
Towards a Not-Too-Rational Macroeconomics

Author(s): Axel Leijonhufvud

Source: *Southern Economic Journal*, Vol. 60, No. 1 (Jul., 1993), pp. 1-13

Published by: Southern Economic Association

Stable URL: <http://www.jstor.org/stable/1059926>

Accessed: 01-12-2015 00:13 UTC

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Southern Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *Southern Economic Journal*.

<http://www.jstor.org>

Towards a Not-Too-Rational Macroeconomics*

AXEL LEIJONHUFVUD

Center for Computable Economics

University of California, Los Angeles

I. Introduction

At UCLA we are establishing a Center for Computable Economics. I have been very much involved in this effort. This may surprise you. That a non-mathematical macro/monetary economist should become an enthusiast for this project is one of those things that “do not compute.” But then, as I will explain, we take a special interest in things that do not compute.

One may come to Computable Economics by many intellectual routes. I will trace my own, not because it makes a particularly edifying story, but because it will tell you what kind of problems I hope we can make progress on by developing the field of Computable Economics. But before we get to that, I had better explain what I mean by Computable Economics.

II. Computable Economics

The computer is now being used in a wide variety of fields to model and to explore the properties of complex dynamic systems. We believe that this approach has a big potential payoff also in economics.

In macroeconomics, to take an example close to my heart, the last ten or fifteen years have seen the almost total abandonment of static in favor of dynamic models. Dynamical systems, however, have to have a very simple structure if one is to obtain closed form solutions. The core of this modern macroliterature consists of representative agent (or social planner) models, where the motion of the entire system is given by the solution to a single optimization problem. It is possible to go a bit beyond the representative agent and introduce, for instance, some asymmetry of information. But analytical methods will not take us very far. The properties of more complex systems can *only* be investigated through computer simulations.

Computable Economics will not only mean the study of more complex systems; it will also bring a different orientation towards the modelling of the elements of those systems, that is towards the representation of individual behavior.

My friend, Daniel Heymann, once remarked that practical men of affairs, if they know any-

*Distinguished Guest Lecture, 1992 Annual Meeting of the Southern Economic Association, Washington, D.C., November 22, 1992. An earlier version of this lecture was given at the 1992 Aalborg University-UCLA Summer School in Computable Economics. My borrowings below from a forthcoming book with Daniel Heymann [15] are evidence of my debt to him. I am particularly grateful to my colleague, Kumaraswamy Velupillai, who has taught me all I know about computability, complexity and related matters. That I do not know more than I do is not his fault!

thing about economics, often distrust it because it seems to describe *the behavior of incredibly smart people in unbelievably simple situations*. Now, non-economists often fail to understand that standard economic theory is useful in a myriad ways, despite its unrealistic assumptions about people's cognitive capabilities, *because* the interaction of ordinary people in markets very often does produce the incredibly smart result. When it does, it can be a convenient short-cut to model the social interaction process as if it was planned (and policed) by a representative agent or social planner possessed of rather superhuman abilities.

The defense of our craft that I have just sketched is in the best UCLA tradition, going back to Armen Alchian's classic "Uncertainty, Evolution, and Economic Theory" [1]. But Alchian was advocating a method very much at variance with the one that dominates macrotheory today, a method

. . . which treats the decisions and criteria dictated by the economic *system* as more important than those made by the individuals in it. By backing away from the trees—the optimization calculus by individual units—we can better discern the forest of impersonal market forces [1, 209].¹

Efficiency, in Alchian's theory, stems less from the *ex ante* rational planning of typical economic agents than from the *ex post* elimination through competition of ill-adapted modes of behavior.²

We might start, then, by asking *how believably simple people cope with incredibly complex situations*. If we knew a bit about that, we could then go on to study the conditions under which market interaction will and will not configure the complex system into that incredibly smart allocational pattern. Because, of course, social interaction does not always produce the perfectly rational result. Sometimes, as James Tobin once said, "the invisible hand is nowhere to be seen." Ordinary people also interact to produce booms and busts in real estate, credit crunches and bank panics, great depressions and hyperinflations—and much other misery besides.

What we should aim for is to model "systems that function pretty well most of the time but sometimes work very badly to coordinate activities" [21; 23].

The true descendants of Alchian in more recent times are Jack Hirshleifer [16], Richard Day [8] and Nelson and Winter [30]. But standard economic theory has not taken this tack. It proceeds instead from the foundational postulate that people are rational in the sense that they will act in the manner most appropriate to the situation. Rationality is postulated here in the service of the explanatory strategy that Latsis [18] has termed "situational determinism". Situational determinism is so called because one seeks to explain or predict the behavior of an individual or organization from the external situation *alone*. Bounded rationality is banished in the hope that so doing will guarantee a unique prediction from the given external situation. Unbounded rationality makes "internal" questions of how decisions are reached uninteresting at best. Situational determinism, therefore, is an important fortress in the boundary defenses of economics. It has allowed the economists, most particularly, to ignore developments in the cognitive sciences. By the same token, if we give it up, the boundary defenses come down. We then have to face the frightening prospect that people in other disciplines may have something to teach us!

