other reviews (Klepper 2008).

2. Research and Development Productivity and Internal Organization

What makes an organization more or less innovative? For a long time, the literature on this topic focused solely on the relationship between innovation and firm size, following Schumpeter's (1942) conjecture that large firms had an inherent advantage at innovative functions relative to smaller enterprises. From the vantage point of twenty-first–century economists, the Schumpeterian hypothesis has not stood the test of time. It turns out to have been the intellectual by-product of an era that saw large firms and their industrial laboratories (e.g., IBM, AT&T, and DuPont) replace the independent inventors who accounted for a large part of innovative activity in the late nineteenth century and early twentieth (Lamoreaux and Sokoloff 2009).

At an anecdotal level, the Schumpeterian hypothesis does not accord with casual empiricism, as in several new industries (medical devices, communication technologies, semiconductors, software), industry leadership is firmly in the hands of relatively young firms whose growth was largely financed by public equity markets (e.g., Boston Scientific, Cisco, Intel, and Microsoft). But more systematic tests of Schumpeter's argument have been inconclusive too. A vast empirical literature, reviewed by Cohen and Levin (1989), has failed to uncover a robust relationship between firm size and innovation inputs.

What explains the relative failure of the empirical program started by Schumpeter's mid-century prophecy? An important reason is institutional change in the form of the emergence of vibrant markets for ideas that heralded a new division of innovative labor between large and small firms (Teece 1986; Gans and Stern 2000; Arora et al. 2002; Gans et al. 2002). If would-be innovators have the ability to cooperate with incumbent firms to appropriate the returns from technological innovation, regressing research and development (R&D) expenditures on firm sales will not be very illuminating.

Long before the emergence of modern ideas markets, some of Schumpeter's contemporaries, however, had some misgivings about the ability of large firms to generate technological innovations. For instance, Jewkes (1958: 55) argued:

> It is erroneous to suppose that those techniques of large-scale operation and administration which have produced such remarkable results in some branches of industrial manufacture can be applied with equal success to efforts to foster new ideas. The two kinds of organization are subject to quite different laws. In the one case the aim is to achieve smooth, routine, and faultless repetition, in the other to break through the bonds of routine and of accepted ideas. So that large research organizations can perhaps more easily become self-stultifying than any other type of large organization, since in a measure they are trying to organize what is least organizable. The director of a large research institution is confronted with what is perhaps the most subtle task to be found in the whole field of administration.

This section is devoted to fleshing out Jewkes's intuition by reviewing work by economists and other social scientists studying how other organizational characteristics impinge on the process of technological innovation. For organizing the discussion, we distinguish between studies of that focus on the *intensive margin* (the degree of effectiveness achieved by firms as they attempt to organize their innovative activities) from those that focus on the *extensive margin* (whether firms are able to innovate at all in the face of technological transitions).

2.1. Beyond Firm Size: Opening the Innovation Black Box

Perhaps the earliest indication that the internal organization of firms deserved more attention than it had traditionally received from economists came from empirical studies of the innovation production function in the often-cited National Bureau of Economics Research conference volume edited by Zvi Griliches (1984), *R&D, Patents, and Productivity*. Pakes and Griliches (1984) found that when estimating patent production functions, the magnitude of the coefficient on R&D investment fell drastically in the within-firm dimension of the data. Similarly, Scott (1984) found that fixed firm effects explained about 50% of the variance in R&D intensity. Of course, firm effects constitute a measure of economists' ignorance, and since 1984 a sizable body of research has emerged to explain the magnitude and relative stability over time of these firm effects. The literature has focused on three broad features of organizations: (1) the design of incentive systems; (2) firms' abilities to manage spillovers of knowledge, with a particular emphasis on the causes and consequences of job mobility among engineers and scientists; and (3) firms' choice of organizational structure, including (but not limited to) acquisition behavior.

2.2. Motivating Innovation: The Role of Incentive Systems

Firms provide incentives to their employees in all realms of economic life. Why would the design of incentive systems for innovative tasks differ from that appropriate for "humdrum" tasks? Holmström (1989) provides a number of reasons that make the provision of incentives for innovation a difficult task. First, innovation projects are risky and unpredictable; second, they are long-term and multistage; third, it might not be clear ex ante what is the correct action for the agent to take; and finally, they tend to be idiosyncratic and difficult to compare to other projects.

To this list, we would add two additional characteristics: innovators tend to bring to their labor a certain degree of intrinsic motivation, and innovation is very often a team activity.

Manso (2011) is the first economist to formalize the trade-off between the exploration of untested actions and the exploitation of well-known actions, which has long been a focus of organization theorists (March 1991). He uses a class of decision problems called "bandit problems," in which the agent does not know the true distribution of payoffs of the available actions. Exploration of new untested actions reveals information about potentially superior actions but is also likely to waste time with inferior actions. Exploitation of well-known actions ensures reasonable payoffs but may prevent the discovery of superior actions. Embedding the bandit problem into a traditional principal-agent model, Manso focuses on the features of incentive schemes that encourage exploration. He finds striking departures from the standard pay-for-performance contracts that are optimal when the principal is focused on eliciting effort for known actions. The principal should tolerate early failure and reward long-term success; she should provide some job security to the agent and should provide feedback on performance to the agent.

Ederer and Manso (2011) devise a laboratory experiment in which subjects are randomly assigned to incentive plans, including (1) a flat wage, (2) a standard pay for performance contract, and (3) a pay for future performance contract. They find that forward-looking incentives result in more exploration and higher profits. These results suggest that appropriately designed incentives do not lead to a crowding-out effect, contrary to earlier arguments that incentives are likely to undermine creativity (e.g., Amabile 1996).

To test the potency of exploration incentives to stimulate innovative activities in a real-world setting, one of the main difficulties is to find agents who are at risk of receiving either exploration- or exploitation-type incentives. Azoulay et al. (2011) focus on alternative funding mechanisms for elite biomedical scientists that appear to provide just this kind of variation. The lifeblood of academic life sciences in the United States is provided by so-called R01 grants awarded by the National Institutes of Health. These grants are relatively short term (the modal length is 3 years), and renewal is both competitive and not forgiving of failure. In addition, brilliant scientists can be appointed as investigators of the Howard Hughes Medical Institute. HHMI investigators continue to be housed in the same institutions but labor under a very different set of incentives. They are explicitly told to "change their fields" and are given the resources, time (5-year renewable appointments, with a lax first review and a 2-year phase down in case of termination), and autonomy (they do not need prior authorization to change the direction of their research) to accomplish this.

Compared with a set of equally eminent scientists at baseline, the authors find that the program appears effective in boosting the rate of production of particularly creative scientific papers—those in the top percentile of the citation distribution. The impacts are even larger for other outcomes, such as the production of trainees who go on to win early-career prizes or election to the National Academy of Sciences.

In the absence of random assignment, the conclusions must remain tempered, despite the care taken to match control scientists with treated scientists, based on observable characteristics. But taken as a whole, these two papers certainly suggest that tolerance for early failure and rewards for good performance over the long term are effective ways to motivate agents engaged in idiosyncratic, nonrepetitive tasks. Of course, many questions remain to be addressed. For example, how does one design incentives to attract the most creative agents? Also, in a setting with multiple agents, what is the optimal balance between individual and team incentives to motivate exploration? In particular, can exploration-type incentives be effective when the principal must rank the agents she oversees to determine a prize, such as in a promotion tournament? These are important questions for future research.

Another generic theme is that of incentive balance, when several tasks compete for the agent's attention, but the principal's ability to infer effort from output is much higher for one set of tasks relative to another set. The main result from this literature is that it may be optimal to provide low-powered incentives in such situations, to avoid distorting the agent's allocation of effort toward the task that is easier to meter. Cockburn et al. (1999) illustrate the implications of this principle in the setting of drug discovery R&D. Pharmaceutical firms would like their researchers to generate many useful patents, but if the scientists are not connected to sources of knowledge located outside the firm, their creativity will eventually run dry. Conversely, if the scientists solely focus on staying on the frontier of science, the firm will find itself running a (possibly very good) biology department but will have little tangible output (actual new drugs) to show for its R&D investment. Cockburn et al. show that pharmaceutical firms resolve this tension by keeping incentives for basic and applied output in balance. On the one hand, scientists will be rewarded and promoted based on their individual standing in their scientific subfield, using traditional academic criteria, such as publications and citations. On the other hand, research teams will be rewarded on the basis of their applied output. In practical terms, a group that generates more useful patents than expected in a given year will see its budget increase the following year. Cockburn et al. (1999) find that firms tend to adopt both or neither of these practices, rather than adopting one without the other.

Innovation settings differ from more traditional ones because employee-technologists often work on something because they find it personally rewarding. Intrinsic motivation is important, because it has implications for job design and wage setting. In a fascinating survey of life scientists on the job market, Stern (2004) finds that pharmaceutical firms who give their scientists freedom to choose their own projects and otherwise remain connected to open science pay 15% less on average, holding scientist ability constant. This finding provides a rationale for firms to harness their employees' preferences in the way they organize research activities, in addition to the crowding-out effect of extrinsic incentives often stressed by psychologists (Amabile 1993).

Of course, firms need to balance the wage savings associated with employee autonomy against the probability that some of the knowledge generated internally will leak out to their competitors. As it turns out, the management of knowledge spillovers from innovative activities is a topic that has also received considerable attention.

