



Edmund Chattoe (1998) 'Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes?'

Journal of Artificial Societies and Social Simulation vol. 1, no. 3,
<<http://jasss.soc.surrey.ac.uk/1/3/2.html>>

To cite articles published in the *Journal of Artificial Societies and Social Simulation*, please reference the above information and include paragraph numbers if necessary

Received: 15-Feb-98 Accepted: 13-Jun-98 Published: 30-Jun-98

Abstract

This paper attempts to illustrate the importance of a coherent behavioural interpretation in applying evolutionary algorithms like Genetic Algorithms and Genetic Programming to the modelling of social processes. It summarises and draws out the implications of the Neo-Darwinian Synthesis for processes of social evolution and then discusses the extent to which evolutionary algorithms capture the aspects of biological evolution which are relevant to social processes. The paper uses several recent papers in the field as case studies, discussing more and less successful uses of evolutionary algorithms in social science. The key aspects of evolution discussed in the paper are that it is dependent on relative rather than absolute fitness, it does not require global knowledge or a system level teleology, it avoids the credit assignment problem, it does not exclude Lamarckian inheritance and it is both progressive and open ended.

Keywords:

evolutionary algorithms, genetic programming, social evolution, selectionist paradigm

Introduction

1.1

There is now a significant body of work in social science (though predominantly in economics) which can legitimately be referred to as evolutionary. This paper focuses on simulations which use various kinds of evolutionary algorithms^[1] to represent social and particularly economic processes. There are now enough simulations of this kind that an "existence proof" has clearly been provided. It may now be appropriate to step back a little and try to provide a more general rationale for models of this kind. That is the object of this paper. An attempt to find a coherent interpretation for evolutionary modelling serves three purposes. The first is self evident. It is easier to pursue a research programme that has been set in some sort of order, albeit provisional and descriptive. The second is that evolutionary models have so far been carelessly criticised or ignored by orthodox economists and other social scientists.^[2] Many of the traditional criticisms rest on misinterpretations or take advantage of ambiguities in the interpretations used by researchers in the field. A coherent interpretation may also serve to improve the quality of subsequent debate. Finally, the process of building an interpretation reveals weak spots in the current simulations and may suggest both immediate solutions and areas for further long term research.

1.2

The [second section](#) of the paper describes the evolutionary process in biology as it is typically represented by social scientists. We also draw out some of the implications of evolution that are relevant to what follows. The [third section](#) summarises how these evolutionary ideas have been implemented in evolutionary algorithms of various kinds. The [fourth section](#) summarises the "standard" interpretation of evolutionary simulations which has emerged and the way in which it relates to concepts like procedural rationality in economics. The [fifth section](#) focuses on a number of specific simulations, including our own work in progress, to refine this interpretation, pointing out its weaknesses and suggesting areas for further development. The [sixth section](#) concludes.

1.3

A synthetic paper of this kind is inevitably both derivative and speculative. I have provided a quantity of references and footnotes that would normally be excessive in the hope of pre-empting some of the "traditional" criticisms. My goal is not to lay claim to very many of the basic ideas presented in this paper, only to putting them into a logical order and providing an extended critique which I hope will be provocative but helpful. (In addition to criticisms of prevailing approaches, I have also provided descriptions of alternatives wherever possible.) I am also trying to address two audiences and provide something of interest to both. The first audience are those who have not given too much thought to the possibilities of evolutionary models and the second are those who actually work in the field. Obviously keeping both groups interested is a rather delicate balancing act. This raises a final tension in my objectives, between current practice and future possibility. Although the paper focuses on examples from economics, the importance of a correct interpretation of evolution applies to modelling involving evolutionary algorithms in *all* the social sciences, even though in some it is barely used at present. [Runciman \(1998\)](#) discusses the profound lack of enthusiasm for evolutionary models in sociology and this attitude has proved so pervasive that no simulations (at least analogous to those described here) yet appear to exist in the discipline. Rather than providing a review (an exhaustive but essentially reactive description of existing work), this paper attempts to develop a thesis with appropriate illustrations. The thesis is that the coherent interpretation of social evolution (what Runciman calls the "selectionist paradigm") is both possible and desirable. For this purpose, the most appropriate illustrations are the best developed models which currently exist. However, it is hoped that in future models based on a coherent interpretation will increasingly be developed and applied in new areas. This paper is intended as a small contribution to that larger project.



Evolutionary Biology and Evolutionary Models

2.1

This section defines social evolutionary processes, summarises the mechanisms of evolutionary biology as they are typically used by social scientists and draws out their implications.

The Meaning Of "Evolutionary" In Social Science

2.2

The term "evolutionary" has been much misused and reviled in social science. Four major interpretations can be distinguished. Firstly, there are Social Darwinist theories ([Hofstadter 1955](#), [Jones 1980](#)), which are also associated with the ideas of eugenics. Such theories were often mere political rhetoric, and were so imprecise they could easily be used to support or oppose the same policy. They now have little more than historical interest since they made no attempt to justify the analogies they proposed between individuals, states or firms and individuals or species in biology.^[3] In some cases, they appear to have been based on misunderstandings of what Darwin and subsequent evolutionary biologists actually said. Secondly, there are Sociobiological theories ([Wilson 1975](#), [Wilson 1978](#)). These attempt to explain certain behaviours, traditionally regarded as

social, in genetic terms. Theories of this kind are not ignored here because they are of poor quality, but because it is generally assumed that there is a realm of distinctively social behaviour. Sociobiology may alter the size of this realm and force us to reconsider the explanations of certain behaviours at the margin, but it does not seem likely that it will swallow up the core of social science. Thirdly, there are theories which simply deal with change, including change in the very long term. Institutional economics has often been described as evolutionary in this sense ([Hodgson et al. 1994](#)). Our reasons for disregarding theories of this kind are more tentative. As with Social Darwinism, *some* of these theories use evolution in a sense that is more rhetorical than substantial ([Petr 1982](#)). They stress the (subjectively) functional nature of certain institutions and practices without saying much about how functionality might arise or how functional and non functional practices might be distinguished within a scientific framework. This is not to say that institutional economics has not made an extremely valuable contribution to unorthodox economics, nor that some models originating in the field are not evolutionary in the sense defined in this paper. The main reason for excluding most models of this kind is semantic and pragmatic. Institutional, historical or dynamic economics are perfectly appropriate terms to describe most of these models. By contrast, the models discussed in this paper would be almost impossible to describe in any other way. The final sense of evolutionary, and the one we shall use in this paper, refers to models or simulations which propose a detailed and coherent analogy between a biological evolutionary process and a social process.^[4] The biological evolutionary process chosen is usually that referred to as the Neo-Darwinian Synthesis ([Darwin 1968](#), [Fisher 1930](#), [Ridley 1985](#), [Edey and Johanson 1989](#)) which combines Darwin's theory of natural selection with Mendelian genetics.

The Neo-Darwinian Synthesis And Its Implications

2.3

For the purposes of this paper, this evolutionary theory can be summarised quite briefly. The objects of the theory are individual living creatures, which each consist of a *genotype* (a certain arrangement of genes) and a *phenotype* (the rest of the physical body which was "constructed" by that genotype during the process of gestation.) This construction process involves the genes "encoding" the synthesis of the enzymes and other molecules required to orchestrate the growth of an individual from a single cell to the point of birth. The environment, which consists of other individuals and physical structures, *selectively retains* individuals depending on the "performance" of their phenotypes.^[5] This performance involves persisting for long enough to reproduce, thus perpetuating the successful genotype. Two other important mechanisms take place during the reproductive process.^[6] In sexual reproduction, there is *genetic mixing*, the genes of the offspring contains some of the genes of *both* parents. In addition, the normal duplication of genes results in *mutation* through "copying errors", though this mutation can also result from external processes such as radiation. The functions of these two processes can be seen as the *generation* of novel combinations of genetic material (that may persist better in the environment than those which currently exist) and the *maintenance* of novelty. (We can envisage a situation in which, with a stable environment, genetic convergence will result. Mutation ensures that there is always a minimum of genetic diversity so that a subsequent change in the environment does not immediately result in the extinction of entire populations.)

2.4

For the discussion that follows, this process has several important and interdependent implications:

- Evolutionary fitness is based on relative and not absolute success. There is no requirement for an "invisible hand" or any global knowledge to exist within the system. Nor is there any requirement for a "minimum standard" of performance before the evolutionary process can begin. Each individual interacts only with other individuals (and physical processes) and persists or ceases accordingly.^[7]

- Evolution avoids the credit assignment problem. Individuals persist or cease on the basis of all their strengths and weaknesses taken together. It is not necessary for them to be capable of making attributions about the particular causes of success or failure.^[8]
- Evolution neither requires nor implies an external teleology. Although individuals may have goals and drives, themselves evolved as responses to the environment, the system as a whole does not have a particular principle for selective retention. (This follows from the local nature of the evolutionary process. There is nothing for such a principle to inhere in.) In not requiring a teleology, evolution avoids an infinite regress of explanation and is self-consistent. However little we know about the actual structure of the first organisms, there is no point at which we are theoretically obliged to postulate a conceptual discontinuity in the evolutionary process.
- Evolution is an open ended process in two senses. Because the role of genes is to encode chemical processes, rather than directly "construct" a phenotype by a fixed mechanism, a completely new sequence of genes can result in a completely new enzyme with previously unknown effects. Similarly, at a higher level of description, the phenotypic abilities that individuals have are not fixed: a stick can be used as a digging tool or a club and neither use is strictly determined by the structure of the arm or hand.
- Evolution is progressive. This is not meant in a teleological sense, but it follows from open endedness and relative success. Initially, organisms would simply have come into existence and ceased again. If one was able to make copies of itself ([Dawkins 1989](#)), then that would be the only one to persist beyond the life span of a single individual. Once persistence was ensured, with all persisting organisms being copies of the original, then properties such as diversity might have developed additional persistence value. Once diversity existed, then genetic mixing could have persistence value and so on.^[9]
- Evolution does not exclude Lamarckian inheritance. It appears the Lamarckian inheritance - inter-generational genotypic transmission of phenotypic characteristics acquired during a single lifetime^[10] - has not evolved in biological systems except for some fiercely contested cases among the simplest organisms ([McDougall 1927](#)). It is also fairly easy to see why this might be. There is simply no mechanism by which additional arm muscles might be "re-encoded" back into the genes, because they are not directly encoded there in the first place.^[11] All that is encoded is the synthesis of enzymes and other chemicals that will, taken together, result in a particular set of muscles and other bodily features but *only through a dynamic process of construction*. The completed body does not retain any representation of this construction process, so it cannot be "reverse engineered". However, this is a contingent fact about genes. There is nothing conceptually incoherent about Lamarckian inheritance as part of an evolutionary framework. Indeed the existence of cultural transmission suggests that if the material of the "genotype" is human language, rather than genes, Lamarckian inheritance has indeed evolved and may now dominate genetic inheritance ([Moravec 1989](#)).

2.5

It is clear that these implications have considerable appeal for unorthodox economics and address serious problems with the orthodox view. The absence of system level teleology and global knowledge is appropriate for emergent phenomena like markets. Absolute success measures are inappropriate^[12] when agents are reflexive and operating in environments where their actions affect the choices available to others. (It is now becoming clear that these conditions obtain in almost all social processes of any interest.) With the addition of Lamarckian inheritance, evolution is neither incompatible with rational (directed) behaviour nor with agents pursuing individual or collective goals.^[13] Finally, from the institutional perspective, it is encouraging that we can see the current state of social processes as a conceptually smooth progression right back to the most basic forms of social organisation.^[14] Economics has often been accused of "naturalising" and rendering ahistorical both the attitudes of the affluent middle class and the economic arrangements of late

capitalism. By emphasising process, as institutional economics does for example, the evolutionary approach opposes this tendency. In addition to these direct implications, the evolutionary approach also has potentially appealing rhetorical implications. It is fundamentally pluralistic and dynamic without being teleological. It stresses the descriptive and provisional nature of our understanding and places the understanding of the social scientist in a non privileged position *within* the evolutionary process. Finally, while it does not exclude deliberate or rational action, it situates this action in a context of profound environmental complexity and uncertainty. This context does not arise mysteriously, but naturally from the autonomy of actors and of the physical processes in the environment.

2.6

Furthermore, in challenging some of the assumptions of orthodox economics in a constructive way, the evolutionary approach moves it somewhat nearer to the perspectives of the other social sciences. Although sociology, economics and psychology will inevitably continue to differ in their views on the relative importance of social and individual factors, conscious and unconscious processes, rationality and impulse, decision and influence, it is much easier to integrate these insights within a broadly cognitivist framework which views agents (and social scientists) as having limited, diverse and provisional models of the world. It is in part the extremely strong (and occasionally dogmatic) assumptions of individualism, extreme rationality and a positivist treatment of quantitative and strictly behavioural data that have (perhaps unfairly) become identified with economics and effectively separated it from the other social sciences.

2.7

In the next section, we move away from biology to a discussion of the way that evolutionary ideas from biology have been used in the development of evolutionary algorithms.



Evolutionary Algorithms

3.1

In this section we will focus on three representative kinds of evolutionary algorithms: Genetic Algorithms ([Holland 1975](#), [Holland 1992](#), [Goldberg 1989](#)), Genetic Programming ([Koza 1992a](#), [Koza 1994](#)) and Classifier Systems ([Holland 1975](#), [Holland 1992](#), [Holland et al. 1986](#), [Goldberg 1989](#)).^[15] In each case, the algorithm involves generating an initial population of individuals and then operating on that population iteratively so that over time the "quality" of the individuals persisting in the population tends to increase. The operations involved and the structure of the individuals in the population are what make the techniques different. Genetic Algorithms (GA) and Genetic Programming (GP) differ mainly in the structure of individuals, while Classifier Systems (CS) not only have a different individual structure from the other two techniques, but different operations too. In the same way that we have abstracted from the intricacies of biological evolution, we shall try to abstract from the implementation details of evolutionary algorithms. Not only are there many different kinds of evolutionary algorithms, but there are many different variants of each kind and many parameter values that may be controlled in each variant. We shall raise these issues only when necessary. For a more detailed overview, we refer the reader to [Goldberg \(1989\)](#), [Koza \(1992a\)](#) and the references in [Nissen \(1993\)](#).

