

602 LECTURE NOTES IN ECONOMICS
AND MATHEMATICAL SYSTEMS



Norman Ehrentreich

Agent-Based Modeling

The Santa Fe Institute Artificial Stock Market
Model Revisited

 Springer

Lecture Notes in Economics and Mathematical Systems

602

Founding Editors:

M. Beckmann
H.P. Künzi

Managing Editors:

Prof. Dr. G. Fandel
Fachbereich Wirtschaftswissenschaften
Fernuniversität Hagen
Feithstr. 140/AVZ II, 58084 Hagen, Germany

Prof. Dr. W. Trockel
Institut für Mathematische Wirtschaftsforschung (IMW)
Universität Bielefeld
Universitätsstr. 25, 33615 Bielefeld, Germany

Editorial Board:

A. Basile, A. Drexler, H. Dawid, K. Inderfurth, W. Kürsten

Norman Ehrentreich

Agent-Based Modeling

The Santa Fe Institute
Artificial Stock Market
Model Revisited

 Springer

Dr. Norman Ehrentreich
RiverSource Investments, LLC
262 Ameriprise Financial Center
Minneapolis, MN 55474
USA
norman.x.ehrentreich@ampf.com

ISBN 978-3-540-73878-7

e-ISBN 978-3-540-73879-4

DOI 10.1007/978-3-540-73878-7

Lecture Notes in Economics and Mathematical Systems ISSN 0075-8442

Library of Congress Control Number: 2007937522

© 2008 Springer-Verlag Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Production: LE-TeX Jelonek, Schmidt & Vöckler GbR, Leipzig
Cover-design: WMX Design GmbH, Heidelberg

Printed on acid-free paper

9 8 7 6 5 4 3 2 1

springer.com

Meinen Eltern gewidmet.

Foreword

When the original Santa Fe Institute (SFI) artificial stock market was created in the early 1990's, the creators realized that it contained many interesting new technologies that had never been tested in economic modeling. The authors kept to a very specific finance message in their papers, but the hope was that others would pick up where these papers left off and put these important issues to the test. Tackling the complexities involved in implementation has held many people back from this, and many parts of the SFI market remain unexplored. Ehrentreich's book is an important and careful study of some of the issues involved in the workings of the SFI stock market.

As Ehrentreich's book points out in its historical perspective, the SFI market was intended as a computational test bed for a market with boundedly rational learning agents replacing the standard setup of perfectly rational equilibrium modeling common in economics and finance. These agents exhibit reasonable, purposeful behavior, but they are not able to completely process every aspect of the world around them. This can be viewed much more as a function of the complexity of the world, rather than the computational limitations of agents. In a financial world out of equilibrium, optimal behavior would require knowledge of strategies being used by all the other agents, an information and computational task which seems well out of reach of any trader. The SFI market's main conclusion was that markets where agents were learning might not converge to traditional simple rational expectations equilibria. They go to some other steady state in which a rich set of trading strategies survives in the trader population. In this steady state the market demonstrates empirical signatures that are present in most financial time series.

This book is an excellent reference to both the learning, and empirical literature in finance. It stresses the difficult empirical facts that are out of reach of most traditional financial models including persistent volatility and trading volume, and technical trading behavior. However, Ehrentreich's main mission is to dig deeply into the SFI market structure to understand what is actually going on. Computational economic models can often be explored at three levels. There is sort of a big picture level where concepts such as rational expectations and bounded rationality are explored. There is also the very low level where researchers discuss the nuts and bolts of different modeling languages. In between these sits a region where many of the computational learning technologies are implemented. This is where technologies such as genetic algorithms, classifier systems, and neural networks drive much of what is going on. This is Ehrentreich's area of exploration, and it is critically important to agent-based modelers since one needs to know the sensitivity of the higher level results to changes in the learning structures used beneath them.