1. To the passage quoted, Alchian appends a note advocating ". . . reverting to a Marshallian type of analysis combined with the essentials of Darwinian evolutionary natural selection." Compare the notes on Marshall below.

2. Recent work by Gode and Sunder [11] following Gary Becker [4] rather than Alchian, does away with the "trees" altogether to get a clearer picture of the "forest": they demonstrate that "zero intelligence traders" disciplined by simple budget constraints and interacting under double auction rules will succeed in extracting a very high proportion of the maximum attainable surplus. It is particularly noteworthy that Darwinian evolutionary mechanisms play no part in generating this result.

Situational determinism is implemented through various optimizing techniques, where the only “internal” factor is the criterion function to be optimized. Optimizing models do not allow a natural characterization either of decision-making under incomplete knowledge (ignorance) or of adaptive dynamic processes.

This research program has never lacked external critics. Now it has run into trouble on terms internal to itself. Problems have surfaced both on the level of individual behavior and on that of systemic behavior:

First, certain decisions can be shown to be uncomputable (undecidable). This means, among other things, that individual behavior in such cases cannot be predicted from a description of the external situation alone. (Of course, computable economics will not solve uncomputable problems either—but it will enable us to determine where the frontier between the computable and the uncomputable runs).

Second, dynamic general equilibrium models have been found generally to have multiple equilibria. Thus, uniqueness is not guaranteed even in large numbers (competitive) cases. Even in principle, complete information on the “fundamentals”—on tastes, technologies, and initial resource endowments—will not uniquely determine expectations and so will not suffice to determine a unique time-path for such a system. Thus, “unbounded rationality” will not buy us “situational determinism” after all. The main attraction of this unpalatable assumption is gone, therefore. To my mind, moreover, the multiple equilibria throw doubt on the behavior assumptions that produce them.

Game theorists take to computable methods far more readily than general equilibrium and macroeconomists.³ They are used to dealing with problems where the right equilibrium concept is not obvious and where rationality postulates will not buy you a convincing answer [5; 9]. It is also in this branch of our field that the value of studying complex dynamics on the computer is being most rapidly proven. Even so, it is notable that much of the most exciting work in computable game theory is being done by non-economists.⁴ Another group that should make natural allies (for much the same reasons) are the experimentalists. Computable economics is in effect a brand of experimental economics done with artificially intelligent agents.⁵

Bottoms Up

Where do computers come in, you may ask? So far, I have been trying to persuade you that, in economics, we have gotten the relationship between the system and its elements—that is, between the economy and its individual agents—backwards.

A sideways glance at what is going on in other fields can sometimes help one’s perspective. The field of Artificial Intelligence is in the grips of a controversy between those who advocate a “top down” and those who favor a “bottom up” approach. The “top down” approach relies on the sheer crunching power of a centralized processor eventually to replicate whatever human intelligence can do. The “bottom up” approach, or “distributed AI”, relies on interacting networks of relatively simple processors and attempts to make neural nets evolve that, by parallel processing, will handle tasks far beyond the capacities of the components.

Neoclassical general equilibrium theory is, in these terms, quintessentially “top down”. That is why, in the absence of externalities, it reduces to the optimal solution of a social planner’s

3. See, however, Sargent [35].

4. Some recent examples are the papers by Glance and Huberman [10], Lomborg [28] and Novak and May [31].

5. Compare the views of Holland and Miller [17].

problem. There is little purpose to economists choosing sides and doing battle as if these two approaches were mutually exclusive across the entire discipline. But to get a handle on such ill-coordinated processes as high inflations or deep depressions, for instance, we may do better to view the system from the “bottom up.” I propose the following conjecture:

CONJECTURE. *The economy is best conceived of as a network of interacting processors, each one with less capability to process information than would be required of a central processor set to solve the overall allocation problem for the entire system.*

Developments in computer science promise to be helpful in eventually implementing a research program along such lines. The first generation of parallel computers were programmed to rely on a central processor to “Gosplan” the allocation and scheduling of tasks to the subordinate processors. Recent work at Xerox laboratories [38] has been directed towards realizing truly decentralized distributed processing. Their SPAWN program operates in effect on “market” principles, making the work flow to the processors currently showing the lowest opportunity cost.

Pursuing the parallel computing metaphor would mean that we would have to concern ourselves with the limits to the ability of agents to process information, to manage complexity, and to coordinate their activities in a complex and perhaps unstable environment. The trouble of course is that, from the standpoint of standard theory, this means venturing into rough country where the economist will have to do without many of the accustomed comforts of home.

Economic theory, the way we teach it, consists largely of a hierarchy of more or less standard decision problems, ordered from the two-dimensional blackboard illustrations up to the infinite dimensions of Arrow-Debreu. Much intuition is poured on students at the two-dimensional level to convince them that they choose bread and sausage by equating their marginal rate of substitution to the relative price. Once hooked they are taken by stages, no further intuition supplied, to infinite dimensional problems. Those that lose their faith somewhere on the way, fall prey to the Darwinian laws of course grading.

