# Towards an Automata approach of Institutional (and Evolutionary) Economics

Philip Mirowski

and

Koye Somefun

University of Notre Dame

September 1998

* This is a draft of a paper to be presented at Boston at the fifth international conference of the Society for Computational Economics. The authors would like to thank the participants at SIAC and audiences at the Universities of Melbourne, Monash, Modena, and Notre Dame for their comments.

## 0. Introduction

A computational approach towards economics potentially enriches economic science beyond increasing available mathematical techniques. Computational economics (CE) can foster a viable and rich institutional economics that encourages both mathematical rigor and historical relevance while avoiding the mechanical aspects of conventional neoclassical theory. Here we begin such an approach by regarding markets as computational entities or literal automata, where 'automata' refers to the formal notion of a computational device developed in computability theory (a branch of formal logic). Already there is a large literature --in particular in experimental economics and finance-- dealing with concepts similar to our suggested automata approach. In the case of the experimental literature, numerous experiments have been conducted that analyze the dependence of economical performance on the market institution. Similarly in the financial literature, as a consequence of the ongoing automation of markets, it has become an issue to analyze the relevance of different market designs. In other words, both situations treat economic performance as dependent on the context given by the market institutions. This paper takes this approach one step further by actually perceiving markets as computational entities.

We begin by introducing the reader to what we call "an agent based approach". This literature, perhaps best represented by the work of Alan Kirman, but found throughout the *avant garde* of the profession, tends to characterize agents as flawed automata or limited computational entities. While there is much to admire in this program, we maintain that invoking a gestalt reversal which regards markets as computational devices has considerable advantage. Next, we discuss some of von Neumann's work on automata theory and its relevance for economic theory. The format in which we choose to do this is to juxtapose von Neumann's work on automata with Simon's work –on what we will call—simulacra. Finally, we illustrate the proposed automata apporach towards markets by providing an example that shows how markets can be encoded as automata.

## 1. Waiting for a little Spontaneous Order

The effect of the computer upon modern developments in economic theory is a topic still in its infancy. All manner of novel and imaginative research programs owe their genesis to the spread of the computer throughout the whole gamut of postwar sciences: Artificial Intelligence, Artificial Life, bounded rationality, cognitive science, information theory, nonlinear dynamics, genetic algorithms, simulation exercises, and so forth. Situated in the midst of this eruption of innovations, it is more than a little difficult to discern the main lines of development inherent in present trends, in part because nascent formations get conflated and confused with other parallel trends, such as the relative health or debilitation of the neoclassical orthodoxy, changing cultural images (evolution, maximum entropy) of the iconic instantiations of social order, sufficient grasp of the implications of intellectual innovations in other sciences, and so on. We broach this issue, not because we believe we can supply answers to these big questions in this paper (we can't!), but rather because we have been struck by the diversity and incompatibility of theoretical exercises which have been suggested at the Santa Fe Institute, the Society for Computational Economics meetings, and at the Ancona conference. Too often, the computer in its protean manifestations has provided yeoman service as the lone common denominator which permits all manner of theorists to pronounce on where economics is going in one others' presence, without taking into account the importance of rival programs, empirical commitments, or scientific constraints. The purpose of this paper is to lay out two very different trends in modern theory in stark outline, both owing their very conceptual essence to the computer, and yet each individually imposing an entirely alternative framework upon the economic phenomenon. Both champion the point of departure of their new economics upon acknowledgement of heterogeneous interactive economic agents, and both can be described as descending from John von Neumann's invention of the theory of automata; but beyond that, they are as different as night and day. We shall dub these alternatives 'agents as flawed computers' versus 'markets as evolving computational entities'.

Although there are many eminent representatives of the first category, we shall take Alan Kirman's lecture "The Emergency of Market Organization" as an admirable summary of many of the objectives of this program.[1] Kirman begins by diagnosing the failure of the neoclassical Walrasian orthodoxy in achieving an understanding of market processes. First in the bill of indictment is the penchant for treating all rational agents as fundamentally alike; next, there is the confusion of rationality with omniscience; and then there is the problem that "too much interdependence destroys the possibility of moving from the micro level to the macro level". The prescription to cure these ills is to keep the orthodox notion of market equilibrium more or less entact, but to loosen up on the specification of individual rationality. Of necessity, this requires

1 This is the lecture presented as the plenary at the Ancona conference 1998.

devoting greater theoretical attention to the psychological makeup of the economic agents. Kirman suggests that this implies paying closer attention to formats of information processing; this might mean more realistic specifications of the modification of individual expectations, or else the resort to imitation of other agents to solve short-term calculational shortcomings, or it may just mean treating the agent more like a computer handicapped with various computational limitations. These modelling choices find their counterparts in techniques such as cellular automata, spin glasses on lattices, stochastic graph theory, and other formalisms pioneered primarily in computer science, but also in the physical sciences. Taking into account the heterogeneity of agent rationalizations and giving due consideration to agent learning does raise the issue of the existence at the macro level of regularities with anything approaching a law-like character. Kirman believes that acknowledgment of the cognitive limitations of individuals leads inexorably to an appreciation of the role of institutional structures in market coordination. As he himself has practiced what he preaches, this involves looking at the sequential character of markets, their closing rules, distributions of realized price dispersions, and other phenomena often overlooked by neoclassical theory.

Although we applaud each and every one of the improvements which Kirman urges over the neoclassical orthodoxy, we do feel impelled to question the extent to which the project is firmly grounded in issues of computation and the externally given identity of individual intentionalities. There is also the troubling question of what the theoretical acknowledgement of the fundamental heterogeneity of agents is intended to achieve. We fully admit that people really do diverge dramatically in terms of their notions of rationality and their cognitive capacities; but the deeper question is why and how this should matter for economics.

As is well known, the Walrasian project took as its givens tastes, technologies and endowments, because it sought to reduce all causal relations to those putatively 'natural' determinants. The very notion of "equilibrium" as appropriated from physics makes little sense without this very important proviso (Mirowski, forthcoming). The first three generations of neoclassicals did not seek to explain psychological regularities, since they wished to treat their postulates concerning utility as universal abstractions of human rationality. They were not bothered by the fact that they were formally treating everyone as identical; indeed, for them, this was a vindication of their scientific aspirations (however much it might clash with the ideology of individual freedom of choice, or indeed, any notion of free will). The purpose of "equilibrium" was to portray the market as using prices to balance the natural givens to an exquisitely precise degree: hence the sobriequet of "marginalism".

Once economists relinquish this particular framework in the interests of allowing for agent heterogeneity, this project tends to lose its rationale. Indeed, this is one way of understanding the significance of the Sonnenschein/Mantel/Debreu theorems. There it has been demonstrated that the project of Walrasian equilibrium largely 'works' when everyone is identical, but generally comes acropper when individual differences are analytically acknowledged. The "no-trade" theorems likewise show that with full rationality but the possibility of differential information, rational agents would abjure participation in the market due to strategic considerations. Indeed, these and other developments have been some of the major motivations behind the new-found fascination with older traditions of bounded rationality, finite automata playing games, machine learning and the like (Mirowski, forthcoming).

Nevertheless, there persist deeper contradictions adhering to attempts to import computational and psychological considerations into the orthodox equilibrium program. The most apparent incongruity is that most exercises which are promoted as advocating agent homogeneity in fact rarely live up to their billing. Instead, agents are frequently equipped with only the most rudimentary algorithmic capacities, perhaps differing only by a parameter or two, and then the model is deployed to demonstrate certain aggregate stochastic regularities. While this does reveal a different approach to bridging the micro-macro divide, it does not begin to acknowledge the analytical significance of true diversity in the population; nor does it lead to guidelines for a theoretically informed empiricism about the nature and importance of agent diversity. What generally happens is that economic theorists end up strenuously avoiding cognitive science and behavioral psychology in the interests of producing tractable equilibrium models. In practice, the ritual adherence to methodological individualism turns out to consist largely of empty gestures.

The second contradiction of this newer work is that, while most of its findings are based to a greater or lesser extent upon formalisms developed in the sciences most heavily influenced by the computer, these economic models predominantly avoid any incorporation of formal theories of computation. For instance, agents may be treated as relatively simple automata; but there is no consideration of why certain limitations are imposed (say, limited memory capacity) whereas other obvious limitations are transcended (say, infinite precision in computation). Fundamental barriers to algorithmic rationality such as the halting problem are rarely if ever addressed. This channels the work of modelling away from analytics and towards *simulation*, as discussed below. A nagging weakness of the program then becomes the palpable absence of any widely agreed-upon criteria of what would constitue a good or superior simulation.

The third contradiction of this trend is that it never stops to consider why the heterogeneity of agents is so very important for economics, as opposed to the conundra of late Walrasianism. Here we might suggest that the major difference between a science predicated upon physics-- be it classical or statistical mechanics-- and one informed by biology, is that the latter recognizes the central significance of heterogeneity as allowing for the possibility of selection, and therefore evolutionary change. Formal acknowledgement of agent heterogeneity is rather unavailing if the outcome of the modelling exercise is for everyone to end up as effectively identical, as frequently happens in the learning and 'evolutionary' game literature. A deeper lesson to be learned from the persistence of heterogeneity is that it is fundamentally opposed to the very notion of equilibrium bequeathed from the physical sciences. Heterogeneity of entities or organisms, when coupled with various selection mechanisms, maintains a diverse population which then displays a capacity in the aggregate to adapt to multiple fluctuating environments.

