

Economists have, of course, always recognized the dominant role that increasingly knowledge plays in economic processes but have, for the most part, found the whole subject of knowledge too slippery to handle.

—E. Penrose, *The Theory of the Growth of the Firm*

The economics of knowledge as a discipline should not be confused with the economics of research, for its main focus is not the formal production of technological knowledge; nor should it be seen as the economics of innovation, for it is not centered exclusively on the study of the conditions, modalities, and effects of technological and organizational change. It should also not be likened to the economics of information, since the object of the economics of knowledge is knowledge (and not information) as an economic good. Its field of analysis covers the properties of that economic good governing its production and reproduction as well as the historical and institutional conditions (such as information technology or patent rights) determining its treatment and processing in a decentralized economy.

Scope of the Economics of Knowledge

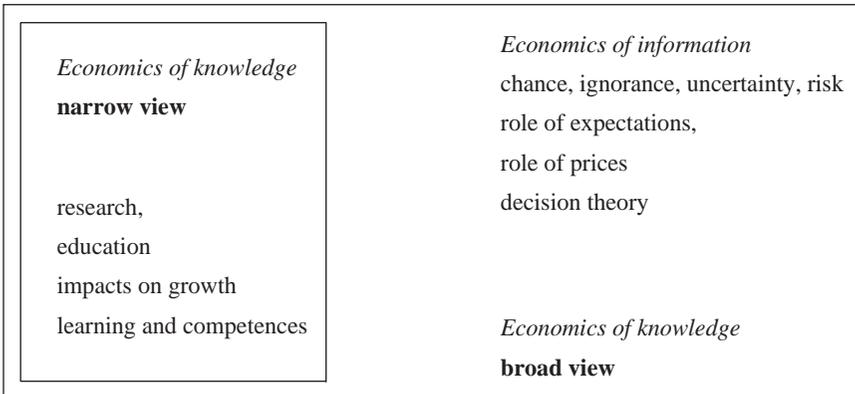
Some Modern Precursors

Apart from historical figures—Smith, Marx, and Schumpeter who all dealt with knowledge, its creation and division, its use and appropriation—the latter-day pioneers in the general economics of knowledge (i.e., not confined to science and technology) are unquestionably Simon, Hayek, Arrow, and Machlup. Simon (1982) has studied numerous subjects pertaining to the economics of knowledge, such as the role of memorization in the learning process, and can be considered as the

real precursor of the economics of information technology. Hayek (1945) examined problems posed by the mass dissemination of knowledge and the impossibility of transferring knowledge to a central planning agency. Arrow, in two seminal articles published in the same year (1962a, b), developed the economics of knowledge creation that was to lay the foundations for two main strands of research (on problems of allocating resources to the creation of knowledge, defined as a public good, and on endogenous technical change).

Machlup's work covers a vast domain. Its dimensions are the consequence of an extremely broad conception of the economics of knowledge, encompassing the economics of information, in particular, and consequently theoretical problems of decision making. Defining information as "a certain type of knowledge," Machlup (1984) is naturally led to extend the economics of knowledge to include not only an analysis of information sectors and industry, an examination of the production of new knowledge, and a study of mechanisms of skills acquisition and transfer, but also an exploration of the vast domain of economic theory of choices and expectations in situations of uncertainty and incomplete information. In this respect Machlup's approach is similar to that of Hayek who uses the terms *knowledge* and *information* interchangeably, especially when studying the role of the pricing system as a mechanism in the communication of information. For Richardson (1960) the problem is similarly that of the availability of technological information for improving the coordination of activities in the market. All these authors see human decision making as being at the heart of economics, and the presence or absence of knowledge and information as factors that crucially determine the conditions in which decisions are made. There is no real difference between knowledge and information, which means that the scope of the economics of knowledge is defined very broadly (a quick look at the seventeen subject groups listed by Machlup (1984, chap. 10) gives an idea of just how broad it is).

A more restrictive conception of the economics of knowledge excludes problems of economic choice in situations of incomplete and uncertain information and focuses more specifically on what I would call "expertise"—namely, knowledge. Here, knowledge is above all a cognitive capacity, which is what distinguishes it clearly from information. This conception was developed in France, in particular, by J. L. Maunoury whose book *Economie du Savoir*, published in 1972, was unquestionably the precursor. Maunoury focused essentially on the system of production and acquisition of knowledge, of which

**Figure 1.1**

The scope of the discipline (narrow and broad view)

research and education are the two mainstays, and on the relationship between this system and economic growth.

