THE EVOLUTION OF KNOWLEDGE

Brian J Loasby

DRUID Conference

Aalborg

12-15 June 2001

11 May 2001

Department of Economics
University of Stirling
Stirling
Scotland
FK9 4LA

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

Theme

Economic development, of which technological change is clearly an important part, may be very appropriately represented as an evolutionary process, which requires an approach substantially different from that of mainstream economics. The growth of knowledge, including the knowledge embodied in 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. However, there are major differences between the two processes 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: 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, the effective 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; and 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.

Evolution, rationality and knowledge

Let me begin by specifying what I understand by the term 'evolution': it is a process, or cluster of processes, which combines the generation of novelty and the selective retention of some of the novelties generated. This basic definition is sufficient to distinguish evolutionary theories from theories that are clearly not evolutionary, while allowing us to identify distinctive kinds of evolutionary theory; in this paper I intend to do both, first differentiating evolutionary from non-evolutionary economics and then arguing that evolutionary economics has important differences from the biological model, which may however provide a useful complement to it.

Neo-Darwinians claim that the only alternative to their 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 Adam Smith's economics), and it is deemed sufficient for standard economic analyses of technical change, although these analyses trivialise change. 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. That is why they are so successful in designing their plans that no revision is necessary.

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. It is true that in both neoDarwinian theory and in standard economics selection is determined by consequences; but whereas in neoDarwinian theory mutations must be introduced into the environment in order to discover what their consequences are, in standard economics the consequences of any contemplated action can be correctly deduced in advance (if necessary, as a probability distribution). Nothing can happen that was not provided for. The criteria of rigorous theorising in economics require the set of possible actions and the set of possible states of the world to be complete, and known to be so; information is problematic only when access to this information is costly, and even then it is optimally seleced, essentially by choosing the basis and fineness of its partitioning. Agents can 'learn' only in the form of obtaining increasingly accurate estimates of the likelihood of each possibility by a procedure that is fully specified in advance; nothing will ever induce agents to envisage new possibilities. Time is incorporated into models as an additional dimension of goods and information sets, to the exclusion of any analysis of an economy which develops in time. Thus explanation by design, incorporating efficient ex-ante selection, 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. (For an extensive treatment, see Vromen 1995.) 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. Evolutionary forces in economic systems are so effective that evolutionary arguments are superfluous in economic reasoning, 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. 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).

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'. 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). As experienced by humans, uncertainty results from the interaction of two fundamental factors. The factor that is sometimes recognised by economists may be called Herbert Simon's problem, though this is better defined as 'bounded cognition' than as 'bounded rationality', which unfortunately lends itself to interpretation as maximisation in the presence of information costs, whereas bounded cognition includes both the limitations of our logical powers and the need to impose, rather than deduce, simplifications in our representations and short-cuts in decision-making. 'Making full use of all relevant information', far from being a definition of rational behaviour, is rarely an option. What is even more fundamental, but totally ignored in both general equilibrium and game-theoretic versions of rational choice theory, is David Hume's problem: there is no way of demonstrating the truth of any general empirical proposition, either by deduction, for there is no way of ensuring the truth of the premises, or by induction, for there is no way of proving that instances not observed would correspond to those that have been observed.

The interaction of Simon's and Hume's problems is most acute in situations of complexity. The need for simplification is generally recognised, but what is not generally recognised, but can be proved by logical reasoning (Marengo 2000), is that a logically impeccable simplification must begin with a full representation of the complex system – which Simon and Hume have both shown is unobtainable. (Hayek (1952, p. 185) makes this argument in *The Sensory Order*, to which we shall refer later). Our only recourse is to start at the other end, with a simple view, and to add complexity as we are impelled to do so and to the extent that we are capable of doing so. That our systems of thought rely on connections that we have invented, or adapted from their inventors, and not those which have been optimally chosen from a complete set, is, I suggest, a fundamental principle of evolutionary economics (see Potts 2000). Human knowledge is a self-organising system.

Since 'information', far from being self-contained truth, acquires meaning from its context, the implications of complexity are highly pervasive. Recourse to 'science' does not resolve these difficulties, for science also depends on 'the linkage of known empirical phenomena into a wider network of accepted – or at least potentially acceptable – "facts" and concepts' (Ziman 2000, p. 291). Of course, 'the production of such linkages is the main business of research', and scientists have developed elaborate procedures for testing them, using logical reasoning to identify testable implications of candidate hypotheses; but there is no scientific proposition which is in principle beyond challenge, and any investigation achieves precision at the focus of attention only by 'ignoring the lack of definition as we approach the edges of the image' (Ziman 2000, p. 291).

