Making Sense of Making Artefacts:

Reflections on

John Ziman (ed.) *Technological Innovation as an Evolutionary Process*

Cambridge: Cambridge University Press, 2000

Brian J Loasby

26 April 2001

Department of Economics
University of Stirling
Stirling FK94LA
Scotland

Tel +44 1786 467470
Fax +44 1786 467469

e-mail b.j.loasby@stir.ac.uk

[M]an is always free to reconstruct what he may not deny

(Kelly 1963, p. xii)

The process which resulted in this book 'started out with a simple question: could the obvious analogy between technical innovation and biological evolution be developed from a "metaphor" into a "model"?'(p. 312). This question was explored in a series of meetings by an interdisciplinary group of scholars brought together for the purpose by John Ziman, and the results of this exploration have now been organised into a 'provisional reply' (p. xv). What follows is primarily an attempt, stimulated by the original question and the reply (the details of which are largely implicit as well as provisional) to present an alternative approach to the analysis of technological innovation, with the aim of contributing to the evolution of thought about this remarkable feature of human history.

Consequently I shall not be attempting a conventional review of the book, and am very happy to recommend Keith Pavitt's review, with most of which I strongly agree. However, some readers may find it useful to have an outline of the contents, and so I shall begin with that; others may prefer to go directly to my own argument in Part II, which does not depend in any essential way on Part I. As a coda to the main argument, I offer in Part III a brief personal history as an illustration of some of the themes of that argument.

Part I. Theme and contents of the book

A fair impression of the scope and methods of this enquiry is given by the chapter titles. The first section, 'Evolutionary thinking', begins with an examination by John Ziman of the issues that are raised by a search for 'Evolutionary models for technological change'; this is followed by an exploration of the biological comparator in two chapters, 'Biological evolution: processes and phenomena' by Eva Jablonka and John Ziman, and 'Lamarckian inheritance systems in biology: a source of metaphors and models in technological evolution' by Eva Jablonka; a major theme of this exploration is that biological evolution includes much more than natural selection among the results of 'blind mistakes' (p. 16) in copying DNA. 'Selectionism and complexity' by John Ziman provides a review of technological systems and a preview of some of the problems of mapping biological reasoning on to them, leading to chapters by Joel Mokyr on 'Evolutionary phenomena in technological change', and by Richard Nelson on 'Selection criteria and selection processes in cultural evolution theories', both of which emphasise the importance of knowledge and of the cognitive, organisational and cultural criteria and mechanisms of selection in shaping technological innovation.

In the second section, 'Innovation as a cultural practice', Alan Macfarlane and Sarah Harrison seek to explain the contrast through much of the second millenium between the development of non-human power in Europe and the increasing dependence on human power in Japan in a path-dependent and incrementalist account of 'Technological evolution and involution: a preliminary comparison of Europe and Japan'. The chapter is nicely complemented by Gerry Martin's study of the Japanese sword as an exemplar of 'Stasis in complex artefacts'. David Turnbull explains in 'Gothic tales of spandrels, hooks and monsters: complexity, multiplicity and association in the explanation of technological change' how a '*distributed* design process' (p. 115) led to many innovations in cathedral building, and Paul David uses the results of computer simulation to provide a focus for his discussion of 'Path dependence and varieties of learning in the evolution of technological practice'.

The third section, 'Invention as a process' begins with a historical investigation by W. Bernard Carlson of 'Invention and evolution: the case of Edison's sketches of the telephone', followed by David Perkins' consideration of 'The evolution of adaptive form' through biological evolution and human invention in rugged (what Perkins calls 'Klondike') fitness landscapes, and Walter G. Vincenti's practical engineering approach to 'Real-world variation-selection in the evolution of technological form: historical examples'. Joan Solomon then offers an educationalist perspective on 'Learning to be inventive: design, evaluation and selection in primary school technology', and Geoffrey Miller discusses genetic algorithms under the title of 'Technological evolution as self-fulfilling prophecy'.

The theme of the fourth section, 'Institutionalized innovation', is the ordering of technological progress. Edward Constant claims that 'Recursive knowledge and the evolution of technological knowledge' are linked by the embodiment of knowledge in standard practices which survive selection, and Rikard Stankiewicz develops and applies 'The concept of "design space" ' to the evolution of standard practices within technological regimes. In 'Artefact <--> activity: the coevolution of artefacts, knowledge and organization in technological innovation', James Fleck insists on the importance of interdependence between the three, but argues that specific organisational forms are relatively weakly connected to the artefacts-activity couple (thus avoiding the three-body problem), before Gerard Fairtlough provides the only sustained treatment in this book of 'The organization of innovative enterprises'.

The final section, 'Technological change in a wider perspective', includes two chapters which suggest unexploited contrasts in the motivation for and the uses of technological change. Edward Constant effectively illustrates 'The evolution of war and technology' but does not discuss the political or social dimension, whereas Janet Davies Burns argues, in 'Learning about technology in society: developing liberating literacy', for education as a means of enabling people to make effective choices of better technologies, on the apparent assumption that such choices would be predominantly benign in both motivation and effect. The section concludes with 'An end-word' by all contributors, which readers might wish were longer than four and a half pages.

Part II. Reflections on Ziman

I encourage anyone interested in understanding technological innovation to read this symposium, whether or not they consider it helpful to think of innovation in evolutionary terms. As so often when thoughtful authors are drawing on their specific expertise, the value of these chapters extends far beyond their responses to the question that prompted them, although one may usefully observe both the diverse ways in which the theme has influenced the content and also the varied relationships between the content of each chapter and the conclusions that are drawn by its author. John Ziman has been very successful in organising the generation of variety within his selected population of contributors, which includes both complementary and alternative elements, thereby producing an example of part of an evolutionary process. Every reader's response will be influenced by that reader's 'absorptive capacity'; some account of the evolution of my own absorptive capacity is given in Part III, and it has contributed to the views that I will now summarise.

I agree entirely with the underlying principle of the book that technological innovation and biological evolution have in common the generation of variety and selection from this variety; and the cumulative effect of both processes typically (though not invariably) is an increasing differentiation of function matched by a closer integration between functions. But there are also major differences in the ways in which variety is generated and selected and in the content of differentiation and integration. In addition, ex-ante as well as ex-post selection is an essential feature of technological innovation, and the processes of variety generation and selection, far from being sharply differentiated, are deeply entwined (as Pavitt notes in his review): the incubation of a new artefact or method of production involves frequent rejection of candidate variants that leads directly to new design, and users are shapers as well as selectors. Furthermore, selection criteria in technological innovation, and in the knowledge (both theoretical and practical) that supports it, include emotional and aesthetic as well as 'rational' elements; even the rationality is often of a kind that fails to meet the (emotional and aesthetic) criteria of orthodox modern economists. What perhaps most distinguishes technological from biological evolution is that it rests on the organisation of knowledge, which is itself supported by the organisation of the process of generating, testing, and modifying knowledge. Underpinning all these activities, of course, are the biologically-evolved capabilities and motivations of human beings, and it seems to me that an understanding of these capabilities and motivations, rather than transferable models, is the prime contribution that evolutionary biology can make to the study of technological innovation. I will develop these arguments in more detail throughout my essay.

