

Markets Come to Bits:

Evolution, Computation and Markets in Economic Science

Philip Mirowski*

Draft 1.3

November 2004

* Dept. of Economics and Policy Studies, 400 Decio Hall, University of Notre Dame, and Visiting Scholar, International Center for Advanced Studies, New York University. Please direct all correspondence to Mirowski.1@nd.edu. I should like to thank the audiences at Carnegie-Mellon, University of Stirling and the Universidad Republica de Montevideo, Uruguay, and more particularly, Matthias Klaes, Duncan Foley, Robert Axelrod, and an anonymous associate editor for their comments.

ABSTRACT:

Based upon my previous historical work, I attempt to isolate and identify what appears to be a profound shift in the conception of the economy in recent economic research, focusing on five areas: mechanism design, zero-intelligence agents, ‘market microstructure’, engineering economics and artificial intelligence. The shift identified concerns treating markets as diverse algorithms, and will have profound effects upon the conceptual frames used to address the economy. Rather than deal in vague imponderables, in the paper we proceed sketch the emergent outlines of the implicit alternative program of an evolutionary computational economics constructed from the theory of automata which situates the problematic existence of diverse market species at the very center of the research agenda, and not, as happens all too frequently, to relegate it to the margins of modern economic thought. The laws that are sought under the new paradigm are laws of the markets, *not* laws of human nature.

1. The Emergence of Markets Theory

I have argued in (Mirowski, 2002) that the trajectory of the orthodoxy began the 20th century primarily as the oft-acknowledged theory of static allocation, patterned upon classical mechanics, but that during World War II its path got deflected by events and personalities (too numerous to recount here) towards an altogether different conception of its core doctrine, one that might be summarized as recasting the economic agent as an information processor. It goes without saying that the wartime development of the computer and its subsequent diffusion into nearly every sphere of intellectual discourse had quite a bit to do with what has been the most significant reorientation of the economics discipline in the last century, one that has nowhere near yet exhausted its promise. Nevertheless, it is the premise of the present paper that the sciences never rest content; and that further developments in the computational and biological sciences are portending another deflection of the central tendency of microeconomics, which, if successful, will transmute once more the very quiddity of economics. Because we are living in the early stages of the emergence of the new tradition, this paper cannot be constructed in an historical idiom, as was my earlier book. Rather, it is an attempt to describe the stark outlines of the new analytic vision, point out some ways in which it has become manifest in the last few decades, and suggest some incipient lines of development.

In a crude and inadequate manner of speaking, the shift which I think I detect in modern microeconomics is one which is becoming less and less interested in the ‘correct’

specification of the economic agent and her cognitive capacities, and is increasingly concerned with the formal specification of markets as evolving computational algorithms. The reader may be tempted to reject this distinction out of hand: At minimum, the neoclassical tradition has always taken the nature of markets as the central province of economics, has it not? Yet that notion would be premature, as some high-profile economists have noted:

“It is a peculiar fact that the literature on economics...contains so little discussion of the central institution that underlies neoclassical economics—the market.” (North, 1977, p.710) “Although economists claim to study the market, in modern economic theory the market itself has even a more shadowy role than the firm” (Coase, 1988, p.7). Arrow and Hahn’s *General Competitive Analysis* asserts in passing that it takes the “existence of markets...for granted” (1971, p.348).

In fact, a judicious and unbiased overview of the history of the first century of neoclassical economics would confirm that it had been much more fascinated with the status and nature of *agents* than with the structure and composition of markets. Most of the time, the concept of the market was treated as a general synonym for the phenomenon of exchange itself, and hence rendered effectively redundant (Rosenbaum, 2000). Even in the few instances when key thinkers in the tradition felt they should discuss the actual sequence of bids and asks in their models of trade – say, for instance, Walras with his *tâtonnement* and his *bons*, or Edgeworth with his recontracting process – what jumps out at the economic historian is the extent to which the sequence of activities posited therein had little or no relationship to the operation of any actual contemporary market.¹ Mid-20th century attempts to develop accounts of price dynamics were, if anything, even further removed from the increasingly sophisticated diversity of market formats and structures and the actual sequence of what markets accomplish.² Whilst there would be many ways to account for this incongruous turn of events, the condition we shall opt to stress here was the strong dependence of the neoclassical tradition upon *physics* to provide the respected paradigm of scientific explanation. Not only had energy physics provided the

¹ A symptom of the general oblivion to market structures is the urban myth about Walras being inspired by the Paris Bourse. A good historian such as Walker (2001) makes short work of this fairy tale. The only claim that a real-world market anywhere near approximated the actual sequence of events in Walras’ *tâtonnement* of which I am aware is a description of the operation of the post-war London bullion price-fixing ring. See Jarecki (1976).

² The essential disconnect between theories of market dynamics and any empirical sensibility with regard to process is revealed by the historical discussions in Weintraub (1991) and Schinkel (2001).

original agent formalism of optimization over a utility field in commodity space (Mirowski, 1989); it also supplied the background orientation to which law-governed explanations were presumed to conform. The strong reductionism inherent in modern physics suggested that all agents would of necessity exhibit some fundamental shared characteristics (viz., “rationality”) and therefore, for modeling purposes, should be treated as all alike. Furthermore, any differences in market structures where the agents congregated would be treated as second-order complications (viz., perfect competition vs. monopoly) or else collapsible to commodity definitions (‘the’ labor market; ‘the’ fish market), and therefore “The Market” came to be modeled as a relatively homogeneous and undifferentiated entity. Whether justified as mere pragmatic modeling tactic (for reasons of mathematical tractability) or a deeper symmetry bound up with the very notion of the possibility of existence of “laws of economics,” market diversity was effectively suppressed, as one can still observe from modern microeconomics textbooks.

It shall be our claim that the post-1980 weakening of the cultural dominance of physics as the prime exemplar of scientific explanation, and its gradual displacement by the sciences of computation and evolutionary biology, have opened up the conceptual space for an economics which has become less fixated upon agency and more concerned to theorize the meaning and significance of a diversity of (small-m) markets. In the same way we now appreciate that neither biology nor computation can be fully reduced to physics, the incipient vision of markets as evolving computational entities will not itself be reducible to the prior neoclassical tradition. Indeed, one objective of the present paper is to highlight the unacknowledged divergences of this literature from neoclassical precepts, and to elevate to consciousness the ways in which the novel orientation prompts heretofore unasked questions to be broached and answered. Whether the new program has sufficient momentum and gravitas to eventually constitute a viable rival to the neoclassical program is something we leave to the reader to judge.

Precisely because we are dealing with a shift from a period when ‘the market’ has been left implicit and undefined to an era in which markets are becoming the center of attention, it may well be prudent to circumvent the errors and omissions of previous generations and start off with a definition of this contentious term. Perhaps because of the legacy of oblivion, it seems that better attempts at definition are to be found outside of

economics proper, either in the legal literature, or in the technical engineering literature.³ For the purposes of our present argument, we shall define a market as a formal automaton, in the sense of the field of mathematics pioneered by John von Neumann, and now taught as basic computational theory.⁴ Intuitively, we shall characterize a particular market as a specialized piece of software, which both calculates and acts upon inputs, comprised of an integrated set of algorithms that perform the following functions:

- Data dissemination and communications, plus rules of exclusion
- Order routing through time and space
- Order queuing and execution
- Price discovery and assignment
- Custody and delivery arrangement
- Clearing and settlement, including property rights assignment
- Record-keeping

Once one gathers together all the various functions in one place for scrutiny, it may begin to dawn upon the spectator the extent to which the abstract portrait of exchange (derived from physics) as the simple motion of a point through commodity space (not to mention, as we shall discuss below, a continuous topologically connected space of the reals) served to obscure rather than illuminate the broad gamut of market functions. Further, it hindered consideration of their dynamics of *change*. As one legal authority has written,

Approaches which ignore the inner structure of exchanges, such as those viewing an exchange solely as a trading system (as in the microstructure finance economics), a reduced form production function (as in neoclassical economics), an impersonal investment guided by managers (as in organization theory), are not rich enough in detail or subtlety to be able to understand the nature and conduct of exchanges. (Lee, 1998, p.8)

2. Where the Auction Is

³ The following definition is derived from: Domowitz (1993), Lee (1998) and Wellman (2004).

⁴ See Davis et al (1994) and Khoussainov & Nerode, (2001). The von Neumann legacy is covered in Mirowski (2002, chap.3).

The treatment of markets as evolving computational entities has not, by and large, been a self-conscious academic movement, either from within the economics profession or without. From within the profession, it has tended to be complementary to the widespread practice of modeling the agent as an algorithmic and strategic reasoning machine, that is, as resembling a computer in many respects. Once the agent came to be modeled as a computer program, it turned out to be a very small step to recast agent interaction as a larger, more encompassing computer program, perhaps with the individual agents as sub-routines. Much of game theory has tended to adopt that format; so too did many of the early attempts to automate subject interaction in experimental economics. The quest to model markets in such disciplines as Artificial Intelligence also set out from initial concentration on the production of artificial agents; only later did the algorithmic nature of markets become a topic of research in its own right. However insensibly and unself-consciously it may have begun, there now exists at least five areas of the literature that have managed to develop a fair corpus of work concerning markets as computational entities. They are (in rough chronological order of appearance):

A] The theoretical literature on “mechanism design” in general and auction design in particular;

B] The offshoot of the experimental economics literature known as the “Zero-Intelligence Agent” canon;

C] The “market microstructure” literature within finance;

D] The modern “engineering economics” literature; and

E] The Artificial Intelligence literature dealing with markets.

The odd aspect of this situation is the distressingly limited extent to which any writer situated in one of these communities seriously references work within any of the other communities. It is almost as if their disciplinary locations served to trump what by all accounts would be shared concerns. While not exactly mired in splendid isolation, the relative insularity of these groups may have presented one obstacle to full emergence of the paradigm of evolutionary computational markets. In the interests of encouraging dialogue, we provide a brief précis of the relevant work in each of those areas.

2.1] Mechanism Design

The field of mechanism design is generally dated from the appearance of Hurwicz (1960). In the postwar development of the Walrasian tradition at the Cowles Commission, two problems were viewed as immediately pressing: the provision of a plausible dynamics to complement the celebrated existence proofs of equilibrium, and the countering of Hayek's arguments against the possibility of planning in a market economy.⁵ One attempt to address both simultaneously was the research program of Leonid Hurwicz, which started out as an attempt to refine the notion of a "decentralized economy" and explore the properties of convergence to Walrasian equilibrium. Realizing that earlier claims that Walras had shown that agents only needed to "know" their own preferences and prices were baseless, Hurwicz set out to formalize the relative autonomy of the agent (which later mutated into the notion of 'incentive compatibility') and the informational decentralization (which was parsed out as 'privacy' and 'anonymity' of the agent) in a model of the passing of 'messages' and their consequences for 'action', all superimposed upon the canonical neoclassical model promoted by Cowles. Since this was an era in which agnosticism about the cognitive nature of the agent was rife, these models were not interpreted as amendments of agent 'rationality' so much as they were purported to provide characterizations of different generic allocation mechanisms; hence the self-description of this school as being concerned with "mechanism design". As Hurwicz wrote, "Instead of being a *given*, the mechanism becomes the *unknown of the problem*. In this spirit, the economist can view the problem as one of designing a mechanism maximizing certain social desiderata, such as efficiency, equity and freedom – subject to behavioral and informational constraints" (1981, p.300).