The epistemology of science is based on uncertainty (Ziman 2000, p. 328), and scientific journals publish many papers which are later dismissed as embodying false hypotheses or inadequate tests; yet over time scientific methods have produced a remarkable growth of knowledge which may reasonably be treated as reliable (Ziman 1978). Reliability is a consequence of appropriate and distinctive operating practices, based on the Mertonian norms of communalism, universalism, disinteredness, originality, and scepticism – which, it should be noted, are responses, not always explicit, to the impossibility of final validation, and to the corresponding dependence on intersubjective appraisals to produce structures of knowledge that are cognitively objective (Ziman 2000, pp.303, 316). These norms are not reducible to any conventional notion of rationality, though they do embody more than an echo of Adam Smith's ([1759] 1776a) *Theory of Moral Sentiments*. Yet although it is possible to make meaningful distinctions between science and what Ziman calls 'life-world knowledge', science has evolved from earlier attempts to make sense of the human situation, and continues to rely on some knowledge which has not been formally brought within the corpus of scientific analysis (Ziman 2000, p. 297). This permeable boundary and the continuing resemblances across it offer some prospect for an investigation by social scientists of the behaviour of people who are not scientists, but who are not rational maximisers either.

Responses to uncertainty include various kinds of coping strategies, including strict observance of routines, contingent application of rules (or of Simon's rather wider category of 'decision premises'), and taking precautions such as the building of reserves, including reserves of goodwill from earthly or heavenly powers. But they also include responses of a very different kind; for the obverse of uncompletable knowledge is the scope for imagination. Uncertainty can be exploited as well as endured. Thus people may generate novel options and imagine new contingencies, make novel selections among alternative hypotheses and embody some of them in artefacts and many in institutions which form the premises for decision-making; on observing the outcomes they may select among them, according to the theories by which they impute causality and their criteria for what is desirable (which are not merely the criteria that drive biological processes, though the relationship between the two kinds of criteria is a legitimate subject for investigation). These ex-post selections may lead to the generation of further hypotheses, sometimes (especially in technological innovation) in a closely-coupled way. Such evolutionary processes are likely to be 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.

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 or the development of human knowledge, though it does encourage us to postulate stable genetic characteristics in the human population over periods which are by comparison extremely brief. 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. That the only reliable baseline in general equilibrium theory is the set of general equilibrium prices is the fundamental reason why there can be no theory of rational choices in disequilibrium – a proposition advanced by Richardson (1960) and Leijonhufvud (1968); that this proposition has not been universally recognised is indicated by the prospectus, which I recently received, for a three-day programme for ‘finance professionals’ which claims to show how uncertainty, which is differentiated from risk, can be exploited by calculating Arrow-Debreu prices. The absence of a reliable baseline also explains why many game-theoretic models yield multiple solutions. 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 particular combinations of persistence and change on innovative sequences.

Neo-Darwinians insist that in any evolutionary process there should be only one unit of selection: genes and ‘memes’ are assigned to distinct processes. I see no reason to accept this principle for evolutionary economics: 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 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

For our foundation model of the growth of knowledge as an evolutionary process we cannot do better than Adam Smith's psychological theory of the emergence and development of science, which he illustrated by the history of astronomy (Smith [1795] 1980). Whether consciously or not, Smith's theory combines two distinctive themes from his friend David Hume. The first, which we have already encountered, is the impossibility of proving empirical truths. The second is the rejection of the supremacy of reason, which 'alone can never produce any action' (Hume 1978, p. 414) and therefore 'is, and ought only to be, the slave of the passions, and can never pretend to any other office than to serve and obey them' (Hume 1978, p. 415). Thus rational choice is an inadequate explanation for behaviour, because neither the empirical premises nor the objectives of behaviour can be logically derived. Standard economics scorns to explain why economic agents should be so keen on maximisation: it is not obvious why people who are equally content with every point on the highest attainable indifference curve should be unwilling to tolerate any point which is fractionally below it. Maximisation therefore fails, both as a method – because the premises of any calculation are always problematic, and as an objective – because it lacks adequate motivation.