Choosing between these two conceptions is difficult. Finding one's way between the very broad definition of the economics of knowledge, encompassing the economics of information and theory of choice, and the narrower definition consisting essentially of analyzing education and research, is no simple matter—especially since the economics of knowledge in a narrow sense has expanded since Maunoury's day. It now includes not only deliberate forms of knowledge production and acquisition, corresponding to the main education and research institutions, but also the vast domain of learning processes that describe increasingly numerous situations in which expertise is produced in the framework of "regular" production and use of goods and services. By extension, this economics of knowledge encompasses the notion of competence and the capacity to learn (Garrouste 2001).

The definition of the scope of the discipline (figure 1.1) depends on one's conception of knowledge and information, which I now consider more closely.

Exploring the Black Box of Knowledge

For a long time economic analysis equated knowledge with information. Based on this amalgam, economic analysis adopts a particular approach to knowledge information—namely, the universe can be described by a finite (but very large) set of states to which probabilities

can be assigned (Laffont 1989). Knowledge improves when the probability of a particular state is estimated more accurately. Knowledge can therefore be expressed by a vector of probabilities relating to a predetermined set of states. Of course there is a huge practical advantage in adopting this type of approach, but it still does not enable economists to grasp phenomena as important as learning and cognition.

In my conception, knowledge has something more than information: knowledge—in whatever field—empowers its possessors with the capacity for intellectual or physical action. What I mean by knowledge is fundamentally a matter of cognitive capability. Information, on the other hand, takes the shape of structured and formatted data that remain passive and inert until used by those with the knowledge needed to interpret and process them. The full meaning of this distinction becomes clear when one looks into the conditions governing the reproduction of knowledge and information. While the cost of replicating information amounts to no more than the price of making copies (i.e., next to nothing, thanks to modern technology), reproducing knowledge is a far more expensive process because cognitive capabilities are not easy to articulate explicitly or to transfer to others: “we can know more than we can tell” (Polanyi 1966, 4). Knowledge reproduction has therefore long hinged on the “master-apprentice” system (where a young person’s capacity is molded by watching, listening, and imitating) or on interpersonal transactions among members of the same profession or community of practice. These means of reproducing knowledge may remain at the heart of many professions and traditions, but they can easily fail to operate when social ties unravel, when contact is broken between older and younger generations, and when professional communities lose their capacity to act in stabilizing, preserving, and transmitting knowledge. In such cases, reproduction grinds to a halt and the knowledge in question is in imminent danger of being lost and forgotten.

Therefore, the reproduction of knowledge and the reproduction of information are clearly different phenomena. While one takes place through learning, the other takes place simply through duplication. Mobilization of a cognitive resource is always necessary for the reproduction of knowledge, while information can be reproduced by a photocopy machine.

As observed by Steinmueller (2002a), by failing to differentiate between knowledge and information, economics—a discipline that often has an imperialistic attitude toward the other social sciences—

has, quite surprisingly, left a vast field open to other disciplines. This field consists of the subjects “learning” and “cognition,” two central themes in my conception of knowledge.

A further complication is the fact that knowledge can be codified—so articulated and clarified that it can be expressed in a particular language and recorded on a particular medium. Codification involves the exteriorization of memory (Favereau 2001). It hinges on a range of increasingly complex actions such as using a natural language to write a cooking recipe, applying industrial design techniques to draft a scale drawing of a piece of machinery, creating an expert system from the formalized rules of inference underlying the sequence of stages geared to problem solving, and so on. As such, knowledge is detached from the individual, and the memory and communication capacity created is made independent of human beings (as long as the medium upon which the knowledge is stored is safeguarded and the language in which it is expressed is remembered). Learning programs are then produced that partially replace the person who holds and teaches knowledge.

