

599 LECTURE NOTES IN ECONOMICS
AND MATHEMATICAL SYSTEMS

Andrea Consiglio
(Editor)

Artificial Markets Modeling

Methods and Applications

 Springer

Lecture Notes in Economics and Mathematical Systems

599

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. Drexl, H. Dawid, K. Inderfurth, W. Kürsten, U. Schittko

Andrea Consiglio
(Editors)

Artificial Markets Modeling

Methods and Applications

With 84 Figures and 28 Tables

 Springer

Professor Andrea Consiglio
Department of Statistics and Mathematics 'Silvio Vianelli'
University of Palermo
Viale delle Scienze - Ed. 13
90128 Palermo
Italy

Library of Congress Control Number: 2007932226

ISSN 0075-8442

ISBN 978-3-540-73134-4 Springer Berlin Heidelberg New York

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.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2007

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

SPIN 12078806 88/3180YL - 5 4 3 2 1 0 Printed on acid-free paper

Preface

Agent-based computational models, generally named “Artificial economics” (AE), represent a new methodological approach where economies, and more generally social structures, are modeled as evolving systems consisting of heterogeneous interacting agents with some degree of cognitive skills. Assuming a precise mechanism that regulates the interaction among different agents, this approach allows through simulation to compute numerically the aggregate behavior of the economy and to discover the regularities emerging from the micro-behavior of the agents. The AE approach has provoked a great deal of academic interest among social scientists because it represents an alternative to both the fully flexible but not computable and testable descriptive models, and the logical consistent but highly simplified analytical models. With AE the researcher retains much of the flexibility of pure descriptive models in the specification of the interaction structure and the individual behavior, while having the precision and consistency imposed by the computer language. The methodology opens up new avenues for analyzing decentralized, adaptive, emergent systems. The use of computer simulations provides an experimental format allowing free exploration of system dynamics, and, at the same time, the opportunity to check the various unfolding behaviors for plausibility. An early use of agent-based models was by R.M. Axelrod in his research on the evolution of cooperation. He employed extensive computational simulations to study individual strategic behavior in the iterated prisoner’s dilemma. This work has stimulated a new approach to game theory based on computational ideas. The research on complex adaptive systems has received a great impulse starting from the mid-eighties with the foundation of the Santa Fe Institute, a non-profit institution specifically devoted to understand the basic principles of human and natural systems, following a multi-disciplinary approach and using computer-based modeling. A new field of scientific inquiry, called Artificial Life (AL), has emerged with the aim to study biology by attempting to synthesize biological phenom-

ena such as life, evolution, and ecological dynamics within computers. This approach has led to wider ideas such as complexity, evolution, auto-organization, and emergence that have influenced social scientists. The initial attempts to mix computational methods and social sciences include pioneering AE work in finance, specifically the “Santa Fe Artificial Stock Market Model” of W.B. Arthur, J.H. Holland, B. LeBaron, R.G. Palmer, and P. Taylor. This model, based on bounded rationality and inductive reasoning, has led to a new generation of agent-based computational models aimed to reproduce stock market dynamics and to explain financial market puzzles. Recently, there has been a surge of interest in studying social interaction, the process by which people form and transmit ideas and information. The emergence of this new topic has been driven by the recognition that understanding the formation and dynamics of social networks may represent the missing element to uncover the functioning of complex systems such as asset markets. Agent-Based Computational Economics, with its intrinsic multidisciplinary approach, is gaining increasing recognition in the social sciences. The methodology is now widely used both to compute numerically analytical models and to test them for departures from theoretical assumptions, and to provide stand-alone simulation models for problems that are analytically intractable.

This book collects a selected range of refereed papers that have been gathered in five sections, each of them devoted to one the following topics:

- Macroeconomic Issues
- Market Mechanisms and Agents Behavior
- Market Dynamics and Efficiency
- Analysis of Economic and Social Networks
- Methodological Issues and Applications

The first section includes papers using an agent-based approach to give micro-foundations to macro-economic analyses. The second section is dedicated to papers developing agent-based computational models aimed to investigate the dynamics of financial markets in order to understand their properties. In this section relevant issues such as the fairness of different trading mechanisms and the evaluation of the performance of technical trading are analyzed. The third section is devoted to models simulating the process of market adjustment towards equilibrium. The section covers different interesting applications spanning from the introduction of an option market, to a model with endogenous costly information acquisition. The fourth section is devoted to papers investigating networks formation and evolution with applications to

the labor market and to the R&D industry. Finally, the last section includes more methodological contributions and some applications such as a model of the venture capital market where the quality of the investment projects is only imperfectly available and venture capitalists play the function of screening high-quality investments.

