Barriers and Bounds to Rationality: Essays on Economic Complexity and Dynamics in Interactive Systems

by Peter S. Albin

Edited with an Introduction by Duncan K. Foley

Preface

by Peter S. Albin

In the social sciences, generally, and in the literature of economics, particularly, "structure" and "complexity" are poorly-defined concepts which are invested with weight, depth and magical significance and presumed to be well understood by the reader. Skimming texts and journals, one sees "complexity" connoting singly or in a melange: "ramified specialization," "intractability," "strategic interaction," "uncontrollability," (or "controllability through intricate maneuver,") "cognitive depth," "cognitive breadth," "non-probabilistic uncertainty," the "obstructions of detail and organization," (or the "challenge" of same). Equal ambiguity surrounds the term "structure". Yet, "complexity" can have precise meaning. It does within the cognitive and informational sciences disciplines — most notably, within automata theory, the study of models of computer architectures, computation and formal languages. The disciplined application of this meaning is transferable and therein is the subject of this book: rigorous analysis of an advanced economy's connective and supportive structures and the informational, evolutionary, and adaptive processes that occur within them. In setting out an automata-theoretic design for the study of complex economic structures, I will play several roles beyond that of a researcher compiling findings from a program of study. The role of text expositor is necessary, since automata-theoretic methods — although now coming into vogue — have not previously been systematized for the economist reader. The need for codification is especially great when it comes to complexity assessment. Here, I act as advocate as well as pedagogue. There are actually a number of candidate technical approaches to complexity measurement and each has distinctive advantages in particular problem areas. I believe, however, that one schema best suits the requirements of economic analysis and I will press strongly for it. This approach, Noam Chomsky's original complexity hierarchy for formal languages as extended to dynamic systems by Steven Wolfram and elaborated in some details by me, encompasses automata as the building blocks of system architecture, data structures, levels of decision-making competence, and dynamics. It disentangles the various connotations of "complexity" alluded to in the previous paragraph and avoids one dangerous intellectual trap. "Complexity" is not so simple a notion that it can be captured by a single index labeled "randomness," "entropy," "size of a computer program," "dimension," "computer time," or any of a number of parameters of a graph. To be sure, the Chomsky-Wolfram synthesis can subsume certain of these indices in certain appropriate contexts. Its value lies, however, in its power to codify the full range of informational properties intrinsic to advanced economies and their dynamics[1]. This position is strenuously maintained throughout the book.

My advocacy of automata-theoretic resolutions to economic complexity puzzles is of long standing and the reader may wish to know the history of this commitment. It originated when I first encountered "The Game of Life," John Conway's recreational version of John von Neumann's system of self-reproducing automata. "Life" went on then for the most part within the emerging raffish subculture of academic computer hackers. For me, "Life" revealed a way to gain hands-on experience with a model system that could literally build for itself the capacity to communicate and calculate while accomplishing astonishing organizational transformations. The configurations built autonomously by "Life" demanded to be seen as structural institutions formed with hidden, but seemingly inevitable logic, from the simplest of building blocks. They expressed in the purest sense attributes of complexity but they were too featureless to model an actual economy. I played for a while and then went back to the source.

Von Neumann's creation (1966), a constructive proof posthumously completed by his associate Arthur Burks, demonstrated the assembly of simple parts within a pre-existing framework of interaction paths to form a coherent whole that possessed essentially unlimited powers of fabrication and computation. The notion of "complexity threshold" was critical in von Neumann's reasoning. In overall size and number of differentiated parts the composite had to exceed scale bounds. If it did not, the system would fail to accomplish one or more stipulated functions: building parts and scaffolding; putting together the replica according to blueprints and instructions; checking (computing) the accuracy of the construction, and equipping the clone with a copy of the construction instructions and blue prints before sending it on its independent (self-reproductive) way. As I saw it, the von Neumann system literally represented factory organization, supporting communications, degrees of hierarchy in control, and roundabout-

[1]The Chomsky-Wolfram schema defines the *qualitative hierarchy* of a system as its position within a four-step hierarchy. In formal computational terms, the steps correspond to (1) computers without memory, (2) computers with memory of fixed size, (3) "fixed program" computers with memories that can expand with time and the demands of a problem, (4) "universal" computers which can replicate the behavior of any fixed-program computer. In the original Chomsky formulation the four complexity levels pertain to the symbol-processing capacities of an artificial or natural language and to the language-recognition capacity of an artifical intelligence. These capacities translate directly into decision-making capacities of economic agents — modeled, for example, as game-playing automata. In terms of dynamics the first three levels correspond to discrete deterministic systems with capacities to generate trajectories equivalent, respectively, to those with limit-point, limit-cycle and strange attractors in continuous dynamic systems. The fourth level corresponds to irregular "structure-changing" dynamics with no strict analogue in continuous dynamic systems but many associations in developmental and adaptive economic systems.

ness in dynamics[2]. The highest von Neumann complexity threshold, that for self-reproduction, seemed to suggest the attainment of a critical frontier in technical development, the capacity to spawn new viable technologies. Subsidiary threshholds formed a declining sequence of subsidiary stages of economic development: systems capable of importing techniques and adapting them to local conditions; systems capable of running imported "turnkey" facilities but without adaptation; and systems unable to support industry.

