1 Markets for Technology: Why Do We See Them, Why We Don't See More of Them, and Why We Should Care

1.1 Objective

It is now a commonplace that we live in a knowledge economy. Like all cliches, this one is also wrong insofar as it suggests that earlier economies did not rely upon knowledge. If there is something different about the economic system that has characterized the majority of industrialized countries over the last two and a half centuries, it is arguably the increased importance of scientific and technological knowledge for economic activity. According to Simon Kuznets (1969), the distinguishing characteristic of modern economic growth has been the systematic application of science to economic ends. Understanding how scientific and technological knowledge is produced and applied to economic goals is the key to understanding the process of modern economic growth.

This book studies the nature and working of markets for technology, namely for intermediate technological inputs—and the implications for business and public policy. Although a wealth of scholarly research in economics exists on this subject, there is very little on how a market in knowledge would function, other than the appreciation that such markets would be characterized by a number of imperfections. Similarly, there is little guidance in the management literature on how managers should behave when markets for technology are present. This neglect is understandable. Although markets for technology have existed for a long time, with the advent of the corporate Research and Development (R&D) laboratory, firms began to develop their own technology. Drawing on the idea of imperfections in technology markets, Nelson (1959) provided the first rationalization for why and under what conditions firms would invest in R&D, an idea elegantly generalized by Arrow (1962a). But perhaps even more powerful has been the firm's-eye view of twentieth-century American economic growth offered by Chandler (1990). For Chandler, the systematic application of science takes place within the more organized confines of the firm, as the production of new knowledge is combined with its application through mutually complementary investments in research, manufacturing, and marketing.

This vision of knowledge creation integrated with knowledge use has become inadequate for understanding economic growth in the twenty-first century. Over the past ten to fifteen years, there has been a rapid growth in a variety of arrangements for the exchange of technologies or technological services, ranging from R&D joint ventures and partnerships, to licensing and cross-licensing agreements, to contracted R&D. Although we lack comprehensive empirical measures of the increase in such arrangements over time, all the available evidence suggests that the trade in technologies is more common than it was in the past. For instance, the industry cases collected in Mowery (1988) suggest that since the 1980s there has been an increase in the number of collaborative ventures among firms, especially of those involving R&D and technology. Grindley and Teece (1997) point to the increasing use of technology licensing by companies such as IBM, Hewlett-Packard, Texas Instruments, and AT&T during the 1990s. A number of firms and software products have emerged to help firms manage their patent portfolios. Firms specializing in the creation of new technology are now an important part of the industrial landscape in many technology-intensive industries. Finally, we have seen the development of electronic and online market places where technologies can be bought and sold.

Many practitioners and scholars have noted these trends (e.g., Rivette and Kline 1999; Teece 1998). However, what is still lacking is a thorough and systematic understanding of how markets for technology arise, how they work, their limits and implications for public policy and corporate strategy. This book is a step in that direction. First, we look at the role of industry structure, the nature of knowledge, and intellectual property rights and related institutions that facilitate the development of markets for technology. Second, we ask what the implications of such markets are for the boundaries of the firm, division of labor in the economy, industry structure, and economic growth. Third, we build on this discussion to draw implications for public policy and corporate strategy. We combine theoretical perspectives from economics and management and draw upon several rich data sources to exemplify and substantiate the theoretical points.

We do not wish to suggest that in-house R&D in corporations will be supplanted by externally conducted R&D. Rather, we want to understand the conditions under which technology can be traded, be it by established firms or by firms specializing in the production of technology. In addition to the diffusion of technology, such transactions could play an important role in fostering innovation. This is the case when the developers of the technology lack the resources necessary to commercialize the technology. Without the prospect of being able to capitalize on their innovations by trading them, many small technology-based firms would not invest in creating new and useful technologies.