It is our contention that the issues of computation, heterogeneity of agents and concepts of evolution have been given cogent interpretations in economics in the era since World War II, but that economic theorists have remained relatively deaf to these discussions. It seems that intellectual clarification of the issues is not a matter of spontaneous order, but rather that economists sometimes have to have their memories jogged by a little history. We would suggest that the two major incarnations of the computational approach in economics can be associated with the names of Herbert Simon and John von Neumann, respectively. While the former is most closely associated with appeals to "bounded rationality" in economic theory (Conlisk, 1996), and the latter is honored as the progenitor of game theory, neither doctrine sufficiently represents or captures the manner in which these figures strove to combine computation and evolution into a single framework. A survey of their complete corpus, which in both cases encompasses numerous works outside the narrow ambit of economics, reveals that each has provided a framework for a formal theory of evolution which abstracts away the 'wet' details of biology; both are predicated upon the computer as both exemplar and instantiation of how evolution is thought to work. Yet, even though they both began from very similar aspirations, their respective frameworks are so different that they can stand as contrasting rival research programs. Indeed, von Neumann sought to subsume evolution under a general theory of "automata" (1966), whereas Simon believed it could be treated under the heading of the 'sciences of the artificial' (1981), or what we shall call here a theory of "simulacra".

1.1 *Simon's Simulacra versus von Neumann's Automata.*

Before venturing into a more detailed discussion of what is entailed with a computational understanding of markets we will use this subsection to give a more explicit account of Simon's theory of simulacra and von Neumann's theory of automata. Since this might be a fruitful way to deepen our understanding of an evolutionary economics along modern information processing lines, which --we think-- should be based upon a computational understanding of markets.

It is not widely appreciated that after the publication of the *Theory of Games and Economic Behavior*, von Neumann did very little further on games, instead devoting most of his prodigious intellectual energies in the last decade of his life to the development of the computer and the theory of automata.[2] This theory of automata was to claim as its subject any information-processing mechanism which exhibited self-regulation in interaction with the

[2] On this seachange, see (Mirowski, forthcoming, chap.3). Von Neumann was very concerned with the question of the extent to which human reasoning could be replaced by mechanisms; as one set of commentators put it, he "thought that science and technology would shift from a past emphasis on the subjects of motion, force, energy and power to a future emphasis on the subjects of communication, organization, programming and control" (Aspray & Burks, 1987, p.365).

environment, and hence resembled the newly-constructed computer. Beyond any pedestrian exploration of parallels, von Neumann envisioned this theory as a province of *logic*, beginning with Shannon's theory of information; it would proceed to encompass the formal theory of computation by basing itself on Alan Turing's theory of the universal calculation machine. Experience had shown that information processors could be constituted from widely varying substrata, all the way from vacuum tubes to the McCulloch-Pitts neuron to mechanical analogue devices. Hence it was the task of a theory of automata to ask: what were the necessary prerequisites, in an abstract sense, for the self-regulation of an abstract information processor? Once that question was answered, the theory would extend Turing's insights into this realm to inquire after the existence of a "universal" constructor of information processors. Biology would make an appearance at this juncture, since the question could be rephrased as asking: under what formal conditions could a universal constructor reconstruct a copy of itself? The logical problems of self-reference initially highlighed by Godel were then brought to the fore. What was it about this system which endowed it with the capacity to resist degradation or "noise" in successive rounds of self-reproduction? Interestingly enough, von Neumann thought the solution to the paradox was a function of the introduction of the irreversible passage of time.[3] Once the conditions for successful repeated reproduction were stated, then the theory of automata would address itself to the theory of evolution. If an automaton could reproduce itself in a second automaton of equal complexity, what, if any, further conditions were required for the same automaton to produce an 'offspring' of greater complexity than itself?

Even from this extremely truncated description, it should be possible to appreciate that von Neumann sought to distill out the formal logic of evolution in a theory of sweeping generality. In this theory, very simple micro-level rule-governed structures interact in mechanical,

---

[3] Von Neumann wrote: "There is one important difference between ordinary logic and the automata which represent it. Time never occurs in logic, but every network or nervous system has a definite time lag between the input signal and output response... it prevents the occurence of various kinds of more or less vicious circles (related to 'non-constructivity,' 'impredicativity' and the like) which represent a major class of dangers in modern logical systems" (in Aspray & Burks, 1987, p.554).

and even possibly random, manners. Out of their interactions arise higher-level regularities generating behaviors more complex than anything observed at the lower micro-levels. The characterization of relative "complexity" is predicated upon the information-processing capacities at each level of the macro-structure. The ability to maintain information transmission relative to noise in a structure of reproduction is a characteristic attribute of an automaton; the appearance of enhanced abilities to operate at higher levels of complexity is the hallmark of evolution. Although intended as a general theory of both living and inanimate entities, most of the formalisms are expressed within the framework of the computer and a theory of computation. Von Neumann justified this dependence upon the computational metaphor because, "of all automata of high complexity, computing machines are the ones we have the best chance of understanding. In the case of computing machines the complications can be very high, and yet they pertain to an object which is primarily mathematical and which we understand better than most natural objects" (1966, p.32). For precisely this reason, von Neumann did *not* believe a theory of automata should be predicated upon the actual organic architecture of our brains. The theory of automata was *not* intended as a surrogate for a theory of human psychology. If anything, von Neumann personally sought a theory of the genesis and maintenance of *organizations*. Furthermore, in contrast to modern orthodox economics, he deemed that his own theory of games played no role in his nascent theory of automata, essentially because it was not sufficiently firmly grounded in logic and in computer architectures.

Herbert Simon's theory of simulacra is also intimately related to the computer, but in a manner orthogonal to that of von Neumann. Beginning as a theorist of administration and management, his experience with computers at RAND led him to develop a general theory of systems which could apply indifferently to individual minds and organizations, and prompted him to become one of the progenitors of the field of Artificial Intelligence. He also found inspiration in the ideas of Alan Turing, but by contrast with von Neumann, he did not ground his theory in mathematical logic, but instead in the idea of the "Turing test". For Turing the problem of defining intelligence threatened to bog down in endless philosophical disputes, and so he proposed to cut the Gordian knot by suggesting that a machine was 'intelligent' if it could not be distinguished from a live human being after a suitable interval of interaction with a qualified interlocutor. Simon follows Turing in asserting that the simulation of a mind or an organization using simple modular and hierarchical protocols is "good enough" for understanding how that entity behaves, if it tracks the behavior of the entity so closely that the simulacrum cannot be distinguished (within some error margin) from the original.

If mere simulation were all there were to his theory of systems, then it would be difficult to credit it with much novelty; but Simon went further in insisting that the very structure of computer programs tells us something very important about the structure of the world. Simon maintains that the modularity of programs in conventional computer architectures[4] mirrors an important fact about our methods of dealing with out own cognitive limitations. If human memory, attention span, and computational capacities are all finite, and control and communication activity between humans is likewise limited, then humans must have developed mechanisms to circumvent these limitations. The primary instrumentality for overcoming these obstacles is to sort and isolate phenomena into decomposable modular structures, which are then reintegrated into a whole through levels of hierarchical interactions. This is the context of the theory of bounded rationality. The resemblance, as he admits, is to the flow chart of a computer program. At the individual level, behavior is not so much substantively as procedurally "rational", breaking down problems into smaller sub-problems which can be attacked using heuristics and rules-of-thumb. At the level of the organization, repetitive problems are dealt with at lower levels, with more vexing unique problems passed up through chains of command and well-specified lines of authority.

Hierarchies of imperfectly decomposable subsystems are for Simon the primary index of "complexity": the more interconnected are the modules, the more complex the phenomenon. Complexity is significant because, given human limitations, it is the major reason certain classes of problems lay beyond our ken. However, it also becomes the bridge to Simon's theory of

---

[4] Which exclude any sorts of newer parallel computational architectures. It is for this reason that Simon is often categorized as a proponent of 'older' AI, characterized by sequential symbol processing, treating thought as "problem solving".

evolution. He maintains (1973; 1981, pp.202) that evolution, or at minimum natural selection, is "speeded up" when organic systems make use of relatively autonomous sub-components. Organisms, like information processors, are asserted to be solving problems; and if those problems are divided up along modular lines, then the organism itself can become a hierarchical structure, dealing with more complex problems than those addressable by its lower-level sub-components.[5] However, this is not intended as an explanation of the actual physiological structure of the brain, much less the morphological layout of the mammal or the organization chart of the M-form corporation, although it is primarily published in psychology journals. Because Simon believes that simulations do provide access to the understanding of complex phenomena, we gain insight into the general problem of information processing by building simulacra of problem-solving exercises.

