# 3

# *Economic and Sociological Explanations of Technological Change*

This chapter seeks to identify tools to overcome the cleavage between economic and sociological analyses of technological change. It draws on the tradition of "alternative economics" deriving from Herbert Simon. A more implicit debt is to Marx's critique of political economy, and an explicit, but of necessity highly tentative, attempt is made to argue that the sociology of scientific knowledge might be brought to bear on the economist's discussion of the unmeasurable uncertainty (rather than quantifiable risk) of technological change.

I am painfully aware of many places where I shall stray into areas where I am ignorant. There may well be answers to the questions I ask and a relevant literature of which I am unaware. It may be that, as a sociologist, I have misunderstood what economists mean. In some places I suspect, though I am not certain, that I am calling for the bringing of coals to Newcastle. If any of this is true, I would be most grateful for both pardon and enlightenment. Unless we take the risk of revealing our ignorance, interdisciplinary bridges will not be built.

In studies of technology, the gap between economic and sociological explanations is pervasive. Economic analyses are often based upon assumptions sociologists regard as absurd, while sociological writing often almost ignores the dimension of cost and profit in its subject matter. Though there are thinkers who provide rich resources for transcending the gap (despite their considerable differences, Karl Marx and Herbert Simon are the two central ones), it is far more common to find economic and sociological studies, even of the same topic, existing in separate conceptual universes.[1]

In the first section of the chapter I contrast neoclassical economics, particularly its assumption of profit maximization, with the alternative economics associated with Simon and more recently developed by Richard Nelson and Sidney Winter. I then go on to discuss possible applications of that alternative view to a false dichotomy sometimes found in labor-process studies, to pricing behavior in the computing industry, and to the setting of research and development budgets.

Next I examine the idea of a "technological trajectory" or a "natural trajectory" of technology, found in the work of Nelson and Winter and other recent contributors to the economics of technology. I argue that, although persistent patterns of technological change do exist, there is a crucial ambiguity in their description as "natural," and that a different understanding of them would help bridge the gap between economic and sociological explanations.

In the next section I discuss another way of bridging the gap, one again loosely in the tradition of Simon but in practice little pursued: the "ethnoaccountancy" of technological change (that is, the empirical study of how people actually reckon financially about technology, as distinct from how economic theory suggests they should reckon).

Finally, I turn to the topic of uncertainty and the construction of the economic. Despite their "thing-like" character, economic relations are never wholly self-sustaining and self-explaining. Whereas this point is normally argued in the large (Marx justifies it by an examination of the evolution of capitalism), technological innovation demonstrates it on a smaller scale. As is well known, the inherent uncertainty of radical innovation makes economic calculation applicable only *ex post*, not *ex ante*—that is, once networks have stabilized, not before. This makes radical innovation a problem for orthodox economics, but it points, I argue, to the relevance here of the sociology of scientific knowledge.

### *Neoclassical and Alternative Economics*

It is convenient to begin with our feet firmly on the economic side of the gap. The neoclassical economics of production technology is crystalline in its explanations. Although the full neoclassical structure is dauntingly complex, its central pivot is simple and clear: firms choose production technology so as to maximize their rate of profit.

Unfortunately, that clarity is purchased at too high a price. The notion of maximization at the heart of the neoclassical structure is incoherent, at least as a description of how firms do, or even could, behave. Perhaps the most cogent statement of why this is so comes from Sidney Winter:

> It does not pay, in terms of viability or of realized profits, to pay a price for information on unchanging aspects of the environment. It does not pay to review constantly decisions which require no review. These precepts do not imply merely that information costs must be considered in the definition of profits. For without observing the environment, or reviewing the decision, there is no way of knowing whether the environment is changing or the decision requires review. It might be argued that a determined profit maximizer would adopt the organization form which calls for observing those things that it is profitable to observe at the times when it is profitable to observe them: the simple reply is that this choice of a profit maximizing information structure itself requires information, and it is not apparent how the aspiring profit maximizer acquires this information, or what guarantees that he does not pay an excessive price for it.[2]