I would like to thank all the members of the Scientific Committee for their invaluable effort in refereeing more than 60 papers:

- Frédéric AMBLARD - Université de Toulouse 1, France
- Gérard BALLOT - ERMES, Université de Paris 2, France
- Bruno BEAUFILS - LIFL, USTL, France
- Olivier BRANDOUY - CLAREE, USTL, France
- Charlotte BRUUN, Aalborg University, Denmark
- José Maria CASTRO CALDAS - ISCTE, DINAMIA, Portugal
- Silvano CINCOTTI - University of Genova, Italy
- Christophe DEISSENBERG - GREQAM, France
- Jean-Paul DELAHAYE - LIFL, USTL, France
- Wander JAGER - University of Groningen, The Netherlands
- Marco JANSSEN - Arizona State University, USA
- Philippe LAMARRE - LINA, Université de Nantes, France
- Marco LI CALZI - University of Venice, Italy
- Michele MARCHESI - University of Cagliari, Italy
- Luigi MARENGO - St. Anna School of Advanced Studies, Pisa, Italy
- Philippe MATHIEU - LIFL, USTL, France
- Nicolas MAUDET - Université Paris 9 Dauphine, France
- Akira NAMATAME - National Defense Academy, Japan
- Paolo PELLIZZARI - University of Venice, Italy
- Denis PHAN - Université de Rennes I, France
- Juliette ROUCHIER - GREQAM, France
- Elpida TZAFESTAS - National Technical University of Athens, Greece
- Murat YIDILZOGLU - IFREDE-E3i, Université Montesquieu Bordeaux IV, France
- Stefano ZAMBELLI - Aalborg University, Denmark

A special thank goes to Valerio Lacagnina and Annalisa Russino who helped me for the preparation of this volume.

A final acknowledgment goes to the Ministero dell'Università e della Ricerca that, under the PRIN04 project "*Models for the price dynamics of financial securities: institutional aspects and behavioral assumptions in a agent-based framework*", contributed to the publishing expenses.

Contents

Part I Macroeconomic Issues

- 1 Beyond the Static Money Multiplier: In Search of a Dynamic Theory of Money**
Michele Berardi 3
- 2 Macroeconomic Effects of the Interest Rate Level: Growth and Fluctuations in an Economy with Bank Capital Adequacy Standards**
Gianfranco Giulioni 17
- 3 Monetary Policy Experiments in an Artificial Multi-Market Economy with Reservation Wages**
Marco Raberto, Andrea Teglio, Silvano Cincotti 33

Part II Market Mechanisms and Agents Behavior

- 4 Testing Double Auction as a Component Within a Generic Market Model Architecture**
Julien Derveeuw, Bruno Beaufils, Philippe Mathieu, Olivier Brandouy 47
- 5 A Conceptual Framework for the Evaluation of Agent-Based Trading and Technical Analysis**
Olivier Brandouy, Philippe Mathieu 63
- 6 Which Market Protocols Facilitate Fair Trading?**
Marco LiCalzi, Paolo Pellizzari 81

Part III Market Dynamics and Efficiency

- 7 An Artificial Economics View of the Walrasian and Marshallian Stability**
Marta Posada, Cesáreo Hernández, Adolfo López-Paredes 101
- 8 The Performance of Option–Trading Software Agents: Initial Results**
Omar Baqueiro, Wiebe Van der Hoek, Peter McBurney 113
- 9 Studies on the Impact of the Option Market on the Underlying Stock Market**
Sabrina Eccia, Mario Locci, Michele Marchesi 127
- 10 On Rational Noise Trading and Market Impact**
Florian Hauser 141