In our view this line of study constituted a fresh approach to one of the oldest problems in the discipline, the division of labor and its effect on productive efficiency. We opened up the black boxes of control technology and job design at least some of the way to uncover largely unsuspected tradeoffs between productivity growth and short-term labor costs. The terms of these tradeoffs relate to organizational features which can be represented by complexity attributes of computer architectures representing different principles of job design. Our subject was today's automated work floor, but the terms of the problem were defined in the *Wealth of Nations*.

As Adam Smith argued in his famous comparison of the pin factory with the craftsman pin-maker, systematic organization can facilitate enormous productivity gains. Clearly the organization of work in hierarchic or roundabout ways permits maximum power, energy, mechanical advantage and specialized learning to be applied to simple separable operations. The skills of Smith's craftsman represented a type of generalized complexity embodied in the experience and knowledge of an isolated individual. However, physiology and neural capacity place natural limits on what one worker can do at one time and the balance of capacities within a representative individual does not necessarily match the mix of capacities required for efficient production. The factory solution is, in some sense, the substitution of organizational "complexity external to the individual" for natural "complexity within the individual" to obtain the balancing advantages of specialization and rationalization.

It happened that at the time of my immersion in the von Neumann world of self-replicating abstract automata I was also trying to find a way around the many discrepancies between what one observes happening on the factory floor and what is representable by the parameters of a production function[3]. It came to me that with suitable modification of von Neumann's system it would be entirely practical to model production from the bottom up, using as data literal engineering specifications of individual pieces of hardware and playing

[2]My initial view was that von Neumann must have had economic applications in mind when he was originating his cellular automaton model — the connections did seem that striking. However, von Neumann's co-workers during the 1940's and 1950's insisted that his thinking followed quite different lines and that he saw his computational forms as pure mathematical objects with some biological referents (personal and telephone conversations with Morgenstern, Burks, Goldstyne, and Whitman, among others). Absent a metaphysics of subconscious influence, von Neumann's understandings of economic theory did not enter this, his last creation.

[3]This work resulted in a series of papers on complexity and job design, which the interested reader may consult: among them, Albin, 1978, 1982a, 1984, 1985, Albin and Weinberg, 1983, and Albin, Hormozi, Mourgios and Weinberg, 1984.

the decision requirement of the technique against the skills and capacities of the work force. With a computer (itself part of the schema) reproducing decisions of managers at the margins of adjustment, the system could calculate detailed microeconomic reactions to changing local factor availabilities and skill levels, while adjusting to external factor and product prices. Finding another way to write a production function is hardly news. The point of the exercise was in its implications for dynamic projection. The model would have the same intrinsic scale and informational complexities as its real world referent. The attributes played a determining role in the choice of appropriate production techniques and in developmental structural change that depended for its productivity effects upon encounters with system complexity threshholds.

In sum, my associates and I came to view the automata-theoretic methodology as providing a constructive and practically insightful way to capture the informational dimension in economic analysis, i.e., the processes, institutions, infrastructure and logical structure that characterize decision-making and control the knowledge-based subsystems which are the distinguishing characteristics of the contemporary economy.

Fritz Machlup suggested in 1962 that "information and knowledge services" and "information and knowledge machines" be classified as distinct sectors of economic activity. Since 1962 the value-added and employment weights of these activities have increased many times over. These can be no doubt that this development is of great quantitative importance; however, interesting divisions exist on whether the change is of historical qualitative significance and on the distinctiveness — if not, uniqueness — of its embodiment, the computer, as an economic good. On the one hand, battalions of technologists and social commentators assert that economic life and productive activity have changed dramatically in recent years and that linked revolutions in information use, information-oriented occupations, and information-handling technologies are central in ongoing transformations.[4] Economists, however, have been prepared, by and large, to see continuity, incremental adjustment and smoothed processes in their observations of technological change. The computer has put on weight[5] within sectoral, national, and trade accounts; but this growth does not necessarily mean that the computer has played a distinctive causal role — let alone, a revolutionary one. Vision within the discipline of economics sees the similarities with past adoptions of technique rather than a transforming change in regime. The computer can enter the conventional analysis as a faster widget-maker.

There exists, simultaneously, within the policy domain, substantial practitioner disenchantment with the relevance and effectiveness of the standard solutions of applied economics which once appeared to have such power and

[4]The sociologist James Beniger (1986) provides the definitive citations list on the perspective of revolutionary change — a classification and tabulation of diagnoses of information-related structural change in major post-War works describing social, technological and economic transformations.