This critique of neoclassical economics draws most importantly upon the work of Herbert Simon. It has been elaborated by Winter, by his collaborator Richard Nelson, and by a goodly number of other economists. Its logic seems inescapable.[3] Furthermore, Simon and his intellectual descendants do not simply highlight the central incoherence haunting neoclassical economics' formidable apparatus of production functions, isoquants, and the like. They provide a different vision of economic activity. In this alternative economics, actors follow routines, recipes, and rules of thumb while monitoring a small number of feedback variables. As long as the values of these variables are satisfactory ("satisficing" is Simon's famous replacement for "maximizing"), the routines continue to be followed. Only if they become unsatisfactory will they be reviewed. But the review will not be an unconstrained evaluation of the full universe of alternatives in search of the best; it will be a local search, given direction by the perceived problem in need of remedy and using heuristics (which are rather like routines for searching).

This intellectual tool kit offers a bridge toward sociological analysis as it is conventionally understood. Routines can be entrenched for a variety of organizational reasons, and different parts of a firm typically follow different routines and different heuristics of search. Since in this perspective there is no longer any ultimate arbiter of routines (such as profit maximization), firms become political coalitions rather than unitary rational decision makers. The actual behavior of a firm may represent a compromise between different and potentially contending courses of action.[4]

Intrafirm processes are not, of course, ultimately insulated from what goes on outside the firm. That outside is a "selection environment" favoring certain routines over others. Nelson and Winter, especially, draw an explicit parallel with evolutionary biology, seeing routines as akin to genes, being selected for or against by their environment. This environment is not just "the market"; it includes other institutional structures as well. It is not necessarily or even generally stable, nor is it simply external and "given." One particular firm may be able to alter its environment only slightly (although some firms patently alter it more than slightly), but the behavior of the ensemble of firms is in large part what constitutes the environment.[5]

This "alternative economics" promotes a subtle change in ways of thinking, even in areas where its relevance is not apparent. Take, for example, David Noble's justifiably celebrated, empirically rich study of the automation of machine tools in the United States. Noble frames his most general conclusion in terms of a dichotomy between profit and capitalists' control over the work force:

> It is a common confusion, especially on the part of those trained in or unduly influenced by formal economics (liberal and Marxist alike), that capitalism is a system of profit-motivated, efficient production. This is not true, nor has it ever been. If the drive to maximize profits, through private ownership and control over the process of production, has served historically as the primary means of capitalist development, it has never been the end of that development. The goal has always been domination (and the power and privileges that go with it) and the preservation of domination.[6]

This analytical prioritization of the sociological[7] over the economic cannot be correct: a firm or an industrial sector that pursued control at the expense of profit would, unless protected from competition, shrink or die. Much of the American industrial sector studied by Noble did indeed suffer this fate, in the period subsequent to the one he examined, at the hands of the Japanese machine-tool manufacturers, who were equally capitalist but who, in their organizational and technological choices, were less concerned with control over the work force. Arguably it was only the protection offered by military funding (a factor to which Noble rightly gives considerable emphasis) that allowed American machine-tool manufacturers to follow the technological strategy they did.

The temptation to counterpose profit and domination, or economics and sociology, arises, I would suggest, from the way our image of economics is permeated by neoclassical assumptions. The alternative economics associated with Simon allows us to make analytical sense of capitalists who are profit oriented (as any sensible view of capitalists must surely see them) without being profit maximizers. The urge to achieve and maintain control over the work force is not an overarching imperative of domination, overriding the profit motive; it is a "heuristic"[8] with



deep roots in the antagonistic social relations of capitalist society. When facing technological choices, American engineers and managers, in the period studied by Noble, often simplified production technology decisions by relying on an entrenched preference for technological solutions that undercut the position of manual labor. Noble quotes a 1968 article by Michael Piore that was based on an extensive survey of engineers: "Virtually without exception, the engineers distrusted hourly labor and admitted a tendency to substitute capital whenever they had the discretion to do so. As one engineer explained, 'if the cost comparison favored labor but we were close, I would mechanize anyway.'"[9]