The practical significance of these twin deficiencies are examined (without reference to Hume) in Chester Barnard's (1938) lecture on 'Mind in Everyday Affairs', which ought to be read by every microeconomic theorist and every policy analyst. Let me just note Barnard's (1938, p. 308) observation that the insufficiency of reason when action is required 'is probably why it is difficult to make correct decisions without responsibility'. In this, as in many other ways, organisation changes behaviour – which is the underlying theme of Barnard's book; more generally, as Potts (2000) argues, the performance of any system cannot be reduced to the performance of its elements but is influenced by the particular nature of the connections between them. Now if maximisation is insufficient to deal with what are represented in economics as allocation problems, it can hardly serve to explain either the decision to seek new knowledge or the process of seeking it. Indeed, the search for novelty presents another challenge to rationality, for 'no kind of reasoning can give rise to a new idea' (Hume 1978, p. 164). It is therefore no accident that explanations of entrepreneurship go beyond maximisation: Kirzner's entrepreneurs are alert, and Schumpeter's have powerful non-pecuniary incentives.

Evolutionary economics, being concerned with the development of knowledge through purposeful behaviour, requires attention to the motivation for generating new ideas; and this is the first element that Smith provides. He begins his exposition by identifying three passions of wide application, arguing that people are disturbed by the unexpected, dismayed by the inexplicable, and delighted by schemes of thought that resolve the inexplicable into plausible generalisations, before claiming that, 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 in Smith's theory is the sequence that is inspired by this complex motivation: the generation of novelty and the ex-ante 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 of good explanations of economic phenomena; and we are now identifying the connecting principles of Smith’s own explanation of the growth of knowledge.

Smith’s third element, which is already implicit in the reference to notions of good order, is the link between emotion and aesthetics. 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 often of major significance. Sometimes aesthetic appeal is a major objective; 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.

The fourth element in Smith’s proto-evolutionary theory is his proposition that 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. These powerful motivations, together with the underlying similarity in human mental, emotional and aesthetic processes which underpin them, are foundational principles 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 both to technological evolution and any adequate understanding of organisational behaviour.

However, because by Hume’s argument invented principles, however widely accepted, are not proven truth – even, Smith ([1795] 1980) 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 of ex-post selection, though it may be the evidence rather than the explanatory system which is rejected; when the evidence is recalcitrant, and satisfactory adaptation is despaired of, a new search for connecting principles begins. This is the fifth element, which renews the evolutionary sequence.

The sixth element in Smith's system is the evolution of the evolutionary process itself. 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 deserves recognition as a seventh element of Smith's evolutionary theory.

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 that a Malthusian struggle to survive would result in the differentiation of species (Raffaelli 2001). The other basic elements in Smith's account of the development of knowledge by motivated trial, error, amendment and diffusion understandably did not. We may therefore suggest that Smith provides a better basis for evolutionary economics than biological models; we may also observe that different analytical systems, focussing on different patterns of connections, may be most effective in developing different kinds of knowledge, thus explaining the value of speciation among academic disciplines.

The differentiation of knowledge, and of kinds of knowledge, is a condition of progress in human society. However it has its opportunity costs, of which I will mention two. 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. 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). This is 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

The double-edged character of uncertainty is the focus of Frank Knight's *Risk, Uncertainty and Profit* (1921), and of much of Shackle's work. 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 therefore be made by a standard procedure which can be demonstrated to be optimal, but which cannot be a source of sustainable profit. For conditions of uncertainty, however, no demonstrably optimal procedure can be devised; we must act in the space between optimality and randomness.

But 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).

Knight observes that in a world without 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 (as does Barnard) 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. (The biological basis for this reliance is explained in the final section of this paper.) For differently-formulated problems, we tend to rely on different contexts of similarity, and for new problems we experiment with new connections that define new contexts: that is how the division of labour leads to differentiated knowledge. If we succeed in making new connections which constitute new knowledge, these connections will provide us with new rules and routines, releasing cognitive capacity for new applications, as in Penrose's (1959) conception and use of 'the receding managerial limit'.

By insisting on contexts of incomplete similarity Knight preserves uncertainty at the core of his analysis. 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? However this ambiguity, though often a problem and indeed a source of error, occasionally catastrophic, is also an opportunity for the generation of ideas by making new combinations (a principle enunciated by both Smith and Schumpeter). 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) emphasis on the need to perceive, by non-logical means, how resources may be directed towards 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. 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 perceptions of the applicability of both theoretical and practical knowledge to novel contexts.

Category-based judgements of possibility, guided by the institutions of individual cognition, a research group, or an organisation, lead and direct 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. 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. 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).

Nightingale (2000, p. 352) observes that 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. Knight's principle of supposedly-relevant similarities, exemplified by scientific and social scientific theories, design trajectories, recognised good practice, 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 are themselves conjectures which need to be reinforced, modified, and sometimes superseded by ex-post selection in order to achieve successful outcomes.