Genetic Algorithms

3.2

In a GA, each individual is represented as a string of numbers which are often binary. These individuals are typically generated randomly.^[16] The three main operators applied to the population are *reproduction*, whereby more copies are made probabilistically of fitter individuals, *crossover*, whereby two individuals are chosen probabilistically on the basis of fitness, "cut" at the same point

and the tails of the strings "flipped", and *mutation*, whereby a single position in an individual is changed randomly to another syntactically acceptable value. (If the numbers in the string are binary for example, then 1 is flipped to 0 and vice versa.) Fitness is defined by a function, called the *fitness function*, which maps all syntactically acceptable strings into a corresponding fitness. Crossover corresponds to one process of genetic mixing in sexual reproduction, while mutation serves the same diversity maintenance function as it does in biological evolution. While the technicalities are rather complex, it is intuitively clear why the population fitness gradually improves. Reproduction results in a higher proportion of fitter individuals in the population. It is these that probabilistically form the subjects for subsequent crossover and mutation. Unfit individuals which result will not persist, but any that are better than the current individuals will in turn come to dominate the population. (This is where relative performance is important.) Even if the initial population is randomly generated, there will always be one individual that performs the tiniest bit better than the others and this will be enough to start the evolutionary process off.

3.3

While the similarity between a GA and biological evolution is fairly clear, there are at least two important differences even at a conceptual level. Firstly, the existence of a fitness function provides both an "invisible hand" and a teleology for the system. The fitness function allows the fitnesses of all individuals in the population to be compared and both proportionate and rank based fitness functions rely on this "global" knowledge. Whatever properties of an individual the fitness function indicates as "desirable" are the properties that will become more prevalent in the population. It is this which provides the teleology for the system. Secondly, no distinction is made between the genotype and phenotype of the individual. The fitness function is able to evaluate an individual directly on a genotypic basis, rather than through some phenotype to which that genotype gives rise.

3.4

Both of these differences can be explained by the "standard interpretation" of the GA as an *instrumental* function optimiser.^[17] If the point of a GA is to serve as a tool allowing the programmer to find the optimal value of some fixed function, then global knowledge and a designed fitness function are both appropriate and necessary to that task. Rather than modelling natural selection, the programmer is attempting husbandry, incorporating an exogenous fitness function into the program in the same way that animal breeders measure their success by the amount of extra meat on a new strain of cow relative to the current herd.^[18] However, for the purposes of this paper, we are interested in the extent to which evolutionary algorithms are capable of serving as *descriptive* models of social processes. (In such cases, the assumptions of global knowledge and a shared teleology are probably not appropriate.) In the same way that we have described evolutionary models in general as those which specify a coherent analogy with biological evolution, so we can envisage simulations based on coherent analogies with evolutionary *algorithms*. It is to these analogies we turn in [section four](#), but for now, we continue with the technical description of evolutionary algorithms.

Genetic Programming

3.5

GP also typically makes use of the operators of reproduction, crossover and mutation. However, the individuals in the population are *programs* in some (simplified) programming language. A commonly used language consists of the set of possible trees for a specified set of operators (internal nodes) and terminals (leaf nodes).^[19] For example, the set of operators could be $\{+ | - | * | /\}$ and the set of terminals $\{0 | 1\}$. This language would produce a set of trees involving simple mathematical expressions of arbitrary depth.

3.6

Again, it is clear that this technique has a simple instrumental interpretation and GP has indeed been used for pattern induction and function estimation ([Koza 1992b](#)). However, there has also been a significant increase in representational richness in at least two senses that prove to be relevant for descriptive models of social processes. Firstly, the tree representation allows for individuals of arbitrary depth.^[20] Secondly, the interpretation of a GP tree is decomposable in a way that a GA string is not. If the operators are chosen to have "obvious" meanings, so that + refers to the addition operation, for example, then GP trees can be parsed by a human almost as easily as by a computer. Decomposability is demonstrated by the fact that an individual can be presented with a syntactically correct GP tree fragment like (+ (+ 3 5) ...) and parse it to (+ 8 ...). By contrast a GA string fragment (... 0 1 0 ...) makes no sense unless one knows its exact position in the string and even then it is necessary to refer to the coding of the fitness function.^[21] This decomposability is not purely a matter of comprehensibility, it also corresponds to a specific problem with the GA as an instrumental optimiser. Although there is no technical reason why a GA string should not encode an arbitrarily complex structure, the more complex it gets in practice, the greater the danger that optimisation will be dramatically impaired by the unequal importance of different string positions in producing fit solutions. Instead of truly parallel search, the GA will tend to optimise the most important positions first and then progressively optimise the less important ones. (This is the problem of *epistasis* [Goldberg 1989](#): 46-48).^[22]

3.7

Despite this increase in representational richness, GP techniques are still used predominantly in an instrumental way ([Koza 1992a](#)). They are set up to assume global knowledge, no phenotype/genotype distinction and the exogenous teleology of the fitness function. However, the richer representation possible in Genetic Programming has drawn attention to the fact that even for instrumental applications:

1. It might be appropriate to evaluate the fitness of individuals on the basis of repeated program outcomes and that these outcomes need not be separable;
2. Terminals and operators could just as easily be interpreted as *actions* in a simulated environment.

3.8

An obvious example of the former is evolving a program to generate random numbers ([Koza 1991](#)). In this case, the fitness must measure the statistical properties of *sequences* of numbers and can do nothing with single instances. As an example of the latter, consider evolving a program to serve as a controller allowing a robot to perform a simple task like finding litter and picking it up ([Koza 1992c](#)). In this case, the "output" of the GP can correspond to an *action* in a simulated environment like "Turn Left" or "Pick Up" rather than simply generating a number. (For instrumental applications, the programmer can *interpret* a number as indicating an action, but for descriptive simulations avoiding an external teleology, it is required that the interpretation be encoded in the program itself.^[23]) In addition, we are now making a concrete distinction between genotype and phenotype. The program is the genotype, but the execution process and actual outcome of executing the program in the simulated (or real)^[24] world involves the phenotype. Trivially, if the indicated action is "Move Forward" but the robot is facing a wall, the phenotypic outcome will be a dented robot! Although both of these applications are still instrumental, it is much shorter step, at least in the latter case, from asking how we might get a robot to optimise its performance of certain tasks given its abilities, to asking what behaviour we might expect to observe in an agent with certain abilities, when it is exposed not to teleological selection, but to selection by the physical environment, by other agents, or a combination of the two.^[25] We consider various answers to this question in [section five](#).^[26]

3.9

The final class of evolutionary algorithms we will discuss in this section are Classifier Systems (CS).^[27] In a CS, the individuals consist of IF ... THEN ... rules (condition-action pairs).^[28] The architecture of the CS is rather more complicated than those of the GA and GP. It begins when suitably encoded input from the environment is added to the "message list" of the CS. This input may satisfy the conditional part of one or more rules in the "rule base". Each rule has an associated strength, which reflects its past success at contributing to the generation of "appropriate" output. Eligible rules compete in an "activation auction", a probabilistic process in which fitter rules are better able to make high "bids" for activation on the basis of past performance. Once it has been established which rules are to be activated, they are copied to the message list and the combination of inputs and rules is "executed". (The combination of inputs and rules may generate outputs to the message list which in turn make new rules eligible.) As well as outputs to the message list, some rules will presumably generate outputs to the environment and these are also carried out. After this, the message list is purged. The outputs to the environment generate environmental feedback, which is used to alter the weights on rules.^[29] Finally, a GA is used from time to time to weed out lightweight rules in the rule base and generate new ones probabilistically from those with greatest weights.

3.10

Despite its apparent complexity, the CS actually functions in a rather similar way to the GA and GP. The population of the rule base should gradually improve in quality as time passes. The activation auction corresponds to the probabilistic application of genetic operators like crossover and reproduction to determine which individuals are actually "used" for tasks. The weights on rules can be seen as an alternative to duplicating fitter rules through reproduction in the GA or GP population. (One rule with a weight of 1 and four with weights of 0 are equivalent to a population of any size consisting only of multiple copies of the first rule. The analogy is not perfect. In the latter case, the rule diversity has really been lost, while in the former case the rules are still "available" if the weights should change. We will return to this point.) The "environment" can either be some instrumental task like summarising a set of cases, or a descriptive simulation of a real environment. As with the GP, the existence of rules which generate outputs to the message list as well as the environment allows for output behaviours of arbitrary complexity. The most important difference between the CS and GA/GP is that the CS explicitly produces strategies that consist of *sets* of rules which deal with a particular environment collectively. By contrast, even though GA and GP produce diverse populations, the idea is that only the fittest individual in that population will be "used" at any one time.^[30] This difference supports descriptive uses of CS and raises questions about the flexibility of the responses of simple GA and GP to systematic environmental variation. Although a single program in a GP *can* cover all the cases in an arbitrary rule base, the resulting program will be very large, unwieldy and perhaps difficult to interpret.^[31] In addition, the majority of the program will be redundant most of the time. We will return to this issue in [section five](#).

3.11

To sum up, it can be seen that even in their most simplified forms, evolutionary algorithms are a creditable abstraction of the biological evolutionary process. However, some caution has to be exercised in using them as models of social processes. In particular, their origin as instrumental optimisation techniques may be incompatible with desirable features of biological evolution like the absence of global knowledge and external teleology. In sections four and [five](#) of this paper we will focus on a number of specific simulations to see what they reveal about the potential and limitations of evolutionary algorithms as representations of *social* processes.

The "Standard Interpretation" Of Evolutionary Simulations

4.1

Although the simulations differ in detail, it is possible to present a "standard interpretation" of genuine evolutionary simulations which forms the basis for discussion in this section. ([Section five](#) attempts to generalise and develop this interpretation.) The standard interpretation draws on several different strands of thought from unorthodox economics. Two of these will be addressed now and others will emerge as the discussion proceeds. As has already been remarked, many of the insights which follow are taken for granted in sociology and psychology.

4.2

Probably the most important ideas underpinning these simulations are that rationality is genuinely bounded and procedural ([Simon 1972](#), [Simon 1976](#), [Simon 1978](#), [Simon 1981](#)). Unfortunately, until relatively recently, the idea of bounded rationality seems to have been reincorporated into orthodox theory as a generalisation of standard economic rationality rather than an ontologically different theory of human decision.^[32] For example, the observation that decision making is time consuming and cognitively difficult formed the basis for models in which agents were supposed to optimise their allocation of time to the decision process ([Winston 1987](#)). The same logic was used to transform the observation that individuals and organisations tend to use "rules of thumb" into an optimising theory ([Baumol and Quandt 1964](#)). Simon himself is responsible for two of the most damaging criticisms of this approach.^[33] In the first place, it involves an infinite regress of explanation. Does the agent include time spend deliberating about how long to deliberate as part of the calculations for the deliberation? Is the rule of thumb chosen using a rule of thumb? If so, how is *that* rule of thumb chosen? In each case, what appears *prima facie* to be an explanation actually involves postponing any real grounding for the theory. Lest it be thought that this is no more than the general problem of reductionist science, Simon further points out that the regress is not leading to explanations of equal or increasing simplicity and power, which would be helpful,^[34] but to explanations of increasing difficulty, which is not. It is intuitive that the rule of thumb for *choosing* rules of thumb must be at least minimally more complex than the rules which it must choose between since it must be capable of not only "executing" those rules but evaluating them. A related point is that information requirements increase for each level of regress. If the rules of thumb contain any parameters which must be tuned for effective use, the rule of thumb for choosing must have *its* parameters tuned to take account of different values for the rules of thumb it is supposed to decide between.^[35]

4.3

The solution is obvious, but extremely different from the orthodox view. Instead of rationality consisting of a "disembodied observer" who is able to look down on mental states from whatever height is required to get the complete picture and act "rationally", to some extent, rationality *must* be "embedded" in the unfolding of those mental states to avoid an infinite regress. The metaphor of the observer is a useful one. The fact that a human observer is embedded in the world certainly does not preclude them having some choice over which direction to look in, though parts of any view will be obscured by trees, fog or whatever. The observer can also change their point of view by walking or going up in a balloon, for example, and this may clear some obscurities, though it will almost certainly introduce others. However, at the very least, any observer is typically prevented from observing the state of their own forehead! If a mirror is used, this observation is possible, but at the expense of any view of what is going on behind the mirror. This example is a purely heuristic one. It shows that any embedded system has at least one "blind spot" in addition to its obscurities, but says nothing about how big it is. The nature and scale of obscurities and blind spots in the mental process is an open question, but we can be almost certain that rational choice theory, which postulates no genuine blind spots at all, is incorrect. Furthermore, the existence of even the smallest blind spot

changes the whole nature of the decision process, since any decision can no longer be absolutely known to be rational. This difficulty is more serious than that of mere probabilistic risk because there can also be no knowledge of the degree of uncertainty induced by the blind spot, since its "size" and "scope" cannot be determined either. [\[36\]](#)

4.4

The implication of this embeddedness is that at some level, as Simon argued, rationality can only consist of "unselfconsciously" following procedures. Two procedures can be compared, but they are compared by something that is itself a procedure and cannot be rationally justified beyond a certain point. [\[37\]](#) ([Wittgenstein 1953](#) makes the same point in his discussion of "language games".) One cannot step outside the whole mental process and "straighten it out". [\[38\]](#) At best, from any particular mental point of view, one can make comparisons and attempt to remove inconsistencies. (Always allowing that changing things to remove one local inconsistency may introduce another somewhere else. [\[39\]](#))

4.5

This interpretation of bounded, embedded or procedural rationality strongly suggests a computational approach. Although a computer program can modify the state of its variables, and in some cases even its own code, at any point it must take the variables and state of the program as given when executing a particular instruction. Thus, in addition to the general advantages of computer simulation ([Doran and Gilbert 1994](#), [Chattoe 1996](#)), there is also a conceptual similarity between programs and mental processes.