Any significant technological change (such as the automation of machine tools) involves deep uncertainty as to future costs and therefore profits—uncertainty far more profound than the quotation from Piore's work implies. Relying on simple heuristics to make decisions under such circumstances is perfectly compatible with giving high priority to profit: there is simply no completely rational, assuredly profit-maximizing way of proceeding open to those involved. Analyzing the decisions taken under such circumstances in terms of heuristics rather than imperatives opens up a subtly different set of research questions about the interaction of engineers' culture with the social relations (including the economic relations) of the workplace, and about the different heuristics found under different circumstances (including different national circumstances).

Existing attempts to give empirical content to the ideas of the alternative economics have, however, naturally been more traditionally "economic" than that sort of investigation. Pricing behavior is perhaps the most obvious example.[10] Prices do typically seem to be set according to simple, predictable rules of thumb. Even in the sophisticated U.S. high-performance computer industry, what appears to have been for many years the basic rule is startlingly simple: set the selling price at three times the manufacturing cost.[11] Of course, much more elaborate sets of procedures have evolved (along with the specialist function of the pricing manager). These procedures, however, still seem likely to be comprehensible in the terms of the alternative economics, and indeed open to research (although, perhaps through ignorance, I know of no published study of them). Cray Research, for example, traditionally set its supercomputer prices according to a well-defined financial model whose relevant rule is that from 35 to 40 percent of the proceeds of a sale should cover manufacturing cost plus some parts of field maintenance, leaving a 60 or 65 percent overhead.[12] Discounting and different ways of determining manufacturing cost make such rules, even if simple in form, flexible in application; I would speculate, however, that understanding them is an essential part of understanding the computer industry, and that they are by no means accidental, but (like the control heuristic) have deep roots. It would, for example, be fascinating to compare pricing in the Japanese and American computer industries. There is certainly some reason to think that, in general, Japanese prices may be set according to heuristics quite different from those that appear prevalent in the United States.[13] If this is correct for computing, it is unlikely to be an accidental difference; it is probably related to the considerable differences in the organizational, financial, and cultural circumstances of the two computer industries.

Similarly, it has often been asserted that large firms determine their total research and development (R&D) budgets by relatively straightforward rules of thumb.[14] At Cray Research, for example, the R&D budget is set at 15 percent of total revenue.[15] On the other hand, some recent British evidence suggests that matters are not always that straightforward,[16] and there seem likely to be many other complications, such as the significance of the definition of expenditure as R&D for taxation and for perception of a firm's future prospects. Here too, however, empirical investigation inspired by the alternative economics might be most interesting.[17]

### *Trajectories*

What, however, of the *content* of R&D, rather than its quantity? Perhaps the most distinctive contribution in this area of recent work within the tradition of alternative economics is the notion of the technological trajectory, or the "natural trajectory" of technology.[18]

That there is a real phenomenon to be addressed is clear. Technological change does show persistent patterns, such as the increasing mechanization of manual operations, the growing miniaturization of microelectronic components, and the increasing speed of computer calculations. Some of these patterns are indeed so precise as to take regular quantitative form. For example, "Moore's Law" concerning the annual doubling of the number of components on state-of-the-art microchips, formulated in 1964, has held remarkably well (with at most a gradual increase in doubling time in recent years) from the first planar-process transistor in 1959 to the present day.[19]



The problem, of course, is how such persistent patterns of technological change are to be explained. "Natural" is a dangerously ambiguous term here. One meaning of "natural" is "what is taken to follow as a matter of course"—what people unselfconsciously set out to do, without external prompting. That is the sense of "natural" in the following passage from Nelson and Winter: "The result of today's searches is both a successful new technology and a natural starting place for the searches of tomorrow. There is a 'neighborhood' concept of a quite natural variety. It makes sense to look for a new drug 'similar to' but possibly better than the one that was discovered yesterday. One can think of varying a few elements in the design of yesterday's successful new aircraft, trying to solve problems that still exist in the design or that were evaded through compromise."[20] The trouble is that "natural" has quite another meaning, connoting what is produced by, or according to, nature. That other meaning might not be troublesome did it not resonate with a possible interpretation of the mechanical[21] metaphor of "trajectory." If I throw a stone, I as human agent give it initial direction. Thereafter, its trajectory is influenced by physical forces alone. The notion of "technological trajectory" can thus very easily be taken to mean that once technological change is initially set on a given path (for example, by the selection of a particular paradigm) its development is then determined by technical forces.