Part IV Analysis of Economic and Social Networks

- 11 A Note on Symmetry in Job Contact Networks**
Andrea Mario Lavezzi, Nicola Meccheri 157
- 12 Innovation and Knowledge Spillovers in a Networked Industry**
Jose I. Santos, Ricardo del Olmo, Javier Pajares 171
- 13 Heterogeneous Agents with Local Social Influence Networks: Path Dependence and Plurality of Equilibria in the ACE Noiseless Case**
Denis Phan 181
- 14 Economy-Driven Shaping of Social Networks and Emerging Class Behaviors**
Philippe Caillou, Frederic Dubut, Michele Sebag 195
- 15 Group Effect, Productivity and Segregation Optimality**
Raúl Conejeros, Miguel Vargas 209
- 16 The Grass is Always Greener on the Other Side of the Fence: The Effect of Misperceived Signalling in a Network Formation Process**
Simone Giansante, Alan Kirman, Sheri Markose, Paolo Pin 223

Part V Methodological Issues and Applications

17 Market Selection of Competent Venture Capitalists	
<i>David Mas</i>	237
18 A Binary Particle Swarm Optimization Algorithm for a Double Auction Market	
<i>Calogero Vetro, Domenico Tegolo</i>	249
19 Better-Reply Strategies with Bounded Recall	
<i>Andriy Zapechelnyuk</i>	259

List of Contributors

Baqueiro Omar

Department of Computer Science
University of Liverpool, UK
omar@csc.liv.ac.uk

Beaufils Bruno

Laboratoire dInformatique
Fondamentale de Lille
USTL, France
beaufils@lifl.fr

Berardi Michele

School of Social Sciences
University of Manchester, UK
Michele.Berardi@manchester.ac.uk

Brandouy Olivier

Lille Economie et Management
USTL, France
olivier.brandouy@univ-lille1.fr

Caillou Philippe

LRI
Université Paris Sud, France
caillou@lri.fr

Cincotti Silvano

DIBE-CINEF
University of Genova, Italy
cincotti@dibe.unige.it

Del Olmo Ricardo

University of Burgos, Spain
rdelolmo@ubu.es

Conejeros Raúl

Department of Biochemical
Engineering
Catholic University of Valparaíso,
Chile
rconej@ucv.cl

Derveeuw Julien

Laboratoire dInformatique
Fondamentale de Lille
USTL, France
derveeuw@lifl.fr

Dubut Frederic

LRI
Université Paris Sud, France
dubut@lri.fr

Ecce Sabrina

DIEE
University of Cagliari, Italy
sabrina.ecce@diee.unica.it

Giansante Simone

CCFEA
University of Essex, UK
sgians@essex.ac.uk

XIV List of Contributors

Giulioni Gianfranco

Department of Quantitative
Methods and Economic Theory
University of Chieti-Pescara,
Italy
g.giulioni@unich.it

Hauser Florian

Department of Banking and
Finance
University School of Management,
Innsbruck, Austria
florian.hauser@uibk.ac.at

Hernández Cesáreo

INSISOC
University of Valladolid, Spain
cesareo@insisoc.org

Kirman Alan

GREQAM
Université d'Aix Marseille, France
kirman@ehess.univ-mrs.fr

Lavezzi Andrea Mario

Department of Studies on Politics,
Law and Society
University of Palermo, Italy
lavezzi@unipa.it

LiCalzi Marco

Dept. Applied Mathematics and
SSAV
University of Venice, Italy
licalzi@unive.it

Locci Mario

DIEE
University of Cagliari, Italy
mario.locci@diee.unica.it

López-Paredes Adolfo

INSISOC
University of Valladolid, Spain
adolfo@insisoc.org

Marchesi Michele

DIEE
University of Cagliari, Italy
michele@diee.unica.it

Markose Sheri

CCFEA
University of Essex, UK
scher@essex.ac.uk

Mas David

ERMES-CNRS
Université Panthéon-Assas Paris
II, France
David.Mas@u-paris2.fr

Mathieu Philippe

Laboratoire d'Informatique
Fondamentale de Lille
USTL, France
mathieu@lil1.fr

McBurney Peter

Department of Computer Science
University of Liverpool, UK
peter@csc.liv.ac.uk

Meccheri Nicola

Department of Economics,
University of Pisa, Italy
meccheri@ec.unipi.it

Pajares Javier

University of Valladolid
pajares@eis.uva.es

Pellizzari Paolo

Dept. Applied Mathematics and
SSAV
University of Venice, Italy
paolop@unive.it

Phan Denis

GEMAS UMR, 8598 CNRS &
University of Paris IV
Sorbonne CREM UMR, 6211
CNRS & University of Rennes 1,
France
dphan@msh-paris.fr

Pin Paolo

Economics Department Univer-
sity of Venice,
and ICTP, Trieste, Italy
pin@unive.it

Posada Marta

INSISOC
University of Valladolid, Spain
posada@insisoc.org

Raberto Marco

DIBE-CINEF
University of Genova, Italy
raberto@dibe.unige.it

Santos Jose I.