Ron Heiner [13; 14] makes the observation that, as we ascend this hierarchy of decision-problems in economics, we always upgrade the competence of the imagined decision-maker so that at each stage it is fully adequate to the added complexity. This leaves out of economic theory any and all questions of what behavior to expect when the complexity of the environment increases relative to what the agent can routinely handle. Standard theory implies, in effect, that added complexity is always matched by commensurably more sophisticated decision strategies on part of the typical agent. This implication, Heiner maintains, is *false*.

In his own “theory of reliable interactions” (my term for it), increasing complexity will beyond some point cause agents to simplify their strategies instead. In environments that demand too much of their information gathering and processing capabilities, they will behave conservatively, “take fewer chances”, and use safe thumb-rules of behavior. This may involve ignoring potentially relevant information that a more competent agent would always utilize to optimize. By not trying “to be too smart”, Heiner’s agent restores a *reliable* relationship between actions and their utility-relevant consequences.

The point goes beyond the individual’s ability to precalculate the outcomes of his own actions. The behavioral principle that Heiner models is also his “origin of predictable behavior.” We observe people’s behavior in a low-dimensional space. If they were actually adapting to events at every margin in a much larger space, we might find their behavior in the observed subspace largely unpredictable, perhaps even incomprehensible. The simplified behavior-patterns, the routines and rules, make our actions predictable to each other—as they have to be for complex coordinated processes to be feasible.

Heiner's work demonstrates that it is possible to theorize in a systematic manner about behavior even when agents are not able to cope "optimally" with the full complexity of the system. Leland [27] is another persuasive example.

*Algorithmic Man*⁶

Herbert Simon has long advocated turning away from the *substantive rationality* of economics to the *procedural rationality* common to the other behavioral sciences. Economists have by and large been resistant to the suggestion because it has not been clear to them how procedural rationality is to be implemented in the context of general economic theory. We propose to take the first step from substantive to procedural rationality in a particular way.

The rational economic man of standard theory is supposed to solve decision-problems many of which are *not computationally feasible* ("NP hard") and some of which are *uncomputable* in principle. By insisting on computational feasibility as a criterion that economic behavior descriptions should—as far as possible—satisfy, one does not restrict the cognitive capabilities that agents may be assumed to possess very much. The departure from existing theory is not so drastic as to lose contact with well-established results. It imposes a new discipline and leads one naturally to algorithmic representations of both decision rules and learning procedures (including expectations formation). The rule for Computable Economics modelling will be that you may assume as much 'rationality' on the part of decision-makers as you want—as long as you can also specify a corresponding implementable algorithm by which they could make their decisions.

Algorithms are, of course, procedures and the focus on decision procedures of this kind avoids the traditional sharp distinction in economics between "given" preferences and preference formation,⁷ between technology and technological change, between equilibrium and learning or adapting. Historically, economics has evolved trying to generalize static models so as to cope with more or less dynamic problems. The computational approach lends itself naturally to the characterization of dynamic processes, some of which will converge to stationarity if not perturbed.

The paradigm of the *algorithmically rational man* forces one to face up to the irreducible limitations of any "logic machine" model of man, namely, Gödelian undecidability. It means that rational behavior is not always inherently predictable. When it is not, the theorist is not entitled to assume without further ado that agents can make unique inferences about each others' intentions. The uncomputability issue thus directs our attention to the institutional structures and behavioral conventions that emerge in society to make it possible for people to interact with reasonable confidence in the predictability of the results of their actions. In the recent literature, sundry institutions are seen as arrangements to control opportunistic behavior. To the need to curb such excesses of rationality, one should add the need for arrangements that provide an environment simple enough for moderate rationality to suffice.

Studying dynamical systems with the aid of the computer means doing a form of *inductive* mathematics. This is not what we are used to and the foundations are also to be found in branches of mathematics other than those we require of our graduate students today, i.e., in computability theory, complexity theory, and algorithmic information theory [6].

The economic man that populates our models has been created in the image of his creators: he *deduces*. Since this is his nature, his creators have to supply him with all the necessary premises required for the correct deduction of the optimal course of action. Our methodological

6. . . . a few millennia later than Neolithic Man.

7. A more radical computable formulation asserts "the irrelevance of primitive assumptions of preference structures for the analysis of predictable behaviour." [34, 420].

compulsion, on the one hand, to introduce often impossible knowledge assumptions and, on the other, not to admit cognitive limitations both stem from this method of generating predictions of behavior.

Now, in computable economics (the way I envisage it) we start from the recognition that we are dealing with dynamical systems too complex for our own powers of deduction. It is only natural to put the agents with which we populate our toy universes in the same position. Not only must the typical agent cope with an incredibly complex environment armed with only “bounded rationality”, he must also make inferences about that environment on the basis of incomplete information.