4.6

A second related theme, which is suggested by the absence of rational justification for procedures, is that agent models of the world are something quite distinct from the world itself ([Chattoe 1994](#), [Moss and Edmonds 1994](#)). Agents do not, as orthodox economics sometimes suggests, all have knowledge of the correct model of the world or even of the same model with divergent opinions about its parameters. [\[40\]](#) It is only in the context of an evolutionary theory, which assumes diversity, that the enormity of this assumption becomes clear. Everyday experience overwhelmingly suggests that agents differ not only in their preferences, but in their calculative techniques, their views on achieving goals and so on. Situations where beliefs or goals based on anything other than direct perception are widely shared are not the norm, but a fascinating field for social research.

4.7

The analogy between this view of agency and the description of biological evolution in the [second section](#) is fairly clear. Actors, which may be individuals, families, social groups, firms or governments, can be viewed as having both a genotype, which consists of their more or less *persistent* mental processes and operating procedures, and a phenotype, which consists of all the other "physical" aspects of their functioning, including the impact which they can have on the environment. The genotype is distinct from the phenotype, but "constructs" it, since we assume that actors are intentional. [\[41\]](#) Phenotypes, and through them genotypes, experience selective retention by the environment. In addition, the processes of genetic mixing and mutation also have their analogues. Genetic mixing occurs during the process of learning by imitation and mutation can result either from incorrect copying of a genotype or random (undirected) experimentation. (These processes will be analysed in more detail in [section five](#).)

4.8

The implications of biological evolution sketched out in the [second section](#) are also congruent with this analogy. In the great majority of cases, "fitness" in social situations must be measured in relative and not absolute terms. The complexity of the environment (consisting of the autonomous and interacting behaviour of diverse individuals as well as purely physical processes) means that credit

assignment for particular modifications to the genotype is extremely hard to implement. This affects the type of learning and adaptation mechanisms we can expect to observe. The *ceteris paribus* requirement (to which orthodox economics often appeals) can seldom even be maintained in the genotype of a single individual. In particular, the refinement of cognitive models involves the acquisition and organisation of new knowledge. This means that cognitive agents cannot easily "step into the same river twice" as far as their perception of the environment is concerned. This in turn means that agents will perform new actions, based on new inferences, and influence the environment in new ways, requiring other agents to modify their practices accordingly. This gives a "direction" to the social process and removes the costless reversibility which is typically assumed in orthodox economic models.^[42]

4.9

This leads on to the fact that social processes are also both open ended and progressive. Simple arrangements for buying and selling, fighting or building settlements gradually diversify and increase in complexity as knowledge is created and organised by agents. Over time, genuinely new products, institutions and technologies are incorporated into the social process.^[43] Open endedness reminds us that economic systems, like biological ones, do not imply or require an external teleology, even though it is almost a matter of definition that autonomous agents should have *internal* goals. This point needs some clarification, since in some contexts, particularly the interaction of firms in a market, the absence of an external teleology is somewhat obscured by "market rhetoric".^[44] It is certainly true that there are many laws and customs, more or less well observed, designed to regulate the function of markets and other social institutions, but not that they reflect any coherent or agreed plan to support "efficiency" or indeed any unitary goal, nor that they can be expected to ensure any particular social behaviour such as profit maximisation in competitive firms.^[45] Indeed, the structure of institutions like the market often reflects some of the interests of the very "players" who are supposed to be "regulated" by them.^[46] Although firms will adapt to the structure of the market place, this structure is not coherent, binding or permanent enough to constitute a teleology, any more than it would be valuable to talk about other individuals constituting an external teleology.^[47]

4.10

Finally, although we have so far focused on a pure evolutionary process, we are not obliged to exclude various mechanisms of Lamarckian inheritance in social systems, because it is much easier to see how phenotypic effects can be re-encoded into the genotype when the material of the genotype is everyday language rather than DNA. This process of re-encoding is precisely what we mean by directed or "rational" adaptation, as opposed to evolutionary adoption. It is important to stress both that evolutionary models are *not* intended to exclude rational or deliberate behaviour and that the evolutionary view of bounded rationality is actually straightforwardly compatible with many forms of learning. Many critics of evolutionary models ([Penrose 1952](#)) accuse them of asserting that human behaviour is either random or genetically determined. This is simply incorrect. These same critics also fail to recognise that the technical reasons why Lamarckian inheritance cannot occur in biological systems are not necessarily binding on social systems. Finally, while it is true that evolutionary models can easily incorporate learning as Lamarckian inheritance, attempts to incorporate evolutionary behaviour into "rational" models have so far been much less successful. Two of these attempts, the GA models of Arifovic and the use of replicator dynamics in game theory, will be discussed subsequently. On the basis of the foregoing discussion, two of the most important general questions raised by evolutionary models are:

1. How much rationality is actually required for the level of organisation we observe in social processes?
2. How can social evolutionary processes *interact* with rational decision?

4.11

On the basis of the evolutionary models discussed in this paper, the answer to the first question is considerably less than orthodox economics has previously suggested. The answer to the second question is much less clear cut and one of the reasons for attempting a coherent interpretation of evolutionary models is precisely to make the boundaries between evolutionary adoption and rational adaptation clearer. (If economics is guilty of over-playing the role of deliberate action, then perhaps sociology and social psychology are guilty of under-playing it.) Before moving on to a detailed discussion of models in the [next section](#), the last purpose of this section is to examine two interdependent aspects of the analogy between biological and economic systems that are traditionally regarded as problematic.

4.12

These aspects are the relationship between the measurement of individual "lifetimes" and the objects of evolutionary selection. For biological evolution, in the absence of Lamarckian inheritance, there is a unique relationship between a genotype and a corresponding phenotype. Once we allow Lamarckian inheritance, this is no longer true and a particular individual can no longer be identified with a fixed genotype. However, this does not, as Penrose suggests in the context of industrial evolution, raise any difficulty with the "identity" of firms, nor are we obliged to treat each change in the genotype as the creation of a "new" firm.^[48] Any such argument fails to understand the definition of Lamarckian inheritance. It is precisely a change in the genotype *during the lifetime of the individual*. One of the things that is often understated in evolutionary modelling is the empirical integrity of the objects of selection which typically form the basis for models. It is trivial that individuals display the properties of organisms: autonomy, internal structure, a well defined boundary and mechanisms of self-maintenance. However, these properties are also very much in evidence in organisations and social groups of all kinds, and particularly in firms with actual or potential competitors. The boundary is maintained by the distinction between employees and others, the use of "non disclosure" agreements for departing employees, the importance of tacit knowledge and "on the job training" and by more intangible phenomena like "company pride". The combination of a correct interpretation of Lamarckian inheritance and the observed integrity of individuals and groups resolves the apparent difficulty. The objects of selection really are organisations and individuals and their lifetimes really are the periods for which they persist. This resolution does however raise a subsidiary issue, which is that very few social processes appear to take place over the real life span of groups or individuals.^[49] If this were actually so, then the role of evolutionary mechanisms would of necessity be considerably weakened. The answer is that although we have stressed Lamarckian inheritance as a way of fitting *directed* learning into an evolutionary framework, there is no conceptual reason why it shouldn't equally apply to the effects of undirected learning within the individual lifetime. All that remains is to distinguish more clearly what we mean by directed and undirected learning and that is a task postponed until later in the next section.



Focusing And Developing The Standard Interpretation

5.1

Ironically, but perhaps not surprisingly, two of the best known kinds of evolutionary models are the least satisfactory from the point of view of the foregoing discussion. This is because they attempt to reincorporate evolutionary ideas into an orthodox economic framework without due regard for their logical coherence.

How Not To Use Evolutionary Algorithms

5.2

The first set of models is GA based and has been developed by Arifovic and others ([Arifovic 1990](#), [Arifovic 1994](#), [Arifovic 1995](#), [Arifovic 1996](#), [Arifovic and Eaton 1995](#), [Bullard and Duffy 1994](#)). In these models, pricing or other decisions are coded as GA strings which then undergo the usual operators of reproduction, crossover and mutation, along with an additional *election* operator which will be explained shortly. [Arifovic \(1990\)](#) draws attention to two different interpretations of the GA, which will be useful in what follows. The first could be called the *mental interpretation*, in which GA strings represent different possibilities in the mind of a single agent or organisation. Under this interpretation, individuals carry out the operators on their own mental contents, using the fitness function to pick the string which actually generates an action in any particular period. The second is the *population interpretation*, which is the one that Arifovic favours and uses subsequently. Under this interpretation, each individual consists of a single string and the GA consists of a *population* of individuals. The first difficulty with these models is the criteria by which they should be judged. The behavioural interpretation provided is cursory and it appears that the objective of these models is to simply to obtain or "explain" convergence to the equilibria indicated by analytical models. In particular, [Arifovic \(1990\)](#) is quite specific about introducing the election operator to deal with what she describes as the "problem" of non-convergence.

5.3

For now, we will assume that Arifovic intends her model to be interpreted descriptively. The interpretation of crossover as learning by imitation and mutation as copying error or random experimentation are fairly straightforward. However, the absence of a genotype/phenotype distinction and the choice of crossover as the mixing operator have important implications. It appears that one agent can look directly into the head of another and perceive what it intends to do.^[50] Furthermore, crossover imposes a lot of "structure" on the imitation process, ensuring that it always produces new strings that are syntactically correct. Thus agents can not only read each others minds, but do so with such clarity that they always imitate fully formed "concepts" despite the highly structured encoding of the GA string. (This structure may also cause the problems with mutation that lead to the introduction of the election operator. Mutation at some positions may have far more effect on the fitness of the string than at others. This is undesirable purely from the point of view of design.) The reproduction operator produces larger numbers of fitter strategies in the population. This operator can be interpreted in two ways, but both are problematic. The first is that the environment imposes the removal of less fit firms and, by implication, the increase in the population of fitter firms. However, while it is easy enough to see why unfit firms might be removed by bankruptcy, it is much harder to say what properties replacement firms would have, and in particular how they would be able to imitate the practices of only fitter firms. (The argument usually given is that new firms adopt market "best practice" but this assumption needs more support. In particular, if this option is available to new firms, why is it not open to current firms performing poorly?) Nothing is said about how the time scales implied by this interpretation can be reconciled with the speed of environmental change in real markets. The other possible interpretation is that reproduction represents imitation of the whole strategy of another firm on the basis of fitness. The two difficulties with this interpretation are firstly, that observing utility is even more implausible than observing intentions^[51] and even if we grant this as a possibility, secondly, there seems to be an strong implication of global knowledge to ensure that people imitate in proportion to the overall distribution of fitness. Another possibility is that firms could "learn" the fitness distribution in the market, but it is not clear that they could do this fast enough in a dynamic situation. (Both interpretations of the reproduction operator also suffer from the absence of a clear genotype/phenotype distinction, since it is not obvious whether we are talking about the success and failure of real firms or of the strategies adopted by those firms.)

5.4

The interpretation of the election operator is similarly ambiguous. This operator is applied to the process of crossover. Both the crossover products are assigned a *potential fitness*, an expectation of how they are likely to perform on the basis of their performance with past data. If the potential fitness of either crossover products exceeds the actual fitness of the "parents", then it will replace

that parent in the population. Full crossover is far harder to justify than a "half crossover" (imitation) and without that, there will not be a pair of crossover products or parents. We *can* apply the same algorithm to the decision as to whether or not to incorporate an imitated section of string from another agent, but it is still not clear what it is about the organisation of the environment that ensures that agents are presented with imitation opportunities in proportion to their fitness. (Again, a real phenotype might provide a solution, with firms tending to imitate those displaying similar patterns of *outward* behaviour or observable features like size and age. This is the approach taken in [Chattoe and Gilbert 1997](#).) In fact, it seems that if agents are so well informed about fitness and the intentions of other agents, they should be able to decide for themselves whether or not to imitate, regardless of the state of the other firms they are presented with. Thus election is in danger of rendering reproduction redundant and turning the GA into a hill-climbing algorithm with all the attendant drawbacks such as non global convergence.