University of Burgos, Spain
jisantos@ubu.es

Sebag Michele

LRI
Université Paris Sud, France
sebag@lri.fr

Teglio Andrea

DIBE-CINEF
University of Genova, Italy
teglio@dibe.unige.it

Tegolo Domenico

Department of Mathematics and
Applications
University of Palermo, Italy
tegolo@math.unipa.it

Van der Hoek Wiebe

Department of Computer Science
University of Liverpool, UK
wiebe@csc.liv.ac.uk

Vargas Miguel

Diego Portales
University School of Business,
Chile
miguel.vargas@udp.cl

Vetro Calogero

Department of Mathematics and
Applications
University of Palermo, Italy
vetro@math.unipa.it

Zapechelnyuk Andriy

Center for Rationality
the Hebrew University, Israel
andriy@vms.huji.ac.il

Macroeconomic Issues

Beyond the Static Money Multiplier: In Search of a Dynamic Theory of Money

Michele Berardi

University of Manchester
School of Social Sciences
`Michele.Berardi@manchester.ac.uk`

1.1 Introduction

Though we all live in a monetary economy where credit money plays a fundamental role, the process through which money is created in the economy is largely neglected by modern macroeconomic theory. A common approach maintains that the process starts with an exogenous increase in the monetary base made by the central bank, and that this, through a fixed multiplier, gives rise to a proportional increase in the amount of money in the economy. The multiplier is usually taken as constant in this process, at least on short time scales, and most importantly, independent from the money creation process itself. The result is essentially a static, aggregate theory, with very poor behavioral micro-foundations, that completely neglects the *process* through which money is generated in an economy.

As a consequence of this representation, money is taken to be exogenously determined and its quantity explained through changes in the monetary base magnified proportionally by the fixed multiplier. Unfortunately, this theory is not able to provide any insights about the process that generates money in a credit economy, apart from assuming that changes in the monetary stock are originated by central bank interventions, and proportional to them. It misses completely the idea that money is created and destroyed endogenously, through the interactions of the many actors (mainly banks, households and firms) participating in the monetary and credit markets.

An important drawback of the traditional theory, as represented by the static multiplier,¹ is that it does not allow for a proper theory

¹ We dub the traditional multiplier as static, to emphasize its lack of attention to the dynamics involved in the process of money creation.

of endogenous money creation that many economists think would be necessary.² Presenting the whole process of money creation as a pure deterministic response of the monetary stock to an exogenous change in the monetary base is deeply misleading. In the words of Goodhart (1984), the standard multiplier theory of money creation is “... such an incomplete way of describing the process of the determination of the stock of money that it amounts to misinstruction”.

In modern economies, where the central bank wants to control the interest rate, money is necessarily endogenous to the system as the policymaker must provide enough monetary base so that the equilibrium interest rate on the market is the desired one. Though this fact is often recognized even in standard macroeconomic textbooks, then an exogenous and fixed multiplier is still considered to be the link between the monetary base and the amount of money available in the economy. It is completely neglected the fact that the ratio between these two aggregates can vary according to the behavior of the system and must not be assumed fixed *a priori*.³

In this work we take a narrow perspective regarding the creation of money in a credit economy and focus our attention only on its *process*. In particular, our analysis should help explain the short term variability in the amount of money, for the part that can be imputed to the volatility in the multiplier.⁴ Our work does not try to analyze the determinants of the behavior of banks and households but puts emphasis on the heterogeneity of the actors involved in the monetary and credit market and tries to provide a better understanding of the dynamics of the process of money creation, stripped down to its mechanics and deprived of any behavioral content. Still, we believe that this approach can provide useful insights and help build a more comprehensive theory of money in a credit economy.

² Post-Keynesian economists, in particular, have long argued about the need of an endogenous theory of money, one that recognizes the fact that the financial system is able to generate monetary liabilities in response to real sector's needs. But also on the other side of the macroeconomics spectrum (see, e.g., Kydland and Prescott, 1990) there is support for the view of endogenous money.