In Daniel Heymann’s and my forthcoming book, *High Inflation* [15], we use the following illustration in a discussion about expectations formation:

Learning what behavior to expect from the government is more akin to pattern recognition than it is to sampling from some known distribution. IQ tests often contain questions that ask the respondent to fill in the next few numbers in some numerical sequence: 1, 2, 3, 1, 2, 3, 1, x, y, Mathematicians often express irritation over questions of this type because they do not have uniquely correct answers to be deduced. The next number could be any number, since the simple and obvious pattern might be a component of a more complex one or be a random occurrence within a larger whole that shows no pattern at all. The psychologists have a point nonetheless: the ability to recognize patterns is an essential aspect of human intelligence. It is essential, moreover, exactly because it allows the agent to make sense of incomplete information where no uniquely correct inference is to be drawn.

Agents may identify simple regularities in the behavior of the government only to see them “violated” the next time around. The knowledge that information relevant to the pattern is always incomplete means that agents always have less than complete confidence in the inferences they have drawn and recognize the limitations of their ability to predict future policy action. When an observation is drawn that does not “fit” the previously inferred pattern, the actual pattern is seen to be more complex than anticipated—how much more complex is not to be known. This in turn means that whatever the length of the “string” of past such observations happens to be, the agent comes to realize that it is less informative than he had thought. . . .

“The search for simplicity in the growth of complexity is the exercise of reason” [34, 432]. Rissanen’s theory of stochastic complexity [32; 33] offers a rational approach to IQ tests of the type just referred to. His MDL (*minimum description length*) principle is a criterion for the most economical description of the regularities to be found in “strings” of data. We should not look at Rissanen’s work as “only” providing a new statistical foundation to the econometricians; it also offers theorists a way to populate their models with agents that learn by *induction*.

Many economic decisions are obviously based on induction from incomplete data. Recent work in artificial intelligence on problems of pattern recognition shows how such behavior patterns can be captured. Neural net models are capable of representing decisions based on incomplete data (and on fuzzy logic as well). Neural networks and other artificial intelligence algorithms are opening up new approaches to the crucial problem of expectations formation in economics. John McCall’s recent work [29] maps out a program for developing a new dynamic economics that would draw on recent advances in neuroscience and artificial intelligence in the analysis both of individual behavior and of the complex system.

III. Why Computable? A Personal View

My own path to an interest in Computable Economics starts in the mid-sixties with two interwoven problems, namely (i) the relationship between Keynesian and “Classical” economics and

(ii) the problem of the “microfoundations of macroeconomics.” Let me take the second one first.

The phrase is now dated. It refers to a set of problems much debated in the years around 1970, but not today. If they are now taken as settled—the settlement bearing the label of New Classical Economics—it is in large part because the frame of reference has shifted and the questions are differently understood today.

The central question that was given that label was *how to unify economic theory*? In the 1960s, microtheory showed us a well-ordered universe where competitive markets coordinated optimal individual decisions quite perfectly. Macroeconomics showed us all manner of things going wrong: persistent involuntary unemployment, thrift-paradoxes, accelerator-multiplier systems forever oscillating and incapable of homing in on their equilibrium time-paths—and sometimes even the vision of social betterment that could be had by the simple expedient of printing paper money. Microeconomics was superbly rational; macro was not.

Now some of that 1960s macro was nonsense—dangerous nonsense in so far as people sought to make public policy on that basis. (Today, we can proudly say, our nonsense is very different from what it was back then!) At that same time, however, my generation took it for granted, I think, that aggregate outcomes of social interaction could easily be “bad.” The tension between individual purposive rationality and episodic purposeless irrationality at the level of the system was something that social theory had somehow to resolve. It could not be ignored or defined away.

The tension was to be resolved, I thought, by modelling systems where individual agents “did the best they could” but where things could and sometimes did “go wrong” in the coordination of their activities. (To resolve it by postulating that the system behaved as a single optimizing agent would have seemed nonsensical to us then—and to show you how little I have learned, I am pretty much of the same opinion today).

The task could not be to “rationalize” all of received macrotheory along such lines. Much of it could not be trusted. The talk about “microfoundations” began, I believe, with the simple idea to use the simplest and empirically most robust elements of ordinary price theory—“demand curves slope downward”—to sort through macrotheory and see what pieces could be relied on.

The image from price theory was of a complex system where, as I said, agents did the best they could and “market mechanisms” forced them to obey “the law of supply and demand” and thus coordinate activities. Keynesian macroeconomics as of 1960 had various ways of wreaking havoc with this by proposing, for instance, that people did not respond to price incentives in all dimensions of commodity space or that certain prices did not obey the law of supply and demand.

In a book of 25 years ago, I tried to challenge all that and I did so by trying to enlist Keynes on my side, arguing that “Keynesian economics” had taken the wrong track away from where it had begun. So I gave an interpretation of the Economics of Keynes where (i) agents do their best to maximize utility or profit, (ii) price incentives are always effective; (iii) all prices do respond to market forces, (iv) and, in principle, a coordinated equilibrium exists.