<u>Evolution, rationality and knowledge</u>

Despite the brevity of the 'end-word', which perhaps indicates a laudable desire to avoid closing down any aspect of an enquiry into the relationship between evolutionary processes in technology and in biological organisms, the implications of the symposium for the applicability of 'the biological model' to technological innovation are clear and important. I wholeheartedly agree with the consensual endorsement (p. 313) of 'the explanatory power of “evolutionary reasoning” in a very wide variety of contexts’. However, the quotation marks in the original suggest a wide variety of interpretations of ‘evolutionary reasoning’ across these contexts, supporting the inference, which is left implicit in the final chapter but had already been stated by Mokyr (p. 58), that biological evolution is just one context and that the methods which are appropriate there have no valid claim to privilege in others. That inference I readily accept, and will later support with other arguments; there are enough similarities between biological and technological development to justify a generic label and an overarching concept, but enough differences to require separate treatment. This speciation of evolutionary processes has its distinctive virtues to complement those of inclusiveness that result from the shared basic notions of variation and selection.

However, the advantages of the speciation of evolutionary processes that became apparent to many of the contributors in the course of their work raises a puzzle about the initial decision to approach evolutionary reasoning from the perspective of a biological model, especially since the co-ordinator of this enterprise has such a profound understanding of the growth of scientific knowledge as a special case of general human knowledge, which – for very good reasons, as we shall see – is highly amenable to evolutionary reasoning. It would not have occurred to me to ask the question that prompted this book, although I shall argue in the final section of this Part that biological evolution provides an important basis for our understanding of human cognitive processes, not least those which generate technological innovation. Now problems are defined by differences (Pounds 1969), as I learnt 35 years ago, and the recognition of problems is a prerequisite of search – above all of a search for knowledge, including the search for routines that embody significant ‘know-how’, or productive capabilities, as in Nelson and Winter’s (1982) evolutionary theory; I therefore thought that it might be helpful to some readers to focus on this difference of approach between the motivation of this book and my own as a problem worth investigating, before going on to explore the evolution of knowledge (of ways of ‘making sense’) as a means of analysing technological innovation. This will, I hope, incidentally serve to illustrate ‘the wide variety of contexts’ in which some kind of evolutionary reasoning may be helpful; any wider exploration of this variety I leave to others, pleading the advantages of the division of labour in the search for knowledge, which will be argued in more detail later. The tendency to variation within a community generates a comparative advantage which enlarges the opportunity set of that community.

I will try to shape my investigation in a way that illustrates the influence of particular evolutionary sequences on the generation of variety in ways of thinking about problems, an issue that is surely relevant to an understanding of technological innovation. Keith Pavitt (1998) recently published an article on technological progress to which he gave the sub-title ‘What Adam Smith tells us and Joseph Schumpeter doesn’t’; I was tempted to give this paper the title ‘What Adam Smith tells us and Charles Darwin doesn’t’, but that would be unfair to Darwin, who was less prescriptive than neo-Darwinians about the sources of variety, and also recognised the significance of Smith’s work, as we shall see. That title also suggests a criticism of the organising principle of this book which I have no wish to make, even though I would not have chosen it. There is no one best way.

However, I do wish to make a strong case for the relevance of Smith who, as we shall see later, gives us three fundamental principles of technological evolution, only one of which (for good reason) has been incorporated into biological theory.

Since the evocation of Smith suggests that I shall be taking an economic perspective, it is essential to emphasise that my perspective is not that of mainstream economics, fairly broadly construed. This differentiation between ways of thinking among economists, both over time and at a particular time, defines another set of problems, about which I have written on several occasions (e.g. Loasby 1978, 1989, 1990, 1996, 1998, 1999b and c), and which has important similarities (not examined in this paper) to the processes of thinking that guide the evolution of technology; but it is appropriate to focus briefly on mainstream economics, because that allows us to identify a fundamental issue in conceptualising evolution, whatever the context. Several of the contributors remind us of the neo-Darwinian claim that the only alternative to the Darwinian explanation of life-forms is the now-discredited explanation by design. However, explanation by design, in the form of equilibria of rational choosers, is the foundational principle of standard economics (though not, as relatively few economists recognise, of Smith's economics), and it is deemed sufficient for standard economic analyses of technical change, although it does not satisfy the economists who have contributed to this book. An evolutionary process within the economics profession, including directed variation, internal as well as external selection, and tribal behaviour (on the details of which and its consequences for the content of their subject very few economists have publicly reflected), has led by incremental adaptation to a style of modelling that relies on what outsiders may consider to be an extreme – and even irrational – form of rationality: all agents base their choices on the correct model of the economy, which includes (usually by implication) the correct model of every other agent's behaviour.

This analytical system assumes that selection is highly efficient, and takes place before the event; there is thus no room for any kind of process that might reasonably be called 'evolutionary' – or indeed anything that might be called a process in the usual sense of the word. (Many economists explicitly refuse to consider modelling anything but equilibria, even when claiming to explain growth and change.) The criteria of rigorous theorising in economics require the set of possibilities to be complete, and known to be so, allowing for 'learning' only in the form of obtaining increasingly accurate estimates of the likelihood of each possibility by the application of a procedure that is fully specified in advance. Time is incorporated into the model as an additional dimension of goods and information sets, to the exclusion of any analysis of an economy which develops in time. Any deficiencies in the outcomes of economic activity, in relation to the best possible allocation of the resources available, are therefore attributable only to inappropriate incentives, which themselves are typically the result of elements of monopoly, missing markets (e.g. for public goods) or asymmetric information which permits opportunistic gains at the expense of others; and it is therefore not surprising that 'today economists can define their field more broadly, as the analysis of incentives in all social institutions' (Myerson 1999, p. 1068).

Non-economists might be amazed by the phrase ‘more broadly’, but Myerson is so impressed by the expansion of economic analysis into non-market applications that he fails to register the restriction of analytical content to the effects of incentives, which are themselves narrowly construed. This is a splendid illustration of the principle that problems are not recognised if differences are not perceived, and it also suggests how problem-generating perceptions which might lead to new knowledge may be crowded out (a phenomenon which is not unknown in the history of technology). In an analytical system which relies on the equilibria of optimising agents, and is constrained by the demands of internal coherence, the information on which choices are based is problematic only when access to this information is costly, and even then it is optimally selected, essentially by choosing the basis and fineness of its partitioning. Thus the set of innovative opportunities is known, though perhaps in somewhat coarse partitions; innovation is a problem of incentives, and perverse incentives may justify public funding of some kinds of innovation. Explanation by design is at the core of economics, which might therefore be considered in direct conflict with biological principles.