1.2 Markets for Technology: Scope of the Analysis and a Typology

1.2.1 A Tentative Definition

Technologies come in very different forms, and it is difficult to provide a general definition that would satisfactorily encompass all interesting cases of technology trade that we, or the reader, could think up. For instance, technology can take the form of "intellectual property" or intangibles (e.g., a software program, or a design), be embodied in a product (e.g., a prototype, or an instrument, like a chip designed to perform certain operations), or take the form of technical services. We will not attempt to define *technology.* Rather, we treat technology as an imprecise term for useful knowledge rooted in engineering and scientific disciplines, but also drawing from practical experience from production.

Our task is further complicated by the fact that technological knowledge can exist in many forms, where the distinction between physical products and technology is not always easy to make. Some forms present no real difficulty. Transactions involving blueprints, designs, formulae, or flowcharts are clearly part of the market for technology. In general, when the right to produce something or the knowledge of how to do so are separated from the thing itself, there is a clear line between the market for the thing itself and the market for the technology used to create it.

But technology can also be embodied in physical artifacts. For instance, a new method for rapidly screening biological compounds may be embodied in a chip that performs the screening. In this case, by purchasing the good, one also purchases the technology. Increasingly, firms are embodying their technology in software programs. Once again, the purchase of the software brings with it the right to use the embodied knowledge—and hence, software is commonly licensed rather than sold outright.

Our general criterion is to look at whether the cost of developing the knowledge embodied in the artifact significantly exceeds the cost of creating the artifact. Another way to make our distinction is to say that in a market for technology, the suppliers have a great deal of autonomy in designing and developing the good, as compared to suppliers that produce according to detailed specifications by their clients. In this case, the value associated with the design and conceptualization of the product would considerably exceed the value to the buyer that is associated with the mere outsourcing of the manufacturing operations.

This criterion is the hardest to apply in the case where knowledge is embodied in a software program. For instance, knowledge about how to test and debug a microprocessor design may be most effective if embodied in a software program. Although in principle the two—the knowledge and the program—are separate, as a practical matter they may be very closely linked because the software makes the knowledge operational and accessible to a much larger group of users. Having the required knowledge of microprocessor design is clearly a prerequisite for developing the software in question but is by no means sufficient, and considerable ingenuity and thought may go into the development of the software itself. In other words, we cannot hope to eliminate all the "gray" areas. Nor is it necessary. All that is needed is that there be enough of a "core" of transactions where the principal focus of the trade is knowledge rather than a physical artifact.

We also use the term *market* in a broad sense. Strictly speaking, market transactions are arms-length, anonymous, and typically involve an exchange of a good for money. Many, if not most, transactions for technology would fail one or the other of these criteria. Often these transactions involve detailed contracts and may be embedded in a technological alliance of some sort. Although the specific form of the transaction may affect the outcome in subtle ways, we shall ignore many of these subtleties in an effort to focus on the issues common to these transactions, such as the role of specialized technology suppliers, the role of intellectual property rights, and the nature of demand.

A final clarification is that we shall ignore some relevant forms of technology exchange. The trend toward acquisition of small, technology-based companies has become an important phenomenon in recent years. Insofar as they are driven by the need to acquire external technology, outright acquisitions should be included in the market for technology. However, acquisitions encompass not only existing technology, but also the capability and competence to develop new technologies. The issues surrounding the acquisition of technological capability are different from those pertaining to the acquisition of technology. Therefore, we exclude corporate mergers and acquisitions from our analysis. We also disregard another channel through which technological knowledge moves across firm boundaries—the movement of people. Neither omission is indicative of the importance of the phenomenon. Yet, ambitious as this book is, to include inter-firm movement of engineers and researchers would be unworkable.[1]

1.2.2 Markets for Technology and for Innovation

One can distinguish between markets for existing technologies and markets for technologies still being developed. Put differently, our definition of markets for technology covers both "current" and "futures" markets. Both share a number of common features, but there are also some interesting differences that we shall discuss in various chapters.

The U.S. Department of Justice, in its *Antitrust Guidelines for the Licensing of Intellectual Property* (U.S. Deptartment of Justice 1995) makes a similar distinction. The guidelines distinguish between markets for "goods," markets for "technology," and markets for "innovations." Markets for technology are markets for "intellectual property that is licensed and its close substitutes—that is, the technologies or goods that are close enough substitutes significantly to constrain the exercise of market power with respect to the intellectual property that is licensed." Markets for innovation include arrangements in which the parties involved agree to conduct activities, jointly or independently, leading to future developments of technologies that will be exchanged (or jointly owned) among them. This is typically the market for contract R&D and technological joint ventures and collaborations (U.S. Deptartment of Justice 1995, 6).