We should think, I suggested, of a system that differed from the GE model only in that agents do not know the equilibrium price vector to start with but have to find it. *The economy should be looked at as a machine that has to “compute” the equilibrium.*

Now, there existed a “story,” sketched by Walras, of how an economic system could solve the system of interdependent demand and supply equations by iteration in prices. I personified this algorithm as “Walras’s auctioneer” in an attempt to make it concrete to readers and make them realize that *it would not in general work as presupposed*. Following up on the original insight of Robert Clower [7], I argued that Keynes’s theory implied that this procedure for finding the equilibrium price vector could fail—and that this was his “revolutionary” contribution.

In the late 1960s, of course, no one had heard of “parallel” or “distributed” as opposed to centralized processing. But today we can look at the matter in those terms. Then, the argument that “there is no Walrasian auctioneer” says, in effect, that there is no central processor programmed to solve the problem. The computations to determine the allocation of resources are made in parallel at two hierarchically separate levels: by individual agents and by markets. The array of markets runs algorithms that iterate on the basis of *effective excess demands*, not “notional” excess demands. The elements of the vector of effective excess market demands do not always have the same sign throughout as the corresponding “notional” elements. When the signs differ on some elements, the “parallel computer” will end up with a different answer—a *Keynesian* answer—from that of the hypothetical “centralized computer.”

What Kind of Theory?

Rather soon thereafter, I came to realize [21] that by my interpretation of Keynes, the system would fail much too easily and too often to find a reasonably coordinated state—and that a system that bad at self-regulation was not a very plausible product of Darwinian evolution. The income-constrained process had spending too tightly geared to current income. Reconsidering the Keynesian model with buffer stocks of liquid assets added yields a more plausible picture: The system now has a “corridor” of stability and it is only for displacements that take it outside the corridor that it will exhibit serious effective demand failures.

This system will *work quite well under most conditions but will work badly under some (extreme) conditions*. Qualitatively, this is as it should be, I think. All other self-regulating systems that we know in nature or from engineering have *bounded homeostatic capabilities*. Macroeconomists ought to wed themselves to models that work “for richer, for poorer, in sickness or in health.” What we have inherited, however, are models of just two types of systems—those that work without fail, and those that always fail to work.⁸

Since that time, I have been particularly interested in “out-of bounds” behavior, i.e., in what happens in economies under extreme conditions: hyperinflations, great depressions, transformation from socialism. In a non-experimental field, one should pay particular attention to observational ‘outliers’ and economics is still very largely a non-experimental field. But such extreme episodes can teach us better than anything else *how crucial are the things we take for granted*.

The Theory of Markets—and Marshall

My book on Keynes ended with a long section, influenced by reading W. Ross Ashby and Norbert Wiener at the time, advocating a “cybernetic” approach to macro. By this I meant systems built up from components governed by negative feedback controls. The components, of course, would be the system’s markets.

A little-known paper by Richard Goodwin [12] contains the suggestion that we are now pursuing: “. . . it seems entirely permissible to regard the motion of an economy as a process of computing answers to the problems posed to it.” In this paper, Goodwin models a market as a single servo-mechanism regulating price by an iterative procedure.

8. The most recent multiple equilibrium models are adopting a somewhat similar perspective. Now, however, the context is growth rather than coordination and the contrasts drawn are between rapid growth and poverty-traps rather than between equilibrium and far-from-equilibrium states.

I built a little model of a single product market with separate controls for output rate and price [20; 22]. Masanao Aoki [2] later investigated it in some detail. The Walrasian story iterated only in the price. But the market for a produced good has to have two servo-mechanisms—one regulating price, the other output. Such a Marshallian market might be represented by two differential equations: (1) the rate of change of prices as a function of excess demand, and (2) the rate of change of output as a function of excess supply price. It is obvious that this little system of two *coupled oscillators* may very easily generate complex dynamics. Marshall tamed the potential chaos, which he did not have the mathematical tools to handle, with the assumptions about relative adjustment speeds underlying his market-day, short-run and long-run period analysis. Those assumptions constrained his market process to go to a stationary attractor. But I do not think he trusted those assumptions. When I realized what the dynamics of that model were like, I thought I understood Marshall's (in)famous distrust of the mathematics of comparative statics.

In the early 1970s I was quite fascinated by my (subjective) discovery that while Marshall's economics may be neoclassical—indeed, it has the best claim to that label if we take it very literally—it *is not based on choice theory* and that, therefore, it might offer an escape from the teleological statics of general equilibrium theory into a brand of process analysis that macroeconomics could utilize. Marshall's agents do not pick optimal points *ex ante* from given opportunity sets. Instead, they obey simple feedback-based *decision rules* in less than completely known environments. His producers increase output if their supply price is below market price; his consumers increase consumption if their demand price exceeds the market price.

Recently, I have come back to my [22] treatment of Marshall's consumer in connection with Heymann's and my work on high inflations. I want to use it here because Marshall's theory of demand makes a perfect illustration of computable economic theory with boundedly rational agents.