2.3. Spillovers of Knowledge and Internal Organization

In his famous 1962 essay, Kenneth Arrow focused economists' attention on the nonrival nature of knowledge and the attendant disclosure problem. Ever since, economists have been focused on firms' ability to actively manage knowledge spillovers. The starting point for the discussion is that knowledge is not as public a good as initially thought, and therefore, distance between the transmitter and the receiver of knowledge matters in determining the extent of spillovers. Jaffe et al. (1993) observed that despite their ethereal nature, knowledge spillovers might leave a paper trail in the form of citations to prior art recorded in patents. By constructing a large dataset of patents and matching the location of inventors for both cited and citing patents, they documented that these citation patterns exhibit a pronounced degree of localization.[1]

However, physical proximity is not necessarily the only, or even the most relevant, concept of distance that economists should consider when attempting to estimate the magnitude of knowledge spillovers. To the three dimensions of physical space, it seems worthwhile to add technological space: advances in the state of knowledge in one particular technological area are likely to spur relatively more developments in technological areas that are thematically related. Jaffe (1986) was the first to construct a measure of technological distance between firms by using the distribution of firms' innovative efforts across patent classes, and he found that the R&D productivity of a focal firm was indeed positively correlated with the R&D of technological neighbors.

Social distance might also impinge on the extent of spillovers. Here the network metaphor is particularly apt, whether one thinks about degrees of separation between scientists in a

1. As documented by Thompson and Fox-Kean (2005), this exercise is fraught with difficulties because of the difficulty of finding nonciting patents that can serve as a control group and capture the agglomeration patterns that are due to other factors than knowledge spillovers.

coauthorship network, of the relationship between a corporate parent and its subsidiaries, or about the vertical linkages between buyers and suppliers. Adams and Jaffe (1996), focusing on the chemical industry, showed that the effects of parent firm R&D on plant-level productivity decrease with both geographic and technological distance between the research lab and the plants. Moreover, spillovers from technologically related firms are significant in magnitude but are diluted by the firms' own R&D intensity. This article stands out in the vast empirical literature, because it simultaneously attends to the three concepts of distance mentioned above.

In a similar vein, Azoulay et al. (2010) estimate the magnitude of spillovers generated by the untimely death of 112 academic superstars in the life sciences on their coauthors' research productivity. They find a lasting and significant decline in publication output, which is surprisingly homogeneous across a wide range of coauthor characteristics. But they do not observe a differential effect of a prominent collaborator's death for co-located coauthors. They interpret their findings as providing evidence that part of the scientific field embodied in the "invisible college" of coauthors working in a particular area dies along with the star.

A second fundamental idea in the study of knowledge spillovers is that of *absorptive capacity:* absorbing spillovers from other firms requires doing research yourself (Cohen and Levinthal 1989, 1990). Using a dataset that combines the Federal Trade Commission's line of business information with Levin et al. (1987) survey data, these authors find that spillovers from input suppliers can be absorbed with less R&D investment than spillovers from government and university labs. In the setting of pharmaceutical drug discovery, Cockburn and Henderson (1998) document strong correlations between the extent of ties between firms and academic science (mostly through coauthorship of scientific articles) and these same firms' research productivity. Moreover, building absorptive capacity has implications for the organization of R&D activity. Henderson and Cockburn (1994) use detailed data at the research program level from ten pharmaceutical firms and show that firms that adopted the practices of open science to motivate and reward scientists increased their research productivity (as measured by important patents per dollars invested) significantly during the 1980s.

A more recent literature stresses that there are other ways—beyond performing basic R&D—for firms to absorb outside knowledge. In a fascinating case study of the semiconductor industry, Lim (2011) shows that IBM's competitors were able to quickly imitate its design of copper interconnects by working closely with equipment suppliers, collaborating with some key academic labs, and hiring newly minted PhDs and postdoctoral fellows in the relevant scientific subfields.

2.4. The Mobility of Engineers and Human Capital Externalities

In her seminal study of the computer and semiconductor cluster in Silicon Valley, Annalee Saxenian (1994) drew the attention of economists to the high rates of job hopping among engineers. She ascribed the sustained success of Silicon Valley firms to the rapid movement of technical professionals among firms in the region and pointed out that an accident of Californian legal history precluded the enforcement of noncompete agreements in the state. Since then, an ever-expanding empirical literature has documented more systematically that job hopping is indeed a source of knowledge spillovers, and that these spillovers tend to be geographically localized (Almeida and Kogut 1999; Fallick et al. 2006).

What is less clear from this evidence, however, is whether these spillovers correspond to true externalities or whether there exist mechanisms through which the labor market can internalize the effects. Pakes and Nitzan (1983) present a model in which engineers bond themselves to their firms by accepting lower entry wages. Moen (2005), using matched employee/employer data from Norway, compares the wage-experience profiles of engineers and nontechnical white collar workers and finds that they are steeper for the former group of employees than for the latter. This suggests that the potential externalities associated with labor mobility are partially internalized in the labor market.

Our knowledge of the practices used by firms to manage outgoing and incoming knowledge spillovers is still fragmentary. The best-known evidence probably pertains to the role of scientific superstars documented by Zucker et al. (1998). These authors document a robust correlation between the rates of founding of new biotechnology firms and proximity to scientists who are leaders in the relevant subfields of biology. In the same spirit, Lacetera et al. (2004) show that pharmaceutical firms are more likely to change the organization of their drug discovery operations toward the open science model after hiring academic superstars. In this line of research, the direction of causality has not been clearly established, for a number of reasons. First, high-quality data on employee mobility is hard to come by. A popular approach has been to rely on patent or publication data to ascertain the movement of employees across firms, but it is obviously problematic to base mobility on potentially endogenous output measures (Almeida and Kogut 1999; Rosenkopf and Almeida 2003). Second, even if mobility could be measured with less error, there is always the possibility that scientists' or engineers' movements in and out of firms are driven by expectations regarding future firm productivity. A productive avenue for mobility research is to gather much more detailed data on individuals to extract exogenous variation in exposure to talent. Life cycle events, such as marriage, birth of children, and becoming "empty nesters" might provide such variation. An alternative is to focus instead on policy changes that affect rates of mobility among technical personnel. Marx et al. (2009) exploit Michigan's inadvertent reversal of its noncompete enforcement legislation to demonstrate that noncompetes decrease inventor mobility by 34%. Although this effect is not by itself surprising, the study's design can be thought of as providing a plausible first stage to study the effect of engineers' mobility on productivity.

2.5. Organizational Structure

In their efforts to broaden the range of organizational characteristics that impinge on the innovation process, economists and other social scientists have focused on various aspects of organizational structure. We focus here on three of the most salient: the choice between centralized and decentralized R&D activities, the effect of takeover and acquisitions on innovation, and the design and composition of R&D project teams.

2.5.1. Corporate R&D Laboratories versus Decentralized Innovation

Managerial scholars have long debated the relative merits of centralized R&D labs and decentralized R&D activities (Rosenbloom and Spencer 1996). Centralized R&D labs can potentially engage in nonlocal search activities and focus on long-term projects, but they run the risk of slowly evolving into ivory towers. Decentralizing R&D investment into divisions prevents employee-technologists from losing sight of market imperatives, but it might result in a slow degradation of technological and scientific competencies. Some recent studies appear to support the view that centralized R&D more effectively supports wide-ranging innovative activities.

Using a sample of 71 large research-intensive corporations, Argyres and Silverman (2004) show that firms with centralized R&D labs generate more highly cited patents than do firms with decentralized or hybrid R&D structures. Lerner and Wulf (2007) combine information about structure with data on incentives provided to corporate R&D managers among Fortune 500 firms. They find that long-term incentives (stock options and restricted stock) result in patents than are more heavily cited, but this relationship is driven solely by firms with centralized R&D.

Together, these studies suggest that aligning the incentives of corporate R&D staff with those of the firm as a whole can mitigate the risk that research output loses relevance, with important caveats. First, patent citations are far from an ideal measure of the relevance of the research to the firm's product markets; sales from innovative products would be a more appropriate metric. Second, the choice of structure is potentially endogenous, and the studies provide suggestive conditional correlations, not estimates of causal effects. The relative virtues of centralized and decentralized R&D structures remain a relatively open question.

2.5.2. Takeovers, Mergers, and Acquisitions: Much Ado about Nothing?

In the 1970s and 1980s, the rise of Japan as an industrial giant stirred a debate in the United States and Europe about the effects of takeovers and merger activity on the rate and direction of R&D undertaken by firms. This prompted several studies probing the link between corporate restructurings and the intensity of R&D among U.S. manufacturing firms. Although the zeitgeist that provided the impetus for these studies has changed radically in the past 20 years, two papers by Bronwyn Hall stand out as particularly informative (Hall 1990, 1991).

As usual, the devil lies in the details of the econometric exercise, but in both cases, the effects she uncovers are of a relatively small magnitude—certainly smaller than what the hyperbolic statements of policymakers at the time would have implied. Hall finds that public firms involved in acquisitions where both partners are in the manufacturing sector have roughly the same pattern of R&D spending as the sector as a whole, and that the acquisition itself does not cause a reduction in R&D activity. Moreover, the target's R&D capital seems to be valued more highly by the acquirer than by the stock market itself (Hall 1991). In a second paper, Hall also studies the effects of leveraged buyouts. Here, she finds a more pronounced negative effect of increased leverage on R&D intensity, but because the firms taken over as a prelude to leveraged buyouts tend to be in less R&D-intensive segments, it seems unlikely that the leveraged buyout wave of the 1980s and early 1990s had a significant negative effect on R&D expenditures (Hall 1990).

2.5.3. Performance of R&D Project Teams: The Role of Social Structure

Employee-technologists are very often organized in project teams. Economists have long studied the provision of incentives in teams, including peer pressure (Kandel and Lazear 1992) and collective rewards (Che and Yoo 2001). Here we draw attention to a distinct attribute of R&D teams that has been shown to play a role in explaining their performance: the structure of the social networks to which team members belong.