Roughly speaking, the distinction between the market for technology and the market for innovation as defined by the United States Department of Justice, corresponds to the distinction between transactions for the use and diffusion of technology on one hand, and transactions for the creation of new technology on the other. In addition to contract research, technology licensing, and R&D joint ventures of various kinds, transactions for the creation of new technologies also include the sale or licensing of research tools and transactions for research tools, as well as other types of technical services. Therefore, for this book, the market for technology includes transactions involving full technology packages (patents and other intellectual property and know-how), and patent licensing. Also included are transactions involving knowledge that is not patented and perhaps not even patentable (e.g., software, or many nonpatented designs) but excluding standard software site licenses.

1.2.3 The Division of Innovative Labor

Transactions in the market for technology can be classified in another way, depending on whether they involve "horizontal" transactions among established producers or "vertical" transactions between specialized firms that do not compete. Horizontal transactions (e.g., licensing and technology joint ventures), especially between firms in an industry and particularly at the international level, have been the focus of much of the literature on this subject (e.g., Teece 1977; Contractor 1981; Caves, Crookell, and Killing 1983; Mowery 1988; Anand and Khanna 2000).[2] However, vertical markets, where the technology is supplied to the downstream firms or industries by an upstream sector of specialized technology producers with no stake in the downstream operations, have become increasingly important in several high-tech industries. Further, the development of these vertical markets constitutes a division of labor in the innovation process itself and thus is closely linked to a much older and more powerful set of economic ideas.

As Smith (1776) and Stigler (1951) pointed out, an input produced under increasing returns is supplied more efficiently by a specialized upstream supplier that serves many firms, rather than by the individual downstream companies. Thus, division of labor is more extensive in larger markets. Young (1928) added a dynamic dimension to this analysis. As the more efficient production of the input lowers its unit cost, users are induced to invest, and the demand for the input increases. In turn, this increase in market size further expands the division of labor.

But while extensive specialization and division of labor mark many economic activities, this has not usually been the case in the production of technology. As noted earlier, R&D and technology have been integrated in large firms for many years. As we shall see in chapter 4, economists like Nelson and Winter (1982), and other scholars who built on their work, provided an economic justification for this pattern. The production of technology is a cumulative process based on tacit knowledge and expertise, requiring extensive interactions among the groups and individuals involved. These interactions can be realized more effectively when the individuals or groups belong to the same organization. In essence, there are transaction costs, both static and dynamic, in the exchange of technological knowledge across organizations, which may offset the advantages of a division of labor.

As discussed more fully in later chapters, especially chapters 3 and 6, there are two potential productivity benefits from a division of innovative labor. The first is specialization according to comparative advantage. If firms specializing in research are more efficient at developing new knowledge, while others, possibly those experienced in production and marketing, are more efficient at exploiting the new knowledge, then a division of innovative labor promotes innovation and productivity growth. A second potential benefit arises due to the increasing returns associated with new knowledge (David 2000; Romer 1990). This point can be clarified by looking briefly at what is happening in markets where an active division of innovative labor is taking place. A recent trend in biotechnology, software, and semiconductor sectors is the growth of firms specializing in the production of research methods and "tools" that can be used for several applications. For instance, many biotech companies have developed general-purpose technologies for drug discovery. These include rapid screening of chemical compounds, combinatorial chemistry techniques, and automated tools to assess the relationships between genes and diseases. Specialized tool developers can spread the fixed cost of development over many users. In contrast, a tool developed by a downstream user is applied far more narrowly. From a social point of view, each user developing its own tools means that the fixed costs of development are incurred several times over.

Our distinction between the division of labor and the division of "innovative" labor is meant only to emphasize that the division of labor we focus upon concerns technologies rather than specific products. The properties and implications of a classical division of labor apply to the division of innovative labor as well. However, insofar as the division of innovative labor deals with knowledge and technologies rather than material goods, there are additional factors that affect the form and extent of the division of labor.