In retrospect, it is surprising the extent to which all this discussion of alternate 'mechanisms' managed to keep well segregated from the consideration of anything we might consider today as actual market formats. Even more germane to our present concerns is the observation that these models managed to avoid any real machines: in their execution, they elided any serious consideration of formalizations of communications theory, network theory, linguistics, or any other topic within

⁵ Both of these imperatives are discussed at length in Mirowski (2002), as are subsequent statements in this section concerning the treatment of computational issues in this tradition. For a detailed survey of the early computational aspects of mechanism design from Hurwicz to Reiter, see Lee (2004, chap.3).

computational science contemporary with those developments. It seems that the first phase of the mechanism design literature was much more concerned to rectify perceived flaws in the 1960s-vintage general equilibrium model; for instance: Was there a dynamic story which might prove superior to Walrasian tâtonnement? How should the model be revised to incorporate nonconvexities and indivisibility of goods? Could one also satisfy other welfare objectives other than Pareto optimality, such as ‘fairness’? Many Cowles-affiliated researchers sought to treat these questions as though they could themselves be treated as trade-offs: for instance, trading some ‘freedom’ for better convergence, or trading off higher ‘costs’ of the mechanism for more ‘fairness’ (Hurwicz, 1986; 1994). This welfarist perspective was rendered more openly apparent in the work of Stanley Reiter (1977). However, both the presumption of an uncontentious welfare orientation and the quest for general principles of Walrasian dynamics and stability of equilibrium faltered in the 1970s for reasons we cannot cover here; it might have seemed that as a consequence the nascent literature would then lose its *raison d’être*. That did not happen, for two unrelated reasons. First, the independent transfer of allegiance of microeconomic orthodoxy from Walrasian models to Nash game theory, which gathered momentum in the later 1970s, actually bolstered this literature through the introduction of a previously absent *strategic* dimension to the question of mechanisms as information systems. This branch of the literature became embroiled in the pursuit of strategy-proof allocation mechanisms which would supposedly induce agents to emit truthful signals due to incentive compatibility (Ledyard, 1987), and rapidly became conflated with principal-agent and other problems of asymmetric information, all of which drew attention further and further away from the original aim of formalizing alternative *market* mechanisms. Perhaps one reason these researchers turned away from markets *per se* was that recourse to the epistemic interpretation of Nash equilibrium required that agents be super-humanly informed, and thus it appeared incongruous to maintain the pretense that the allocation mechanism itself was kept in the dark about their preferences and intentions. This turn to Nash constituted a major input into the subsequent rise of the specialized literature on auction theory.⁶

⁶ While many might regard the auctions literature as an ‘engineering’ phenomenon (as described in section 2.4 below), we would also stress its provenance out of the mechanism design literature. A sketch of the

The second development, more pertinent to the present paper, was the slow recognition that the mechanism design literature had been re-inventing the wheel, in that many of its themes had already been addressed within computer science, and in particular, the theory of automata. This trend began with the paper by Edward Ames (1983), which translated the Hurwicz mechanism into automata theory. In a way that was possibly unanticipated, the initial metaphor of ‘mechanisms’ was recast into the explicit notation of formal sequential machines. This program has subsequently been pursued to a higher degree of abstraction in Mount and Reiter (2002). This latter work openly declares its commitment to the original neoclassical program, but finds it has to dispense with one of the cardinal precepts of the school: “it is not appropriate to separate the person from the machine... there is a need for a model of computation that facilitates analysis of computational problems to be solved by a combination of human and machine resources” (2002, p.5). This artifice, we argue, illustrates the transformation that we maintain has been imperceptibly happening throughout much of microeconomics: as computational themes loomed larger, the dominant role of the rational human agent was diminished. What foiled this literature from venturing even further down the path to an evolutionary computational economics was an unwillingness on the part of its champions to subject the neoclassical agent to a full complexity audit with regard to their capacity to even compute a fully transitive and complete set of preferences, either as a static proposition or through some theory of “learning”, much less to perform feats of inference whose complexity outstrip the capacity of the standard model of the Turing Machine. The repeated recourse to Pareto superiority to rank mechanisms, and indeed, appeals to conventional ‘welfare’ itself as an appropriate category to attribute to human/machine agglomerations, begs the entire question of the proper location of the issue of computational limits to mechanism design.

2.2] Experimental Economics and zero-intelligence Agents

The role of experimental economics in the shift to a market-centered theory of computational economics is a story that encapsulates more twists and turns than a

development of auction theory that points to its roots in mechanism design can be found in Nik-Khah (2004).

season's episodes of the *Sopranos*. The movement is very diverse, but we shall restrict ourselves to illustration of one strand found in the career of Vernon Smith.⁷ Smith, as is well known, produced his first experiments in reaction to Edward Chamberlin's Harvard classroom exercises that purported to show that free competition would not produce efficient exchange outcomes. Smith bypassed psychological experimentation (which had tended to refute neoclassical models of agency) in favor of market simulations predicated upon rules derived from stock exchange manuals. His first experimental article ([1962] in 1991) demonstrated that in the presence of what became known as 'double auction' rules, prices and quantities rapidly converged to supply/demand equilibria, even in the absence of any knowledge of the theory or data on the part of his subjects (usually university undergraduates). His early success led Smith to explore a range of other market formats instantiated as alternative rule structures in his laboratories; in conjunction with his colleagues, he found that no other type of market regularly produced what they considered to be superior outcomes, namely, rapid convergence to predicted price and quantity equilibria and near-full realization of the pre-defined consumer and producer surplus in his experimental setups. (In contrast with Pareto optimality, both objectives were computable by construction.) One lesson he drew from this line of research was a thesis he dubbed the 'Hayek Hypothesis' ([1982] in 1991): "Strict privacy together with the trading rules of a market institution are sufficient to produce competitive market outcomes at or near 100% efficiency," that is, independent of the cognitive abilities or status of the agents involved. Another lesson grew out of the move to automate his experimental laboratory through integration of computer technology in the mid-1970s. In his own words:

Science is driven more fundamentally by machine builders, than either the theorists or experimentalists... Computer/communication and imaging technologies are driving experimental economics in new directions. Both will marginalize extant research on individual decision. When Arlington Williams programmed the first electronic double auction (e-commerce in the lab) in 1976, it changed the way we thought about markets, much as the internet is changing the way people think about doing business. Circa 1976 we thought going electronic would merely facilitate experimental control, data collection and record keeping. What we discovered was altogether different: computerization vastly expanded the message space within which economic agents could communicate at vanishingly small transactions cost. This

⁷ More elaborate accounts of Smith's work can be found in Lee (2004) and Miller,(2002).

enabled Stephen Rassenti to invent the first computer-assisted combinatorial auction market.... Lab experiments became the means by which heretofore unimaginable market designs could be performance tested” (2001, p.428).

So early experience with the computer automation of experiments, which meant in Smith’s case experience with computer simulation of various market operations, prepared the way for Smith-style experimentalists to be among the first researchers with the requisite skills to program and implement never-before imagined variants of electronic markets. The computerization of experimentation also had more than a little to do with Smith’s predisposition to focus upon the double auction format to the exclusion of other forms: more than almost any other species of market, it was amenable to full reduction to an algorithm. (Other formats, such as the dealer/market maker, were less immediately susceptible to reduction to a comprehensive set of abstract rules.⁸) These specializations in the coding of markets and favorable inclinations towards auctions, of course, will explain the preponderance of experimental economists to be found occupying the nascent field of ‘engineering economics’. But more to the point, it also encouraged a more direct research initiative on the part of others even less committed to the neoclassical program than Vernon Smith himself to develop an analytical rationale concerning the relative independence of the algorithmic market from the neoclassical agent.

While there were a number of conceptual problems with Smith’s enunciation of his “Hayek hypothesis,” the most nagging was the query: If the success of the double auction format was truly not dependent upon the cognitive states of the experimental subjects, then what did account for it? This question was posed and answered brilliantly by two Carnegie researchers, Dan Gode and Shyam Sunder. In a now-classic paper (1993), they compared the experimental outcome of a double auction setup using human subjects with a fully automated setup that replaced the subjects with random-number generators, which they dubbed ‘zero-intelligence agents’ [henceforth ZI]. While there is still substantial dispute over the interpretation of their results,⁹ it appeared that brainless

⁸ Indeed, serious attempts to bring the dealer/market maker format into the lab only began with Krahnert & Weber (2001).

⁹ See Cliff (2001); Walia (2002); and Mirowski (2002, pp. 551-560) for various attempts to draw out the implications.

programs produced nearly indistinguishable results with regard to convergence and efficiency compared to Smith's human subjects. The computational insight of the 'zero-intelligence' exercise was that human cognitive capacities could be zeroed out under controlled circumstances (thanks to the prior automation of experimental economics) in order to explore the causal capacities of markets conceived as free-standing automata. The predictable regularities of the double auction in experimental settings (and perhaps "in the wild", as experimentalists were wont to say) should be attributed to their algorithmic structures, and not to any psychological predispositions of their participants; as Sunder himself put it, "a science of markets need not be built from the science of individual behavior... Market institutions may be society's way of dealing with human cognitive limitations... Efficiency of markets is primarily a function of their rules" (Sunder, 2004). In Gode and Sunder's subsequent work, they deconstructed the auction down to its algorithmic parts and subjected each to the ZI trader test in order to further explore the sources of the efficacy of the double auction. In a review of the ZI and other work on the double auction for the Santa Fe Institute, another experimentalist, Dan Friedman, proposed a market species taxonomy that resembled a phylogenetic tree (in Friedman & Rust, 1993, p.8), but did not follow up on the evolutionary implications.

2.3] Market microstructure in finance

The interest in diverse market structures within finance seems to be the consequence of a different set of converging forces than those that prompted developments in mechanism design and experimental economics, which may go some distance in explaining the fact that the program has kept relatively aloof from those developments. The tempo and extent of the computerization of stock markets has outpaced that of any other species of markets, and has remained a contentious issue ever since Reuters began to offer its "Stockmaster" service in 1964, and continued with the National Association of Securities Dealers launch of its Automatic Quotation system (NASDAQ) in 1972 (Wells, 2000), and the opening of the first fully automated electronic exchange, the Toronto CATS in 1977. However, the idea that certain aspects of market rules or algorithms might prove faulty, to the dire extent of precipitating collapse, impressed itself with stark insistence upon participants in the aftermath of the crash of

1987, and from thenceforth the self-identified “market microstructure” literature began to explode.¹⁰ Early attempts to ground the models of market activities in the rational expectations economics literature soon faltered (Admati, 1991), and were replaced by models seeking to capture the rivalry between mechanical crossing systems and a dealer/market maker who provided liquidity and various other services. This fascination was itself the partial consequence of an attempt by Congress to mandate the eventual appearance of a National Market System for stocks in 1975 (Miller, 2002, p.264) – the imposition of the neoclassical image of the monolithic unified Market by fiat -- and a subsequent series of attacks by the SEC on dealers for what was deemed ‘noncompetitive behavior’. In the interim, the more highly computerized derivatives markets came to dwarf their prosaic securities predecessors in volume and volatility (MacKenzie & Millo, 2003). Once they had embarked down this path of exploring the impact of process on outcomes, the finance researchers encountered problems for which no existing microeconomic theory had sufficiently prepared them. “The perfect market hypothesis, that trades have no impact on prices, has been strongly rejected” (Biais, Glosten & Spratt, 2003). Should all floor trading be swept into the dustbin of history by electronic systems? Was the dealer an obstacle to efficient operation of the market, or rather, an indispensable component? Was the apotheosis of the elimination of ‘transactions costs’ a unified national limit order book, or something else altogether? Did the fad for privatization of regional and national exchanges render them vulnerable to the centralization of financial markets in just a few metropolitan world centers? These conundrums led to deeper, seemingly ontological issues, for which no one was adequately prepared by their training. To wit: What was being presumed to be the real or fundamental objective of a securities market? Was it merely to systematically organize order flow, or provide a portfolio of implicit options, or to facilitate cross-markets arbitrage, or perhaps to embody and convey information to further ‘price discovery’, to reduce price volatility? Or was it something more, something closer to Keynes’ beauty contest, a vast asymmetric

¹⁰ For surveys, see O’Hara (1995); Madhavan, (2000); and Biais, Glosten & Spratt (2003). The interaction of the crash of 1987 with finance theory in general is nicely covered in MacKenzie (2005).

information game, where the whole point of the exercise was to defeat the opponent?¹¹ Models of microstructure began to diverge not only in minor variations in rules (order routing, execution, clearing) but also in major differences over what the market was conceived of accomplishing.

The fact of closer familiarity with and attention to a subset of real markets in real-time operation generally meant that the microstructure literature was much less inclined to imperiously pronounce a single market format uniformly more ‘efficient’ or preferable than its rivals, or at least that part of the literature not emanating from within economics departments. Even those who propounded the optimality of the periodic call auction were distressed to admit its relative failure in catching on in any modern financial environment. One small offshoot of the literature began to even model auctions as modular computer programs (Domowitz & Wang, 1994), in order to explore whether continuous or periodic auctions had different effects on price volatility and bid-ask spread. Arguments were broached to the effect that more information was not always better for market operation, that ‘transparency’ often had unintended consequences, denials that finer price gradations or ‘ticks’ always led to greater market stability, and the irony that stupid mechanical ‘noise traders’ were often identified as the indispensable ingredient for a successful market, even though their very existence “makes it impossible to conduct any welfare analyses or to compare different market structures, since [they] prevent accounting for the impact of market structure on noise trading” (Biais, Glosten & Spratt, 2003). But if noise traders effectively thwarted the reduction of microstructure theory to a reflection of the desires of the rational agent, they were certainly amenable to the modeling of individual market structures as diverse automata. This trend was only magnified by the waves of physicists *manqué* and other engineers recruited by Wall Street firms to program automated trading algorithms and concoct ever more baroque derivatives. The market microstructure literature found itself repeatedly returning to what seemed the key question: What were the ultimate reasons “for the wide diversity of trading mechanisms across assets” (Madhavan, 2000, p.251), not to mention within the compass of a single asset?

¹¹ The fundamental inability of the rational actor paradigm to address these issues became most apparent in the so-called “no-trade theorems”; the mere fact of entering a bid or ask supposedly exposed you to all manner of exploitation by knowledgeable participants.

2.4] Economists as engineers

One of the great challenges for intellectual historians of the future will be to explain how it came to be that a professional academic orthodoxy that had eschewed most considerations of the specificity of markets (as explained in Section 1 above) then neatly negotiated a 180° turn, and managed to convince a broad array of outsiders that they possessed special expertise to construct all manner of actual usable markets, tailor-made for their narrowly specified purposes. The novelty of this development can be savored by noting that “engineering economics” in the 1960s-80s designated someone concerned with planning of allocation, usually in some branch of industry, and was familiar with materials science and other engineering topics,¹² whereas now by contrast it denotes someone who designs markets, possibly codes them as computer programs, and perhaps even patents his algorithms as ‘intellectual property’. The terminology of ‘engineering’ is used advisedly, since engineers do not restrict their activities to a single methodology or theoretical tradition.