Economists have contributed very little to this vibrant area of social science research. In light of the relevance of the results to the rate of technological innovation inside organizations, we chose to include a bare-bones exposition of the key idea, partly in the hope of stimulating economists' interest in the topic. Sociologists characterize social networks as cohesive or closed when many redundant contacts exist among network members. In a closed network, information flows quickly from one individual to another. In contrast, social networks can be rich in structural holes, that is, when some individuals can broker ties between subgroups that would have otherwise no opportunity for communication (Burt 1992). A central trade-off arises because network closure among team members makes coordination easier—which is beneficial in the execution phase of a project. At the same time, assembling in a team individuals whose social worlds do not overlap provides learning benefits, because it ensures that team members will be exposed to diverse sources of information (Burt 2004). In a fascinating paper, Reagans et al. (2004) provide evidence from a contract research organization that sheds light on how organizations can resolve this tension. Using the fact that employees participate in several distinct teams over the period during which they collected the data, these authors show that the most effective project teams have both high internal density (i.e., strong ties exist among team members) and high external range (i.e., team members have ongoing relationships with employees belonging to a diverse set of other project teams).

Of course, an economist might be skeptical of the empirical results for a variety of standard reasons. We might rightly wonder about simultaneity or omitted variables, though the authors have designed the study while attending to these concerns. More importantly, the outcome variable—time to project completion—is not ideal, especially as other research has found that employees in brokerage positions extract higher wages from their employers. Nonetheless, the effects the authors uncover are important in magnitude and statistically significant. Given the explanatory power of social network theories, it would seem that economists ignore these findings to their own detriment.

2.6. The Extensive Margin: When Established Firms Fail to Innovate at All

Rather than ask whether large firms invest more in R&D relative to small ones, another line of research focuses on evaluating theoretically and empirically the costs and benefits of market incumbency.

One branch of the literature provides an explanation for the dominance of entrepreneurial firms in new industries. Aron and Lazear (1990) present a model in which new firms pursue high-variance strategies and hence are more likely to introduce new products. Prusa and Schmitz (1994) test their argument by examining the introduction of new software programs. The authors suggest that new firms appear to be more effective at creating new software categories, whereas established firms have a comparative advantage in extending existing product lines. It should be noted that models like that of Aron and Lazear pertain to innovation generally, not necessarily to technological innovation.

A distinct stream of research explores the conditions under which market incumbents can adapt in the face of technological change. Here there has been a closer connection between theoretical work and empirical studies. The theoretical debate is well known and has been reviewed elsewhere (e.g., Reinganum 1989). In brief, models of technology races make sharply different predictions about the innovation incentives of incumbents and entrants depending on (1) the nature of uncertainty (i.e., can investment shift the timing of the new product's introduction?) and (2) whether the innovation is radical or incremental (i.e., is the new product so obviously superior that demand for the old product disappears?) (Gilbert and Newberry 1982; Reinganum 1983).

Because these features are hard for researchers to measure, empirical work has attempted to test the implications of racing models rather than compare the models directly. Lerner (1997), in his empirical examination of the disk drive industry, is a good example of this approach. For each product generation, he finds that firms who followed the leader in the previous generation display the greatest propensity to innovate. Interestingly, the same pattern does not exist in situations where strategic incentives are muted, as is the case when the firm is a division within a vertically integrated firm. Lerner interprets these findings as providing support for Reinganum's model of technological racing, but it is not clear that innovation in the disk drive industry should be construed as radical. In particular, the old and new generations of products tend to coexist for a period of time, because they typically appeal to distinct customer segments.

Christensen (1997) provides a different interpretation of the same phenomenon. According to his analysis, the new product generation initially poorly serves the needs of the incumbent's customers, but its quality eventually catches up. This presents an ideal situation for entrants to exploit incumbents' blind spots, which stem from their single-minded focus on existing customers. Christensen goes on to recommend that established firms systematically create "skunkworks" to incubate their next generation of products, far from the paralyzing influence of the old business.

Economists will naturally be skeptical of such a blanket recommendation. The disk drive industry is characterized by a weak appropriability regime (patents can be invented around), and it would be hazardous for an entrant to try to license an innovation to an established firm. The conclusions may be therefore highly contingent on particular industry characteristics. Nonetheless, the reasons behind the apparent bias of incumbents in favor of home-grown technology is a recurrent theme of the managerial literature, and its persistence and effects are worthy of attention by economists.

Other strands of the managerial literature put forth a conception of radical and incremental innovations based on supply-side considerations: A radical innovation is one that destroys, or at least does not build on, the technological capabilities of incumbent firms (Tushman and Anderson 1986). In her study of the photolithographic industry, Henderson (1993) notes that it is entirely possible for innovations to be incremental in an economic sense but radical in an organizational sense, and she provides evidence that incumbents have particular difficulty adapting to such architectural changes. Using qualitative evidence, she explores the mechanisms that might explain the incumbents' lack of success in commercializing architectural innovations. She advances two types of explanations. The first focuses on dysfunctional incentives and stilted communication channels; the second emphasizes behavioral biases that affect the managers of incumbent firms. One lesson from recent advances in behavioral corporate finance is that cognitive biases can magnify the salience of agency problems in firms (Baker et al. 2007). The study of internal capital markets—especially in the presence of informed but biased R&D managers—seems a very promising area for future research. So far, the growing literature on internal capital markets (reviewed by Gertner and Scharfstein, this volume) has been largely focused on resource allocation in diversified publicly listed conglomerates. Although attention has increasingly been paid to the role of divisional boundaries on R&D investment and productivity (e.g., Seru 2010), further progress in this area will probably require access to finer-grained, project-level data.

Before moving on to other issues, we can venture a few conclusions. First, most of what economists know about technological innovation and organizations does not stem from the mining of traditional datasets. On the contrary, some of the more lasting insights come from the collection of original data at a level that is usually finer grained than the whole organization. Case studies also inform the state of economists' knowledge on these topics. Second, the empirical literature has been rather casual in its treatment of endogeneity issues, though there are of course exceptions (e.g., Stern 2004). What is missing are more systematic attempts to estimate the causal effect of various features of firms' internal organization on innovative outcomes.

3. Contracting and Innovation

One of the most fertile areas of research in recent years has been the relationship between innovation and contracting. Thanks to a variety of theoretical work, we now better understand the critical importance of contracting in the innovation process and the ways in which contract structure can affect the innovations being undertaken.

3.1. Why Is Contracting Important for Innovation?

One of the dominant features of many high-technology industries has been a reliance on contracting. In many cases, start-up firms, rather than going head-to-head with established incumbents, will choose to license technologies to their peers. Such industries as telecommunications and biotechnology have been profoundly shaped by this business model.

It is natural to wonder what the rationale for this approach in technology-intensive industries is. Gans and Stern (2000) suggest that there are several reasons firms will choose a cooperative strategy, including avoidance of thinner profit margins and duplicative investments associated with head-to-head competition. They argue that a cooperative strategy is more likely for start-ups in three cases: when intellectual property rights are stronger (which puts the start-up firm in a better bargaining position with potential licensors), when intermediaries who can facilitate and reduce the costs of such transactions are present, and when established firms have made expensive investments (e.g., in a sales force) that the start-up would need to duplicate. The authors show that these patterns not only hold theoretically but also in a survey of the commercialization strategies of more than 100 start-up firms.

This rationale for cooperative commercialization is compelling, but efforts to contract on technology often encounter many challenges. In many instances, academic technology managers or corporate executives seeking to commercialize early-stage technologies have encountered real difficulties. They may find investors unwilling to invest the time and resources to examine early-stage technologies, or offering only modest payments in exchange for large stakes in innovations that the scientists, technology transfer officers, and company executives believe to be quite valuable.

Much of this reluctance stems from the information problems that surround innovations, particularly in technologically advanced industries, which often scare off potential investors. The work of Guedj and Scharfstein (2004) highlights the importance of these information problems. The paper documents the extent of agency problems among young biotechnology firms, which emerges from a comparison with established pharmaceutical companies. The authors show that the young firms—particularly those with large cash reserves—are more likely than their more established peers to push drug candidates forward in clinical trials, but the success rate in these trials is much lower. The evidence points to an agency problem, in which managers of single-product early-stage firms are unwilling to drop the development of their only viable drug candidates, and the difficulty that outside shareholders have in monitoring innovative activities.

The consequences of these difficulties have been illustrated by Shane (1995) and Majewski (1998), who examine the decisions of biotechnology firms to raise capital through public markets and alliances. Through their different methodologies, these authors show that firms turn to alliance financing when asymmetric information about the biotechnology industry is particularly high. During these periods—which are measured through such proxies as the variance of the returns of biotechnology securities—firms are likely to delay the time until their next equity issuance and to rely on alliances rather public offerings as a source of external financing. The authors argue that the greater insight on the part of the pharmaceutical company into the nature of the biotechnology firm's activities allows it to make successful investments at times when uninformed public investors are deterred by information problems.

3.2. Contract Structure and Innovation

Much academic interest in recent years has surrounded the question of how firms contract for innovations and the implications of these contracts. The seminal work of Aghion and Tirole (1994) has inspired a variety of investigations.

This work builds on the tradition, beginning with Grossman and Hart (1986) and Hart and Moore (1988), that depicts incomplete contracting between a principal and an agent. A typical assumption is that it is impossible for the two parties to write a verifiable contract, enforceable in a court of law, that specifies the effort and final output of the two parties. This is because there are many possible contingencies, not all of which can be anticipated at the time the contract is drafted. Due to this nonverifiability problem, these models suggest that it is optimal for ownership of the project to be assigned to the party with the greatest marginal ability to affect the outcome. This party, who will retain the right to make the decisions that cannot be specified in the contract ex ante, should also receive any surplus that results from the project. Because of this incentive, the party will make the decisions that maximize—or come close to maximizing—the returns from the project.

Aghion and Tirole (1994) adapt this general model to an R&D alliance between two firms. In their basic model, the authors assume that the research unit is without financial resources of its own, cannot borrow any funds, and has no ability to commercialize the innovation itself. As a result, it turns for financing to a customer, a firm that may intend to use the product itself or to resell it to others but cannot make the discovery independently. The success of the research project is an increasing function, though at a decelerating rate, of both the effort provided by the research unit and the resources provided by the customer.

Developing a contract between the two parties is challenging. Even though the ownership of the product can be specified in an enforceable contract, and the resources provided by the customer can also be so specified, uncertainty precludes writing a contract for the delivery of a specific innovation. Similarly, an enforceable contract cannot be written that specifies the level of effort that the research unit will provide.