5.5

The other class of well known evolutionary models are those in game theory using the so called replicator dynamics.^[52] This is a rapidly expanding field ([Weibull 1995](#), [Vega-Redondo 1996](#)) and there is not space to provide a detailed survey here. The main point that is relevant to the current discussion is that the behavioural rationale provided for these models bears even less cursory inspection than that in the Arifovic models. (Further evidence for this is provided by the fact that, as far as I know, there have never been attempts to fit game-theoretic models of this kind to any data or even to suggest social processes from which these data might be collected.) Replicator dynamics originated in theoretical biology, where they were intended to model shifts in the proportions of genes or species of differing fitness. In extremely simple terms, a gene or species will increase (decrease) its share in the population if its own fitness is greater (less) than the average fitness and the change will be proportional to the difference between the average and the particular gene/species fitness. All models of this kind face the same difficulty. Replicator dynamics is based on the assumption of an underlying biological process which justifies its mathematical form.^[53] For example, it is assumed that species display exponential population growth when they are not resource constrained. In applying such models to social processes, the results will only be useful if the same assumptions apply. Without exception, replicator dynamics models justify themselves with reference to discursive behavioural arguments about imitation and observation errors and not by reference to the other underlying assumptions of the original model. This produces a number of problems:

1. How do individuals get information about average fitness, or even observe the fitnesses of other individuals? Usually these models simply substitute an unexplained rational theory of imitation for a rational theory of individual game play.
2. How do individuals correlate game strategies with decision processes? In far too many models, this problem is assumed away. Either the individual just *is* their strategy, or else individuals are both telepathic and know the full set of possible decision processes. The latter assumption is required to ensure only legal mutations.
3. Biological models assume competition for scarce resources like food and shelter which directly impact on reproductive success through physical mechanisms. The possibility of making correlations between certain independently observable and measurable skills (like foraging rates) and reproductive success makes the concept of selection by reproductive success both non tautologous and potentially falsifiable, though this is very hard work outside the laboratory. In games, the resource "emerges" only in game play and its only value to the player is that it enhances the likelihood of strategy propagation. While a strategy like "co-operate" is observable and can be correlated with reproductive success (though with the reservation that this success must actually be correlated with all possible strategy combinations), this correlation is meaningless because the only "measure" of the strategies performance is precisely the same "utility" which determines that success.^[54]

5.6

What lessons can be learnt from these unsatisfactory models? The first is that the interpretation of evolutionary models *does* matter on at least two levels. On the methodological level, very few unorthodox economists (or non economists come to that) subscribe to the "as if" argument for modelling, which states that it is irrelevant whether a model is behaviourally plausible as long as it predicts adequately. In any event, it has no force in areas like evolutionary game theory where predictive models are not available. It is not enough for those who believe in procedural rationality to know *what* happens, they also want to know *how* it happens. In that sense, an interpretation is just a more complete specification of the *process* by which evolution could take place in a social rather than biological context. On the practical level, although the Arifovic models *are* used to make limited sense of experimental data, the behavioural interpretations used are in the main either inconsistent or implausible. This could be avoided if the behavioural interpretation of the model was more than a rhetorical defence of the selected formalism. This brings us on to the second lesson, which is that a working knowledge of the history and implications of evolutionary biology and evolutionary algorithms appears to be important in their sensible application to social science.^[55]

Simulations Based On The Mental Interpretation

5.7

In this section I will discuss a series of models which appear to achieve a more plausible behavioural interpretation in the use of evolutionary algorithms. Simulations using the mental interpretation and based on Genetic Programming have been produced by [Chattoe and Gilbert \(1996\)](#), [Chattoe and Gilbert \(1997\)](#), [Dosi et al. \(1994\)](#), [Edmonds \(1997\)](#) and [Edmonds and Moss \(1996\)](#) among others.^[56] These models have been applied to the behaviour of households learning to budget, firms setting prices in a market and modelling individuals who face an environment where having the same world view as everyone else is undesirable. Despite these differences, we can identify two sets of common questions which must be addressed by each of these models.

The Interpretation of the Agent Decision Process

5.8

All the questions in this group depend on how much detail can we provide about the "mental process" corresponding to an internal evolutionary algorithm. It is clear that for any problem, an individual or organisation will have a number of candidate solutions which seem feasible. The reproduction operator corresponds to a process by which some candidate solutions become "more favoured" in the population. The crossover operator corresponds to a process in which one candidate solution is modified by the incorporation of a sub tree of material from another. The mutation operator corresponds to a process by which some part of a candidate solution is replaced by a *novel* sub tree or "innovation". All three operators require more detailed interpretation.

The Interpretation of Reproduction

5.9

How is the process of reproduction represented in the state of the population? We have already discussed in passing that the weights assigned to Classifier System rules can be seen as corresponding to the duplication of fitter strategies in a traditional GA. However, the behavioural meaning of strategy duplication is much less clear than that for some sort of weighting or "confidence measure" attached to particular strategies. In particular, there are empirical features of human cognitive abilities which do not correspond well to the duplication interpretation. The most important of these is the existence of long term memory. Although it is true that humans can hold rather few strategies "in mind" at any one time, they are extremely good at remembering strategies

which they have already implemented or observed elsewhere. Thus the traditional GA or GP architecture, which implies parallel processing of operators applied to many candidate solutions is probably not appropriate for models of individual decision making.^[57] A behaviourally plausible alternative exists, in the shape of the GENITOR algorithm (Whitley 1989). In this algorithm, candidate strategies are permanently ranked by fitness. In each generation, instead of parallel processing, strategies are chosen proportionate to their ranking and a single operator is applied.^[58] Offspring of crossover and mutation are returned to the population at the "appropriate" positions indicated by their fitness. They either replace the strategies with the nearest fitness or are inserted into the ranking so that the least fit strategy "falls off the bottom". This approach has three advantages. Firstly, it explicitly avoids duplicates if strategies replace those with the nearest fitness or can easily be made to do so under the insertion strategy. The "importance" of strategies is represented by their position in the ranking rather than by the existence of duplicates and insertion does not require the re-evaluation of fitnesses for all other strategies.^[59] Secondly, this ranking also remains relatively stable because operators are applied one at a time. (The flipping from strategy to strategy displayed by non converged traditional GA or GP is not self-evidently plausible as a representation of human behaviour.) Finally, the algorithm relies on the "active consideration" of no more than two strategies at any one time, all other strategies are simply stored in memory. However, there is a drawback to this approach, for which detailed discussion will be postponed. In an instrumental GA, the fitness function is known and fixed so new strategies can be "placed" in the ranking immediately. It is not clear what should be done to them if fitness evaluation requires interaction with the environment. We will return to this point.

The Interpretation of Crossover

5.10

Is crossover to be proportional to some measure of fitness? If so, what measure? In the GENITOR algorithm we have already discussed, crossover is more likely to occur among solutions with higher ranking. In an organisation, this could be interpreted as reflecting the number of people preferring and arguing for a particular strategy. In an individual, it could reflect the fact that better strategies are more likely to "spring to mind" or remain in the forefront of memory. One thing to bear in mind is that once we become interested in descriptive models, it is more appropriate to consider what is realistic than what produces rapid convergence and this applies just as much to the parameters of the simulation as it does to the selection of the evolutionary algorithm. (We may even become interested in distinctive *failures* of convergence if these can be identified empirically.) Simulations are a very useful technique for comparing different possibilities in this context. This point is just as relevant to the assumption of strictly proportional fitness as it is to the size of the population of candidate solutions discussed in the last section. Just how much do we lose if the population is 5 rather than 50?^[60] How much slower would convergence be if solutions were only ranked as good, average and poor or if solutions were ranked with probabilistic accuracy?

5.11

If the exogenous fitness function used in the Arifovic and replicator dynamics models is regarded as imposing too much structure on the evolutionary process, then at least two alternatives exist. The first appears not to have received much attention. This would involve each individual using an internal fitness function which is nevertheless regarded as being "infallible" by that individual. Even if individuals have identical internal fitness functions, they avoid the problem of external teleology, but this approach also allows for the fact that agents may simply have different goals.^[61] In addition, it is possible to investigate how the environment selects fitness functions at the same time as it selects strategies which are fitter relative to those fitness functions it selects.^[62] One of the reasons why the "internal fitness function" approach has not received much attention is that it goes against an important assumption of most bounded rationality models, which is simply that agents

cannot accurately rank the fitnesses of strategies *a priori* except in trivial environments. Instead, they can only evaluate strategies in use and make comparisons on the basis of what has already happened. This is the basis of the [Dosi *et al.* \(1994\)](#) simulation, in which firms use cumulated profits to choose GP strategies for price setting on a probabilistic basis, with a minimum probability for previously untried strategies arising from crossover or mutation on the current population. (Depending on how mutation is defined, it is capable of generating whole new strategies.) Strategies which fail on any one of three conditions will be replaced, either with new strategies or crossovers/mutations of current strategies. The first failure condition is that the probability for a strategy being picked should fall below a minimum value. The second is that a strategy should generate a negative profit. The third is that it should indicate a price outside a rather broad range of "sensible" prices which runs between 0 and some large positive value. This model suggests a number of points. Firstly, its functioning is rather similar to the GENITOR algorithm in that the methods for generating replacement strategies tend to preclude duplication and force out the worst strategies disproportionately. Secondly, there is negligible calculation involved in the decision process. All that is required is a running total of cumulated profits for each current strategy and maximum and minimum values for sensible prices. Finally, the model draws attention to the distinction between syntactic (genotypic) and semantic (phenotypic) evaluation of strategies. It is possible to reject a strategy because it generates what is believed to be a "silly" price without reference to what happens when it is tried out in the market. By contrast, trying it is the only way that a strategy can be found to be unprofitable. Although we have already cast doubt on the plausibility of internal fitness functions which are capable of ranking strategies, it is much more plausible that there should be internal *filtering* processes that cut down the size of the strategy space based on general domain knowledge ([McCain 1992](#)).

The Interpretation of Mutation

5.12

Does mutation correspond to "transcription error" or experimentation? This is a relatively minor point, but it affects the fitness attributed to new strategies. In the Dosi *et al.* model, new strategies resulting from crossover are assigned the average of the cumulated profits of both parents. (This is the nearest that firms come to doing any calculation.) If mutation is interpreted as transcription error, then the mutant should be attributed the same cumulated profit as its parent since by definition such errors would be corrected if they were noticed. This could have unfortunate effects if a mutant was much poorer than a parent which had already cumulated a large profit. (In fact, this is one potential drawback with the Dosi *et al.* model, that firms have an infinite memory for cumulated profits, thus leading to possible lock in. On the other hand, as we have already remarked, evolutionary algorithms which produce too much strategy vacillation may be empirically implausible.^[63]) By contrast, if mutation corresponds to experimentation then all mutants should be treated as untried strategies or perhaps receive the parental fitness penalised according to some distance metric between parent and offspring.

5.13

This section suggests that with care, coherent interpretations can be produced for all the traditional genetic operators. These can be fitted together within the framework of the GENITOR algorithm, with some modifications to allow for the fact that strategies cannot be evaluated for actual fitness without trial. Strategies remain ranked by fitness (cumulated profit for example) at all times. New strategies are assigned a minimum fitness (or probability of being chosen) reflecting their untried status. Crossover products receive the average fitness of their parents. Mutation products are regarded as untried.^[64] The parents are chosen probabilistically on the basis of their ranking as is the strategy which is actually to be implemented. Tried strategies that yield negative profits or produce "silly" prices may still remain on the list as a mechanism to prevent their being reintroduced, but they will have no possibility of being chosen as parents or of actually being

implemented. (Firms will reject crossovers, mutations and novel strategies if they have already been tried within institutional memory or are currently "under discussion".) This algorithm has a number of desirable features. Firstly, the probability distribution for strategy implementation need not be the same as the probability distribution to select parents for the genetic operators and neither needs to be equivalent to the institutional memory of the firm. For example, the implemented rule could just be the top ranked one at all times. This would make firm behaviour very stable. For equivalence to the Dosi *et al.* model, parent selection would take place from the set of all strategies which had actually been tried and achieved a cumulated profit greater than the minimum for untried strategies while implementation would take place from the set of all strategies including the untried ones. Secondly, the evolutionary process is "driven" by the firm's view of its own institutional memory, without this making it a rational choice process. While the GENITOR algorithm performs one genetic operator every period, a firm will only attempt to identify new strategies when it is unhappy (in some sense to be defined) with the set which it already has. If the firm always wants to entertain four possibilities, then initially it will have four untried strategies, but as it tries them, they will get ranked. Only strategies which perform so poorly that they are probabilistically less likely to be chosen than untried strategies (which are assumed to have a fixed minimum probability of being chosen reflecting their "curiosity" or "novelty") will encourage the firm to cast about for new ones. This also produces a measure of realistic stability into the behaviour of the firm. Thirdly, this algorithm not only incorporates the coherent interpretation of the operators, but takes into account earlier comments about the distinction between institutional memory and active consideration, the behavioural implausibility of duplicate strategies, the limits on what firms can know and calculate and the possibilities for syntactic reasoning about strategies even when semantic reasoning is not possible. It should be made clear that this algorithm is only being proposed for further investigation rather than advocated as the *correct* model of evolutionary decision. It does however seem to have a number of desirable features which may make it worthy of further consideration.

5.14

In the next section, we turn from discussion of the decision processes of the agent to their interaction with other agents and physical processes in the environment.

The Interpretation of the Environment

5.15

So far, individuals under the mental interpretation have just been concerned with evaluating strategies they have generated "internally". Depending on the nature of these strategies, it is also necessary to develop a coherent interpretation of the individual's relationship with the environment of other individuals and physical processes.