However, matters are not so simple. Some economists have suggested that models of rational choice equilibria may be regarded simply as convenient instruments for prediction, and that we may have considerable confidence in their predictive value because even a complete absence of rationality, in a market setting, will result in convergence on outcomes which are similar to those of well-informed optimisation. Ex-ante and ex-post selection, it is claimed, are close to being observationally equivalent; and ex-ante selection is easier (and more elegant – one should not overlook the aesthetics of rationality) to model. Selection processes in market systems are so effective that evolutionary arguments are thought to be superfluous in economics, including the economics of technological innovation, which becomes essentially a race to attain a technology that is known to exist.

The most thoughtful version of this argument, by Alchian (1950), made no claims for the optimality of ex-post market selection operating on non-rational behaviour, but simply for an average response in the appropriate direction to any change in circumstances, which was as much as Alchian thought could reasonably be hoped for. (Endogenous change was not considered.) Other economists have been less cautious in asserting that market selection can duplicate the results of rationality – in striking parallel to the claims, once widely thought to be both irrefutable and significant, that central planning and perfect competition (under appropriate conditions) deliver identical outcomes. It is noteworthy that some neo-Darwinians are similarly inclined to consider the outcomes of biological evolution to be structures and behaviour which are very close to optimality, and some may even argue that they have better grounds than economists for this claim because biological selection has had much longer to produce such outcomes (Maynard Smith 1996, p. 291). Thus it is not surprising to find game-theoretic equilibria being invoked as a short-cut to the terminal points of evolutionary biological processes.

Paul Krugman (1999) warned members of the European Association for Evolutionary Political Economy in 1996 that evolutionary biologists were increasingly attracted to equilibrium modelling, implying that a distinctively ‘evolutionary’ economics was a chimera, and even that biologists might be looking to modifications of standard economics for their own models. The contributors to this symposium seem to feel no attraction to equilibrium modelling, even in explaining the reliance on human power in historical Japan (Chapter 7) and the absence of change 'in form, function and manufacturing process for over 700 years' of the Japanese sword (Chapter 8), where I would have thought that the foundational idea of equilibrium – as the result, sometimes self-destroying, of an emergent balance of forces – would be highly appropriate. Game theory is mentioned twice (pp. 50, 120) but never used, whereas Fairtlough, in the only chapter written directly from industrial experience, explains the value of multiple scenarios as a stimulus to changing ways of thinking, and therefore to the creation of variety (pp. 275-6); and there are fewer claims to optimality of any kind than warnings of the difficulty of deciding how optimality might be appropriately defined in each context – the most striking examples being found in Constant's chapter on war, where determining what is to count as success and failure turns out to be much less straightforward than one might assume.

There is no doubt that these two conceptions, of natural selection from random mutations, and of optimal choice from known opportunity sets, both facilitate the construction of closed and (apparently) completely specified models which meet fashionable criteria of 'rigour', although, as Ziman (p. 42) points out, randomness, like more complex probability distributions, assumes a careful – and correct (but how do we know?) – definition of the relevant search space. Their popularity may therefore be explained by a combination of ex-ante and ex-post selection by and of practitioners. Unfortunately, however, neither is a good match to the problems of human activity, of which technological innovation is a prominent example. The fundamental difficulty with rational choice theory is its untenable assumption about human knowledge, as Frank Knight (1921) pointed out 80 years ago; and the fundamental difficulty about neo-Darwinian explanations of human activity, as Edith Penrose (1952) insisted, is that it ignores human purpose. Human action is often the result of human design; but human design is inherently fallible, however secure its logic, since it is based on knowledge that is usually incomplete or erroneous. This has long been recognised. 'Purposes mistook, fall'n on the inventors' heads' is the stuff of tragedy – and of comedy too; on the other hand many of the desirable features of society, though the consequences of human action, were not consciously intended by anyone. Technical change, like most human activities, lies in the interval between optimal choice and chance variation, and by opting for either, or both, of these models (which we might think of as corner solutions in the space of theoretical principles) we exclude at the outset the possibility of understanding what is happening, and not least – though this topic will not be directly addressed in this paper – of understanding the selection processes within academic disciplines.

Evolutionary processes in human societies need not, and I suggest should not, exclude rationality in the broad sense of acting for good reasons. What is essential is uncertainty: the absence of any procedure for decision-making that is known to be correct (Knight 1921), which often extends to the absence of any means of ensuring that all possibilities have been identified (Shackle 1972). In the presence of uncertainty, the generation of alternative hypotheses (some of which may be embodied in artefacts) and selection among them, which may lead to the generation of further hypotheses, is likely to be an effective means of progress (though not always of improvement in terms of human welfare). Uncertainty seems to be pervasive enough to justify an evolutionary approach to the growth of academic, technological and everyday knowledge, but an approach which is significantly different from the biological model. The evolutionary growth of knowledge provides the theme for the remainder of this paper.

Neo-Darwinian evolution requires stability in both the selection environment and in the genotypes which are subject to selection; it also requires genetic mutation to provide new variants from which to select. This dual genetic requirement can be satisfied only if the chances of a defective copy are extremely small but not zero, and that in turn requires neo-Darwinian evolution to be not only incremental but extremely slow. This doesn't look like a good model for technological change, though it does encourage us to postulate stable genetic characteristics in the human population over periods which are by comparison extremely brief – a postulate which underlies the final section of the argument presented in this Part. What it does have in common with technological innovation is the importance of a reliable baseline; without this neither ex-ante nor ex-post selection can be significant. A more promising hypothesis is Schumpeter's (1934) proposition that purposeful innovation depends on the baseline for decision making that is provided by a stable economic environment, which major innovations are certain to destroy; his theme, developed at substantial length (especially in Schumpeter 1939), that technological innovation generates real business cycles, has attracted little attention among either mainstream economists or contemporary Schumpeterians. Schumpeter himself was careful to avoid any association with Darwin's ideas, which had fallen out of favour at the time – the early 20th century – that he was developing his own. The general principle that I would emphasise is that theories of innovation should explain what does not change as well as what does, and the effect of this balance on particular processes; this requirement is certainly met in the accounts of change in this book, but its importance is nowhere made explicit.

In his review, Pavitt discusses the problem of identifying the unit of selection, but seems inclined to accept the neo-Darwinian principle that in any evolutionary process there should be only one such unit. I do not: techniques, artefacts and firms are all relevant, and so too are institutions, organisational arrangements, and bodies of knowledge, including know-that, know-why, know-how and know-who. The essential requirement is to distinguish, at each stage of analysis, between the elements and the connections that remain stable and the elements and connections that change. This combination varies according to time and circumstance; and there is no simple hierarchy. Sometimes established elements are assembled into a novel architecture; sometimes a modular architecture facilitiates quasi-independent developments. Stability in the direction of technological change, as Pavitt notes in his review, is likely to encourage variation within that trajectory and also variation in the combination of techniques to produce artefacts. Decomposition and recombination are important principles both in technological innovation and in the study of technological innovation – as in other kinds of knowledge.