Aghion and Tirole (1994) consider two polar cases: when the research unit has the ex ante bargaining power and when the customer does. When the research unit has the bargaining power, the ownership of the research output will be efficiently allocated. If the marginal impact of the research unit's effort on the innovative output is greater than the marginal impact of the customer's investment, then the research unit will receive the property rights. If not, the research unit will transfer ownership to the customer in exchange for a cash payment. This result is similar to that of Grossman and Hart (1986).

When the customer has the bargaining power, a different pattern emerges. If it is optimal for the customer to own the project, it will retain the project. If, however, the total amount of value created would be greater were property rights to be allocated to the research unit, the ideal outcome will not be achieved, because the cash-constrained research unit will not have enough resources to compensate the customer. As a result, an inefficient allocation of the property rights occurs, with the customer retaining the rights to the invention.

This work has inspired a variety of empirical investigations. Lerner and Merges (1998) examine the determinants of control rights in a sample of 200 alliances between biotechnology firms and established pharmaceutical firms. This setting is particularly auspicious, because it exhibits wide variation in contractibility (e.g., early-stage versus late-stage projects) and bargaining power (e.g., does the start-up have other products in development or is it essentially a single-product firm?). They analyze the share of 25 key control rights allocated to the financing firm by regressing the assigned number of rights on independent variables denoting the project stage and financial conditions, as well as on controls for a variety of alternative explanations. The results are generally consistent with the framework developed by Aghion and Tirole (1994): the greater the financial resources of the R&D firm, the fewer control rights are allocated to the financing firm; in contrast, there is no evidence that early-stage projects are associated with more control rights being assigned to the research-based firm.

Lerner and Malmendier (2010) expand on this line of research. They point out that the key variable in Lerner and Merges (1998) is a bit of a grab bag. Among these 25 decision rights, which ones really matter? They also point out that Aghion and Tirole's (1994) canonical model should not be taken too literally. Rather, they ask what central incentive problem the proper assignment of these rights should alleviate. Based on insights gained from conversations with practitioners, they conjecture that pharmaceutical firms are especially worried that their biotechnology partners will inappropriately cross-subsidize other projects using the payments received as part of a specific drug development project. They go on to develop a simple model in the spirit of Aghion and Tirole but tailored to the specifics of the context at hand. This leads them, in turn, to focus on the allocation of a specific contractual clause: a termination right coupled with broad transfer of intellectual property and associated payments. The model makes some specific predictions: (1) such clauses should be more frequently observed when the direction of research is less contractible (e.g., when the alliance does not center on a "lead compound" with a specific chemical formula); and (2) this correlation should be much weaker, or even go away, when the biotechnology firm is not financially constrained (e.g., when the biotechnology firm has low net income). These sharp predictions are borne out in the empirical analysis and appear robust to a number of carefully examined alternative explanations.

Whereas Aghion and Tirole (1994) do not explicitly depict a role for the public market, Lerner et al. (2003) argue that variations in the availability of public financing will affect the bargaining power of R&D firms. During periods when public financial markets are readily accessible, these firms may be able to finance projects through either public equity issues or alliances. But during periods when equity issues are more difficult, R&D firms may have few alternatives to undertaking alliances. In the latter periods, it is also reasonable to assume that the R&D firm's bargaining power will be considerably reduced.[2]

The authors examine these patterns in biotechnology, noting that equity financing in the industry has undergone dramatic variations over the years. These shifts have been largely in the nature of industrywide shocks. The authors show that in periods when public equity financing is readily available, the agreements are more likely to grant key control rights to the R&D firm. This pattern—consistent with the theory of Aghion and Tirole (though, of course, with other theories as well)—holds even after controlling for variations in the quality of the technology in the agreement. Lerner et al. (2003) then examine whether the agreements are successful in terms of the progress of the product under development. Alliances that grant the bulk of the control to the R&D firm are more successful, an effect that is more pronounced in weak financing markets, as Aghion and Tirole (1994) predict.[3]

Finally, the authors examine the likelihood of renegotiation. If it would have maximized innovative output to assign control to the small biotechnology company, but this allocation of control was precluded by financial market conditions, there should then be a distinct pattern in renegotiations. In particular, when financing conditions for biotechnology firms improve, the agreements that assign the bulk of the control to the financing firm should be disproportionately renegotiated. The empirical results are consistent with this pattern.[4]

Robinson and Stuart (2007) take a different approach, examining 125 alliances involving genomics research at young biotechnology firms. The authors consider a variety of characteristics of these contracts, including the use of equity, extent of up-front payments, contractual provisions employed, provisions regarding termination, and length of the agreement. Although it is difficult to do justice to so many disparate analyses in a brief summary, the key findings are the resemblance between strategic alliances and venture capital contracts[5] and the importance of contractual provisions that are difficult to verify. The authors interpret the latter conclusion as

2. As noted above, the posited relationship between the strength of the public equity market and the bargaining power of the R&D firm is not explicitly modeled in the Aghion-Tirole model. Under the plausible assumption that in the subset of parameters where the R&D firm chooses an alliance instead of public equity issue, the strength of the market still affects the R&D firm's bargaining power in choosing the terms of the alliance, an extension to their model could deliver this result. The claim that control rights are more likely to be transferred to the financing firm during periods of diminished public equity market activity is also supported by theoretical work by Aghion and Bolton (1992) and Holmström and Tirole (1997).

3. Of course, if we assume that having the R&D firm control the alliance is inevitably efficient, then it is hard to distinguish the Aghion-Tirole hypothesis from one where the lack of financial resources leads the R&D firm to make important concessions. The results regarding the renegotiations of alliances, however, seem less consistent with this alternative hypothesis.

4. See Elfenbein and Lerner (2003) for another example of an analysis of these issues.

5. This point has also been raised in the work of legal scholars, such as Bernard Black, Ronald Gilson, and Robert Merges. For one example, see Merges (1995).

raising doubts about the substantial literature on incomplete contracts most closely associated with Grossman, Hart, and Moore.

3.3. Spin-offs and Firm Boundaries

A consequence of the ability of start-up firms to enter into contracts with established firms is that industry structure may change. The prospect of garnering profits through licensing—or alternatively, through successful head-to-head competition—may lead individuals to leave established firms and begin new concerns. A variety of work has examined these changes to industry structure and the consequences for innovation.

Much of this literature highlights the difficulties that firms face when managing multiple projects in a single organization. This literature has been largely motivated by the diversification discount puzzle: the empirical observation that diversified firms appear to trade at prices that are significantly below those of comparable portfolios of specialist firms.

Much of the literature has highlighted the presence of agency problems in firms. Jensen (1986), for instance, suggests that CEOs have a tendency to use the cash flows from the business to overinvest in unprofitable projects. Scharfstein and Stein (2000) argue that the presence of agency problems at both the CEO and divisional manager level may lead to firms overinvesting in weaker projects and neglecting stronger ones, particularly when there is a substantial divergence in the projects and when the CEO has low-powered incentives.[6]

Several works have explicitly examined the trade-off between undertaking an innovative project in a major corporation and in a start-up firm. Typically, they assume that an employee makes a valuable discovery while working for an established firm. These models usually assume that substantial informational asymmetries preclude the established firm from learning (at least initially) about the project's prospects. In some cases, the difficulties of contracting and bargaining lead to an independent firm being established. Wiggins (1995) predicts that projects that are high risk, take longer to develop, and are less capital intensive will be more likely to be pursued in independent firms. Although employing a somewhat different theoretical setup, Anton and Yao (1994) similarly conclude that more radical innovations will be more likely to be pursued in start-ups.[7]

These issues are captured nicely in a model by Hellmann and Perotti (2011), who explore the challenges that an inventor who has developed an incomplete idea faces. They argue that if the would-be entrepreneur can find a third party (e.g., a venture capitalist) to evaluate his idea, he will be able to validate his concept and decide whether to invest more resources into it. But the third party may be tempted to report to the inventor that the idea is worthless, and then steal the idea and commercialize it herself. They argue that the likelihood a third party will respond truthfully depends on the extent to which her skills are complementary with, or simply identical to, those of the inventor.

6. Although the work is not cast in this manner, the findings of Steven Klepper highlighting the importance of disagreements among individuals as a driver of new firm formation (e.g., Klepper and Thompson 2005) may be seen as consistent with these arguments.

7. Dix and Gandelman (2007) and Hellmann (2007) are two more recent examinations of these questions.

Alternatively, the entrepreneur may take the idea to his employer. Although this approach may not be as efficient in some respects (the firm may not have as much expertise as a third party), Hellmann and Perotti argue that firms can more readily develop a reputation for trustworthiness and honesty. They argue that these commitments are least likely to be honored when the invention is very valuable: in these instances, individual investors are more likely to try their luck in the treacherous open market.

Acemoglu et al. (2003) take a somewhat broader view of these questions. They depict a world where firms can engage in two activities: adopting existing technologies from the world's best innovators and developing innovations themselves. They argue that managerial skill is relatively unimportant in the adoption process but is critically important to successful innovation. As a result, in countries where technology is more advanced—and hence, the returns from imitation are lower—it becomes critically important to select the right managers. Thus, they suggest, as countries approach the technological frontier, technology-intensive industries should become increasingly Darwinian, with younger firms, shorter relationships among firms, and more weeding out of managers.

On the empirical side, Gompers et al. (2005) contrast two views of the spin-off process. In one view, individuals already working for entrepreneurial firms—particularly those already backed by venture capitalists and located in hotbeds of venture capital activity—may find launching their own venture less daunting than others might for a number of reasons. These reasons include: they have already been exposed to a network of suppliers of labor, goods, and capital, as well as a network of customers (Saxenian 1994); they have already learned by doing how to establish an entrepreneurial firm; or individuals with a higher taste for risky activities may have already found their way to entrepreneurial firms, consistent with models of the sorting processes (Jovanovic 1979; Holmes and Schmidt 1990; Gromb and Scharfstein 2002).

