

THE ECONOMY AS AN EVOLVING COMPLEX SYSTEM II

EDITED BY

W. Brian Arthur
Steven N. Durlauf
David A. Lane



CRC Press
Taylor & Francis Group

A CHAPMAN & HALL BOOK

A PROCEEDINGS VOLUME IN THE

SANTA FE INSTITUTE STUDIES IN THE SCIENCES OF COMPLEXITY

ABP

**THE ECONOMY AS AN
EVOLVING COMPLEX
SYSTEM II**



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

THE ECONOMY AS AN EVOLVING COMPLEX SYSTEM II

Editors

W. Brian Arthur

Santa Fe Institute
Santa Fe, New Mexico

Steven N. Durlauf

Department of Economics
University of Wisconsin at Madison
Santa Fe Institute
Santa Fe, New Mexico

David A. Lane

University of Modena
Italy

Proceedings Volume XXVII

Santa Fe Institute
Studies in the Sciences of Complexity

Advanced Book Program



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

A CHAPMAN & HALL BOOK

Director of Publications, Santa Fe Institute: Ronda K. Butler-Villa
Production Manager, Santa Fe Institute: Della L. Ulibarri
Publication Assistant, Santa Fe Institute: Marylee Thomson

First published 1997 by Westview Press

Published 2018 by CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

CRC Press is an imprint of the Taylor & Francis Group, an informa business

Copyright © 1997 Taylor & Francis Group LLC

No claim to original U.S. Government works

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>


ISBN 13: 978-0-201-32823-3 (pbk)
ISBN 13: 978-0-201-95988-8 (hbk)

This volume was typeset using T_EXtures on a Macintosh IIsi computer.

About the Santa Fe Institute

The *Santa Fe Institute* (SFI) is a private, independent, multidisciplinary research and education center, founded in 1984. Since its founding, SFI has devoted itself to creating a new kind of scientific research community, pursuing emerging science. Operating as a small, visiting institution, SFI seeks to catalyze new collaborative, multidisciplinary projects that break down the barriers between the traditional disciplines, to spread its ideas and methodologies to other individuals, and to encourage the practical applications of its results.

All titles from the *Santa Fe Institute Studies in the Sciences of Complexity* series will carry this imprint which is based on a Mimbres pottery design (circa A.D. 950–1150), drawn by Betsy Jones. The design was selected because the radiating feathers are evocative of the outreach of the Santa Fe Institute Program to many disciplines and institutions.





Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Santa Fe Institute Series List

Lecture Notes Volumes in the Santa Fe Institute Studies in the Sciences of Complexity

Volume I:

John Hertz, Anders Krogh, Richard G. Palmer, editors: Introduction to the Theory of Neural Networks

Volume II:

Gerard Weisbuch: Complex Systems Dynamics

Volume III:

Wilfred Stein and Francisco J. Varela, editors: Thinking About Biology

Volume IV:

Joshua M. Epstein: Nonlinear Dynamics, Mathematical Biology, and Social Science

Volume V:

H.F. Nijhout, Lynn Nadell, and Daniel L. Stein, editors: Pattern Formation in the Physical and Biological Sciences

Proceedings Volumes in the Santa Fe Institute Studies in the Sciences of Complexity

Volume I:

David Pines, editor: Emerging Synthesis in Science

Volume II:

Alan S. Perelson, editor: Theoretical Immunology, Part One

Volume III:

Alan S. Perelson, editor: Theoretical Immunology, Part Two

Volume IV:

Gary D. Doolen, senior editor: Lattice Gas Methods for Partial Differential Equations

Volume V:

Philip W. Anderson, Kenneth J. Arrow, and David Pines, editors: The Economy as an Evolving Complex System

Volume VI:

Christopher G. Langton, editor: Artificial Life

Volume VII:

George I. Bell and Thomas G. Marr, editors: Computers and DNA

Volume VIII:

Wojciech H. Zurek, editor: Complexity, Entropy, and the Physics of Information

Volume IX:

Alan S. Perelson and Stuart A. Kauffman, editors: Molecular Evolution on Rugged Landscapes: Proteins, RNA, and the Immune System

Volume X:

Christopher G. Langton, Charles Taylor, J. Doynce Farmer and Steen Rasmussen, editors: Artificial Life II

Volume XI:

John A. Hawkins and Murray Gell-Mann, editors: The Evolution of Human Languages

Volume XII:

Martin Casdagli and Stephen Eubank, editors: Nonlinear Modeling and Forecasting

Volume XIII:

Jay E. Mittenthal and Arthur B. Baskin, editors: The Principles of Organizations in Organisms

Volume XIV:

Daniel Friedland and John Rust, editors: The Double Auction Market: Institutions, Theories, and Evidence

Volume XV:

Andreas S. Weigend and Neil A. Gershenfeld, editors: Time Series Prediction: Forecasting the Future and Understanding the Past

Volume XVI:

George J. Gummerman and Murray Gell-Mann, editors: Understanding Complexity in the Prehistoric Southwest

Volume XVII:

Christopher G. Langton, editor: Artificial Life III

Volume XVIII:

Gregory Kramer, editor: Auditory Display: Sonification, Audification, and Auditory Interfaces

Volume XIX: George A. Cowan, David Pines, and David Meltzer, editors: Complexity: Metaphors, Models, and Reality

Volume XX:

David Wolpert, editor: The Mathematics of Generalization

Volume XXI:

P.E. Cladis and P. Palffy-Muhoray, editors: Spatio-Temporal Patterns in Nonequilibrium Complex Systems

Volume XXII:

Harold J. Morowitz and Jerome L. Singer, editors: The Mind, the Brain, and Complex Adaptive Systems

Volume XXIII: Bela Julesz and Ilona Kovacs, editors: Maturational Windows and Adult Cortical Plasticity

Volume XXIV:

Joseph A. Tainter and Bonnie Bagley Tainter, editors: Evolving Complexity and Environmental Risk in the Prehistoric Southwest

Volume XXV:

John B. Rundle, Donald L. Turcotte, and William Klein, editors: Reduction and Predictability of Natural Disasters

Volume XXVI:

Richard K. Belew and Melanie Mitchell, editors: Adaptive Individuals in Evolving Populations

Volume XXVII:

W. Brian Arthur, Steven N. Durlauf, and David A. Lane, editors: The Economy as an Evolving Complex System II

Volume XXVIII:

Gerald Myers, editor: Viral Regulatory Structures and Their Degeneracy

Santa Fe Institute Editorial Board
December 1996

Ronda K. Butler-Villa, *Chair*
Director of Publications, Santa Fe Institute

Prof. W. Brian Arthur
Citibank Professor, Santa Fe Institute

Dr. David K. Campbell
Chair, Department of Physics, University of Illinois

Dr. George A. Cowan
Visiting Scientist, Santa Fe Institute and Senior Fellow Emeritus, Los Alamos
National Laboratory

Prof. Marcus W. Feldman
Director, Institute for Population & Resource Studies, Stanford University

Prof. Murray Gell-Mann
Division of Physics & Astronomy, California Institute of Technology

Dr. Ellen Goldberg
President, Santa Fe Institute

Prof. George J. Gumerman
Center for Archaeological Investigations, Southern Illinois University

Prof. John H. Holland
Department of Psychology, University of Michigan

Dr. Erica Jen
Vice President for Academic Affairs, Santa Fe Institute

Dr. Stuart A. Kauffman
Professor, Santa Fe Institute

Dr. Edward A. Knapp
Visiting Scientist, Santa Fe Institute

Prof. Harold Morowitz
Robinson Professor, George Mason University

Dr. Alan S. Perelson
Theoretical Division, Los Alamos National Laboratory

Prof. David Pines
Department of Physics, University of Illinois

Dr. L. Mike Simmons
700 New Hampshire Avenue, NW, Apartment 616, Washington DC 20037

Dr. Charles F. Stevens
Molecular Neurobiology, The Salk Institute

Prof. Harry L. Swinney
Department of Physics, University of Texas



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Contributors to This Volume

- Anderson, Philip W., Joseph Henry Laboratories of Physics, Badwin Hall,
Princeton University, Princeton, NJ 08544
- Arthur, W. B., Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501
- Blume, Lawrence E., Department of Economics, Uris Hall, Cornell University,
Ithaca, NY 14853
- Brock, William A., Department of Economics, University of Wisconsin at Madison,
Madison, WI 53706
- Darley, V. M., Division of Applied Sciences, Harvard University, Cambridge, MA
02138
- Durlauf, Steven, Department of Economics, University of Wisconsin at Madison,
Madison, WI 53706 and Santa Fe Institute, 1399 Hyde Park Road, Santa
Fe, NM, 87501
- Geanakoplos, John, Cowles Foundation, Yale University, 30 Hillhouse Avenue,
New Haven, CT 06520
- Holland, John H., Department of Computer Science and Engineering, University
of Michigan, Ann Arbor, MI 48109 and Santa Fe Institute, 1399 Hyde Park
Road, Santa Fe, NM 87501
- Ioannides, Yannis M., Department of Economics, Tufts University, Medford, MA
02155
- Kauffman, Stuart A., Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM
87501
- Kirman, Alan P., G.R.E.Q.A.M., E.H.E.S.S. and Université d'Aix-Marseille III,
Institut Universitaire de France, 2 Rue de la Charite, 13002 Marseille,
FRANCE
- Kollman, Ken, Department of Political Science and Center for Political Studies,
University of Michigan, Ann Arbor, MI 48109
- Krugman, Paul, Department of Economics, Stanford University, Stanford, CA
95305
- Lane, David, Department of Political Economy, University of Modena, ITALY
- LeBaron, Blake, Department of Economics, University of Wisconsin, Madison,
WI 53706 and Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM
87501
- Leijonhufvud, Axel, Center for Computable Economics, and Department of
Economics, UCLA, 405 Hilgard Avenue, Los Angeles, CA 90095
- Lindgren, Kristian, Institute of Physical Resource Theory, Chalmers University of
Technology and Göteborg University, S-412 96 Göteborg, SWEDEN
- Manski, Charles F., Department of Economics, University of Wisconsin at
Madison, Madison, WI 53706
- Maxfield, Robert, Department of Engineering and Economic Systems, Stanford
University, Stanford, CA 95305

Miller, John H., Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213

North, Douglass C., Department of Economics, Washington University, St. Louis, MO 63130-4899

Padgett, John F., Department of Political Science, University of Chicago, Chicago, IL 60637

Page, Scott, Division of Humanities and Social Sciences, California Institute of Technology 228-77, Pasadena, CA 91125

Palmer, Richard, Department of Physics, Duke University, Durham, NC 27706 and Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501

Shubik, Martin, Yale University, Cowles Foundation for Research in Economics, Department of Economics, P.O. Box 208281, New Haven, CT 06520-8281

Taylor, Paul, Department of Computer Science, Brunel University, London, UK

Tesfatsion, Leigh, Department of Economics, Heady Hall 260, Iowa State University, Ames, IA 50011-1070

Acknowledgment

The conference at which these papers were presented was sponsored by Legg Mason, whose support we gratefully acknowledge. Over the years the Santa Fe Institute's Economics Program has benefited from the generosity of Citicorp, Coopers & Lybrand, The John D. and Catherine T. MacArthur Foundation, McKinsey and Company, the Russell Sage Foundation, and SFI's core support. We thank Eric Beinhocker, Caren Grown, Win Farrell, Dick Foster, Henry Lichstein, Bill Miller, John Reed, and Eric Wanner, not only for their organizations' financial support but for the moral and intellectual support they have provided. Their many insights and suggestions over the years have greatly bolstered the program. We also thank the members of SFI's Business Network, and the many researchers who have taken part in the program. George Cowan took a chance early on that an economics program at the Institute would be a success. We thank him for his temerity.

One of the pleasures of working at the Santa Fe Institute is the exemplary staff support. In particular we thank Ginger Richardson, the staff Director of Programs, and Andi Sutherland, who organized the conference this book is based on. We are very grateful to the very able publications people at SFI, especially Marylee Thomson and Della Ulibarri.

Philip Anderson and Kenneth Arrow have been guiding lights of the SFI Economics Program since its inception. Their intellectual and personal contributions are too long to enumerate. With respect and admiration, this book is dedicated to them.