Although both the "automata" and "simulacra" approaches intersect at numerous points, from dependence upon computer metaphors to algorithmic implementation to shared concern with complexity and evolution, it will be important for our subsequent argument to pay close attention to the critical ways in which they diverge. Von Neumann's automata are squarely based upon the abstract theory of computation for their formal basis, whereas Simon's simulacra usually avoid all reference to the formal theory of computation. Von Neumann regarded computational intractability as a significant component of any theory of automata, whereas Simon appears to believe that heuristic search through hierarchies 'solves' the problem of computational intractability (1981,p.35). Von Neumann did not regard the standard sequential-architecture of his computers as an adequate or even approximate model of the mind; Simon has predicated his entire career on the thesis that computer simulations of psychology are "good enough".[6] Von Neumann tended to stress the processes of interaction between the automaton and

[5] One notable way in which Simon diverges from von Neumann is that the latter saw his theory as explaining how automata could triumph over the second law of thermodynamics, whereas Simon declines to make any statements about the possible relationship between entropy and evolution (1981, p.204).

[6] There is some evidence that Simon pursued the computer/ brain analogy in his early work *because* von Neumann had warned against taking it too seriously. On this incident, see (Sent, forthcoming).

its environment as comprising the logic of its operation, while Simon has tended to blur the environment/organism distinction in his own writings, perhaps because his notion of hierarchy 'internalizes' interactions as representations of problems to be solved within the processor.

Although both von Neumann and Simon have been notorious in their disdain for orthodox neoclassical economic theory over the course of their careers, it is our impression that it has been Simon's simulacra approach which has attracted the bulk of attention and elaboration by economists relative to von Neumann's theory of automata, even among those who find themselves out of sympathy with the neoclassical orthodoxy.[7] While the reasons are undoubtedly many and varied, it might be conjectured that Simon's theory of bounded rationality appeared to hold out the promise that it could potentially be reconciled with the general neoclassical approach to economic theory, especially ignoring his doctrines concerning hierarchies. After all, Simon appears to remain resolutely methodologically individualist, treating market coordination largely as a problem of individual cognition; his simulacra approach resonates with the orthodox position that people act 'as if' they were rational maximizers, eschewing any commitment to actual empirical psychology. Indeed, much modern effort has been expended recasting bounded rationality as itself the outcome of a constrained optimization with scarce cognitive resources.[8] The end product has not resulted in much in the way of distinctly evolutionary propositions or theories, either because of the relatively superficial treatment of evolution in Simon's own writings, or else because of the tendency of neoclassical theorists to access a conception of selection resembling earlier notions of a convergence to a fixed equilibrium or teleological imperative.

We should like to suggest that the options for the possible development of an

---

[7] See, for example (Hodgson, 1993; Egidi, 1991; Conlisk, 1996).

[8] (Conlisk, 1996) surveys this literature. He is commendable in that he points out the paradox of self-reference which bedevils this project: How can someone carry out a constrained maximization to decide that a constrained maximization was not 'worth it'? This is a subset of the larger problem of positing neoclassical 'costs' of information: Who sets the prices of the
price system, and who sets the prices of the prices of...?
The contrast with von Neumann is stark: he directly confronts the logical problem of Godellian self-reference.

evolutionary economics along modern information-processing lines would gain in clarity if Simon's simulacra were juxtaposed with von Neumann's automata. In particular, if Simon fosters a project which involves building little simulated problem solvers who internalize evolution, von Neumann initiated an intellectual project which constructs markets as algorithms which then evolve in the direction of increased computational complexity.

## 2. The Computational Approach to Markets

> "The economic system can be viewed as a gigantic computing machine which tirelessly grinds out the solutions of an unending stream of quantitative problems" (Leontief, 1966, p.237).

While comparisons of markets to computers are thick on the ground in the history of economics, explicit explications of the operation of markets in formal computational terms are much less common. Here we briefly endeavor to lay out the prerequisites of the treatment of markets as automata in the spirit of von Neumann.

### *2.1 What a computational understanding of markets entails*

One of the main effects of the neoclassical school on economic thought was a stress on the desires of the agent to the detriment of consideration of the mechanics of individual market operations. Yet recent developments in a number of seemingly unrelated areas-- experimental economics, finance, incentive structure designs, automated trading devices, law and economics-- have increasingly focused upon the definition of a market as a set of rules which facilitate the conduct of exchange and the conveyance of information between buyers and sellers. This shift in perspective is salutary from the viewpoint of a computational approach, since it permits the reconceptualization of markets as a set of procedures which terminate in a given goal or output.

The paramount goal from the viewpoint of neoclassical economists has been some version of the notion of allocative efficiency. Whereas the Walrasian approach tended to define a state of Pareto optimality for the 'generic' market devoid of procedural specification, research in the area of experimental economics began to raise the question of differential degrees of allocative efficiency in different types of market rule-structures. In particular, (Smith, 1991) has claimed that that the Double Auction market (DA) dominates most other market formats (such as, say, the sealed-bid auction) in producing high allocative efficiency in controlled experiments. Abstaining for the moment from accepting Smith's particular characterization of the goal or end-state of the market, and postponing a detailed description of the DA market, his work does propose the valuable idea of *ranking different market procedures according to their relative success in achieving a particular end-state*. This realization has been carried further in the work of Gode & Sunder (1993,1997), and for that reason, we shall opt to translate their work into more explicit computational terms in section 3 below. The valuable insight of Gode & Sunder is that it is possible to produce a ranking of differential market procedures by abstracting away from the cognitive capacities of the market participants, at least as a first approximation. Where we shall diverge from Gode & Sunder is that we show that the categorical arrangement of market procedures (or, perhaps better, 'institutions') in some hierarchy can be generalized for any given goal or end-state, and that the principle of categorization is provided by computational theory.

Hence, in the computational approach, particular market institutions possess a certain computational capacity independent of the computational capacity of the agents participating in the market, and this capacity can be deduced from an enumeration of the rules that constitute the specific market. Interesting examples of this approach can be found in Miller (1986). He employs first order logic in order to give a mechanical description of a DA market. Given an input of type of agent (buyer, seller), type of action (bid, ask) and quantitative proposed price, the set of rules generates an outcome, namely, an allocation of goods. Miller demonstrates that modelling of the sorts of optimality conditions favored by neoclassical economics goes well beyond this mechanical specification, requiring, for instance, the use of second-order logic.

Thus, there exists substantial precedent for attributing a certain computational capacity to types of markets predicated upon the set of market rules that describe a repetitive procedure. In the theory of computation, a procedure which contains a finite set of instructions or rules is called an algorithm. An algorithm may be described as a finite sequence of instructions, precisely expressed, that -- when confronted with a question of some kind and carried out in the most literal-minded way-- will invariably terminate, sooner or later, with the correct answer (Lewis &

Papadimitriou, 1981, p.36). We would argue that the authors cited above, and indeed many other economists, are conceptualizing markets as algorithms without being fully aware of the implications of that activity. In particular, the notion of market as algorithm is entirely separable from whatever one conceives as the *purpose* of the market. For Smith, the output of the algorithm is a particular proportional realization of his definition of pre-existent consumers' surplus. For others, as in the incentive compatibility literature, it may be some notion of allocative efficiency conditional upon precise specification of agent preferences. Simpler end-state conditions might alternatively be posited, such as the clearing of a market in a certain time frame, or the output of a set of prices obeying the no-arbitrage condition.

If the individual market rules meet the criteria for an algorithm, then this constitutes the formal content of the widespread impression that the market system is a giant computer. The central lesson for an evolutionary economics is that multiple criteria for end-states justify the existence of multiple types of market/algorithms, and that these can be arrayed along a continuum of computational capacities.

*2.2 Hierarchies and classes of automata*

Automata theory is a familiar framework within the larger theory of computation, permitting a more formal specification of the informal notion of an algorithm. There exist a number of introductions to the theory of automata and computation (Lewis & Papadimitriou, 1981; Davis et al, 1994); they are the sources for our brief summary characterization below. We shall define an automaton as a restricted and abstract formal model of a computer. Algorithms can be processed on abstract automata of various capacities and configurations.

The standard theory of computation proposes a hierarchy of automata of various capacities (in increasing order) to handle strings of inputs: finite automata, pushdown automata, and Turing machines. All automata possess a finite set of internal states (including an initial state), a well-defined finite alphabet, an input (and) output device, and a transition function which carries the automaton from one state to its successor. A deterministic automaton has only one successor state for each given active state, whereas a nondeterministic automaton may have more than one successor state. The primary array of relative computational capacities of the hierarchy of automata are determined by the amount and kind of memory to which the machine has access. A Turing machine has no restriction on its memory, in the sense that there always exists the opportunity to expand its capacity. Pushdown automata also have unlimited memory, but is restricted to the process of stacking data -- last in first out -- and finite automata lack any storage device. Turing machines occupy the pinnacle of the computational hierarchy because, ignoring for the moment issues of efficiency in computation, a Turing machine can simulate the operation of any other machine, and therefore, in an abstract sense, all Turing machines are equivalent in computational capacity. "Church's Thesis" states that because Turing machines can carry out any computation that can be successfully prosecuted on any other type of automata, the Turing machine captures the intuitive content of the notion of an algorithm.