In standard neoclassical economics, an agent's activity vector is the result of just one single choice. The Slutsky consumer chooses a basket of n consumer goods; if the economist puts him in a temporal context, he will obediently choose an nT -dimensional timepath; if faced with the uncertainty of c possible states of nature per time-period, he will unerringly precalculate his contingent paths through this nTc -dimensional jungle, unfazed by the multiplication of margins at which first and second-order conditions have to be checked. In each instance, there is just a single decision.⁹

If Slutsky demand theory were more concerned with actual consumer behavior, it might be considered a weakness that it has trouble explaining “shopping.” Alfred Marshall's demand theory may have other weaknesses, but his *computable consumer* can go shopping—making up her mind on what to buy as she goes along. Marshall's consumer is able to break down her Slutsky cousin's horrendous decision problem into a sequence of manageable pieces. Her main trick is knowing the marginal utility of money, but it helps a lot that she has a cardinal, additive utility function.

The conceptual experiment goes as follows. The consumer receives her income as periodic money payments. Into the present period she carries with her the memory of the “final” utility of the last shilling spent in the latest pay-period. This historical magnitude she treats as a constant.

9. This modern conception of rational choice may be contrasted, for instance, to that of Pareto who was willing to postulate a rationality in the economic realm that he did not expect elsewhere because he conceived of economic behavior as consisting largely of *frequently repeated actions* the outcomes of which could be checked against intentions in a direct and reliable manner. Pareto's agents *learn to be rational* through trial-and-error. They will fit perfectly into Alchian's evolving world. The modern Slutskyite's consumption choice is not such a repeated action—it is a once-in-a-lifetime choice.

She has a marginal utility function for each separate consumer good. This marginal utility, however, is a subjective magnitude that cannot be meaningfully communicated to others. To express demand for the good it has to be made subject to the “measuring rod of money”. The number of utils anticipated from consuming the third package of tea divided by the constant number of utils attached to a shilling equals the number of shillings she is willing to pay for a third package of tea. The consumer’s decision-rule is: If the demand-price for a good, calculated thusly, exceeds the market price: Buy! Knowing “the value of money” (to herself, in utils), she will thus be able to weigh it against market prices of goods in simple pairwise comparisons, make consecutive “shopping” decisions, and still end up at the optimal point on her n -dimensional budget-constraint.

A rule-of-thumb decision strategy of this sort spares one’s mental health from contemplating all marginal rates of substitution in n -space before buying a cup of coffee. It is in the nature of rules-of-thumb, however, that they may introduce error. Using the rule, a Marshallian consumer can set out to spend her money without knowing beforehand all the goods that may be offered for sale, or all the prices, or even her own tastes in all dimensions of the commodity space. Each area of ignorance, naturally, becomes a potential source of allocational error.

Suppose, in particular, that the consumer spends a major portion of her income on commodities whose prices have not changed before discovering—too late—that other goods, which she had been used to consuming, have become significantly more expensive. This will mean that, when the budget is exhausted, various marginal rates of substitution are out of line with relative prices. Such a discovery teaches her in effect that she has been calculating her demand-prices with the help of a measuring rod that has “stretched”—i.e., on the basis of a subjective evaluation of what money is worth that, though learned in the past, no longer applies. She must then “recalibrate” the marginal utility of money, as best she can, before setting out to spend next month’s wages, and so on. In times of rapid inflation, therefore, this behavior pattern relies on learning of a type subject to higher than average depreciation—and this depreciation is one of the social costs of inflation.

Marshall’s “biological vision” was of a complex system in real time consisting of innumerable agents, all subject to birth-and-death processes, constantly adapting by gradient procedures towards equilibria that keep shifting. His mode of theory construction has its own limitations and deficiencies of which I would single out two:

i) His basic supply-and-demand apparatus operates totally by feedback, i.e., by adapting based on evaluating the results of behavior in the immediate *past*. It is completely backward-looking. If modern theory has gone off the teleological deep end, Marshallian theory is entirely too evasive about expectations. It isn’t much use if the subject is asset markets, for instance. With the help of Masanao Aoki, I worked out a Marshallian model of investment in which forward-looking expectations are revised on the basis of feedback [3]. If the lag in feedback is minimal and Darwinian pressures strong, the expectations of the “surviving fittest” will converge on rational expectations.

ii) Simple “cellular automata” agents constantly crawling up their profit- or utility-mountains will not do well if their habitat is full of non-convexities. Increasing returns in production means indivisibilities are present and gradient rules will not help you cope with integer constraints. Marshall did not succeed in providing a simple, convincing description of the behavior of firms operating under increasing returns matching the behavior rule followed by his increasing marginal cost producers. His theory of producers’ behavior, Sir John Hicks used to maintain, dealt with a world that was already disappearing around the turn of the century. Manufacturing in the 20th

century has been predominantly a story of decreasing cost firms operating as price-setters rather than price-takers.

Of course, a hundred years later, we are doing no better. We are if anything more addicted to models where people only choose quantities and no one decides prices. I can offer myself as a warning example, having dealt with increasing returns and the division of labor in a couple of papers [25; 26]—in which the representative firm's short-run decision problem is not formalized or solved. It is a problem requiring combinatorial optimization methods and finding an effective algorithm that has a plausible behavioral interpretation will be a great challenge for Computable Economics.¹⁰

Optimal Choice versus Adaptive Behavior

As I have explained, I find Marshall interesting as a neoclassical economist who does *not* put his theory on choice-theoretical foundations. At the risk of some exaggeration—a risk I have bravely decided to take—let me suggest that choice theory has been “a snare and a delusion” for us all.