Thus, as Nightingale (2000, p. 337) points out, the significance of recent advances in medical knowledge is that they have created new contexts of similarity or institutional frameworks for innovation, which have enabled pharmaceutical companies to refine their search for new compounds and to reduce the costs of search. However they do not allow the specification of a safe and effective drug to be deduced from a definition of desired effects, and so 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 should 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 whose knowledge is differently ordered. 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.' The organisation of each firm privileges a particular set of connections between people who each tend to rely on somewhat idiosyncratic contexts of similarity, thus producing, at individual, sub-unit and firm level, somewhat distinctive responses to information and generating somewhat different conjectures. (The 'administrative framework' is crucial to Penrose's (1959) account of a firm's development.) Different forms of organisation are required in order to accommodate diverse combinations of similar and complementary capabilities (Richardson 1972) while permitting some variety within each combination as a consequence of differences in temperament, associations and experience which define manageable problems (Marshall 1920, p. 355). 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). As Young maintained, increasing returns are not a property of equilibrium, but of a process of developing new patterns of connections within an economy; but we may extend Young's argument from changing the distribution of activities between firms to the distribution within each firm and indeed to the changing patterns of knowledge within each person.

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; it is a product as well as a precondition of decisions. Knowledge lies in the particular connections between elements, rather than the elements themselves; this is a concept foreign to microeconomics, in which connections are assumed to be complete except when the absence of a particular connection is identified as a source of market or organisational failure. Since technological innovation is an expression of the development of human knowledge, especially of knowledge how, an understanding of human knowledge provides a basis for understanding 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 the combination of uncertainty – the unlistability of possibilities and the absence of any procedure, known to be correct, for assessing and evaluating those possibilities which are listed – and the evolved characteristics of human cognition that both warns of the likelihood of failure and creates the alluring prospect of extraordinary success, as well as explaining our reliance on institutions. Progress in both knowledge and technology depends on the diversity of individual initiative, but also on the relationships, formal and informal, between individuals; for every one of us, as well as for the communities to which we belong, knowledge depends on 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.

A Biological Foundation

Although the absence of human purpose in the variety-generation process is sufficient reason for rejecting the neo-Darwinian biological model as a basis for analysing evolutionary processes in social and economic systems, nevertheless the model may reasonably be applied to the biological processes which have shaped the human beings whose behaviour we wish to study. Indeed, models of biological evolution may be very helpful in understanding the cognitive characteristics of biological creatures who are capable of producing true novelty, and yet are so dependent on rules and routines; 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 a shift of attention 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 results of a 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. As previously noted, this facility for imposing similarities and a readiness to accept similarities which have been imposed by others, provided that they can be assimilated to some kind of association which is already familiar, are the psychological characteristics which underpin Smith's ([1795] 1980) account of both the creation and the spread of new knowledge, which includes a recognition of the obstacles to that spread among those for whom no such assimilation is possible – or in other words, who lack the relevant absorptive capacity.

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 predecessor species 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.

A generalised capacity for making patterns provides the potential for much more flexibility than a generalised pattern of behaviour – provided, of course, that there is sufficient motivation to create new patterns when appropriate. The importance of pattern-making activity suggests an obvious role for aesthetic sensibility as a motivator which has tended to enhance evolutionary success, and which has both stimulated and shaped scientific as well as other forms of knowledge. Smith's discussion of motivation is thus not only an essential feature of his theory, which cannot be adequately represented by the standard conception of incentives in modern economics, but an essential feature of satisfactory models of biological evolution. The differentiation of function between networks in a single brain 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, an improvement that is more easily achieved by this means than by incremental adaptations towards a general-purpose logical processor.

The neural structures of each person's brain are the product of two selection processes: genetic selection provides the basic architecture and some of the connections, and the process of making sense of the world which begins at birth creates a particular network of connections that imposes order on events. It is this imposed order, not the events themselves, that constitutes experience. 'We conceive a person's processes as operating through a network of pathways rather than as fluttering about in a vast emptiness. The network is flexible and it is frequently modified, but it is structured and it both facilitates and restricts a person's range of actions' (Kelly 1963, p. 49). The development of individual knowledge and skills is subject to path-dependency: the network supports routines of behaviour and rules for conceptualising and resolving problems, and it strives to preserve itself, sometimes by denying the validity of information. This is necessary for each individual, even when acting in isolation, despite the obvious dangers; for without firm anchors no intelligent variation is possible. This is another manifestation of the importance of a persistent baseline; some concept of stability, to which we might give the label 'equilibrium', is necessary to understand change.