If Nelson and Winter incline to the first meaning of "natural," Giovanni Dosi—whose adoption of the notion of trajectory has been at least equally influential—can sometimes[22] be read as embracing the second. To take two examples:

> "Normal" technical progress maintains a momentum of its own which defines the broad orientation of the innovative activities.
>
> Once a path has been selected and established, it shows a momentum of its own.[23]

A persistent pattern of technological change does indeed possess momentum, but never momentum *of its own.* Historical case-study evidence (such as Tom Hughes's study, rich in insights, of the trajectory of hydrogenation chemistry) can be brought to bear to show this, as can the actor-network theory of Michel Callon, Bruno Latour, John Law, and their colleagues.[24] I shall argue the point rather differently, drawing on an aspect of trajectories that is obvious but which, surprisingly, seems to not to have been developed in the literature on the concept[25]—namely, that a technological trajectory can be seen as a self-fulfilling prophecy.

Persistent patterns of technological change are persistent in part because technologists and others believe they will be persistent.

Take, for example, the persistent increase in the speed of computer calculation. At any point in time from the mid 1960s to the early 1980s there seems to have been a reasonably consensual estimate of the likely rate of increase in supercomputer speed: that it would, for example, increase by a factor of 10 every five years.[26] Supercomputer designers drew on such estimates to help them judge how fast their next machine had to be in order to compete with those of their competitors, and thus the estimates were important in shaping supercomputer design. The designer of the ETA10 supercomputer told me that he determined the degree of parallelism of this machine's architecture by deciding that it must be 10 times as fast as its Cyber 205 predecessor. Consulting an expert on microchip technology, he found that the likely speedup in basic chips was of the order of fourfold. The degree of parallelism was then determined by the need to obtain the remaining factor of 2.5 by using multiple processors.[27]

Although I have not yet been able to interview Seymour Cray or the designers of the Japanese supercomputers, the evidence suggests similar processes of reasoning in the rest of mainstream supercomputing (excluding massively parallel architectures and minisupercomputers).[28] Where possible, speed has been increased by the amount assumed necessary by using faster components, while preserving the same architecture and thus diminishing risks and reducing problems of compatibility with existing machines. When sufficiently faster components have not been seen as likely to be available, architectures have been altered to gain increased speed through various forms of parallelism.

The prophecy of a specific rate of increase has thus been self-fulfilling. It has clearly served as an incentive to technological ambition; it has also, albeit less obviously, served to limit such ambition. Why, the reader may ask, did designers satisfice rather than seek to optimize? Why did they not design the fastest possible computer (which is what they, and particularly Seymour Cray, have often been portrayed as doing)? The general difficulties of the concept of optimization aside, the specific reasons were risk and cost. By general consensus, the greater the speed goal, the greater the risk of technological failure and the greater the ultimate cost of the machine. Though supercomputer customers are well heeled, there has traditionally been assumed to be a band of "plausible" supercomputer cost, with few machines costing more than $20 million. If designers did not moderate their ambitions to take risk and



cost into account, their managers and financiers would.[29] The assumed rate of speed helps as a yardstick for what is an appropriately realistic level of ambition.