The SFI market uses two learning mechanisms extensively: the genetic algorithm, GA, and classifier system. Both of these are developments of John Holland, one of the SFI market coauthors. The GA is a type of general evolutionary learning mechanism, and it is used in both computer science and economics. Its properties have been studied, but it is still not completely understood. In computer science it is often studied in difficult optimization problems. These are problems with well defined objectives, and are quite different from the more open ended co-evolutionary problems in economics where agents are competing with each other. The classifier system is an interesting learning structure that allows agents to dynamically find relevant states in the world around them. For example, actions might be conditioned on whether a stock is currently priced above a certain multiple of dividends. The classifier has the power to endogenously slice up a stream of empirical information into states of the world. Very few learning mechanisms are able to do this. With this generality comes a lot of model complexity, and many implementations of the classifier seem computationally unwieldy. They also involve many implementation questions that need to be explored.

In several chapters Ehrentreich explores some of the more important aspects of the SFI classifier implementations. He shows that the SFI classifier is sensitive to certain design characteristics. Under different assumptions about evolution the classifier system behaves very differently from the original SFI model. Ehrentreich carefully modifies

and explores his own operation on mutating trading strategies. Using this modified mutation causes a situation in which the SFI market is much more likely to converge to the rational expectations equilibrium, and the rich technical trading dynamic does not emerge. The results in the original SFI market are clearly sensitive to how mutation is implemented. The book goes on to do a comparative study between mutation operators. A key issue is how many technical trading related rules are evolved, and whether the system is likely to generate lots of technical rules by chance in the evolutionary process. The modified mutation operator does not generate many of these rules, so they never really get a foot in the door of trading activity. The SFI structure facilitates their formation, but it is possible this could be driven more by genetic drift than selection. The original SFI studies never really answered these questions, and it only looked at trading strategy formation in an indirect level by looking at aggregate numbers. This was a clear weakness. Ehrentreich does some careful checks to see if technical rules are adding value at the agent level. It appears that they are, so many of the SFI indirect conclusions are sound.

The dynamics of wealth was never part of the original SFI market. It is an interesting omission that the SFI market never really considered long term wealth in a serious way in its implementation. This is strange since many arguments about efficient markets thrive on the relative dynamics of trader wealth. Ehrentreich concludes that this is a complex problem, and there may be difficulties with some of the other studies that try to tag a wealth dynamic onto the SFI market. In my opinion this is one of the biggest limitations of the actual SFI market structure.

This book is an important piece of work for understanding the dynamics of models with interacting learning agents. I think researchers in the future will find it critical in helping them to understand the dynamics of evolutionary learning models. Most importantly, it sets an important standard for doing careful internal experiments on these markets and the learning mechanisms inside them.

Brandeis University, Waltham, MA
September 2007

Blake LeBaron

Preface

The road of science is filled with surprises. When embarking on a scientific journey, we probably have a specific destination in mind, but we never know whether the road will take us there nor what places we may encounter along the way.

This trip was no exception. Before anyone starts reading this travelogue, I think that I should briefly mention a few places that I visited, but decided to pass over while writing this book. I originally aimed at converting the well-known Santa Fe Institute Artificial Stock Market (SFI-ASM) into a two stock version to study portfolio decisions of individual investors. My early forays into this unknown territory yielded some interim results, but until now they are still waiting to be further examined.

Instead, my road took a sudden and unexpected turn. One of the most important findings of the original SFI-ASM was the emergence of technical trading for faster learning speeds. Yet a thorough analysis of the agent's learning algorithm suggested that this might have been caused by an ill-designed mutation operator. For a couple of years, many tests confirmed this supposition. For instance, even though technical trading rules emerged in the original SFI-ASM, they were rarely acted upon. Most importantly, though, was that agents with an alternative mutation operator discovered the homogeneous rational expectations equilibrium, a result that found immediate approval by neoclassically inclined economists.

I traveled a long way down this road. Since I considered the existence of technical trading to be an empirical fact of financial markets, I tried to unearth the necessary ingredients to reintroduce it into my model. Nothing that I devised, neither social learning nor explicit herding mechanisms, succeeded in that endeavor. There was, however,

another surprise waiting behind the supposedly final turn of my journey. One newly designed test showed a slight superiority of technical trading rules in the original model. A side-trip all the way down to population genetics finally proved that my agents were committing a mistake by deciding to ignore technical trading rules. Again, parts of my prior research were discarded, and a new chapter was written explaining why I and previous researchers went wrong in interpreting the simulation results. I hope that this chapter will prove most useful for any research involving genetic algorithms. My prior belief that technical trading was an artificially introduced model artifact had also caused me to visit some previous studies about wealth levels. I was able to show that the SFI-ASM was not designed to address any questions related to wealth. Fortunately, this part was unaffected by the breakdown of the initial motivation to look into the wealth generation process.