When knowledge is differentiated from information, economic problems relating to the two can be distinguished. Where knowledge is concerned, the main economic problem is its reproduction (problem of learning), while the reproduction of information poses no real problem (the marginal cost of reproduction is close to nothing). The economic problem of information is essentially its protection and disclosure, that is, a problem of public goods. However, the codification of knowledge creates an ambiguous good. This good has certain properties of information (public good) but its reproduction as knowledge requires the mobilization of cognitive resources.

Example of a Paradox Resolved by the Distinction between Knowledge and Information Paradoxically, in view of the enormous advances in information and communication technologies, many trades and professions are experiencing a crisis regarding the transmission of expertise and knowledge, both vertically between masters and apprentices, and horizontally between experienced practitioners (see OECD 1999a on the case of education and health). The paradox disappears, however, when one distinguishes between problems of reproduction of knowledge and those of transmission of information. Moreover, in all the occupations concerned it is shortcomings in the social networks (which previously played the part of transmitting and

building on expertise—see chapter 4) that explain these problems of reproduction of knowledge, and there is nothing obvious about offsetting these weaknesses with new communication technologies.

What Is the Meaning of the French (German and Spanish) Distinction between *Connaissance* (“Kenntnis,” “conocer”) and *Savoir* (“Wissen,” “saber”)? The French language offers a distinction between *savoir* and *connaissance* that has no real equivalent in English, though it can be conveyed by adding the qualifier *certified*. Certified knowledge (“*savoir*”) means knowledge that has been legitimized by some institutional mechanism (be it scientific peer review or any kind of rituals and belief systems in oral societies). Other forms of knowledge (“*connaissance*”) also enable action (knowing how to do the gardening) but have not been put through the same tests as certified knowledge. What separates the two has less to do with a contrast between the scientific and nonscientific than with whether or not the knowledge has been subjected to institutional testing: “gardening knowledge” is reliable, wide-ranging, and relatively decontextualized, but each gardener has his or her own local (and locality-specific) knowledge. Yet the economics of knowledge does not preclude either form, meaning that it is not devoted solely to the analysis of formal production of “certified knowledge.”

Narrowing the Scope of the Discipline In view of this conception of knowledge and information, I now turn away from the economics of information and decision theory and focus essentially on knowledge in the strict sense of the word (as a cognitive capacity). I am thus opting for a narrow conception of the economics of knowledge, although the field I wish to study—research, learning processes, positive externalities, problems of coordination of innovative activities, and codified and tacit knowledge—is vast and covers many areas as yet unexplored.

Economists’ Difficulties Concerning the Economics of Knowledge

Categories that No Longer Fit

To apprehend knowledge, economists constructed a “comfortable world” in which only some agents, institutions, and sectors were specialized in the production of knowledge. R&D laboratories at the corporate level and “knowledge industries” at the level of the economy

were the main categories of a “world” that excluded a large part of all activities and agents, considered not to be stakeholders in the economics of knowledge.

Economists’ Comfortable World With regard to innovation in enterprise, economists reduced knowledge production to the function of R&D, defined as the activity specifically devoted to invention and innovation. This representation can be credited with the generation of the immensely useful and extensive collection of data and production of statistics at the international level. But analysis of R&D covers only a small part of all innovation and knowledge production activities.

Economists similarly delimited a number of sectors in the economy, specifically devoted to the production and manipulation of knowledge and information. Machlup (1962), rightly seen as the founder of this tradition, studied the economic importance of the knowledge-based economy, identified as a specialized sector and consisting primarily of activities relating to communication, education, the media, and computing and information-related services. This statistical frame of analysis generated abundant research, commissioned mainly by the OECD. Despite significant methodological variations, all these studies were grounded in the same basic logic of defining a specialized sector covering all activities related to the production and processing of information.

Representations were therefore produced to deal with problems of indicators and quantification on the basis of stabilized information and knowledge and skillfully used measurement tools. But the price to pay is high: representations formed in this way fail, to a large extent, to explain knowledge-based economies.

From R&D to Learning Processes Of course all knowledge produced in a firm cannot be attributed to formal research activities. Depending on the sector and the firm, the share of formal research in knowledge production can range from “huge” to “minute.” Other major activities can also play a part.

First, design and engineering play an important role in the growth of knowledge. This role has been clearly identified by Vicenti (1990) who shows that design is an essential locus for the autonomous production of knowledge. The articulation between research and design then raises a series of important questions, since the idea of

autonomous production of knowledge implies that design and technology are not subordinate to science and R&D (they are not “applied science”).