³ These issues are somewhat related to the debate between verticalists and horizontalists that was popular in the 1970s. For a detailed exposition and analysis of the two positions, see Moore (1988).

⁴ Moore (1988) shows that variations in the monetary base can explain only about 40% of the variability in the M1 aggregate on a monthly base, while this proportion raises to about 65% with quarterly data and to 90% over horizons of one year. Over short time horizons, therefore, a lot of variability in M1 is left unexplained by the standard theory.

1.2 Models of money creation

1.2.1 The static multiplier

Standard macroeconomic theory explains the amount of money available in an economy starting from the monetary base (H), which is composed of currency held by the public (CU) and reserves held by the banking sector (R).⁵ The money multiplier is simply derived as the ratio between the monetary base provided by the central bank and a monetary aggregate (M), composed of currency (CU) and deposits (D):⁶

$$H = CU + R \quad (1.1)$$

$$M = CU + D, \quad (1.2)$$

from which, dividing everything by D and defining $cu = CU/D$, $re = R/D$, it follows that

$$m = \frac{M}{H} = \frac{1 + cu}{cu + re}. \quad (1.3)$$

The standard money multiplier represents therefore an aggregate characteristic of the economy, with essentially no behavioral content. Nevertheless, the ratios re and cu are often taken to represent agents' individual preferences, assumed constant and homogeneous. The whole approach is essentially static and neglects completely the process through which money is created.

1.2.2 A dynamic version of the multiplier

We present here a different way to obtain the multiplier: instead of using ratios of aggregate quantities, we consider the dynamic process that unfolds through monetary and credit transactions. We start with an increase in monetary base, in the form of an increase in funds available to the public. Suppose we are in a situation where households have exactly the proportion of cash/deposits (cu) that they wish, and banks have the proportion of reserve/deposit (re) that they want to hold. Therefore households will split the additional funds they receive

⁵ It is customary not to distinguish between households and firms, and consider them as an aggregate entity (the public). We will follow here this simplification as well.

⁶ In this work we will refer to a generic monetary aggregate M , which could be understood as $M1$ in US or Europe.

between deposits and cash, in the proportion cu . Banks in turn will keep a fraction (re) of the additional deposits they receive as reserves and use the rest to extend new loans (L) to the public, who will split them again into cash and deposits, and the process continues.⁷

From the definitions above, we get that at each step i of the process:⁸

$$CU_i = \frac{cu}{1 + cu} L_i \quad (1.4)$$

$$D_i = \frac{1}{1 + cu} L_i \quad (1.5)$$

$$L_{i+1} = (1 - re) D_i \quad (1.6)$$

which lead to

$$M_i = \left(\frac{1 - re}{1 + cu} \right)^i M_0 \quad (1.7)$$

and therefore

$$m = \frac{\sum_{i=0}^{\infty} M_i}{M_0} = \sum_{i=0}^{\infty} \left(\frac{1 - re}{1 + cu} \right)^i = \frac{1}{1 - \frac{1-re}{1+cu}} = \frac{1 + cu}{cu + re}, \quad (1.8)$$

where M_0 is the original increase in monetary base, in the currency component. This alternative derivation of the static multiplier shows its micro-foundations when the behavioral parameters cu and re are constant and homogeneous. But once we introduce heterogeneity in those individual parameters, the system changes significantly its behavior.

To better analyze the importance of heterogeneity, the aggregate description for the process (1.4)-(1.8) must be replaced with a distributed one, where each single bank and household are represented and explicitly considered. This implies that in general a closed form solution for the multiplier will not exist, and computer simulations will be used to gain insights into the behavior of the system.

⁷ The following restrictions apply: $0 \leq re \leq 1$, $cu \geq 0$.

⁸ Here CU_i is the additional amount of cash available at time i with respect to time $i - 1$, not the total cash available at time i . The same for the other variables here used.