Furthermore, the development of new networks embodying new knowledge typically relies on exaptation – the use, often with some modification, of existing structures for new purposes. What novelties are possible for any person at any point of time depends on pre-existing structures and the history of past adaptations; but these constraints – the evolved institutions of the brain – are rarely sufficient to be of much help in predicting novelty, except in a negative sense. Innovation may require the breaking of some established connections, but the new connections must be adequate substitutes for the old in forming complementary relationships with some preserved patterns. The new ideas and the old may be incommensurable in the straightforward sense of not being partitions of a single structure of knowledge, but there is no absolute break. The emphasis on discontinuity in Kuhn's (1962, 1970) account of paradigm change and Schumpeter's (1934) invocation of entrepreneurial vision is, in my view, more misleading than helpful; on the contrary 'good continuity' (Schlicht 2000) is important, and Schlicht has drawn attention to the importance of aesthetic factors in determining what is good continuity – which over a long period may resemble Wittgenstein's rope.

The development of an architecture of the brain which facilitated the creation of neural networks necessarily preceded the emergence of conscious thought, which did not displace these networks of unarticulated 'knowledge how'. It is therefore necessarily true that we know more than we can tell, and that codification must always rest ultimately on tacit knowledge. (That is not to deny the value of codification.) 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 by exaptation, conscious thought was similarly built upon connections rather than logical processes; Hayek explains why the connections of science may not correspond to the evidence of our senses.

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'; this 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.

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, is the biological precondition of modern economic systems; for 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. The fostering of differences in interest and situation by both formal and informal organisation encourages and shapes the development of these capabilities. 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 uncertainty and imagination; indeed we may agree with Rizzello (1999, p. xv) that '[t]he economics of the mind is the economics of creativity, uncertainty and complexity'.

Individual cognition, including the development of knowledge which is new to that individual, is governed, though not determined, by a dense network of rules and familiar relationships, many of them partly or wholly tacit. When these rules and relationships are shared within a community, we call them 'institutions', and many of the rules and relationships on which each of us relies are indeed institutions in this sense; but it is important to recognise that their origin, as a class of phenomena, lies not in the management of interactions but in the requirements of effective individual cognition. Indeed it is in this fundamental cognitive requirement that we can discover the possibility, as well as the incentive, for developing the institutions that guide interactions.

The fundamental principles on which human brains are organised gives some prospect of understanding at least aspects of other people's behaviour; and Heiner (1983) suggested that our dependence on rules which could not be precisely adapted to specific situations was what made possible the interpretation of other people's behaviour without knowledge of these specific situations. Choi (1993) has argued, drawing on Smith's *Theory of Moral Sentiments*, to which we have already referred, that this possibility of interpretation may allow us to conduct vicarious experiments by observing others and imitating apparently successful performance. This would lead to shared routines and some shared rules for choosing between routines, even when actions were quite independent. Sometimes these shared routines and rules would form a natural basis for concerted action, and when they did not the experience of shared routines might encourage the notion that co-ordination problems might be resolved by a new form of sharing: thus exaptation seems to be a promising clue to the explanation of institutions in the popular sense, and perhaps in particular to the general acceptance, usually tacit, of the idea that organisations are incubators of institutions. Among the significant organisations which depend on both individual and shared routines and rules are research communities, firms and markets. Without such institutions, economic evolution as we have experienced it would not be possible.

References

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

Barnard, Chester A. (1938) *The Functions of the Executive*, Cambridge MA: Harvard University Press.

Choi, Young B. (1993) *Paradigms and, Conventions:Uncertainty, Decision Making and Entrepreneurship*, Ann Arbor: University of Michigan Press.

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.

Heiner, Ronald A. (1983) ‘The origin of predictable behaviour’, *American Economic Review*, 73, pp. 560-95.

Hume, David (1978) *A Treatise on Human Nature*, ed. L. A. Selby-Bigge, 2nd edn. Revised by P.H. Nidditch, Oxford: Clarendon 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.

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

Leijonhufvud, Axel (1968) *Keynesian Economics and the Economics of Keynes*, New York and London: Oxford University Press.

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.

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. 3rd edn. Oxford: Oxford University Press, 1995.

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

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

Richardson, George B. (1960) *Information and Investment*. Oxford: Oxford University Press.

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

Rizzello, Salvatore (1999) *The Economics of the Mind*, Cheltenham and Northampton MA: Edward Elgar.

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.

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.

Vromen, Jack J. (1995) *Economic Evolution*, London: Routledge.

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

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

Ziman, John M. (2000) *Real Science*, Cambridge: Cambridge University Press.