In the case of supercomputers, all those involved are agreed that increased speed is desirable. Similarly, all those involved with chip design seem to assume that, other things being equal, increased component counts are desirable. Trajectories are self-fulfilling prophecies, however, even when that is not so. Take the "mechanization of processes previously done by hand." Though analyzed as a natural trajectory by Nelson and Winter,[30] it has of course often seemed neither natural nor desirable to those involved—particularly to workers fearing for their jobs or skills, but sometimes also to managements disliking change, investment, and uncertainty. A powerful argument for mechanization, however, has been the assumption that other firms and other countries will mechanize, and that a firm that does not will go out of business. Increasing missile accuracy is a similar, if simpler, case: those who have felt it undesirable (because it might make attractive a nuclear first strike on an opponent's forces) have often felt unable to oppose it because they have assumed it to be inevitable and, specifically, not stoppable by arms control agreements. Their consequent failure to oppose it has been one factor making it possible.

The nature of the technological trajectory as self-fulfilling prophecy can be expressed in the languages of both economics and sociology. As an economist would put it, *expectations* are an irreducible aspect of patterns of technological change. The work of Brian Arthur and Paul David is relevant here, although it has, to my knowledge, largely concerned either/or choices of technique or standard rather than the cumulative, sequential decisions that make up a trajectory. In an amusing and insightful discussion of the almost universal adoption of the inferior QWERTY keyboard, David writes:

> Intuition suggests that if choices were made in a forward-looking way, rather than myopically on the basis of comparisons among currently prevailing costs of different systems, the final outcome could be influenced strongly by the expectations that investors in system components—whether specific touch-typing skills or typewriters—came to hold regarding the decisions that would be made by the other agents. A particular system could triumph over rivals merely because the purchasers of the software (and/or the hardware) expected that it would do so. This intuition seems to be supported by recent formal analyses of markets where purchasers of rival products benefit from externalities conditional upon the size of the compatible system or "network" with which they thereby become joined.[31]

Actors' expectations of the technological future are part of what make a particular future, rather than other possible futures, real. With hindsight, the path actually taken may indeed look natural, indicated by the very nature of the physical world. But Brian Arthur's "nonergodic," path-dependent models of adoption processes are vitally helpful in reminding us of ways in which technologies devoid of clear-cut, initial, intrinsic superiority can rapidly become irreversibly superior in practice through the very process of adoption.[32]

The sociological way of expressing essentially the same point is to say that a technological trajectory is an *institution.* Like any institution, it is sustained not through any internal logic or through intrinsic superiority to other institutions, but because of the interests that develop in its continuance and the belief that it will continue. Its continuance becomes embedded in actors' frameworks of calculation and routine behavior, and it continues because it is thus embedded. It is intensely problematic to see social institutions as natural in the sense of corresponding to nature (although that is how they are often legitimated), but institutions do of course often become natural in the sense of being unselfconsciously taken for granted. The sociological work most relevant here is that of Barry Barnes, who has argued that self-fulfilling prophecy should be seen not as a pathological form of inference (as it often was in earlier sociological discussions), but as the basis of all social institutions, including the pervasive phenomenon of power.[33]

My claim is not the idealist one that all prophecies are self-fulfilling. Many widely held technological predictions prove false. Not all patterns of technological change can be institutionalized, and it would be foolish to deny that the characteristics of the material world, of Callon and Latour's "nonhuman actors," play a part in determining the patterns that do become institutionalized. One reason for the attractiveness of the notion of a natural trajectory to alternative economics is that the latter field has been reacting not against technological determinism (as has much of the sociology of technology), but against a view of technology as an entirely plastic entity shaped at will by the all-knowing hands of market forces.[34]

I entirely sympathize with the instinct that technology cannot be shaped at will, whether by markets or by societies. The risk, however, of expressing that valid instinct in the notion of natural trajectory is that it may actually deaden intellectual curiosity about the causes of persistence in patterns of technological change. Although I am certain this is not intended by its proponents, the term has an unhappy resonance