The evolution of ideas and capabilities

It was only after the publication of *Choice, Complexity and Ignorance* (Loasby 1976), the origins of which are explained in Part III, that my attention was drawn to the significance of Adam Smith's earliest surviving major work. His friend David Hume had demonstrated that there was no way of proving the truth of any general empirical proposition, either by deduction, for there was no way of ensuring the truth of the premises, or by induction, for there was no way of proving that instances not observed would correspond to those that had been observed. Hume's response had been to turn to the manageable question of how people came to accept certain empirical propositions as true; and Smith followed Hume's example by producing a psychological theory of the emergence and development of science which, as I claimed earlier, provides three essential elements of evolutionary explanations in human society, and illustrated it by the history of astronomy (Smith [1795] 1980).

The first element is the motivation for generating new ideas. The evolution of human knowledge, both theoretical and practical, though unpredictable in any detail, is driven by purpose, and this is often an emotional rather than a rational force. Smith argues that people are disturbed by the unexpected, dismayed by the inexplicable, and delighted by schemes of thought that resolve the inexplicable into plausible generalisations. In the absence of any assured procedure for attaining correct knowledge, these are the motives which 'lead and direct philosophical enquiries'. They are a long way from the incentives that economists seek to model, but perhaps not so far from some of the incentives that shape the behaviour of technologists, and of economists also.

The second element is the sequence that is inspired by this complex motivation: the generation of novelty and the selection processes which guide its adoption or rejection. People try to invent 'connecting principles' which will afford a basis for collecting phenomena into categories and link each category with an explanation which is credible enough to 'soothe the imagination'. The 'equalizing circle' in Ptolemaic geometry and the rule that 'when one body revolved round another it described equal areas in equal times' in Kepler's system are examples that Smith ([1795] 1980, pp. 61, 90) uses of the resolution of difficulties by appealing to general principles of motion that appear congenial to prevailing notions of good order; rational choices based on rational expectations are widely accepted principles for explaining economic phenomena; and we are now identifying the connecting principles of Smith's own explanation of the growth of knowledge.

Smith gives particular attention to the importance of aesthetic criteria both in guiding conjectures, for example in the ideas of Copernicus and Kepler, and in encouraging their acceptance, notably in discussing the rhetorical appeal of the Newtonian system, which in his *Lectures on Rhetoric* exemplifies Smith's ideal method of 'giving an account of some system' (Smith 1983, p. 146). Aesthetic influences in the natural sciences and in economics (signalled earlier by the reference to the elegance of rational choice equilibria) are occasionally recognised but rarely explored (see Schlicht 2000); aesthetic influences on the design of artefacts are an unacknowledged theme of this book. In cathedral building aesthetic appeal is a major objective, duly reflected in the templates to which Turnbull draws attention; but of particular interest in an exploration of evolutionary processes is the extent to which aesthetic criteria are also surrogates for effective performance; bridges and aircraft are obvious examples, and the flawed design of the Millennium footbridge in London, which causes it to sway so disconcertingly in use that it has been closed, and for which no simple or cheap remedy has been identified, is a current illustration that surrogacy should not be assumed.

Connecting principles which seem to work well are widely diffused, because of the human readiness to look for guidance from others who seem to know better, and because of our desire to act, and indeed think, in ways that merit the approval of others. This is a foundational principle of Smith's ([1759] 1976a) *Theory of Moral Sentiments*, which is itself an essential element in Smith's complex account of social organisation, and applicable to technological evolution. (For an excellent discussion of the impact of social approval and disapproval on technological development, effectively linking general principles to detailed histories, see Pool 1997.) However, because invented principles, however widely accepted, are not proven truth (even, Smith notes, when these principles have been invented by Newton) they are liable eventually to be confronted with unexpected phenomena which they cannot be adapted to explain. This disjunction between evidence and established means of explanation defines a pressing problem; when satisfactory adaptation is despaired of, a new search for connecting principles begins.

The third element in what we might retrospectively call Smith's evolutionary theory is the process by which this basic human activity generates first an increasingly distinct category of knowledge which comes to be called 'scientific' and subsequently a progressive differentiation between sciences that we might now label speciation. The consequent differences both of focus and of criteria for acceptable categories and acceptable explanations generate a greater variety of more precisely-defined problems and consequently accelerate the growth of science. This, it is worth noting, is the context in which the effects of the division of labour first appear in Smith's ([1795] 1980) surviving work (though publication did not occur until after his death); in this sequence it seems almost natural, and therefore a source of pleasure, that the division of labour is invoked in the *Wealth of Nations*, not as the best way to make the most of differentiated skills – which was a very old idea – but as the chief instrument of the growth of productive knowledge (Smith [1776] 1976b). Since this is easily the most important idea in economics – the co-ordination problem which normally receives priority among economists would be trivial without the continuous generation of new knowledge and new artefacts – it is worth recognising the process by which it came about.

Smith's prime 'connecting principle' of the division of labour was applied to physiology in 1827 (Milne-Edwards 1827) and this application in turn contributed to Darwin's vision of the reason why a Malthusian struggle to survive should result in the differentiation of species (Raffaelli 2001). The other two basic elements in Smith's account of the development of knowledge by motivated trial, error, amendment and diffusion understandably did not. From the point of view of technological innovation, therefore we may suggest that Smith encompasses Darwin; hence the temptation to give this paper a different title.

The differentiation of knowledge is a condition of progress in human society. However it has its opportunity costs, of which two are especially important in understanding technological innovation. One is that differences in the structure of understanding, and in the criteria for good theory and good practice, may create substantial obstacles to the integration of knowledge across disciplines or between technological fields, as well as obstacles to the integration of technological and non-technological perceptions of the value of any particular innovation, which is the primary focus of Pool's book. A special, but not uncommon, case of such differences in perceptual structure is that between sensory perception and scientific categorisation: 'events which to our senses may appear to be of the same kind may have to be treated as different in the physical order, while events which physically may be of the same or at least of a similar kind may appear as altogether different to our senses' (Hayek 1952, p. 4). Theoretical developments may not map readily onto recursive practice, and know-how may resist usable codification. The desire to assuage the discomfort of this apparent contradiction led Hayek to construct an evolutionary account of the development of *The Sensory Order*, to which we shall refer later.

The other opportunity cost of the differentiation of knowledge is the neglect of potentially crucial interdependencies. 'When the compass of potential knowledge as a whole has been split up into superficially convenient sectors, there is no knowing whether each sector has a natural self-sufficiency…Whatever theory is then devised will exist by sufferance of the things that it has excluded' (Shackle 1972, pp. 353-4). Charles Suckling (see Part III) saw this as a key issue in the management of innovation, as in many other fields. Unanticipated technological disasters are frequently traceable to unjustified assumptions (usually unconscious, but not always so) about the sufferance of something excluded from the processes of design, testing, or operator training. The Millennium footbridge already mentioned is an exemplary demonstration.