Alternatively, individuals become entrepreneurs because the large bureaucratic companies for which they work are reluctant to fund their entrepreneurial ideas. As discussed above, established firms may be incapable of responding to radical technological changes that upset the established ways of organizing their businesses, or these hierarchical organizations may have a hard time assessing such investment opportunities (Stein 2002; Berger et al. 2005). After systematically examining which publicly traded firms had employees depart to start new venture-backed firms, Berger et al. and Stein conclude the findings appear to be more consistent with the view that entrepreneurial learning and networks are critical factors in the creation of venture-backed firms.

Azoulay (2004) examines instead the question of how the innovative process affects firms' vertical integration decisions. Pharmaceutical companies have long outsourced at least part of the process by which new drugs are evaluated (known as clinical trials) to specialists in these procedures. He argues that these trials typically are dominated by one of two tasks, routine compilation of data on the drugs' effectiveness and more fundamental analyses of the pharmaceuticals' workings.

The analysis suggests that firms respond to this variation. They will choose to undertake the project in-house, or alternatively outsource it, depending on the nature of activities that will dominate the trial process. In the more cutting-edge projects involving the production of new knowledge, they are more likely to use their own employees. A key driver of this choice, Azoulay

(2004) argues, is the difficulty of measuring contractor performance when undertaking truly innovative activities: the difficulty of designing appropriate incentives for innovation leads to the choice of firm boundaries.

3.4. New Organizational Structure and Innovation

Another area of considerable research interest has been the growth of new organizational structures to promote innovation. These have many intriguing features, which will well reward ongoing research.

Perhaps the most intriguing of these topics is the open source process of production and innovation, which seems very unlike what most economists expect. Private firms usually pay their workers, direct and manage their efforts, and control the output and intellectual property thus created. In an open source project, however, a body of original material is made publicly available for others to use, under certain conditions. In many cases, anyone who makes use of the material must agree to make all enhancements to the original material available under these same conditions. This rule distinguishes open source production from, say, material in the public domain and "shareware." Many contributors to open source projects are unpaid. Indeed, contributions are made under licenses that often restrict the ability of contributors to make money on their own contributions. Open source projects are often loosely structured, with contributors free to pursue whatever area they feel most interesting. Despite these unusual features, recent years have seen a rise of major corporate and venture capital investments into open source projects.

Economics research into open source has focused on two issues. The first has been the motivation of contributors. Lerner and Tirole (2002) argue that the standard framework of labor economics can be adapted to capture activity in the open source environment. Even if there are no short-run monetary returns from working on open source projects, they argue that participation can have important signaling benefits in the long run. The paper highlights the importance of programmers' desire to signal their quality—that is, the desire to impress prospective employers and financiers, as well as obtain peer recognition—as a spur to contributing to open source projects. The presence of these signaling incentives will lead to more success for open source projects where contributions are more visible to the relevant audience (e.g., peers or employers) and where the talent of the contributor is better discerned from his or her contributions. These observations lead to a series of predictions about the likely success and structure of open source projects.

The empirical evidence, particularly the survey work of Hann et al. (2004), is largely consistent with the belief that individual contributors to open source projects do ultimately benefit financially from their participation in these projects. The results suggest that the sheer volume of contributions to the Apache project have little impact on salary. But individuals who attain high rank in the Apache organization enjoy wages that are 14–29% higher, regardless of whether their work directly involves the Apache program.

The second issue on which economists have focused concerns the legal rules under which open source projects operate. The licenses differ tremendously in the extent to which they enable licensors and contributors to profit from the code that is contributed.

Lerner and Tirole (2005) argue that permissive licenses, where the user retains the ability to use the code as she sees fit, will be more common in cases where projects have strong appeal to the community of open source contributors—for instance, when contributors stand to benefit considerably from signaling incentives or when the licensors are well trusted. Conversely, restrictive licenses, such as the General Public License, will be commonplace when such appeals are more fragile. Lerner and Tirole also examine the licenses chosen in 40,000 open source projects. The authors find that, consistent with theory, restrictive licenses are more common for applications geared toward end users and system administrators. Similarly, projects whose natural language is not English—whose community appeal may be presumed to be much smaller—are more likely to employ restrictive licenses.

But many issues posed by open source are not unique to this setting. Open source can be seen as at the end of a spectrum of technology-sharing institutions. Many of these other institutions have encountered similar conflicts. Leaders of patent pools, for instance, have had to deal with the conflicting goals of the potential members, sometimes reconciling their disparate goals successfully and in other cases failing to overcome these gaps. To cite another example, the challenges of enlisting cooperation from commercial firms while guarding against opportunistic behavior are familiar to leaders of standard-setting organizations. These institutions have lengthy and well-documented track records; for example, the first patent pool dates back to the 1850s.

The modern study of patent pools has its origins in the work of Shapiro (2001), who uses Cournot's (1838) analysis to point out that patent pools raise welfare when patents are perfect complements and harm it when they are perfect substitutes. Although this observation is a useful first step in the antitrust analysis of patent pools, patents are rarely perfect complements or perfect substitutes. Indeed, antitrust authorities sometimes are unsure about the relationships among the patents. Moreover, frequently observed features of pools, such as provisions demanding independent licensing, cannot be analyzed in a setting where only the polar cases exist.

Lerner and Tirole (2004) build on this work, analyzing the strategic incentives to form a pool in the presence of current and future innovations that either compete with or are complementary to the patents in the pool. The authors begin with a very stylized model (though one that allows the full range between the two polar cases of perfectly substitutable and perfectly complementary patents) and then consider progressively more realistic scenarios. A major focus of the analysis is the process through which competition authorities examine patent pools. Recent antitrust doctrines are that only "essential patents" should be included in pools and that patent owners retain a right to license their invention separately from the pool (known as "independent licensing"). Among other conclusions, the paper highlights the effectiveness of the demands for independent licensing: a pool is never affected by the possibility of independent licensing if and only if the pool is welfare enhancing.

Lerner et al. (2008) empirically focus on patent pools. They construct a sample of 63 pools established between 1895 and 2001 and determine the structure of these agreements. They then analyze the determinants of the features of these agreements and highlight the extent to which the structure does appear to be consistent with theoretical predictions.

These areas are certainly not the only ones of interest to researchers in terms of new organizational forms and innovation. Open source and academia have many parallels. The most obvious parallel relates to motivation. As in open source, the direct financial returns from writing academic articles are typically nonexistent, but career concerns and the desire for peer recognition provide powerful inducements. In recent years, academic institutions have begun more aggressively experimenting with new organizational forms, whether in the hopes of more thoroughly diffusing their discoveries (e.g., the Biological Resource Centers; Furman and Stern 2011) or to more effectively profit from their intellectual property.

Another fascinating and important area is the study of standard-setting bodies. This topic is important because how open a standard is can critically affect its evolution. The rapidity with which the standard is adopted and the incentives to innovate may be shaped by this decision. For instance, the Internet today runs on a nonproprietary architecture largely because the Internet Engineering Task Force in its early years had a strict policy of only incorporating technology where the developer agreed to license it on reasonable and nondiscriminatory terms. Had the Standard-Setting Organization had a more permissive policy (as indeed they adopted in the mid-1990s), the development of the Internet may have been very different (Bradner 1999). Although there have been some initial looks at standard-setting bodies and how they affect innovation (Lerner and Tirole 2006; Chiao et al. 2007; Greenstein and Stango 2007; Simcoe 2012), much more needs to be done.

4. The Financing of Innovation

Another cluster of relevant research has been in finance literature. Although the relationship between finance and innovation has been studied in many contexts not particularly relevant to organizational economics—for instance, work on the market reaction to the announcement of R&D projects and the value relevance of R&D expenditures in accounting statements—other work is very germane.

In this section, we highlight two relevant bodies of work. The first of these stresses the importance of capital constraints and the financing of innovation. The second focuses on financial intermediaries—particularly venture capitalists—and their role in fomenting innovation.[8]

Another source of uncertainty surrounds the ability of capital constraints—or limits on the ability to raise external financing—to affect the innovation process. Even if a firm has a great idea, it may be unable to raise the capital to market it.

4.1. Capital Constraints and Innovation

The examination of financial constraints has been an important topic in corporate finance. In practitioner accounts and theoretical models, the state of financial markets may limit a firm's ability to raise outside capital. This inability to raise capital may stem from information gaps or an inability to reach a satisfactory contract on the outcomes of a venture. In this environment, good projects may find it impossible to get financing on reasonable terms if the firm does not have enough internal funds.

8. Even here, there is other work we do not discuss. One example is work on R&D financing organizations, such as those typically seen in the biotechnology industry (Beatty et al. 1995).

A variety of empirical studies of financial constraints explore the investment behavior of firms and its sensitivity to changes in internally generated funds. Fazzari et al. (1988) find that firms with low or no dividend payout ratios are more likely to have investment that was sensitive to changes in free cash flow. The authors interpret their results as demonstrating that capital constraints likely affect companies that do not pay dividends, as they forego investment when internal cash is not available. In a similar vein, Lamont (1997) looks at companies that have oil-related production and nonoil-related businesses. He finds that investments in the nonoil-related businesses are dramatically affected by swings in the world price of oil. This is true even though the firm's nonoil businesses were largely uncorrelated with the prospects for their oil businesses. He interprets this result as suggesting the capital constraints limit the ability of firms to raise outside financing.

Since Arrow (1962), it has been understood that not only investments in physical goods but also innovation should be affected by an inability to raise capital. The substantial information problems surrounding R&D projects make it difficult to raise external capital to finance them. As a result, firms with promising projects may be unable to pursue them.

The sensitivity of innovation to capital constraints has been corroborated in a number of studies. For instance, Himmelberg and Petersen (1994) look at a panel of small firms and show that the sensitivity of R&D investment to cash flow seems to be considerably greater than that of physical investment. This suggests that the problems discussed above are if anything more severe for innovation as opposed to more traditional capital investments.