5.16

The first issue arising in this context is a general one, concerning the open endedness of biological and social evolution and the way this compares with the functioning of evolutionary algorithms. *Prima facie*, it may appear that the problem with evolutionary algorithms is that they have a fixed set of terminals and operators while biological evolution does not. In fact, this is not the case. Within certain limits, like the coining of new scientific terms, any human language also has a more or less fixed set of components, though it is obviously far larger than the typical GP language. It is the *combination* and *structuring* of these components that can accommodate both the richness of sonnets and the pragmatism of operating instructions. However, neither richness nor pragmatism can be attributed purely to the combinations of words themselves. Instead, it results from some comparison between the words and some other aspect of reality. In the case of operating instructions, the words are compared with perceptions of the thing to be operated and feedback on whether implementing the instruction produces the desired result. In the case of a sonnet, the words are "compared" with memories of summer days and sweet spring showers to see if they evoke

sensations appropriate to what the poet is talking about. In each case, it is the *interplay* of genotypic and phenotypic aspects that gives rise to richness. In an instrumental GP, the fitness function just summarises the environment. In a descriptive GP, open endedness arises not just from enlarging the GP vocabulary, but equally from enriching the environmental consequences of phenotypic behaviour to which it gives rise. For example, instead of the terminals "pick up box", "pick up cone" and "pick up sphere" in a simple blocks world, we might have the terminal "pick up nearest" which, when applied, has consequences which depend on what the nearest object is. Of course, the environmental interpretation of "pick up nearest" is not immediately accessible to the agent who carries out the operation. In fact, for any sort of "pick up" terminal, it consists of the "physics" of the environment. All the individual observes is that "pick up nearest" usually works very well when the nearest thing is a pen, but results in a hernia when the nearest thing is a safe! It is on this basis that the individual may start to organise their perceptions and make inferences. (For example learning not to attempt "pick up nearest" on "big things".) This view of the relationship between agent and environment is described as qualitative physics and is essential to the existence of genuine novelty in the environment.^[65] In fact, as far as we can tell, the environment is not completely open ended but follows "natural" or "social" laws. The fact that nitric acid mixed with glycerine produces nitroglycerine was "known" to the physical environment but not to Alfred Nobel, who blew himself up several times finding out! Thus, building open ended evolutionary algorithms is, perhaps paradoxically, a matter of increasing the sophistication of the environment and not of the agent.^[66]

5.17

This observation has an immediate application to the issue of what individuals are supposed to be able to infer or observe about each other. In the first place, it is important to distinguish *behavioural* imitation (where one firm simply imitates the price set by another in the last period for example) from *cognitive* imitation (where one firm is actually able to incorporate some aspect of another firm's decision process into its own). More research is needed in the extent to which firms are capable of making inferences about each other and actually do so.^[67] On one hand, a great deal of information about a firm is extremely hard to observe, even if it is not actively kept secret. Against this can be set the incentives for success and the considerable skill which humans have at making inferences. For example, one firm may infer that another is testing a new production process secretly by observing that it has started purchasing a certain sort of raw material in small quantities.^[68] It may make *this* inference by observing the delivery van of a specialist chemical company. In the light of these difficulties, transcription errors are a far more likely explanation of mutation than deliberate experimentation during the process of imitation. It is also clear that the syntactic "interpretability" of GP trees is important not just for the simulator, but also in ensuring that a firm "knows" when it has obtained a coherent piece of "intelligence", even before it can place a value on it. By contrast, GA models have to assume syntax preserving operators like crossover so that firms don't have to address the issue of interpretation. The qualitative physics perspective is relevant both in the process of making inferences as part of imitation - inference involves "filling in the gaps" using correlations that have been acquired from past experience - and in trying to assess who it would be appropriate to imitate. Again, more research could perhaps be done on what knowledge about firms can be regarded as available in the market, but we can be sure that even if the profits accruing to individual strategies are not available, firms will be attempting to build up proxies that define their competitors in terms of age, share dividends, recent expansion or recruitment, firm size and so on.^[69] All of these proxies can be defined in relative and not absolute terms.

5.18

The final difficulty in interpreting the environment comes in representing what individuals believe about the effects of their actions. Arifovic assumes by omission that firms don't have any beliefs about this and perhaps, since the only option they have is to set their own prices, this is appropriate. However, it may raise a problem with the election operator since this requires a firm to assume

market stability in two interdependent respects. Firstly, the profit accruing to a strategy in the previous period must be positively correlated with its profit in the present period. Secondly, the impact which one firm has on the environment must not be so great that the profit accruing to it, had it used the strategy it is now testing for election, would have been completely different to the one it actually observed for the strategy it *did* use. These problems are avoided in the Arifovic models by the extreme simplicity of the environment and the effective "invisible hand" of the fitness function, but in real markets such assumptions of stability might not be appropriate. The same problem occurs with more force in a GP model of the El Farol Bar ([Arthur 1994](#)) developed by [Edmonds \(1997\)](#). In this model, each agent is keen to visit the bar only if it is not too crowded. Agents make use of mental models which allow them to decide about whether to go to the bar each night, based on previous experience. The preferences of the agents require the evolution of mental model diversity, since if *everyone* believes the bar will not be too crowded, they will all go and make it too crowded. In [Edmonds \(1997\)](#), despite the fact that agents are evaluating sophisticated GP strategies which take account of the past actions and utterances of other individuals, they perform evaluations on the assumption that what *they* do has no impact on the actions of those other agents! The kind of comparisons which it is appropriate for individuals to make will depend on the particular structure of the model, but it is important to devise comparisons which are both consistent and plausible. Again, the qualitative physics approach suggests that individuals will be attempting to *identify* stable comparisons which can then be used as parts of their decision process. We may choose to simplify the model by allowing only "appropriate" comparisons, on the basis of "external" observations of the model properties, but we should not ignore the issue altogether.

Simulations Based On The Population Interpretation

5.19

Population based models, which attribute a single strategy to each individual are obviously more appropriate to modelling norms or habits rather than sets of strategies which have any rational component in their selection. Many of the issues connected with models based on the population interpretation, such as imitation and the significance of the genotype/phenotype distinction have already been discussed. One obvious area for further research is models which combine evolution at the individual and population levels.^[70] At their simplest, however, individuals may only consist of a single action such as "co-operate" or "defect".

5.20

It should not be thought that replicator dynamics is the only technique for representing population level models. Apart from the reservations already expressed, it is straightforward to build models which more closely approximate to a coherent social evolutionary process. For example, we can view the payoff to games in terms of "energy" rather than "utility".^[71] This is used up at a steady rate and any agent reaching a zero energy level is removed from the population. (This approach is used in the "sugarscape" described by [Epstein and Axtell 1996](#).) Depending on whether there is more to agents than particular game strategies, performance based imitation can take place on the basis of proxies like age and involve behavioural imitation or inference of the genotype with imitation and mutation. It is not yet known whether models of this kind would arrive at different equilibria than the replicator dynamics models and this remains an unexplored field of study.

Syntactic And Semantic Evaluation

5.21

In addition to developing coherent interpretations of the decision processes of individuals and their interactions with the environment, there is also a more general issue of interpretation which needs discussion. This is the relationship between directed and undirected learning and the role of syntactic and semantic analysis of strategies. Syntactic analysis involves being able to say

something about the value of a strategy purely on the basis of its syntax. An example was provided by a GP tree which indicates a negative price. This is obviously unsatisfactory and does not need to be tried out in the market. By contrast, semantic analysis involves being able to say something about the "meaning" of a GP tree, for example that the reason why a cost-plus pricing strategy is failing to show a profit is that the definition of the cost element fails to take account of depreciation in the form of intermittent costs for machine replacement. The obvious difficulties with semantic analysis have meant that descriptive GP models have ignored it so far. This has meant that in some sense, learning by alterations to GP strategies has been largely undirected. Although whole strategies are chosen according to past performance, the components which make them up are not selected on this basis, or at best only very indirectly. (Firms must implicitly assume that sub trees of better firms - where better is defined by their criteria on whether or not to imitate - are themselves on average likely to be better, but this is much weaker than being able to identify and imitate only fitter sub trees.) By contrast, directed learning involves being able to evaluate the choices one is presented with in some way and pick the better ones directly.

5.22

Although the development of models which include semantic analysis of GP strategies is a major project, one possible direction is indicated by the discussion of qualitative physics in the previous section.^[72] In the same way that what is lacking in open ended GP models is not more operators and terminals but a richer environment, what is lacking for the firm to perform semantic analysis is not some "special" GP language but adequate *internal* structure to the model of the firm.^[73] For example, if "costs" are categorised as actions which result in a reduction in the current account, then a comparison of this set with the set of terminals in the current GP strategy will reveal the absence of any terminal connected with machine replacement costs. Semantic analysis of strategies thus presupposes that the operators and terminals possess additional attributes which give them semantic content.

5.23

Semantic analysis also raises two other interdependent issues. The first is that figuring out what a GP tree means can also be a taxing problem for the simulator as well as for the individual that is using it. Some of this difficulty can be removed by sensible choice of operators and terminals and avoidance of very large trees.^[74] (Combinations of operators and terminals which are easy to interpret individually but hard to interpret as a tree may provide *prima facie* evidence against those combinations occurring in real decision processes given our evolved linguistic competence. Mixing logical and arithmetical operators in the same tree provides an obvious example.) However, some difficulty is bound to remain and this results from the fact that GP trees are single entities. They can be contrasted with CS rule bases in which each rule is simple to understand but the combined effect can be very complex, particularly if rules are able to alter the state of the message list so that new rules become eligible to fire. There is no simple resolution for this problem, which is no less than the long running knowledge representation debate from Artificial Intelligence ([Brachman and Levesque 1985](#), [Sloman 1985](#)). However, research into ADFs may also provide some suggestions as to how useful sub trees can be identified and propagated more effectively.^[75]

Conclusions

6.1

This paper has had to draw in a wide range of ideas and examples to present a coherent framework for social evolution. However, to conclude effectively, it is possible to decompose the resulting framework into three constituents: the process of biological evolution, the instantiation of the evolutionary process in an evolutionary algorithm and the implications of both of these constituents for evolutionary models in social science.

6.2

On the first count, it matters very much whether one has an effective understanding of evolutionary biology in producing models of social evolution, both in recognising the parallels between social and biological evolution and, equally importantly, acknowledging their differences. In this paper, I have focused particularly on three aspects of this understanding. Firstly, the importance of a proper understanding of Lamarckian inheritance, stressing that its absence in biology is only a contingent fact rather than a theoretical or logical necessity. Lamarckian inheritance not only fits smoothly into the *social* evolutionary framework but, correctly interpreted, resolves difficulties about the identity of units of selection and the likely speed of evolutionary change. Secondly, there is a need to avoid spurious teleological reasoning. Although it is true that social institutions, including the market, exert some selection pressure on individuals and organisations, it is not at all clear how strong that pressure is and whether it should be seen as congruent with any desirable social goals. Finally, it is important to be clear about the behavioural interpretation of processes of "genetic transfer" essential to evolutionary models, so that agents are assumed to perform actions that they seem likely to be capable of in the real world. There seems little point in moving from rational models of decision to equally implausible rational models of imitation.

6.3

On the second count, it is equally important to understand the strengths and weaknesses of evolutionary algorithms and to recognise their many variants. The majority of these have been developed with instrumental objectives in mind so caution must be used in applying them to social processes. This is particularly true of the fitness function which imposes an external teleology on the instrumental GA in a way that may not be appropriate for descriptive models. However, on a more positive note, the quest for rapid optimisation techniques has led to many creative uses of the biological analogy in generating variant algorithms which have not yet been fully exploited in models of social evolution. One illustration is provided by the discussion of the GENITOR algorithm, a less well known alternative to the traditional "Holland type" GA, which is argued to be a more behaviourally plausible representation of an evolutionary decision process at least under the mental interpretation. Two other obvious areas of research which have remained unexplored to date are the possibility of inducing reusable subroutines in the decision process using the ADF approach in GP ([Koza 1994](#)) and attempts to implement any semantic (rather than purely syntactic) modification processes for decision trees, based on some background knowledge in the agent. It also surprising, given the relative ease of interpretation for single rules compared to GA strings or GP trees, that genuinely social models based on CS have not been more extensively developed. (It is also possible that new developments in the modelling of social evolution will inspire new instrumental approaches.)

6.4

On the final count, despite the emphasis on economic models in this paper, it is clear that with certain reservations, evolutionary modelling is appropriate to the other social sciences too. The most important dimension of difference between the disciplines is the degree of deliberation assumed to be involved in social action and the role of others in the decision process. However, it is clear that imitation and mutation processes are if anything even more appropriate for representing mechanisms of social influence or conformity. Since most descriptive models of social evolution dispense with the global fitness function, they can be used equally well for rather economic approaches involving a well defined utility function as for more social approaches based on reference points or aspiration levels. It has already been remarked that the more "cognitivist" approach is capable of drawing the different social science disciplines more closely together, at least in principle. This suggests a whole new field of modelling in which attempts are made to *integrate* the insights of the different disciplines within a simulation framework. (One example is provided by [Chattoe and Gilbert 1997](#)). Another relevant area of research that has not really been touched on in this paper is the use of *neural networks* to model the evolution of social processes. Like the use of evolutionary algorithms, this research is mainly concentrated in unorthodox economics ([Margarita](#)

1992, [Beltratti et al. 1996](#)), but there are also simulations which can be seen as modelling "low level sociality" and could be assigned to economics, sociology, anthropology or even archaeology ([Parisi 1997](#), [Pedone and Parisi 1997](#)). The reason for neglecting these approaches is not that neural networks are an inappropriate way of representing agent competencies, that depends very much on the behaviour concerned and on disciplinary preferences for models of conscious deliberate action or instinctive, behaviourist and socially contextual response. Rather, the difficulty is that there is little analysis of how the evolved neural networks do what they do. By contrast, far from being a black box, some accessible and conscious (though plainly not fully rational) mental process is widely believed to be what makes social action distinctively social. Be this as it may, neural network models draw attention to the fact that social evolution does not need to imply a rational or cognitivist view of agents.

6.5

In addition to these individual observations about the three components of the synthesis, an overarching conclusion can tentatively be drawn: that a coherent and complete interpretation of social evolution can at least be sketched. Furthermore, it is possible to move the debate on from the conceptual or methodological level (where some of the misunderstandings originated with contemporaries of Darwin) to the implementation of some concrete models which have been illustrated in this paper. Applications of evolutionary algorithms in new areas can do nothing but sharpen the debate further.