Perhaps, I ought to explain that a bit(?) When the context is social, the choice paradigm needs to be supplemented with a competitive equilibrium or other coordinating assumption so that opportunity sets are clearly defined. When you then extend the application of the paradigm to intertemporal planning, the result is a view of the world as a Clockwork—a machine whose future path is predetermined once wound up at the origin of time. This is a profoundly unsatisfactory image of an economic system and there is no real escape from it. [Many economists do, of course, regard stochastic clockworks as providing an escape, but I do not so regard them]. My disenchantment with this brand of microfoundations—spelled out in my Marshall Lectures in 1974 but never published—was such that, to be frank, I drifted out of the professional mainstream from the mid-70s onwards, as intertemporal optimization became all the rage.

As a final illustration of the difference between the two approaches, let me take a problem from my recent work with Heymann on high inflations. To my mind, the core of the literature on inflation theory portrays a system that remains “much too rational” no matter how high the inflation rate; the main if not the only consequences of inflation are the sundry distortions due to the so-called “inflation tax.” Basically, this view of the matter is in the Clockwork tradition.

Now, one of the things that happen in inflations is that *markets disappear*. In the quite moderate U.S. inflation of the 1970s, the 30-year bond market disappeared and the 30-year fixed rate mortgage market just about disappeared. In High Inflations—by which we mean those where people quote inflation rates per month rather than per annum—the longest maturities for nominal contracts that survive are often shorter than 12 months and sometimes no longer than 6 weeks. The questions are: Why do these intertemporal markets disappear? What difference does it make?

If inflationary policies introduce new uncertainties, Clockwork theory would make us expect new markets to emerge to cope with the added contingencies. But if they disappear instead of multiplying, it is not clear that the Clockwork is impaired thereby. Our standard *incredibly smart agents will most likely find the situation unbelievably simple*: they will just replace equilibrium prices in the missing markets with the rational expectations of the same prices and proceed to draw up their intertemporal consumption plans exactly as if the markets were really there. The question then becomes whether it really matters to them that some of the planned trades cannot be executed already at $t = 0$. I am not suggesting that the answer to that one is a simple

10. For an attack on it, see Scarf [36].

No. There is by now a sizeable technical literature on missing markets, investigating the conditions under which people can or cannot get around imposed constraints not to transact in parts of the commodity space. What is less clear is why such clever people choose to labor under these constraints.

To understand what is actually going on, I strongly believe, one must abandon this entire mode of theory construction and rethink the matter from Alchian's evolutionary perspective. Here *believably simple people face incredible complications* and, finding themselves unable to precalculate the consequences, give up trading in most future markets. New externalities¹¹ appear where price-interaction has withered away. As coordinating mechanisms disappear, imperfect decision-makers no longer face the same Darwinian pressures to adapt. Potential gains from trade are left unexploited. Various inefficient practices survive. Resources fail to find their highest valued uses.¹²

The difference between the two approaches matters. The rationally expectant optimizers of today's standard theory do not need market interaction to teach them how to attain the efficient social outcome. Alchian's imperfect decision-makers do. But an Alchian market-process is not an aggregate of mutually consistent optimal decisions. So, it cannot be modelled in the standard way. But I believe we can do it in the computable way.

11. Actually, some addition to our terminology may be needed here. A "missing market" gives rise to a 'true' externality (I think we can agree) when the reason for its absence is that property rights have not been created in the relevant good. This, however, does not fit the "disappearing" intertemporal markets. In this case price-interaction is currently not taking place (and misallocations can occur due to this lack of communication) but it is known that it will take place in the future. People will do their best to *anticipate* what that price will be once a market makes its appearance. "Quasi-externality" is hardly euphonius—but some such term is needed.

12. In a sense, of course, survival still goes to the fittest—but it is a tautological sense because inflation changes the definition of "fitness" [14].