*2.3 How Automata diverge from conventional economics*

Because of the preoccupation of orthodox economics with the characterization of the market as a manifestation of what some agent or agents *think* about it, it may require some stretch of the imagination to realize that the evolutionary automata approach (at least initially) pushes the cognitive states of the agents to the margins and focuses upon the mechanics of the processing of messages. In this framework, a market accepts well-formed sentences as inputs (orders, bids, asks, reservation schedules), grinds through the set of states implied by those messages, and produces output messages (realized prices, quantities). These messages may correspond to actions or activities (conveyance of goods, payments of currency, assumption of debts, etc.), but then again, *they may not*. One might regard this analytical separation as a projection of the standard distinction between syntax and semantics; but it proves much more far-reaching in this framework. Critical assumptions about the topology of commodity space, the 'smoothness' of preferences, independence of irrelevant alternatives, and all the rest of the standard armamentarium of the mathematical economist play no role here. A well-known result in computational theory suggests that an abstract Turing machine can be realized in a myriad of underlying physical configurations of automata; thus the physical attributes of economic activity (such as technologies or 'endowments') can be readily placed in a different class of theoretical abstractions, those effectively removed from general considerations of market operation and efficiency, but of course relevant to specific historical circumstances and social structures. Instead of the orthodox habit of imagining an atemporal generic phenomenon called a "market" (or "human rationality") fully and equally present in all of human history, the automata approach posits an abstract category of information processor which then evolves into variant formats of *plural* markets depending upon local circumstances and some generic notions of *computational* complexity and efficiency.

### *2.4 Some Immediate Implications of the Automata Approach for an Evolutionary Economics*

Already at this very early stage, the theory of evolutionary automata bears very specific economic content. Turing machines may (ideally) possess infinite "tape" or memory, but they are restricted to a finite number of internal states and a finite alphabet. The motivation behind this inflexible requirement is that we are enjoined to adhere to a "constructivist" approach to mathematics, showing how an answer is arrived at deploying prosaic sets of rules without appeal to intuition or insight.[9] For technical reasons, it follows that our machine must restrict itself to a discrete alphabet, or when calculating, restrict its operations to the set of countable numbers (e.g., natural numbers). Far from being a nuisance, this restriction embodies an empirical generation about the history of markets: prices have always and invariably been expressed as rational numbers (i.e., ratios of natural numbers), and further, they have been denominated in monetary units which are discrete and possess an arbitrary lower bound to the smallest possible negotiable value. The appeal of the mathematical economist to the real orthant has no correspondence to economic history. This empirical regularity is not an artifact of "convenience" or some notion of costs of calculation; it is a direct consequence of the algorithmic character of markets.[10]

---

[9] The constructivist approach to mathematics is best contrasted with the "Bourbakist" approach which has dominated work on Walrasian general equilibrium in the postwar period. For a nice discussion of this distinction, see (Velupillai, 1996); the history of the Bourbakist incursion into economics is described in (Weintraub & Mirowski, 199x).

[10] One could also pursue a similar inquiry into the interplay of market evolution and the invention and definition of quantitative commodity identities along the lines of (Kula, 1986), an inquiry we must bypass here. Nevertheless, the requirement that one specify the algebra over which computations are performed dictates an

A salutary influence of the computational approach to markets is that it forces the analyst to be much more precise in specification of *how* the market operates, by demanding the enumeration of the sequence of steps that carries the automaton from its initial state to the final output of a sequence of prices and quantities. It is no longer acceptable to build a model simply to identify a supposed equilibrium configuration, leaving aside the question of the 'dynamics' of putative convergence until a later exercise. Indeed, it was only with the dual pressures of the automation of real-time markets, combined with the need to algorithmically specify the details of the computerized experimental protocols in the promulgation and standardization of the nascent experiemental economics, that the economics profession was induced to confront the fact that there exist plural structures of market institutions, and that differences in configurations might lead to differential price-quantity outcomes. Although there has been a pronounced tendency to focus attention upon the "double auction" market (one of which we shall be equally guilty), probably due to its resonance with certain neoclassical models, it has now become commonplace to admit that microeconomics should provide a taxonomy of market forms, along the lines of that in Figure I ((Friedman & Rust, 1993, p.8).

---

intrindically *monetary* theory of markets, in sharp contrast with the real/nominal dichotomy of neoclassical microeconomic theory.

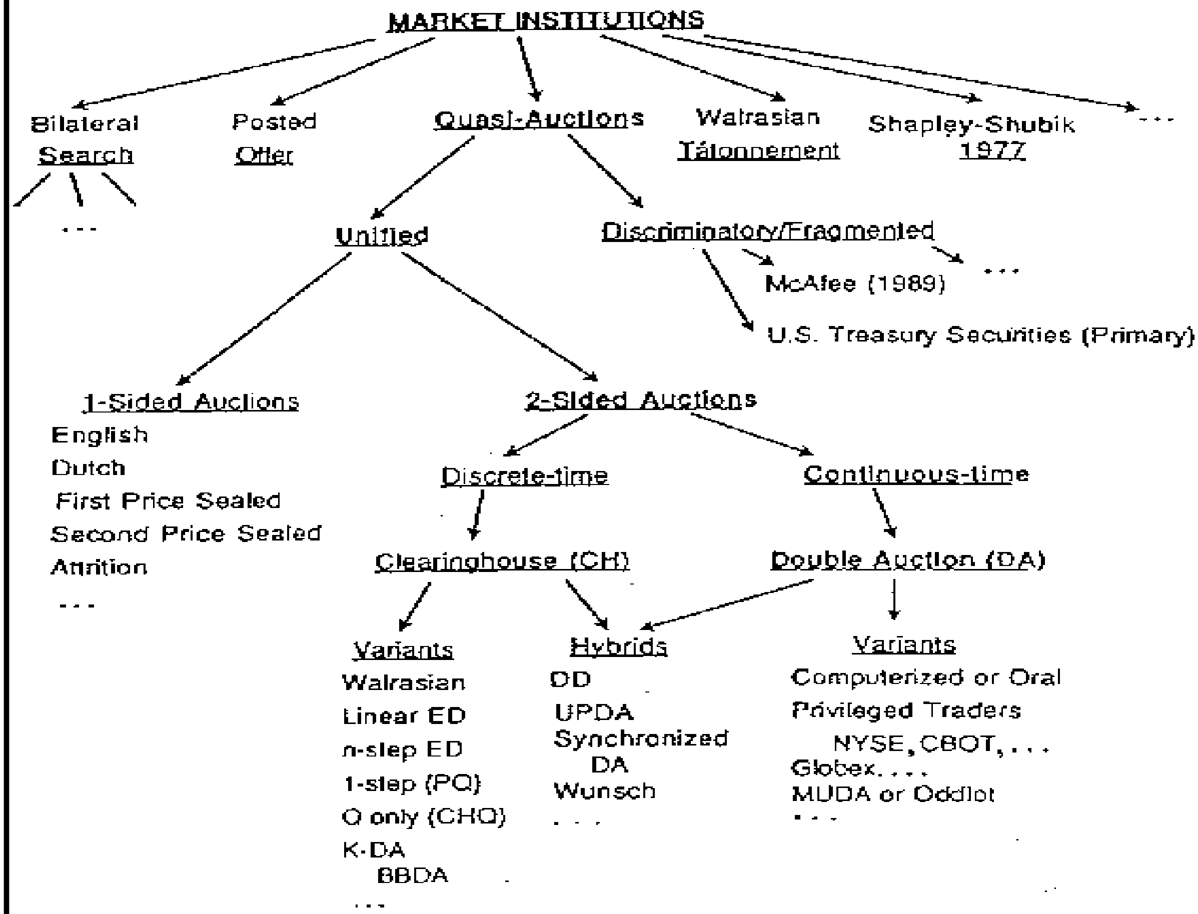


FIGURE 1 A Family Tree for the Double Auction.

It is noteworthy that the diagram in Figure I resembles a phylogenetic tree, a device commonly used to demonstrate descent with modification in evolutionary biology. We do not claim that this particular diagram captures any such evolutionary phylogeny-- indeed, orthodox economics possesses no means by which to judge whether one abstract form could or could not descend from another, much less the historical curiosity to inquire whether it actually happened.

Our major thesis is that the automata approach does supply the wherewithal to prosecute this inquiry. Once the algorithm which characterizes a particular market format is identified and represented as a specific automata, then it becomes possible to bring von Neumann's project back into economics. First, from examination of the algorithm, one would enquire whether and under what conditions the algorithm halts. This would include questions about the conditions under which the algorithm is regarded as arriving at the "correct" answer. Is the desideratum of the algorithm to "clear" the market in a certain time frame? Or is it simply to provide a "public order book" in which information about outstanding bids and offers is freely and accurately available to all? Or, alternatively, is it to produce a simple quantifiable halt condition, such as the absence of arbitrage possibilities within a specific time frame? Or is it constructed to meet a certain specific welfare criterion? It is of paramount importance to keep in mind that the objectives to be attained by market automata are multiple, some potentially complementary and some in conflict. The mix of objectives is geographically and temporally variable: the first prerequisite of an evolutionary situation.