A long journey with such detours was certainly not easy. I could not have arrived at the final destination without the tremendous support and encouragement that I have found along the way. Above all, I wish to thank my parents Werner and Ellinor Ehrentreich, for without them, I would not have had the opportunity to embark on this journey. I would also like to thank Reinhart Schmidt for letting me choose my destination and for giving me the freedom to follow my own path. Among the numerous friends, colleagues, and conference participants who have contributed in many ways are Manfred Jäger, Ulrike Neyer, Ralf Peters, Martin Klein, Heinz-Peter Galler, Joseph Felsenstein, Alan Kirman, and James Stodder. Of course, this book would not have been finished without the contributions by Blake LeBaron. Not only did he play a major role in the creation of the model that I set out to extend, then critiqued, and finally confirmed, he also often helped and clarified many questions that I was pondering. Many thanks also go to Lars Schiefner, Doris Storch, and Klaus Renger, especially for their help during the final stages of this project. Last, but not least, I thank Tanya Novak for her patience and help, especially for her proofreading. Nonetheless, I absolve her from all remaining mistakes and typos and credit them to my cats, Zina and Francesco, who stubbornly insisted on their input by jumping on the keyboard.

I now hope that the reader will find it useful to visit the places that I have found worthwhile to mention in this book.

Contents

Part I Agent-Based Modeling in Economics

1	Introduction	3
2	The Rationale for Agent-Based Modeling	5
2.1	Introduction	5
2.2	The Representative Agent Modeling Approach	7
2.2.1	Avoiding the Lucas-Critique	8
2.2.2	Building Walrasian General Equilibrium Models ..	9
2.2.3	Representative Agents and the Fallacy of Composition	10
2.2.4	Expectation Formation in Markets with Heterogeneous Investors	11
2.3	Rational Expectations and Disequilibrium Dynamics....	13
2.4	The Economy as an Evolving Complex Adaptive System	14
2.5	Some Methodological Aspects of Agent-Based Simulations	16
3	The Concept of Minimal Rationality	19
3.1	Introduction	19
3.2	Economic, Bounded, and Situational Rationality	21
3.3	Situational Analysis, Minimal Rationality, and the Prime Directive	24
3.4	Minimal Rationality and the Phillips-Curve	26
4	Learning in Economics	29
4.1	Introduction	29
4.2	Definitions of Learning	29
4.3	Rationality-Based Learning Models	31
4.4	Biologically Inspired Learning Models	32

4.4.1	Learning Through Replicator Dynamics	34
4.4.2	Learning Through Genetic Algorithms	36
4.4.3	Learning Through Classifier Systems	46
5	Replicating the Stylized Facts of Financial Markets ..	51
5.1	Efficient Markets and the Efficient Market Hypothesis ..	51
5.1.1	Definitions	51
5.1.2	Random Walks or Martingales?	53
5.1.3	Tests for Market Efficiency	55
5.2	Stylized Facts of Financial Markets	56
5.2.1	Non-Normal Return Distributions	56
5.2.2	Volatility Clustering of Returns	60
5.2.3	High and Persistent Trading Volume	64
5.2.4	Existence of Technical Trading	65
5.3	Alternative Market Hypotheses	70
5.3.1	The Fractal Market Hypothesis	70
5.3.2	The Coherent Market Hypothesis	71
5.3.3	The Adaptive Market Hypothesis	72
5.3.4	The Interacting-Agent Hypothesis	75
5.4	Agent-Based Computational Models of Financial Markets	75
5.4.1	Allocative Efficiency with Zero-Intelligence Traders	76
5.4.2	Models with a Random Communication Structure	79
5.4.3	Models of Chartist-Fundamentalist Interactions...	83
5.4.4	Many-Strategy Models with Learning	85