Engineers often use different levels of incommensurable models to test their design formulations throughout the design process. Improvements in these methods of testing come from advances in several areas including theoretical or analytical modeling, physical testing and computational modeling... This cycle of formulation and reformulation takes place in an evolutionary manner as experience in identification and testing of unanticipated externalities and failures accumulate. (Subrahmanian & Talukdar, 2003).

Here is where matters become contentious, since this literature is clearly divided between those who insist, generally without direct documentation, that economic engineers have successfully designed modern auctions and other markets by means of the *application* of neoclassical theory (usually, but not exclusively, Nash game theory), and those who come closer to admitting that their experience with engineering rarely validates the existing toolbox.¹³ Our intention is not to contend with the former through the provision of historical evidence in the present context (see, however, Nik-Khah, 2004), although we would suggest this commonplace piece of folk-wisdom, repeated

¹² For a textbook example of this now-vanishing creature, consult (Blank & Tarquin, 1976).

¹³ For examples of the former, see Vulkan (2003); Binmore (2000); Dash et al, (2003). Examples of the latter are Roth (2002); Salant (2000) and Phelps et al (2004)

even in introductory textbooks, requires much greater scrutiny. Rather, we seek to briefly indicate how those economic engineers who do acknowledge the ambiguities in the process of the construction of markets are regularly pushed in the direction of an evolutionary computational economics.

A celebrated paper that makes this point is Alvin Roth's "The Economist as Engineer" (2002). As he writes therein, "markets don't always grow like weeds – some of them are hothouse orchids" (p.1373). Although this is probably only metaphorical flourish, it is noteworthy that neither Roth nor anyone else who indulges in such biological language (Vulkan, 2003, p.98) seems to evince any curiosity concerning the fact that their tradition is utterly devoid of any theory of the 'natural' growth of market forms. It would almost be tantamount to asserting a capacity to conjure life in a test tube without any recourse whatsoever to biology. Nevertheless, Roth does concede that, when it came to construction of a mechanism to allocate hospital internship slots to medical school graduates, "None of the available theory... could be directly applied to the complex [intern] market... The only theoretical parts of [his] book that applied directly to the medical market were the *counterexamples*" (2002, p.1372). The engineer was consequently forced to become a *bricoleur*, observing that, "in the service of design, experimental and computational economics are natural complements to game theory" (p.1342). Roth reports numerous incidents where he was led to attend to the computational complexity of various proposals, not by theory, but by pragmatic concerns. Furthermore, in his own experience as much as in the FCC spectrum auctions, "experiments were constructed not to test specific hypotheses...but rather as 'proof of concept' demonstrations" (p. 1369). The proof of the pudding was not in the achievement of any given orthodox definition of 'efficiency', but rather in the satisfaction of the various client groups. The pole star of 'informational efficiency', the guiding light of formal mechanism design, was unavailing in these applied contexts, since, as another specialist has written, "The concept of informational efficiency is fraught with difficulty... There is no clear basis in the economics literature for concluding that informational efficiency in markets does obtain, and if so to what extent" (Lee, 1998, p.248). No practicing engineer could long pretend that there was any such thing as an 'optimal auction,' much less an optimal generic 'mechanism'.

One notable deficiency with many of the self-identified engineers within the economics profession is that they did not cast their net very widely in pursuit of conceptual resources with which to ply their trade. A perspicuous counter-example to this generalization is Ross Miller, who started out with training in economics and Smith's experimental tradition, but augmented it with expertise in computational theory, and then went to work in the real world of financial markets. Well before any of our other cadres, Miller foresaw that it was more promising to treat the market itself as algorithmic, instead of the agent (1986); he later argued that if orders have to be executed in an all-or-nothing manner, the law of one price might be violated in 'efficient' settings (1996); and more recently he has maintained that, "Any market mechanism that blindly relies on textbook economic theory certainly cannot be considered intelligent" (2002, p.236). Because his bailiwick is finance, he shares with the experimentalists and market microstructure schools a fascination with the automated double auction, but diverges from them in concern for what it might mean to endow the DA with the computational attributes of a 'smart market'. "Smart markets" are an engineering term of art that refers to forms of automation that are constructed to assist traders in their activities. While not avowedly evolutionary, Miller's orientation permits him to ask where modern enthusiasm for the engineering of financial smart markets might lead. Taking seriously the proposition that institution of a national system of electronic open limit order books in stock markets would imply excess arbitrage opportunities, given that any placement of a limit order is the provision of a potentially 'free' option to someone else, he proposes that a truly unified system of auctions would require automatic self-arbitrage between options markets and share markets. These 'smart markets' would rationalize the 'front-running' of traders, something that now only happens on a sporadic and arbitrary basis, by automated creation and sale of options instantaneously adjusted to the existing distribution of portfolios of shareholdings. A fully integrated electronic system would then finally deconstruct the last remaining tissue of distinctions between the financial markets and the gambling industry:

One can even imagine taking the toxic waste from asset-backed securities and recycling it back into the plethora of casinos that have blossomed across the American landscape... it may make little difference to gamblers whether the numbers and fruit-like objects that appear on their displays come from a microprocessor based random

number generator or are leftovers from the kitchen of Wall Street's financial engineers (2002, p.235).

The purpose of citing this chilling *Matrix*-like vision of a dystopian future is not to endorse its likelihood, but instead to suggest that the expanding phenomenon of engineering markets for hire will eventually itself need to be incorporated into a evolutionary computational theory of markets, if it is to venture beyond the naïve prognostications of sweetness and light that surround the conventional 'auction design' literature. A theory of the consequences of the large-scale engineering of markets only makes sense when embedded within a larger theory of the 'natural history' of market evolution. If Miller's vision seems unduly pessimistic, that can only be attributed to the fact that he has fallen into the habit of regarding all markets as minor variations on homogeneous auctions, rather than keeping in view the variegated motley of species that is revealed in a proper natural history of markets.

2.5] Artificial intelligence and automated markets

The incursion of specialists in artificial intelligence into the precincts of economics may seem an unexpected turn of events, although a case could be made that they share a genealogy with certain branches of economics.¹⁴ Nevertheless, the artificial intelligentsia found themselves lured into confrontation with questions of economics by the Internet bubble of the 1990s. Initially, it seemed that their long-standing concern with agency and its simulation was a natural match for the commercialization of the Internet. As all manner of commodities were being promoted for sale online, it rapidly became apparent that there was room for automation of various aspects of the commercial process. Difficulties in finding the exact commodity you were looking for and conducting price comparisons, not to mention negotiating the exotic processes of bidding and asking online, suggested that a little bit of automated intelligence might come to the rescue. Precisely because AI had enjoyed a tradition of constructing simulations of agents of one form or another dating from the 1950s, it seemed natural that the earliest forays into economics assumed the format of forging little 'autonomous artificial agents'.¹⁵ For a

¹⁴ See, for instance, Mirowski (2004; 2002, pp.456-469) and Sent (2001).

¹⁵ See Wagner (2000) and Jennings et al, (1998). Given the obsession with agency in neoclassical economics, it should come as no surprise that the acceptable face of AI within the economics profession

while, AI researchers such as Jeffrey Kephart and Pattie Maes were making the sorts of hyperbolic predictions about the future of webbots that one tended to associate with an earlier generation of AI boosters.¹⁶ While the ‘shopbot’ approach still boasts its advocates, it seems that this wave of enthusiasm for the computational automation of agency has been played out, at least for the present,¹⁷ and replaced by a more solidly supported initiative which seeks to apply computational expertise to the construction of market algorithms.

One of the most important representatives of this subtle shift within AI is Michael Wellman. Wellman started out as an enthusiast for the shopbot approach, but had the wonderful idea of convening the (now annual) Trading Agent Competition at the University of Michigan, a tournament in which shopbot engineers could submit their programs to compete with each other to achieve some prespecified set of goals.¹⁸ What Wellman and his colleagues at the AI Lab had to do was to program the framework – ie., the ‘market’ – within which the competition was to take place. In order to provide a flexible platform upon which to run future competitions, Wellman sought to produce a generic configurable auction program that would modularize interchangeable components (Wellman 2004, p.9). While he may have been frustrated in achieving this goal, he and his colleagues were brought face to face with the absence of any general theory of markets in the course of their endeavors. While not exactly hostile to neoclassical theory, he did notice that there was a tendency for economists to be a little too fixated upon Walrasian mechanism design and financial markets, and the prevalence of the double auction in particular (2004, p.14). He and his colleagues therefore sought to provide a more comprehensive taxonomy of auction design space, opening up for consideration

still assumes the format of “Agent-based Computational Economics.” See, for instance, the extensive website maintained by Leigh Tesfatsion at www.econ.iastate.edu/tesfatsi/, or Tesfatsion (2003); Potts, (2000); and Axtell (2002).

¹⁶ “Humans are on the verge of losing their status as the sole economic species on the planet. Software agents will [soon] represent –and be—consumers, producers and intermediaries” (Kephart, 2002). There seems to be no immediate danger of this a half decade later.

¹⁷ Partly due to legal restraints; see Kaplan (2001); but more germane was the bursting of the Internet bubble. One illustration is the discontinuance of IBM’s Institute for Advanced Commerce, an incubator for artificial agency. Visit the prematurely frozen website at www.research.ibm.com/iac.

¹⁸ The latest round of TAC is described in Wellman et al (2004) and Zhang et al. (2004). More elaborate description and documentation can be found at <http://www.sics.se/tac>. News updates can be gleaned from the newsletter of the ACM Special Interest Group on e-Commerce at www.acm.org/sigs/sigecom/exchanges.

algorithmic issues of information dissemination and clearing policy, going well beyond the more conventional mechanics of the bid-ask mechanism (Wurman et al, 2001). As one can readily appreciate from perusal of their Table VIII (p.335), the combinatorial explosion of different ‘types’ of auctions would transcend the usual evocation of a small menu of mechanisms.

The study of market design has subsequently diffused rapidly to all the major departments of computer science in America by the turn of the century.¹⁹ In many instances these scholars tended to reprise some themes from earlier mechanism design or engineering economics; but in many cases they have ventured well beyond the economists. For instance, it has devolved to these and other AI researchers to complement this approach with the more concerted application of computational complexity theory to auction design. For instance, Sandholm & Suri (2001) have shown that, in the limited case of the clearing mechanism of call auctions, the attempt to impose discriminatory pricing in a two-sided auction is *NP*-complete, although uniform price auctions are polynomial in time. Conitzer and Sandholm (2002) show that, were one to try and formalize a generic mechanism design program that sought to satisfy different preference specifications for each individual setup, then Bayes-Nash equilibrium reduces to the knapsack problem, which is *NP*-complete. One suspects if they had extended their consideration to the constitution of preference functions, they would find that the problem of non-computability (and not just computational complexity) extends all the way down to the root of their initial assumptions concerning agent intentionality; although in practice, AI researchers have been happy to work with higher-level primitives, such as given piecewise linear demand curves. Other researchers have explored the use of genetic algorithms to search over a space of possible variant auction mechanisms in order to arrive at some notion of an improved mechanism design, although their findings have been hampered by an inadequate parameterization of the full range of species of auctions (unable to encompass even the taxonomy in Wurman et al, 2001), and an undue dependence upon zero-intelligence agents to reduce their domain to

¹⁹ See the syllabi posted on the web of the following: Tuomas Sandholm at Carnegie Mellon, <http://www-2.cs.cmu.edu/~sandholm/cs15-892.html> ; Joan Feigenbaum at Yale, <http://zoo.cs.yale.edu/classes/cs455> ; David Parkes at Harvard, <http://www.eecs.harvard.edu/~parkes/cs286r/emd.html> ; Subhash Suri at Univ. of California-Santa Barbara, <http://cs.ucsb.edu/~suri/ecom.html> ; and Peter Stone at Univ. of Texas, www.cs.utexas.edu/users/ptstone/courses/395tfall03 . All sites last accessed June 2004.

a closed form.²⁰ Whatever the present obstacles which hamper their attempts to feel their way through the economic landscape, it is interesting to see that the AI researchers have essentially traversed the same trajectory from initial stress on agent models to treating different species of markets as algorithms, in much the way the other groups have done.

3. Markets constructed and markets Evolving

The previous section has documented a widespread phenomenon of diverse academic groups who seem united in their ambitions to “make markets better” by resorting to various computer metaphors or formal theories of computation. Each of these groups have traced a more or less similar trajectory, from beginning with computational models of the generic economic agent in a generic ‘Market’, and ending up with something resembling a model of diverse market forms (possibly facing a diversity of agent types) based upon computational notions. In some instances, the portrait of the diversity of market species is an entry point to introduce some further notions of economic evolution. Because this change in mindset has happened slowly, and in an insensible and unself-conscious manner, it seems many of its proponents have been unaware that it involves a serious contradiction. The problem is this: What could it mean to make markets “better” if the change in conceptual orientation is taken to its logical conclusions?