W. Brian Arthur, Steven N. Durlauf, and David A. Lane



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Contents

Introduction <i>W. B. Arthur, S. N. Durlauf, and D. Lane</i>	1
Asset Pricing Under Endogenous Expectations in an Artificial Stock Market <i>W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, and P. Tayler</i>	15
Natural Rationality <i>V. M. Darley and S. A. Kauffman</i>	45
Statistical Mechanics Approaches to Socioeconomic Behavior <i>S. N. Durlauf</i>	81
Is What Is Good for Each Best for All? Learning From Others in the Information Contagion Model <i>D. Lane</i>	105
Evolution of Trading Structures <i>Y. M. Ioannides</i>	129
Foresight, Complexity, and Strategy <i>D. Lane and R. Maxfield</i>	169
The Emergence of Simple Ecologies of Skill <i>J. F. Padgett</i>	199

Some Fundamental Puzzles in Economic History/ Development <i>D. C. North</i>	223
How the Economy Organizes Itself in Space: A Survey of the New Economic Geography <i>P. Krugman</i>	239
Time and Money <i>M. Shubik</i>	263
Promises Promises <i>J. Geanakoplos</i>	285
Macroeconomics and Complexity: Inflation Theory <i>A. Leijonhufvud</i>	321
Evolutionary Dynamics in Game-Theoretic Models <i>K. Lindgren</i>	337
Identification of Anonymous Endogenous Interactions <i>C. F. Manski</i>	369
Asset Price Behavior in Complex Environments <i>W. A. Brock</i>	385
Population Games <i>L. E. Blume</i>	425
Computational Political Economy <i>K. Kollman, J. H. Miller, and S. Page</i>	461
The Economy as an Interactive System <i>A. P. Kirman</i>	491
How Economists Can Get ALife <i>L. Tesfatsion</i>	533
Some Thoughts About Distribution in Economics <i>P. W. Anderson</i>	565
Index	567

W. B. Arthur,* S. N. Durlauf, and D. Lane†**

*Citibank Professor, Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501

**Department of Economics, University of Wisconsin at Madison, 53706 and Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM, 87501

†Department of Political Economy, University of Modena, ITALY

Introduction

PROCESS AND EMERGENCE IN THE ECONOMY

In September 1987, twenty people came together at the Santa Fe Institute to talk about “the economy as an evolving, complex system.” Ten were theoretical economists, invited by Kenneth J. Arrow, and ten were physicists, biologists, and computer scientists, invited by Philip W. Anderson. The meeting was motivated by the hope that new ideas bubbling in the natural sciences, loosely tied together under the rubric of “the sciences of complexity,” might stimulate new ways of thinking about economic problems. For ten days, economists and natural scientists took turns talking about their respective worlds and methodologies. While physicists grappled with general equilibrium analysis and noncooperative game theory, economists tried to make sense of spin glass models, Boolean networks, and genetic algorithms.

The meeting left two legacies. The first was a volume of essays, *The Economy as an Evolving Complex System*, edited by Arrow, Anderson, and David Pines. The

other was the founding, in 1988, of the Economics Program at the Santa Fe Institute, the Institute's first resident research program. The Program's mission was to encourage the understanding of economic phenomena from a complexity perspective, which involved the development of theory as well as tools for modeling and for empirical analysis. To this end, since 1988, the Program has brought researchers to Santa Fe, sponsored research projects, held several workshops each year, and published several dozen working papers. And, since 1994, it has held an annual summer school for economics graduate students.

This volume, *The Economy as an Evolving Complex System II*, represents the proceedings of an August 1996 workshop sponsored by the SFI Economics Program. The intention of this workshop was to take stock, to ask: What has the complexity perspective contributed to economics in the past decade? In contrast to the 1987 workshop, almost all of the presentations addressed economic problems, and most participants were economists by training. In addition, while some of the work presented was conceived or carried out at the Institute, some of the participants had no previous relation with SFI—research related to the complexity perspective is under active development now in a number of different institutes and university departments.

But just what *is* the complexity perspective in economics? That is not an easy question to answer. Its meaning is still very much under construction, and, in fact, the present volume is intended to contribute to that construction process. Indeed, the authors of the essays in this volume by no means share a single, coherent vision of the meaning and significance of complexity in economics. What we will find instead is a family resemblance, based upon a set of interrelated themes that together constitute the current meaning of the complexity perspective in economics.

Several of these themes, already active subjects of research by economists in the mid-1980s, are well described in the earlier *The Economy as an Evolving Complex System*: In particular, applications of nonlinear dynamics to economic theory and data analysis, surveyed in the 1987 meeting by Michele Boldrin and William Brock; and the theory of positive feedback and its associated phenomenology of path dependence and lock-in, discussed by W. Brian Arthur. Research related to both these themes has flourished since 1987, both in and outside the SFI Economics Program. While chaos has been displaced from its place in 1987 at center stage of the interest in nonlinear dynamics, in the last decade economists have made substantial progress in identifying patterns of nonlinearity in financial time series and in proposing models that both offer explanations for these patterns and help to analyze and, to some extent, predict the series in which they are displayed. Brock surveys both these developments in his chapter in this volume, while positive feedback plays a central role in the models analyzed by Lane (on information contagion), Durlauf (on inequality) and Krugman (on economic geography), and lurk just under the surface of the phenomena described by North (development) and Leijonhufvud (high inflation).

Looking back over the developments in the past decade and the papers produced by the program, we believe that a coherent perspective—sometimes called

the “Santa Fe approach”—has emerged within economics. We will call this the complexity perspective, or Santa Fe perspective, or occasionally the process-and-emergence perspective. Before we describe this, we first sketch the two conceptions of the economy that underlie standard, neoclassical economics (and indeed most of the presentations by economic theorists at the earlier 1987 meeting). We can call these conceptions the “equilibrium” and “dynamical systems” approaches. In the equilibrium approach, the problem of interest is to derive, from the rational choices of individual optimizers, aggregate-level “states of the economy” (prices in general equilibrium analysis, a set of strategy assignments in game theory with associated payoffs) that satisfy some aggregate-level consistency condition (market-clearing, Nash equilibrium), and to examine the properties of these aggregate-level states. In the dynamical systems approach, the state of the economy is represented by a set of variables, and a system of difference equations or differential equations describes how these variables change over time. The problem is to examine the resulting trajectories, mapped over the state space. However, the equilibrium approach does not describe the *mechanism* whereby the state of the economy changes over time—nor indeed how an equilibrium comes into being.^[1] And the dynamical system approach generally fails to accommodate the distinction between *agent-* and *aggregate-*levels (except by obscuring it through the device of “representative agents”). Neither accounts for the emergence of new kinds of relevant state variables, much less new entities, new patterns, new structures.^[2]

To describe the complexity approach, we begin by pointing out six features of the economy that together present difficulties for the traditional mathematics used in economics:^[3]

DISPERSED INTERACTION. What happens in the economy is determined by the interaction of many dispersed, possibly heterogeneous, agents acting in parallel. The action of any given agent depends upon the anticipated actions of a limited number of other agents and on the aggregate state these agents cocreate.

^[1]Since an a priori intertemporal equilibrium hardly counts as a mechanism.

^[2]Norman Packard’s contribution to the 1987 meeting addresses just this problem with respect to the dynamical systems approach. As he points out, “if the set of relevant variables changes with time, then the state space is itself changing with time, which is not commensurate with a conventional dynamical systems model.”

^[3]John Holland’s paper at the 1987 meeting beautifully—and presciently—frames these features. For an early description of the Santa Fe approach, see also the program’s March 1989 newsletter, “Emergent Structures.”

NO GLOBAL CONTROLLER. No global entity controls interactions. Instead, controls are provided by mechanisms of competition and coordination among agents. Economic actions are mediated by legal institutions, assigned roles, and shifting associations. Nor is there a universal competitor—a single agent that can exploit all opportunities in the economy.

CROSS-CUTTING HIERARCHICAL ORGANIZATION. The economy has many levels of organization and interaction. Units at any given level—behaviors, actions, strategies, products—typically serve as “building blocks” for constructing units at the next higher level. The overall organization is more than hierarchical, with many sorts of tangled interactions (associations, channels of communication) across levels.

CONTINUAL ADAPTATION . Behaviors, actions, strategies, and products are revised continually as the individual agents accumulate experience—the system constantly adapts.

PERPETUAL NOVELTY. Niches are continually created by new markets, new technologies, new behaviors, new institutions. The very act of filling a niche may provide new niches. The result is ongoing, perpetual novelty.

OUT-OF-EQUILIBRIUM DYNAMICS. Because new niches, new potentials, new possibilities, are continually created, the economy operates far from any optimum or global equilibrium. Improvements are always possible and indeed occur regularly.

Systems with these properties have come to be called *adaptive nonlinear networks* (the term is John Holland’s⁵). There are many such in nature and society: nervous systems, immune systems, ecologies, as well as economies. An essential element of adaptive nonlinear networks is that they do not act simply in terms of stimulus and response. Instead they anticipate. In particular, economic agents form expectations—they build up models of the economy and act on the basis of predictions generated by these models. These anticipative models need neither be explicit, nor coherent, nor even mutually consistent.

Because of the difficulties outlined above, the mathematical tools economists customarily use, which exploit linearity, fixed points, and systems of differential equations, cannot provide a deep understanding of adaptive nonlinear networks. Instead, what is needed are new classes of combinatorial mathematics and population-level stochastic processes, in conjunction with computer modeling. These mathematical and computational techniques are in their infancy. But they emphasize the *discovery* of structure and the *processes* through which structure *emerges* across different levels of organization.

This conception of the economy as an adaptive nonlinear network—as an evolving, complex system—has profound implications for the foundations of economic

theory and for the way in which theoretical problems are cast and solved. We interpret these implications as follows:

COGNITIVE FOUNDATIONS. Neoclassical economic theory has a unitary cognitive foundation: economic agents are rational optimizers. This means that (in the usual interpretation) agents evaluate uncertainty probabilistically, revise their evaluations in the light of new information via Bayesian updating, and choose the course of action that maximizes their expected utility. As glosses on this unitary foundation, agents are generally assumed to have common knowledge about each other and rational expectations about the world they inhabit (and of course cocreate). In contrast, the Santa Fe viewpoint is pluralistic. Following modern cognitive theory, we posit no single, dominant mode of cognitive processing. Rather, we see agents as having to cognitively structure the problems they face—as having to “make sense” of their problems—as much as solve them. And they have to do this with cognitive resources that are limited. To “make sense,” to learn, and to adapt, agents use variety of distributed cognitive processes. The very categories agents use to convert information about the world into action emerge from experience, and these categories or cognitive props need not fit together coherently in order to generate effective actions. Agents therefore inhabit a world that they must cognitively interpret—one that is complicated by the presence and actions of other agents and that is ever changing. It follows that agents generally do not optimize in the standard sense, not because they are constrained by finite memory or processing capability, but because the very concept of an optimal course of action often cannot be defined. It further follows that the deductive rationality of neoclassical economic agents occupies at best a marginal position in guiding effective action in the world. And it follows that any “common knowledge” agents might have about one another must be attained from concrete, specified cognitive processes operating on experiences obtained through concrete interactions. Common knowledge cannot simply be assumed into existence.

STRUCTURAL FOUNDATIONS. In general equilibrium analysis, agents do not interact with one another directly, but only through impersonal markets. By contrast, in game theory all players interact with all other players, with outcomes specified by the game’s payoff matrix. So interaction structures are simple and often extreme—one-with-all or all-with-all. Moreover, the internal structure of the agents themselves is abstracted away.^[4] In contrast, from a complexity perspective, structure matters. First, network-based structures become important. All economic action involves interactions among agents, so economic functionality is both constrained and carried by networks defined by recurring patterns of interaction among agents. These network structures are characterized by relatively sparse ties. Second, economic action is structured by emergent social roles and by socially supported procedures—that is,

^[4]Except in principal-agent theory or transaction-costs economics, where a simple hierarchical structure is supposed to obtain.

by institutions. Third, economic entities have a recursive structure: they are themselves comprised of entities. The resulting “level” structure of entities and their associated action processes is not strictly hierarchical, in that component entities may be part of more than one higher-level entity, and entities at multiple levels of organization may interact. Thus, reciprocal causation operates between different levels of organization—while action processes at a given level of organization may sometimes be viewed as autonomous, they are nonetheless constrained by action patterns and entity structures at other levels. And they may even give rise to new patterns and entities at both higher and lower levels. From the Santa Fe perspective, the fundamental principle of organization is the idea that units at one level combine to produce units at the next higher level. ^[5]

WHAT COUNTS AS A PROBLEM AND AS A SOLUTION. It should be clear by now that exclusively posing economic problems as multiagent optimization exercises makes little sense from the viewpoint we are outlining—a viewpoint that puts emphasis on process, not just outcome. In particular, it asks how new “things” arise in the world—cognitive things, like “internal models”; physical things, like “new technologies”; social things, like new kinds of economic “units.” And it is clear that if we posit a world of perpetual novelty, then outcomes cannot correspond to steady-state equilibria, whether Walrasian, Nash, or dynamic-systems-theoretical. The only descriptions that can matter in such a world are about transient phenomena—about process and about emergent structures. What then can we know about the economy from a process-and-emergence viewpoint, and how can we come to know it? Studying process and emergence in the economy has spawned a growth industry in the production of what are now generally called “agent-based models.” And what counts as a solution in an agent-based model is currently under negotiation. Many of the papers in this volume—including those by Arthur et al., Darley and Kauffman, Shubik, Lindgren, Kollman et al., Kirman, and Tesfatsion—address this issue, explicitly or implicitly. We can characterize these as seeking emergent structures arising in interaction processes, in which the interacting entities anticipate the future through cognitive procedures that themselves involve interactions taking place in multilevel structures.