Economists have also studied the impact of debt on R&D spending. The classic study in this mold, Hall (1990), shows that firms that take on more debt and increase their leverage tend to reduce R&D spending. Similar conclusions are offered by Greenwald et al. (1992).

Thus, even if a firm develops an invention, it may not have the resources to commercialize it. The difficulty of raising capital stemming from the information problems surrounding the proposed innovation may preclude the development of innovations that would be otherwise successful.

4.2. Financial Intermediaries and Innovation

A second relevant stream of the finance literature on innovation has examined the relationship between intermediaries and innovation. Much of the attention here has focused on venture capital organizations, that is, independently managed, dedicated capital focusing on equity or equity-linked investments in privately held high-growth companies.

Venture capital has attracted extensive theoretical scrutiny. These works suggest that venture capitalists promote innovation by mitigating agency conflicts between entrepreneurs and investors. The improvement in efficiency might be due to the active monitoring and advice that is provided (Hellmann 1998; Marx 2000; Cornelli and Yosha 2003), the screening mechanisms employed (Chan 1983), the incentives to exit (Berglöf 1994), the proper syndication of the investment (Admati and Pfleiderer 1994), or investment staging (Sahlman 1990; Bergemann and Hege 1998).

Theorists have suggested a variety of mechanisms by which venture capital may affect innovation, but the empirical record is more mixed. It might be thought that establishing a relationship between venture capital and innovation would be straightforward. For instance, one could look at regressions across industries and time to determine whether, controlling for R&D spending, venture capital funding has an impact on various measures of innovation. But even a simple model of the relationships among venture capital, R&D, and innovation suggests that this approach is likely to give misleading estimates.

Both venture funding and innovation could be positively related to a third unobserved factor, the arrival of technological opportunities. Thus, there could be more innovation at times when there was more venture capital, not because the venture capital caused the innovation, but rather because the venture capitalists reacted to some fundamental technological shock that was sure to lead to more innovation. To date, only two papers have attempted to address these challenging issues.

The first of these papers, Hellmann and Puri (2000), examines a sample of 170 recently formed firms in Silicon Valley, including both venture-backed and nonventure firms. Using questionnaire responses, the authors find empirical evidence that venture capital financing is related to product market strategies and outcomes of start-ups. They find that firms that are pursuing what they term an "innovator strategy" (a classification based on the content analysis of survey responses) are significantly more likely to obtain venture capital and to obtain it more quickly. The presence of a venture capitalist is also associated with a significant reduction in the time taken to bring a product to market, especially for innovators. Furthermore, firms are more likely to list obtaining venture capital as a significant milestone in the lifecycle of the company compared to other financing events.

The results suggest significant interrelations between investor type and product market dimensions, and a role for venture capital in encouraging innovative companies. Given the small size of the sample and the limited data, Hellmann and Puri (2000) can only modestly address concerns about causality. Unfortunately, the possibility remains that more innovative firms select venture capital for financing, rather than venture capital causing firms to be more innovative.

In contrast, the second paper, Kortum and Lerner (2000), examines whether these patterns can be discerned on an aggregate industry level rather than on the firm level. These authors address concerns about causality in two ways. First, they exploit the major discontinuity in the recent history of the venture capital industry: in the late 1970s, the U.S. Department of Labor clarified the Employee Retirement Income Security Act, a policy shift that freed pensions to invest in venture capital. This shift led to a sharp increase in the funds committed to venture capital. This type of exogenous change should identify the role of venture capital, because it is unlikely to be related to the arrival of entrepreneurial opportunities. They exploit this shift in instrumental variable regressions. Second, they use R&D expenditures to control for the arrival of technological opportunities that are anticipated by economic actors at the time but are unobserved by econometricians. In the framework of a simple model, they show that the causality problem disappears if they estimate the impact of venture capital on the patent-R&D ratio, rather than on patenting itself.

Even after addressing these causality concerns, the results suggest that venture funding does have a strong positive impact on innovation. The estimated coefficients vary according to the techniques employed, but on average a dollar of venture capital appears to be three to four times more potent in stimulating patenting than a dollar of traditional corporate R&D. The estimates therefore suggest that venture capital, even though it averaged less than 3% of corporate R&D from 1983 to 1992, is responsible for a much greater share—perhaps 10%—of U.S. industrial innovations in this decade.

Some of the most interesting theoretical work in recent years has focused not on the question of whether venture capitalists spur innovation but rather on the societal consequences of the relationship between venture-backed entrepreneurship and innovation. Landier (2006) presents a model in which entrepreneurial ventures succeed or fail on the basis of ability and luck.[9] He argues that as the venture progresses, the entrepreneur is likely to learn about the probable eventual success of the venture, but that the decision to continue or abandon it will not be the same in all environments. In particular, the decision depends critically on how expensive it would be to raise capital for a new venture from investors after a failure. In this setting, Landier shows, multiple equilibria can arise. If the cost of capital for a new venture after a failure is not very high, entrepreneurs will be willing to readily abandon ventures, and failure is commonplace but not very costly. Alternatively, if the cost of capital for failed entrepreneurs is high, only extremely poor projects will be abandoned. Thus, societies may differ dramatically in the prevalence of experimentation in high-risk, innovative ventures.

5. Conclusion

In this chapter, we review the key features of the literature on innovation and organizational structure. We highlight the key areas where work as been undertaken, as well as the limitations of the literature to date.

Stepping back, two metathemes emerge across the strands of literature reviewed above. The first is that there is now a much more fruitful back-and-forth between theoretical and empirical work than was the case at the inception of the field. The patent race literature, for example, was only loosely motivated by empirical observations and did not (at least initially) necessarily provide much grist for the empiricist's mill (e.g., Cockburn and Henderson 1994). Similarly, much of the early empirical work on the relationship between firm size and innovation was atheoretical.

In recent years, these two approaches have been moving more closely together. Consider, for example, the response to the rise of an active market for early development projects in the biopharmaceutical sector. These changes have made salient—in theorists' minds—the possibility of appropriating at least part of the returns from innovation on the ideas market, rather than on the product market. In addition, they have stimulated both theoretical work and empirical investigations that test these theories' specific predictions, either in different papers (Gans and Stern 2000; Gans et al. 2002) or sometimes in a single paper combining theory and evidence (Lerner and Malmendier 2010). These developments bode well for future research in this area.

The second theme pertains to the style of empirical work practiced by applied economists studying innovation and organizations. Although aware of the perils involved in sweeping generalizations, it seems clear to us that we are lagging behind other subfields of microeconomics in our treatment of endogeneity and unobserved heterogeneity. It is possible that our subject matter provides inherently fewer opportunities to exploit natural experiments, relative to the type of questions studied by health, development, or labor economists. An alternative explanation is that we have collectively been too reluctant to place identification at the center of our preoccupations. This reluctance is not always for the worse. Researchers often face trade-offs between

9. See also Gromb and Scharfstein (2002) for a thoughtful theoretical analysis that touches on many of these issues.

the importance of the problem they tackle and the degree of confidence they can achieve in providing answers to this problem.

At the same time, we suggest that putting relatively more emphasis on the establishment of causal relationships—as opposed to the careful documentation of conditional correlations—would be a welcome development. One hallmark of the fields of technical change and organizational economics is that researchers often laboriously produce the data that they analyze rather than simply consuming existing sources of administrative data, such as the Current Population Survey or the Panel Survey of Income Dynamics. Thus, in the quest for exogenous sources of variation, we face potentially fewer constraints than do researchers in other subfields of economics, provided that we channel more effort into addressing identification problems at the time of a project's inception. Stern's (2004) study of compensating differentials in the entry-level science labor market is an exemplar of this particularly creative genre. More studies in that vein on the traditional concerns of the field (e.g., the estimation of knowledge spillovers) would enable us to step back and sort out what we can confidently claim to know from what is at most informed speculation.

REFERENCES

Abramowitz, Moses. 1956. "Resources and Output Trends in the U.S. since 1870." Occasional Paper 52, National Bureau of Economic Research, New York.

Acemoglu, Daron, Philippe Aghion, and Fabrizio Zilibotti. 2003. "Vertical Integration and Distance to Frontier." *Journal of the European Economic Association* 1: 630–638.

Adams, James D., and Adam B. Jaffe. 1996. "Bounding the Effects of R&D: An Investigation Using Matched Establishment-Firm Data." *RAND Journal of Economics* 27: 700–721.

Admati, Anat R., and Paul Pfleiderer. 1994. "Robust Financial Contracting and the Role for Venture Capitalists." *Journal of Finance* 49: 371–402.

Aghion, Phillipe, and Patrick Bolton. 1992. "An Incomplete Contract Approach to Financial Contracting." *Review of Economic Studies* 59: 473–494.

Aghion, Phillipe, and Jean Tirole. 1994. "On the Management of Innovation." *Quarterly Journal of Economics* 109: 1185–1207.

Almeida, Paul, and Bruce Kogut. 1999. "The Localization of Knowledge and the Mobility of Engineers in Regional Networks. *Management Science* 45: 905–917.

Amabile, Teresa M. 1993. "Motivational Synergy: Toward New Conceptualizations of Intrinsic and Extrinsic Motivation in the Workplace." *Human Resource Management Review* 3: 185–201.

———. 1996. *Creativity in Context*. Boulder: Westview Press.

Anton, James J., and Dennis A. Yao. 1994. "Expropriation and Inventions: Appropriable Rents in the Absence of Property Rights." *American Economic Review* 84: 190–209.

Argyres, Nicholas S., and Brian S. Silverman. 2004. "R&D, Organization Structure and the Development of Corporate Technological Knowledge." *Strategic Management Journal* 25: 925–958.

Aron, Debra J., and Edward P. Lazear. 1990. "The Introduction of New Products." *American Economic Review* 80(papers and proceedings): 421–426.

Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella. 2002. *Markets for Technology: The Economics of Innovation and Corporate Strategy*. Cambridge, MA: MIT Press.