The contradiction unfurls by noticing that many of these analysts frequently make some distinction in passing between ‘markets in the wild’ and their own constructed entities; but then they simply *opt to ignore the non-constructed markets*. For all their ambitions to build markets for their clients, they tend to behave as though this activity could take place in a vacuum. But the question cannot be avoided: Why isn’t the pristine ‘natural’ situation ‘in the wild’ one that cannot be improved upon? Many of the above authors tend to finesse this question by recourse to some conventional notion of “welfare”, ranging from Pareto optimality to postulating specific agent objective functions. But then their program gets snagged on the other horn of the dilemma. Insofar as the conceptual trend is to downplay the specification of the agent, or at minimum, to acknowledge such a range of diversity of agents that their aims cannot be reduced to a

²⁰ See, for instance, Byde (2004); Cliff (2001); Walia (2002); Walia et al, (2003); and Phelps et al, (2002).

single objective function, then it becomes apparent that this supposed escape route is closed to these analysts. We have already observed the symptoms of this contradiction in our summaries in Section 2. For instance, what can be the meaning of Pareto optima for a ‘machine-human agglomeration’ in the modern mechanism design literature? Or: What does ‘consumer surplus’ signify for a population of ZI traders? Or: how is it possible to appeal to neoclassical welfare improvements for ‘noise traders’? Or: What is the import of satisfaction of the client who commissions the engineered mechanism when it rapidly meets political opposition by other groups impacted by the new-forged ‘mechanism’? Welfare economics has always subsisted uneasily in the history of neoclassical economics, but in these literatures it meets its Waterloo. Each and every one of these dubious phenomena can be found in the literatures cited above, and yet, the authors persist in behaving as though the fruits of their labors should be treated as enjoying uncontroversial status as a better mousetrap.

The shift identified in this paper in each of these research programs in the direction of treating markets as diverse algorithms will have profound effects upon the conceptual frames used to address the economy. Although the automatic tendency will be to try and recast those frames as consistent with everything that has gone before in economic science, the track record so far has not been promising. When the template of a successful scientific model shifts from physics to biology and computational science, the repertoire of practices will tend to be revised from top to bottom: from the mathematics employed, to the forms of empiricism encouraged, to the tools put to use, to the kinds of questions asked and answers tendered.²¹ Rather than deal in vague imponderables, in the rest of this paper we shall sketch the emergent outlines of an alternative program of an evolutionary computational economics constructed from the foundations up in order to situate the existence of diverse market species at the very center of the research agenda, and not, as happens all too frequently, to relegate it to the margins of modern economic thought. The laws which will be sought are laws of the markets, *not* laws of human

²¹ This raises the issue of what can be hoped for from the pursuit of an ‘economics of complexity’, a question posed in Rosser (1999). Most of the literature cited therein is characterized by the disappointing track record of various attempts to directly appropriate models from physics, and then bend them to the description of economic variables. This constitutes a major source of continuity in the history of neoclassical economics (Mirowski, 1989; 2002; Thompson, 2004). Here, by contrast, we abjure the physics, and attempt to build models which openly relate markets to computational entities, which then evolve according to economic criteria.

nature—something which we believe economists are ill-suited and ill-equipped to pursue, whether they exist or not.

Although the following quote is taken from a description of a global system of meteorological scientific observation, it could apply equally well to the global system of markets:

Standards are socially constructed tools. They embody the outcomes of negotiations that are simultaneously technical, social and political in character. Like algorithms, they serve to specify exactly how something will be done. Ideally, standardized processes and devices always work in the same way, no matter where, what or who applies them. Consequently, some elements of standards can be embedded in machines or systems. When they work, standards lubricate the construction of technological systems and make possible widely shared knowledge. In practice, few standards can be specified as algorithms. Therefore, most standards also involve discipline on the part of the human participants, who are notoriously apt to misunderstand and resist. As a result, maintaining adherence to a standard involves ongoing adjustments to people, practices and machines. (Edwards, 2004, p.827)

4. The theory of Markomata

The culmination of the trends identified above is to formally model the diversity of market forms by making use of the mathematical theory of automata, and then to taxonomize and organize these entities by means of a theory of evolution. The abstract theory of computation seems well-suited to encompass the diverse (and open-ended) roster of functions performed by the range of extant market forms: data dissemination, order routing, order execution, price and quantity output, delivery, clearing and settlement. A half-century of experience with computers has taught us that they are not simply or solely calculators or language-recognition devices (although that is the idiom that has been prevalent in their formalization), but protean command-control-communication devices, the consequences of which often outstrip the intentions of their builders. Although experience with markets has extended back through incomparably more vast stretches of history, the realization that markets are equally command-control-communications prostheses has been stymied up until now by the century-old predilection to pattern market models upon physical machine systems (Mirowski, 1989). This tendency within economics has not only prompted recourse to physical mathematics (the calculus, field theory, Euclidean space), but also physics-envy aspirations to a Theory of Everything in which all markets were but minor variations on a canonical

model of The Market occupied by The Agent. A theory of markets based upon automata theory provides a welcome propaedeutic, in that it codifies the fact that there is no ur-model or über-machine to which the blooming, buzzing profusion of markets can be reduced. Furthermore, since the theory of automata is independent of the nature of the substrate upon which they may be physically realized, the program is amenable to portrayal of markets as composed solely of humans, or human-machine hybrids, or indeed, entirely of machines in the format of modern computers. Since this constitutes a major departure within the history of economic thought, we shall require a new term to refer to such entities; thus we coin the neologism “markomata”.

What role is played by abstract mathematical theory in such a research program? First and foremost, it provides an analytical framework of permissions and prohibitions of what can and cannot be done by specified classes of markomata. This immediately dispels the misapprehension promoted by Friedrich Hayek and his followers that all markets everywhere are indifferently effective purveyors and processors of information. Secondly, it reveals how diverse markomata can be arrayed in hierarchies for the purposes of further analysis: hierarchies of computational capacity, hierarchies of language recognition, and hierarchies of computational complexity. This insistence upon the diversity of markomata explains why mathematical expression starts off with the theory of automata, and does not immediately commence with the theory of Turing Machines, as the icon of the maximum degree of computational capacity, as suggested by the Church-Turing Thesis.²² The economic rationale for the distinction is that the theory of Turing Machines ignores limitations of space and time in the process of calculation, whereas the theory of automata immediately takes them into account. Nevertheless, anything that cannot be computed on a Turing Machine (henceforth, ‘Turing non-computable’) will be treated as subsisting outside the realm of science from the vantage point of the theory of markomata.²³ Third, even though the theory of automata serves in the first instance as a taxonomizing device for markomata, we shall argue it also permits

²² The Church-Turing Thesis identifies effectively computable functions with recursive functions, or equivalently with functions computable by Turing Machines. For further explication, see Davis et al, (1994, pp.68-9); and Cotogno (2003).

²³ This includes the mathematical specification of agent maximization over infinite preference sets or continuous utility functions. See Mirowski (2002, pp.427-435) and Prasad (2004). By implication, this rules out any welfare appeals to Pareto optimality as well.

the postulation of certain abstract theoretical generalizations about the market system in its totality. It is there we will begin to enunciate ‘laws of economics’ that, admittedly, bear little resemblance to the laws enunciated by neoclassical theory, but nonetheless display a marked kinship with the sorts of developments in the five movements discussed above in Section 2.

Finally, some economists have demurred at the following sections: Where is “the model” which this paper propounds? Old habits indeed die hard, even when one is unaware of their provenance. Modern biologists don’t ask for “the model” of evolution any more; nor do computer scientists cite ‘the model’ of ‘the computer’.²⁴ In order to deal with phenomena that are intrinsically diverse and always undergoing metamorphosis, they have renounced the Cold War ambition to find that Bourbakist mother structure to which all scientists within the disciplinary bailiwick must pledge their troth. Since the first commandment of this program is that “Thou shalt not reify The Market,” then readers looking for a canonical model are bound to be disappointed. There are only specific formalisms intended to capture the salient features of specific markets, all couched in the mathematics of the theory of automata. In any event, Section 2 explicitly cites several individual exercises in the formalization of specific markomata that already exist in the economics literature.

4.1] Markomata defined

The most rudimentary description of a market begins with the notion of a finite automaton. A finite automaton \mathfrak{M} defined over an alphabet $\alpha = \{\alpha_1, \dots, \alpha_m\}$ with states $\theta = \{\theta_1, \dots, \theta_n\}$ is given by a function T called a transition function which maps each pair (θ_i, α_j) into a state θ_k ; a subset of states $\Theta = \{\theta_k\}$ called final accepting states causes \mathfrak{M} to halt. A finite automaton can be thought of as an extremely limited computing device with no external memory capacity but a single working tape, which it can read only once. After reading a symbol on the tape, it either accepts or rejects it, depending upon the state that the device is in; it then enters the next state prescribed by the transition function. If

²⁴ For the situation in biology, one might consult Depew & Weber (1995), Kay (2000) and Oyama et al. (2001); for the situation in computer science, see Mahoney (1997).

the transition function F maps an existing state of \mathfrak{M} into more than one state, then it is called a nondeterministic finite automaton (NDF).

Suppose we set out to formalize one function of one simple market as an automaton. In one (arbitrary) initial economic example, the order execution function of a very rudimentary market, such as the posted- or fixed-price market, will be modeled as a nondeterministic finite automaton (Khossainov & Nerode, 2001, pp.40-45). A single unit of the commodity is offered at a single price, where the alphabet concerned is the rational numbers; order execution either matches that number as bid by the purchaser, or is rejected. At this early stage, it is important to note that it is merely the order execution function that is captured by this NDF, and not the entire range of functions potentially performed by any real-world instantiation of the posted-price market. Data dissemination, order routing, clearing, record-keeping, and all the rest might themselves be composed of automata of various degrees of computational capacity; any real-world market is formally characterized by the composition of these component automata; and this begins to reveal the true combinatorial explosion of forms inherent in the theory of markomata.

Even restricting ourselves solely to order matching and execution, the possibilities present in any real-life situation begin to outstrip our capacity to subject them to formal abstraction. Can buyers themselves bid, or only respond to the sellers' ask? Are there multiple buyers/sellers, and can they initiate/respond in real time? Can they react to one another, as well as to the opposing side of the market? Can they communicate through channels other than the order execution algorithm? The explosion is partially mitigated by subjecting markomata to the computational and complexity hierarchies propounded within automata theory. The first, and most important, computational hierarchy is known in computer science as the "Chomsky hierarchy" (Davis et al., 1994, pp.327-9). It relates the complexity of the language recognized to the memory capacity of the class of automata deployed.²⁵ It is summarized for the order execution function in Table I below.

²⁵ More elaborate definitions of each class of automaton can be found in Mirowski (2002, pp. 88-92), Taylor (1998), and of course, Davis et al. (1994).

Table 1: Markomata Hierarchy of order execution

Automaton type	Recognizes lang.	Memory	Markomata
Finite \mathfrak{M}	Regular	None	Posted-price
Pushdown	Context-free	Pushdown stack	Sealed bid
Linear bounded	Context sensitive	Finite tape	Double auction
Turing Machine	Recursively enumerable	Infinite tape	None

One implication of the Chomsky hierarchy is that some problems, which are unsolvable at the lower levels of computational capacity, can be shown to be solvable at the higher levels. Furthermore, there exist some problems that cannot be solved even at the most powerful level of the hierarchy; some strings are Turing non-computable on the Turing Machine. However, the hierarchy is inclusive, in the sense that the more powerful automaton can perform all the calculations of the automaton lower down in the hierarchy, because it can *simulate* the operation of machines of lesser computational capacity. This leads to the important notion of ‘markomata simulation’.

The idea of one markomata simulating the operation of another is quite familiar to market practitioners, even though it has been absent up until now in economic theory. For instance, the futures market for red no.6 wheat ‘simulates’ the spot market for red no.6 wheat, in the sense that it can perform the same operations, augmented by other related operations, in the course of ‘tracking’ the wheat market. Likewise, the dealer-organized wholesale market ‘simulates’ the posted-price markets of the retailer, while superimposing other functions. In an abstract computational sense, the futures market ‘encapsulates’ the model of the spot market within its own algorithms. This would be the case even if the futures markets were operated as a double auction, whereas the spot markets were operated as a sealed-bid auction. The theory of computation informs us that certain specific market forms can simulate other market forms *as long as* they are composed of markomata of greater or equal computational capacity. The reason that the

markomata hierarchy does not collapse down to a single flat uniformity is that more computationally complex markets situated higher in the Chomsky hierarchy perform other functions over and above those performed by the markets that they simulate: for instance, futures markets may seek to arbitrage price discrepancies as well as track the spot markets in their purview. Foreshadowing section 5, it would appear that markomata capable of advanced simulation arise chronologically later than their simpler relatively self-contained predecessors.