Second, any activity involving the production or use of a good (or service) can generate learning and hence knowledge production. In other words, in many activities knowledge production is not the goal but may nevertheless occur. Knowledge is a by-product of the activity of production or use. This is where we find the well-known forms of “learning by doing” and “learning by using,” concepts formulated by Arrow (1962b) and Rosenberg (1982), respectively. These studies progressively revealed that this type of learning process occupies an essential place in the economics of knowledge. It became more and more evident that certain types of strongly “motivated” and explicitly cognitive learning had economic effects that could go much further than just the consequence of doing one’s job better by repeating the same actions. But measuring knowledge produced by learning is difficult.

From Specialized Sectors to the Entire Economy Eliasson (1990) developed an important innovation when he broke away from Machlup’s tradition which basically defined a specialized sector encompassing all activities related to the production and processing of knowledge, and measured its contribution to the gross domestic product (GDP). Eliasson considers that knowledge production and information processing are located in all economic activities, including in low technology-intensive sectors. In other words, the advent of the knowledge-based economy is manifested less in the continuous expansion of a specialized sector than in the proliferation of knowledge-intensive activities throughout all sectors of the economy. But here again measurement is complicated. I return to this point in the next chapter.

Both approaches—the analysis of either a specialized sector or of the generalization of knowledge-intensive activities throughout the entire economy—have their pros and cons. However, using the former exclusively may produce serious policy failures. For example, if the conclusion that net job creation takes place only in knowledge-intensive parts of the economy were interpreted in the framework of an approach that reduced the knowledge-based economy to a specialized sector, it could lead to bad choices being made in education policy.

Unobservable Phenomena and Problems of Measurement

Yet traditional categories—R&D in the corporate world and the information sectors in the national economy, which, as mentioned earlier, could not contain all knowledge-producing activities—had a big advantage. They provided a way of measuring by facilitating the identification of knowledge-intensive activities. This in itself is ample justification for the category, because most phenomena relating to knowledge are largely unmeasurable. Apart from the question of the definition of knowledge, mentioned earlier, the main problems involved in measurement include:

1. Elements of knowledge are heterogenous. No comparison can possibly be made between the invention of writing and the discovery of a new distant star.
2. Knowledge is largely unobservable. The observation of knowledge (and especially tacit knowledge; see chapter 4) seems simply impossible. The most distinctive feature of tacit knowledge is its incorporation in thoughts and deeds, and its invisibility, even for those who possess it and use it “automatically.” Knowledge appears only when it is expressed and written and when it becomes possible to attach a property right to it. Yet tacit knowledge is constantly being reconstituted, so that a vast world remains perpetually invisible.
3. There is no stable model that can be used to convert inputs (into the creation of knowledge) and outputs (economic effects). There is no stable formula such as the one used *ceteris paribus* to link an increase in the quantity of steel to growth in car production. Knowledge, unlike classic capital goods, has no fixed capacity in terms of impact of an additional quantity on the economy. Depending on the prevailing spirit of initiative, the situation of competition or the social organization, a new idea can trigger huge change or have no effect (see Quah 1999, who thus explains China’s technological stagnation from the fourteenth century onwards). Thus, there is no production function that can be used to forecast, even approximately, the effect that a unit of knowledge will have on economic performance. Conversely, it is very difficult to impute an economic effect to particular knowledge. Effects of externality and cumulativity do not make it possible to identify with any certainty an element of knowledge as being behind a particular improvement in the economy. Or else that imputation is at a very general level (e.g., “information technology is at the origin of a particular effect on the economy”).

4. Finally, measuring stocks, already difficult in the case of physical capital, becomes an impossible undertaking in the case of knowledge. How could the composition of a stock be defined? What should be selected or rejected in this vast domain encompassing practical, intellectual, and spiritual knowledge: knowledge of perpetual value and significance, and knowledge of fleeting importance; knowledge which is important for many and that which is valued by very few?