Implications of uncertainty for cognition and the growth of knowledge

The double-edged character of uncertainty is the focus of Frank Knight's *Risk, Uncertainty and Profit* (1921). Knight restricted the concept of risk to situations in which both the set of possibilities and the probability distribution over this set are known, either by argument *a priori,* as in calculating the expected results of throwing dice, or by statistical analysis of appropriate evidence. Choices under risk may be made by a standard procedure which can be demonstrated to be optimal, and no-one can gain a sustainable advantage in making such choices (except by forcibly preventing anyone else from making them). But when neither basis for calculation is adequate, no demonstrably optimal procedure can be devised; we are faced with uncertainty, which must be handled in some other way – in the space between optimality and randomness. Knight hints that both methods of assigning probabilities to supposedly risky situations may themselves be subject to uncertainty; this would be consistent with Hume's arguments about the unprovability of all empirical knowledge.

However, if uncertainty creates difficulties, it also creates opportunities for imagination – as in Smith's psychological theory: indeed, Knight argues that it is a necessary condition for entrepreneurship and profit – and also for the firm, which provides shelter for those who are unwilling to cope with uncertainty in person and prefer the conditional security offered by entrepreneurs. The opportunities perceived by Knight are to be found both within the economic system and in the corpus of economic theory, where it is appropriate to cite the (very different) ideas that economic interaction might be formally analysed as a game between hyper-rational players or that a firm might be conceived, not as a production function or a nexus of contracts but as a pool of resources, of uncertain applicability, within an administrative framework (Penrose 1959). .

In this paper we need not follow Knight into his particular application (though it is hardly irrelevant to the analysis of technological innovation) pausing only to note that most economists assimilate Knight's category of uncertainty to risk by the invocation of subjective probability, sacrificing the opportunities in Knight's analysis for theoretical development 'in order to preserve the coherence of the ideas of the imagination' (Smith ([1795] 1980, p. 77). Smith would have understood this response very well, even though it is unlikely that he would have approved; it may also help to explain the difficulties noted by Pavitt (1998 and in his review) in adapting organisational practices and power relationships to technological requirements. Instead we may focus on the broader opportunities for improving our understanding of the growth of human knowledge, and its technological manifestations. An opportunity which Knight himself fails to develop is his observation that in the complete absence of uncertainty 'it is doubtful whether intelligence itself would exist' (Knight 1921, p. 268): this locates the role of intelligence squarely in the space between optimal choice or optimal design and random activity, and in doing so warns us not to identify intelligence with logical operations. Niels Bohr's rebuke was blunter: 'You are not thinking; you are merely being logical' (Frisch 1979, p. 95). This dissociation of intelligence from logic underlies Knight's (p. 241) observation that '[m]en differ in their capacity by perception and inference to *form correct judgements* as to the future course of events in the environment. This capacity, moreover, is far from homogenous'; moreover, individuals differ in their capacity to change, and learning takes time (Knight 1921, p. 243). Knight is talking about the effect of the division of labour on the development of differentiated intelligence, though without reference to Adam Smith.

Knight (1921, p. 206) is also unconsciously close to Smith in arguing that 'in order to live intelligently in our world…we must use the principle that things similar in some respects will behave similarly in certain other respects even when they are very different in still other respects': in other words, we rely on 'connecting principles' of association and causation – together with 'the sufferance of the things that [they have] excluded' – in developing our own ideas and in adapting other people's. What similarities we emphasise and what differences we ignore depend both on our perception of problems and our own evolved connecting principles, or those of our discipline, profession, or organisation. That is why it is not surprising, given the prestige of evolutionary biology and the absence of an agreed tightly-specified model of technological innovation, that Ziman's group formulated their research agenda in the way that they did, nor that I followed a different path (outlined in Part III); drawing on different authors in an attempt to deal with differently-formulated problems, we relied on different contexts of similarity. That is how the division of labour leads to differentiated knowledge.

Fleck (p. 255) complains that a 'focus purely upon knowledge...makes the evolutionary problem very tough. It is difficult to put boundaries around an idea'. Why, indeed, should we assume that, within the categories that we invent, the similarities dominate the differences, while between these invented categories the reverse applies? (Compare Kelly 1963, pp. 53-5). However, ambiguity, like uncertainty of which it is a special case, is both a problem and an opportunity for the generation of ideas by making new combinations (a principle enunciated by both Smith and Schumpeter); thus the difficulty of putting boundaries around an idea is a major enabler of innovation. The difficulty of putting boundaries around the capabilities of any individual or organisation, and the consequent ambiguity of their range of application, is a prominent theme in Nelson and Winter's (1982) theory, and underlies Penrose's (1959) brilliant exploitation of her distinction between resources and productive services.

Ambiguities of both capabilities and 'knowledge that' also explain why diffusion, typically across different contexts of similarity, as my former colleague Frank Bradbury frequently reminded us, is often both unexpectedly difficult and also a major contributor to the content of innovation. The use of metaphor, which has played no small role both in technological innovation and in the attempt to understand it, illustrates the point; abstract thought relies on language which originated in metaphor – indeed the terms 'abstract' and 'metaphor' originate in Latin and Greek metaphors. The Smith-Knight principle of sufficient similarity is fundamental both to making sense and to making artefacts. If, in Shackle's (1979, p. 26) beautiful phrase, innovation begins with 'the imagined, deemed possible', both what is imagined and the judgement of its possibility rest on the exploitation of ambiguity. What is deemed possible, and – even more – what is deemed capable of being made possible, depends upon judgements about the applicability of both theoretical and practical knowledge to novel contexts.

Category-based judgements of possibility guide the innovation process; but because they are possibilities, not specific predictions – and because the judgements are themselves subject to error – they cannot, as some writers on corporate strategy assert, allow us to deduce a successful course of action from the specification of a desired final state. Reverse engineering may allow us to reconstruct the process of manufacturing an existing artefact; but a succesful artefact is a resolved ambiguity (or cluster of ambiguities), and we have the evidence of its resolution to guide our reconstruction. As Perkins (p. 160) somewhat hesitantly reports, it may also be possible to simulate the path to an achieved scientific discovery, for there is always retrospectively a pathway to current knowledge; but such success does not provide a procedure for deducing fresh knowledge, because there are many divergent pathways from established ideas and many ways of linking ideas, and which path seems worth following depends on conjectured contexts of similarity. Connections have to be invented – even in physics, where the pervasiveness of a particular class of connections is a requirement of good theory. (On the significance of the treatment of connections in the analysis of systems, see Potts 2000). The development of science may be presented to students as a logical progression; but the logic is typically available only in retrospect. There is no better example of this than the centuries-old search for a proof of Fermat's last theorem (Singh 1997).

As Nightingale (2000, p. 352) observes, Bradshaw's paradox, that we 'need to know the biological results before we can decide on the appropriate space to represent our compounds' (Nightingale 2000, p. 337) applies to the whole innovation process; indeed the optimal decomposition of any complex problem can be discovered only by solving the problem (Marengo 2000). Knight's principle of supposedly-relevant similarities, exemplified by scientific and social scientific theories, design trajectories, recognised good practice (as explained in the chapters by Constant and Stankiewicz on recursive knowledge and design spaces), and many other institutional aids to cognition, enable us to do far better (most of the time) than random speculation; but all these forms of ex-ante selection need to be reinforced, modified, and sometimes superseded by ex-post selection in order to achieve successful outcomes.