A description of an approach to economics, however, is not a research program. To build a research program around a process-and-emergence perspective, two things have to happen. First, concrete economic problems have to be identified for which the approach may provide new insights. A number of candidates are offered in this volume: artifact innovation (Lane and Maxfield), the evolution of trading networks (Ioannides, Kirman, and Tesfatsion), money (Shubik), the origin and spatial distribution of cities (Krugman), asset pricing (Arthur et al. and

^[5]We need not commit ourselves to what constitutes economic “units” and “levels.” This will vary from problem context to problem context.

Brock), high inflation (Leijonhufvud) persistent differences in income between different neighborhoods or countries (Durlauf). Second, cognitive and structural foundations for modeling these problems have to be constructed and methods developed for relating theories based on these foundations to observable phenomena (Manski). Here, while substantial progress has been made since 1987, the program is far from complete.

The essays in this volume describe a series of parallel explorations of the central themes of process and emergence in an interactive world—of how to study systems capable of generating perpetual novelty. These explorations do not form a coherent whole. They are sometimes complementary, sometimes even partially contradictory. But what could be more appropriate to the Santa Fe perspective, with its emphasis on distributed processes, emergence, and self-organization? Here are our interpretations of the research directions that seem to be emerging from this process:

COGNITION. The central cognitive issues raised in this volume are ones of interpretation. As Shubik puts it, “the interpretation of data is critical. It is not what the numbers are, but what they mean.” How do agents render their world comprehensible enough so that “information” has meaning? The two papers by Arthur, Holland, LeBaron, Palmer, and Tayler and by Darley and Kauffman consider this. They explore problems in which a group of agents take actions whose effects depend on what the other agents do. The agents base their actions on expectations they generate about how other agents will behave. Where do these expectations come from? Both papers reject common knowledge or common expectations as a starting point. Indeed, Arthur et al. argue that common beliefs cannot be deduced. Because agents must derive their expectations from an imagined future that is the aggregate result of other agents’ expectations, there is a self-reference of expectations that leads to deductive indeterminacy. Rather, both papers suppose that each agent has access to a variety of “interpretative devices” that single out particular elements in the world as meaningful and suggest useful actions on the basis of the “information” these elements convey. Agents keep track of how useful these devices turn out to be, discarding ones that produce bad advice and tinkering to improve those that work. In this view, economic action arises from an evolving ecology of interpretive devices that interact with one another through the medium of the agents that use them to generate their expectations.

Arthur et al. build a theory of asset pricing upon such a view. Agents—investors—act as market statisticians. They continually generate expectational models—interpretations of what moves prices in the market—and test these by trading. They discard and replace models if not successful. Expectations in the market therefore become endogenous—they continually change and adapt to a market that they create together. The Arthur et al. market settles into a rich psychology, in which speculative bubbles, technical trading, and persistence of volatility emerge. The homogeneous rational expectations of the standard literature become a special case—possible in theory but unlikely to emerge in practice. Brock presents

a variant of this approach, allowing agents to switch between a limited number of expectational models. His model is simpler than that of Arthur et al., but he achieves analytical results, which he relates to a variety of stylized facts about financial time series, many of which have been uncovered through the application of nonlinear analysis over the past decade.

In the world of Darley and Kauffman, agents are arrayed on a lattice, and they try to predict the behavior of their lattice neighbors. They generate their predictions via an autoregressive model, and they can individually tune the number of parameters in the model and the length of the time series they use to estimate model parameters. Agents can change parameter number or history length by steps of length 1 each period, if by doing so they would have generated better predictions in the previous period. This induces a coevolutionary “interpretative dynamics,” which does not settle down to a stable regime of precise, coordinated mutual expectations. In particular, when the system approaches a “stable rational-expectations state,” it tends to break down into a disordered state. They use their results to argue against conventional notions of rationality, with infinite foresight horizons and unlimited deductive capability.

In his paper on high inflation, Leijonhufvud poses the same problem as Darley and Kauffman: Where should we locate agent cognition, between the extremes of “infinite-horizon optimization” and “myopic adaptation”? Leijonhufvud argues that the answer to this question is context dependent. He claims that in situations of institutional break-down like high inflation, agent cognition shifts toward the “short memory/short foresight adaptive mode.” The causative relation between institutional and cognitive shifts becomes reciprocal. With the shrinking of foresight horizons, markets for long-term loans (where long-term can mean over 15 days) disappear. And as inflation accelerates, units of accounting lose meaning. Budgets cannot be drawn in meaningful ways, the executive arm of government becomes no longer fiscally accountable to parliament, and local governments become unaccountable to national governments. Mechanisms of social and economic control erode. Ministers lose control over their bureaucracies, shareholders over corporate management.

The idea that “interpretative devices” such as explicit forecasting models and technical-trading rules play a central role in agent cognition fits with a more general set of ideas in cognitive science, summarized in Clark.² This work rejects the notion that cognition is all “in the head.” Rather, interpretive aids such as autoregressive models, computers, languages, or even navigational tools (as in Hutchins⁶) and institutions provide a “scaffolding,” an external structure on which much of task of interpreting the world is off-loaded. Clark² argues that the distinctive hallmark of in-the-head cognition is “fast pattern completion,” which bears little relation to the neoclassical economist’s deductive rationality. In this volume, North takes up this theme, describing some of the ways in which institutions scaffold interpretations of what constitutes possible and appropriate action for economic agents.

Lane and Maxfield consider the problem of interpretation from a different perspective. They are particularly interested in what they call attributions of functionality: interpretations about what an artifact does. They argue that new attributions of functionality arise in the context of particular kinds of agent relationships, where agents can differ in their interpretations. As a consequence, cognition has an unavoidable social dimension. What interpretations are possible depend on who interacts with whom, about what. They also argue that new functionality attributions cannot be foreseen outside the particular generative relationships in which they arise. This unforeseeability has profound consequences for what constitutes “rational” action in situations of rapid change in the structure of agent-artifact space.

All the papers mentioned so far take as fundamental the importance of cognition for economic theory. But the opposite point of view can also be legitimately defended from a process-and-emergence perspective. According to this argument, overrating cognition is just another error deriving from methodological individualism, the very bedrock of standard economic theory. *How* individual agents decide what to do may not matter very much. What happens as a result of their actions may depend much more on the interaction structure through which they act—who interacts with whom, according to which rules. Blume makes this point in the introduction to his paper on population games, which, as he puts it, provide a class of models that shift attention “from the fine points of individual-level decision theory to dynamics of agent interaction.” Padgett makes a similar claim, though for a different reason. He is interested in formulating a theory of the firm as a locus of transformative “work,” and he argues that “work” may be represented by “an orchestrated sequence of actions and reactions, the sequence of which produces some collective result (intended or not).” Hence, studying the structure of coordinated action-reaction sequences may provide insight into the organization of economic activity, without bringing “cognition” into the story at all. Padgett’s paper is inspired by recent work in chemistry and biology (by Eigen and Schuster³ and by Fontana and Buss,⁴ among others) that are considered exemplars of the complexity perspective in these fields.

STRUCTURE. Most human interactions, even those taking place in “economic” contexts, have a primarily social character: talking with friends, asking advice from knowledgeable acquaintances, working together with colleagues, living next to neighbors. Recurring patterns of such social interactions bind agents together into networks.^[6] According to standard economic theory, what agents do depends on their values and available information. But standard theory typically ignores where values and information come from. It treats agents’ values and information as exogenous and autonomous. In reality, agents learn from each other, and their values may be influenced by others’ values and actions. These processes of learning

[6]There is a voluminous sociological literature on interaction networks. Recent entry points include Noria and Eccles,⁷ particularly the essay by Granovetter entitled “Problems of Explanation in Economic Sociology,” and the methodological survey of Wasserman and Faust.⁸

and influencing happen through the social interaction networks in which agents are embedded, and they may have important economic consequences. For example, one of the models presented in Durlauf's paper implies that value relationships among neighbors can induce persistent income inequalities between neighborhoods. Lane examines a model in which information flowing between agents in a network determines the market shares of two competing products. Kirman's paper reviews a number of models that derive economic consequences from interaction networks.

Ioannides, Kirman, and Tesfatsion consider the problems of how networks emerge from initially random patterns of dyadic interaction and what kinds of structure the resulting networks exhibit. Ioannides studies mathematical models based on controlled random fields, while Tesfatsion works in the context of a particular agent-based model, in which the "agents" are strategies that play Prisoner's Dilemma with one another. Ioannides and Tesfatsion are both primarily interested in networks involving explicitly economic interactions, in particular trade. Their motivating idea, long recognized among sociologists (for example, Baker¹), is that markets actually function by means of networks of traders, and what happens in markets may reflect the structure of these networks, which in turn may depend on how the networks emerge.

Local interactions can give rise to large-scale spatial structures. This phenomenon is investigated by several of the papers in this volume. Lindgren's contribution is particularly interesting in this regard. Like Tesfatsion, he works with an agent-based model in which the agents code strategies for playing two-person games. In both Lindgren's and Tesfatsion's models, agents adapt their strategies over time in response to their past success in playing against other agents. Unlike Tesfatsion's agents, who meet randomly and decide whether or not to interact, Lindgren's agents only interact with neighbors in a prespecified interaction network. Lindgren studies the emergence of spatiotemporal structure in agent space—metastable ecologies of strategies that maintain themselves for many agent-generations against "invasion" by new strategy types or "competing" ecologies at their spatial borders. In particular, he compares the structures that arise in a lattice network, in which each agent interacts with only a few other agents, with those that arise in a fully connected network, in which each agent interacts with all other agents. He finds that the former "give rise to a stable coexistence between strategies that would otherwise be outcompeted. These spatiotemporal structures may take the form of spiral waves, irregular waves, spatiotemporal chaos, frozen patchy patterns, and various geometrical configurations." Though Lindgren's model is not explicitly economic, the contrast he draws between an agent space in which interactions are structured by (relatively sparse) social networks and an agent space in which all interactions are possible (as is the case, at least in principle, with the impersonal markets featured in general equilibrium analysis) is suggestive. Padgett's paper offers a similar contrast, in a quite different context.

Both Durlauf and Kirman explore the emergence of geographical segregation. In their models, agents may change location—that is, change their position in a social structure defined by neighbor ties. In these models (especially Durlauf's),

there are many types of agents, and the question is under what circumstances, and through what mechanisms, do aggregate-level “neighborhoods” arise, each consisting predominantly (or even exclusively) of one agent type. Thus, agents’ choices, conditioned by current network structure (the agent’s neighbors and the neighbors at the sites to which the agent can move), change that structure; over time, from the changing local network structure, an aggregate-level pattern of segregated neighborhoods emerges.

Kollman, Miller, and Page explore a related theme in their work on political platforms and institutions in multiple jurisdictions. In their agent-based model, agents may relocate between jurisdictions. They show that when there are more than three jurisdictions, two-party competition outperforms democratic referenda. The opposite is the case when there is only one jurisdiction and, hence, no agent mobility. They also find that two-party competition results in more agent moves than does democratic referenda.

Manski reminds us that while theory is all very well, understanding of real phenomena is just as important. He distinguishes between three kinds of causal explanation for the often observed empirical fact that “persons belonging to the same group tend to behave similarly.” One is the one we have been describing above: the behavioral similarities may arise through network interaction effects. But there are two other possible explanations: *contextual*, in which the behavior may depend on exogenous characteristics of the group (like socioeconomic composition); or *correlated effects*, in which the behavior may be due to similar *individual* characteristics of members of the group. Manski shows, among other results, that a researcher who uses the popular linear-in-means model to analyze his data and “observes equilibrium outcomes and the composition of reference groups cannot empirically distinguish” endogenous interactions from these alternative explanations. One moral is that nonlinear effects require nonlinear inferential techniques.

In the essays of North, Shubik, and Leijonhufvud, the focus shifts to another kind of social structure, the institution. North’s essay focuses on institutions and economic growth, Shubik’s on financial institutions, and Leijonhufvud’s on high-inflation phenomenology. All three authors agree in defining institutions as “the rules of the game,” without which economic action is unthinkable. They use the word “institution” in at least three senses: as the “rules” themselves (for example, bankruptcy laws); as the entities endowed with the social and political power to promulgate rules (for example, governments and courts); and as the socially legitimized constructions that instantiate rules and through which economic agents act (for example, fiat money and markets). In whichever sense institutions are construed, the three authors agree that they cannot be adequately understood from a purely economic, purely political, or purely social point of view. Economics, politics, and society are inextricably mixed in the processes whereby institutions come into being. And they change and determine economic, political, and social action. North also insists that institutions have a cognitive dimension through the aggregate-level “belief systems” that sustain them and determine the directions in which they change.

North takes up the question of the emergence of institutions from a functionalist perspective: institutions are brought into being “in order to reduce uncertainty,” that is, to make agents’ worlds predictable enough to afford recognizable opportunities for effective action. In particular, modern economies depend upon institutions that provide low transaction costs in impersonal markets.