Arrow, Kenneth. J. 1962. "Economic Welfare and the Allocation of Resources for Innovation," in Richard R. Nelson (ed.), *The Rate and Direction of Inventive Activity: Economic and Social Factors.* Princeton, NJ: Princeton University Press, pp. 609–625.

Azoulay, Pierre. 2004. "Capturing Knowledge within and across Firm Boundaries: Evidence from Clinical Development." *American Economic Review* 94: 1591–1612.

Azoulay, Pierre, Joshua Graff Zivin, and Jialan Wang. 2010. "Superstar Extinction." *Quarterly Journal of Economics* 125: 549–589.

Azoulay, Pierre, Joshua Graff Zivin, and Gustavo Manso. 2011. "Incentives and Creativity: Evidence from the Academic Life Sciences." *RAND Journal of Economics* 42: 527–554.

Baker, Malcolm, Richard Ruback, and Jeffrey Wurgler. 2007. "Behavioral Corporate Finance: A Survey," in Espen Eckbo (ed.), *Handbook of Corporate Finance: Empirical Corporate Finance*. Amsterdam: Elsevier/North-Holland, pp. 145–187. [Formerly Working Paper 10863, National Bureau of Economic Research, Cambridge, MA.]

Beatty, Anne, Philip G. Berger, and Joseph Magliolo. 1995. "Motives for Forming Research & Development Financing Organizations." *Journal of Accounting and Economics* 19: 411–142.

Bergemann, Dirk, and Ulrich Hege. 1998. "Venture Capital Financing, Moral Hazard, and Learning." *Journal of Banking and Finance* 22: 703–735.

Berger, Allen N., Nathan H. Miller, Mitchell A. Petersen, Raghuram G. Rajan, and Jeremy C. Stein. 2005. "Does Function Follow Organizational Form? Evidence from the Lending Practices of Large and Small Banks." *Journal of Financial Economics* 76: 237–269.

Berglöf, Erik. 1994. "A Control Theory of Venture Capital Finance." *Journal of Law, Economics, and Organization* 10: 247–267.

Bloom, Nick, and John Van Reenen. 2010. "Why Do Management Practices Differ across Firms and Countries?" *Journal of Economic Perspectives* 24(1): 203–224.

Bradner, Scott. 1999. "The Internet Engineering Task Force," in Chris DiBona, Sam Ockman, and Mark Stone (eds.), *Open Sources: Voices from the Open Source Revolution*. Sebastopol, CA: O'Reilly Media, pp. 47–52.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics* 117: 339–376.

Burt, Ronald S. 1992. *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

———. 2004. "Structural Holes and Good Ideas." *American Journal of Sociology* 110: 349–399.

Chan, Yuk-Shee. 1983. "On the Positive Role of Financial Intermediation in Allocation of Venture Capital in a Market with Imperfect Information." *Journal of Finance* 38: 1543–1568.

Che, Yeon-Koo, and Seung-Weon Yoo. 2001. "Optimal Incentives for Teams." *American Economic Review* 91: 525–541.

Chiao, Benjamin, Josh Lerner, and Jean Tirole. 2007. "The Rules of Standard Setting Organizations: An Empirical Analysis." *RAND Journal of Economics* 38: 905–930.

Christensen, Clayton. 1997. *The Innovator's Dilemma*. Cambridge, MA: Harvard Business School Press.

Cockburn, Iain, and Rebecca Henderson. 1994. "Racing to Invest? The Dynamics of Competition in Ethical Drug Discovery." *Journal of Economics and Business Strategy* 3: 481–519.

———. 1998. "Absorptive Capacity, Coauthoring Behavior, and the Organization of Research in Drug Discovery." *Journal of Industrial Economics* 46: 157–182.

Cockburn, Iain, Rebecca Henderson, and Scott Stern. 1999. "Balancing Incentives: The Tension between Basic and Applied Research." Working Paper 6882, National Bureau of Economic Research, Cambridge, MA.

Cohen, Wesley M., and Richard Levin. 1989. "Empirical Studies of Innovation and Market Structure," in Richard Schmalensee and Robert Willig (eds.), *Handbook of Industrial Organization,* volume II. Amsterdam: North-Holland, pp. 1059–1107.

Cohen, Wesley M., and Daniel A. Levinthal. 1989. "Innovation and Learning: The Two Faces of R&D." *Economic Journal* 99: 569–596.

Cohen, Wesley M., and Daniel A. Levinthal. 1990. "Absorptive Capacity: A New Perspective on Learning and Innovation." *Administrative Science Quarterly* 35: 128–152.

Cornelli, Francesca, and Oved Yosha. 2003. "Stage Financing and the Role of Convertible Debt." *Review of Economic Studies* 70: 1–32.

Cournot, Augustin. 1838. *Recherches sur les Principes Mathématiques de la Théorie des Richesses.* Paris: Hachette.

Dix, Manfred, and Néstor Gandelman. 2007. "R&D Institutional Arrangements: Start-up Ventures versus Internal Lab." *Manchester School* 75: 218–236.

Ederer, Florian P., and Gustavo Manso. 2011. "Is Pay-for-Performance Detrimental to Innovation?" SSRN working paper. Available at Social Science Research Network: http://ssrn.com/abstract=1270384 or http://dx.doi.org/10.2139/ssrn.1270384.

Elfenbein, Dan, and Josh Lerner. 2003. "Ownership and Control Rights in Internet Portal Alliances." *RAND Journal of Economics* 34: 356–369.

Fallick, Bruce, Charles A. Fleischmann, and James B. Rebitzer. 2006. "Job Hopping in Silicon Valley: Some Evidence Concerning the Micro-Foundations of a High Technology Cluster." *Review of Economics and Statistics* 88: 472–481.

Fazzari, Steven M., R. Glenn Hubbard, and Bruce C. Petersen. 1988. "Financing Constraints and Corporate Investment." *Brookings Papers on Economic Activity: Microeconomics* 1: 141–205.

Furman, Jeffrey L., and Scott Stern. 2011. "Climbing Atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research." *American Economic Review* 101: 1933–1963.

Gans, Joshua S., and Scott Stern. 2000. "Incumbency and R&D Incentives: Licensing the Gale of Creative Destruction." *Journal of Economics and Management Strategy* 9: 485–511.

Gans, Joshua S., David H. Hsu, and Scott Stern. 2002. "When Does Start-up Innovation Spur the Gale of Creative Destruction?" *RAND Journal of Economics* 33: 571–586.

Gilbert, Richard J., and David M. G. Newberry. 1982. "Preemptive Patenting and the Persistence of Monopoly." *American Economic Review* 72: 514–526.

Gompers, Paul, Josh Lerner, and David Scharfstein. 2005. "Entrepreneurial Spawning: Public Corporations and the Formation of New Ventures, 1986–1999." *Journal of Finance* 60: 577–614.

Greenstein, Shane, and Victor Stango (eds.). 2007. *Standards and Public Policy.* Cambridge: Cambridge University Press.

Greenwald, Bruce, Michael Salinger, and Joseph E. Stiglitz. 1992. "Imperfect Capital Markets and Productivity Growth." Working paper, Bellcore, Morristown, NJ; Boston University; and Stanford University, Stanford, CA.

Griliches, Zvi (ed.). 1984. *Patents, R&D, and Productivity*. Chicago: University of Chicago Press.

Gromb, Denis, and David Scharfstein. 2002. "Entrepreneurship in Equilibrium." Working Paper 9001, National Bureau of Economic Research, Cambridge, MA.

Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Lateral and Vertical Integration." *Journal of Political Economy* 94: 691–719.

Guedj, Ilan, and David Scharfstein. 2004. "Organizational Scope and Investment: Evidence from the Drug Development Strategies and Performance of Biopharmaceutical Firms." Working Paper 10933, National Bureau of Economic Research, Cambridge, MA.

Hall, Bronwyn H. 1990. "The Impact of Corporate Restructuring on Industrial Research and Development." *Brookings Papers on Economic Activity: Microeconomics* 1990: 85–124.

———. 1991. "The Effect of Takeover Activity on Corporate Research and Development," in Alan J. Auerbach (ed.), *Takeovers: Causes and Consequences.* Chicago: University of Chicago Press, pp. 69–87.

Hann, Il-Horn, Jeff Roberts, Sandra Slaughter, and Roy Fielding. 2004. "An Empirical Analysis of Economic Returns to Open Source Participation." Working paper, Carnegie-Mellon University, Pittsburgh, PA.

Hart, Oliver D., and John Moore. 1988. "Incomplete Contracts and Renegotiation." *Econometrica* 56: 755–785.

Hellmann, Thomas. 1998. "The Allocation of Control Rights in Venture Capital Contracts." *RAND Journal of Economics* 29: 57–76.

———. 2007. "When Do Employees Become Entrepreneurs?" *Management Science* 53: 919–933.

Hellmann, Thomas, and Enrico Perotti. 2011. "The Circulation of Ideas in Firms and Markets." *Management Science* 57: 1813–1826.

Hellmann, Thomas, and Manju Puri. 2000. "The Interaction between Product Market and Financing Strategy: The Role of Venture Capital." *Review of Financial Studies* 13: 959–984.

Henderson, Rebecca. 1993. "Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Industry." *RAND Journal of Economics* 24: 248–270.

Henderson, Rebecca, and Iain Cockburn. 1994. "Measuring Competence—Exploring Firm Effects in Pharmaceutical Research." *Strategic Management Journal* 15: 63–84.

Himmelberg, Charles P., and Bruce C. Petersen. 1994. "R&D and Internal Finance: A Panel Study of Small Firms in High-Tech Industries." *Review of Economics and Statistics* 76: 38–51.

Holmes, Thomas J., and James A. Schmidt, Jr. 1990. "A Theory of Entrepreneurship and Its Application to the Study of Business Transfers." *Journal of Political Economy* 98: 265–294.