Second, the analyst would rate the basic computational capacity of the specific market format relative to the previously identified objective. Is a simple finite automata, or something more powerful, approaching the capacity of a Turing machine? If it qualifies as the latter, can it then be arrayed in order of the complexity of the inputs it is expected to handle? Additionally, one could compare and contrast automata of the same complexity class by invoking standard measures of time complexity in the theory of computation-- can the "worst case" computation be carried out in polynomial time (Garey & Johnson, 1979)? Given a suitable complexity index, one would then proceed to tackle von Neumann's question, namely, under what set of conditions could a market of a posited level of complexity give rise to another market form of equal or greater complexity? In what formal sense is market evolution possible?

*2.5 The Problem of Reproduction*

It may be that it is here, at the idea of a specific market structure "giving rise" to another, that economic intuition may falter. What could it mean for a market automaton to "reproduce"? Here is where the abstract computational approach comes to dominate an "embodied" conception of markets. Market institutions spread in an extensive or embodied manner by simple replication of their rules, say, at another geographic location. This does not qualify as von Neumann reproduction, since it was not the market algorithm that produced the copy of itself. Market automata "reproduce" when they can imitate the abstract operation of other markets as a part of their own algorithm, incorporating a simulation of the operation of the specific market format into their own, different market format.

The simplest example of this "universal" capacity is markets for derivatives. When agents trade in the futures market for grain contracts, they are attempting to simulate the outputs of a different market, namely, that of the spot market for actual grain. It is of special importance to recognize that the spot market (say, an English auction) can and frequently does operate according to distinctly divergent algorithms than does the futures market (say, a double auction); hence one automaton is "reproducing" an altogether different automaton. It is the appearance of this "self-referential" aspect of automata that creates the possibility of a hierarchy of computational complexity, and hence "evolution" in the von Neumann sense. Market forms may "spread" relative to one another (say, fixed-price storefronts replace itinerant haggling peddlers), and as such may be subject to a particularly crude kind of "selection"; but since nothing profoundly novel arises from this process, there is no phylogeny and no evolution as such. The system only displays a phylogeny and therefore a distinct arrow of time when market automata begin to emulate the operation of other such automata, resulting in calculations of ever-higher complexity.

In this evolutionary automata approach, many economic terms undergo profound redefinition. Markets are no longer environments within which agents operate; it is within the ecosystem of the multiform diversity of agents and cultures in which markets calculate and evolve. Since the first prerequisite of an evolutionary theory of 'natural selection' is that the entity which encapsulates the inheritance principle displays greater inertia than the surrounding environment, for the first time economists may escape the Lamarckian indictment which has bedeviled the efforts of theorists like Nelson and Winter (1982). Moreover, "market failure" no longer indicates some disjunction between an imaginary optimum of utility and the equilibrium identified by the analyst; it now refers to a real and easily identifiable phenomenon, that where a market algorithm fails to halt within the parameters of operation of the specific automaton. Since it is a theorem of the theory of computation that there is no general algorithm for deciding whether or not a Turing machine will halt, market failure is understood to be an endemic and inescapable fact of life. Examples of such phenomena are market "crashes", price free-falls, markets unable to conduct arbitrage operations, markets incapable of conveying order information to other markets, cascading shortfalls and macroeconomic contractions. "Efficiency" likewise becomes decoupled from any prior specification of the desires of the individuals involved. An 'efficient market' now becomes an automaton that can handle a wide diversity of messages emanating from people with differing beliefs, desires, cognitive skills and cultural backgrounds as inputs, and produce price and quantity outputs meeting fixed prior desiderata (market clearing, arbitrage, etc.) in finite, and preferably polynomial, time. Furthermore, the age-old theme of the benefits of the division of labor, so vanishingly present within the Walrasian tradition, enjoys a revival within the computational tradition. There need be no presumption that market automata are restricted to be sequential symbol processing devices; as markets necessarily become more interconnected, accepting as inputs the price and quantity outputs of other market automata, it becomes helpful to regard clusters of market automata as constituting a distributed processing device, mimicking connectionist architectures (Barbosa, 1993). This concatenation of multiple processors at a lower level into a novel entity at a macro level may mirror the transitions between levels of organization found in the biological record (Smith & Szathmary, 1995).

3. **A Concrete Example of the Institutional Automata Approach**.

We have argued that the evolutionary automata approach (at least initially) should relegate the cognitive states of the agents to a subordinate status and focus upon the mechanics of the processing of messages and information. Elements of this research project can already be found in the economic literature (e.g., Miller (1996), DeVany & Walls(1996), Cason & Friedman(1997), and Gode & Sunder(1993;1997)). To give some idea of the extent to which this literature can be synthesized and expanded within the automata approach suggested in the previous section, we here opt to translate the work of Gode&Sunder (henceforth G&S)-- specifically, parts of their 1997 paper-- into our automata framework. By so doing, we do not intend to endorse their claim that Marshallian consumer surplus is the single correct index of allocative efficiency, nor do we necessarily agree with their assertions concerning discovery of the attributes of success true for all markets, whatever their structure. Our objective is rather to provide an illustration of how one goes about formulating an abstract market structure as an automaton, and to reveal (in a manner they did not) that there exists a computational hierarchy along which different markets can be ranked, as a prelude to full prosecution of von Neumann's program of formalizing evolution as the temporal unfolding of increased complexity.

We can identify at least three aspects of the work of G&S that render it suitable to be encompassed by our suggested evolutionary automata approach.

[1]The analytical distinction between market (automata) and the agents' cognitive skills is an issue already rendered salient by the experimental economics literature. This distinction was initiated by the work of Smith (1991), which reports strikingly high allocative efficiency of experiments conducted in a DA market. G&S (1993) resorted to simulated Zero Intelligence (ZI) traders submitting bids and asks in a computerized market setting explicitly to "zero out" all cognitive considerations and highlight the primary determinants of this standard finding of high allocative efficiency in the DA market. They show that almost all of this efficiency is due to the algorithmic aspects of the market alone, i.e., the market closing rules and the restriction that the bid/asks are only effective when they conform to a budget constraint.11 Thus they make explicit

11 It is not completely obvious that the latter is part of the market automata under consideration. But what is important is that it is not necessarily part of the agents cognitive skill.

the key analytical distinction between market automata and environment (i.e., agents' cognitive skills) which we claim is central to the automata approach and permits our distinctive conception of evolution.

[2] The 1993 paper also shows that ZI simulations with budget constraints mimic the results of human experiments fairly closely. This suggests the idea that certain market formats are relatively robust in performance relative to different environments consisting of people of diverse cognitive skills. This idea echoes our discussion (in sections 2.3-4) of different markets possessing different levels of computational capacity arrayed in order of their robustness (here provisionally defined as stability of outputs in different environments) as a first step towards the introduction of von Neumann's project into economics.

[3] The 1997 paper begins to deal with a hierarchy of different market rules in an explicit ZI environment, with the idea of quantifying incremental improvements in the efficiency measure and attributing them to specific rules. This work takes the largest strides toward treating different markets as different types of automata, although G&S do not motivate it in this manner. It also begins to entertain more explicitly the idea of arraying the different markets along some continuum of complexity, and testing that continuum against a controlled environment, thus opening up the possibility of a von Neumann-style definition of evolution.

Our exercise assumes the following format. G&S provide models of three classes of market rules: their so-called "Null Market" (random matching of good to buyer); a simple first-price sealed bid auction; and a synchronized double auction (DA). Real markets are much more complicated than such sparse characterizations, so the abstractions of G&S provide us with a convenient opportunity to encode this much simpler set of rules as automata. We show that their definitions of sealed-bid and DA fall into different classes of computational capacity, viz., different types of automata. Our initial contribution to an evolutionary approach is to show that a DA can mimic a sealed-bid auction but that the converse is not true. This intransitivity has not previously been a topic of commentary in the economics literature, and can only become obvious when both markets are redefined in an automata framework. The DA requires extra memory capacity that makes it impossible for the sealed bid market to mimic the DA market. This suggests that, for certain given objectives, the DA is formally a "descendant" of the sealed bid market, and can therefore handle more complex information processing tasks. This formal result can then be brought to bear in the future on a "natural history" of markets, to research the possible ways in which these particular market forms have evolved relative to one another.

3.1 *The work of G&S.*

In a sequence of papers G&S (1993;1997) address the question why the experimental economics literature reports high allocative efficiency for DA markets. Allocative efficiency-- the ratio of the actual to the potential gains from trade-- is one of the criteria used by the experimental literature to analyze the working of markets in general, and the DA in particular. Allocative efficiency in this definition is high if the consumers who value a good the most manage to buy their units from the lowest cost producers.