1.2.3 Introducing heterogeneity

In a heterogeneous setting, each bank has its own reserve/deposit ratio and each household its own currency/deposit ratio. If we assume that each agent (bank or household) is linked to only one agent of the other type, so that the flow of money is never split into different streams, it is then possible to express the multiplier (for a unitary increase in the monetary base) as

$$m_d = 1 + \sum_{i=0}^{\infty} \left(\prod_{j=1}^i \frac{1 - re_j}{1 + cu_j} \right), \quad (1.9)$$

where the index i refers to a “round” in the process (i.e., household i deposits money in bank i ; bank i extends a loan to household $i + 1$, who will deposit money into bank $i+1$). A bank or household can be activated in more than one round during the process, as the index does not identify an agent uniquely, only the action of an agent.

We can see that if $re_z = 1$, or $cu_z = \infty$, for some generic z , then the terms in (1.9) for $i \geq z$ are all zero, because agent z acts as an absorbing state in the system and interrupts the multiplicative process of money creation. This implies that heterogeneity is important, and can not be simply averaged out. In fact, the value of the multiplier computed with (1.9) is different from the one we would obtain by using averages of all the reserve/deposit and currency/deposit ratios:

$$m_a = \frac{1 + \frac{1}{n} \sum_{h=1}^n cu_h}{\frac{1}{n} \sum_{h=1}^n cu_h + \frac{1}{k} \sum_{b=1}^k re_b}, \quad (1.10)$$

where k is the number of banks and n the number of households in the economy. Here indexes represent individual banks and households. Under homogeneity ($\forall b, re_b = re$; $\forall h, cu_h = cu$), (1.8) = (1.9) = (1.10). But with heterogeneous agents, this is not in general true, as it can be seen from a simple experiment. We create 100 different economies, each characterized by 1000 banks and 1000 households with randomly drawn individual ratios and derive the empirical cumulative distribution function (cdf) for the dynamic multiplier computed using (1.9) and for the one computed using averages as in (1.10). As can be seen in Figure 1.1, the average multiplier m_a varies over a restricted range of values, as much of the variability is washed out by the averaging.




Fig. 1.1. Empirical CDF of average (dotted line) and dynamic (solid line) multipliers.

When the behavioral parameters are heterogeneous, the value of the dynamic multiplier depends, among other things, on the position where the process starts (for an exogenous intervention, where the CB “drops” the monetary base). The system is in fact path dependent and the order by which agents take part in the process becomes relevant. This is confirmed by our simulations when we compute the dynamic multiplier 1000 times for the same economy, each time changing the order by which agents are activated. Results show that the multiplier can vary over a wide range of values, *for the same economy*, depending on the order by which agents take part in the process.⁹

The standard way to represent the multiplier is therefore misleading, as in that representation the coefficients re and cu are not really behavioral parameters, as it may appear by their definitions, but simply ratios of aggregate quantities.

Note then that equation (1.9) is valid only when all the money remains in a unique stream and never gets split into different branches. If we allow each agent (bank or household) to be connected with more than one counterpart, we then need to keep track of all the streams

⁹ In one of the experiments that we ran, the dynamic multiplier showed a distribution of values in the interval 1-2.5. Of course m_a was instead constant (and equal to 1.06).

of money that get generated, and the analytic formula becomes intractable.

1.2.4 Monetary network

We therefore build an artificial economy and try to gain some insights into the process of money creation by means of simulations. We abstract from any considerations involving the real side of the economy and only model the structure of monetary and credit transactions, considering different possible network topologies at the base of the system and their impact on the multiplicative process.

The network composed of banks and households is a bipartite network, where edges exist only between nodes belonging to different classes. In the process that we describe, each node (bank or household) receives some money from its incoming links, keeps part of it (as reserves or cash holdings) and passes along the rest through the outgoing edges. We can uniquely define each node by its ratio of reserve/deposit or currency/deposit, and build two matrices, one for the links from banks to households (where the edges of this network represent the flow of credit that banks extend to households), and one for the links from households to banks (where the edges represent the flow of deposits from households to banks).

We will consider three different network topologies and try to understand how they impact on the size distribution of the multiplier: a random graph, a regular graph and star graph. Other topologies of course could be considered (e.g., small-world á la Watts and Strogatz (1998) or scale-free á la Albert and Barabasi (2002)), but we restrict for now to these more common structures.

We start by considering a random network, where banks and households are assigned random behavioral ratios (cu_h and re_b)¹⁰ and are randomly linked to each other. The system is composed of 5 banks and 100 households, with each bank receiving money from and extending loans to a random number of households. We simulate 100 economies and compute for each the average and the dynamic multiplier. In Figure 1.2 we show the distributions (as histograms) of the two measures. We can see that the variability in the dynamic multiplier is much higher than in the average one, where the part due to heterogeneity gets washed out.