Table I above suggests that some forms of automata may be mapped into different formats of order execution familiar from the literatures of experimental economics and market microstructure. While the posted price format possesses no memory capacity and therefore qualifies as a finite automaton, a sealed bid auction requires the comparison of a submitted bid to an ordered array of previously entered bids stored in a memory, and therefore qualifies as one of a number of k-headed pushdown automata (Mirowski, 2002, p.571). Sealed bid order execution requires an ordering of submitted bids, which can be captured by a first-in first-out memory stack: hence the ‘pushdown’. The standard double auction requires even more prodigious memory capacity, given that sequences of bids and asks stored in different identifiable memory locations must be retrieved and compared, and therefore should exhibit the computational capacity of (at least) a linear bounded automaton. Table I also suggests that no extant markomata has the power of a Turing Machine. This will be one of the first empirical predictions of evolutionary computational economics, explained below in Section 4.4.3.

4.2] Resource-constrained markomata

All the considerations in Section 4.1 related the order execution function of markomata to standard Chomsky classes of computational capacity, without any regard for the magnitude of the resources required in order to carry out the computation. If there is an impossibility result for certain classes of problems on a pushdown automaton,²⁶ then

²⁶ See, for instance, (Davis et al, 1994, p.300), where it is shown that pushdown automata cannot determine whether a grammar is ‘ambiguous’. These kinds of distinctions will become important in the markomata-based critique of the work of Rust (1997) and others who deny the relevance of Turing computability to economics. This critique will have to be postponed for a venue where familiarity with markomata can be taken for granted to a greater degree than the present situation.

no amount of resource augmentation will enable that automaton to circumvent the obstacle. However, *within* a particular class of automaton, there is a metric along which computational speed can be augmented by greater resources. In particular, there is a body of theory which gauges the tractability of a computation by comparing the growth in the time required (usually in abstract computational ‘steps’ required in a worse-case scenario) given a unit increase of the size of the problem (usually, the size of the input string). This is called the theory of “computational complexity” within computer science (Davis et al, 1994, chap. 15). The theory of complexity is defined relative to a particular automaton, namely, the Turing Machine; but there is nothing in principle that prevents it from being extended down the Chomsky hierarchy to other automata as well. While we omit the details here, problems that grow in computational time as a polynomial function of their input are considered tractable, whereas those that grow as an exponential function of size of input, or those which are designated ‘NP-complete’, are considered intractable.²⁷

Here, the notion of a ‘tradeoff’ between ease of communication or control and the cost of the market coordination, given in such an intangible interpretation in the neoclassical mechanism design literature, is rendered concrete and determinate in the theory of markomata. Allocational mechanisms based upon inherently intractable algorithms (like most linear programming solution algorithms, which are NP-complete) do not prove durable in the world of real markets. Indeed, one of the explanations for the widespread prevalence of certain low-Chomsky capacity markomata (such as the posted-price mechanism), even in the face of their supposed inferiority with regard to welfare or convergence criteria, can be traced to their high degree of computational tractability. Multi-unit multi-commodity markets with strong complementarities scale in a more reliable and cheap manner in posted-price markomata than they do in the double auction, something that might be regarded as one of the more durable lessons of the accumulated history of the spectrum auctions.

²⁷ The designation “NP-complete” is standard terminology for a class of algorithms that grow exponentially in resource requirements as the size of input grows, but whose solution can be verified in polynomial (P) time. Verification of the conjecture that $P \neq NP$ is one of the great outstanding problems of computer science.

Within the theory of markomata, the theory of computational complexity serves to short-circuit a flaw that bedevils most ‘transactions-cost’ explanations of markets vs. hierarchies (or the state or...) in the ‘new’ institutional economics. One cannot appeal to ‘transactions costs’ of routing social coordination through a particular market format without specifying in what form of market the costs were determined in the first place. Suppose that one posited that the pecuniary costs of use of the posted-price framework in a land market were being set elsewhere (say, a registry) in a double auction; the problem of infinite regress rears its head when one inquires what set of prices underpinned the determination that those costs be imputed within the double auction. Most appeals to transactions costs repress this problem of specification of the meta-market that sets the costs and parameters of market operations. The theory of markomata conveniently avoids that confusion by rooting its explanation of resource cost of particular market forms ‘outside’ of the web of existing markets in the theory of computational complexity. Once that lesson is absorbed, then analysis can begin to grasp that markomata are able to exert some forms of social control, while direct state regulation exerts other, only partially commensurate, control functions (Salant, 2000). The politics of this theory will resist the collapse into economics one finds, for instance, in the social choice and public choice schools.

4.3] The teleology of markomata

When one observes the operation of an isolated markomata, it rapidly becomes apparent that the diversity of their operational characteristics puts paid to the commonplace notion that there is a single terminus towards which price and quantity are predestined to converge. If there happened to be pre-existent fixed schedules of demand and supply in the guise of reservation bids and production costs, then indeed the double auction would effectively guarantee arrival at a preset point of price and quantity in an isolated frame, as the work of Gode and Sunder reveals. But those reservations and costs are themselves functions of the operation of other markomata, most of which are not organized as double auctions.²⁸ Thus the system as a whole exhibits no tendency to move

²⁸ In a curious way, this observation merely recapitulates a complaint of the Cowles Commission against adherents of Marshallian partial equilibrium: we cannot “move along” a demand schedule without knowing

towards any ‘equilibrium’ (a term borrowed from physics in any event); rather, individual markomata do serve to achieve very specific local functions and objectives, often discussed in the experimental and microstructure literatures. The Dutch or descending clock auction promotes the clearing of a market in a fixed specific time frame. The posted price market reduces personal interaction in the marketplace to a relative minimum. Dealer-mediated markets often provide liquidity to a target clientele. The computerized limit order book provides a public record in real time in the form of an accessible order book. The double auction market helps reduce the immediate opportunities for profitable arbitrage of the commodities sold. The sealed-bid limits the transparency of the identities of prospective buyers to each other. The posted-price market leaves open vast opportunities for arbitrage, but manages to withstand most efforts on the part of buyers to ‘game’ the rules of the market to their own advantage. The roster of objectives served by markomata of differing stripes is effectively limitless; this will serve as an important axiom for the foundations of evolutionary economics in section 5 below.

Because the same physical commodity can be and often is sold through different markomata, sometimes even within the same spatiotemporal coordinates, and as experimental economics reveals, different markomata display different price and quantity profiles, it follows that there can be no such lemma as the ‘law of one price’ in computational economics. This is expanded up in Section 4.4.5 below. It follows that there can exist no ‘law of supply and demand’ at the aggregative level, although for pragmatic purposes it may be thought to exist for certain individual markomata. If there might be a universal terminus toward which all automata tend, it is toward their internally defined ‘halting conditions’. But even here, one can easily overstate the predictable mechanical character of market automata. It is a theorem of computational theory that:

There is no algorithm that, given a program in the language $L(\alpha)$ and an input to that program, can determine whether or not the given program will eventually halt on the given input. (Davis et al, 1994, p.68)

The undecidability of the halting problem bears direct relevance for the ambitions of an evolutionary computational economics. The impossibility theorems of computational theory do not belie the construction of specific markomata for attainment

what happens to all other prices. In markomata theory, the space of all these virtual displacements is only partially explored by the ecology of interlinked market forms.

of specific targeted functions (since this is the practitioner's notion of the 'predictability' of the market); they merely prohibit the economist from making any ironclad predictions about the inevitable outcomes of the price system as a whole. As individual markomata become increasingly networked, their computational powers become increasingly complex, and transcendental guarantees that a particular market format will continue to operate as it has done in the past are repeatedly falsified.

In markomata economics, the very notion of 'market failure' thus assumes an entirely different meaning. When a markomata fails, it appears unable to halt. Prices appear to have no floor (or ceiling, in the case of hyperinflation), and the communication/coordination functions of the market break down. Hence there exists the phenomenon of 'circuit-breakers', which make eminent good sense in a computational economics (even as they are disparaged in neoclassical finance theory). Earlier generations of market engineers had apprehended the need for a manual override when there were 'bugs' in the system. And as any software engineer knows, one never entirely banishes all bugs from real-world programs. Markomata, therefore, never can become reified as the apotheosis of rationality.

4.4] Some empirical implications of markomata Theory

Far from subsisting as a mere metaphor or parable, the theory of markomata bears some immediate empirical implications. The theory proves its worth through its consequences, and not by appeal to abstruse techniques or mode of expression. Here we simply list a few implications in order to demonstrate the proposition that the theory of computation does permit law-like statements, even as it acknowledges the irreducible diversity of market forms.

4.4.1] The price system operates Exclusively with rational numbers.

As is well known, there exist an infinity of Turing non-computable numbers contained within the set of the real numbers. Given that most markomata throughout history were comprised of forms that possessed a computational capacity falling well short of a Turing Machine, this would have presented an insuperable barriers to the regular operations of markets, were they deemed to work in terms of the real orthant, an assumption that is nearly ubiquitous in neoclassical economics. Luckily, the historical

record indicates otherwise: every single known market in history has used the rational numbers (or else integers) as their ‘alphabet’. This direct implication of computational theory suggests that markets have always been digital rather than analog, and furthermore, that the discrete character of monetary units is an intrinsic rather than accidental property of market systems. The fascination with the reals in neoclassical economic theory is yet another accidental artifact of its origins in physics, and as such is eminently disposable. Markomata always operate with discrete enumerable alphabets predicated upon the local monetary unit.

4.4.2] Commodity space doesn’t exist.

The existence of multiple diverse automata deployed to buy and sell the same physical ‘thing’ opens up the corresponding possibility that the quantification of commodities need not be unique, nor even isomorphic across markets. The algorithmic construction of how things are sold includes how they are packaged and presented, which in turn is a function of the protocols for recruitment and retention of a clientele. While the drive to impose standardization is certainly one aspect of the growth and spread of markets, there always exists a counter-trend to individuate and differentiate that which may have previously been treated as uniform and homogeneous.²⁹ Hence not only are the units of the commodity persistently in flux, but so too are the ‘attributes’ to which numbers are attached. Therefore, it is not the physical commodity quantities that supply the ‘natural’ numerical backdrop to the market process, but rather the markomata which define what will count as a quantity. This is not a secondary complication which may be safely ignored by economists, but instead a major locus where various participants strive to ‘bend the rules’ in their favor.

This axiom has a number of empirical consequences. Since the dichotomy between a real vs. monetary economy depends crucially upon the independent existence of commodity space, it becomes untenable in markomata theory. Once the obsession with agency is quashed, the location of the elusive real economy ‘behind’ appearances loses its urgency. Index-number theory sheds most of its rationale, as do many exercises in long-term intertemporal comparisons. The formal specifications of markomata are inseparable

²⁹ See the discussion of the rise of different systems of measurement in Mirowski (2004b).

from the local formal definitions of the commodities. Furthermore, this has direct consequences for the topological assumptions to be put to work in models of evolution, as we shall discover in section 5 below.

4.4.3] No markomata possesses the power of a Turing machine

The Turing Machine stands as the most powerful class of computational automata, at least up to the present (Cotogno, 2003). It is distinguished from the weaker classes of automata by its infinite tape, or working space. Since no extant market appears to possess the analogous capacity of an infinite working memory, then it follows directly that no markomata possesses the computational capacity of a Turing Machine.

Some readers might object that no laptop possesses an infinite hard drive either, but that does not stop the Turing Machine from being a useful abstract model for actual physical computers. It is true that computers that come equipped with effectively expandable memories sufficiently approximate the power of a Turing Machine, in that under most circumstances they are capable of emulating the operation of any other machine. Therein lies the reason for the ubiquity of computers in modern life as command/control/ communication devices. Yet there is a good reason not to extend the same dispensation to markomata. The rationale is that if some markomata did possess the computational capacity of a Turing Machine, it would be capable of simulating the operation of any other market, and the way would be open to the collapse of the diversity ecology of markomata to a uniform homogeneity of market form. Since we have never have witnessed such a homogeneous economy in economic history, and do not observe it now (even with the innovation of computer-assisted smart markets), we deduce that no functioning markomata has ever attained the power of a Turing machine.

4.4.4] Commodification of markomata increases system instability

It is interesting to observe how few markomata have themselves been rendered commodities prior to the 20th century. That is not to say that geographic marketplaces were not subject to property regulations, or that merchant traders did not organize themselves into firms, whose equity might pass from hand to hand. Rather, it is to note that the concept of fungibility of the entire algorithmic structure of the market, ranging over the entire set of functions identified in Section 1, was essentially absent until very

recently. Only in the last few decades have there been whole self-governing exchanges, predominantly in the area of finance (Lee, 1998, chap.3), which have ‘demutualized’ and sold shares in their enterprises. There is a circularity inherent in this situation that may be worrisome from a computational vantage point. It has now become unexceptional to list the shares of, say, the Stockholm Stock Exchange on the Stockholm Stock Exchange. In the present terminology, this opens up the possibility for a markomata to ‘sell itself’. In computational theory, this rather resembles the possibility of self-reference in the specification of an algorithm, and evokes the standard setup for Cantor’s diagonalization argument, which is the stock technique for the construction of non-computable numbers (Davis et al, 1994, pp.88-90). Translated into the theory of markomata, this increases the vulnerability to the construction of an undecidable calculation, which stands as the paradigm of ‘market failure’ as defined in Section 4.3 above.