G&S (1993) report a market experiment in which human traders are completely replaced by "zero-intelligence" (ZI) agents. These ZIs are in practice simulation programs that randomly choose their bids and asks subject to market rules. The ZIs generate an allocative efficiency close to 100% once their bids and asks are forced to conform to their budget constraints. Thus, by introducing ZIs, G&S are able to "zero out" all cognitive considerations and highlight the primary causes of the standard finding of experimental economists that the DA market achieves high level of allocative efficiency. In G&S (1997) they further elaborate upon this striking result. This paper shows that starting with a base situation of almost no market rules (their Null Market) and then recursively adding procedural rules to the simulation will improve its expected allocative efficiency. The trading environment remains specified as ZI traders that choose randomly subject to market rules, whereas the market rules determine the bid/ask range within which traders choose subject to a uniform probability distribution. They show that recursively adding market rules will further limit the possible price range in which the good can be traded by extramarginal traders, consequently increasing the expected allocative efficiency. The work of G&S demonstrates that almost all of the observed allocative efficiency is due to the algorithmic aspects of the market, independent of cognitive considerations.

The work of G&S (1997) therefore seems to justify the conclusion that much of the computational capacity required to conduct trade is incorporated in certain repetitive market rules, or in our language, the market itself can be interpreted as an algorithm. To render this assertion more precise we will use several of the rules analyzed by G&S to provide a sketch of the automaton to which the markets--that are constituted by these rules--can be encoded.

Note that throughout this analysis we will make the simplifying assumption that in every trade round only one unit is under consideration by the agents. This is an artifact of the translation of G&S's exercise. Additionally, bids and asks are encoded in a "convenient" way, which for our purpose is an unary notation (e.g., bid=10 is represented as a string of 10 ones denoted as I). Thus, for purposes of illustration, we limit the construction of the machine to the encoding of the market rules. This assumption is justified by the fact that we are only interested in comparing the computational complexity of markets constituted by certain commensurate rules, i.e., all markets are assumed to have the same auxiliary capacity to encoded bids and ask in an unary input string.

3.2 *Null Market Encoded on a Finite Automaton*.

The base case of our analysis will be the Null market described in Gode & Sunder (1997, p.612); this market is deficient of rules. The machine that represents this market is exceedingly trivial, something which the name already seems to imply. The only thing that the machine needs to do is recognize if both a bid and ask are submitted; after this condition is satisfied, it will arbitrarily halt or read the next available input. This requires so little computational capacity that this market can be modeled by a nondeterministic finite automata. A nondeterministic finite automaton is a quintuple *M=(Q, Σ, Δ_, Ä,* $q_0$*, F)*, where *Q* is a finite set of which its elements are called states; $q_0$ is a distinguished element of *Q* called the start state; _ denotes an alphabet; *FfQ* is the set of final states, and Ä, the transition relation, is a finite subset of Qx_$^*$*xQ*. (_$^*$ is the set of all strings--including the empty string--over _.) The rules according to which the automaton *M* picks its next state are encoded into the transition relation. For example, if *(q,a,q=)*0Ä, *M* may read of *a* and move to state *q=*. Note, that this automata is called nondeterministic because Ä does not (necessarily) uniquely determines the machine=s next configuration. Given that a machine is in state *q*, Ä can have, for example the instructions: read of an arbitrary number of *a*=s, *a*0_, and go to state *q*= or do not read of any element of the input string and go to state *q*==. The essential feature of finite automata is that it is a language recognizer. It reads of an input string, ó0_* and if after reading of the whole string it reaches a final state q0F then the input sting is considered to be accepted. The language accepted by the machine is the set of strings it accepts.

We can encode the Null market onto a finite automaton $M_0$ by using only four states where $Q=\{ q_0, q_1, q_2, q_3\}$, _={a, b, I}, $F=\{q_3\}$ and the essential transitions in Ä are ($q_0$, *a*, $q_1$), ($q_1$,*b*, $q_3$), ($q_0$, *b*, $q_2$), and ($q_2$, *a*, $q_3$). We adopt the convention that the symbol A*a*@ proceeds the unary input if it is an ask and A*b*@ if it is a bid. Now, $M_0$ constitutes an and-gate that ensures, that at least one bid and offer are submitted before the machine halts. In constructing the automaton we ignored details highlighted by G&S, such as what mechanism determines the probability that a transaction takes place. By extending its design, and introducing probabilistic matrices that encode the probability of a certain transition-taking place, we can simulate the interaction of Null market. To do so we need to add an extra element to the machine such that we have the sextuple $M_0$==(Q, _, Ä, $q_0$, *F*, *{A(x)}*), where only *A(x)* differs from the original machine $M_0$. *A(x)* is a finite set containing $|\Sigma|$ square stochastic matrices of the order $|Q|$ where $x \in \Sigma$ and a matrix *A(x)* is stochastic if all the entries of its row vectors are greater than or equal to 0 and sum up to 1. $M_0$= is a probabilistic finite automaton. We first introduced $M_0$ in order to draw an explicit distinction between an algorithmic analysis of market rules and the more conventional approach that begins with behavior without specification of context.

3.3 *A Sealed Bid Auction Encoded on a Two-Tape Pushdown Automaton*.

In the previous subsection we completed our sketch of the automaton of the Null market, $M_0$. Before continuing we want to emphasis that automaton $M_0$ has very little computational capacity as its only real computation-- with respect to calculating a market price-- is the and-gate which ensures that at least one bid and ask are submitted before a final state is reached. The first thing we need to introduce in order to encode a market with a more substantial computational capacity on a machine *M* is to introduce a more sophisticated halting condition. A halting condition gives the circumstances under which a machine *M* reaches a final state and ceases computation. In the previous case, the halting condition was that at least one bid and ask need to be submitted. Henceforth, the condition for a machine to halt will be that a submitted bid is greater than or equal to a submitted ask. This apparently simple condition already requires a machine substantially more powerful than a finite automata. The reason for this is that the machine needs to be able to accept the following types of string,

$L=\{I^m b I^n a : m \$ n\}$,

where *a,b,I*θ_ and $I^m b$ encodes a bid of magnitude *m* in denary notation (e.g., $I^m b$ where *m*=10 gives a bid of $10). Thus *L* defines the language that *M* needs to accept. First note that we do not put any a priori restrictions on the size of a bid and ask, this implies that *L* is an infinite language. Therefore we can show that *L* cannot be accepted by a finite automaton, using for instance one of the *pumping theorems* for Finite Automata (Lewis & Papadimitriou 1981, chapter 2). The intuition here is that, since finite automata have no storage capacity, they can only handle strings of arbitrary length that have very limited complexity.

In order to encode this halting condition onto a machine we need to introduce appropriate storage capacity. This can be done by restricting the data structure to a stack--last in first out--i.e., we can use a pushdown automaton to encode the halting condition. A pushdown automaton is a quintuple *M=(Q, _, _,* Ä, $q_0$*, F)*, where *Q* is a finite set of states, $q_0$θ*Q* is the start state, _ is an alphabet (the input symbols),_ is an alphabet (the stack symbols), *FfQ* is the set of final states, and Ä, the transition relation, is a finite subset of (Qx_*x_*)x(Qx_*). Intuitively, if ((q,a,ù),(q=,è))θ Ä then whenever *M* is in state *q* with ù on top of the stack it may read *a* from the input tape replace ù by è on the top of the stack and enter state *q=*. The essence of *M* is that it stores an encoded ask, and if a bit is submitted starts crossing out I=s. If the stack is emptied before the complete bid string is read it will reach a halting state *q*θ*F* and halts (after reading the remaining element of *b*).

We will now add the following two rules to the halting condition (taken from Gode &

Sunder (1997, p.610)):

1. *Binding contract rule: bids and asks are binding, i.e., buyers must pay what they bid; and sellers must sell at what they ask.*

2. *Price priority rule: higher bids dominate lower bids, and lower asks dominate higher asks*.

Rule 1 enjoins the market participant to bid in accordance with their budget constraint, and rule 2 makes it more likely that buyers with the highest redemption value and sellers with the lowest cost will trade. These two rules, in combination with the halting condition, can give rise to a Sealed-Bid market.12 Two essential features of this institution are: the one-sidedness of the auction, dictating that for a given fixed supply only bids are submitted; and the restriction that traders have no knowledge of or influence upon other bids being submitted.

For simplicity, we will consider a single-unit sealed-bid auction (only one unit is offered). This institution can be encoded on a pushdown automaton, where only the *minimum ask*-- the price below which the supplier is not willing to trade--needs to be stored in the automaton. The automaton will halt as soon as a bid exceeds this minimum price. Given this very simple structure of the auction rule 2 is automatically satisfied by the halting condition, because the machine halts as soon as a bid exceeds the minimum ask. It is not equally straightforward if the binding contract rule (rule 1) should also be encoded on the market automaton. A possibility might exist that traders are enjoined to keep a deposit accessible by the auction mechanism so that the credibility of every bid can be checked. However, it is more likely that rule 1 is enforced outside of the immediate framework of the auction, by either the law, or some other separate institution such as the banking system, gilds of individual traders, credit-rating services, or any combination of these structures. Henceforth, we will assume that rule 1 is enforced outside the market institution under consideration. Hence the most important aspect of a pushdown automaton $M_1$-- the aspect that gives rise to a single-unit sealed-bid auction-- is the halting condition.