Thus, as Nightingale points out, there is no prospect of using recent advances in medical knowledge to deduce the specification of a safe and effective drug from a definition of desired effects. Problems are defined by differences, but the search for solutions is guided – sometimes in wrong directions – by the perception of similarities with existing solutions. The significance of these advances, as Nightingale emphasises, is that they have created new 'contexts of similarity' (Nightingale 2000, p. 337) which have enabled pharmaceutical companies to refine their search for new compounds and to reduce the costs of search; but reliance on these contexts has not reduced the number of candidate compounds that it is thought necessary to screen, and 'there is little evidence that this is translating into improved performance' (Nightingale 2000, p. 351). Moreover, we should remember Knight's warning that if there were to be a standard procedure for attaining optimal outcomes, no-one could expect to make a sustainable profit from its use. Detailed agreement on the best way to organise research is more likely to reduce than enhance the profits of pharmaceutical businesses, unless they can also erect barriers against rivals. Diversity remains a general condition both for profit and the growth of knowledge; and the effect of system diversity on development is a basic evolutionary theme (Pavitt 1998, p. 439).

Diversity, especially when based on different ways of connecting perceptions, phenomena and ideas, entails significant problems of co-ordination between individuals and between groups; these problems are not ignored by Ziman and his fellow-authors, but they are underplayed. Alfred Marshall (1920, pp. 138-9) offers some helpful advice. 'Organisation aids knowledge; it has many forms, e. g. that of a single business, that of various businesses in the same trade, that of various trades relatively to each other, and that of the State providing security for all and help for many.' Different forms are required in order to accommodate diverse combinations of similar and complementary capabilities (Richardson 1972), each providing a context of similarity which permits some variety within each combination as a consequence of differences in temperament, associations and experience which define manageable problems (Marshall 1920, p. 355), and allows people to draw on vicarious experience both to evaluate and to modify their ideas. Moreover, the most effective forms, both of internal structure and external relationships, change over time, largely as a consequence of their own effects; this theme was most forcefully expounded by Allyn Young (1928).

Knowledge itself is organisation, produced by trial and error, and always subject to challenge, including changes in its form and relationships to other bodies of knowledge. As many examples in this book illustrate, an artefact is a social construction, and so is any particular piece of knowledge; in general, both might have been somewhat different from what actually exists. That does not mean that we can construct any *reliable knowledge* (Ziman 1978) or reliable artefacts that we like, for of all the artefacts and knowledge structures that are conceivable, only a very small proportion will actually work. Organisation aids knowledge by guiding us towards this small proportion, though it can also lead us astray. That is why we need appropriate organisation, procedures, motivation and imagination both for selection and the continued generation of variety.

A biological foundation for the characteristics of human cognition

Since the development of artefacts is an expression of the development of human knowledge, especially of knowledge how, an understanding of the development of human knowledge deserves consideration as a basis for understanding the process of technological innovation – not least because the power and fallibility of human imagination and human calculation seem to correspond to the remarkable successes and myriad failures of technology. It is uncertainty – the unlistability of possibilities and the absence of any procedure, known to be correct, for assessing and evaluating those possibilities which are listed – that both warns of the likelihood of failure and creates the alluring prospect of extraordinary success. Progress in both knowledge and technology therefore depends on the diversity of individual initiative, but also on the relationships, formal and informal, between individuals; for every one of us knowledge consists of the organisation of categories and the relationships between them, and the organisation of people into categories and relationships, if appropriately managed, aids the development and use of knowledge in society.

I have drawn on economists whose names are much better known than their work, even by economists, to present an account of the growth and organisation of knowledge which I believe is applicable to technological innovation, and I have argued that models of biological evolution are not particularly useful for this purpose. However, models of biological evolution may be very helpful in understanding the cognitive characteristics of the biological creatures who produce technological innovation, and it is therefore appropriate to consider briefly the consistency of the argument that I have presented with current biological understanding of evolved human capabilities. This consistency has become much clearer with the shift of emphasis from the simple artificial intelligence model of the human brain as a serial, logical processor in favour of a conception of multiple neural networks which appears entirely compatible with the Smith/Knight theory that the growth of intelligence is driven by the imperative of coping with situations that are not amenable to logical solutions. That this intelligence, as Smith and Knight believed, relies more on connecting principles than on formal logic is suggested by the wide range of experimentation by psychologists with versions of the Wason test, in which subjects are asked to identify evidence which is relevant to the refutation of a simple proposition. These experiments have produced abundant evidence of very poor performance when the test is presented in the most abstract form, in which the underlying logical structure should be most apparent, and far better performance when the test is presented in contexts which are more complex but with which the subjects are familiar. The human brain appears to recognise similarities much more readily than logical implications.

By considering the environment in which the human brain has evolved, it is possible to trace a plausible biological pathway to a brain with such characteristics. The evolutionary success of our predecessors – over many millions of years – was promoted by rapid identification of threats and opportunities, closely linked to effective and specific responses to each; and identification and response rested on the close co-ordination of sensory impressions and physical activities. In the early stages of animal evolution, locally-appropriate networks were genetically programmed – as is still true of many of the neural networks that regulate human activity; from this apparently secure basis later mutations produced programmes for the development of new networks in response to new threats and opportunities. This differentiation of function between networks is a straightforward application, recognised by Milne-Edwards (1827), of Adam Smith's principle of improved performance as a consequence of the division of labour, more easily achieved by this means than by incremental adaptations towards a general-purpose logical processor.

The development of an architecture of the brain which facilitated the creation of neural networks necessarily preceded the emergence of conscious thought; Hayek's (1952) account of the formation of our sensory order, formulated at the outset of his career, is a remarkable anticipation of this model of evolutionary psychology. By the normal rules of biological evolution, conscious thought was similarly built upon contexts of similarity rather than logical processes; Hayek explains why these may be different contexts. This sequence, from connections between impressions and actions to connections between ideas of impressions and actions, including the imagination of possible connections, was conjectured in Alfred Marshall's (1994) early paper 'Ye Machine', which predates his interest in economics but may well have had some influence on his understanding of economic processes, as argued by Tiziano Raffaelli (2001), who is primarily responsible for recognising the significance of this work. The ability to construct logical sequences is a relatively recent and relatively weak development, almost an 'artificial' form of intelligence, and its effectiveness depends on the prior creation of appropriate categories, as has been repeatedly – and sometimes spectacularly – demonstrated. Of the early computer manufacturers, IBM alone created the categories which enabled it to identify a market; its astonishing record of success subsequently trapped it within this categorisation, and the policies which followed logically from it, when a very different categorisation became necessary for effective reasoning.