Part II The Santa Fe Institute Artificial Stock Market Model Revisited

6	The Original Santa Fe Institute Artificial Stock Market	91
6.1	Introduction	91
6.2	The Marimon-Sargent Hypothesis and the SFI-ASM	92
6.3	An Overview of SFI-ASM Versions	93
6.4	The Basic Structure of the SFI-ASM	94
6.4.1	Trading Rules and Expectation Formation	95
6.4.2	Learning and Rule Evolution	99
6.4.3	Other Programming Details and Initialization of Model Parameters	101
6.5	The Homogeneous Rational Expectations Equilibrium ..	102
6.6	The Marimon-Sargent Hypothesis Refined	103
6.7	Simulation Results of the SFI-ASM	104

6.7.1	Time Series Behavior	104
6.7.2	Forecast Properties	106
6.8	A Potential Problem: A Biased Mutation Operator	108
7	A Suggested Modification to the SFI-ASM	113
7.1	Introduction	113
7.2	An Unbiased Mutation Operator	114
7.3	Simulation Results with the Modified SFI-ASM	115
7.3.1	Trading Bit Behavior	115
7.3.2	Time Series Properties	118
7.4	Robustness of the Zero-Bit Solution	121
7.4.1	Stochastic versus Periodic Dividends and the Classifier System	121
7.4.2	Dependence on Other Parameter Values	122
7.4.3	Generalization or Consistent Trading Rules?	123
8	An Analysis of Wealth Levels	127
8.1	Introduction	127
8.2	Wealth Levels in the SFI-ASM: An Economic(al) Explanation	128
8.3	Previous Studies Based on Wealth Levels in the SFI-ASM	129
8.3.1	Financial Markets Can Be at Sub-Optimal Equilibria	129
8.3.2	Technical Trading as a Prisoner’s Dilemma	131
8.4	Wealth Levels in the SFI-ASM: Alternative Explanations	133
8.4.1	Risk-Premium, Taxation, and Two Benchmark Wealth Levels	133
8.4.2	Average Stock Holdings and Wealth Levels	135
8.4.3	Activated Rules and Rule Selection	139
8.5	A Verdict on Wealth Analysis in the SFI-ASM	145
9	Selection, Genetic Drift, and Technical Trading	147
9.1	Introduction	147
9.2	Technical Trading and the Aggregate Bit Level	148
9.3	The Zero-Bit Solution: Some Disturbing Evidence	150
9.4	Random Genetic Drift in Genetic Algorithms	152
9.5	The Neutralist–Selectionist Controversy	154
9.6	Fitness Driven Selection or Genetic Drift?	157
9.6.1	Selection or Genetic Drift in the Modified SFI-ASM?	157

9.6.2	Selection or Genetic Drift in the Original SFI-ASM?	159
9.6.3	Genetic Drift, Fitness Gradient, and Population Size	161
9.7	The Effect of Mutation on Genetic Drift	162
9.7.1	Genetic Drift, Mutation, and Crossover Only in SFI Agents	162
9.7.2	Genetic Drift, Mutation, and Crossover Only in Bit-Neutral Agents	166
9.7.3	An Equilibrium Analysis of Genetic Drift and Mutation	166
9.7.4	A Final Assessment of the Two Mutation Operators	171
9.8	Detection of Emergence of Technical Trading	172
9.8.1	Predictability in the Price Series	172
9.8.2	Trading Bits and Fitness Values	173
9.8.3	Equilibrium Bit Frequencies and Bit Fixations ...	176
9.9	An Evolutionary Perspective on Technical Trading	177
10	Summary and Future Research	181
Appendix	187
11.1	Timing in the Stock Market	187
11.2	Fundamental and Technical Trading Bits	189
References	195
Index	227

Agent-Based Modeling in Economics

Introduction

What I cannot create, I do not understand.
Richard P. Feynman

In addition to deduction and induction, simulation is sometimes seen as a third methodology for doing research. Even though simulation does not prove theorems, it can enhance our understanding of complex phenomena that have been out of reach for deductive theory. Tesfatsion defines *agent-based simulations* as the computational study of economies that are modeled as evolving systems of autonomous interacting agents [421]. In the last decade, they have become a widely accepted tool for studying decentralized markets.