One can summarize the foregoing discussion in the form of a simple typology shown in table 1.1. The table also provides a canonical example, taken from a commercial database, of each cell.

Table 1.1
A Simple Typology of Markets for Technology

	Existing Technology	Future Technology or Component for Future
Horizontal market/ Transactions with actual or potential rivals	Union Carbide licensing Uniopol polyethylene technology to Huntsman Chemicals	Sun licensing Java to IBM; R&D joint ventures between rivals
Vertical market/ licensing to nonrivals	Licensing of IP Core in semiconductors	R&D joint ventures; Affymax licensing combinatoric technology to pharmaceutical firm

1.2.4 Markets for Technology: Why We Should Care

Markets for technology promote the diffusion and efficient use of existing technology and can enhance the rate of technological advance by providing additional incentives to invest in research and development. In particular, they can encourage firms to specialize in the production of technology.

Companies, particularly large companies, often develop technologies that they do not commercialize. In many cases, there could be other companies that could profitably use these technologies. Often there are strategic reasons for not licensing unused technologies, including the fear of creating new competitors or of cannibalizing existing markets. Often the reason is different. Technology contracts are thought to be inefficient, and the returns from licensing inadequate to offset other costs. As a result, firms have tended to ignore the option of licensing their technologies.

Technology licenses, and especially international licenses, do exist. However, there is anecdotal evidence that the licensing market is less developed than socially desirable. For instance, a recent study by British Technology Group (BTG), a consulting firm, found that large companies in the United States, Western Europe, and Japan ignore a substantial fraction of their patented technologies, which could be profitably sold or licensed (British Technology Group 1998). Moreover, the study found that companies fail to license not because licensing is unattractive, but simply because they do not take this possibility into account. Similarly, the European Union estimated that 20 billion U.S. dollars are spent every year in Europe to develop new products or ideas that have already been developed elsewhere.[3] Well-functioning markets for technology can improve efficiency by reducing duplicative R&D and by matching technology producers and users.

One objective of this book is therefore to understand the factors that induce established companies to license their technologies and become active suppliers in the market. In particular, we examine how industry structure and competition affect these incentives. Moreover, there are important managerial implications that flow from dealing with technologies as "products." As Grindley and Teece (1997) have noted, this may require different modes of managing the firms; in particular, it requires a different approach to the management of intellectual capital compared with the use of technology merely as an input for the company's final products. Even competitive strategy may change substantially when there are well-functioning markets for technology. For example, in industries like chemicals, extensive technology licensing among established producers has contributed to increased competition in many product markets (Arora and Gambardella 1998).

Technology markets are also a precondition for the existence of specialized technology suppliers operating in vertical markets. Specialization and division of labor is a powerful determinant of industry and economic growth. But specialized suppliers can also act as a mechanism for knowledge transfer that resembles technological spillover across firms, a subject that has attracted a great deal of attention from economists. While spillovers may reduce the private incentives to do R&D, they increase the social returns to R&D and technological investments, and therefore are another source of technological diffusion and growth. There is a large empirical literature that attempts to measure the extent and impact of such spillovers on economic measures of performance such as productivity (Griliches 1979, 1984; Jaffe 1986; Coe and Helpman 1995; Cohen and Levinthal 1989).

However, some so-called spillovers may in fact be market-mediated transfers of knowledge. One thesis of this book is that the intermediation of an upstream sector of technology suppliers can be a powerful mechanism through which spillovers can take place. This recognition that spillovers are not simply "in the air" suggests that they do not arise merely because of geographical agglomeration but require well-defined institutions to work. Moreover, benefiting from spillovers may well require the development of cooperative links or other types of relationships with upstream technology suppliers. Nathan Rosenberg's work provides a compelling historical account of this process. For instance, Rosenberg (1963) describes how specialized machine tool suppliers absorbed and improved the metal-working technology first developed in armament manufacture, and made this improved technology available to other sectors that followed, including sewing machines and bicycles.