Moreover, from a theoretical point of view, serious problems of additivity appear when we want to measure the stock of an entire society (or social group). In the economy of tangible goods, this problem of addition is governed by laws which link a prototype to various scales of mass-produced products, or an original to a small or large number of copies. But knowledge defies both of these laws. There is neither a prototype nor an original, so that the notion of an additional unit is meaningless. It is as if one were trying to measure a stock of flames. Each neighbor can take fire from the others without reducing the size of the fire of the person who had it first. Thus, in a sense, when knowledge appears it is potentially available to all. There is no difference between the situation in which one theorem of Pythagoras exists and one in which a billion such theorems exist. Yet we cannot consider that anyone in the world has the means or opportunity to have access to this element of knowledge. It is knowledge that is useful to some, useless to others, and an impenetrable mystery to others still. We thus arrive at the notion of the absorptive capacity (or learning capacity) of a society, the importance of which is variable for each type of knowledge and probably brings us closer to the measurement of stock.

Finally, the depreciation of knowledge is governed by a wide variety of "laws" (forgetfulness, obsolescence), and it seems that no one rule can adequately account for it (Machlup 1984).

It is possible to observe and measure the resources allocated to knowledge production activities (primarily R&D spending) as well as the results of these activities expressed either in the form of specific outputs (patents, publications, software, new products) or of economic variables, thought to be related to the production of new knowledge. The difficulties mentioned earlier disappear to a large extent when we measure contributions to knowledge (R&D, human resources, patents, and publications) and the product of knowledge (social and private outputs, innovation).

Table 1.1

Framework for Indicators in the Economics of Knowledge; Application to the Health Sector

Category	Concepts	Indicators
Inputs	Person-years, equipment-years	Expenditures
	Organizational capacity	Use of particular organizational practices
Outputs	Ideas, discoveries	Papers, prizes
	New products	Patents, new drug applications
Outcomes or impacts	Broad advance of human knowledge	Papers, citations, expert evaluations
	Improvements in health status and length of life	Outcome studies, life expectancy
	Reduction in healthcare expenditures	Outcome studies, statistical analyses of healthcare expenditures
	Economic output	Revenue growth, revenue from new products, profitability
	Productivity improvements	Productivity studies

Source: Jaffe (1999).

But this is proximation, which does not directly measure knowledge. Very recent and extremely sophisticated studies have therefore tried to measure flows of knowledge (Jaffe and Trajtenberg 1996) or even the degree to which certain knowledge is fundamental (Henderson, Jaffe, and Trajtenberg 1998). In order to do so the authors use what is observable, namely, patents and citations.

These indicators, summarized in table 1.1 for the health sector, are therefore necessary. Yet they illuminate only a small fraction of all economic activity in a sector producing or exploiting new knowledge.

For many sectors (e.g., education), the part of the economics of knowledge that remains unknown is far greater than the part that is known. That is generally the case with sectors in which R&D plays a relatively small role compared to multiple learning experiences that are difficult to grasp. As A. Carter (1996) put it, the indicators in table 1.1 basically shed light on the tip of the iceberg only. That is why use and interpretation of these indicators for exploring and measuring the economics of knowledge always require the economist to have a certain degree of faith.

But if we cannot measure knowledge itself, why not add up the values of knowledge-related transactions (the method usually applied in many cases when there is no clearly defined unit of output)? Unfortunately, our market institutions face daunting problems when a price has to be set for knowledge. The reasons are interesting:

- the seller—by selling knowledge—does not lose anything; knowledge is acquired definitively, even if it is shared or sold afterward;
- the buyer does not need to buy the same knowledge several times, even if it is to be used several times;
- the buyer cannot really assess the value of knowledge without actually acquiring it.

For these reasons (the first two of which will be considered in chapter 5 because they express the “nonrival” property of knowledge) the prices fixed are unique and specific and can never be used as consistent and reliable indicators. Insofar as prices have to be determined, they can vary widely from one transaction to the next. A huge proportion of knowledge is not traded in the framework of monetary transactions; it is accumulated in firms, other organizations, and actor networks without any value being attributed to it.