A major advantage is that agent-based models allow the removal of many restrictive assumptions that are required by analytical models for tractability. For instance, all investors could be modeled as heterogeneous with respect to their preferences, endowments, and trading strategies.

Among the numerous agent-based simulations of financial markets [247, 268, 81], the Santa Fe Institute Artificial Stock Market (SFI-ASM) is one of the pioneering models and thus, probably the most well-known and best studied. It was created by a group of economists and computer scientists at the Santa Fe Institute in New Mexico to test whether artificially intelligent agents would converge to the homogeneous rational expectations equilibrium or not. The original SFI-ASM has been described in a series of papers [331, 11, 330, 244].

In agent-based simulations, replication of existing models from scratch is an important, but often neglected step. Axelrod emphasizes that without this outside confirmation, possible erroneous results

based on programming errors or misinterpretation of simulation data can go undetected [13]. Skepticism towards computational models is sometimes voiced since their results are seen as counterintuitive and incomprehensible because the computer program remains an impenetrable “black box”. Judd considers the black box criticism to be more a result of poor exposition of computational work and a lack of sufficient sensitivity analyses. Since third party replications need to open that black box widely, that criticism can effectively be addressed [208].

This book starts by presenting the reasons that led to the adoption of the agent-based programming approach in economics. Since one of these reasons, the desire to replace the representative agent approach with heterogeneous agents, leads to a break down of rational expectation formation, chapter 3 contrasts several concepts of agent rationality. When modeling less than fully rational agents, researchers have to equip their agents with learning algorithms which are discussed in chapter 4.

The replication of selected stylized facts of financial markets through agent-based simulations of financial markets is the focus of chapter 5. It starts by introducing the Efficient Market Hypothesis (EMH). Much of our empirical knowledge about financial markets, often summarized as stylized facts, stems from attempts to either prove or disprove the EMH. Several competing market hypotheses that strive to better explain the stylized facts than the EMH are subsequently discussed. In the final part of chapter 5, a selection of agent-based models of financial markets is briefly introduced.

The main part of this book analyzes a particular Java-replication of the SFI-ASM model which was originally programmed in Objective-C. The goal is to assess whether the SFI-ASM’s result of emergent technical trading is robust to changes in the model design. To this end, the simulation results are supplemented by theoretical analyses of certain model features. The replication results of the reprogrammed Java-version are presented in chapter 6. A Markov chain analysis of the original mutation operator delivers the motivation to develop an alternative mutation operator in chapter 7. Because of the differences in simulation results, chapter 8 needs to reexamine the model structure with respect to wealth accumulation. The final chapter analyzes and compares the two mutation operators in more detail. By interpreting the insights gained from the theoretical analysis of the two mutation operators, chapter 9 concludes by offering an explanation as to why the validity of the EMH is entirely consistent with the simultaneous existence of technical trading rules.

The Rationale for Agent-Based Modeling

Imagine how hard physics would be if electrons could think.

Murray Gell-Man¹

2.1 Introduction

The ascent of powerful and affordable microcomputers and the availability of huge economic data sets have sparked the development of a rather recent branch in economic research. The rapidly growing field of computational economics is a broad concept and encompasses many different areas. An eclectic and not exhausting overview of all these different activities² can be found in the “*Handbook of Computational Economics*” by Amman et al. [8]. Since a precise definition of computational economics is not being offered by its editors, Riechmann goes as far as to suggest that every economist who uses a computer for more than mere typewriting engages in computational economics [356].

But if we emphasize the aspect of computability of economic problems, i.e., problems that allow for numerical results, the roots of computational economics could easily be dated back before computers were actually used by economists.³ Several contributions in this handbook

¹ Attributed to Gell-Mann, 1969 Nobel laureate in physics and co-founder of the Santa Fe Institute, cited by Page [329].

² Among those are, for instance, the numerical computation of Nash equilibria, deterministic and stochastic simulations, or numerical dynamic programming problems.