[5]Economists working in a Schumpeterian tradition (e.g., Freeman, Clark and Soete, 1982; Dosi, 1981; Nelson and Winter, 1982; Best, 1990) know to exempt themselves from this criticism. I will argue subsequently that the modern information technologies have distinctive features even from the perspective of Schumpeterian dynamics.

range. Evocations of policy impasse abound:[6] two are germane here. The first is specific to labor markets. Outside the discipline, computerization is often viewed as a prime cause of stagnating general employment growth and deteriorating job quality. Although the factual basis for the claim is controversial and the causal reasoning of the critics is frequently muddied, the perception of computerization as an economic problem untouched by economists remains. That perception is reinforced in the debates over macroeconomic policy that are closest to determining action. In this economists' domain, technological unemployment and labor-displacing automation are seen most often variously as: (1) transient phenomena, (2) invisible as symptoms, or (3) dismissed as untreatable with the prescriptions at hand.[7] The second domain of policy concern is more conjectural. It involves a hypothesis that transformations of production, finance and consumption — largely around new computation and communications technologies — have promoted new patterns of international dependence and have altered the dynamic response of the domestic economy in ways that invalidate old rules for guidance, stabilization, and control. It may[8] be possible to explain many of these changes, retrospectively — even to capture them as parameter variations in standard formulations. But, repeated failures and omissions in prospective analysis suggest a disjunction between the economic system and the economists' system on matters that count.

I happen to be an economist who believes that transformations in informational technology and information use have affected the functioning of the economy in determining ways. I also share in a view that many of the analytical tools used by economists are poorly adapted to the study of informational phenomena and that some are intellectual blinders. The criticism applies directly to descriptive formulations but extends to the theoretical core of the discipline itself. A fundamental flaw of the dominant scholarly paradigm is its mechanistic treatment of information-handling, signalling, decision-making and knowledge institutions.[9] One is equipped with tools to measure resources committed to information and cognitive activities but not their systemic effects. One can footnote a market study with the observation that increased interdependency and shortened lags have increased the complexity of the decision tasks faced by agents but can say little more of substance since consensus on a rigorous

[6]I will not attempt to sort out the roles and possible influences of energy shocks, deregulation, disasters, multinational competition, monetarist or supply-sider zeal, financial innovation, debt crises and many other factors said to change the rules of the stabilization game.

[7]Economists who engage these and related matters must do so at one remove from the prevailing "non-selective" perspective on policy. Structuralist perspectives appear in works of Leontief and Duchin (1985), Bluestone and Harrison (1982), Freeman (1986), and Piore (1984) among others.

[8]Although the whole must be inferred, elements of the hypothesis can be formalized. For example, if the production of informational machines and their use increases the number of sectors subject to increasing returns the stability properties of the (general-equilibrium) system are dramatically worsened, exacerbating tendencies toward stagflation (Heal, 1986). The "financial-fragility" restrictions on the scope of monetary policy (Minsky, 1986) are particularly severe within a Eurocurrency environment and hardly eased where financial trading is automated. (The preceding sentences were written prior to Oct. 17, 1987.)

[9]This theme is prominent in works drawing on Kornai (1971).

definition of ”complexity” is lacking.

However, I believe that better tools are available. In effect, economics, the analytical discipline of decision and choice has remained peculiarly untouched by advances that have transformed the pragmatic arts of decision-making and information management. This book is about ways in which the economist can draw techniques from the cognitive and computational sciences to advance the analysis of informational phenomena that impact on the functioning of actual economic systems and their theoretical counterparts. Its range is the *informational dimension*: the Machlup activities, their special efficiency and cost properties, their systemic and dynamic implications, and analytical and mathematical forms that can represent them within economic theory.

Exploring the information dimension

The core of my argument is a claim that informational technologies and processes have critical properties which are imperfectly captured by models in the mechanistic analytical tradition but which are well defined when building blocks of the economy are represented as abstract automata. (Real economies contain agents who have adapted to these information technologies, while economists on the whole have not.) Abstract automata, mathematical models of computing devices and informational processes, have obvious referents in literal hardware installations. Additionally, automaton modelling of economic structures can be highly informative when the identified parts and links of an economic unit studied correspond to conduits for information flow, barriers to information flow, decision-making entities, stores of data or knowledge, and technologies devoted to decision making, computation and control. The real power of automaton-based methods as a foundation for economic methodology, however, lies in their formal attributes. These include as elements: resource-dimensioned measures of computational requirements; a hierarchy of rigorously defined complexity levels which serves to classify the information content of economic activities; rich dynamics which cover within a single family of models the widest possible variety of “stylized facts” pertaining to system growth, fluctuation and historic development; and natural representation of system-building from component parts. The essays reprinted here explore these attributes both in the abstract and as they associate with properties of real and model economies.

Model or metaphor

Some delicate questions of exposition and style should be aired at this point. Reasoning which associates a referent system with an image system can range from scientific model-building, through heuristic exploration of a detailed analogy, to intuitive play with loose metaphor. Associative reasoning of any type may contribute to analysis; however, an appeal to an association can only communicate understandings where there is a fund of shared experience with and appreciation of the image system. Certainly we, as economists, understand one another when we speak among ourselves of “equilibrating market processes” even if there is no physical market place, when the goods involved are amorphously defined, and while fluctuations abound in the empirical referent. Common sense, common mathematical training, and a shared background in working with exemplar cases help to bind what others might see as a loose metaphor into a working model appropriate to the context. With their rich connections to pure science, the cognitive disciplines and mathematics, automaton formulations also invite associative translations of all types. Unfortunately, few economists have developed the easy familiarity with the computer — as an object of study, as opposed to an object of use — that would make a demonstration of association or correspondence between economic forms and automata instantly interesting, and meaningful. My tasks, then, are to demonstrate the necessity of pursuing the association at the level of rigorous model building, to provide easy access to the more powerful tools and to develop exemplar analyses.