Shubik takes a different approach. His analysis starts from his notion of strategic market games. These are “fully defined process models” that specify actions “for all points in the set of feasible outcomes.” He shows how, in the context of constructing a strategic market game for an exchange economy using fiat money, the full specification requirement leads to the logical necessity of certain kinds of rules that Shubik identifies with financial institutions. Geanakoplos’ paper makes a similar point to Shubik’s. Financial instruments represent promises, he argues. What happens if someone cannot or will not honor a promise? Shubik already introduced the logical necessity of one institution, bankruptcy law, to deal with defaults. Geanakoplos introduces another, collateral. He shows that, in equilibrium, collateral as an institution has institutional implications—missing markets.

Finally, in his note concluding the volume, Philip Anderson provides a physicist’s perspective on a point that Fernand Braudel argues is a central lesson from the history of long-term socioeconomic change. Averages and assumptions of agent homogeneity can be very deceptive in complex systems. And processes of change are generally driven by the inhabitants of the extreme tails of some relevant distribution. Hence, an interesting theoretical question from the Santa Fe perspective is: How do distributions with extreme tails arise, and why are they so ubiquitous and so important?

WHAT COUNTS AS A PROBLEM AND AS A SOLUTION. While the papers here have much to say on cognition and structure, they contain much less discussion on what constitutes a problem and solution from this new viewpoint. Perhaps this is because it is premature to talk about methods for generating and assessing understanding when what is to be understood is still under discussion. While a few of the papers completely avoid mathematics, most of the papers do present mathematical models—whether based on statistical mechanics, strategic market games, random graphs, population games, stochastic dynamics, or agent-based computations. Yet sometimes the mathematical models the authors use leave important questions unanswered. For example, in what way do equilibrium calculations provide insight into emergence? This troublesome question is not addressed in any of the papers, even those in which models are presented from which equilibria are calculated—and insight into emergence is claimed to result. Blume raises two related issues in his discussion of population games: whether the asymptotic equilibrium selection theorems featured in the theory happen “soon enough” to be economically interesting; and whether the invariance of the “global environment” determined by the game and interaction model is compatible with an underlying economic reality in which rules of the game undergo endogenous change. It will not be easy to resolve the

inherent tension between traditional mathematical tools and phenomena that may exhibit perpetual novelty.

As we mentioned previously, several of the papers introduce less traditional, agent-based models. Kollman, Miller, and Page discuss both advantages and difficulties associated with this set of techniques. They end up expressing cautious optimism about their future usefulness. Tesfatsion casts her own paper as an illustration of what she calls “the alife approach for economics, as well as the hurdles that remain to be cleared.” Perhaps the best recommendation we can make to the reader with respect to the epistemological problems associated with the process-and-emergence perspective is simple. Read the papers, and see what you find convincing.

REFERENCES

1. Baker W. "The Social Structure of a National Securities Market." *Amer. J. Sociol.* 89 (1984): 775–811.
2. Clark, A. *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press, 1997.
3. Eigen, M., and P. Schuster. *The Hypercycle*. Berlin: Springer Verlag, 1979.
4. Fontana, W., and L. Buss. "The Arrival of the Fittest: Toward a Theory of Biological Organization." *Bull. Math. Biol.* 56 (1994): 1–64
5. Holland, J. H. "The Global Economy as an Adaptive Process." In *The Economy as an Evolving Complex System*, edited by P. W. Anderson, K. J. Arrow, and D. Pines, 117–124. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. V. Redwood City, CA: Addison-Wesley, 1988.
6. Hutchins, E. *Cognition in the Wild*. Cambridge, MA: MIT Press, 1995.
7. Noria, N., and R. Eccles (Eds.) *Networks and Organizations: Structure, Form, and Action*. Cambridge, MA: Harvard Business School Press, 1992.
8. Wasserman, W., and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge, MA: Cambridge University Press, 1994.

W. Brian Arthur,† John H. Holland,‡ Blake LeBaron,* Richard Palmer,* and Paul Tayler**

†Citibank Professor, Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501

‡Professor of Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109 and Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501

*Associate Professor of Economics, University of Wisconsin, Madison, WI 53706 and Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501

*Professor of Physics, Duke University, Durham, NC 27706 and Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501

**Department of Computer Science, Brunel University, London

Asset Pricing Under Endogenous Expectations in an Artificial Stock Market

We propose a theory of asset pricing based on heterogeneous agents who continually adapt their expectations to the market that these expectations aggregatively create. And we explore the implications of this theory computationally using our Santa Fe artificial stock market.^[1]

Asset markets, we argue, have a recursive nature in that agents' expectations are formed on the basis of their anticipations of other agents' expectations, which precludes expectations being formed by deductive means. Instead, traders continually hypothesize—continually explore—expectational models, buy or sell on the basis of those that perform best, and confirm or discard these according to their performance. Thus, individual beliefs or expectations become endogenous to the market, and constantly compete within an ecology of others' beliefs or expectations. The ecology of beliefs coevolves over time.

Computer experiments with this endogenous-expectations market explain one of the more striking puzzles in finance: that market traders often believe in such concepts as technical trading, “market psychology,” and bandwagon effects, while

^[1]For a less formal discussion of the ideas in this paper see Arthur.³

academic theorists believe in market efficiency and a lack of speculative opportunities. Both views, we show, are correct, but within different regimes. Within a regime where investors explore alternative expectational models at a low rate, the market settles into the rational-expectations equilibrium of the efficient-market literature. Within a regime where the rate of exploration of alternative expectations is higher, the market self-organizes into a complex pattern. It acquires a rich psychology, technical trading emerges, temporary bubbles and crashes occur, and asset prices and trading volume show statistical features—in particular, GARCH behavior—characteristic of actual market data.

1. INTRODUCTION

Academic theorists and market traders tend to view financial markets in strikingly different ways. Standard (efficient-market) financial theory assumes identical investors who share rational expectations of an asset's future price, and who instantaneously and rationally discount all market information into this price.^[2] It follows that no opportunities are left open for consistent speculative profit, that technical trading (using patterns in past prices to forecast future ones) cannot be profitable except by luck, that temporary price overreactions—bubbles and crashes—reflect rational changes in assets' valuations rather than sudden shifts in investor sentiment. It follows too that trading volume is low or zero, and that indices of trading volume and price volatility are not serially correlated in any way. The market, in this standard theoretical view, is rational, mechanistic, and efficient. Traders, by contrast, often see markets as offering speculative opportunities. Many believe that technical trading is profitable,^[3] that something definable as a "market psychology" exists, and that herd effects unrelated to market news can cause bubbles and crashes. Some traders and financial writers even see the market itself as possessing its own moods and personality, sometimes describing the market as "nervous" or "sluggish" or "jittery." The market in this view is psychological, organic, and imperfectly efficient. From the academic viewpoint, traders with such beliefs—embarrassingly the very agents assumed rational by the theory—are irrational and superstitious. From the traders' viewpoint, the standard academic theory is unrealistic and not borne out by their own perceptions.^[4]

While few academics would be willing to assert that the market has a personality or experiences moods, the standard economic view has in recent years

[2] For the classic statement see Lucas,³⁴ or Diba and Grossman.¹⁶

[3] For evidence see Frankel and Froot.¹⁹

[4] To quote one of the most successful traders, George Soros⁴⁷: "this [efficient market theory] interpretation of the way financial markets operate is severely distorted... It may seem strange that a patently false theory should gain such widespread acceptance."

begun to change. The crash of 1987 damaged economists' beliefs that sudden price changes reflect rational adjustments to news in the market: several studies failed to find significant correlation between the crash and market information issued at the time (e.g., Cutler et al.¹²). Trading volume and price volatility in real markets are large—not zero or small, respectively, as the standard theory would predict^{32,44,45}—and both show significant autocorrelation.^{7,21} Stock returns also contain small, but significant serial correlations.^{18,33,39,48} Certain technical-trading rules produce statistically significant, if modest, long-run profits.¹⁰ And it has long been known that when investors apply full rationality to the market, they lack incentives both to trade and to gather information.^{23,24,36} By now, enough statistical evidence has accumulated to question efficient-market theories and to show that the traders' viewpoint cannot be entirely dismissed. As a result, the modern finance literature has been searching for alternative theories that can explain these market realities.

One promising modern alternative, the noise-trader approach, observes that when there are “noise traders” in the market—investors who possess expectations different from those of the rational-expectations traders—technical-trading strategies such as trend chasing may become rational. For example, if noise traders believe that an upswing in a stock's price will persist, rational traders can exploit this by buying into the uptrend, thereby exacerbating the trend. In this way positive-feedback trading strategies—and other technical-trading strategies—can be seen as rational, as long as there are nonrational traders in the market to prime these strategies.^{13,14,15,46} This “behavioral” noise-trader literature moves some way toward justifying the traders' view. But it is built on two less-than-realistic assumptions: the existence of unintelligent noise traders who do not learn over time that their forecasts are erroneous; and the existence of rational players who possess, by some unspecified means, full knowledge of both the noise traders' expectations and their own class's. Neither assumption is likely to hold up in real markets. Suppose for a moment an actual market with minimally intelligent noise traders. Over time, in all likelihood, some would discover their errors and begin to formulate more intelligent (or at least different) expectations. This would change the market, which means that the perfectly intelligent players would need to readjust *their* expectations. But there is no reason these latter would know the new expectations of the noise-trader deviants; they would have to derive their expectations by some means such as guessing or observation of the market. As the rational players changed, the market would change again. And so the noise traders might again further deviate, forcing further readjustments for the rational traders. Actual noise-trader markets, assumed stationary in theory, would start to unravel; and the perfectly rational traders would be left at each turn guessing the changed expectations by observing the market.

Thus, noise-trader theories, while they explain much, are not robust. But in questioning such theories we are led to an interesting sequence of thought. Suppose we were to assume “rational,” but nonidentical, agents who do not find themselves in a market with rational expectations, or with publicly known expectations. Suppose we allowed each agent continually to observe the market with an eye to

discovering profitable expectations. Suppose further we allowed each agent to adopt these when discovered and to discard the less profitable as time progressed. In this situation, agents' expectations would become endogenous—individually adapted to the current state of the market—and they would cocreate the market they were designed to exploit. How would such a market work? How would it act to price assets? Would it converge to a rational-expectations equilibrium—or would it uphold the traders' viewpoint?

In this chapter we propose a theory of asset pricing that assumes fully heterogeneous agents whose expectations continually adapt to the market these expectations aggregatively create. We argue that under heterogeneity, expectations have a recursive character: agents have to form their expectations from their anticipations of other agents' expectations, and this self-reference precludes expectations being formed by deductive means. So, in the absence of being able to deduce expectations, agents—no matter how rational—are forced to hypothesize them. Agents, therefore, continually form individual, hypothetical, expectational models or “theories of the market,” test these, and trade on the ones that predict best. From time to time they drop hypotheses that perform badly, and introduce new ones to test. Prices are driven endogenously by these induced expectations. Individuals' expectations, therefore, evolve and “compete” in a market formed by others' expectations. In other words, agents' expectations coevolve in a world they cocreate.

The natural question is whether these heterogeneous expectations coevolve into homogeneous rational-expectations beliefs, upholding the efficient-market theory, or whether richer individual and collective behavior emerges, upholding the traders' viewpoint and explaining the empirical market phenomena mentioned above. We answer this not analytically—our model, with its fully heterogeneous expectations, is too complicated to allow analytical solutions—but computationally. To investigate price dynamics, investment strategies, and market statistics in our endogenous-expectations market, we perform carefully controlled experiments within a computer-based market we have constructed, the SFI Artificial Stock Market.^[5]

The picture of the market that results from our experiments, surprisingly, confirms both the efficient-market academic view and the traders' view. But each is valid under different circumstances—in different regimes. In both circumstances, we initiate our traders with heterogeneous beliefs clustered randomly in an interval near homogeneous rational expectations. We find that if our agents very slowly adapt their forecasts to new observations of the market's behavior, the market converges to a rational-expectations regime. Here “mutant” expectations cannot get a profitable footing; and technical trading, bubbles, crashes, and autocorrelative behavior do not emerge. Trading volume remains low. The efficient-market theory prevails.

If, on the other hand, we allow the traders to adapt to new market observations at a more realistic rate, heterogeneous beliefs persist, and the market self-organizes

[5] For an earlier report on the SFI artificial stock market, see Palmer et al.³⁸

into a complex regime. A rich “market psychology”—a rich set of expectations—becomes observable. Technical trading emerges as a profitable activity, and temporary bubbles and crashes occur from time to time. Trading volume is high, with times of quiescence alternating with times of intense market activity. The price time series shows persistence in volatility, the characteristic GARCH signature of price series from actual financial markets. And it shows persistence in trading volume. And over the period of our experiments, at least, individual behavior evolves continually and does not settle down. In this regime, the traders’ view is upheld.

In what follows, we discuss first the rationale for our endogenous-expectations approach to market behavior; and introduce the idea of collections of conditional expectational hypotheses or “predictors” to implement this. We next set up the computational model that will form the basic framework. We are then in a position to carry out and describe the computer experiments with the model. Two final sections discuss the results of the experiments, compare our findings with other modern approaches in the literature, and summarize our conclusions.