³ Nagurney [316], as well as Kendrick [217], mention the early contributions to computational economics in the 1950ies by Koopmans, Samuelson, or Solow. In finance, the work by Markowitz about efficient portfolio diversification needs to be mentioned at this point [288, 289].

focus on algorithms and numerical methods for finding Nash-equilibria or the solutions to dynamic nonlinear systems of equations, yet some of them were developed even before computers were. The access to computer technology just reduced the sometimes prohibitive computational costs of these algorithms and allowed them to be used for practical purposes.

Within computational economics, the field of agent-based modeling or simulation (ABM, ABS), sometimes also called microscopic simulation (MS, Levy et al. [248]) or agent-based computational economics (ACE, Tesfatsion [420]), seems currently to be the most rapidly growing discipline. This is acknowledged, for instance, by the appearance of a second volume of the “*Handbook of Computational Economics*” which will be solely dedicated to this approach [209].

In the ABS approach, model economies are built from the bottom up, i.e., they consist of many autonomous and interacting agents. It had its first breakthrough with the influential models by Schelling [375, 376] about endogenous neighborhood segregation.⁴ These models are populated by two types of agents who only care about the composition of their own small neighborhood. In particular, they do not tolerate more than a certain fraction of agents of the other type in their vicinity, however, they do not care about integration or segregation on the city level. Unsatisfied agents are allowed to move to a neighborhood that they are happy with. Schelling showed that for a wide range of the agents’ neighborhood tolerance parameter, an initially integrated city emerges to an almost completely segregated entity.⁵ Thus, Schelling demonstrated how stable macrobehavior may emerge from strictly local motives on the agents’ level, a macrobehavior that would be hard to predict by exclusively looking at individual motives. Nowadays, this phenomenon is known as emergence, a key concept of the theory of complex adaptive systems.

The agent-based approach is currently considered by many researchers as the latest revolution in economic methodology. However,

⁴ A comparison of the different versions by Schelling is given by Pancs and Vriend [332]. The most important distinction is the dimensionality of neighborhoods. In his 1969 model, a neighborhood is defined only in one dimension, i.e., agents populate a line, while his later models use two dimensional lattices. A two-dimensional web-based simulation example in NetLogo by Wilensky [436] can be found at <http://ccl.northwestern.edu/netlogo/models/Segregation>.

⁵ Pancs and Vriend [332] extend the framework by asking how individual preferences may alter the outcome. Surprisingly, the results are very robust to changes in preferences. Even if all agents strictly prefer perfect integration, neighborhood segregation will still occur.

in order to understand and properly evaluate this shift in methodology, Bankes points out that one should not focus on the progress in the computer sciences which made this whole development possible, but on the inappropriateness of traditional methods which made it necessary [22].

The deficiencies of analytical modeling can only be sought in its assumptions. Thus, the first part of this chapter discusses some of these and contrasts them with the requirements of agent-based modeling. The main culprits from an agent-based point of view are the widespread assumption of representative agents and that of rational expectations. First, it is obvious that fictitious representative agents in macroeconomic models are incapable of generating emergent phenomena. Financial markets would also be characterized by the absence of any trading activity. Secondly, fully rational and perfectly informed agents have essentially no ability to exercise free choice. Despite some suggestive rhetoric, individuals can follow only one, i.e., the rational course of action [235].

The agent-based modeling approach, on the other hand, requires neither of these two assumptions. The ability to cope with heterogeneous and boundedly rational agents makes it a perfect tool to study decentralized markets. Instead of a reductionist approach, agent-based models treat the economy as an evolving complex adaptive system consisting of many heterogeneous and interacting agents. The development of artificial financial markets has thus become a major application for the agent-based paradigm.⁶

2.2 The Representative Agent Modeling Approach

The notion of representative agents appeared in the economic literature already in the late 19th century. Edgeworth used the term ‘*representative particular*’ [104, p. 109], while Marshall introduced a ‘*representative firm*’ in his *Principles of Economics* [290].⁷ However, only after Lucas [261] had published his article about econometric policy evaluation—the famous Lucas-critique—they became the dominant macroeconomic

⁶ According to [227], scientific paradigms share two essential characteristics: First, their achievements must have enough novelty to attract a permanent group of scientists away from competing modes of scientific activity. Secondly, their open-endedness must allow for addressing many different kinds of problems. Whether the agent-based modeling approach already is or might become the next paradigm in the economic sciences is left to the reader to be answered.