---

12 Note that above rules in combination with that halting condition ensures that only voluntary trade will occur, hence satisfying all three conditions for a Sealed Bid market stated by Gode&Sunder.

Appendix I presents the essential features of a pushdown automaton that accepts the language $L=\{I^m b I^n a:m\$n\}$. That is the halting condition-- a bid has to exceed an ask before the machine reaches a final state-- is encoded on this machine. Henceforth we will call this automaton the halting machine, $M_H$. This halting machine can be best portrayed as a subroutine of the machine $M_1$. Additionally, $M_1$, needs to have an extra stack to store the minimum ask (below which the supplier is not willing to trade). Every time a new bid is placed, this minimum ask is copied to (the stack of) $M_H$, after which $M_H$ determines if the most recent bid satisfies the halting condition. If this is the case $M_1$ halts and a trade takes place. Thus the single-unit sealed-bid auction can be encoded on a pushdown automaton $M_1$ that has two stacks of arbitrary size. (See appendix II for the pseudo-code.)

Note that $M_1$ is a machine with minimum computational complexity upon which such an auction can be encoded. The computational capacity of a pushdown automaton with two stacks is larger in magnitude than an automaton with one stack. This is due to the fact that elements stored earlier have to be deleted before elements that are stored later are accessible. In other words, to read an ask in order to compare it with a bid, it is necessary for machine $M_H$ to delete the stored ask while checking if the halting condition is satisfied. Therefore $M_1$ needs an additional stack in order to compute the subroutine $M_H$ an arbitrary number of times.

The complexity of the sealed bid auction discussed above can be increased by adding one more rule that needs to be encoded on the machine (Gode & Sunder 1997, p.610):

3. *Accumulation rule: the highest bid (and the lowest ask if it is a double auction) are chosen only after all bids (and asks) have been collected*.

Rule 3 ensures that the market no longer automatically clears as soon as a bid exceeds a ask. Now it will be the case that, under a certain regime, all bids (and asks if it is a double auction) submitted within a certain time frame are collected before the market clears. To encode the machine $M_2$ with the additional rule 3 requires the machine--in addition to the functions it inherit from $M_1$ --to store the highest bid that has been registered of the input string so far. Additionally, the machine now needs to keep track of who submitted the stored bid and ask. Implicitly, the machine needs to keep track of the order in which bids and ask are scanned. Although, $M_2$ is a more complex machine than $M_1$ it can still be encoded on a two tape pushdown automaton (see appendix II). We can be more precise with our notion of complexity by arguing in this instance that in the worst case scenario, it will take $M_2$ longer to halt than $M_1$.

3.4 *The Double Auction Encoded on a Three-Tape Pushdown Automaton.*

The last rule we will introduce in this paper is:

4. *Double Auction rule: buyers can bid as well as sellers can ask.*

Rules 1 to 4 give rise to a rudimentary double auction (DA). To encoded the DA on a machine $M_3$ it is necessary to deploy an automaton of a different computational complexity class. It is impossible for the DA to be encoded on a two-tape pushdown automaton that stores both the highest bid and ask and simultaneously keeps track of the number of bids and ask that preceded a stored bid or ask (in order to know who submitted the bid or ask).13 For simplicity, we will still assume that only one good per trading round is traded/offered, but now both bids and asks can be submitted. $M_3$ works in an almost identical fashion to $M_2$; only now an extra stack (stack 2) is needed in order to store ask and bids separately. $M_H$ and $M_1$ are again subroutines; only now, an additional subroutine $M'_1$ is introduced that checks if a submitted ask is below the currently stored ask (and $M_1$ checks if an bid is above the stored bid). $M_1$ operates on stack 1 and $M'_1$ operates on stack 2. Submissions are ordered so that the lowest ask until that point is stored on stack 2 and the highest bid is stored on stack 1. (See appendix III for further details.)

---

13 It is possible to encode a DA on a two-tape push-down automaton that only keeps track of the number of bids that preceded a stored bid and the number of asks that preceded the stored ask (instead of the total number of bids and asks that precede a stored bid/ask). This requires more external computational capacity in order to determine the identity of those who submitted a winning bid and ask than, for instance, the sealed-bid market encoded on $M_2$, making a commensurate comparison between the two essentially impossible.

3.5 *Concluding Remarks*.

For the purpose of illustration of our discussion of a von Neumann-style definition of evolution as a progression through increasing degrees of complexity, we accessed the definitions by Gode & Sunder of a Null market baseline, a sealed-bid market and a DA market. We demonstrated that all three "markets" can be characterized by automata of differing computational complexities. The Null baseline can be described as a simple finite automata; the sealed bid market can be encoded on a two tape push down automata; and the double auction market needs at least the computational capacity of a three tape pushdown automata, primarily because both sides of the market are active. As a direct corollary of this result, it follows that it is (in general) impossible for the sealed-bid market to mimic the operation of the DA, but the converse is not true since a sealed-bid market can be encoded as a subroutine of the DA.

The implications of this exercise for an evolutionary economics are immediate and striking. Gode & Sunder have argued that these three market formats underwrite a progressively greater degree of allocative efficiency, admittedly gauged by their single criterion. We have demonstrated that these three formats display an increasing degree of computational complexity when viewed as automata, independent of the cognitive capacities of the market participants. Thus, given some specific goals or criteria for success, it is now possible to provide a formal characterization of the hierarchy of the diversity of market forms in terms of their complexity, defined relative to that specific goal. Further, "reproduction" is now given an unambiguous interpretation as one market format mimicking the operation of another. With the recognition of multiple goals and their attendant complexity hierarchies, *for the first time* there exists the outlines of a formal economic model of the modern conception of evolution as a dynamic selection of information processors which is not itself a metaphorical projection of the attributes of biological entities. It is a revival of the project of von Neumann, not that of Darwin. It is an "evolutionary economics" where the stress is on the noun, not the adjective.

In contrast to previous competing versions of evolutionary economics, the computational approach has one historical trend in its favor, which suggests that it will eventually transcend mere academic interest. Already, automated markets and artificial agents are playing an ever-increasing role in real-world economic transactions (Anon, 1997; Miller, 1996). Specialists are employed today to program the automata we have described. Market participants will not have to stretch their imaginations to conceptualize the automata approach to economics, for they will increasingly find it all around themselves in their everyday activities.

**Appendix I:**

*The halting condition.*

The halting condition that a bid needs to exceeds an ask can be captured by the language $L=\{I^m b I^n a : m \geq n\}$.

This language can be accepted by a pushdown automaton $M_H=(Q, \Sigma, \Gamma, \Delta, q_0, F)$, where *Q is a finite set of states,* $q_0 \in Q$ *is the start state,* $\Sigma$ *is an alphabet (the input symbols),* $\Gamma$ *is an alphabet (the stack symbols),* $F \subseteq Q$ *is the set of final states, and ?, the transition relation, is a finite subset of* $(Q x \Sigma^* x \Gamma^*) x (Q x \Gamma^*)$.

The essential transitions of a (nondetermenistic) pushdown automaton that accepts this language are

1. $((q_0,a,e),(q_1,a))$ push (add) *a* on top of the stack
2. $((q_1,I,e),(q_1,I))$ push (add) *I* on top of the stack
3. $((q_1,b,e),(q_2,e))$ switch states
4. $((q_2,I,I),(q_2,e))$ pop (replace) I from the stack
5. $((q_2,I,\#),(q_f,e))$ reach final state $q_f$, since bid>ask,

where all 5 transitions are an element of $\Delta$, the transition relation. Transition 1 recognizes an ask and puts the machine in state $q_1$ the state in which an ask is stored. The storage is completed by transition 2, which has the ones, I, stored. (Remember that *I...I* denotes the size of bids and asks in unary notation.) Transition 4 checks if a bid exceeds the stored ask by crossing out ones and if a blank symbol--denoted by #--is reach first on the stack then the bid is bigger then the ask, hence the machine reaches the final state $q_f$ (transition 5).

**Appendix 2.**

The encoding of the one-unit sealed-bid market.

Machine $M_1$ gives rise to the sealed-bid market with only the price priority rule. $M_1$ is a pushdown automaton with two stacks, 0 and 1, respectively and $M_H$ as a subroutine. The pseudo-code of $M_1$ is the following.

1. *Copy next input string on input tape (i.e., minimum ask) to stack 1.*

2. *While* $M_H$ *has not reached the halting state* ***do***

*3.* *Have* $M_H$ *operate on the next string (bid) of the input tape.*

4. *end*

3. ***If*** $M_H$ *halts in final state* ***then*** $M_1$ *will also halt.*

Thus $M_1$ halts as soon as subroutine $M_H$ detecting a bid that exceeds the minimum ask. The last bid given to $M_1$ gives the price for which the good offered will be traded. This bid is stored on stack 1 and the order of this bid is stored on stack 0. (Implicitly we assume that the order in which bids are submitted—i.e., the number of bids and ask that preceded this bid--uniquely determines the market participant that submitted the bid.)