Modeling Knowledge

It is toward growth models that endogenize technological change that we naturally turn to evaluate the capacity of neoclassical theory to solve problems of the economics of knowledge. Two aspects of the modelling of endogenous growth are relevant here. First, in these models firms benefit from R&D investments because they are able to control at least part of the resulting productivity growth or product improvement. Second, markets are assumed not to be perfectly competitive. This makes it possible to obtain a market equilibrium in conditions of increasing returns (generated by the production of knowledge; see chapter 3). The endogenization of technological progress in these models was completed by the construction or deduction of other phenomena, for example, creative destruction that captures the process of depreciation of older technologies when new ones appear; or externalities derived from R&D and education. Finally, in many of these models the rate of investment in new plants and equipment affects the regularity of the growth rate. These models therefore afford many angles from which to study why and how growth rates

differ in time and from one country to the next (see Aghion and Howitt 1998, for an overview).

Of course this short presentation hardly does justice to the richness of this research. It helps, however, to show that these studies served to bring formal theoretical work on economic growth closer to what Abramovitz (1989) called the immediate determinants of growth. Yet many other aspects of the economics of knowledge, of the utmost importance in explaining the determinants of growth, are still overlooked or considered only superficially. Nelson (1994) identifies three other issues:

1. Knowledge itself, the vehicle of externalities, is always represented in models of endogenous growth in the form of a written expression, a manual, a computer program, in short, a set of codified instructions which provide access to immediate and free exploitation of the technology. This is of course a huge simplification, with disastrous consequences on our understanding of knowledge-based economies (Dosi 1996). A large share of knowledge does not appear in the form of codified instructions; it is tacit and naturally excludable, which sharply reduces the dimension of externalities.
2. The firm remains a black box. Given the public knowledge infrastructure and the opportunities to invest in private technologies, firms choose their strategies to maximize profits, taking into account market conditions. But mastery of a new technology or new knowledge is an extremely complex process that each firm will succeed in to a greater or lesser degree, depending on its organization and forms of management and strategy. Economists of innovation, as well as specialists of corporate history and management, use the term *corporate capability* or *corporate competence* to convey these different aspects (Dosi, Teece, and Winter 1992). Yet very few economists of endogenous growth seem prepared to take into consideration the diversity in firms' capacities to innovate as a key element explaining economic growth.
3. Finally, the corporate environment, apart from the market, plays an essential role which, once again, is seldom recognized in endogenous growth models. Many aspects of that environment are determining factors in economic growth, including relations with universities, the quality of the intellectual property rights (IPR) system or of the functioning of the financial market, and laws governing the labor market. The concept of a national innovation system (Carlsson and Stankiewicz 1991; Foray and Freeman 1992; Lundvall 1992; Nelson 1993; Edquist

1997) helps to explain those clusters of institutions which, at the national level, strongly influence firms' innovation strategies and performance. This concept is more relevant than ever at a time of knowledge-based economies.

This brief review is intended primarily to highlight the importance, for economic research, of constant dialogue and mutual attentiveness between the formal theory of growth and what is called *appreciative theories*. Recent work by Keely and Quah (1998) on formal theory shows just how fruitful such dialogue can be.

Economic Issues

In order to understand better the "economics of knowledge" I broadly outline the general problems of the discipline. It starts with the analysis of the peculiar properties of knowledge as an economic good and proceeds to the normative analysis of resource allocation mechanisms in the field of knowledge production and distribution and, more generally, socioeconomic institutions that can be relied upon to produce, mediate, and use knowledge efficiently.

I simply point out some features which are problematic, not only because they make it difficult to observe and measure knowledge but because they complicate the issues of building efficient mechanisms of resource allocation in both static and dynamic worlds.

Knowledge Creation

New Knowledge Stems from a Discovery or Invention Much knowledge is produced by invention, that is, it does not exist as such in nature and is "produced" by man. Other types of knowledge stem from discoveries, that is, the accurate recognition of something which already existed but which was concealed. Invention is the result of production; discovery the result of revealing. This distinction, although it may seem vague in many cases (the hammer is an invention but the use of the first hammer, an appropriately shaped stone, was probably a discovery), has many implications for the economics of knowledge. In terms of incentives, one can claim an intellectual property right on an invention, not a discovery. One can patent a new machine but one cannot patent a fresh water spring even if one has "discovered" it. As a result, recurrent debate on the nature of novelty in certain disciplines

such as mathematics—is it an invention or a discovery?—has extensive economic implications.