2. WHY INDUCTIVE REASONING?

Before proceeding, we show that once we introduce heterogeneity of agents, deductive reasoning on the part of agents fails. We argue that in the absence of deductive reasoning, agents must resort to *inductive* reasoning, which is both natural and realistic in financial markets.

A. FORMING EXPECTATIONS BY DEDUCTIVE REASONING: AN INDETERMINACY

We make our point about the indeterminacy of deductive logic on the part of agents using a simple arbitrage pricing model, avoiding technical details that will be spelled out later. (This pricing model is a special case of our model in section 3, assuming risk coefficient λ arbitrarily close to 0, and gaussian expectational distributions.) Consider a market with a single security that provides a stochastic payoff or dividend sequence $\{d_t\}$, with a risk-free outside asset that pays a constant r units per period. Each agent i may form individual expectations of next period’s dividend and price, $E_i[d_{t+1}|I_t]$ and $E_i[p_{t+1}|I_t]$, with conditional variance of these combined expectations, $\sigma_{i,t}^2$, given current market information I_t . Assuming perfect arbitrage, the market for the asset clears at the equilibrium price:

$$p_t = \beta \sum_j w_{j,t} (E_j[d_{t+1}|I_t] + E_j[p_{t+1}|I_t]). \quad (1)$$

In other words, the security’s price p_t is bid to a value that reflects the current (weighted) average of individuals’ market expectations, discounted by the factor

$\beta = 1/(1+r)$, with weights $w_{j,t} = (1/\sigma_{j,t}^2)/\sum_k 1/\sigma_{k,t}^2$, the relative “confidence” placed in agent j ’s forecast.

Now, assuming intelligent investors, the key question is how the individual dividend and price expectations $E_i[d_{t+1}|I_t]$ and $E_i[p_{t+1}|I_t]$, respectively, might be formed. The standard argument that such expectations can be formed rationally (i.e., using deductive logic) goes as follows. Assume *homogeneous* investors who (i) use the available information I_t identically in forming their dividend expectations, and (ii) know that others use the same expectations. Assume further that the agents (iii) are perfectly rational (can make arbitrarily difficult logical inferences), (iv) know that price each time will be formed by arbitrage as in Eq. (1), and (v) that (iii) and (iv) are common knowledge. Then, expectations of future dividends $E_i[d_{t+k}|I_t]$ are by definition known, shared, and identical. And homogeneity allows us to drop the agent subscript and set the weights to $1/N$. It is then a standard exercise (see Diba and Grossman¹⁶) to show that by setting up the arbitrage, Eq. (1), for future times $t+k$, taking expectations across it, and substituting backward repeatedly for $E[p_{t+k}|I_t]$, agents can iteratively solve for the current price as^[6]

$$p_t = \beta^k \sum_{k=1}^{\infty} E[d_{t+k}|I_t]. \quad (2)$$

If the dividend expectations are unbiased, dividend forecasts will be upheld on average by the market and, therefore, the price sequence will be in rational-expectations equilibrium. Thus, the price fluctuates as the information $\{I_t\}$ fluctuates over time, and it reflects “correct” or “fundamental” value, so that speculative profits are not consistently available. Of course, rational-expectations models in the literature are typically more elaborate than this. But the point so far is that if we are willing to adopt the above assumptions—which depend heavily on homogeneity—asset pricing becomes deductively determinate, in the sense that agents can, in principle at least, logically derive the current price.

Assume now, more realistically, that traders are intelligent but heterogeneous—each may differ from the others. Now, the available shared information I_t consists of past prices, past dividends, trading volumes, economic indicators, rumors, news, and the like. These are merely qualitative information plus data sequences, and *there may be many different, perfectly defensible statistical ways*, based on different assumptions and different error criteria, to use them to predict future dividends.^{1,30} Thus, there is no objectively laid down, expectational model that differing agents can coordinate upon, and so there is no objective means for one agent to know other agents’ expectations of future dividends. This is sufficient to bring indeterminacy to the asset price in Eq. (1). But worse, the heterogeneous price expectations

[6]The second, constant-exponential-growth solution is normally ruled out by an appropriate transversality condition.

$E_i[p_{t+1}|I_t]$ are also indeterminate. For suppose agent i attempts rationally to deduce this expectation, he may take expectations across the market clearing Eq. (1) for time $t + 1$:

$$E_i[p_{t+1}|I_t] = \beta E_i \left[\sum_j \{w_{j,t+1} (E_j [d_{t+2}|I_t] + E_j [p_{t+2}|I_t])\} | I_t \right]. \quad (3)$$

This requires that agent i , in forming his expectation of price, take into account his expectations of *others'* expectations of dividends and price (and relative market weights) two periods hence. To eliminate, in like manner, the price expectation $E_j[p_{t+2}|I_t]$ requires a further iteration. But this leads agents into taking into account their expectations of others' expectations of others' expectations of future dividends and prices at period $t + 3$ —literally, as in Keynes'²⁷ phrase, taking into account “what average opinion expects the average opinion to be.”

Now, under homogeneity these expectations of others' expectations collapsed into single, shared, objectively determined expectations. Under heterogeneity, however, not only is there no objective means by which others' dividend expectations can be known, but attempts to eliminate the other unknowns, the price expectations, merely lead to the repeated iteration of subjective expectations of subjective expectations (or, equivalently, subjective priors on others' subjective priors)—an infinite regress in subjectivity. Further, this regress may lead to instability: If investor i believes that others believe future prices will increase, he may revise his expectations to expect upward-moving prices. If he believes that others believe a reversion to lower values is likely, he may revise his expectations to expect a reversion. We can, therefore, easily imagine swings and swift transitions in investors' beliefs, based on little more than ephemera—hints and perceived hints of others' beliefs about others' beliefs.

Under heterogeneity then, deductive logic leads to expectations that are not determinable. Notice the argument here depends in no way on agents having limits to their reasoning powers. It merely says that given differences in agent expectations, there is no logical means by which to arrive at expectations. And so, perfect rationality in the market can not be well defined. Infinitely intelligent agents cannot form expectations in a determinate way.

B. FORMING EXPECTATIONS BY INDUCTIVE REASONING

If heterogeneous agents cannot deduce their expectations, how then do they form expectations? They may observe market data, they may contemplate the nature of the market and of their fellow investors. They may derive expectational models by sophisticated, subjective reasoning. But in the end all such models will be—can only be—hypotheses. There is no objective way to verify them, except by observing their performance in practice. Thus, agents, in facing the problem of choosing appropriate predictive models, face the same problem that statisticians

face when choosing appropriate predictive models given a specific data set, but no objective means by which to choose a functional form. (Of course, the situation here is made more difficult by the fact that the expectational models investors choose affect the price sequence, so that our statisticians' very choices of model affect their data and so their choices of model.)

In what follows then, we assume that each agent acts as a market "statistician."^[7] Each continually creates multiple "market hypotheses"—subjective, expectational models—of what moves the market price and dividend. And each simultaneously tests several such models. Some of these will perform well in predicting market movements. These will gain the agent's confidence and be retained and acted upon in buying and selling decisions. Others will perform badly. They will be dropped. Still others will be generated from time to time and tested for accuracy in the market. As it becomes clear which expectational models predict well, and as poorly predicting ones are replaced by better ones, the agent learns and adapts. This type of behavior—coming up with appropriate hypothetical models to act upon, strengthening confidence in those that are validated, and discarding those that are not—is called *inductive reasoning*.^[8] It makes excellent sense where problems are ill defined. It is, in microscale, the scientific method. Agents who act by using inductive reasoning we will call inductively rational.^[9]

Each inductively rational agent generates multiple expectational models that "compete" for use within his or her mind, and survive or are changed on the basis of their predictive ability. The agents' hypotheses and expectations adapt to the current pattern of prices and dividends; and the pattern of prices changes to reflect the current hypotheses and expectations of the agents. We see immediately that the market possesses a *psychology*. We define this as the collection of market hypotheses, or expectational models or mental beliefs, that are being acted upon at a given time.

If there were some attractor inherent in the price-and-expectation-formation process, this market psychology might converge to a stable unchanging set of heterogeneous (or homogeneous) beliefs. Such a set would be statistically validated, and would, therefore, constitute a rational-expectations equilibrium. We investigate whether the market converges to such an equilibrium below.

[7]The phrase is Tom Sargent's.⁴² Sargent argues similarly, within a macroeconomic context, that to form expectations agents need to act as market statisticians.

[8]For earlier versions of induction applied to asset pricing and to decision problems, see Arthur^{1,2} (the *El Farol* problem), and Sargent.⁴² For accounts of inductive reasoning in the psychological and adaptation literature, see Holland et al.,²⁵ Rumelhart,⁴¹ and Schank and Abelson.⁴³

[9]In the sense that they use available market data to learn—and switch among—appropriate expectational models. Perfect inductive rationality, of course, is indeterminate. Learning agents can be arbitrarily intelligent, but without knowing others' learning methods cannot tell a priori that *their* learning methods are maximally efficient. They can only discover the efficacy of their methods by testing them against data.

3. A MARKET WITH INDUCED EXPECTATIONS

A. THE MODEL

We now set up a simple model of an asset market along the lines of Bray⁹ or Grossman and Stiglitz.²⁴ The model will be neoclassical in structure, but will depart from standard models by assuming heterogeneous agents who form their expectations inductively by the process outlined above.

Consider a market in which N heterogeneous agents decide on their desired asset composition between a risky stock paying a stochastic dividend, and a risk-free bond. These agents formulate their expectations separately, but are identical in other respects. They possess a constant absolute risk aversion (CARA) utility function, $U(c) = -\exp(-\lambda c)$. They communicate neither their expectations nor their buying or selling intentions to each other. Time is discrete and is indexed by t ; the horizon is indefinite. The risk-free bond is in infinite supply and pays a constant interest rate r . The stock is issued in N units, and pays a dividend, d_t , which follows a given exogenous stochastic process $\{d_t\}$ not known to the agents.

The dividend process, thus far, is arbitrary. In the experiments we carry out below, we specialize it to an AR(1) process

$$d_t = \bar{d} + \rho(d_{t-1} - \bar{d}) + \varepsilon_t, \tag{4}$$

where ε_t is gaussian, i.i.d., and has zero mean, and variance σ_ε^2 .

Each agent attempts, at each period, to optimize his allocation between the risk-free asset and the stock. Assume for the moment that agent i 's predictions at time t of the next period's price and dividend are normally distributed with (conditional) mean and variance, $E_{i,t}[p_{t+1} + d_{t+1}]$, and $\sigma_{i,t,p+d}^2$. (We say presently how such expectations are arrived at.) It is well known that under CARA utility and gaussian distributions for forecasts, agent i 's demand, $x_{i,t}$, for holding shares of the risky asset is given by:

$$x_{i,t} = \frac{E_{i,t}(p_{t+1} + d_{t+1} - p(1+r))}{\lambda \sigma_{i,t,p+d}^2}, \tag{5}$$

where p_t is the price of the risky asset at t , and λ is the degree of relative risk aversion.

Total demand must equal the number of shares issued:

$$\sum_{i=1}^N x_{i,t} = N, \tag{6}$$

which closes the model and determines the clearing price p —the current market price—in Eq. (5) above.

It is useful to be clear on timing in the market. At the start of time period t , the current dividend d_t is posted, and observed by all agents. Agents then use this information and general information on the state of the market (which includes the historical dividend sequence $\{\dots d_{t-2}, d_{t-1}, d_t\}$ and price sequence $\{\dots p_{t-2}, p_{t-1}\}$) to form their expectations of the next period's price and dividend $E_{i,t}(p_{t+1} + d_{t+1})$. They then calculate their desired holdings and pass their demand parameters to the specialist who declares a price p_t that clears the market. At the start of the next period the new dividend d_{t+1} is revealed, and the accuracies of the predictors active at time t are updated. The sequence repeats.

B. MODELING THE FORMATION OF EXPECTATIONS

At this point we have a simple, neoclassical, two-asset market. We now break from tradition by allowing our agents to form their expectations individually and inductively. One obvious way to do this would be to posit a set of individual-agent expectational models which share the same functional form, and whose parameters are updated differently by each agent (by least squares, say) over time, starting from different priors. We reject this in favor of a different approach that better reflects the process of induction outlined in section 2 above. We assume each agent, at any time, possesses a multiplicity of linear forecasting models—hypotheses about the direction of the market, or “theories of the market”—and uses those that are both best suited to the current state of the market and have recently proved most reliable. Agents then learn, not by updating parameters, but by discovering which of their hypotheses “prove out” best, and by developing new ones from time to time, via the genetic algorithm. This structure will offer several desirable properties: It will avoid biases introduced by a fixed, shared functional form. It will allow the individuality of expectations to emerge over time (rather than be built in only to a priori beliefs). And it will better mirror actual cognitive reasoning, in which different agents might well “cognize” different patterns and arrive at different forecasts from the same market data.