Formal models of informational activities

In a nutshell, many informational and decision-making processes of economic importance require computation or equivalent data handling and organization. Accordingly, it is a natural step to tie an informational process or function to a reference computing device which has the capacity to handle the process or solve the function. The referencing is bidirectional: given a process, it is appropriate to ask for a measure of the computational resources required to accommodate it: "How elaborate a personal computer (PC) is needed to solve this production-control algorithm in real time?" Or, given computational resources, one can inquire into their capacities: "How large a regression program can be fitted on this PC?"

A microcomputer with restricted memory, however, represents only one level in a hierarchy pertaining to computational capacities or the implied requirements of informational functions. These levels, which are frequently called "indices of system complexity," register threshold combinations of scale, differentiation, and interconnection of computational resources. Although in some instances it is possible to trade off one aspect of complexity for another within a qualitative level, the levels themselves are absolute. The microcomputer has capabilities beyond 12 or 1200 or N calculators wired together. And, as one already suspects, there are qualitative levels beyond a memory-restricted microprocessor with a fixed budget for diskettes.

Complexity

The matter of qualitative complexity level is critical because at the heart of many informational processes of economic interest are transformations and functions — virtual or implied — which encounter complexity thresholds in significant ways. For example, firms in the United States and Japan install the same microprocessors to control the same production equipment but assign decision responsibilities to workers according to different principles. In a formal translation of the observation, abstract automata are used to model the actual combinations of human and machine resources committed to decision making and control in the two locations. Study of the derived forms indicates that one approach to work organization leads to a combination of decision resources that crosses a critical complexity threshold, whereas the other does not. In context, the threshold implies the capacity to perform the "knowledge functions" required for continuous up-grading and reconfiguration of the production system. Thus, an objectively determined complexity threshold registers as a precondition for "learning-by-doing," and "adaptiveness". On this interpretation, the analysis permits strong inferences as to rates of productivity change in the two nations. Astute commentators have drawn similar conclusions from other indicators but only via *ad hoc* means.

In other explorations of the informational dimension, we examine standard expectational and decision-making assumptions by translating the assumptions into implied calculations on the data. The computing procedures implied by the theory associate with formal automata. In some cases complexity thresholds are inescapable: seemingly plausible anticipational assumptions require for their fulfillment volumes of computational resources that are beyond the mass of the physical universe. For example, "rational expectation" assumptions entail each actor in a system having a model of the decision-making apparatus of every other actor. It may also be the case that problems which are seeming well specified are at intrinsic complexity levels that permit no rigorous inferences as to solvability.

The power of complexity classification as a tool of theoretical analysis comes out most dramatically in the study of highly interactive systems modeled as N-person non-negotiated games. In the many-person repeated prisoners' dilemma analyzed in Chapter 6 below, for example, there are a multitude of possible behavioral strategies. Yet the search for a solution is relatively painless. First, one can readily demonstrate that a cooperative solution can not be sustained in any dynamic system formed from player strategies if the overall complexity of the system falls below that of the highest of four complexity classes. This complexity threshold wipes out most candidates including those analogous to "tit-for-tat" and other popular ones that fare well in two-person tournaments. Of the remainder in the sparse highest class an effective cultivator of cooperative behavior turns out to be — gratifyingly so, I may add — the "Game of Life."

"Life" can be implemented by the most simple-minded of bounded-rationality agents to yield payoffs close to those of uniform cooperation and offers self-enforced security against wiseguy defectors. In effect, "Life" players rely on the built-in complexity of a network of social interactions. The robustness of the solution can be demonstrated, but only through simulation. The complexity theory that guides one to an effective strategy bars definitive proof of its optimality.

In these cases and in several others, patterns of organization, interconnection, and differentiation matter as much as does "scale", defined as an absolute measure of hardware and time resources devoted to computation. This means that we can not treat an expansion of informational resources uncritically as "investment" or as a "substitution" for ordinary labor or capital, even though the installation may be viewed as such by the decision makers involved and may have many investment-like or substitution-like qualities. Informational functions carry within themselves the potential for profound change in complexity level — qualitative variation or structural change, as it were — with no apparent variation in economic measures of conventional inputs. To get the analysis right it is necessary to open up the black boxes of technical, institutional, and organizational givens.

Inside some black boxes

As economists we are accustomed to working on one side of a line that separates us from the technologist, the programmer, the engineer and the production manager. Ordinarily we would be quite comfortable with a study that took as givens the terms of a tradeoff between instantaneous labor cost and longer-term productivity change (the main *economic* variables in the work-organization application) and used these data to calculate production, cost and employment impacts, changes in comparative advantage, alterations in dynamic properties, etc. Unfortunately, the necessary tradeoff data are not available in ready-to-use form nor have we suitable qualitative information on function properties. Furthermore, the actors who design computers, program them for production control, alter job assignments, and build communications systems are generally aware of only the most narrow economic implications of their work; so it would be blind faith to expect far-reaching optimal adjustment. In short, if we are to pursue the major implications of the ongoing revolutions in computing technology and work organization, it is necessary by default to model and evaluate these production systems at the level of fine technical detail — at least until we can get a fix on parameters and critical indicators.