⁷ A discussion of the origins of representative agents can be found in [172].

approach. Today’s representative agent models are characterized by an explicitly stated optimization problem of the representative agent, which can be either a consumer or a producer. The derived individual demand or supply curves are then, in turn, used for the corresponding aggregate demand or supply curves.

A series of papers in which the rationale for using representative agent models is convincingly set forth does not exist according to [172]. From the “*large set of introductions, paragraphs, and parenthetical asides*”, however, Hartley identifies several motives for their use. They were thought to avoid the Lucas critique, to provide rigorous microfoundations to macroeconomics, and to help build powerful Walrasian general equilibrium models.

2.2.1 Avoiding the Lucas-Critique

Before [261], macroeconomic models were often defined in terms of three vectors: y_t being the set of endogenous variables, x_t the vector of exogenous forcing variables, and ϵ_t the set of random shocks. Fixed parameters are subsumed in a vector θ .

$$y_{t+1} = F(y_t, x_t, \theta, \epsilon_t). \quad (2.1)$$

Lucas, however, criticized that it is likely that some of the the parameters contained in θ change due to a shift of policy regime λ . Aggregate quantities and prices might react differently than predicted since agents may change their behavior in a way which is not captured in the aggregate equations. For instance, agents could adapt their expectations about future inflation rates or, in the case of rational expectations, change them even before an anticipated policy shift is implemented. An attempt to exploit a potential trade-off between unemployment and the inflation rate through an expansionary monetary policy may thus be foiled. Taking the Lucas-critique into account, equation (2.1) should be rewritten as

$$y_{t+1} = F(y_t, x_t, \theta(\lambda), \mu, \epsilon_t), \quad (2.2)$$

where $\theta(\lambda)$ contains regime dependent parameters, while μ is thought to consist of truly invariable taste and technology parameters.

While Lucas offered no solution to this fundamental problem, representative agent models were soon to be thought of offering an easy escape from it. Going beyond simple aggregate relationships and analyzing the economy at a deeper level than before, macroeconomists pretended to know the structural equations from which the aggregate

supply and demand curves are derived [371]. If a policy change is announced or implemented, the representative agent simply recalculates his optimization problem, given his objective function and budget constraints. This approach also satisfied the desire for a microfoundation in the new classical sense since behavior is derived from a utility maximization problem.⁸

Hartley argues though that it is impossible to identify truly invariable taste and technology parameters [172]. Acknowledging this, economists should realize that the Lucas critique imposes a standard that no macroeconomic model can probably ever fulfil. Since representative agent models suffer from the same deficiencies as old style Keynesian macroeconomic models, the justification for their use is greatly undermined.

2.2.2 Building Walrasian General Equilibrium Models

The desire to build Walrasian general equilibrium models of the economy provides another strong motivation for using representative agents. Not only are economists interested in the existence of equilibria within these types of models, they should furthermore be unique and stable. The Arrow-Debreu framework as the modern embodiment of Walrasian models, however, is far too complex to be solved for millions of heterogeneous consumers and firms. Using a representative agent instead makes it easy to find the competitive equilibrium allocation for a model economy [370].

However, for Walrasian models to be true, their structural assumptions have to be true. False structural assumptions lead to false conclusions.⁹ Since real individuals are obviously heterogeneous with respect to their preferences and cognitive abilities, the assumption of a representative agent cannot be considered structural since it is not true. Thus, it must be superficial, i.e., irrelevant to the underlying structure of the economy. When using representative agents instead of heterogeneous individuals in a Walrasian model, the modeler must believe

⁸ According to [172], there is another view of what constitutes an appropriate microfoundation. Keynesian models, for instance, backed up their macroeconomic relationships with some explanatory story which is thought to be entirely sufficient.