In the expectational part of the model, at each period, the time series of current and past prices and dividends are summarized by an array or information set of J market descriptors. And agents' subjective expectational models are represented by sets of *predictors*. Each predictor is a *condition/forecast* rule (similar to a Holland classifier which is a condition/action rule) that contains both a market condition that may at times be fulfilled by the current state of the market and a forecasting formula for next period's price and dividend. Each agent possesses M such individual predictors—holds M hypotheses of the market in mind simultaneously—and uses the most accurate of those that are *active* (matched by the current state of the market). In this way, each agent has the ability to “recognize” different sets of states of the market, and bring to bear appropriate forecasts, given these market patterns.

It may clarify matters to show briefly how we implement this expectational system on the computer. (Further details are in Appendix A.) Suppose we summarize the state of the market by $J = 13$ bits. The fifth bit might correspond to “the price has risen the last 3 periods,” and the tenth bit to “the price is larger than 16 times dividend divided by r ,” with 1 signaling the occurrence of the described state, and 0 its absence or nonoccurrence. Now, the condition part of all predictors corresponds to these market descriptors, and thus, also consists of a 13-bit array, each position of which is filled with a 0, or 1, or # (“don’t care”). A condition array matches or “recognizes” the current market state if all its 0’s and 1’s match the corresponding bits for the market state with the #’s matching either a 1 or a 0. Thus, the condition (#####1#####) “recognizes” market states in which the price has risen in the last 3 periods. The condition (#####0####) recognizes states where the current price is not larger than 16 times dividend divided by r . The forecasting part of each predictor is an array of parameters that triggers a corresponding forecasting expression. In our experiments, all forecasts use a linear combination of price and dividend, $E(p_{t+1} + d_{t+1}) = a(p_t + d_t) + b$. Each predictor then stores specific values of a and b . Therefore, the full predictor (#####1#####0####)/(0.96,0) can be interpreted as “if the price has risen in the last 3 periods, and if the price is not larger than 16 times dividend divided by r , then forecast next period’s price plus dividend as 96% of this period’s.” This predictor would recognize—would be activated by—the market state (0110100100011) but would not respond to the state (0110111011001).

Predictors that can recognize many states of the market have few 1’s and 0’s. Those more particularized have more 1’s and 0’s. In practice, we include for each agent a default predictor consisting of all #’s. The genetic algorithm creates new predictors by “mutating” the values in the predictor array, or by “recombination”—combining part of one predictor array with the complementary part of another.

The expectational system then works at each time with each agent observing the current state of the market, and noticing which of his predictors match this state. He forecasts next period’s price and dividend by combining statistically the linear forecast of the H most accurate of these active predictors, and given this expectation and its variance, uses Eq. (5) to calculate desired stock holdings and to generate an appropriate bid or offer. Once the market clears, the next period’s price and dividend are revealed and the accuracies of the active predictors are updated.

As noted above, learning in this expectational system takes place in two ways. It happens rapidly as agents learn which of their predictors are accurate and worth acting upon, and which should be ignored. And it happens on a slower time scale as the genetic algorithm from time to time discards nonperforming predictors and creates new ones. Of course these new, untested predictors do not create disruptions—they will be acted upon only if they prove accurate. This avoids brittleness and provides what machine-learning theorists call “gracefulness” in the learning process.

We can now discern several advantages of this multibit, multipredictor architecture. One is that this expectational architecture allows the market to have potentially different dynamics—a different character—under different states or circumstances. Because predictors are pattern-recognizing expectational models, and so can “recognize” these different states, agents can “remember” what happened before in given states and activate appropriate forecasts. This enables agents to make swift *gestalt*-like transitions in forecasting behavior should the market change.

Second, the design avoids bias from the choice of a particular functional form for expectations. Although the forecasting part of our predictors is linear, the multiplicity of predictors conditioned upon the many combinations of market conditions yield collectively at any time and for any agent a nonlinear forecasting expression in the form of a piecewise linear, noncontinuous forecasting function whose domain is the market state space, and whose accuracy is tuned to different regions of this space. (Forecasting is, of course, limited by the choice of the binary descriptors that represent market conditions.)

Third, learning is concentrated where it is needed. For example, $J = 12$ descriptors produces predictors that can distinguish more than four thousand different states of the market. Yet, only a handful of these states might occur often. Predictor conditions that recognize states that do not occur often will be used less often, their accuracy will be updated less often and, other things being equal, their precision will be lower. They are, therefore, less likely to survive in the competition among predictors. Predictors will, therefore, cluster in the more visited parts of the market state space, which is exactly what we want.

Finally, the descriptor bits can be organized into classes or information sets which summarize fundamentals, such as price-dividend ratios or technical-trading indicators, such as price trend movements. The design allows us to track exactly which information—which descriptor bits—the agents are using or ignoring, something of crucial importance if we want to test for the “emergence” of technical trading. This organization of the information also allows the possibility setting up different agent “types” who have access to different information sets. (In this chapter, all agents see all market information equally.)

A neural net could also supply several of these desirable qualities. However, it would be less transparent than our predictor system, which we can easily monitor to observe which information agents are individually and collectively using at each time.

4. COMPUTER EXPERIMENTS: THE EMERGENCE OF TWO MARKET REGIMES

A. EXPERIMENTAL DESIGN

We now explore computationally the behavior of our endogenous-expectations market in a series of experiments. We retain the same model parameters throughout these experiments, so that we can make comparisons of the market outcomes using the model under identical conditions with only controlled changes. Each experiment is run for 250,000 periods to allow asymptotic behavior to emerge if it is present; and it is run 25 times under different random seeds to collect cross-sectional statistics.

We specialize the model described in the previous section by choosing parameter values, and, where necessary, functional forms. We use $N = 25$ agents, who each have $M = 100$ predictors, which are conditioned on $J = 12$ market descriptors. The dividend follows the AR(1) process in Eq. (4), with autoregressive parameter ρ set to 0.95, yielding a process close to a random walk, yet persistent.

The 12 binary descriptors that summarize the state of the market are the following:

- 1–6 Current price \times interest rate/dividend $> 0.25, 0.5, 0.75, 0.875, 1.0, 1.125$
- 7–10 Current price $>$ 5-period moving average of past prices (MA),
10-period MA, 100-period MA, 500-period MA
- 11 Always on (1)
- 12 Always off (0)

The first six binary descriptors—the first six bits—reflect the current price in relation to current dividend, and thus, indicate whether the stock is above or below fundamental value at the current price. We will call these “fundamental” bits. Bits 7–10 are “technical-trading” bits that indicate whether a trend in the price is under way. They will be ignored if useless, and acted upon if technical-analysis trend following emerges. The final two bits, constrained to be 0 or 1 at all times, serve as experimental controls. They convey no useful market information, but can tell us the degree to which agents act upon useless information at any time. We say a bit is “set” if it is 0 or 1, and predictors are selected randomly for recombination, other things equal, with slightly lower probabilities the higher their specificity—that is, the more set bits they contain (see Appendix A). This introduces a weak drift toward the all-# configuration, and ensures that the information represented by a particular bit is used only if agents find it genuinely useful in prediction. This market information design allows us to speak of “emergence.” For example, it can be said that technical trading has emerged if bits 7–10 become set significantly more often, statistically, than the control bits.

We assume that forecasts are formed by each predictor j storing values for the parameters a_j, b_j , in the linear combination of price and dividend, $E_j[p_{t+1} +$

$d_{t+1}|I_t] = a_j(p_t + d_t) + b_j$. Each predictor also stores a current estimate of its forecast variance. (See Appendix A.)

Before we conduct experiments, we run two diagnostic tests on our computer-based version of the model. In the first, we test to see whether the model can replicate the rational-expectations equilibrium (r.e.e.) of standard theory. We do this by calculating analytically the homogeneous rational-expectations equilibrium (h.r.e.e.) values for the forecasting parameters a and b (see Appendix A), then running the computation with all predictors “clamped” to these calculated h.r.e.e. parameters. We find indeed that such predictions are upheld—that the model indeed reproduces the h.r.e.e.—which assures us that the computerized model, with its expectations, demand functions, aggregation, market clearing, and timing sequence, is working correctly. In the second test, we show the agents a given dividend sequence and a calculated h.r.e.e. price series that corresponds to it, and test whether they individually learn the correct forecasting parameters. They do, though with some variation due to the agents’ continual exploration of expectational space, which assures us that our agents are learning properly.

B. THE EXPERIMENTS

We now run two sets of fundamental experiments with the computerized model, corresponding respectively to slow and medium rates of exploration by agents of alternative expectations. The two sets give rise to two different *regimes*—two different sets of characteristic behaviors of the market. In the slow-learning-rate experiments, the genetic algorithm is invoked every 1,000 periods on average, predictors are crossed over with probability 0.3, and the predictors’ accuracy-updating parameter θ is set to $1/150$. In the medium-exploration-rate experiments, the genetic algorithm is invoked every 250 periods on average, crossover occurs with probability 0.1, and the predictors’ accuracy-updating parameter θ is set to $1/75$.^[10] Otherwise, we keep the model parameters the same in both sets of experiments, and in both we start the agents with expectational parameters selected randomly from a uniform distribution of values centered on the calculated homogeneous rational-expectations ones. (See Appendix A.) In the slow-exploration-rate experiments, no non-r.e.e. expectations can get a footing: the market enters an evolutionarily stable, rational-expectations regime. In the medium-exploration-rate experiments, we find that the market enters a complex regime in which psychological behavior emerges, there are significant deviations from the r.e.e. benchmark, and statistical “signatures” of real financial markets are observed.

We now describe these two sets of experiments and the two regimes or phases of the market they induce.

[10] At the time of writing, we have discovered that the two regimes emerge, and the results are materially the same, if we vary *only* the rate of invocation of the genetic algorithm.

THE RATIONAL-EXPECTATIONS REGIME. As stated, in this set of experiments, agents continually explore in prediction space, but under low rates. The market price, in these experiments, converges rapidly to the homogeneous rational-expectations value adjusted for risk, even though the agents start with nonrational expectations. In other words, homogeneous rational expectations are an attractor for a market with endogenous, inductive expectations.^[11] This is not surprising. If some agents forecast differently than the h.r.e.e. value, then the fact that most other agents are using something close to the h.r.e.e. value, will return a market-clearing price that corrects these deviant expectations: There is a natural, if weak, attraction to h.r.e.e. The equilibrium within this regime differs in two ways from the standard, theoretical, rational-expectations equilibrium. First, the equilibrium is neither assumed nor arrived at by deductive means. Our agents instead arrive inductively at a homogeneity that overlaps that of the homogeneous, theoretical rational expectations. Second, the equilibrium is a stochastic one. Agents continually explore alternatives, albeit at low rates. This testing of alternative explorations, small as it is, induces some “thermal noise” into the system. As we would expect, in this regime, agents’ holdings remain highly homogeneous, trading volume remains low (reflecting only variations in forecasts due to mutation and recombination) and bubbles, crashes, and technical trading do not emerge. We can say that in this regime the efficient-market theory and its implications are upheld.

THE COMPLEX OR RICH PSYCHOLOGICAL REGIME. We now allow a more realistic level of exploration in belief space. In these experiments, as we see in Figure 1, the price series still appears to be nearly identical to the price in the rational-expectations regime. (It is lower because of risk attributable to the higher variance caused by increased exploration.)

On closer inspection of the results, however, we find that complex patterns have formed in the collection of beliefs, and that the market displays characteristics that differ materially from those in the rational-expectations regime. For example, when we magnify the difference between the two price series, we see systematic evidence of temporary price bubbles and crashes (Figure 2). We call this new set of market behaviors the rich-psychological, or complex, regime.

This appearance of bubbles and crashes suggests that technical trading, in the form of buying or selling into trends, has emerged in the market. We can check this rigorously by examining the information the agents condition their forecasts upon.

[11] Within a simpler model, Blume and Easley⁵ prove analytically the evolutionary stability of r.e.e.




FIGURE 1 Rational-expectations price vs. price in the rich psychological regime. The two price series are generated on the same random dividend series. The upper is the homogeneous r.e.e. price, the lower is the price in the complex regime. The higher variance in the latter case causes the lower price through risk aversion.




FIGURE 2 Deviations of the price series in the complex regime from fundamental value. The bottom graph shows the difference between the two price series in Figure 1 (with the complex series rescaled to match the r.e.e. one and the difference between the two doubled for ease of observation). The upper series is the h.r.e.e. price.

Figure 3 shows the number of technical-trading bits that are used (are 1's or 0's) in the population of predictors as it evolves over time. In both sets of experiments, technical-trading bits are initially seeded randomly in the predictor population. In the rational-expectations regime, however, technical-trading bits provide no useful information and fall off as useless predictors are discarded. But in the complex regime, they bootstrap in the population, reaching a steady-state value by 150,000 periods. Technical trading, once it emerges, remains.^[12]

Price statistics in the complex regime differ from those in the rational-expectations regime, mainly in that kurtosis is evident in the complex case (Table 1) and that volume of shares traded (per 10,000 periods) is about 300% larger in the complex case, reflecting the degree to which the agents remain heterogeneous in their expectations as the market evolves. We note that fat tails and high volume are also characteristic of price data from actual financial markets.