6.6

Finally, however, it seems appropriate to conclude a speculative paper such as this not with a definitive conclusion but rather with two unanswered questions, one empirical and one theoretical:

1. Just how much teleology is there in social organisations like markets? On one hand, it is clear that markets do not enforce profit maximisation. On the other, it seems to be constitutive of a competitive market that firms must satisfy the minimum requirement that they cover costs, at least in the long run. Because, at least theoretically, bounded rationality permits us to envisage a world in which everybody has a completely different cognitive model, we can see the considerable social advantages in agreements to set up institutions binding everyone participating in them to relatively similar behaviour. What can we say about the status of evolutionary processes taking place within these institutions? The answer has to be empirical rather than ideological. Diversity provides some guidance to the evolutionary pressure being exerted, though with the qualification that firms may respond to this pressure by specialising and deliberately occupying different niches. Simulations which explicitly model the processes by which social structures like markets come to be legislated might provide another fruitful field for future research.
2. To what extent should evolutionary algorithms (GA strings, CS rule bases and GP trees) even be *expected* to represent the whole process of decision? In particular, should we be thinking about the development of *hybrid* systems in which the processes of identifying and testing regularities to serve as operators and terminals would be carried out in parallel and using different techniques to the process of evolving good strategies based on those operators and terminals? Neural networks are an obvious technology for identifying patterns in data and could be seen to correspond to the innate (and largely unconscious) human ability to recognise patterns. [\[76\]](#) GP technology can also be used in the same way. In particular, we can envisage a firm devoting its "Statistics Department" to the task of predicting what other firms (or "the market") will do. If these predictions ever attain an adequate quality they will be incorporated into the strategy of the firm that has developed them. (Of course, this may not work for long if the other firms are doing the same!) This self prediction technique can also be applied instrumentally and descriptively to the task of understanding GP trees. The task of a subsidiary GP would be to predict the behaviour of a complex GP tree, either using the same language or a simplified one. The subsidiary GP is rewarded for the quality of its prediction

but heavily penalised for tree depth thus enforcing simplification. Instrumentally, this technique may be added to the automated syntactic simplification traditionally practised on GP trees, where ($* 1 (...)$) is replaced by (...) for example ([Koza 1992a](#)). Descriptively we can imagine firms continuously trying to identify predictable sub trees in their strategies using both syntactic and semantic techniques.

Acknowledgements

I would like to thank Bruce Edmonds and Luigi Marengo for helpful discussions of their models which are described here ([Edmonds 1997](#), [Dosi et al. 1994](#)). Any errors of interpretation remain my responsibility. My thanks are also due to the editor and two anonymous referees for their helpful comments.

Notes

¹ This term includes, but is not restricted to, Genetic Algorithms, Evolutionary Strategies, Genetic Programming and Classifier Systems. For further useful discussion see [Nissen \(1993\)](#).

² Orthodox is used here in the sense defined by [Nelson and Winter \(1982\)](#).

³ There are, however, some fascinating exceptions ([Bagehot 1887](#)).

⁴ The earliest paper satisfying this definition is generally agreed to be [Alchian \(1950\)](#).

⁵ Although genetic "fitness" is typically associated with numbers of live offspring, the tendency to reproduce as often as possible may be an effect as much as a cause of evolutionary selection, which merely requires persistence. Although the first species that began to reproduce rapidly would gain an advantage against species which did not, no species could gain any further advantage once rapid reproduction was universal.

⁶ This is a simplification. There are actually several processes by which genetic mixing and mutation take place, corresponding to various things which can happen to the chromosomes ([Weaver and Hedrick 1992](#): 87-94). In asexual reproduction there is no genetic mixing. However, asexual reproduction is limited to simple creatures and sexual reproduction can itself be seen as an evolved mechanism by which genetic mixing can take place if it is beneficial.

⁷ The validity of the assumption that social behaviour is no more than the sum of individual interactions remains contested in social science. Although economics favours a purely individualistic interpretation, sociologists sometimes appear to argue that norms and other social influences are more than the sum of individual mental contents.

⁸ We can thus see the ability to make these attributions and cultural (rather than genetic) transmission of behaviours as evolved mechanisms which reduce the wastefulness of the evolutionary process. Instead of "throwing away" a whole individual, it is possible for that individual to both identify and modify particular aspects of behaviour that are persistence threatening. We would expect these features to be both a cause and a consequence of increasing complexity. As individuals become more complex and the potential for diversity increases, it becomes more wasteful to "throw away" whole individuals. At the same time, self awareness and cultural transmission dramatically speed up processes of experimentation in relation to the

environment, which were previously restricted to one "try" per lifetime.

⁹ Note that *progressive* increases in complexity are not implied. That depends on whether a species with a given level of complexity persists robustly against all simpler organisms and their subsequent developments. So far, humanity appears to be in this position, but only so far.

¹⁰ The traditional example is that the children of hard working blacksmiths might be born with extra muscles. But equally the children of amputees might be born with missing limbs.

¹¹ Lamarckism arose before genes were discovered. It persisted, rather uneasily, when there still appeared to be a simple mapping between genes and "traits" like blue eyes, but has not survived the discovery of detailed biochemical mechanisms of development.

¹² They may even be meaningless ([Tintner 1941a](#), [Tintner 1941b](#)) or incalculable ([Thomas 1993](#)).

¹³ It will be recalled that the evolutionary process implies nothing about whether competitive or co-operative strategies will tend to persist more effectively ([Kropotkin 1939](#)).

¹⁴ Conceptual smoothness merely implies that we do not need to postulate any uncaused causes or mysterious jumps. It is no more than "good science" writ large. Although we will never know what went through the mind of the person who invented the wheel at the crucial moment, we can be confident that it was something to do with what they knew and had experienced, seen or perhaps been told. Although the discovery of fire may have caused an enormous discontinuity in social practices, there is no conceptual discontinuity implied by its discovery and application.

¹⁵ These were also chosen because they are the techniques used in most of the evolutionary modelling done in social science to date. Other related techniques include Evolutionary Programming ([Fogel et al. 1966](#), [Fogel 1991](#)), Evolution Strategies ([Rechenberg 1973](#), [Schwefel 1981](#)) and Simulated Annealing ([Davis 1987](#)).

¹⁶ Random generation is compatible with the conceptual smoothness of the evolutionary process in obliging the population of individuals to begin in a state of maximum disorder.

¹⁷ This interpretation has however been challenged by [De Jong \(1992\)](#).

¹⁸ This analogy raises an interesting point. It is *assumed* that the meat yield of cows is a purely genetic matter, just as it is assumed that an instrumental GA process can have a fixed fitness function attributed to it, as it clearly can in the simplest cases. If it turns out that meat yields are affected by the "social" organisation of the herd or that the fitnesses of population members actually affect each other, then the instrumental approach is an *empirically* invalid choice.

¹⁹ One reason for this is that the tree representation translates directly into a nested bracket representation suitable for list processing languages such as LISP.

²⁰ There are GA techniques which make use of variable length representations but they are typically unwieldy. See [Harvey \(1992\)](#) for further discussion.

²¹ It is hard to say conclusively that a GP tree is more expressive than a GA string, since any arbitrarily complex encoding for the GA can be hidden away in the fitness function. However, it is intuitive that the GP tree and the fixed "meanings" of its operators *taken together* are likely to be more economical than such an arbitrary encoding.

- ²² It is clear that this is a drawback for an instrumental GA which should ideally solve a problem as quickly as possible. However, it may not be a drawback for a GA used descriptively since it corresponds *prima facie* to human behaviour in focusing on important issues first. In practice, this analogy does not bear close inspection and other unrealistic aspects of the simple GA, such as the exogenous fitness function, are far more damaging to its use as a descriptive model.
- ²³ This can either mean that the program is "self executing" or that a process of interpretation and execution is *internal* to the individual. Either interpretation also requires some distinction between the genotype and phenotype.
- ²⁴ We can distinguish three interpretations of the robot controller example. In the first, both the environment and the controller are "abstract" and exist purely within the GP. In the second, the controller remains "abstract" but there is also an attempt to simulate a real environment in which the controller will evolve. In the third, the GP is actually a program inside a real robot, operating in a real environment.
- ²⁵ Social science seems to give inadequate general attention to differing degrees of the social. Although it rightly has no interest in situations where a single individual interacts purely with a physical environment, it offers little by way of guidance as to how we might expect behaviour to change as we move from agents interacting predominantly with an environment to agents predominantly interacting with each other.
- ²⁶ An interesting paper using this approach, applied to zoology rather than social science, is [Koza et al. \(1992\)](#).
- ²⁷ This description follows [Nissen \(1993\)](#).
- ²⁸ The structure of individuals could be subsumed straightforwardly into a GP, but not the operations which are carried out on those individuals.
- ²⁹ For the technicalities of this process, the reader is referred to [Holland et al. \(1986\)](#). Rules which fire and result in a good outcome "share out" the positive feedback, as do those which result in negative feedback. This sharing discourages bloating of the rule base.
- ³⁰ A GA which generates diverse populations to solve problems collectively is discussed in [Smith et al. \(1992\)](#). It is an interesting question whether the syntax of GP makes diverse populations redundant, impractical or neither.
- ³¹ Although the GP may produce a program that is equivalent to (rule1 AND rule2 AND rule3), it has no endogenous process to ensure hierarchical comprehensibility, so even if the operators are designed for easy interpretation, there is no guarantee that the overall structure will be easy to interpret. Furthermore, as the depth of the GP tree increases, so does the number of equivalent trees. These can be seen as a potential drawbacks of GP representations which may be addressed by current research into program modularisation through ADFs ([Koza 1994](#)).
- ³² The prevailing view that bounded rationality involves no more than applying rational principles to the process of cognition seems both incompatible with the bulk of what Simon wrote and incoherent on closer inspection, for precisely the reasons that Simon gives.
- ³³ In addition to the obvious point that these models are empirically very implausible!

³⁴ An example is provided by the history of atomic theory, where compounds were explained in terms of molecules, composed of atoms, composed of neutrons, protons and electrons, composed of quarks, composed of ...

³⁵ There may also be a signal extraction or credit assignment problem when an agent co-varies rules and meta-rules while trying to make sense of the environment.

³⁶ For a simple illustration, consider an individual who has a gun with a laser sight trained on him by a distant assassin. The assassin is so far away that they cannot be seen directly, but the reflection of the laser sight can. If the reflection can be seen near the target, the rational action is to take cover, otherwise it is to try and move to locate the assassin. Unfortunately, if the assassin trains the laser on some part of the victim's forehead, which is the best way to be sure of a kill, the only way the victim can see it is by holding up a mirror which, we suppose, is of such a shape and size that using it to look for the reflection will block the beam. Even if the blind spot is tiny, the victim will certainly die, despite having a rational decision process, because they have no way of simultaneously or sequentially observing the location of the reflection and the effect which the mirror is having on the beam. Admittedly, this example relies on an assumption about the shape of the mirror, but recall that its object was only to explain why even the smallest blind spot can destroy the possibility for rational action as economics defines it. In fact, one could argue that the shape of the mirror is consistent with defining some part of the forehead as a proper blind spot in the first place. A blind spot is more than somewhere you can't see when not looking in the correct direction! (Even if the victim had a second mirror of the same size and shape, and was very dextrous, they still could not use that mirror to observe the effect the laser was having on the back of the first mirror because it would then be the second mirror blocking the beam.) These ideas are developed further in work on "autopoietic systems" (Varela [1979](#), [1991](#)).

³⁷ However, they can be justified in *other* ways, for example as biologically evolved competencies. Unfortunately, this seems rather an admission of defeat as far as social science is concerned.

³⁸ This raises another interesting issue. Although one agent is hampered by the difficulties of obtaining adequate information about another, it does not suffer from any *conceptually necessary* blind spot in observing other agents. (It is an open question whether the blind spot of the first agent will definitely impair its understanding of the second.) It is truly the case that others may know us better than we know ourselves from a logical point of view!

³⁹ The same logic applies to any models constructed by the social scientist. This view also implies a coherentist rather than positivist notion of truth.

⁴⁰ Note that *common* knowledge of this shared correct model, which enables it to form the basis for rational action, is an even stronger assumption than that agents all merely happen to have the correct model ([Parikh 1990](#)).

⁴¹ There is an obvious but very important difference between the assumption that individuals *do* what they intend to do and that what *happens* is what agents intend. The fact that it is possible to miss this difference is illustrated by the argument between Alchian ([1950](#), [1953](#)) and Penrose ([1952](#), [1953](#)).

⁴² Although we are concentrating on cognitive irreversibility, it is obvious that irreversibility relating to autonomous physical processes ([Georgescu-Roegen 1971](#)) can also be part of the same framework.

⁴³ As [Nelson and Winter \(1982\)](#) have pointed out, orthodox economic theory struggles with genuine

novelty. We will have return to this issue as the same problem appears to beset simple evolutionary algorithms.

⁴⁴ An example is provided by the discussion of "big players" in [Koppl and Langlois \(1994\)](#).

⁴⁵ [Friedman \(1953\)](#) argues that the assumption that firms are profit maximisers can be justified by the fact that the market will tend to eliminate those firms that are not. This pseudo-Darwinian argument is still widely believed, despite being convincingly refuted by [Witt and Perske \(1982\)](#) and [Chiappori \(1984\)](#) among others. Ironically, the error in Friedman's reasoning is one that originates with Herbert Spencer over a hundred years ago!

⁴⁶ Although the abstract theory of the market may imply a universal framework of law standardising the behaviour of firms, a more detailed view reveals a rich structure of compliance, evasion, detection, political action, punishment and resistance. A question for the future is the extent to which more or less generally agreed but not universal practices *can* be said to constitute an external teleology and whether or not different degrees of agreement can be detected in the dynamics of different social processes.

⁴⁷ The minimal condition for market persistence is rather similar to that in biological systems, that the organism should "cover costs". Discussion of the widely held but mistaken belief that evolution produces optimal behaviour can be found in [Hodgson \(1991\)](#).

⁴⁸ There are some extremely interesting and widespread developments in industrial organisation which can be considered in evolutionary terms. The first of these is the existence of firms which develop by merger and buy out rather than pure production. The second is the existence of franchises and chains which really do reproduce branches whose operating practices may or may not be appropriate to their locales. These chains have to "trade off" economies resulting from common practices against loss of sales from local idiosyncrasies. They also have to consider the "carrying capacity" of the environment in siting their new branches.

⁴⁹ We should not perhaps presume on this issue, economic socialisation is still inadequately explored and although long run historical analyses of the nature of the firm do exist, these are largely discursive.