4.4.5] The intractability of complete arbitrage

One attractive feature of the theory of markomata is that it suggests the natural introduction of the theory of graphs and networks into economic theory.³⁰ While automata theory itself can also be expressed through the formalisms of graph theory (Davis et al, 1994, p.240), here we highlight the network conception of different markomata selling the same (or different) commodities being interconnected in the sense that the ‘information outputs’ of one market are treated as ‘inputs’ to another market. Hence the individual markomata are associated with the different vertices V_i of a ‘market graph’ \mathcal{G} , whereas markomata that make use of prices provided by other external markomata will be connected by an edge E_{ij} .³¹ The markomata at vertex V_i may be simulating the operation of the markomata at vertex V_j , as explained in section 4.1 above; or alternatively, it may be simply trading the ‘identical’ commodity in real time. Rather than treat the macro system of markets as one giant integrated automaton (as was the penchant of the mechanism design tradition), we instead opt to approach it as an evolving

³⁰ The first introduction of market graphs into theory conformable with the markomata approach can be dated to Ellerman (1984), and was further discussed in Mirowski (1994). In the neoclassical tradition, it has been traced back even further to the work of Tjalling Koopmans on routing of freighter tankers.

³¹ While the more general setup would specify a ‘directed graph’ with edges which display a specific orientation, as in Ellerman (1984), here we simplify the exposition by only using undirected graphs.

network of individual processors, with an uneven topology of network connections, rather on the model of existing computer networks. Individual markomata at any point in time need not agree upon the price of the commodity (indeed, with futures/spot or wholesale/retail markets, this may be taken as axiomatic), and so therefore define an ‘exchange rate’ along any single edge $w_{ij} = 1/w_{ji}$. Now imagine a fully connected subgraph of \mathcal{G} consisting of the k markomata dealing in the same commodity, called a ‘ k -clique’. Define the k -clique \mathcal{H} as the triple $\{V_i, E_{ij}, w_{ij}\}$, with the number of markomata vertices N and edges $\frac{1}{2}(N-1)N$. Then define $S(l)$ as the set of all possible cycles in \mathcal{H} of length l . The elements of $S(l)$ each define a possible chain of exchanges $w_t[s] = w_{1i} * w_{ij} * w_{jk} \dots w_{z1}$. Clearly, if one finds $w_t[s] \neq 1$, and exceeds some fixed level of transactions costs, then an opportunity exists to realize arbitrage profit.

The search for arbitrage is doubly hampered by (a) the need to identify the set of complete cycles in \mathcal{H} of a given length, and then (b) to find whether $w_t[s]$ exceeds a given positive rational number c and is the maximum within \mathcal{H} . Clearly the magnitude of the set $S(l)$ grows exponentially with length l and markomata N . Consequently, the decision problem of finding arbitrage profits $(w_t[s] - 1) > c$ is reducible to the well-known traveling salesman problem, which is NP-complete (Davis et al, 1994, pp.457-462). This result from the theory of computational complexity does not suggest that the quest for arbitrage profits is impossible (an idea verging on the deranged); only that in worst-case scenarios, the resources required to eliminate arbitrage profit rapidly outstrip any polynomial function of the size of the problem. The lesson that may be drawn from this exercise is that, contrary to Robert Lucas, we may never be fully able to calculate to our complete satisfaction whether or not there are any \$50 bills lying on the sidewalk. The persistence of arbitrage inconsistencies is an endemic fact of life in the theory of markomata.

4.4.6] Small World Networks and the prevalence of fat-tailed distributions in the economy

The graph \mathcal{G} of markomata and their connections defined in section 4.4.5 is of a particularly special sort: it is neither a regular lattice (contrary to neoclassical theory, everything is most emphatically *not* connected to everything else) nor a sparse network characteristic of uncorrelated random graphs. Rather, in general they are “small world networks”, which exhibit much of the local structure of a lattice but possess the average short distances between any two vertices more characteristic of the random graph.³² This is the property of connectedness popularized by the idea that there are no more than ‘six degrees of separation’ between any two people on earth.

Define the ‘degree’ k of a vertex in \mathcal{G} as the numbers of edges that extend out from it. In the graph the mean length (defined as the number of edges that must be traversed) between any two vertices is related to the average degree of the graph $\langle k \rangle$ and its overall size N . For small world networks, the mean length grows much more slowly than it does for any finite-dimensional lattice, and can be approximated by $E(l) \sim \ln N / \ln \langle k \rangle$. This results in graphs with a high degree of clustering, but a small subset of vertices enjoying long distance connections; it also exhibits a distribution of k which possesses ‘fat tails’ and power law dependence. Much of the postwar history of economics has been spent either trying to explain the pervasive appearance of fat-tailed distributions of economic variables such as firm sizes, price changes, income distributions, and so on, or else trying to explain them away as deformed combinations of Gaussian distributions. Here in the theory of markomata networks, the scale-free distribution of many variables arises rather naturally through the interaction of individual markets, because the higher moments of empirical degree distributions diverge in large and growing networks. Rather than arbitrarily superimposing Normal distributions as unexplained ‘noise’ (as in noise traders in the microstructure literature), the fat-tailed distribution of prices can be derived from the architecture of connectivity between markomata.

³² See Watts (1999) and Dorogovtsev & Mendes (2003; 2004). Markomata often are connected in small world topologies because they are frequently restricted by their adaptation to particular local environments (ie. Humans) to only take as inputs information from other markomata similarly so adapted. Everything is not locally connected to everything else in the world of markomata, a point insisted upon by Potts (2000).

5. Evolutionary Computation and Evolutionary Economics

Why has economics not yet become an evolutionary science? A century after Thorstein Veblen posed the question in his own inimitable manner, the question still persists in a poorly understood *demimonde* of heterodoxy. Perhaps the problem is rooted in the history of an economics mired in the conundrums of agency, as suggested in Section 1. Take, for instance, the book with the extremely promising title, *The Evolution of Economic Diversity* (Nicita & Pagano, 2001), a collection of work on ‘evolutionary economics’. Upon perusal, our hopes are immediately dashed by discovering that none of the authors are actually concerned with *economic* diversity; rather, they occupy themselves predominantly with cognitive diversity, studding their models with ‘memes’ or putative mechanisms of inductive learning or routines or even ‘science’. The more perceptive of the lot attempt to discuss technologies or firms, although even they conceive of the evolving entities as somehow persisting outside of the operation of the market. For these economists, agency evolves, but markets apparently never do. Upon further reflection, a second wave of disappointment descends when we realize the mindset prevalent therein is neither one of explanation of economic *diversity*, in that all and sundry seem to believe a good model explains the convergence of an unfortunate situation of initial unexplained diversity (error, confusion, controversy) to a uniform equilibrium. If this attitude had prevailed in biology, then most evolutionary models would have rejoiced in demonstrating that a monoculture was the ultimate *telos* of natural selection, and all life tends towards stasis. Luckily, they do not.

The gross imitation of the mathematical trappings of a particular biological model, or the citation of an animal behavior specialist who happens to have appropriated a neoclassical model for biological purposes, does not suffice to render an entire tradition ‘evolutionary’. Rather, one should strive to attain a clear conception of what precisely it is that constitutes the subject matter of a distinctly evolutionary economics, and what specific phenomena constitutes its *explananda*. This paper argues that the necessary point of departure for any evolutionary economics lies in the phenomenal diversity of market forms, and that the purpose of economics is to explain the ubiquity of change within that diversity.

5.1] The confluence of computation and evolution

The question immediately arises concerning what motives could possibly prompt modern economists to seek to ground an evolutionary orientation in a foundation of explicitly computational mathematics, such as that described in Section 4 above. This is a prudent and legitimate question concerning the future of economics, since biology had managed to nurture the theory of evolution for at least a century without any reference to computational notions, at least up until WWII (Kay, 2000).

One immediate answer that might be tendered is that the last few decades have witnessed an intensified cross-fertilization of biology and computation. Biology has been serving as a major source of inspiration for a number of developments in computational theory (Forbes, 2004), and conversely, the theory of evolution has been increasingly recast as a process of information transmission and retrieval. While this apparent confluence might turn out to be a fortuitous development for the natural sciences, that fact in and of itself still bears no necessary consequences for the outlines of an evolutionary economics. Indeed, there abides a substantial tradition within the economics literature that has argued that Darwinian evolution, while immensely fruitful within biology, has very little bearing upon economic evolution. Rather than becoming embroiled in that particular controversy, it appears there exists an alternative path for justification of the marriage of computation and evolution as the appropriate foundation of an evolutionary economics.

If one accepts that the theory of automata provides a perspicuous language for description of the insights of the multiple traditions coming to acknowledge the pervasive diversity of market formats, such as those that we surveyed in Section 2, then we can justify the appeal of evolutionary themes from the opposite direction. As explained in (Mirowski, 2002, pp. 139-45, 536-9), John von Neumann invented the theory of automata not to provide a pedagogical basis for ‘computer science’ (which is the position it occupies in contemporary textbooks), but rather to formalize a notion of evolution where abstract logical entities, in the process of replicating copies of themselves, might be able to produce offspring of greater computational capacity than they themselves possessed. Von Neumann himself believed this would promote the development of a generic mathematical theory of information processing and communication that would not

depend upon the exact details of biological reproduction or computer architectures, but could provide descriptions of the emergence of novelty out of recombinations of relatively simple modular components. It would reveal how daunting logical problems could be overcome by populations of low complexity algorithms fortified by the power of random variation. This protocol would not be the province of any one science, but was intended to gain purchase throughout the sciences outside of physics. It is this Neumannian vision of an alternative parallel mathematical formalization of (one kind of) evolution that has subsequently inspired the research traditions in evolutionary computation, artificial life, and complex adaptive systems. The resort to automata theory has been thus motivated by perceived failures of older models of evolution predicated either upon mathematical formalisms derived from classical mechanics (eg., deterministic dynamics) or statistical mechanics (eg., Ronald Fisher's population dynamics).³³ Many of these failures can be traced to problematic notions of 'fitness'.

5.1.1] The farrago of fitness

There is now a substantial literature that expresses deep discontent with the mathematical image of evolution as a process of search over an independently defined and given fitness surface. Since this image is essentially isomorphic to the neoclassical posit of a given objective function subject to search for extrema, this literature has direct consequence for commonplace notions of the congruence of optimization with evolution. For instance, it has recently been argued that, "Any attempt to introduce a unitary analogous concept of 'reproductive fitness' into dynamical models as a scalar ordinal, which will explain or predict quantitative changes in the frequency of types, must fail" (Ariew & Lewontin, 2004, p.348). In biology, the attempt to equate fitness with frequency classes of reproduction has served to suppress ecological and demographic details of species that were empirically shown to be critical to understanding the survival and reproduction of demes, not to mention aspects of inter-species interactions. In game theory, 'replicator dynamics' has equally misrepresented the ways in which information processing is not effectively separable from the context in which it is taking place.

³³ Fisher's model, and its perceived failures, is wonderfully covered in Depew & Weber (1995).

Mathematical choices originally justified in the name of tractability often have served to suppress the very aspects of the problem that had caused the inquiry to be situated within the broad purview of evolution in the first place. One role of the computational tradition has been to isolate those aspects of mathematical models that were obstructing truly evolutionary theorizing.

One of the consequences of the rise to prominence within biology of the “information transmission” paradigm of evolution has been the wholesale re-evaluation of the conventional portrait of evolution as a dynamical system traversing an independently-constituted fitness surface. The “traditional theory of ‘dynamical systems’ is not equipped for dealing with constructive processes... it was precisely the elimination of [the transformation of] objects from the formalism that made dynamical systems approaches so successful” (Fontana & Buss, 1996, p.56). When it came to modeling evolution after the fashion of dynamical systems, evolving entities were often treated as black boxes, with variation attributed to some external stochastic process inducing motion on an isotropic phenotype space, with a one-to-one correspondence to a putative additively decomposable genotype, usually motivated by considerations of mathematical tractability. Dissatisfaction with this reduction of change to stasis, especially at the Santa Fe Institute, led in the interim to a fascination with high dimensionality, chaos, determinism indistinguishable from randomness, and other mathematical phenomena all frequently lumped together under the broad tent of ‘complexity theory’. The quest of these researchers was to try and capture real change as the qualitative transformation of entities arising out of quantitative dynamical interactions.

While it has proven much more difficult to abjure all dependence upon mathematical metaphors of motion than anyone had originally imagined, the biologist Walter Fontana and his collaborators have come up with some concrete proposals to explain why the mathematical presuppositions of dynamical systems have presented obstacles to the modeling of biological phenomena neglected by the Fisherian school and propounded by their opponents, the followers of the ‘modern synthesis’, such as punctuated equilibrium, path dependency, irreversibility, and the appearance of real novelty. Briefly, Fontana insists that evolution consists of (at least) two analytical phenomena, selection and development, which must be accorded equal attention in model

construction. Selection can be modeled as motion on a space; but development must take into account the convoluted relationship of phenotypes to genotypes. Conventional treatments of fitness surfaces misconstrue the phenomena because phenotypes cannot be modified directly. The geometry of fitness surfaces “relates phenotypes without taking into account the indirection required to change them, an indirection which runs through the process by which phenotypes arise from genotypes... what is needed is a criterion of *accessibility* of one phenotype from another by means of mutations” (2003, p.13). Hence, mutation has received insufficient appreciation within evolutionary theory because it is better conceived as a structural component of the topology of phenotypic space.