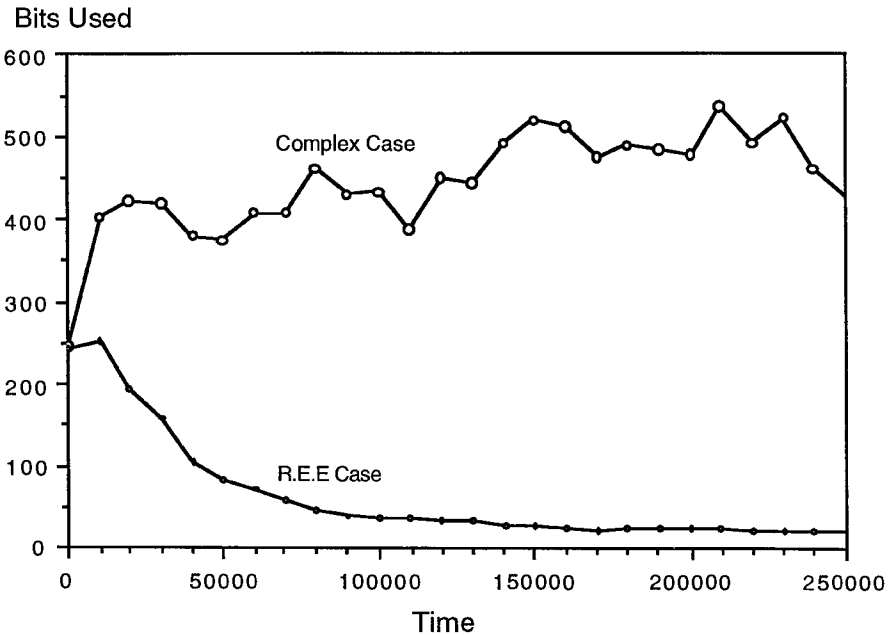


FIGURE 3 Number of technical-trading bits that become set as the market evolves, (median over 25 experiments in the two regimes).

[12] When we run these experiments informally to 1,000,000 periods, we see no signs that technical-trading bits disappear.

TABLE 1 Returns and volume statistics (medians) in the two regimes collected for 25 experiments after 250,000 periods.

	Mean	Std. Dev.	Skewness	Kurtosis ¹	Vol. traded
R.e.e. Regime	0.000	2.1002	0.0131	0.0497	2,460.9
Complex Regime	0.000	2.1007	0.0204	0.3429	7,783.8

¹ Kurtosis numbers are excess kurtosis (i.e., kurtosis -3).

How does technical trading emerge in psychologically rich or complex regime? In this regime the “temperature” of exploration is high enough to offset, to some degree, expectations’ natural attraction to the r.e.e. And so, subsets of non-r.e.e. beliefs need not disappear rapidly. Instead they can become mutually reinforcing. Suppose, for example, predictors appear early on that, by chance, condition an upward price forecast upon the markets showing a current rising trend. Then, agents who hold such predictors are more likely to buy into the market on an uptrend, raising the price over what it might otherwise be, causing a slight upward bias that might be sufficient to lend validation to such rules and retain them in the market. A similar story holds for predictors that forecast reversion to fundamental value. Such predictors need to appear in sufficient density to validate each other and remain in the population of predictors. The situation here is analogous to that in theories of the origin of life, where there needs to be a certain density of mutually reinforcing RNA units in the “soup” of monomers and polymers for such replicating units to gain a footing.^{17,26} Thus, technical analysis can emerge if trend-following (or mean-reversion) beliefs are, by chance, generated in the population, and if random perturbations in the dividend sequence activate them and subsequently validate them. From then on, they may take their place in the population of patterns recognized by the agents and become mutually sustainable. This emergence of structure from the mutual interaction of system subcomponents justifies our use of the label “complex” for this regime.

What is critical to the appearance of subpopulations of mutually reinforcing forecasts, in fact, is the presence of market information to condition upon. Market states act as “sunspot-like” signals that allow predictors to coordinate upon a direction they associate with that signal. (Of course, these are not classic sunspots that convey no real information.) Such coordination or mutuality can remain in the market once it establishes itself by chance. We can say the ability of market states to act as signals primes the mutuality that causes complex behavior. There is no need to assume a separate class of noise traders for this purpose. We can

test this signaling conjecture in further experiments where we “turn off” the condition part of all predictors (by filling them with nonreplaceable #'s). Now forecasts cannot differentiate among states of the market, and market states cannot act as signals. We find, consistent with our conjecture that signaling drives the observed patterns, that the complex regime does not emerge. As a further test of the significance of technical-trading signals, we regress the current price on the previous periods plus the technical indicator (price > 500-period moving average). In the rational-expectations regime, the technical indicator is of course not significant. In the complex regime, the trend indicator *is* significant (with *t*-value of 5.1 for the mean of the sample of 25 experiments), showing that the indicator does indeed carry useful market information. The corresponding test on actual financial data shows a similar result.¹⁰

One of the striking characteristics of actual financial markets is that both their price volatility and trading volume show persistence or autocorrelation. And volatility and volume show significant cross-correlation. In other words, both volume and volatility remain high or low for periods of random length, and they are interrelated. Our inductive market also shows persistence in volatility or GARCH behavior in the complex regime (see Figure 4), with the Chi-square statistic in the Engle GARCH Test significant at the 95% level.^[13] It also shows persistence in trading volume (see Figure 5), as well as significant cross-correlation between trading volume and volatility (see Figure 6). The figures include corresponding correlations for the often-used market standard, IBM stock. (Note that because our time period and actual market days do not necessarily match, we should expect no exact overlap. But qualitatively, persistence in our market and IBM's is similar.) These correlations are not explained by the standard model, where theoretically they are zero.

Why financial markets—and our inductive market—show these empirical “signatures” remains an open question. We conjecture a simple evolutionary explanation. Both in real markets and in our artificial market, agents are constantly exploring and testing new expectations. Once in a while, randomly, more successful expectations will be discovered. Such expectations will change the market, and trigger further changes in expectations, so that small and large “avalanches” of change will cascade through the system. (Of course, on this very short time-lag scale, these avalanches occur not through the genetic algorithm, but by agents changing their active predictors.) Changes then manifest in the form of increased volatility and increased volume. One way to test this conjecture is to see whether autocorrelations increase as the predictor accuracy-updating parameter θ in Eq. (7) in Appendix A is increased. The larger θ is, the faster individual agents “switch” among their

[13] Autocorrelated volatility is often fitted with a Generalized Autoregressive Conditional Heteroscedastic time series. Hence, the GARCH label. See Bollerslev et al.⁷ and Goodhart and O'Hara.²¹




FIGURE 4 Autocorrelation of volatility in rational-expectations and complex regimes, and in IBM daily returns.




FIGURE 5 Autocorrelation of trading volume in the rational-expectations and complex regimes, and in IBM daily returns.




FIGURE 6 Cross-correlation of trading volume with volatility, in the rational-expectations and complex regimes, and in IBM daily returns.

predictors. Thus, the more such switches should cascade. Experiments confirm that autocorrelations indeed increase with θ . Such cascades of switching in time are absorbed by the market, and die away. Hence, our evolutionary market exhibits periods of turbulence followed by periods of quiescence, as do actual markets.^[14]

5. DISCUSSION

To what extent is the existence of the complex regime an artifact of design assumptions in our model? We find experimentally by varying both the model's parameters and the expectational-learning mechanism, that the complex regime and the qualitative phenomena associated with it are robust. These are not an artifact of some deficiency in the model.^[15]

[14] For a discussion of volatility clustering in a different model, see Youssefmir and Huberman⁵⁰; and also Grannan and Swindle.²²

[15] One design choice might make a difference. We have evaluated the usefulness of expectational beliefs by their accuracy rather than by the profit they produce. In practice, these alternatives may produce different outcomes. For example, buying into a price rise on the basis of expectations may yield a different result if validated by profit instead of by accuracy of forecast when "slippage" is present, that is, when traders on the other side of the market are hard to find. We believe, but have not proved, that the two criteria lead to the same qualitative results.

It might be objected that if some agents could discover a superior means of forecasting to exploit the market, this might arbitrage complex patterns away, causing the market again to converge to rational expectations. We believe not. If a clever metaexpectational model was “out there” that might exploit others’ expectations, such a model would, by aggregation of others’ expectations, be a complicated nonlinear function of current market information. To the degree that the piecewise linear form we have assumed covers the space of nonlinear expectational models conditioned on current market information, agents would indeed, via the genetic algorithm, pick up on an approximate form of this superior metamodel. The complex regime owes its existence then not to limitations of forecasting, but rather to the fact that in our endogenous-expectations model market information can be used as signals, so that a much wider space of possibilities is open—in particular, the market can self-organize into mutually supporting subpopulations of predictors. (In fact, in a simpler, analytical model, with a small number of classes of trader whose beliefs adapt endogenously, Brock and Hommes¹¹ find similar, rich, asset-price dynamics.) There is no reason these emergent subpopulations should be in stochastic equilibrium. Indeed, agents may mutually adapt their expectations forever, so that the market explores its way through this large space, and is nonstationary. In some early exploratory experiments, we “froze” successful agents’ expectations, then reintroduced these agents with their previously successful expectations much later. The reintroduced agents proved less successful than average, indicating that the market had evolved and was nonstationary.

It might be also objected that by our use of condition bits in the predictors, we have built technical trading into our model. And so it is no surprise that it appears in the complex regime. But actually, only the possibility of technical trading is built in, not its use. The use of market descriptors is selected against in the model. Thus, market signals must be of value to be used, and technical trading emerges only because such market signals induce mutually supporting expectations that condition themselves on these market signals.

If the market has a well-defined psychology in our model, does it also experience “moods”? Obviously not. But, notice we assume that agents entertain more than one market hypothesis. Thus, we can imagine circumstances of a prolonged “bull-market” uptrend to a level well above fundamental value in which the market state activates predictors that indicate the uptrend will continue, and simultaneously other predictors that predict a rapid downward correction. Such combinations, which occur easily in both our market and actual markets, could well be described as “nervous.”

What about trade, and the motivation to trade in our market? In the rational-expectations literature, the deductively rational agents have no motivation to trade, even where they differ in beliefs. Assuming other agents have access to different information sets, each agent in a prebidding arrangement arrives at identical beliefs. Our inductively rational agents (who do not communicate directly), by contrast, do not necessarily converge in beliefs. They thus retain a motivation to trade, betting ultimately on their powers as market statisticians. It might appear that, because

our agents have equal abilities as statisticians, they are irrational to trade at all. But although their *abilities* are the same, their luck in finding good predictors diverges over time. And at each period, the accuracy of their predictors is fully accounted for in their allocations between the risk-free and risky asset. Given that agents can only act as market statisticians, their trading behavior is rational.

Our endogenous-expectation theory fits with two other modern approaches. Our model generalizes the learning models of Bray and others^{8,42} which also assume endogenous updating of expectations. But while the Bray models assume homogeneous updating from a shared nonrational forecast, our approach assumes heterogeneous agents who can discover expectations that might exploit any patterns present. Our evolutionary approach also has strong affinities with the evolutionary models of Blume and Easley.^{5,6} These assume populations of expectational (or more correctly, investment) rules that compete for survival in the market in a given population of rules, and that sometimes adapt. But the concern in this literature is the selective survival of different, competing, rule types, not the emergence of mutually supportive subpopulations that give rise to complex phenomena, nor the role of market signals in this emergence.

Our inductively rational market, of course, leaves out many details of realism. In actual financial markets, investors do not perfectly optimize portfolios, nor is full market clearing achieved each period. Indeed, except for the formation of expectations, our market is simple and neoclassical. Our object, however, is not market realism. Rather it is to show that given the inevitable inductive nature of expectations when heterogeneity is present, rich psychological behavior emerges—even under neoclassical conditions. We need not, as in other studies,^{20,28} assume sharing of information nor sharing of expectations nor herd effects to elicit these phenomena. Nor do we need to invoke “behaviorism” or other forms of irrationality.⁴⁹ Herding tendencies and quasi-rational behavior may be present in actual markets, but they are not necessary to our findings.

6. CONCLUSION

In asset markets, agents' forecasts create the world that agents are trying to forecast. Thus, asset markets have a reflexive nature in that prices are generated by traders' expectations, but these expectations are formed on the basis of anticipations of *others'* expectations.^[16] This reflexivity, or self-referential character of expectations, precludes expectations being formed by deductive means, so that perfect rationality ceases to be well defined. Thus, agents can only treat their expectations as hypotheses: they act inductively, generating individual expectational models that

[16] This point was also made by Soros⁴⁷ whose term *reflexivity* we adopt.

they constantly introduce, test, act upon, discard. The market becomes driven by expectations that adapt endogenously to the ecology these expectations cocreate.