I have taken it as an obligation to do more than just identify a gap in our knowledge. The papers reprinted here provide tools with which an economist can take apart an informational system — whether a control facility within an actual firm or micro-organization or a formalism implied by pure theory — and determine its capabilities and effectiveness in performing the economic functions ascribed to it. In most instances sample analyses using the tools are carried to definitive conclusions or to a point where the terms of the problem are sufficiently well defined for conventional analytical tools to be applied.

Comparisons over time between older control systems that required mechanical controls and current systems based on digital electronics lead one to distinguish the earlier era as one in which a clockwork control technique drives a clockwork basic machine and delivers productivity improvements at the same rate as the driven facility. The difference in productivity growth rates between the clockwork and digital control techniques suggests the usefulness of applying unbalanced growth analysis (as in my 1978 book and Baumol, 1967) to the setting to identify the associated trends in product and service costs and the related income distribution.

The most important of the technical black boxes is the one enclosing the logic and architecture of computational devices. At a minimum, it has become necessary for economists to have an intuitive understanding of the complexity properties of informational systems equivalent to the physical intuitions which they bring to analyses using models deriving from the study of mechanical systems and equilibrium processes. The closer look at the theory of computational devices will certainly disclose that representation of the informational dimension of economic processes is an important scientific step. Treating the dimension in terms of ordinary investment or substitution can cloud understanding — an assertion which will come to life as we expose the cost barriers and scale properties of actual information technologies.

Informational perspectives on economic analysis

From the foregoing, the reader can anticipate that some economists will have problems in reconciling the approach here with established perspectives within the discipline. Fortunately, there are a few important predecessor works in the literature which can frame the formal analyses and the message of this book.

Anti-equilibrium

Janos Kornai's *Anti-equilibrium* (1971) stands out in this regard. Kornai's now-classic work develops a broad critique of analyses which rely on the presumption of general equilibrium. In doing so, it provides foundations for a general information-based approach to institutional functioning, the formalization of which leads to significant classification concepts such as the "pressure" and "suction" categories of macroeconomies. The concepts have proven to be operational as in Kornai's many studies of micro behaviors in suction economies (e.g., 1982) and to be eminently transferable, as in Weitzman's investigations of alternative incentive structures (1984). It has passed almost unnoticed that Kornai's main arguments are based explicitly on formal properties of automata as models of information systems (1971, p. 51). Kornai had been working within a community of active computer-science theoreticians, so it is not at all surprising that his work took the line it did. It is a measure of his skill as an expositor that his analyses were accepted largely on grounds of theoretical elegance and institutional realism rather than on arguably stronger[10] grounds of automata-theoretic reasoning from informational foundations.

Kornai's critique defines an essential point for epistemological inquiry. If informational properties are foundational and distinctive in the ways he indicates, then the presumptive basis of general equilibrium is impaired and the

[10] Kornai opposed the institutional fact that enterprises rely on highly differentiated information to the presumption of the "general-equilibrium school" that price-type data suffice for the guidance of economic activity. It might seem that one could accept the Kornai critique for "real" economies and at the same type accept the informational presumptions of the general-equilibrium school as appropriate for a meaningful abstract economy. The argument based on formal properties of automata whittles down the domain of "meaningfulness" to a few trivial cases. Kornai alluded to the complexity properties of automaton representations on several occasions, either directly or through references, and gave an intuitive sketch of the full formal critique. Since the formal argument was not presented mathematically it may have been regarded as a heuristic (or metaphor!) and escaped wider attention.

associated welfare conclusions become ephemera at best. However, if informational properties are not both foundational and distinguishable, then market forces can be relied on to dominate the system — distinctive structural forms become ephemera (or local imperfections). In the anti-equilibrium domain of inquiry, "positive" economic analysis is directed towards uncovering persistent institutional regularities; whereas in the "equilibrium-school" domain positive economics is directed towards identifying parameters of processes conditioned on and supporting general equilibrium — as in[11] Reder's description of the "tight priors" of the Chicago school (1982, pp 11-19). It seems impossible to straddle the fence on this matter; a study of special informational properties is hardly meaningful outside the anti-equilibrium framework and such is the position[12] taken here.

Bounded rationality

The information-oriented institutional perspective, of course, has its genesis in Herbert Simon's work and "bounds on rationality" will appear here as a recurrent theme — recall, for example, the construction in which available decision resources are counterpoised against decision requirements imposed by theory. Simon, however, is a protean figure in all of the senses of the allusion; and to grapple with Simon, the economist, it is necessary to be prepared to grasp him as a computer scientist and cognitive theorist. In terms of another image, he has kept the black boxes of computer technology, software design, and artificial intelligence closed in his economic writing, although full understanding of his research program relies on knowledge of their contents as explicated in his work within the cognitive disciplines. Simon has certainly left the keys in plain sight on his citation pages; but except for a few prominent economists — in particular, Kornai again, Montias (1976), Williamson (1975), Nelson and Winter (1982) – they have not been picked up outside the original Carnegie- Mellon community.