Fontana makes reference to topology in a sense not generally used in the literature on evolutionary computation. His contention is not simply the standard complaint that phenotypes are collapsed to genotypes in most fitness surfaces; it is that, without exception, these surfaces are portrayed as exhibiting a specific topology, that of a metric space. This means that there is presumed to exist a well-defined distance metric between any two points of the space, that every element can be reached from every other element, and that motion is reversible on these spaces, because the relation of “nearness” is presumed symmetrical. The evolutionary modelers rarely devote explicit consideration to the nature of the fitness space, however; mostly they just posit a Euclidean vector space for their dynamical systems as though they were second nature. To acquiesce in this practice essentially means subscribing to the doctrine that space has no built-in biases; that you can always get there from here. No wonder mutation comes to resemble a third wheel or an unnecessary appendage.

Fontana proposes that we replace this practice with the posit of a fitness space which possesses less topological structure than a metric space, but whose structure embodies the developmental constraints which link the genotype to the realized phenotype (Stadler et al, 2001; Fontana, 2003). Formally, he suggests an ‘accessibility pretopology’ based upon formal notions of asymmetry of neighborhoods and nearness. In such a pretopology, France can be considered to be ‘near’ Monaco, since a large proportion of Monaco’s boundary borders on France; but conversely, Monaco cannot be said to be ‘near’ France, since only a tiny fraction of France’s boundary borders Monaco. Consequently, it will be easier to leave Monaco for France than it will be to leave France for Monaco. Translated

back into biological terms, the pretopology of the fitness surface captures the amount of ‘neutral’ genetic mutation that is possible without showing up as phenotypic change, as well as incorporating an index of the extent of epistasis within the system. The implications of such a revision of fitness concepts has direct consequences for the conceptualization of evolution:

A population of replicating and mutating sequences under selection drifts on a neutral network of currently best shapes until it encounters a ‘gateway’ to a network that conveys some advantage or is fitness neutral. That encounter, however, is evidently *not* under the control of selection, for selection cannot distinguish between neutral sequences. While similar to the phenomenon of punctuated equilibrium recognized by Gould and Eldridge in the fossil record of species evolution, punctuation in evolving RNA populations occurs in the *absence* of externalities (such as meteorite impact or abrupt climate change), since it reflects the variational properties of the underlying developmental architecture (here: folding).

Fontana’s fundamental point is that treating evolution purely on the paradigm of a physical dynamical system invests ‘too much’ plasticity in the population and too little in the fitness surface; and if the rate of change of the organism is roughly on a par with the rate of change of the environment, then there is no ‘evolution’ *per se*, only standard optimization. What permits true evolution is a reservoir of variability which is not immediately accessible to ‘search’ or selection, but is generated by principles specific to the structure of the phenomena in question—in molecular biology, it is the principles of DNA-RNA transcription and subsequent RNA folding; at the level of individual organic *bauplan* it could be the developmental constraints or ‘spandrels’ of Lewontin and Gould (1978); at the population level it would be the phenomenon of co-evolution. The devil hides in the details of the very notion of ‘continuity’ built into the posit of the fitness surface: “What determines continuity is not the degree to which a modification is incremental, but the degree to which that modification is easy to achieve by virtue of the mechanism underlying the genotype-phenotype relation” (Stadler et al, 2001). Both development and mutation matter fundamentally to evolution because they govern these principles of ease or difficulty of change, and as such determine the pretopology of the fitness surface. They also help explain why it is frequently impossible to ‘work backwards’ to major evolutionary transitions: over time, the population drifts away from

the critical point of accessibility at which there were major regime changes: novelty itself is context-specific.

Biology cannot be reduced to physics by means of the copying the formalisms of physical dynamics; computation will not be reducible to biology by copying the structural interactions of selection (dynamical systems) and development (genotype-phenotype pretopology) found there. The scientists concerned with each class of phenomenon will only begin to comprehend true change within the ambit of their studies when their models incorporate mathematical presumptions of the most basic sort – primitive notions of distance, nearness, continuity, symmetry, computability, and the like – which they have independent reasons to certify are characteristic of the phenomena which are the subject of their inquiries. The more we become concerned with the “sciences of the artificial”, in Herbert Simon’s telling phrase, the more this dictates that we must take the activities of the scientist more directly into account. Fontana seeks to make this point at an abstract level about scientific research:

When we wish to change the behavior of systems, we often have a spatial metaphor in mind, such as going from ‘here to there’, where ‘here’ and ‘there’ are positions in the space of behaviors. But what exactly is the nature of this space? Who brought it to the party? It is a popular fallacy to assume that the *space* of behaviors is there to begin with. This is a fallacy even when all possible behaviors are known in advance. How does this fallacy arise? When we are given a set of concrete or abstract entities of any kind, we almost always can cook up a way of comparing two such entities, thereby producing a definition of similarity (or distance). A measure of similarity makes those entities hang together naturally in a familiar metric space. The fallacy is to believe the so-constructed space is real. It isn’t, because that measure of similarity is not based on available real-world operations, since we cannot act on behaviors directly. We have only system-editors, we don’t have property-editors. Seen from this operational angle, that which structures the space of behaviors is not the degree of similarity among behaviors but a rather different relation: operational *accessibility* of one behavior from another in terms of system-reconfigurations. This brings the mapping from systems to behaviors into the picture. The structure of behavior-space is then induced by this mapping. It cannot exist independently of it. (2003, p.17)

Here is where the initial foundational connection between computational and evolutionary economics is forged. As section 4.2.2 demonstrated, there is no such thing as commodity space; and from the work of Fontana we can come to realize that the ubiquitous dependence upon the Euclidean metric of commodity space was the primary obstacle to the capture of truly *evolutionary* phenomena, such as the intrinsic

irreversibility of economic activities, the significant role of mutation, the advent of real novelty, and the sustenance of true diversity in market operations. To put it starkly: belief in the myth of The Monolithic Market has been unwittingly predicated upon belief in the existence of an independent homogeneous commodity space, and enforced by the properties of symmetry and invariance embodied in that space. (In a phrase: You could always get there from here, so the vehicle didn't matter.) The traditions described above in Section 2 have woken up to the fact that the myth was empirically dubious; computational economics demonstrates in an analytical fashion why no one had previously noticed that it was nevertheless logically entailed by the 'harmless' mathematical assumptions of neoclassical models. Evolution was neutralized by the assumed symmetry of the space of the neoclassical economy.

5.2] Markomata evolution

All that remains is to flesh out the alternative vision of the evolution of markomata. Following von Neumann's suggestion, computational evolutionary economics would concern itself with the evolution of diverse algorithmic market formats within an environment of irreducibly diverse human beings. To a first approximation, the individual 'species' of markomata would be defined according to the algorithmic specifications in section 4.1; 'genera' would be classified according to taxonomies such as the Chomsky hierarchy of computational capacities. "Selection" occurs through humans promoting the differential use and reproduction of specific markomata in distinct spatiotemporal locations. Just as in biology, selection pressures are exerted at various different levels: at the level of the individual component algorithm (say, alternative protocols of order routing), at the level of the individual markomata (say, the conventional hypermarket posted-price markomata), and at the group level (say, the network of markomata related by their computational ability to simulate each other's operations). Some markomata become established in certain limited areas (such as double auctions in finance) because they are perceived to bundle an array of functions deemed particularly suited to the community constituted to promote certain objectives (futures and options, arbitrage, rapid market clearing) whereas others become more evenly distributed across the landscape (such as posted price) due to their relative robustness to a

wide array of situations. “Mutation” is present when humans try to ‘bend the rules’ or otherwise circumvent the invariant operation of particular Markomata. A small amount of mutation is beneficial for the system, since it is a major source of novelty; however, pervasive mutation in the sense of lawlessness or corruption is ultimately harmful to the continued existence of the deme, if not the entire ecosystem.

Following the lead of Fontana, markomata evolve through an abstract genotype-phenotype space which relates an open-ended roster of human objectives to the existing range of algorithmic components of markomata in existence. This space would only exhibit the continuity properties of a pretopology which would permit extensive homologies (features of algorithms which were conserved despite radical changes in function) and ranges of neutral innovation (alterations in algorithms which did not immediately have an impact upon market function, but which would transform the accessibility of one set of market organizations from another) (Stadler et al, 2001). The space would be structured by the computational asymmetry described in section 4.1 that only markomata of a greater computational capacity would be capable of simulation of the operations of markomata of equal or lesser capacity. The networks described in 4.4.6 would also constrain the pretopology.

Many people have encountered difficulty with the figure/ground reversal implicit in the theory of markomata: it is the markets that evolve in an ecology of human beings, rather than the other way around. To many this appears perverse, a violation of the original mandate of the human sciences, the antithesis of a social theory. That might be a reasonable objection if indeed orthodox economic theory had been concerned to do justice to the diversity of human individuality; but that turns out to be a highly debatable proposition (Mirowski, 2002, chap.7). It might be a more prudent policy to leave the theory of evolution of organisms to the biologists, and the theory of cognitive evolution to the cognitive scientists, and have the economists concentrate upon the thing which they are supposed to know best, namely markets. In this sense, the theory of markomata is more faithful to the precepts of methodological individualism than has been the neoclassical program, in the sense that it cherishes the premise that people really are irreducibly different, and that difference will never be completely subsumed under some supposed ‘laws of economic man’. As one cognitive scientist has written:

Traditional economic theory (invoking the substantive rationality paradigm) succeeds whenever individual choice is strongly constrained by social and institutional scaffolding that has *itself* evolved subject to selective pressure to maximize rewards. Outside such highly constrained settings, genuine individual thought plays a greater role, and the psychological irrationalism of the substantive rationality model takes its toll. (Clark, 1997, p.274)

5.3] some empirical questions concerning markomata evolution

A novel research program often entails the concurrent development of new empirical research protocols. Although it is very early days to discern how this might happen, a few suggestions may not be out of place in the present contest.

Taking a page from the history of biology, it will be very difficult to prosecute an evolutionary theory of markomata without the complementary existence of a *natural history of markets*. In a sense, this was the complaint launched in Section 3 against all the nascent traditions of experimental economics, market microstructure, artificial intelligence, and so forth. One might have thought that this would have been a thriving area of economic history, but observations of Douglas North and others indicate otherwise.³⁴ It will be essentially impossible to build a formal theory of market evolution in the absence of a well-documented phylogeny of market taxa in economic history. (This is another reason to postpone the provision of a benchmark markomata model at this early date.)

Once a phylogeny of markomata becomes available, then it will be possible to begin to measure the quantitative distributions of markomata in particular historical settings. For instance, from casual empiricism it seems fairly evident that in the modern world, markomata of relatively low computational capacity on the Chomsky hierarchy tend to make up the bulk of markets which one encounters in daily life. This would correspond to observations in ecology that the preponderance of biomass is comprised of organisms of the lower taxa. While these sorts of empirical distributions are relatively commonplace in the theory of evolution, they contradict the types of arguments found in orthodox economics: double auctions may be regarded as ‘superior’ by various criteria in

³⁴ If I may be permitted to insert a personal note, one of the reasons I have essentially left the field of economic history after beginning my career in that neighborhood was precisely the absence of any serious interest in an economic history of markets in the contemporary profession.

the market microstructure literature, for instance, but outside of Wall Street, they are actually relatively rare.

Nevertheless, the hierarchy of computational capacities of markomata does appear to suggest that there might just exist an arrow of time inherent in economic history: that is, mutation of markomata eventually induces the appearance of markets of greater computational complexity than their predecessors, as well as more elaborate hierarchies of ‘small worlds’ networks. An empirical evolutionary economics would therefore confront the possibility that as a general trend, markets have attained a higher degree of computational complexity throughout time, even though most individual markomata might still operate at a relatively rudimentary level. ‘Development’ then takes on a richer set of connotations than it has previously done under more monolithic notions of a single market regime. Likewise, “globalization” might be regarded in an altogether different light.

References

Admati, Anat. 1991. “The Informational Role of Prices,” *Journal of Monetary Economics*, (28):347-61.

Ames, Edward. 1983. “Automaton and Group Structures in Certain Economic Adjustment Mechanisms,” *Mathematical Social Sciences*, (6):247-60.

Ariew, Andre & Lewontin, Richard. 2004. “The Confusions of Fitness,” *British Journal for the Philosophy of Science*, (55):347-63.

Arthur, Brian; Durlauf, Steven & Lane, David, eds. 1997. *The Economy as an Evolving Complex System II*. Reading: Addison-Wesley.