This distinction is also important in terms of the mode of development of knowledge. If knowledge stems from successive discoveries, there must be constants in research activities—like explorers who discover the same land and write different accounts about it containing common points. If, on the other hand, it stems from inventions, one can expect noteworthy differences, even if a number of socioeconomic forces lead toward the convergence of inventions.

Knowledge Is Often a Joint Product Knowledge is very often produced in a context of activities in which other motivations (the manufacturing of a good or the provision of a service) are predominant. People learn by doing or by using (chapter 3). There is learning-by-doing or learning-by-using because knowledge is not absolute but must be defined in relation to a specific physical context (Tyre and von Hippel 1997). Such a characteristic gives many activities an important potential value in terms of knowledge production and innovation: those activities related, for instance, to the introduction of a novel type of equipment, organization, or method. There are, however, inherent limitations to the production of knowledge in this kind of context. Constraints and limitations are due to the basic tension and conflict between the “doing” aspect (the performance to be achieved at the end of the day) and the “learning” aspect (the experiment that is carried out as a consequence of “doing”). Maximizing learning benefits implies tolerating a certain degree of deterioration of static efficiency.

On Some Properties that Magnify the Social Benefits of Knowledge Creation

Knowledge Is Partially Nonexcludable and Nonrival These properties are investigated in depth in chapter 5 and their welfare economic aspects are discussed in chapter 6. At this point, it is enough to say that making knowledge exclusive and controlling it privately are difficult and costly. Knowledge continuously escapes from the entities producing it. Second, knowledge is nonrival, meaning that economic agents are not rival users of knowledge. Knowledge can theoretically be used by a million people at no additional cost because its use by an additional agent does not imply the production of an additional copy of

that knowledge. This characteristic is a form of nonconvexity or an extreme form of decreasing marginal costs as the scale of use is increased. The aforementioned properties define what is meant by a pure public good and, as such, create a difference between the private and the social return in the domain of knowledge production. Recipients of knowledge largely extend beyond those who have produced it and can be multiplied ad infinitum, both geographically, in space, and historically, in time.

Knowledge Is (Often) Cumulative Those external benefits can be made even stronger in the case of “cumulative” knowledge. It is the attribute of cumulativeness that distinguishes knowledge as consumption capital (enabling people to undertake “final” action: I know how to garden; I know how to paint) and knowledge as an intellectual input (enabling people to create new knowledge and thus to broaden the spectrum of possible future actions). Most knowledge in mathematics is cumulative because it may give rise to new ideas and open new lines of research.

The “Comedy of the Knowledge Commons” Owing to these three properties, the production of knowledge has the potential to create a “combinatorial explosion.” This is a good which is difficult to control and which can be used infinitely, to produce other knowledge which in turn is nonexcludable, nonrival and cumulative, and so on. In many cases knowledge is also deliberately disclosed and organized in order to facilitate its access and reproduction by others. All these processes give rise to the creation and expansion of “knowledge commons.” “Knowledge commons” are not subject to the classic tragedy of commons that describes the case where exhaustible resources (such as a pasture or a shoal of fish) are subject to destruction by unregulated access and exploitation (see chapter 8). Knowledge may be used concurrently by many, without diminishing its availability to any of the users, and will not become “depleted” through intensive use. As Paul David writes (2001, 56), “Knowledge is not like forage, depleted by use for consumption; knowledge is not subject to being “overgrazed” but instead is likely to be enriched and rendered more accurate the more researchers, engineers or craft workers are allowed to comb through it.” The properties of nonexcludability, nonrivalry, and cumulativeness have features akin to quasiinfinite increasing returns. Thus, the commons is not tragic, but comedic, in the classical sense of a story

with a happy outcome (Rose 1986). Managing the “knowledge commons” requires social regulations that are entirely different to the social arrangements used to regulate ecological systems of exhaustible resources.

On Some Properties of Knowledge that Impede the Full Realization of Social Benefits

Social benefits stemming from the full exploitation of the “knowledge commons” are, however, neither obvious nor automatic. New knowledge is most often partially localized and weakly persistent, tacit and sticky, dispersed and divided.