Simon's own "architectural" approach to complexity (Newell and Simon, 1972) finds a common cognitive hierarchy in the ways chessmasters identify pattern in board positions, successful organizations assign coping responsibili-

[11] It is interesting to note that studies of informational processes taken within the anti-equilibrium framework can yield results that might either confirm or invalidate its priors. In the latter case one might reject the approach and fall back to general equilibrium as a limiting case. Within the stacked-deck of the general-equilibrium framework, however, results consistent with the existence of distinct information structures are epiphenomena that can be resisted as "imperfections." The analyses presented in later chapters do not shake any of Kornai's basic arguments for his theoretical critique. In fact, the "positive" results add significant supplementary and confirmatory detail. The automaton approach provides rigorous demonstrations of informational properties which are only alluded to in Kornai's work and the applications extend the scheme to new sets of institutional phenomena.

[12] Given the neat fit with the anti-equilibrium frame, I see no reason to write yet another version of the standard dissident's introduction. Consider as read the invocation of Kuhn and Lakatos and the recital on the impasses and irrelevancies to which standard theory leads. Kornai's critique is sufficiently filled out and extended by those of Nelson and Winter (1982), Simon (1978, 1984), and Elster (1989a, 1989b).

ties, and insightful decision-makers filter information. A guiding model for such hierarchical procedures is the methodology for computational decomposition of large otherwise-intractable systems. Simon is most persuasive when the information presented by the economy for interpretation is inadequate for global analysis but is prestructured so that there is a breakdown into fast processes working within nearly-integral subsystems and slow processes working between subsystems. The more fully these conditions are met, the better a hierarchical approach that separately analyzes subsystems and their linkages can work. Hierarchy also seems to be a natural way for human intelligence and institutions formed by human intelligence to organize their cognitive processes. Even if the technical conditions of near decomposability are not met, the model is still appropriate for a behavioral analysis of actual institutions in realistic frustrating environments such as some bureaucracies. The schema is flexible enough to accommodate the variety of forms studied by Simon, the behaviorist and organization theorist, and yet is tight enough to serve as a foundation for Simon the methodological critic in his dismissal of global optimization and unbounded individual rationality as foundations for pure theory.

There is a surface similarity between Simon's approach and that here: an explicit model for the organization of computations is proposed as a model for social and economic organization. Yet the theme of qualitative complexity levels, so prominent here, is absent in Simon's work. This omission can hardly be an error or an oversight and is best seen as reflecting Simon's judgement that behavioral regularities, e.g., the hierarchical organization of natural cognitive processes, are primary as compared to technical properties of the data. There are epistemological pitfalls along the path I take; but I argue, nevertheless, that the technical properties matter a great deal — particularly in analyses that deal with changing computational technologies. I treat rationality bounds as conditional and variable — endogeneously determined, as it were — and employ different informational architectures at several distinct qualitative complexity levels[13] as templates for the construction of economic models.[14]

[13]It is perhaps premature to bring out the fine points before actually getting into the technical questions. However, any work that proposes a bridge between computer science and the social sciences needs to be positioned *vis a vis* Simon. I see my approach as complementary to his with respect to positive analysis of institutions and reinforcing with respect to his critique of rationality assumptions.

[14]As is well known, the Simon-Kornai perspective is referenced and partially incorporated in many works which are primarily critical including writings by "post-Keynesians" (Eichner, 1991), "institutionalists" (Solo, 1991; Tool, 1986), "behaviorists" (Cyert and March, 1963), "evolutionary Schumpeterians" (Nelson and Winter, 1982), and "rationalists" (Hollis and Nell, 1975). Results deriving from automaton modelling will provide these schools with additional ammunition. However, since my main concern is with applications I have confined critical comments to this chapter and to side notes scattered in the text.

One additional point should be brought up in this context however. A common focus of many critiques of neoclassical and "equilibrium-school" analysis is reliance on mathematical methods which are intrinsically time-reversible and ahistorical. Development, evolution, technical revolution, knowledge accumulation, and qualitative progress are thus extrinsic categories for analysis within the standard paradigms. Computational forms including automata models are intrinsically time-irreversible (time-reversible abstract automata exist, but as special cases) and inherently accommodate such processes as knowledge accumulation. The crossing

Is this economics?

In short, the Simon-Kornai perspective: 1) reduces the system property of actual or potential equilibrium from a prior condition to one of several testable hypotheses that depend upon technical and institutional givens; 2) shifts attention towards organizational behaviors and information usage; 3) calls for explicit specification of information processes and functions; and, 4) calls implicitly for consideration of informational resources.[15]

Is a study in this perspective "Economics"? Obviously not, if the Chicago "tight priors" are taken as definitive. Yet it is economics, in the trivial sense that it represents the "doing of economists" in problem areas that are recognizably economic. I advance narrower and stronger grounds for an affirmative response. Economics is the study of rational choice and the consequences of choice under restrictions imposed by productive technique, a legacy of institutions, and finite resources. By and large the technology of rationality has remained unexamined. This book represents a reconstruction of economic analysis where rationality is subject to the rules and restrictions of informational technology.