with widespread (if implicit) prejudices about the proper sphere of social-scientific analysis of technology—prejudices that shut off particular lines of inquiry. Let me give just one example. There is wide agreement that we are witnessing an information-technology "revolution," or a change of "technoeconomic paradigm" based on information and communication technologies. Of key importance to that revolution, or new paradigm, is, by general agreement, microchip technology and its Moore's Law pattern of development: "clearly perceived low and rapidly falling relative cost"; "apparently almost unlimited availability of supply over long periods"; "clear potential for . . . use or incorporation . . . in many products and processes throughout the economy."[35] Yet in all the many economic and sociological studies of information technology there is scarcely a single piece of published research—and I hope I do not write from ignorance here—on the determinants of the Moore's Law pattern.[36] Explicitly or implicitly, it is taken to be a natural trajectory whose effects economists and sociologists may study but whose causes lie outside their ambit. In Dosi's work on semiconductors, for example, Moore's Law is described as "almost a 'natural law' of the industry," a factor shaping technical progress, but not one whose shaping is itself to be investigated.[37] Until such a study of Moore's Law is done, we cannot say precisely what intellectual opportunities are being missed, but it is unlikely that they are negligible.

### *Ethnoaccountancy*

A revised understanding of persistent patterns of technological change offers one potential bridge over the gap between economic and sociological explanations of technical change. Another potential bridge I would call "ethnoaccountancy." I intend the term as analogous to ethnomusicology, ethnobotany, or ethnomethodology. Just as ethnobotany is the study of the way societies classify plants, a study that should not be structured by our perceptions of the validity of these classifications, ethnoaccountancy should be the study of how people do their financial reckoning, irrespective of our perceptions of the adequacy of that reckoning and of the occupational labels attached to those involved.

Ethnoaccountancy has not been a traditional concern of writers within the discipline of accounting. Their natural concern was with how accountancy ought to be practiced, rather than with how it actually is practiced.[38] Although studies of the latter have been become much more common over the past decade (see, for example, the pages of the journal *Accounting, Organizations and Society*), there has still been little systematic study by accountancy researchers of the ethnoaccountancy of technological change. Sociologists, generally, have not been interested in ethnoaccountancy, again at least until very recently.[39] Compare, for example, the enormous bulk of the sociology of medicine with the almost nonexistent sociology of accountancy.[40] Since the latter profession could be argued to be as important to the modern world as the former, it is difficult not to suspect that sociologists have been influenced by accountancy's general image as a field that may be remunerative but is also deeply boring.

It is somewhat more surprising that economists have ignored the actual practices of accounting, but this appears to be the case. Nelson and Winter suggest a reason that, though tendentiously expressed, may be essentially correct: "For orthodoxy, accounting procedures (along with all other aspects of actual decision processes) are a veil over the true phenomena of firm decision making, which are always rationally oriented to the data of the unknowable future. . . . Thanks to orthodoxy's almost unqualified disdain for what it views as the epiphenomena of accounting practice, it may be possible to make great advances in the theoretical representation of firm behavior without any direct empirical research at all—all one needs is an elementary accounting book."[41]

Ethnoaccountancy most centrally concerns the category of "profit." As noted above, even if firms cannot maximize profit, it certainly makes sense to see them as oriented to it. But they can know their profits only through accounting practices. As these change, so does the meaning, for those involved, of profit. Alfred Chandler's *The Visible Hand*, for example, traces how accounting practices and the definition of profit changed as an inseparable part of the emergence of the modern business enterprise.[42] Unfortunately, Chandler clothes his insightful analysis in teleological language—he describes an evolution toward correct accounting practice and a "precise" definition of profit[43]—and he does not directly tie the changes he documents to changing evaluations of technology.

The teleology has largely been corrected and the connection to technological change forged, albeit in a much more limited domain, by the historian of technology Judith McGaw.[44] Though adequate for the purposes of those involved, accounting practice in early-nineteenth-century U.S. papermaking, she notes, "hid capitalization" and highlighted labor costs, facilitating the mechanization of manual tasks. Though others have not made the same connections McGaw has, it is clear that the



practices she documents were not restricted to the particular industry she discusses.[45]

The general issue of whether accounting practice highlights one particular class of cost, thus channeling innovation toward the reduction of that cost, is of considerable significance. Accounting practices that highlight labor costs might generally be expected to accelerate mechanization. They may, however, be a barrier to the introduction of capital-saving or energy-saving technologies, and many current information-technology systems are regarded as having these advantages.