Knowledge Is Partially Localized and Weakly Persistent Apart from strategic choices of private agents who are inclined to impose exclusivity on their knowledge (through secrecy and intellectual property rights), new knowledge is most often not of general value for the economy because it has been produced in a local context for particular purposes. A large body of literature argues that the production of knowledge is at least partially localized: learning that improves one technology may have little or no effect on other technologies (Atkinson and Stiglitz 1969; Antonelli 1999, 2001). The process toward generalization of knowledge is a very difficult one. It involves, for instance, the creation of theoretical knowledge that can fit in many local situations, or the search for analogic links among fields and disciplines, or the identification of similarities between the professional knowledge of various occupations. However, the degree of standardization and maturation of technology and knowledge can mitigate these difficulties (Cowan et al. 2002).

Moreover, knowledge is weakly persistent. Evidence in the psychological literature show that people forget. If the practice of a task is interrupted, forgetting occurs. Hirsch (1952) found that when performance was resumed after an interruption it was lower than the level achieved prior to the interruption. Moreover, knowledge can be depreciated (through deterioration and obsolescence). Communities that are in possession of it can break up, resulting in the disintegration of their collective knowledge.

New Knowledge Is Tacit and “Sticky” Typically, new knowledge and expertise have a broad tacit dimension, meaning that they are neither articulated nor codified. Tacit knowledge resides in people,

institutions, or routines. Tacitness makes knowledge difficult to transport, memorize, recombine, and learn. Such difficulties can be overcome when the number of people possessing the tacit knowledge is high. In this case there will be a labor market that can be used to transport and transfer tacit knowledge. If the number of people is too small, tacitness increases the risk of “accidental uninvention” and hampers the full exploitation of knowledge. Given tacitness, knowledge is costly to transfer from one site to another in useable form. As von Hippel put it, knowledge is sticky (von Hippel 1994). Stickiness raises a number of issues in terms of the organization of knowledge production, product design, and system integration.

Knowledge Is Dispersed and Divided There is a natural tendency for knowledge to fragment as it becomes subject to more in-depth division and dispersion (Machlup 1984). The division of knowledge stems from divisions of labor and increasing specialization in the field of knowledge production. Its dispersion is related to local situations in which knowledge is produced (a site, a workshop, a laboratory). The result is an extremely fragmented knowledge base, which makes it difficult to form a broad and integrated view of things. This can have disastrous consequences. At the level of global policy making, knowledge that can help resolve a particular problem may exist without being “visible.” It can go unnoticed by the decision maker. Knowledge of the greenhouse effect, for instance, has been in the public domain since 1886, thanks to the study by Svante Arrhenius, but failed to capture the attention of the political system for another hundred years. There is a big difference between the existence of knowledge in some or other place, and its availability to the right people in the right place at the right time. The crux of the matter is knowing how to integrate and organize fragmented, scattered, and thinly spread knowledge.

Conclusion: The Aim of the Economics of Knowledge

The aim of the economics of knowledge is thus to analyze and discuss institutions, technologies, and social regulations that can facilitate the efficient production and use of knowledge. Given the peculiar properties and features of knowledge as an economic good, most of the usual resource allocation mechanisms used in the world of tangible goods do not work properly to maximize knowledge creation and diffusion. In this perspective, the most important institutions are of two kinds: those

which enable economic agents to appropriate the fruit of their intellectual creation, and those that make it possible to preserve, consolidate, and exploit “knowledge commons.” The complexity of the institutional problem derives from the fact that these two objectives are both contradictory and indissociable. Moreover, depending on the nature of the knowledge, the “optimum solution” could vary widely. Thus, this problem is addressed differently in relation to the following three categories:

- knowledge is reducible to “consumption capital” (Machlup 1984)
- it constitutes productive capital (notion of cumulativeness presented earlier)
- it represents a piece of strategic information (notion of “aforeknowledge” developed by Hirshleifer 1971).

In this set of questions, only one agent knows that a particular event is going to occur and that it will change the structure of prices; he can therefore speculate on a given factor. For example, I know that an epidemic is likely to wipe out the entire bee population, so I stock honey.

This question is studied at length in chapter 6. The goal of the economics of knowledge is therefore to develop a framework in order to devise and compare socioeconomic institutions that can be relied upon to create and exploit knowledge in an efficient manner; that is, institutions that can sustain an efficient production and allocation of knowledge of all kinds.