To the genetic endowment of a set of behaviours was therefore added the genetic capability of developing a set of behaviours, out of a very large potential, by selecting connections in response to perceptions of phenomena, together with the emotional impulse to develop particular parts of this potential; and if differences of interest and situation lead members of a population to develop different parts of this potential, then the capabilities of that population may far exceed what even the most gifted individual can attain. It is this characteristic of human cognition that underlies Smith's recognition of the crucial importance of the division of labour as an evolutionary process. It has its own pathology, not least in technological innovation; yet this combination of capabilities and motivation has made possible a non-biological evolutionary process that has operated much faster and encompassed unprecedented categories of applications. These applications may even include manipulation of the genetic evolution that made it possible, thus reversing the hierarchy of causation. One may claim, with Cosmides and Tooby (1994), that the mental capabilities that have resulted from our biological evolution are 'better than rational' for coping with the range of problems that lie between randomness and the economists' concept of rationality. These certainly include the problems of technological innovation.

Part III A personal history

It was as a way of dealing with problems of knowledge, not by reference to biology, that I became interested in concepts which may be categorised as evolutionary, though it was a good many years before I came to apply that label. This difference in 'connecting principles' is my basic explanation for the puzzle which gave direction to this paper, and its title encapsulates this alternative connection to the growth of knowledge. My personal intellectual history may itself be considered as an evolutionary process, incorporating path-dependency without path-determinism, and illustrating the interaction of design, conjecture, selection (both ex-ante and ex-post) and unanticipated consequences in shaping the ways in which we attempt to make both sense and artefacts.

As an undergraduate at Cambridge (UK) I noticed that major disagreements between Cambridge economists appeared to have no effect on the confidence with which conflicting theories were asserted by most disputants, and I observed that the theory of imperfect competition required decision makers to use knowledge which they could not possibly obtain. I gradually realised that there was something not quite right about both the knowledge that economists claimed for themselves and the knowledge that most of them attributed to economic agents. In a subsequent empirical study of location decisions, in which I visited factories and talked to owners and managers – not a common practice among economists either then or now – I learned first to ask how people came to think about moving, and discovered that the main benefit of almost all moves, which turned out to be increased efficiency through improved organisation, was rarely part of the initial motivation, which was predominantly the sharply defined problem of insufficient space to cope with a growing business.

In seeking to explain this decision process and its unanticipated outcome, I borrowed from a colleague, David Clarke, his concept of a decision cycle, which extended Herbert Simon's three-stage sequence of intelligence – search – choice into a five-stage sequence by adding implementation and control, the last being a major source of intelligence for new decisions, and so dynamising the analysis in a way that subsequently proved readily applicable to the management of innovation (as illustrated in the following paragraph). Here we have both ex-ante and ex-post selection, both of which may be fallible, and the generation of variants, typically by sequential search, in response to a particular stimulus, though it did not occur to either Clarke or myself to cast our thoughts in such proto-evolutionary terms – a rather obvious failure in retrospect, as so many failures are. I was thus very receptive to Pounds' (1969) analysis of 'The process of problem finding', which I have already mentioned, and subsequently used it to illuminate my own exposition of a location decision in which operational problems were effectively defined and resolved, leading to a substantial reduction in the workforce, while potentially crippling industrial relation problems resulting from the prospect of large-scale redundancies were avoided by not defining them (Loasby 1973).

I subsequently had the great good fortune (this is a path-dependent story of its own, which is summarised in Loasby 1989, pp. xiii-xiv) to work with people who had been, or still were, engaged in the management of innovation, notably Frank Bradbury and Charles Suckling, and soon picked up such organising principles as 'many starters, few finishers', and 'chemistry is full of surprises'. (I also recall Frank Bradbury's description of the management of a particular project as 'lurching round and round the decision cycle, cannoning off the constraints'). Neither failures nor new ideas could be predicted; yet it was possible to act intelligently in the space between optimal choice and selection from random variations. The key to operating in this space, I gradually realised, was appropriate organisation, both of activities and knowledge; and I also realised that this key had already been identified by many people in different contexts. My tentative conclusion from the location studies that selection was mainly rejection was greatly strengthened, but modified by the recognition that the reasons for rejection could guide redesign. The close interaction between the generation and winnowing of variety, linking ex-ante and ex-post selection, the definition of problems and attempts to resolve them, in technological innovation – so different from the biological model – is apparent in many of the chapters of this book (and appropriately highlighted in Keith Pavitt's review), but it nowhere receives specific attention as a major characteristic.

This sequence of experiences shaped the development of *Choice, Complexity and Ignorance* (Loasby 1976), which (like many technological innovations) turned out to be substantially different from the original idea. Its index, however, contains no entry for 'evolution' – and I compiled the index myself. The sequential, and often cyclical, process of identifying problems, searching for solutions, choosing options, and reacting to outcomes, which constituted a major theme, was related to decision-making, especially about innovation, rather than evolution. It did not occur to me to look for biological analogies, even though plant breeding was one of my examples of directed search with unexpected results. Instead I looked to Popper, and the central idea of the growth of knowledge through conjecture and exposure to refutation, where the objective, though not the content, of conjecture is chosen and the refutation of any conjecture is itself is a matter for decision, because unsatisfactory outcomes can always be blamed on something other than the theory which is supposedly under test. (John Ziman (1978) and I independently cited the example of the missing neutrinos to illustrate the reasoning behind such decisions.) Kuhn's (1962, 1970) idea of a research community, defined by a general agreement on the principles by which work will be guided, was also important, not least because of the resemblance to the decision premises that have been identified by Simon as a response to bounded rationality and a means of providing coherence within a formal organisation. This linkage seems to have crowded out a link to evolutionary biology, as is the way with apparently powerful linkages.

What led me to begin using the term evolution was a conversational remark by Andy Van de Ven, of the University of Minnesota, that one cannot have change unless there is something that does not change; this prompted me to start thinking explicitly about the relationship, not of contrast but of complementarity, between the absence of change and the changes which are thereby made possible, which was implicit in my earlier work, and to label the twin elements of this relationship equilibrium and evolution. Kuhn's (1962, 1970) emphasis on discontinuity which evaded explanation consequently appeared as an obstacle rather than an aid to understanding (in contrast to Smith's explanation of the transition between theoretical systems), though incommensurability, explained in his second edition as an aggregation problem, remained helpful in explaining the diversity of reactions to new ideas. This complementarity provided the framework for my Manchester Lectures of 1990, published as *Equilibrium and Evolution* (Loasby 1991) and have remained at the core of my thinking about evolutionary processes, most recently in thinking about the cognitive basis of human knowledge (e.g. Loasby 1999a). Reasoning is always bounded, though the bounds are not immutable; and 'rationality', in the sense of rigorous logic, is not the only kind of reasoning – in particular it is not sufficient for novelty, of thought or artefact, which depends on making new connections.

References

Alchian, Armen A. (1950) 'Uncertainty, evolution and economic theory', *Journal of Political Economy*, 58, pp. 211-21.

Cosmides, Lena and Tooby, John (1994) 'Better than rational: evolutionary psychology and the invisible hand', *American Economic Review*, 84, pp. 327-32.

Frisch, Otto (1979) *What Little I Remember*, Cambridge: Cambridge University Press.