My first attempts at simulating structure change in systems with high intrinsic complexity were combined with theoretical studies of such changes and the information properties of finely-detailed complex systems. These studies formed the principal subjects of my 1975 book, *The Analysis of Complex Socioeconomic Systems*. The book stressed the potential power of the methodology for structural representation and the generality and scope of the von Neumann forms when suitably translated. In effect it advocated a constructive approach to economic modeling and it sketched in broad outline a research program which is largely realized in these pages. The specifics of this program involved: (1) Working out the practical details of a method of structural description that would reveal critical complexity threshholds; (2) developing a theory of "structural formations," entities which possess capacities for computation and information exchange but are not optimally designed in any architectural or economic sense; (3) developing an associated theory of decentralized, parallel, quasi-autonomous entities; (4) identifying the practical limits on economic decision-making that stem from intrinsic data complexities; and (5) completing a taxonomy of economic agents and institutions organized according to the intrinsic complexity levels of the functions they perform.

---

of qualitative complexity thresholds as can occur in a dynamic sequence may be interpreted as "development" or "structural change" but this is metaphorical usage and subject to the qualifications noted earlier. My own view is that dynamic automata models are rich enough to illuminate aspects of historical processes but are no substitute for historical and institutional analysis. They provide only partial relief with respect to the critical charge of ahistoricity in pure theory.

[15] Of course there are significant exceptions. Economists who have specified automata as economic units without necessarily endorsing a broader informational perspective include: Ames (1983) on information exchange and signalling; Gottinger (1983) on bounded-rationality organizational design; Chenault and Flueckiger (1983) on adaptation; and Rubinstein (1986) on game-playing surogates; Gottinger (1978) and Piccoli (1973) on decision requirements. While Binmore and Dasgupta (1987) examine unrestricted rationality. Warsh (1984) provides an intuitive description of a "complexity perspective" as an emerging tendency in the discipline.

The theoretical side of the research was augumented and shaped by parallel case analyses, participant-observer studies, and technology reconstructions, all aimed at testing the applicability, practicality and scope of the approach. A critical test ground was the work floor where I sought to model and reproduce the decision-complexity content of an entire installation, covering the full range of choices from the most primitive routine tasks, through intricate operational decisions to the deepest optimizations. I recruited a team with credentials in industrial psychology, industrial sociology, industrial engineering, computer science and job design. Together and separately we ran lathes, prowled construction sites, and coded the content of jobs as we attempted them, as they were performed by experts, and as they were co-ordinated (or not) by managers.

Subsequent structured interviews and informal discussions told us to what extent, how, and in what ways the formal complexity properties we identified entered the considerations of workers and managers. The feedback provided gratifying confirmation that the automata representations we used captured essential features of the choices presented by the technology and the thought processes and heuristics employed to resolve them. These were actually thrilling moments. Unprompted, a shovel operator at a New York city excavation site described the hierarchy of decisions involved in doing his job so as to maintain an efficient flow of work for his satellite loaders, dump trucks and drilling crews. The formalisms and complexity orderings we had identified came to life. A lathe operator, asked an innocuous question about what made for a satisfying day's work, replied with a categorization of significant job "complexities," routine "complications" and how what we have came to call "learning-by-doing" rested on his resolution of the former. In short, complexity thresholds mattered a great deal.

The empirical work and the research program, as I had initially conceptualised it, were essentially completed several years ago. However, an important expository element seemed absent. I lacked a persuasive demonstration of the distinctiveness of high-complexity systems — a first-hand means to communicate to the reader the kind of insight I had gained during many hours spent with "Life" and had reinforced in the recapitulations of job content. I feared that unless that insight were cultivated the reader would not see complexity assessment as a central new development but rather as a side aspect or metaphorical adjunct to yet another technical means of modeling bounded-rationality systems. Short of bundling "Life" software with the book, it appeared that I would have to proselytize with second-hand evocations of privileged visions — a sufficient apostolic basis for founding a religion but not for influencing my colleagues.

Coincidentally, at about that time, I was invited to present a paper at a conference on *Dynamical Systems and Cellular Automata* held in Lumigny, outside Marseilles. There, physicists Norman Packard and Steven Wolfram showed crude early versions of their now-celebrated slides of the evolution of one-dimensional cellular automata. My first impression was that they had simply found effective pictorial means to display results that were implicit in the work of Alvy Smith and others in the immediate post-von Neumann generation of automata theorists. But on reflection, I came to understand how much more they had accomplished. The pictures were the final evidence that tied the complexity hierarchy to an equivalent hierarchy for dynamic systems. The supporting calculations and documentation gave definitive support to their conjectures but the pictures told the story.