References

1. Alchian, Armen A., "Uncertainty, Evolution, and Economic Theory," *Journal of Political Economy*, LVIII, 1950, as reprinted in *Readings in Industrial Organization and Public Policy*, edited by R. B. Heflebower and George W. Stocking. Homewood, Illinois: Richard D. Irwin, Inc. 1958.
2. Aoki, Masanao. *Optimal Control and System Theory in Dynamic Economic Analysis*. Amsterdam: North-Holland, 1976.
3. ——— and Axel Leijonhufvud. "The Stock-Flow Analysis of Investment," in *Finance Constraints, Expectations, and Macroeconomics*, edited by Meir Kohn and S. C. Tsiang. Oxford: Clarendon Press, 1988.
4. Becker, Gary S., "Irrational Behavior and Economic Theory," *Journal of Political Economy*, February 1962, 1–13.
5. Binmore, Ken, "Modelling Rational Players." Parts I and II, *Economics and Philosophy*, 1987, 179–214, and 1988, 9–55.
6. Chaitin, Gregory J. *Algorithmic Information Theory*. Cambridge: Cambridge University Press, 1990.
7. Clower, Robert W. "The Keynesian Counter-Revolution," in *The Theory of Interest Rates*, edited by F. H. Hahn and F. Brechling. London: Macmillan, 1965, 103–25.
8. Day, Richard H. "Adaptive Processes and Economic Theory," in *Adaptive Economic Models*, edited by R. H. Day and T. Groves. New York: Academic Press, 1975, 1–38.
9. Friedman, Daniel, "Evolutionary Games in Economics," *Econometrica*, May 1991, 637–66.
10. Glance, N. S. and B. A. Huberman, "Dynamics of Social Dilemmas," *Scientific American*, forthcoming.
11. Gode, Dhananjay K. and Shyam Sunder, "Allocative Efficiency of Markets with Zero Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, February 1993.
12. Goodwin, Richard, "Iteration, Automatic Computers, and Economic Dynamics," *Metroeconomica*, III:1, 1951, 1–7.
13. Heiner, Ronald A., "The Origin of Predictable Behavior," *American Economic Review*, September 1983, 560–595.

14. ———. "Uncertainty, Signal-Detection Experiments, and Modelling Behavior," in *Economics as Process: Essays in the New Institutional Economics*, edited by R. N. Langlois. Cambridge: Cambridge University Press, 1986.
15. Heymann, Daniel and Axel Leijonhufvud. *High Inflation*. Oxford: Oxford University Press, 1993.
16. Hirshleifer, Jack, "Natural Economy versus Political Economy." *Journal of Social and Biological Structures*, 1, 1978, 319–37.
17. Holland, John H. and John H. Miller, "Artificial Adaptive Agents in Economic Theory." *American Economic Review*, May 1991, 365–70.
18. Latsis, Spiro J., "Situational Determinism in Economics." *British Journal for the Philosophy of Science*, 27, 1972, 51–60.
19. Leijonhufvud, Axel. *On Keynesian Economics and the Economics of Keynes*. New York: Oxford University Press, 1968.
20. ———, "Notes on the Theory of Markets." *Intermountain Economic Review*, 1970.
21. ———, "Effective Demand Failures." *Swedish Journal of Economics*, 75, 1973, 27–58, repr. in *idem.*, *Information and Coordination*, New York: Oxford University Press, 1981.
22. ———, "Maximization and Marshall" (1974 Marshall Lectures, Cambridge University), unpublished.
23. ———, "Schools, 'Revolutions', and Research Programmes in Economic Theory," in *Method and Appraisal in Economics*, edited by S. J. Latsis. Cambridge: Cambridge University Press, 1976. Reprinted in *Information and Coordination*, *op.cit.*
24. ———. "Inflation and Economic Performance," in *Money in Crisis*, edited by B. Siegel. San Francisco: Pacific Institute, 1984.
25. ———. "Capitalism and the Factory System", in *Economics and Process: Essays in the New Institutional Economics*, edited by R. N. Langlois. Cambridge: Cambridge University Press, 1986.
26. ———, "Information Costs and the Division of Labor." *International Social Science Journal*, May 1989, 165–76.
27. Leland, Jonathan W. "A Theory of Approximate Expected Utility Maximization." Unpublished.
28. Lomborg, Björn. "The Structure of Solutions in the Iterated Prisoner's Dilemma." University of Copenhagen, Institute of Political Science Working Paper, 1992.
29. McCall, John. "The 'Smithian' Self and Its 'Bayesian' Brain." UCLA Working Paper No. 596, 1990.
30. Nelson, Richard R. and Sidney G. Winter. *An Evolutionary Theory of Economic Change*. Cambridge, Mass.: Harvard University Press, 1982.
31. Novak, Martin A. and Robert M. May, "Evolutionary Games and Spatial Chaos." *Nature*, October 1992.
32. Rissanen, Jorma, "Understanding the 'Go' of It." *IBM Research Magazine*, Winter 1988.
33. ———. "Stochastic Complexity, Information, and Learning." IBM Almaden Research Center, Working Paper 1992.
34. Rustem, Berc and Kumaraswamy Velupillai, "Rationality, Computability, and Complexity." *Journal of Economic Dynamics and Control*, 14, 1990, 419–32.
35. Sargent, Thomas. *Bounded Rationality in Macroeconomics*. Oxford University Press, forthcoming.
36. Scarf, Herbert E. "Mathematical Programming and Economic Theory." Cowles Foundation Discussion Paper No. 930, 1989.
37. Velupillai, Kumaraswamy. "The Computable Approach to Economics." Lecture at the Aalborg University-UCLA Summer School in Computable Economics, 1992.
38. Waldspurger, Carl A., Tad Hogg, Bernardo Huberman, Jeffrey O. Kephart and Scott Stornetta. "SPAWN: A Distributed Computational Economy." Xerox System Sciences Laboratory, Working Paper SSL-89-18, 1990.