- Axtell, Robert. 2002. "Economics as Distributed Computation" Brookings working paper.
- Biais, Bruno; Glosten, Larry & Spratt, Chester. 2003. "Microstructure of Stock Markets," working paper, Carnegie Mellon.
- Biggs, Norman. 1993. *Algebraic Graph Theory*. 2nd ed. New York: Cambridge University Press.
- Binmore, Ken. 2000. "Economic Theory Sometimes Works," ELSE Discussion Paper, University College London.
- Blank, Leland & Tarquin, Anthony. 1976. *Engineering Economy*. New York: McGraw Hill.
- Byde, Andrew. 2004. "Applying Evolutionary Game Theory to Auction Mechanism Design," Hewlett-Packard laboratories.
- Clark, Andy. 1997. "Economic Reason," in J.N. Drobak & J.V. Nye, eds., *The Frontiers of the New Institutional Economics*, pp.269-90, San Diego: Academic Press.
- Cliff, Dave. 2001. "Evolution of a market mechanism through a continuous space of auction types," Hewlett Packard Labs Technical Report HPL-2001-326.
- Conitzer, Vincent & Sandholm, Tuomas. 2002. "Complexity of Market Design"
- Coase, Ronald. 1988. *The Firm, the Market and the Law*. Chicago: University of Chicago Press.
- Cotogno, Paolo. 2003. "Hypercomputation and the Physical Church-Turing Thesis," *British Journal for the Philosophy of Science*, (54):181-223.
- Dash, R.; Jennings, N.; & Parkes, D. 2003. "Computational Mechanism Design: a call to arms," *IEEE Intelligent Systems*, (Nov):40-48.
- Davis, Martin; Sigal, Ron & Weyuker, Elaine. 1994. *Computability, Complexity and Languages*. 2nd ed. San Diego: Morgan Kaufmann.
- Depew, David & Weber, Bruce. 1995. *Darwinism Evolving*. Cambridge: MIT Press.
- Domowitz, Ian. 1993. "A Taxonomy of Automated Trade Execution Systems," *Journal of International Money and Finance*, (12):607-31.
- Domowitz, Ian & Wang, Jianxin. 1994. "Auctions as Algorithms," *Journal of Economic Dynamics and Control*, (18): 29-60.

- Dorogotsev, S. & Mendes, J. 2003. *Evolution of Networks*. Oxford: Oxford University Press.
- Dorogovtsev, S. & Mendes, J. 2004. "The Shortest Path to Complex Networks," in N. Johnson, et al, eds. *Complex Systems and Interdisciplinary Science*. Singapore: World Scientific.
- Edwards, Paul. 2004. "A Vast Machine: Standards as Social Technology," *Science*, (304):827-8.
- Ellerman, David. 1984. "Arbitrage Theory: A Mathematical Introduction," *SIAM Review*, (26):241-261.
- Fontana, Walter. 2003. "The Topology of the Possible," in A. Wimmer & R. Koessler, eds., *Paradigms of Change*.
- Fontana, Walter & Buss, Leo. 1996. "The Barrier of Objects," in John Casti & Anders Karlqvist, eds., *Boundaries and Barriers*. Reading: Perseus.
- Fontana, Walter; Wagner, George & Buss, Leo. 1995. "Beyond Digital Naturalism," in Christopher Langton, ed., *Artificial Life: an Overview*. Cambridge: MIT Press.
- Forbes, Nancy. 2004. *Imitation of Life: how biology is inspiring computing*. Cambridge: MIT Press.
- Friedman, Dan & Rust, John, eds. 1993. *The Double Auction Market*. Reading: Addison-Wesley.
- Gode, Dan & Sunder, Shyam. 1993. "Allocative Efficiency of Markets with Zero-Intelligence Traders," *Journal of Political Economy*, (101):119-137.
- Hurwicz, Leonid. 1960. "Optimality and Informational Efficiency in Resource Allocation Processes," in Kenneth Arrow, et al., eds., *Mathematical Methods in the Social Sciences, 1959*. Stanford: Stanford University Press.
- Hurwicz, Leonid. 1981. "Perspectives in Economics," in G. Horwich & J. Quirk, eds. *Essays in Contemporary Fields of Economics*. West Lafayette: Purdue University Press.
- Hurwicz, Leo. 1986. "On Informationally Decentralized Systems," in C.B. McGuire & Roy Radner, eds., *Decision and Organization*, 2nd ed., Minneapolis: University of Minnesota Press.

- Hurwicz, Leonid. 1994. "Economic Design, Adjustment and Mechanism," *Economic Design*, (1):1-14.
- Jarecki, H. 1976. "Bullion Dealing, Commodity exchange trading and London gold fixing," in Y. Amihud, ed. *Bidding and Auctioning for Procurement and Allocation*. New York: New York University Press.
- Jennings, N.; Sycara, K. & Wooldridge, M. 1998. "A Roadmap of Agent Research and Development," *Autonomous Agents and Multi-agent Systems*, (1):7-38.
- Kaplan, Carl. 2001. "Tough Times for Data Robots," *New York Times Online*, January 12.
- Kay, Lily. 2000. *Who Wrote the Book of Life?* Stanford: Stanford University Press.
- Kephart, Jeffrey. 2002. "Software Agents and the route to the information economy," *Proceedings of the National Academy of Sciences*. (99):7207-13.
- Khoussainov, Bakhadyr & Nerode, Anil. 2001. *Automata Theory and its Applications*. Boston: Birkhauser.
- Krahenen, J. & Weber, M. 2001. "Marketmaking in the Laboratory," *Experimental Economics*, (4):55-85.
- Ledyard, John. 1987. "Incentive Compatibility" in J. Eatwell et al, eds. *Allocation, Information and Markets*. New York: Norton.
- Lee, Kyu-Sang. 2004. *Rationality, Mind and Machine in the Laboratory: A Thematic History of Vernon Smith's Experimental Economics*. Notre Dame Ph.D. thesis.
- Lee, Ruben. 1998. *What is an Exchange?* Oxford: Oxford University Press.
- *Lewontin, Richard & Gould, Stephen. 1978. "The Spandrels of San Marco and the Panglossian Paradigm: A Critique Of The Adaptationist Programme," *Proceedings Of The Royal Society of London, Series B*, Vol. 205, No. 1161, Pp. 581-598.
- MacKenzie, Donald & Millo, Yuval. 2003. "Constructing a Market, Performing Theory: the historical sociology of a financial derivatives exchange," *American Journal of Sociology*, (109): 107-45.
- MacKenzie, Donald. 2005. *An Engine, not a Camera: Finance Theory and the Making of Markets*. Cambridge: MIT Press.

- Madhavan, A. 2000. "Market Microstructure: a survey," *Journal of Financial Markets*, (3):205-258.
- Mahoney, Michael. 1997. "Computer Science," in John Krige & Dominique Pestre, eds., *Science in the 20th Century*. Amsterdam: Harwood.
- Miller, Ross. 1986. "Markets as Logic Programs," pp.129-136 in *Artificial Intelligence in Economics and Management*. Ed. L. Pau. Amsterdam: North Holland.
- Miller, Ross. 1996. "Smart Market Mechanisms: from practice to theory," *Journal of Economic Dynamics and Control*, (20):967-978.
- Miller, Ross. 2002. *Paving Wall Street*. New York: Wiley.
- Mirowski, Philip. 1989. *More heat than Light*. New York: Cambridge University Press.
- Mirowski, Philip. 1994. "Arbitrage, Symmetries and the Social Theory of Value," in Amitava Dutt, ed., *New Directions in Analytical Political Economy*. Aldershot: Edward Elgar.
- Mirowski, Philip. 2002. *Machine Dreams: Economics becomes a cyborg science*. New York: Cambridge University Press.
- Mirowski, Philip. 2004a. "McCorduck's *Machines Who Think* after 25 Years," *AI Magazine*, (24):135-138.
- Mirowski, Philip. 2004b. *Effortless Economy of Science?* Durham: Duke University Press.
- Mount, Kenneth & Reiter, Stanley. 2002. *Computation and Complexity in Economic Behavior and Organization*. New York: Cambridge University Press.
- Nicita, A & Pagano, U. eds. 2001. *The Evolution of Economic Diversity*. London: Routledge.
- Nik-Khah, Edward. 2004. *Designs on the Mechanism*. Notre Dame Ph.D. thesis.
- North, Douglas. 1977. "Markets and other Allocation Systems in History," *Journal of European Economic History*, (6):703-716.
- O'Hara, Maureen. 1995. *Market Microstructure Theory*. Cambridge: Blackwell.
- Oyama, Susan; Griffiths, Paul; & Gray, R. eds. 2001. *Cycles of Contingency*. Cambridge: MIT Press.

Phelps, Steve, et al. 2002. "Co-evolutionary Auction Mechanism Design: a preliminary report," in J. Padget et al, eds., *Agent mediated Electronic Commerce IV*. Berlin: Springer Verlag.

Phelps, Steve, et al. 2004. "Using Genetic Programming to Optimise Pricing Rules for a double auction market,"

Potts, Jason. 2000. *The New Evolutionary Microeconomics*. Cheltenham: Elgar.

Prasad, Kislaya. 2004. "Constructive and Classical Models for Results in Economics and Game Theory," *Metroeconomica*, (55):141-154.

Reiter, Stanley. 1977. "Information and Performance in the (New) Welfare Economics," *American Economic Review*, (67):226-234.

Rosenbaum, Eckehard. 2000. "What is a Market?" *Review of Social Economy*, (58): 455-482.

Rosser, J.B. 1999. "On the Complexities of Complex Economic Dynamics," *Journal of Economic Perspectives*, (13):169-192.

Roth, Alvin. 2002. "Economist as Engineer: Game Theory, Experimentation and Computation as Tools for Design Economics," *Econometrica*, (70):1341-1378.

*Rust, John. 1997. "Dealing with the Complexity of Economic Calculation" Working Paper, Dept. of Economics, Yale University.

Rust, John; Miller, John; & Palmer, Richard. 1993. "Behavior of trading automata in a computerized double auction market," in Daniel Friedman and John Rust, eds., *The Double Auction Market: institutions, theories and evidence*. Reading: Addison-Wesley.

Salant, David. 2000. "Auctions and Regulation," *Journal of Regulatory Economics*, (17):195-204.

Sandholm, Tuomas & Suri, Subhash. 2001. "Market Clearability," Carnegie Mellon working paper

*Schinkel, Maarten. 2001. *Disequilibrium Theory*. Ph.D. thesis, University of Maastricht.

Sent, Esther-Mirjam. 2001. "Sent simulating Simon simulating Scientists," *Studies in the History and Philosophy of Science*, (32):479-500.

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

- Smith, Vernon. 2001. "From Old Issues to New Directions in Experimental Psychology and Economics," *Behavioral and Brain Science*, (24):428-9.
- Stadler, Barbel; Stadler, Peter; Wagner, Gunter; Fontana, Walter. 2001. "The Topology of the Possible: formal spaces underlying patterns of evolutionary change," *Journal of Theoretical Biology*, (213):241-274.
- Subrahmanian, Eswaran & Talukdar, Sarosh. 2003. "Engineering of Markets and Artifacts,"
- Sunder, Shyam. 2004. "Markets as Artifacts," in M. Augier and J. March, eds. *Models of a Man: essays in memory of Herbert A. Simon*. Cambridge: MIT Press.
- Taylor, R.G. 1998. *Models of Computation and Formal Languages*. New York: Oxford University Press.
- Tesfatsion, Leigh. 2003. "Agent-based Computational Economics" *Information Sciences*, (149):263-9.
- Thompson, Grahame. 2004. "Is All the World a Complex Network?" *Economy and Society*, (33):411-424.
- Vulkan, Nir. 2003. *The Economics of E-commerce*. Princeton: Princeton University Press.
- Wagner, Dirk. 2000. "Software Agents take the Internet," *First Monday*, 5:7
- Walia, Vibhu. 2002. *Evolving Market Design*. HP-2002-289, HP Laboratories Bristol.
- Walia, Vibhu; Byde, Andrew & Cliff, Dave. 2003. "Evolving Market Design in Zero-Intelligence Trader Markets," Hewlett Packard Technical Report,
- Walker, Donald. 2001. "A Factual Account of the Functioning of the 19th Century Paris Bourse," *European Journal for the History of Economic Thought*, (8): 186-207.
- Watts, D.J. 1999. *Small Worlds*. Princeton: Princeton University Press.
- Weintraub, E.R. 1991. *Stabilizing Dynamics*. New York: Cambridge University Press.
- Wellman, Michael et al. 2004. "Price Prediction in a Trading Agent Competition," *Journal of Artificial Intelligence Research*, (21):19-36.
- Wellman, Michael. 2004. "Online Marketplaces," in M.P.Singh, ed. *Practical Handbook of Internet Computing*.

Wells, Wyatt. 2000. "Certificates and Computers: the remaking of Wall Street, 1967-71," *Business History Review*, (74):193-235.

Wurman, P.; Walsh, W. & Wellman, M. 1998. "Flexible Double Auctions for Electronic Commerce: theory and implementation," *Decision Support Systems*, (24):17-27.

Wurman, P.; Wellman, M.; & Walsh, W. 2001. "A Parameterization of the Auction Design Space," *Games and Economic Behavior*, (35):304-338.

Zhang, D.; Zhao, K.; Liang, C.; & Huq, G. 2004. "Strategic Trading Agents via Market Modeling," *ACM SIGecon Exchange*, (4):46-55.