We then consider one single economy with a fixed set of parameters (thus fixing the average multiplier) and simulate 1000 different pro-

¹⁰ With re_b and $\frac{cu_h}{1+cu_h}$ uniformly distributed between 0 and 1.




Fig. 1.2. Histograms of average (grey) and dynamic (black) multiplier with a random network of monetary transactions.

cesses of money creation by randomly inject money on different sites. Figure 1.3 shows the empirical cdf of the resulting dynamic multiplier: as it can be seen, the monetary system is path dependent and the final size of the money multiplier depends, among other things, on the position where money is injected into the economy. This means that the multiplier could change even when behavioral ratios for banks' reserves and households' currency remain fixed, an aspect that is completely neglected by the standard theory.

The next topology that we consider is a regular structure, where banks and households are laid down on a bi-dimensional lattice. Each bank is linked to four households, and each household to four banks. Each link is bi-directional, for deposits and loans (though some can have zero weight). We simulate the process of money creation on a lattice composed by 18 banks and 18 households, and show the distribution (histograms) for the average and the dynamic multipliers in Figure 1.4. Compared with the case of a random graph, the variability in the dynamic multiplier is now reduced, as the presence of absorbing states does not disconnect entire regions of the system.

To conclude, we look at the extreme case of a star topology, where all households are linked to one single bank which receives deposits and extends loans to them. We simulate the process of money creation on a structure of this type with 100 households and one bank, and show




Fig. 1.3. Empirical cdf of the dynamic multiplier in a random economy with different paths of propagation.




Fig. 1.4. Histograms of average (grey) and dynamic (black) multiplier with a regular network of monetary transactions.

the results in Figure 1.5. As we can see, the variability in the dynamic multiplier increases again now, because the presence of only one bank makes the whole system dependent on the behavior of that bank.




Fig. 1.5. Histograms of average (grey) and dynamic (black) multiplier with a star network of monetary transactions.

1.2.5 Monetary cascades and the sandpile model: an attempt at perspective

We try to suggest here an alternative but somewhat complementary interpretation of the process through which money is created in a credit economy, viewing it as an avalanche that propagates across the economy through monetary and credit transactions.

An interesting phenomenon that has been studied in physics is that of self-organized criticality (SOC), where a system drives itself on the edge of a critical state, right between stability and instability.¹¹ The classical example is that of the sandpile model developed by Bak et al. (1987).

We think that this interpretation could provide useful insights for the explanation of the process of money creation in a credit economy. If the system operates right on the edge of a critical state, the introduction of new monetary base could have a final effect on the monetary

¹¹ For a review of the concept, see Turcotte (1999).

aggregate that is unpredictable and can vary across a wide range of values.

Suppose that banks try to keep an average reserve/deposit ratio in line with legislation requirements, but take actions and extend new loans only when their individual reserve/deposit ratio reaches a fixed threshold; and that households try to keep an average currency/deposit ratio according to their individual needs/preferences, but take actions and deposit funds into a bank only when their ratio reaches a certain upper bound. So that when banks extend new loans and households make new deposits, they will do it for an amount that exceeds the marginal availability of funds beyond their own threshold.¹² In this way, as time passes, the system could drive itself towards a critical state, on the edge between stability and instability.

Once in this critical state, for each increase in monetary base we could see a final increase in the monetary aggregate M of any size. At times, the process of money creation would end soon, when money reaches an agent that is below its threshold and therefore hoards the additional money he receives; but at times the process could spread out and generate an avalanche, if many nodes involved reach their own threshold and pass along money to others.

This interpretation could provide a good explanation of the variability observed in the multiplier, and if the analogy with the sandpile model is correct, the size of monetary cascades should be distributed according to a power-law.^{13 14}

We now turn to data to see if a power law characterizes the size of the multiplier. In this respect, there are a number of issues to keep in mind. First, the central bank does not "drop" monetary base constantly and regularly in fixed amounts in the economy; secondly, the temporal scale is such that different avalanches may overlap, as there is no guarantee that the time between one central bank intervention and the next is enough for the system to fully respond and adjust to the

¹² Technically, these behaviors prevent the system from reaching a stationary state of equilibrium, where all agents have just the desired reserve and currency ratios and simply pass along any additional funds they receive.