There is also fragmentary but intriguing evidence that the techniques of financial assessment of new technologies used in the United Kingdom and the United States may differ from those used in Japan. In effect, "profit" is defined differently. In the United Kingdom and the United States there is typically great reliance (for decision-making purposes, and also in rewarding managers) on what one critic calls "financial performance measures, such as divisional profit, [which] give an illusion of objectivity and precision [but which] are relatively easy to manipulate in ways that do not enhance the long-term competitive position of the firm, and [which] become the focus of opportunistic behavior by divisional managers."[46] Japanese management accounting, by contrast, is less concerned with financial measurement in this short-term sense. While Japanese firms are patently not indifferent to profit, and are of course legally constrained in how profit is calculated for purposes such as taxation, they seem much more flexible in the internal allocation of costs and the definition of profit. Japanese firms "seem to use [management] accounting systems more to motivate employees to act in accordance with long-term manufacturing strategies than to provide senior management with precise data on costs, variances, and profits."[47]

### *Uncertainty and Closure*

Ethnoaccountancy is one aspect of the much larger topic we might call the construction of the economic. Economic phenomena such as prices, profits, and markets are not just "there"—self-sustaining, self-explaining—but exist only to the extent that certain kinds of relations between people exist. This insight, simultaneously obvious and easy to forget, is perhaps Marx's most central contribution to our topic.[48] Marx devoted the final part of volume 1 of *Capital* to an analysis of the historical emergence of capital as a way of mediating relations between persons. Implicit, too, in Marx's account is the reason why the insight is forgettable. It is not just that capitalism gives rise to a particular type of economic life. Under capitalism, aspects of social relations inseparable in previous forms of society (such as political power and economic relations) achieve a unique degree of separation, giving rise to the "thing-like" appearance of the economic.

One of the fascinations of technological change is that it turns the question of the construction of the economic from a general question about capitalist society into a specific and unavoidable concern. The oft-noted unquantifiable uncertainty of technological change defies the calculative frameworks of economics. Chris Freeman, for example, compares attempts at formal evaluation of R&D projects to "tribal war-dances."[49] He is referring to participants' practices, but it is worth noting that the economists of technological change, in their search for an ancestor to whom to appeal, have often turned to Joseph Schumpeter, with his emphasis on the noncalculative aspects of economic activity, rather than to any more orthodox predecessor.

The issue can usefully be rephrased in the terms of actor-network theory. Radical technological innovation requires the construction of a new actor-network.[50] Indeed, that is perhaps the best way of differentiating radical innovation from more incremental change. Only once a new network has successfully been stabilized does reliable economic calculation become possible.[51] Before it is established, other forms of action, and other forms of understanding, are needed.

Unstabilized networks are thus a problem for economics, at least for orthodox economics. By comparison, their study has been the very lifeblood of the sociology of scientific knowledge.[52] Scientific controversy, where the "interpretative flexibility" of scientific findings is made evident, has been the latter field's most fruitful area of empirical study, and interpretative flexibility is the analogue of what the economists refer to as "uncertainty."[53] The weakness of the sociology of scientific knowledge has, rather, been in the study of "closure"—the reduction of (in principle, endless) interpretative flexibility, the resolution of controversy, the establishment of stable networks.

The economics of technological change and the sociology of scientific knowledge thus approach essentially the same topic—the creation of stable networks—from directly opposite points of view. I confess to what is perhaps a disciplinary bias as to how to proceed in this situation: using tools honed for stable networks to study instability seems to me likely to be less fruitful than using tools honed for instability to study stability.[54] Indeed, attempting the former is where, I would argue, the alternative economists have gone wrong in the concept of technological trajectory.