At this point I developed the "snapshots of complex systems" which both decorate the present text and form the basis of the substantive chapters on complexity measurement and dynamics. The pictures will, I am sure, cultivate the reader's intuitive understanding of the complexity hierarchy, particulary, as it applies to economic data structures. But one does, of course, still have to read the captions. The new work illuminated many aspects of the original program, especially those that dealt with decentralized systems and the essentially parallel architectures which characterize the fields of interaction of many economic processes. The results could even be dramatic as, for example, in a pictorial record of the misfiring of policy attempts to fine-tune an expectations-driven economy.

This brings us almost up to date. I had planned to complete the compilation of these papers into a polished manuscript during a sabbatical stay at Nuffield College. Yet one loose end of exposition remained. I still lacked a compelling — direct, as opposed to inferential — example of an important economic problem whose resolution required the attainment of the highest complexity threshold. I reshelved my manuscript in my tower room and went downstairs to the computer room to pursue a hunch. I had long felt that complexity properties were the key to what Martin Shubik calls the N-person game with "many" players. My focus was on an iterated version of Thomas Schelling's "multiperson prisoner's dilemma" where "many" equated to the many thousands who could be portrayed through parallel computation. It turned out to be simple enough to prove by standard game-theoretical reasoning that no system in the lower levels of the complexity hierarchy could support an equilibrium other than universal defection or an all-or-nothing trigger strategy of the sorts proposed by Roy Radner and James Friedman. But did a self-enforcing cooperative equilibrium exist at the highest level of system capacity? It took some billions of simulated games to demonstrate a close and robust approximation of cooperative equilibrium by a strategy suitable for a bounded rationality player. The strategy is unique and to my astonished delight it turned out to be the "Game of Life." I am sorry to be the one to turn "Life" from an object of intellectual dalliance to an instrument serving a utilitarian purpose, but such is frequently the obligation of the economist.

With the story of this work told, I can at last pass to the most pleasant part of presenting it — the recollection and compilation of the many courtesies received, the guidance and criticism that helped so much to shape the work.

I owe thanks for many valuable criticisms of this work. Eileen Appelbaum has reviewed most of these pages. Others implicated, if only for side comments that made a difference, include: Ed Ames, Luca Anderlini, Brian Arthur, Michael Bacharach, Christopher Bliss, Steven Brams, Ivan Christin, Richard Day, Robin Cowan, Christophe Deissenberg, Gerald Flueckiger, Duncan Foley, Frank Hahn, Geoff Heal, Oskar Itzinger, Janos Kornai, Robert Kuenne, Cigdem Kurdas, John Miller, James Ramsey, Ariel Rubenstein, Andrew Schotter, Willi Semmler, Avner Shaked, Tom Vietorisz and Chuck Wilson. My co-authors Roger Alcaly, Eileen Appelbaum, Charles Bahn, Duncan Foley, Hans Gottinger, Farokh Hormozi, Ross Koppel, Stergios Mourgos and Art Weinberg will surely recognize their influence and perhaps their thoughts. Al Eichner, Hy Minsky, and Aaron Wildavsky, regrettably, cannot acknowlege my thanks for their encouragment and unshaking support for the entire project. On the technical side I have gained much from discussions of automata forms with Per Bak, Norman Margolus, Norman Packard, Thomas Toffoli and Steven Wolfram. I should like to thank for their particular assistance Bob Tater of Automatrix Inc., the intelligence behind CAM-PC, and our contact at Thinking Machines Corporation, David Ray. Cam-6, a Toffoli and Margolus production, and "cell-systems," written by Charles Platt, were used for many of the graphics displays. Hang Kim is owed special thanks for the support given to automata-based structural methodology in the pages of *Mathematical Social Sciences.* To David Warsh, thanks in general. I would like to acknowledge helpful comments from participants in seminars at the Universities of Oxford and Cambridge, the University of California at Irvine, the State University of New York at Albany, Vassar College, the Jerome Levy Economics Institute and New York University. In addition I should like to mention the contributions of Suzanne Hurley, Ariel Rubenstein, my Seminar Group at CUNY Graduate Center, and the artists who helped with this work.

My debts to institutions are also great. I received financial support from the National Science Foundation; the OECD; the US Congressional Office of Technology Assessment; the State of New Jersey, the PSC-CUNY Faculty Research Fund, the John Jay College graduate program, and the Jerome Levy Economic Institute. Workers, managers, and researchers at IBM (Amsterdam), Phillips Electric (Einthoven), Mitsubishi Heavy Industries (Kobe), Brown and Williamson (Louisville) and Reynault (Boulogne) are owed much gratitude. Special thanks are due President Monty Gettler, lathe operators, and forge teams at Nepco along with site managers and heavy-equipment operators at several Tishman excavation locations. I can only regret that research-design considerations forced anonymity on so many who contributed to this study. As is so often the case, ideas flow in new surroundings; leaves spent at the Institute for Advanced Studies in Vienna, the Medis Institute in Munich, the University of Paris (Pantheon), Nuffield College, the Santa Fe Institute, and the University of Bonn proved this to my satisfaction.

Very particular thanks are due to Pat Albin, without whose practical and emotional support this book could not have been completed, Joanna Koldynska–Pekala, who gave me irreplaceable personal assistance, and Ted Yanow, who provided invaluable help in the preparation of the manuscript.

New York, 1991; revised December, 1996.