⁵⁰ Another interpretation is that if firms know that the encoding is common and that the fitness function is one to one, they can work back from what firms did to what their GA string must be, although this raises issues about timing. Perhaps telepathy is a more behaviourally plausible assumption after all!

⁵¹ If we assume sociable agents then intentions *are* transmissible, but it is not clear how accurate measures of utility could be transmitted, even with the most active desire to communicate. Transmission of information about money amounts is obviously not helpful!

⁵² It should be noted that not all game theoretic evolutionary models are based on replicator dynamics.

⁵³ Other differences between economic and biological games are discussed by [Selten \(1993\)](#).

⁵⁴ One could view replicator dynamics as a purely instrumental technique on this basis. Once it is assumed that all individuals are utility maximisers and it is utility which determines propagation, the attainment of game equilibria is a foregone conclusion. If, as many users of replicator dynamics argue, its important contribution is to show *which* equilibrium occurs, then it is no longer possible to

argue that behavioural assumptions about the process are irrelevant.

⁵⁵ One could diagnose the convergence problems in the Arifovic models knowing nothing at all about its interpretation.

⁵⁶ There are a number of models based on GA which avoid one or more of the shortcomings of the Arifovic models, for example [Curzon Price \(1997\)](#), [Lomborg \(1992\)](#) and [Vriend \(1994\)](#). Space considerations preclude their detailed discussion here, but some of the comments on GP models will also raise issues appropriate to those using GA techniques. In what follows, GP terminology will be used for simplicity.

⁵⁷ [Edmonds \(1997\)](#) uses a population of 40! Two additional comments can be made on this observation. The first is that because most of the cost of executing an instrumental GP comes in evaluating strategies, the actual coding of the GP involves deliberately *not* evaluating duplicate strategies, but using a hash table to weed out strategies that have already been tried ([Koza 1992a](#)). Even using this technique cannot avoid the wastefulness of strategies which are semantically equivalent but syntactically different. The second is that the number of strategies which can be "borne in mind" is far greater in a firm, where these would correspond to the views of particular individuals in the firm. The weights of strategies could then correspond to the number of people arguing for a particular view or their importance in the hierarchy.

⁵⁸ Rank based selection has the advantage that selection pressure does not drop as the population converges ([Whitley 1989](#)). It is an interesting question whether this also applies in social situations.

⁵⁹ This is trivial for a static fitness function, but not for a descriptive simulation in which fitness may change.

⁶⁰ This issue is discussed from an instrumental perspective by [Reeves \(1993\)](#).

⁶¹ One interesting possibility is that common knowledge may be mimicked by the fact that individuals tend to project their own models onto others. If, in fact, everybody did have the same model, this would be as good as common knowledge for the purposes of deciding what to do. Of course, it would fail miserably if models differed significantly. Perhaps this is why people have wars over religion rather than food.

⁶² In implementing a version of the [Dosi et al. \(1994\)](#) model, I followed them in assuming that firms selected strategies probabilistically on the basis of cumulative profits. When the strategies simply consisted of fixed prices rather than GP trees, all the firms quickly learnt to set the maximum price they could and rapidly priced themselves out of the market. Although there was an equilibrium with everyone charging the maximum permissible price, it was never observed because it required that all firms initially charged that price and did not deviate. This is an example where the fitness function for firms was not a complete representation of the restrictions on the market and thus prevented any firms from persisting. In this case, firms valued positive profit without limit but placed no value on market share at all.

⁶³ It is also the case that retention of profits *does* make firms far less susceptible to whatever discipline the market imposes and also perhaps more able to influence the terms on which market discipline is applied. This point seems to receive inadequate attention from evolutionary market apologists. Since this escape from discipline reflects the very success of firms, perhaps it is true as Marx suggested that competition, if not capitalism contains the seeds of its own destruction!

⁶⁴ They can be assigned the fitness of their parents providing that firms do not have an infinite

memory for cumulated profit.

⁶⁵ The relevance of qualitative physics is illustrated in a paper by [Sims \(1991\)](#). This describes a simulation in which artificial organisms constructed of rigid blocks and joints are "evolved" to perform simple tasks like "capturing" food. Rather than seeing the properties of joints and blocks as attributes of the organism, as classical Artificial Intelligence might, these attributes are modelled as functions of environmental factors such as "gravity" and "mass". The result is that the behaviour of arbitrary evolved combinations of blocks and joints is always defined.

⁶⁶ This conclusion is also reached in the Artificial Life literature ([Langton 1989](#), [Langton et al. 1991](#), [Langton 1993](#)). Complex environments have another interesting effect. Because evolutionary algorithms are very good at optimisation, they can often produce strategies which, in being self-evidently silly, reveal something about the inadequacy of the environmental specification. Developing the simulation thus becomes a co-evolutionary process, with behaviourally implausible strategies sometimes revealing unrealistic assumptions about the agent and sometimes about the environment.

⁶⁷ This distinction is often badly muddled in replicator dynamics models. If individuals just are their strategies then there is no difficulty, except that this is a very unrealistic view of individuals. If there is a difference between the strategy and the actions it produces, then strong knowledge assumptions are required to avoid worrying about how inferences can be made from actions back to strategies.

⁶⁸ In one firm where I worked, employee research group membership was not listed in the internal phone directory so nobody outside the company could get an overall picture of the amount of research being done in different areas!

⁶⁹ In [Chattoe and Gilbert \(1997\)](#), agents learn how to budget by evolving budgeting plans individually, but they are also able to observe the consumption patterns of other agents. The result is a co-evolution of effective budgeting plans and stratified lifestyles based on income.

⁷⁰ There is an analogy here with the highly instrumentally effective Island Model class of Genetic Algorithms in which relatively small populations evolve in parallel, but transmit their best strategies at random to other populations from time to time ([Gordon and Whitley 1993](#)).

⁷¹ Another way of seeing the issue of profits reducing the competitive pressure on firms is to ask whether profit tends to "evaporate" or not. Individuals can never achieve more than a certain level of energy or wakefulness and this continues to drain away whatever they do.

⁷² This program is proceeding for instrumental GP in the development of ADF techniques ([Koza 1994](#)) which may provide interesting insights which descriptive models can use.

⁷³ There is an obvious role here for the sort of computational organisation theory models devised by Carley and others ([Carley and Prietula 1994](#)).

⁷⁴ Large trees are behaviourally implausible as well as almost impossible to interpret. An additional danger, suggested by the Dosi *et al.* results on price tracking is that a large GP tree may just become a lookup table for the state of the present environment. There has been little research so far on pulling GP trees that are apparently successful out of one environment and putting them in another.

⁷⁵ Another consequence of semantic analysis is the possibility of analogical reasoning. If a sub tree has a "meaning" attached to it, like "total costs" it becomes possible to substitute one sub tree for another in a way that is somewhat directed, as it relies on similarities of meaning. This also applies

at the level of whole GP trees, which may only produce a number as output, but it is the use to which the organisation puts that number which determines its meaning. Thus a tree that solves one environmental problem, may be more likely to solve an analogous problem.

⁷⁶ Work of this kind could make use of research such as that by [Sun and Bookman \(1993\)](#) on the integration of neural and symbolic processing.



References

- ALCHIAN, Armen A. (1950) 'Uncertainty, Evolution and Economic Theory', *Journal of Political Economy*, 58(3), June, pp. 211-222.
- ALCHIAN, Armen A. (1953) 'Comment: Biological Analogies in the Theory of the Firm', *American Economic Review*, 43(3), Part 1, September, pp. 600-603.
- ARIFOVIC, Jasmina (1990) 'Learning by Genetic Algorithms in Economic Environments', Working Paper 90-001, Economics Research Program, Santa Fe Institute, New Mexico, October.
- ARIFOVIC, Jasmina (1994) 'Genetic Algorithm Learning and the Cobweb Model', *Journal of Economic Dynamics and Control*, 18(1), Special Issue on Computer Science and Economics, January, pp. 3-28.
- ARIFOVIC, Jasmina (1995) 'Genetic Algorithms and Inflationary Economies', *Journal of Monetary Economics*, 36(1), August, pp. 219-243.
- ARIFOVIC, Jasmina (1996) 'The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies', *Journal of Political Economy*, 104(3), June, pp. 510-541.
- ARIFOVIC, Jasmina and EATON, Curtis (1995) 'Coordination via Genetic Learning', *Computational Economics*, 8, pp. 181-203.
- ARTHUR, Brian (1994) 'Inductive Reasoning and Bounded Rationality', *American Economic Association Papers*, 84, pp. 406-411.
- BAGEHOT, Walter (1887) *Physics and Politics or Thoughts on the Application of the Principles of 'Natural Selection' and 'Inheritance' to Political Society*, eighth edition (London: Kegan Paul, Trench and Company).
- BAUMOL, William J. and QUANDT, Richard E. (1964) 'Rules of Thumb and Optimally Imperfect Decisions', *American Economic Review*, 64(2), Part 1, March, pp. 23-46.
- BELTRATTI, A., MARGARITA, S. and TERNA, P. (1996) *Neural Networks for Economic and Financial Modelling* (London: International Thompson Computer Press).
- BRACHMAN, R. and LEVESQUE, Hector J. (eds.) (1985) *Readings in Knowledge Representation* (Los Altos, CA: Morgan Kaufmann).
- BULLARD, James and DUFFY, John (1994) 'A Model of Learning and Emulation with Artificial Adaptive Agents', Draft Paper, Research Department, Federal Reserve Bank of St Louis, MO and Department of Economics, University of Pittsburgh, May.
- CARLEY, Kathleen M. and PRIETULA, Michael (eds.) (1994) *Computational Organization Theory* (Hillsdale, NJ: Lawrence Erlbaum Associates).

- CHATTOE, Edmund (1994) 'Some Thoughts on the Methodology of Computational Economics', paper presented at the conference *Computational Economics*, Amsterdam, 8-11 June.
- CHATTOE, Edmund (1996) 'Why Are We Simulating Anyway? Some Answers From Economics', in Troitzsch, Klaus G., Mueller, Ulrich, Gilbert, Nigel and Doran, Jim E. (eds.) *Social Science Microsimulation* (Berlin: Springer-Verlag), pp. 78-104.
- CHATTOE, Edmund and GILBERT, Nigel (1996) 'The Simulation of Budgetary Decision Making and Mechanisms of Social Evolution', paper presented at ISA '96: Fourth International Social Science Methodology Conference, Colchester, Essex.
- CHATTOE, Edmund and GILBERT, Nigel (1997) 'A Simulation of Adaptation Mechanisms in Budgetary Decision Making', in Conte, Rosaria, Hegselmann, Rainer and Terna, Pietro (eds.) *Simulating Social Phenomena* (Berlin: Springer-Verlag), pp. 401-418.
- CHIAPPORI, Pierre-Andre (1984) 'Sélection Naturelle et Rationalité Absolue des Entreprises', *Revue Economique*, 35(1), January, pp. 87-108.
- CURZON PRICE, Tony (1997) 'Using Coevolutionary Programming to Simulate Strategic Behaviour in Markets', *Journal of Evolutionary Economics*, 7, pp. 219-254.
- DARWIN, Charles (1968) *On The Origin of Species by Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life* (Harmondsworth: Penguin).
- DAVIS, Lawrence D. (ed.) (1987) *Genetic Algorithms and Simulated Annealing* (London: Pitman).
- DAWKINS, Richard (1989) 'The Evolution of Evolvability' in Langton, Christopher G. (ed.) *Artificial Life: The Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, Los Alamos, New Mexico, September 1987*, SFI Studies in the Sciences of Complexity Volume VI (Redwood City, CA: Addison-Wesley), pp. 201-220.
- DE JONG, Kenneth (1992) 'Are Genetic Algorithms Optimisers?', in Männer, Reinhard and Manderick, Bernard (eds.) *Parallel Problem Solving from Nature, 2, Proceedings of the Second Conference, Brussels, Belgium, 28-30 September 1992* (Amsterdam: North-Holland), pp. 3-13.
- DORAN, Jim and GILBERT, Nigel (1994) 'Simulating Societies: An Introduction', in Gilbert, Nigel and Doran, Jim (eds.) *Simulating Societies: The Computer Simulation of Social Phenomena* (London: UCL Press), pp. 1-18.
- DOSI, Giovanni, MARENGO, Luigi and VALENTE, Marco (1994) 'Norms as Emergent Properties of Adaptive Learning: The Case of Economic Routines', paper presented at the Royal Economic Society Conference, Exeter, 28-31 March.
- EDEY, Maitland A. and JOHANSON, Donald C. (1989) *Blueprints: Solving the Mystery of Evolution* (Oxford: Oxford University Press).
- EDMONDS, Bruce (1997) 'Gossip, Sexual Recombination and the El Farol Bar: Modelling the Emergence of Heterogeneity', Centre for Policy Modelling Discussion Paper Number CPM-97-31, Manchester Metropolitan University, <http://www.cpm.mmu.ac.uk/cpmrep31.html>.
- EDMONDS, Bruce and MOSS, Scott (1996) 'Modelling Bounded Rationality using Evolutionary Techniques', Centre for Policy Modelling Discussion Paper Number CPM-96-10, Manchester Metropolitan University, <http://www.cpm.mmu.ac.uk/cpmrep10.html>.
- EPSTEIN, J. M. and AXTELL, R. (1996) *Growing Artificial Societies: Social Science from the*

Bottom Up (Cambridge, MA: Brookings Institution/MIT Press).

FISHER, R. A. (1930) *The Genetical Theory of Natural Selection* (Oxford: Clarendon Press).

FOGEL, L. J., OWENS, A. J. and WALSH, M. J. (1966) *Artificial Intelligence Through Simulated Evolution* (New York, NY: John Wiley and Sons).

FOGEL, D. B. (1991) 'The Evolution of Intelligent Decision Making in Gaming', *Cybernetics and Systems*, 22, pp. 223-236.