In sum, if markets for technology were more extensive and more widespread, existing technologies would stand a better chance of being used, and being used more extensively. New technologies would be more likely to be developed, because even if the technology's inventor did not commercialize the technology itself, the inventor could still profit by licensing the technology to others better able to commercialize it. Not only would there be an increase in the rate of inventive activity, but there would also be profound influences on conditions of entry for new firms and new types of firms, on the competitive position of existing firms, and on the structure of the industry itself.

1.3 Structure of the Book

This book is divided into four parts. Part I provides evidence of the existence of markets for technology. Chapter 2 presents evidence from the available literature, systematic data on worldwide technology licensing deals and related technology transactions in recent years. Using the available data on the values of these transactions, we estimate the total value of worldwide technology transfer deals by the granting and receiving of two-digit SIC sectors. This enables us to quantify the extent of the worldwide market for technology, and to gain insight into market mediated inter-sectoral technology flows.

Chapter 3 provides additional qualitative evidence about markets for technology from four high-tech industries: chemicals, software, biotechnology, and semiconductors. We document the development of markets for technology in these sectors, with a special focus on the division of innovative labor in these industries. These cases illustrate many of the issues discussed in later chapters of the book.

Part II focuses on the limitations and determinants of markets for technology. Chapter 4 deals with the "cognitive" limitations to markets for technology. These limits arise from context dependence, the idea that knowledge created in one context is not readily transferred and used in another context (Arora and Gambardella 1994a). This raises the costs of technology transfer, especially in the context of a division of innovative labor, because context dependence makes it difficult to partition the innovation process into independent activities to be assigned to independent actors.

Chapter 4 begins by illustrating the problems involved in partitioning innovation activities, and continues to review the available literature. It first discusses the literature that follows Nelson and Winter (1982) and Teece (1988), who argue that innovation is largely the outcome of organizational routines, and hence is more effectively performed within organizations. Building on this literature and particularly on the work by Eric Von Hippel (1990, 1994) and Kogut and Zander (1992), we argue that industries and technologies differ in the extent to which task-partitioning is possible. We discuss these conditions and the related differences across industries and technologies. We also argue that changes in the technology of technical change itself, and specifically, the growing use of computers, and information technology in research are enhancing the market for technology and division of innovative labor.

Context-dependent knowledge is less likely to be articulated and codified. Put differently, much of what is useful about technology may be tacit, neither codified nor embodied in machinery or equipment. Not only is tacit knowledge costly to transfer for reasons discussed in Chapter 4, its transfer also raises potential contracting problems. Indeed, the existence of tacit knowledge is a commonly advanced as a reason why technology trade may be inefficient. Chapter 5 develops a formal model showing that when tacit know-how is bundled with complementary codified technology inputs, and the latter is protected by patents or other means, simple contracts can accomplish the transfer of tacit knowledge.

Chapter 5 addresses the suggestion in the literature that part of the difficulty in creating markets for technology is that one cannot exchange tacit knowledge through arm's-length contracts. The chapter shows that the problem of contracting for know-how can be overcome by bundling know-how with complementary codified inputs and leveraging the superior enforceability of contracts over the latter. The chapter also provides empirical support using data on 144 technology import agreements by Indian firms during 1950–1975.

Chapter 6 analyzes the tradeoff between increasing returns to knowledge production and the superior ability of users to understand their own needs. Whereas the latter favors each user developing its own technology, the former favors specialized technology suppliers. This tradeoff determines whether an upstream sector of technology specialists will arise. This chapter examines how two different dimensions of demand—its breadth (the number of users) and depth (the average size of each user)—affect the terms of this tradeoff and hence, have very different effects on the division of innovative labor. We draw on the experience of specific industries like biotechnology and software to provide evidence supporting the theoretical reasoning developed in the chapter.