⁹ Hartley contrasts the requirement of true structural assumptions in Walrasian equilibrium models with Friedman's famous dictum that the realism of assumptions is irrelevant [172]. For Friedman, the validity of a theory is based on how well its prediction match reality, no matter how realistic its assumptions.

that the actual world would not look very different from the model if populated by identical clones [172, p. 66].

There is mounting evidence, however, that this is not the case. First of all, representative agent models are usually characterized by a complete absence of trade and exchange in equilibrium [225], one of the most basic activities in a market economy.¹⁰ For instance, in the classical CAPM [389, 252, 313], there is no trade after agents have completed their initial portfolio-diversification.

2.2.3 Representative Agents and the Fallacy of Composition

It has long been known in economics that what is true for individual agents may not hold for the aggregate economy. This phenomenon is called the *fallacy of composition* [61, 172]. Together with its logical counterpart, the fallacy of division, it highlights the tension between micro- and macroeconomics. The economy as a whole is formed by many consumers and firms whose interactions may cause emergent behavior at the macroeconomic level. Correct policy recommendations for individual economic units may not work for the aggregate economy or vice versa. For instance, in times of recession, a profit maximizing firm is likely to lay off workers in order to survive, while a similar action by the government as an aggregate player will aggravate the economic downturn.

Representative agent models usually commit the fallacy of composition by ignoring valid aggregation concerns. Kirman, for instance, provides a graphical example based on [203] in which the representative agent disagrees with all individuals in the economy [225]. Policy recommendations based on the representative agent, a common practice in today's macroeconomics, are illegitimate in this case.

A rigorous treatment of this logical fallacy can be found in the literature on exact aggregation.¹¹ Gorman [160] was the first who derived general conditions under which the aggregation of individual preferences is possible. He showed that aggregate demand is dependent on income distribution unless all agents have identical homothetic utility functions. Only when their Engel curves are parallel and linear, a redistribution of income will leave the aggregate demand unaffected. Other authors [412, 203, 249] derived similar conditions for exact aggregation,

¹⁰ In the literature, there are as many no-trade theorems [306, 12, 422] as attempts to solve this apparent contradiction with economic reality. These attempts are usually characterized by a relaxation of the assumption of strict homogeneity of market participants [83, 230, 428].

¹¹ An introduction to the problem of exact aggregation can be found in [161].

the least restrictive ones given in [249]. In any case, those conditions are still so special that no economist would ever consider them to be plausible [225]. In the unlikely event of them being satisfied, no-trade situations would be the result.

For apparent reasons, the literature on exact aggregation has largely been ignored by representative agent modelers. In summarizing this body of research, Kirman concludes that the reduction of a group of heterogeneous agents to a representative agent is not just an analytical convenience, but is *“both unjustified and leads to conclusions which are usually misleading and often wrong”*. Hence, *“the ‘representative’ agent deserves a decent burial, as an approach to economic analysis that is not only primitive, but fundamentally erroneous”* [225, p. 119].

2.2.4 Expectation Formation in Markets with Heterogeneous Investors

In asset pricing models with homogeneous investors, asset prices typically reflect the discounted expected payoffs and follow a martingale, i.e., today’s expectation of next period’s price $E_t [p_{t+1}]$ just equals the current price p_t [365]. This martingale property of asset prices implies that the expectations of the representative investor satisfy the *law of iterated expectations*: $E_t [E_{t+1}(p_{t+2})] = E_t [p_{t+2}]$, that is, his expectation today of tomorrow’s expectation of future payoffs equals his current expectation of these future payoffs [6]. Since all investors are alike, individual and average expectations of future asset prices coincide.

Yet, when giving up the concept of a representative investor, the knowledge of average expectations of future payoffs becomes important for an individual investor. With differential private information and public information, an individual’s expectation and the average expectation about future payoffs are likely to diverge. Furthermore, for average expectations, the law of iterated expectations is not satisfied anymore. Thus, according to Keynes [218], real financial markets with heterogeneous agents resemble more a beauty contest in which the competitors have to choose the six prettiest faces from a hundred photographs, the winner being the one whose choice most nearly corresponds to the average preferences of the other competitors. Instead of choosing the face that one considers prettiest, participants devote their intelligence to anticipate *“what average opinion expects the aver-*