Holmström, Bengt. 1989. "Agency Costs and Innovation." *Journal of Economic Behavior and Organization* 12: 305–327.

Holmström, Bengt, and Jean Tirole. 1997. "Financial Intermediation, Loanable Funds, and the Real Sector." *Quarterly Journal of Economics* 112: 663–691.

Jaffe, Adam B. 1986. "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value." *American Economic Review* 76: 984–1001.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Quarterly Journal of Economics* 108: 578–598.

Jensen, Michael C. 1986. "Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers." *American Economic Review* 76(papers and proceedings): 323–329.

Jewkes, John. 1958. "The Sources of Invention." *The Freeman: Ideas on Liberty* 58(4): 45–57.

Jovanovic, Boyan. 1979. "Firm-Specific Capital and Turnover." *Journal of Political Economy* 87: 1246–1260.

Kandel, Eugene, and Edward P. Lazear. 1992. "Peer Pressure and Partnerships." *Journal of Political Economy* 100: 801–817.

Klepper, Steven. 2008. "Product Life Cycle," in Steven N. Durlauf and Lawrence E. Blume (eds.), *The New Palgrave Dictionary of Economics,* second edition. London: Palgrave Macmillan. Available at: http://www.dictionaryofeconomics.com/article?id=pde2008_P000331.

Klepper, Steven, and Peter Thompson. 2005. "Spinoff Entry in High-Tech Industries: Motives and Consequences." Working Paper 0503, Florida International University, Miami.

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

Lacetera, Nicola, Iain Cockburn, and Rebecca Henderson. 2004. "Do Firms Change Their Capabilities by Hiring New People? A Study of the Adoption of Science-Driven Drug Research," in Joel Baum and Anita McGahan (eds.), *Business Strategy over the Industry Lifecycle: Advances in Strategic Management,* volume 21. New York: JAI/Elsevier Science, pp. 133–159.

Lamont, Owen. 1997. "Cash Flow and Investment: Evidence from Internal Capital Markets." *Journal of Finance* 52: 83–109.

Lamoreaux, Naomi R., and Kenneth L. Sokoloff. 2009. "The Decline of the Independent Inventor: A Schumpeterian Story?" in Sally H. Clarke, Naomi R. Lamoreaux, and Steven Usselman (eds.), *The Challenge of Remaining Innovative: Lessons from Twentieth Century American Business*. Stanford: Stanford University Press, pp. 43–78.

Landier, Augustin. 2006. "Start-up Financing: Banks vs. Venture Capital." Working paper, Massachusetts Institute of Technology, Cambridge, MA.

Lerner, Josh. 1997. "An Empirical Exploration of a Technology Race." *RAND Journal of Economics* 28: 228–247.

Lerner, Josh, and Ulrike M. Malmendier. 2010. "Contractibility and the Design of Research Agreements." *American Economic Review* 100: 214–246.

Lerner, Josh, and Robert Merges. 1998. "The Control of Technology Alliances: An Empirical Analysis of the Biotechnology Industry." *Journal of Industrial Economics* 46: 125–156.

Lerner, Josh, and Jean Tirole. 2002. "Some Simple Economics of Open Source." *Journal of Industrial Economics* 52: 197–234.

———. 2004. "Efficient Patent Pools." *American Economic Review* 94: 691–711.

———. 2005. "The Scope of Open Source Licensing." *Journal of Law, Economics, and Organization* 21: 20–56.

———. 2006. "A Model of Forum Shopping." *American Economic Review* 96: 1091–1113.

Lerner, Josh, and Julie Wulf. 2007. "Innovation and Incentives: Evidence from Corporate R&D." *Review of Economics and Statistics* 89: 634–644.

Lerner, Josh, Hillary Shane, and Alexander Tsai. 2003. "Do Equity Financing Cycles Matter? Evidence from Biotechnology Alliances." *Journal of Financial Economics* 67: 41g–446.

Lerner, Josh, Marcin Strojwas, and Jean Tirole. 2008. "The Design of Patent Pools: The Determinants of Licensing Rules." *RAND Journal of Economics* 38: 610–625.

Levin, Richard C., Alvin K. Klevorick, Richard R. Nelson, and Sidney G. Winter. 1987. "Appropriating the Returns from Industrial Research and Development." *Brookings Papers on Economic Activity* 3 (special issue on microeconomics): 783–831.

Lim, Kwanghui. 2011. "The Many Faces of Absorptive Capacity: Spillovers of Copper Interconnect Technology for Semiconductor Chips." *Industrial and Corporate Change* 18: 1249–1284.

Majewski, Suzanne E. 1998. "Causes and Consequences of Strategic Alliance Formation: The Case of Biotechnology." PhD dissertation, Department of Economics, University of California, Berkeley.

Manso, Gustavo. 2011. "Motivating Innovation." *Journal of Finance* 66: 1823–1860.

March, James G. 1991. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2: 71–87.

Marx, Leslie M. 2000. "Contract Renegotiation in Venture Capital Projects." Working paper, University of Rochester, Rochester, NY.

Marx, Matt, Debbie Strumsky, and Lee Fleming. 2009. "Mobility, Skills, and the Michigan Non-Compete Experiment." *Management Science* 55: 875–889.

Merges, Robert P. 1995. "Intellectual Property and the Costs of Commercial Exchange: A Review Essay." *Michigan Law Review* 93: 1570–1615.

Moen, Jarle. 2005. "Is Mobility of Technical Personnel a Source of R&D Spillovers?" *Journal of Labor Economics* 23: 81–114.

Pakes, Ariel, and Zvi Griliches. 1984. "Patents and R&D at the Firm Level: A First Report," in Zvi Griliches (ed.), *Patents, R&D, and Productivity*. Chicago: University of Chicago Press, pp. 55–72.

Pakes, Ariel, and Shmuel Nitzan. 1983. "Optimal Contracts for Research Personnel, Research Employment and the Establishment of 'Rival' Enterprises." *Journal of Labor Economics* 1: 345–365.

Prusa, Thomas J., and James A. Schmitz, Jr. 1994. "Can Companies Maintain Their Initial Innovation Thrust? A Study of the PC Software Industry." *Review of Economics and Statistics* 76: 523–540.

Reagans, Ray E., Ezra W. Zuckerman, and Bill McEvily. 2004. "How to Make the Team: Social Networks vs. Demography as Criteria for Designing Effective Projects in a Contract R&D Firm." *Administrative Science Quarterly* 49: 101–133.

Reinganum, Jennifer F. 1983. "Uncertain Innovation and the Persistence of Monopoly." *American Economic Review* 73: 741–748.

———. 1989. "The Timing of Innovation: Research, Development, and Diffusion," in Richard Schmalensee and Robert Willig (eds.), *Handbook of Industrial Organization,* volume I. Amsterdam: North-Holland, pp. 849–908.

Robinson, David, and Toby E. Stuart. 2007. "Financial Contracting in Biotech Strategic Alliances." *Journal of Law and Economics* 50: 559–596.

Rosenbloom, Richard S., and William J. Spencer (eds.). 1996. *Engines of Innovation: U.S. Industrial Research at the End of an Era*. Boston: Harvard Business School Press.

Rosenkopf, Lori, and Paul Almeida. 2003. "Overcoming Local Search through Alliances and Mobility." *Management Science* 49: 751–766.

Sahlman, William A. 1990. "The Structure and Governance of Venture Capital Organizations." *Journal of Financial Economics* 27: 473–524.

Saxenian, Annalee. 1994. *Regional Advantage: Culture and Competition in Silicon Valley and Route 128.* Cambridge, MA: Harvard University Press.

Scharfstein, David S., and Jeremy C. Stein. 2000. "The Dark Side of Internal Capital Markets: Divisional Rent-Seeking and Inefficient Investment." *Journal of Finance* 55: 2537–2564.

Schumpeter, Joseph A. 1942. *Capitalism, Socialism and Democracy*. New York: Harper and Row.

Scott, John T. 1984. "Patents and R&D at the Firm Level: A First Report," in Zvi Griliches (ed.), *Patents, R&D, and Productivity*. Chicago: University of Chicago Press, pp. 233–245.

Seru, Amit. 2010. "Firm Boundaries Matter: Evidence from Conglomerates and R&D Activity." *Journal of Financial Economics,* forthcoming.

Shane, Hillary. 1995. "Asymmetric Information and Alliance Financing in the Biotechnology Industry," in *Three Essays in Empirical Finance in High-Technology Firms.* PhD dissertation, Wharton School, University of Pennsylvania, Philadelphia.

Shapiro, Carl. 2001. "Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting." *Innovation Policy and the Economy* 1: 119–150.

Simcoe, Timothy. 2012. "Standard Setting Committees: Consensus Governance for Shared Technology Platforms." *American Economic Review* 102: 305–336.

Solow, Robert M. 1957. "Technological Change and the Aggregate Production Function." *Review of Economics and Statistics* 39: 312–320.

Stein, Jeremy C. 2002. "Information Production and Capital Allocation: Decentralized versus Hierarchical Firms." *Journal of Finance* 57: 1891–1921.

Stern, Scott. 2004. "Do Scientists Pay to Be Scientists?" *Management Science* 50: 835–853.

Teece, David J. 1986. "Profiting from Technological Innovation: Implications for Integration, Collaboration, Licensing and Public Policy." *Research Policy* 15: 285–305.

Thompson, Peter, and Melanie Fox-Kean. 2005. "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment." *American Economic Review* 95: 450–460.

Tushman, Michael. L., and Philip Anderson. 1986. "Technological Discontinuities and Organizational Environments." *Administrative Science Quarterly* 31: 439–465.

Wiggins, Steven N. 1995. "Entrepreneurial Enterprises, Endogenous Ownership, and the Limits to Firm Size." *Economic Inquiry* 33:54–69.

Zucker, Lynne G., Michael R. Darby, and Marilynn B. Brewer. 1998. "Intellectual Human Capital and the Birth of U.S. Biotechnology Enterprises." *American Economic Review* 88: 290–306.