The latter path, using the tools developed in the study of instability, does, however, require a step back in research on technological change: a return to the "natural history"[55] of innovation of the 1960s and the 1970s, but a return with a different focus, highlighting the empirical study of heuristics, the role of the self-fulfilling prophecy in persistent patterns of technological change, and the ethnoaccountancy of technological change. We need to know more about the structure of the interpretative flexibility inherent in technological change, and about the ways that interpretative flexibility is reduced in practice. How, in the economists' terminology, is uncertainty converted into risk?[56] How, for example, do participants judge whether they are attempting incremental or radical innovation?[57] What is the role of the testing of technologies (and of analogues such as prototyping and benchmarking)?[58] How is technological change "packaged" for the purposes of management—in other words, how is a process that from one perspective can be seen as inherently uncertain presented as subject to rational control? What are the roles here of project proposals, project reviews, and milestones—of the different components of Freeman's "war-dances"? How is the boundary between the "technical" and the "nontechnical" negotiated? What are the determinants of the credibility of technical, and of nontechnical, knowledge claims?

Even if we set aside the fact that technological change is not substantively the same as scientific change, we cannot look to the sociology of scientific knowledge for theories or models that could be applied directly in seeking to answer questions such as these. That is not the way the field has developed. It is more a question of sensitivities, analogies, and vocabularies. Nevertheless, the parallels between closure in science and successful innovation in technology, and between interpretative flexibility and uncertainty, are strong enough to suggest that exploring those parallels may be an important way forward for the study of technological change. In the closure of scientific controversies and in successful technological innovation, an apparently self-sustaining realm (of objective knowledge, of economic processes) emerges, but only as the end product of a process involving much more than either natural reality or economic calculation. Understanding of the one should surely help develop understanding of the other.

### *Conclusion*

I have argued that the alternative economics associated with Simon, Nelson, Winter, and others is more plausible than neoclassical economics, with its incoherent notion of profit maximization. Ideas from the former tradition could help bridge the gap between the economic and the sociological in fields where those ideas have not (to my knowledge) been widely drawn upon, such as labor-process studies. This alternative economics can also fairly straightforwardly be applied to pricing and to firms' overall R&D budgets, although recent empirical work in these areas seems surprisingly sparse.

Applying the alternative economics to the content of R&D is more difficult. The metaphor of "technological trajectory" can mislead. Persistent patterns of technological change do exist, but they should not been seen as "natural" in the sense of corresponding to nature. Nor do they have a momentum of their own. Expectations about the technological future are central to them: they have the form of self-fulfilling prophecies, or social institutions. Conceiving of persistent patterns in this way offers one way of bridging the gap between economic and sociological explanations of technological change.

Another way of bridging the gap is what I have called ethnoaccountancy. Studying how people actually do the financial reckoning of technological change would bring together the economist's essential concern for the financial aspects of innovation with the sociologist's equally justified empiricism. I have suggested that ethnoaccountancy would not be a marginal enterprise, rummaging though the boring details of economic activity, but ought to throw light on central questions such as the practical definition of profit and the relative rate of technological change in different historical and national contexts.

Finally, I have argued that, because of the centrality of uncertainty (or nonstabilized networks) to technological change, the sociology of scientific knowledge, with its experience in the study of the essentially equivalent matter of interpretative flexibility, ought to be of relevance here. Scientists construct stable, irreversible developments in knowledge in a world where no knowledge possesses absolute warrant; out of potential chaos, they construct established truth. Technologists, workers, users, and managers construct successful innovations in a world where technological change involves inherent uncertainty; out of potential chaos, they construct a world in which economics is applicable.

### *Acknowledgments*

The research reported here was supported by the Economic and Social Research Council's Programme on Information and Communication



Technologies. Earlier versions of this chapter were read to the conference on Firm Strategy and Technical Change: Micro Economics or Micro Sociology? (Manchester, 1990) and to a meeting of the ESRC New Technologies and the Firm Initiative (Stirling, 1991). Thanks to Tom Burns, Chris Freeman, John Law, Ian Miles, Albert Richards, Steve Woolgar, and both the above audiences for comments, and particularly to the neoclassical economists for their toleration of a sociologist's rudeness.