FRIEDMAN, Milton (1953) 'The Methodology of Positive Economics', in *Essays in Positive Economics* (Chicago, IL: University of Chicago Press), pp. 4-14.

GEORGESCU-ROEGEN, Nicholas (1971) *The Entropy Law and the Economic Process* (Cambridge, MA: Harvard University Press).

GOLDBERG, David E. (1989) *Genetic Algorithms in Search, Optimization and Machine Learning* (Reading, MA: Addison-Wesley).

GORDON, V. Scott and WHITLEY, Darrell (1993) 'Serial and Parallel Genetic Algorithms as Function Optimizers', in Forrest, Stephanie (ed.) *Proceedings of the Fifth International Conference on Genetic Algorithms, University of Illinois at Urbana-Champaign, 17-23 July 1993* (San Mateo, CA: Morgan Kaufmann), pp. 177-183.

HARVEY, Inman (1992) 'The SAGA Cross: The Mechanics of Recombination for Species for Variable Length Genotypes', in Männer, Reinhard and Manderick, Bernard (eds.) *Parallel Problem Solving from Nature, 2, Proceedings of the Second Conference, Brussels, Belgium, 28-30 September 1992* (Amsterdam: North-Holland), pp. 269-278.

HODGSON, Geoffrey M. (1991) 'Economic Evolution: Intervention Contra Pangloss', *Journal of Economic Issues*, 25(2), June, pp. 519-533.

HODGSON, Geoffrey M., SAMUELS, W. J. and TOOL, Marc R. (eds.) (1994) *The Elgar Companion to Institutional and Evolutionary Economics* (Aldershot: Edward Elgar).

HOFSTADTER, Richard (1955) *Social Darwinism in American Thought* (Boston, MA: Beacon Press).

HOLLAND, John H. (1975) *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence* (Ann Arbor, MI: University of Michigan Press).

HOLLAND, John H. (1992) *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, second edition (Cambridge, MA: The MIT Press/A Bradford Book).

HOLLAND, John H., HOLYOAK, K. J., NISBETT, R. E. and THAGARD, P. R. (1986) *Induction: Processes of Inference, Learning, and Discovery* (Cambridge, MA: MIT Press).

JONES, Greta (1980) *Social Darwinism and English Thought: The Interaction Between Biological and Social Theory* (Brighton: Harvester).

KOPPL, Roger and LANGLOIS, Richard N. (1994) 'When Do Ideas Matter? A Study in the Natural Selection of Social Games', *Advances in Austrian Economics*, 1, pp. 81-104.

KOZA, John R. (1991) 'Evolving a Computer Program to Generate Random Numbers Using the Genetic Programming Paradigm', in Belew, Richard K. and Booker, Lashon B. (eds.) *Proceedings of the Fourth International Conference on Genetic Algorithms, UCSD, San Diego 13-16 July 1991* (San Mateo, CA: Morgan Kaufmann), pp. 37-44.

KOZA, John R. (1992a) *Genetic Programming: On the Programming of Computers by Means of Natural Selection and Genetics* (Cambridge, MA: MIT Press/A Bradford Book).

KOZA, John R. (1992b) 'A Genetic Approach to Econometric Modelling', in Bourguine, Paul and Walliser, Bernard (eds.) *Economics and Cognitive Science, Selected Papers from the Second International Conference on Economics and Artificial Intelligence, Paris, 4-6 July 1990* (Oxford: Pergamon), pp. 57-75.

KOZA, John R. (1992c) 'Evolution and Coevolution of Computer Programs to Control Independently Acting Agents', in Meyer, Jean-Arcady and Wilson, Stewart W. (eds.) *From Animals to Animats, Proceedings of the First International Conference on Simulation of Adaptive Behaviour (SAB 91), Paris, 24-28 September 1991* (Cambridge, MA: The MIT Press/A Bradford Book), pp. 366-375.

KOZA, John R. (1994) *Genetic Programming II: Automatic Discovery of Reusable Programs* (Cambridge, MA: The MIT Press/A Bradford Book).

KOZA, John R., RICE, James P. and ROUGHGARDEN, J. (1992) 'Evolution of Food Foraging Strategies for the Caribbean *Anolis* Lizard Using Genetic Programming', Working Paper 92-06-028, Santa Fe Institute, New Mexico.

KROPOTKIN, Prince Petr (1939) *Mutual Aid: A Factor in Evolution* (Harmondsworth: Penguin).

LANGTON, Christopher G. (ed.) (1989) *Artificial Life: The Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, Los Alamos, New Mexico, September 1987*, SFI Studies in the Sciences of Complexity, Proceedings Volume VI (Redwood City, CA: Addison-Wesley).

LANGTON, Langton, Christopher G., TAYLOR, Charles, FARMER, J. Doyme and RASSMUSSEN, Steen (eds.) (1991) *Artificial Life II: Proceedings of the Workshop on Artificial Life, Santa Fe, New Mexico, February 1990*, SFI Studies in the Sciences of Complexity, Proceedings Volume X (Redwood City, CA: Addison-Wesley).

LANGTON, Langton, Christopher G. (ed.) (1993) *Artificial Life III: Proceedings of the Workshop on Artificial Life, Santa Fe, New Mexico, June 1992*, SFI Studies in the Sciences of Complexity, Proceedings Volume XVII (Redwood City, CA: Addison-Wesley).

LOMBORG, Bjorn (1992) 'Cooperation in the Iterated Prisoner's Dilemma', Papers on Economics and Evolution, Number 9302, edited by the European Study Group for Evolutionary Economics, Max Planck Institute, Evolutionary Economics Unit, Jena, Germany.

MARGARITA, Sergio (1992) 'Genetic Neural Networks for Financial Markets: Some Results', in Neumann, Bernd (ed.) *Proceedings of the 10th European Conference on Artificial Intelligence, Vienna, Austria, 3-7 August 1992* (New York, NY: John Wiley and Sons), pp. 211-213.

MCCAIN, Roger A. (1992) *A Framework for Cognitive Economics* (Westport, CT: Praeger).

MCDUGALL, W. (1927) 'An Experiment for Testing the Hypothesis of Lamarck', *British Journal of Psychology*, 17, pp. 267-304.

MORAVEC, Hans (1989) 'Human Culture: A Genetic Takeover Underway', in Langton, Christopher G. (ed.) *Artificial Life: The Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, Los Alamos, New Mexico, September 1987*, SFI Studies in the Sciences of Complexity Volume VI (Redwood City, CA: Addison-Wesley), pp. 167-199.

MOSS, Scott and EDMONDS, Bruce (1994) 'Modelling Learning as Modelling', CPM Report 3, Centre for Policy Modelling, Manchester Metropolitan University,
<http://www.cpm.mmu.ac.uk/cpmrep03.html>.

NELSON, Richard R. and WINTER, Sidney G. (1982) *An Evolutionary Theory of Economic Change* (Cambridge, MA: Belknap/Harvard University Press).

NISSEN, Volker (1993) 'Evolutionary Algorithms in Management Science - An Overview and List of References', Papers on Economics and Evolution, Number 9303, edited by the European Study Group for Evolutionary Economics, Max Planck Institute, Evolutionary Economics Unit, Jena, Germany.

PARIKH, Rohit (1990) 'Recent Issues in Reasoning about Knowledge', in Parikh, Rohit (ed.) *Theoretical Aspects of Reasoning about Knowledge: Proceedings of the Third Conference (TARK 3)* (San Mateo, CA: Morgan Kaufmann), pp. 3-9.

PARISI, D. (1997) 'What to Do With a Surplus', in Conte, Rosaria, Hegselmann, Rainer and Terna, Pietro (eds.) *Simulating Social Phenomena* (Berlin: Springer-Verlag), pp. 133-151.

PEDONE, R. and PARISI, D. (1997) 'In What Kind of Social Groups can "Altruistic" Behaviors Evolve?', in Conte, Rosaria, Hegselmann, Rainer and Terna, Pietro (eds.) *Simulating Social Phenomena* (Berlin: Springer-Verlag), pp. 195-201.

PENROSE, Elizabeth Tilton (1952) 'Biological Analogies in the Theory of the Firm', *American Economic Review*, 42(4), December, pp. 804-819.

PENROSE, Elizabeth Tilton (1953) 'Rejoinder: Biological Analogies in the Theory of The Firm', *American Economic Review*, 43(3), Part 1, September, pp. 603-607.

PETR, Jerry L. (1982) 'Economic Evolution and Economic Policy: Is Reaganomics a Sustainable Force?', *Journal of Economic Issues*, 16(4), December, pp. 1005-1012.

RECHENBERG, I. (1973) *Evolutionsstrategie: Optimierung Technischer Systeme nach Prinzipien der Biologischen Evolution* (Stuttgart: Frommann-Holzboog).

REEVES, Colin R. (1993) 'Using Genetic Algorithms with Small Populations', in Forrest, Stephanie (ed.) *Proceedings of the Fifth International Conference on Genetic Algorithms, University of Illinois at Urbana-Champaign, 17-23 July 1993* (San Mateo, CA: Morgan Kaufmann), pp. 92-99.

RIDLEY, Mark (1985) *The Problems of Evolution* (Oxford: Oxford University Press).

RUNCIMAN, W. G. (1998) 'The Selectionist Paradigm and Its Implications for Sociology', *Sociology*, 32(1), February, pp. 163-188.

SCHWEFEL, Hans-Paul (1981) *Numerical Optimization of Computer Models* (New York: John Wiley and Sons).

SELTEN, Reinhard (1993) 'Evolution, Learning, and Economic Behaviour', *Games and Economic Behaviour*, 3(1), February, pp. 3-24.

- SIMON, Herbert A. (1972) 'Models of Bounded Rationality', in McGuire, C. B. and Radner, Roy (eds.) *Decision and Organization: A Volume in Honour of Jacob Marshack* (Amsterdam: North-Holland), pp. 161-176.
- SIMON, Herbert A. (1976) 'From Substantive to Procedural Rationality', in Latsis, Spiro J. (ed.) (1976) *Method and Appraisal in Economics* (Cambridge: Cambridge University Press), pp. 129-148.
- SIMON, Herbert A. (1978) 'Rationality as a Process and as a Product of Thought', *American Economic Review*, 68(2), May, pp. 1-16.
- SIMON, Herbert A. (1981) *Models of Bounded Rationality*, two volumes (Cambridge, MA: MIT Press).
- SIMS, Karl (1991) 'Interactive Evolution for Computer Graphics', *Computer Graphics*, SIGGRAPH Conference Proceedings, 25(4), pp. 343-350.
- SLOMAN, Aaron (1985) 'Why We Need Many Knowledge Representation Formalisms', CSRP 052, Cognitive Science Research Paper, University of Sussex.
- SMITH, Robert Elliot, FORREST, Stephanie and PERELSON, Alan S. (1992) 'Searching for Diverse Cooperative Populations with Genetic Algorithms', TCGA Report Number 92002, The Clearinghouse for Genetic Algorithms, Department of Engineering Mechanics, University of Alabama (Tuscaloosa), 19 March.
- SUN, Ron and BOOKMAN, L. (eds.) (1993) *Computational Architectures Integrating Neural and Symbolic Processing* (Boston, MA: Kluwer Academic Publishers).
- THOMAS, J. (1993) 'Non-Computable Rational Expectations Equilibria', *Mathematical Social Science*, 25, pp. 133-142.
- TINTNER, Gerhard (1941a) 'The Theory of Choice Under Subjective Risk and Uncertainty', *Econometrica*, 9(3-4), July-October, pp. 298-304.
- TINTNER, Gerhard (1941b) 'The Pure Theory of Production Under Technological Risk and Uncertainty', *Econometrica*, 9(3-4), July-October, pp. 305-312.
- VARELA, Francisco J. (1979) *Principles of Biological Autonomy* (New York, NY: Elsevier North Holland).
- VARELA, Francisco J., THOMPSON, Evan and ROSCH, Eleanor (1991) *The Embodied Mind: Cognitive Science and Human Experience* (Cambridge, MA: MIT Press).
- VEGA-REDONDO, Fernando (1996) *Evolution, Games, and Economic Behaviour* (Oxford: Oxford University Press).
- VRIEND, Nicolaas J. (1994) 'Self-Organized Markets in a Decentralized Economy', Working Paper 94-03-013, Economics Research Program, Santa Fe Institute, New Mexico.
- WEAVER, Robert F. and HEDRICK, Philip W. (1992) *Genetics*, second edition (Dubuque, IA: William C. Brown).
- WEIBULL, Jörgen W. (1995) *Evolutionary Game Theory* (Cambridge, MA: MIT Press).
- WILSON, Edward O. (1975) *Sociobiology: The New Synthesis* (Cambridge, MA: Harvard

University Press).

WILSON, Edward O. (1978) *On Human Nature* (Cambridge, MA: Harvard University Press).

WINSTON, Gordon C. (1987) 'Activity Choice: A New Approach to Economic Behaviour', *Journal of Economic Behaviour and Organization*, 8(4), December, pp. 567-585.

WITT, Ulrich and PERSKE, Joachim (1982) *SMS - A Program Package for Simulating and Gaming of Stochastic Market Processes and Learning Behaviour*, Lecture Notes in Economics and Mathematical Systems, Number 202 (Berlin: Springer-Verlag).

WHITLEY, Darrell (1989) 'The GENITOR Algorithm and Selection Pressure: Why Rank-Based Allocation of Reproductive Trials is Best', in Schaffer, J. David (ed.) *The Proceedings of the Third International Conference on Genetic Algorithms, George Mason University, 4-7 June 1989* (San Mateo, CA: Morgan Kaufmann), pp. 116-121.

WITTGENSTEIN, Ludwig (1953) *Philosophical Investigations* (Oxford: Basil Blackwell).

[Return to Contents of this issue](#)

© [Copyright Journal of Artificial Societies and Social Simulation, 1998](#)