$M_H$ differs slightly from appendix I, partly because $M_1$ is a double taped push-down automaton. Its transitions have the form $((q,a,\varphi,\sigma),(q',\theta,\rho))\in\Delta$ where $a$ is read of the input tape, φ may be replaced by σ on stack 0 and θ may be replaced by ρ on stack 1. The transitions of $M_H$ are:

1 $((q_i,e,e,I),(q_i,I,e))$ copy I frome stack 1 onto stack 0. (Note that min. ask/bid is stored in reverse order on stack 1.)

2 $((q_i,b,e,a),(q_{i+1},e,b))$ switch states when end of string (stored on stack 1) is reached and start reading next string from input tape.

3. $((q_{i+1},I,I,e),(q_{i+1},e,I))$ simultaneously remove I from stack 0 and add I onto stack 1.

4. $((q_{i+1},I,\#,e),(q_{i+2},e,I))$ bid>min. ask (or bid) stored (#, is the blank symbol).

5. $((q_{i+2},I,e,e),(q_{i+2},e,I))$ copy remainder of input onto stack 1.

6. (($q_{i+2}$,#,e,e),($q_{i+3}$,I,e)) add I to counter stored on stack 0.

7. (($q_{i+3}$,e,e,e),($q_{i+4}$,$,e)) add $ to mark that all I's stored before this $ give the order of bid stored on stack 1(i.e., when the bid was scanned).

8. (($q_{i+4}$,e,e,e),($q_{i+f}$,#,e)) blank symbol is add onto stack 0 to mark end of counter.

9. (($q_{i+1}$,#,I,e),($q_{i+5}$,e,I)) bid<min. ask (or bid) stored.

10. (($q_{i+5}$,e,I,e),($q_{i+5}$,e,I)) copy remainder of stack 0 onto stack 1.

11. (($q_{i+5}$,e,#,e),($q_{i+6}$,I,e)) place I on the counter stored on stack o.

12. (($q_{i+6}$,e,e,e),($q_i$,#,e))

If $M_H$ reaches the halting state $q_{i+f}$, then the subroutine Concatenate will delete the $ stored earlier so that the stack again only contains one $ and every I stacked before $ will again denote the order of the stored bid (i.e., when bid was scanned). Transitions of Concatenate are:

1 (($q_{i+f}$,e,#,e),($q_{i+f+1}$,e,e))

2 (($q_{i+f+1}$,e,$,e),($q_{i+f+2}$,e,$))

3 (($q_{i+f+2}$,e,I,e),($q_{i+f+2}$,e,I))

4 (($q_{i+f+2}$,e,$,e),($q_{i+f+3}$,e,e))

5 (($q_{i+f+3}$,e,e,I),($q_{i+f+3}$,I,e))

6 (($q_{i+f+3}$,e,e,$),($q_{i+f+h}$,$,e))

Machine $M_2$ gives rise to a sealed-bid auction with price priority rule and accumulation rule. $M_2$ is again a pushdown automaton with two stacks, 0 and 1 respectively with $M_1$ as a subroutine. $M_2$ is very similar to $M_1$ only now the machine will only halts if all bids within a certain round are read of. The pseudo-code of $M_2$ is as follows

*Have $M_1$ operate **until** the complete round of bids is read of the input tape.*

2. ***end***.

**Appendix III**.

*Double Auction Market*.

In order to encode the DA on a push down automata $M_3$ it is necessary to have three stacks one additional stack 2. $M_1$ is again a subroutine for the automaton. Additionally, we need to introduce the subroutine $M'_1$ that compares an ask read of the input string with an ask stored on stack 2 storing again the smallest of the two. The transitions of $M'_1$ are very similar to the transitions of $M_1$ given in appendix II. The pseudo-code for the DA is

1. *Have either* $M_1$ *ore* $M'_1$ *operate* ***until*** *the complete round of bids and asks is read of the input tape.*

2. ***If*** *next input string of tape is an ask.*

3. *Have* $M'_1$ *operate on the string (ask) currently scanned.*

4. ***Else*** *have* $M_1$ *operate on the string (bid) currently scanned.*

5. ***end***.

**References**


Anon. 1997. "Intelligent Agents Roboshop," *The Economist*, June 14:72.

Aspray, W. & Burks, A. eds. 1987. *The Papers of John von Neumann on Computing and Computer Theory.* Cambridge: MIT Press.

Badii, R. & Politi, A. 1997. *Complexity: Hierarchical Structures and Scaling in Physics.* Cambridge: Cambridge University Press.

Barbosa, Valmir. 1993. *Massively Parallel Models of Computation.* New York: Horwood.

Cason, Timothy & Friedman, Daniel. 1997. "Price Formation in Single Call Markets," *Econometrica*, (65):311-45.

Conlisk, John. 1996. "Why Bounded Rationality?" *Journal of Economic Literature*, (34):669-700.

Davis, M.; Sigal, R. & Weyuker, E. 1994. *Computability, Complexity and Languages.* San Diego: Academic Press.

Dennett, Daniel. 1995. *Darwin's Dangerous Idea.* New York: Simon & Schuster.

Depew, David & Weber, Bruce. 1995. *Darwinism Evolving.* Cambridge: MIT Press.

De Vany, Arthur and W. David, Walls. "The law of one price in a network: Arbitrage and price dynamics in natural gas city gate markets," Journal of Regional Science, vol. 36, no. 4, 1996, p. 555-570.

Egidi, M. et al. 1991. *Economics, Bounded Rationality and the Cognitive Revolution*. Aldershot: Elgar.

Friedman, D. & Rust, J. eds. 1993. *The Double Auction Market.* Reading: Addison-Wesley.

Garey, M. & Johnson, D. 1979. *Computers and Intractability.* New York: WH Freeman.

Gode, D. & Sunder, S. 1993. "Allocative Efficiency of Markets with Zero-Intelligence Traders," *Journal of Political Economy*, (101):119-137.

Gode, D. & Sunder, S. 1997. "What Makes Markets Allocatively Efficient?" *Quarterly Journal of Economics*, (**):603-30.

Hirshleifer, Jack. 1978. "Natural economy versus political economy," *Journal of Social and Biological Structures*, (2):319-337.

Hodgson, Geoffrey. 1993. *Economics and Evolution.* Ann Arbor: University of Michigan Press.

Kay, Lily. 1997. "Cybernetics, Information, Life," *Configurations*, (5):23-91.

Kula, Witold. 1986. *Measures and Men.* Princeton: Princeton University Press.

Leontief, Wassily. 1966. *Essays in Economics*. Oxford: Oxford University Press.

Lewis, Alain. 1985. "On Effectively Computable Realizations of Choice Functions," *Mathematical Social Sciences*, (10):43-80.

Lewis, Harry & Papadimitriou, Christos. 1981. *Elements of the Theory of Computation.* Englewood Cliffs: Prentice Hall.

Miller, Ross. 1986. "Markets as Logic Programs," in L. Pau, ed., *Artificial Intelligence in Economics and Management.* Amsterdam: North Holland.

Miller, Ross. 1996. "Smart Market Mechanisms: from practice to theory," *Journal of Economic Dynamics and Control*, (20):967-978.

Mirowski, Philip. 1989. *More Heat than Light.* New York: Cambridge University Press.

Mirowski, Philip. forthcoming. *Machine Dreams*. Durham: Duke University Press.

Murphy, M. & O'Neill, L. eds. 1995. *What is Life? The Next 50 Years.* Cambridge: Cambridge University Press.

National Science Board. 1996. *Science and Engineering Indicators, 1996.* Washington: US Government Printing Office.

Nelson, R. & Winter, S. 1982. *An Evolutionary Theory of Economic Change.* Cambridge: Harvard University Press.

Page, Scott E. 1995. "*Two measures of Difficulty,*" California Institute of Technology working paper.

Sent, Esther-Mirjam. forthcoming. "A Simon Who Is Not All that Simple" University of Notre Dame working paper.

Simon, Herbert. 1973. "The Organization of Complex Systems," in H. Patee, ed., *Hierarchy Theory*. New York: George Braziller.

Simon, Herbert. 1981. *The Sciences of the Artificial.* rev. ed. Cambridge: MIT Press.

Smith, John Maynard & Szathmary, Eors. 1995. *The Major Transitions in Evolution.* Oxford: WH Freeman.

Smith, Vernon. 1991. *Papers in Experimental Economics*. New York: Cambridge University Press.

Velupillai, Kumaraswamy. 1996. "The Computable Alternative in the Formalization of Economics," *Kyklos*, (49):251-272.

von Neumann, John. 1966. *Theory of Self-Reproducing Automata.* Urbana: University of Illinois Press.

Weintraub, E.R. & Mirowski, P. 1994. "The Pure and the Applied," *Science in Context*, (7):245-272.

Yockey, Hubert. 1992. *Information Theory and Molecular Biology.* New York: Cambridge University Press.