Hayek, Friedrich A. (1952) *The Sensory Order*, Chicago: University of Chicago Press.

Kelly, George A. (1963) *A Theory of Personality: The Psychology of Personal Constructs*, New York: W.W. Norton.

Knight, Frank H. (1921) *Risk, Uncertainty and Profit*, Boston: Houghton Mifflin.

Krugman, Paul (1999) ‘What economists can learn from evolutionary theory – and vice versa’, in John Groenewegen and Jack Vromen (eds) *Institutions and the Evolution of Capitalism: Implications of Evolutionary Economics,* Cheltenham and Northamption, MA: Edward Elgar, pp. 17-29.

Kuhn, Thomas S. (1962, 1970) *The Structure of Scientific Revolutions*, 1st and 2nd edns, Chicago: University of Chicago Press.

Loasby, Brian J. (1973) *The Swindon Project*. London: Pitman.

Loasby, Brian J. (1976) *Choice, Complexity and Ignorance: An Inquiry into Economic Theory and the Practice of Decision Making*, Cambridge: Cambridge University Press.

Loasby, Brian J. (1978) ‘Whatever happened to Marshall’s theory of value?’, *Scottish Journal of Political Economy*, 25 (1), pp. 1-12.

Loasby, Brian J. (1989) *The Mind and Method of the Economist*, Aldershot: Edward Elgar.

Loasby, Brian J. (1990) ‘The firm’, in John Creedy (ed.) *Foundations of Economic Thought*, Oxford: Basil Blackwell, pp. 212-23.

Loasby, Brian J. (1991) *Equilibrium and Evolution*, Manchester, Manchester University Press.

Loasby, Brian J. (1996) ‘The division of labour’, *History of Economic Ideas*, IV (1-2), pp. 299-323.

Loasby, Brian J. (1998) ‘Co-ordination failure in economic theory: economists in the 1930s’, in Albert Jolink and Philippe Fontaine (eds) *Historical Perspectives on Macroeconomics: 60 Years after the General Theory*, London: Routledge, pp. 72-87.

Loasby, Brian J. (1999a) *Knowledge, Institutions and Evolution in Economics*, London and New York, Routledge.

Loasby, Brian J. (1999b) ‘Marshall’s theory of the firm’, in Roger E. Backhouse and John Creedy (eds), *From Classical Economics to the theory of the Firm: Essays in Honour of D. P. O’Brien*, Cheltenham: Edward Elgar, pp. 175-93.

Loasby, Brian J. (1999c) ‘Edith Penrose’s place in the filiation of economic ideas’, *Economies et Sociétés – Cahiers de l’ISMEA*, XXXIII, 8 *Série Oeconomica*, pp. 103-21.

Marengo, Luigi (2000) ‘Decentralisation and market mechanisms in problem solving’, paper presented to DRUID Conference, Rebild, Denmark 15-17 June.

Marshall, Alfred (1920) *Principles of Economics*, 8th edn., London: Macmillan.

Marshall, Alfred (1994) ‘Ye machine’, *Research in the History of Economic Thought and Methodology, Archival Supplement 4*, Greenwich, CT: JAI Press, pp. 116-32.

Milne-Edwards, Henry (1827) ‘Nerf’, in M. Bory de Saint-Vincent (ed.) *Dictionnaire Classique de l’Histoire Naturelle*, Paris: Rey et Gravier.

Myerson, Roger B. (1999) ‘Nash equilibrium and the history of economic theory’, *Journal of Economic Literature*, 37, pp. 1067-82.

Nelson, Richard R. and Winter, Sidney G. (1982) *An Evolutionary Theory of Economic Change*, Cambridge, MA: Belknap Press.

Nightingale, Paul (2000) ‘Economies of scale in experimentation: knowledge and technology in pharmaceutical R&D’, *Industrial and Corporate Change*, 9, pp. 315-59.

Pavitt, Keith (1998) ‘Technologies, products and organization in the innovating firm: what Adam Smith tells us and Joseph Schumpeter doesn’t’, *Industrial and Corporate Change*, 7, pp. 433-52.

Penrose, Edith T. (1952) ‘Biological analogies in the theory of the firm’, *American Economic Review*, 42, pp. 804-19.

Penrose, Edith T. (1959) *The Theory of the Growth of the Firm*, Oxford: Basil Blackwell.

Pool, Robert (1997) *Beyond Engineering: How Society Shapes Technology*, New York and Oxford: Oxford University Press.

Potts, Jason (2000) *The New Evolutionary Microeconomics: Complexity, competence and Adaptive Behaviour*, Cheltenham and Northampton, MA: Edward Elgar.

Pounds, William F. (1969) The process of problem finding’, *Industrial Management Review*, 11, pp. 1-19.

Raffaelli, Tiziano (2001) ‘Marshall on mind and society: neurophysiological models applied to industrial and business organization’, *European Journal of the History of Economic Thought* (forthcoming).

Richardson, George B. (1972) ‘The organisation of industry’, *Economic Journal*, 82, pp. 883-96.

Schlicht, Ekkehart (2000) ‘Aestheticism in the theory of custom’, *Journal des Economistes et des Etudes Humaines*, 10: 1, pp. 33-51.

Schumpeter, Joseph A. (1934) *The Theory of Economic Development*, Cambridge MA: Harvard University Press.

Schumpeter, Joseph A. (1939) *Business Cycles*, New York and London: McGraw-Hill.

Shackle, George L. S. (1972) *Epistemics and Economics*, Cambridge: Cambridge University Press.

Shackle, George L. S. (1979) *Imagination and the Nature of Choice*, Edinburgh: Edinburgh University Press.

Singh, Simon (1997) *Fermat's Last Theorem*, London: Fourth Estate.

Smith, Adam ([1759] 1976a) *The Theory of Moral Sentiments*, ed. David D. Raphael and Alec L. Macfie, Oxford: Oxford University Press.

Smith, Adam ([1776] 1976b) *An Inquiry into the Nature and Causes of the Wealth of Nations*, ed. Roy H. Campbell, Andrew S. Skinner and W. B. Todd, Oxford: Oxford University Press.

Smith, Adam ([1795] 1980) 'The principles which lead and direct philosophical enquiries: illustrated by the history of astronomy', in *Essays on Philosophical Subjects*, ed. W. P. D. Wightman, Oxford: Oxford University Press, pp. 33-109.

Smith, Adam (1983) *Lectures on Rhetoric and Belles Lettres*, ed. J. C. Bryce, Oxford: Oxford University Press.

Smith, John Maynard (1996) 'Conclusion', in W. G. Runciman, John Maynard Smith and R. I. M. Dunbar (eds) *The Evolution of Social Behaviour Patterns in Primates and Man*, Oxford: Oxford University Press for the British Academy, pp. 291-7.

Young, Allyn (1928) 'Increasing returns and economic progress', *Economic Journal*, 38, pp. 527-42.

Ziman, John M. (1978) *Reliable Knowledge*, Cambridge: Cambridge University Press.