Experiments with a computerized version of this endogenous-expectations market explain one of the more striking puzzles in finance: Standard theory tends to see markets as efficient, with no rationale for herd effects, and no possibility of systematic speculative profit, whereas traders tend to view the market as exhibiting a “psychology,” bandwagon effects, and opportunities for speculative profit. Recently the traders’ view has been justified by invoking behavioral assumptions, such as the existence of noise traders. We show, without behavioral assumptions, that both views can be correct. A market of inductively rational traders can exist in two different regimes: Under a low enough rate of exploration of alternative forecasts, the market settles into a simple regime which corresponds to the rational-expectations equilibrium of the efficient-market literature. Under a more realistic rate of exploration of alternative forecasts, the market self-organizes into a complex regime in which rich psychological behavior emerges. Technical trading appears, as do temporary bubbles and crashes. And prices show statistical features—in particular, GARCH behavior—characteristic of actual market data. These phenomena arise when individual expectations that involve trend following or mean reversion become mutually reinforcing in the population of expectations, and when market indicators become used as signaling devices that coordinate these sets of mutually reinforcing beliefs.

Our endogenous-expectations market shows that heterogeneity of beliefs, deviations from fundamental trading, and persistence in time series can be maintained indefinitely in actual markets with inductively rational traders. We conjecture that actual financial markets lie within the complex regime.

APPENDICES

APPENDIX A: DETAILS OF THE MARKET'S ARCHITECTURE

MODEL PARAMETERS. Throughout the experiments we set the interest rate r to 0.1, and agents' risk-aversion parameter λ to 0.5. The parameters of the dividend process in Eq. (4) are $\rho = 0.95$, $\bar{d} = 10$, $r = 0.1$, $\sigma_\epsilon^2 = 0.0743$. (This error variance value is selected to yield a combined price-plus-dividend variance of 4.0 in the h.r.e.e.)

PREDICTOR ACCURACY. The accuracy, or precision, of agent i 's j th predictor is updated each time the predictor is active, and is recorded as the inverse of the moving average of squared forecast error:

$$e_{t,i,j}^2 = (1 - \theta)e_{t-1,i,j}^2 + \theta[(p_{t+1} + d_{t+1}) - E_{t,i,j}(p_{t+1} + d_{t+1})]^2 \quad (7)$$

with $\theta = 1/75$ in the complex regime, and $1/150$ in the rational-expectations regime.

This accuracy is used in three places. First, if multiple predictors are active, only the most accurate is used. Second, it is part of the fitness measure for selecting predictors for recombination in the genetic algorithm. This fitness measure is defined as

$$f_{t,i,j} = M - e_{t,i,j}^2 - Cs \quad (8)$$

where M is a constant; s is specificity, the number of bits that are set (not #) in the predictor's condition array; and $C = 0.005$ is a cost levied for specificity. The value of M is irrelevant, given tournament rankings.

Third, agents use the error variance of their current predictor for the forecast variance in the demand Eq. (5). (We keep this latter variance fixed between genetic algorithm implementations, updating it to its current value in Eq. (7) at each invocation.)

INITIAL EXPECTATIONS. We initialize agents' expectations in both regimes by drawing the forecasting parameters from a uniform distribution of values centered upon the h.r.e.e. ones. We select a to be uniform (0.7, 1.2) and b to be uniform (-10, 19.002). The variance of all new predictors is initialized in all cases to the h.r.e.e. value of 4.0.

THE GENETIC ALGORITHM. New predictors are generated by updating each agent's predictor set at random intervals, on average every 250 periods or 1,000 periods, depending on the regime, asynchronously across agents. The worst performing (least accurate) 20% of the agent's 100 predictors are dropped, and are replaced by new ones, using uniform crossover and mutation. The agents are initialized by seeding them with random predictors: condition bits are set to 0 or 1 with probability 0.1, otherwise to #. This avoids bias in choosing predictors at the outset, and allows intelligent behavior to bootstrap itself up as the artificial agents generate predictive models that perform better. For the bitstrings, these procedures are standard genetic algorithm procedures for mutation and crossover (uniform crossover is used which chooses a bit at random from each of the two parents). The forecasting parameter vectors are mutated by adding random variables to each individual component. And they are crossed over component-wise, or by taking linear combinations of the two vectors, or by selecting one or the other complete vector. Each of these procedures is performed with equal probability. Crossover on a predictor is performed with probability 0.3 or 0.1 in the rational-expectations and complex regimes, respectively. Individual bits are mutated with probability 0.03. New predictors are brought into the predictor set with variance set to the average of their parents. If a bit has been changed, the new predictor's variance is set to the average of that of all predictors. If this new variance is lower than the variance of the current default predictor less an absolute deviation, its variance is set to the median of the predictors' variance. This procedure gives new predictors a reasonable chance of becoming used.

MARKET CLEARING. The price is adjusted each period by directly solving Eqs. (5) and (6) for p , which entails passing agents' forecasting parameters to the clearing equation. In actual markets, of course, the price is adjusted by a specialist who may not have access to agents' demand functions. But we note that actual specialists, either from experience or from their "books," have a keen feel for the demand function in their markets, and use little inventory to balance day-to-day demand. Alternatively, our market-clearing mechanism simulates an auction in which the specialist declares different prices and agents continually resubmit bids until a price is reached that clears the market.

CALCULATION OF THE HOMOGENEOUS RATIONAL-EXPECTATIONS EQUILIBRIUM. We calculate the homogeneous r.e.e. for the case where the market price is a linear function of the dividend $p_t = f d_t + g$ which corresponds to the structure of our forecasts. We can then calculate f and g from the market conditions at equilibrium. A homogenous equilibrium demands that all agents hold 1 share, so that, from Eq. (5)

$$E_t(p_{t+1} + d_{t+1}) - (1 + r)p_t = \lambda \sigma_{p+d}^2. \quad (9)$$

From the dividend process Eq. (4) and the linear form for the price, we can calculate $\sigma_{p+d}^2 = (1+f)\sigma_e^2$ and $E_t(p_{t+1} + d_{t+1})$ as

$$E_t(p_{t+1} + d_{t+1}) = (1+f) [(1-\rho)\bar{d} + \rho d_t] + g.$$

Noting that the right side of Eq. (9) is constant, we can then solve for f and g as

$$f = \frac{\rho}{1+r-\rho},$$

$$g = \frac{(1+f) [(1-\rho)\bar{d} - \lambda\sigma_e^2]}{r}.$$

Therefore, the expression:

$$E_t(p_{t+1} + d_{t+1}) = (1+r)p_t + \frac{\lambda(2+r)\sigma_e^2}{1+r-\rho} \quad (10)$$

is the h.r.e.e. forecast we seek.

APPENDIX B: THE SANTA FE ARTIFICIAL STOCK MARKET

The Santa Fe Artificial Stock Market has existed since 1989 in various designs (see Palmer et al.³⁸ for a description of an earlier version). Since then a number of other artificial markets have appeared: e.g., Beltratti and Margarita,⁴ Marengo and Tordjman,³⁵ and Rieck.⁴⁰ The Santa Fe Market is a computer-based model that can be altered, experimented with, and studied in a rigorously controlled way. Most of the artificial market's features are malleable and can be changed to carry out different experiments. Thus, the artificial market is a framework or template that can be specialized to focus on particular questions of interest in finance: for example, the effects of different agents having access to different information sets or predictive behaviors; or of a transaction tax on trading volume; or of different market-making mechanisms.

The framework allows other classes of utility functions, such as constant relative risk aversion. It allows a specialist or market maker, with temporary imbalances in fulfilled bids and offers, made up by changes in an inventory held by the specialist. It allows a number of alternative random processes for $\{d_t\}$. And it allows for the evolutionary selection of agents via wealth.

The market runs on a NeXTStep computational platform, but is currently being ported to the Swarm platform. For availability of code, and for further information, readers should contact Blake LeBaron or Richard Palmer.

ACKNOWLEDGMENTS

We are grateful to Kenneth Arrow, Eric Beinhocker, Larry Blume, Buz Brock, John Casti, Steven Durlauf, David Easley, David Lane, Ramon Marimon, Tom Sargent, and Martin Shubik for discussions of the arguments in this chapter, and of the design of the artificial market. All errors are our own.

REFERENCES

1. Arthur, W. B. "On Learning and Adaptation in the Economy." Working Paper 92-07-038, Santa Fe Institute, Santa Fe, NM, 1992.
2. Arthur, W. B. "Inductive Behavior and Bounded Rationality." *Amer. Econ. Rev.* **84** (May 1994): 406-411.
3. Arthur, W. B. "Complexity in Economic and Financial Markets." *Complexity* **1** (1995): 20-25.
4. Beltratti, A., and S. Margarita. "Simulating an Artificial Adaptive Stock Market." Mimeo, Turin University, 1992.
5. Blume, L., and D. Easley. "Learning to Be Rational." *J. Econ. Theory* **26** (1982): 340-351.
6. Blume, L., and D. Easley. "Evolution and Market Behavior." *J. Econ. Theory* **58** (1990): 9-40.
7. Bollerslev, T., R. Y. Chou, N. Jayaraman, and K. F. Kroner. "ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence." *J. Econometrics* **52** (1990): 5-60.
8. Bossaerts, P. "The Econometrics of Learning in Financial Markets." *Econometric Theory* **11** (1995): 151-189.
9. Bray, M. "Learning, Estimation, and Stability of Rational Expectations." *J. Econ. Theory* **26** (1982): 318-339.
10. Brock, W., J. Lakonishok, and B. LeBaron. "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." Working Paper 91-01-006, Santa Fe Institute, Santa Fe, NM, 1991.
11. Brock, W., and C. H. Hommes. "Models of Complexity in Economics and Finance." Mimeo, Department of Economics, University of Wisconsin, Madison, 1996.
12. Cutler, D. M., J. M. Poterba, and L. H. Summers. "What Moves Stock Prices?" *J. Portfolio Mgmt.* **15** (1989): 4-12.
13. De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann. "Noise Trader Risk in Financial Markets." *J. Pol. Econ.* **98** (1990): 703-738.

14. De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann. "Positive Feedback Strategies and Destabilizing Rational Speculation." *J. Fin.* 45 (1990): 379–395.
15. De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann. "The Survival of Noise Traders in Financial Markets." *J. Bus.* 64 (1991): 1–18.
16. Diba, B. T., and H. I. Grossman. "The Theory of Rational Bubbles in Stock Prices." *Econ. Jour.* 98 (1988): 746–754.
17. Eigen, M., and P. Schuster. *The Hypercycle: A Principle of Natural Self-Organization*. New York: Springer, 1979.
18. Fama, E. F., and K. R. French. "Permanent and Temporary Components of Stock Market Prices." *J. Pol. Econ.* 96 (1988): 246–273.
19. Frankel, J. A., and K. A. Froot. "Chartists, Fundamentalists, and Trading in the Foreign Exchange Market." *AEA Papers & Proc.* 80 (1990): 181–185.
20. Friedman, D., and M. Aoki. "Inefficient Information Aggregation as a Source of Asset Price Bubbles." *Bull. Econ. Res.* 44 (1992): 251–279.
21. Goodhart, C. A. E., and M. O'Hara. "High Frequency Data in Financial Markets: Issues and Applications." Unpublished manuscript, London School of Economics, 1995.
22. Grannan, E. R., and G. H. Swindle. "Contrarians and Volatility Clustering." Working Paper 94-03-010, Santa Fe Institute, Santa Fe, NM, 1994.
23. Grossman, S. J. "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information." *J. Fin.* 31 (1976): 573–585.
24. Grossman, S. J., and J. Stiglitz. "On the Impossibility of Informationally Efficient Markets." *Amer. Econ. Rev.* 70 (1980): 393–408.
25. Holland, J. H., K. J. Holyoak, R. E. Nisbett, and P. R. Thagard. *Induction*. Cambridge, MA: MIT Press, 1986.
26. Kauffman, S. *The Origin of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press, 1993.
27. Keynes, J. M. *General Theory of Employment, Interest, and Money*. London: Macmillan, 1936.
28. Kirman, A. "Epidemics of Opinion and Speculative Bubbles in Financial Markets." In *Money and Financial Markets*, edited by M. Taylor. London: Macmillan, 1991.
29. Kleidon, A. W. "Anomalies in Financial Economics: Blueprint for Change?" *J. Bus.* 59 (1986): 285–316 (supplement).
30. Kurz, M. "On the Structure and Diversity of Rational Beliefs." *Econ. Theory* 4 (1994): 877–900.
31. Kurz, M. "Asset Prices with Rational Beliefs." Working Paper 96-003, Economics Department, Stanford University, 1995.
32. Leroy, S. F., and R. D. Porter. "Stock Price Volatility: Tests Based on Implied Variance Bounds." *Econometrica* 49 (1981): 97–113.
33. Lo, A. W., and C. MacKinlay. "Stock Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test." *Rev. Fin. Stud.* 1 (1988): 41–66.