Part III discusses the functioning of markets for technology. Chapter 7 examines the incentives of established producers to license their own technologies. Here we focus on the interaction between the downstream "goods" market and the market for technology. The key insight is that competition in the goods market can induce licensing of in-house technology by established producers. The logic of the argument is straightforward: Incumbent producers have a disincentive to license because licensing increases competition in the downstream goods market and dissipates rents. However, licensing also creates revenue. Although the existing literature often implicitly assumes that the revenue effect is smaller than the rent dissipation effect, this is not always true. We rigorously explore how market share, the extent of product differentiation, and the efficiency of licensing contracts affect the balance between the two. We also introduce a third effect—the role of competition in the market for technology itself. We show how the presence of other technology holders, particularly firms that only supply technology, can create additional incentives for licensing by established incumbents. The chapter also explores the empirical validity of these ideas using data on chemical process technology licenses.

Chapter 8 focuses on the role of an upstream industry of technology suppliers as a vehicle for transmitting investment opportunities across downstream companies and industries. This chapter highlights how technology is transferred internationally as the outcome of a functioning market for technology. The chapter focuses on the story of the specialized engineering firms (SEFs) in the chemical processing industry (discussed in chapter 3). Beginning in the 1930s and continuing into the 1960s, the rapid growth of the chemical industry in the developed countries stimulated the growth of firms that specialized in the design and engineering of chemical plants—the SEFs. Since the 1970s, as a modern chemical industry emerged in the less developed countries (LDCs) the presence of an upstream sector of technology suppliers in the first world proved very valuable. SEFs had already accumulated expertise in plant design and technology, which could be supplied to the chemical firms in LDCs.

We exploit a rich database on investments in chemical processing industries worldwide from 1980–1990. The empirical analysis shows that the greater the number of technology suppliers (SEFs) that operate in the first world, the greater the investments in chemical plants in LDCs. Moreover, the effect of SEFs is greater for LDC firms rather than multinationals. A major contribution of the analysis developed in this chapter is that it identifies an important and understudied mechanism through which technology is made available and through which spillovers take place—notably, the intermediation of an upstream sector of technology specialists.

Part IV examines the implications for corporate strategy and policy. Chapter 9, which focuses on implications for corporate strategy, also links the discussion in the earlier parts of the book to the "resource-based" view of the firm and clarifies how the development of a market for technology affects corporate boundaries and corporate strategy. It documents the growing recognition by established firms of the importance of technology licensing in relation to their overall business. As the sale of technologies becomes a business of its own, some of these companies are organizing internal divisions focusing on licensing and seeking better ways to manage their intellectual capital and patent portfolios. Second, the chapter argues that for many technology-based smaller firms, licensing may be a better strategy than bearing the costs and risks of downstream manufacturing and commercialization. Third, it points out that markets for technology increase the "penalty" of company strategies based on the notorious "not invented here" syndrome. Finally, markets for technology have natural implications for industry structure. Such markets lower entry barriers and reduce concentration.

Chapter 10 addresses institutional changes and policy implications. This chapter argues that the growth of markets for technology is enhanced by the growth of complementary institutions. It discusses some of these institutions, and the ways in which they can reduce the transaction costs involved in trading technology. In addition to standards, and standard-setting bodies, a key consideration in this respect is intellectual property rights. As discussed in chapter 5, much of the literature on intellectual property rights has focused on the extent to which they provide incentives for firms to invest in R&D. However, as property rights, they can be traded, implying that they can facilitate the efficient utilization of innovations. Such trades could also play an important role in inducing innovation, when the developers of the technology are not the firms best able to commercialize the technology. Without the prospect of being able to capitalize on their innovations by trading the property rights protecting the innovation, many small technology-based firms would not invest in creating new and useful technologies.

In other words, intellectual property rights are the means for defining the object of the transaction and the property rights in the markets for technology. That said, intellectual property differs from tangible property in many important ways. Some scholars have argued that these differences make intellectual property rights more prone to "fragmentation." In this chapter we discuss this theory and some possible policy responses. Finally, the chapter notes that one consequence of markets for technology is the possible encouragement of a greater "privatization" of knowledge. This may weaken norms of conduct commonly associated with academic research and undercut the important role of academic research based on open disclosure and information sharing in the generation and the diffusion of knowledge. Chapter 11 concludes the book, summarizing its main issues and discussing further developments in this line of research.