¹³ A feature that is crucial in the sandpile model is the dispersion of the sand involved in the avalanche. In the monetary system, of course, there is no dispersion of money, so that the "pile" of money keeps growing in absolute size, but the relative size with respect to deposits, that is what matters here, remains constant.

¹⁴ While earlier studies of the sandpile model were done using a regular lattice to represent the interactions among sand grains, Goh et al. (2003) study the avalanche dynamics of the sandpile model on a scale-free network with heterogeneous thresholds and find that the avalanche size distribution still follows a power law.

first intervention; third, we have data available at regularly intervals (bi-weekly or monthly), but an avalanche of money may take different lengths of time to reach its full extent at different times; finally, we detrend the multiplier, as its trend is likely to derive from long-run changes in behaviors that we do not try to explain here and want to abstract from.¹⁵ Having all these limitations in mind, we test for the presence of a power law in the size distribution of the multiplier.¹⁶ Figure 1.6 (in a log-log scale) shows the best fit of the estimated Pareto distribution for the right tail (dashed-dotted line) with the vertical dotted line showing the point from which the Pareto distribution has been identified. Out of the 568 observations available (bi-weekly data for US, February 1984–November 2005),¹⁷ only 157 were identified to be distributed according to a power-law, and the estimated coefficient is 2.55.

According to this test, the evidence for a Pareto distribution in the data for the multiplier seems rather weak so far, though we believe that a more careful analysis is required. In particular, it has to be identified the measure that better captures the avalanche style behavior of the system, since the multiplier, suffering from the limitations described above, might be a poor indicator of such a behavior.

1.3 Conclusions

This paper is a tentative contribution in the field of monetary economics and offers a representation of the money creation process in a credit economy that is alternative to the standard one provided by the static multiplier. We have focused our attention on the mechanics of the process, and we have shown the importance of the role played by the heterogeneity of the actors involved and their interactions. An important feature that has been shown here is the path dependence of

¹⁵ The series is detrended using the Hodrick-Prescott filter.

¹⁶ We apply a procedure that first tests for the presence of a Pareto distribution in the data, identifies a region that with a 95% confidence interval follows such a distribution and then applies bootstrapping techniques to find the Hill estimator for the coefficient of the distribution.

¹⁷ We also applied the same procedure to a constructed series for the multiplier, obtained as the ratio between the monetary aggregate M1 and the monetary base, using US monthly data for the period 01/1959-08/2006, with the resulting multiplier then detrended using the HP filter. We obtained similar results in terms of the proportion of data appearing to be Pareto distributed, though the estimate for the coefficient was lower, about 2.25.




Fig. 1.6. Empirical distribution of the detrended money multiplier. The dashed-dotted line indicates the best fit for a power-law.

the system, which implies that position and timing of CB's interventions on the money market will have an impact on their effectiveness. Finally, the structure of the monetary system has been shown to affect the variability of the multiplier and therefore the process of money creation. It is therefore important that some effort be devoted in order to understand the empirical structure of monetary and credit transactions.

The approach we have adopted in this work, we believe, is well suited for supporting a theory of endogenous money, as it does not imply a deterministic and causal relationship between the monetary base and the quantity of money. Emphasis is placed on the monetary and credit transactions, and though we did not try to link these transactions to the economic activity, the two aspects are clearly interrelated.

Our analysis is just an initial step and much road has still to be covered in order to develop a theory that can properly account for the process of money creation, but we hope that our work will stimulate others to join the ride.

References

- R. Albert and A.L. Barabasi. Statistical mechanics of complex networks. *Review of Modern Physics*, 74:47, 2002.

- P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: an explanation of $1/f$ noise. *Physical Review Letters*, 59(4):381–384, 1987.
- K.-I. Goh, D.-S. Lee, B. Kahng, and D. Kim. Sandpile on scale-free networks. *Physical Review Letters*, 91(14):148701, Oct 2003.
- C.A.E. Goodhart. *Monetary theory and practice*. Macmillan, London, 1984.
- F.E. Kydland and E.C. Prescott. Business cycles: real facts and a monetary myth. *Federal Reserve Bank of Minneapolis Quarterly Review*, pages 3–18, Spring 1990.
- B.J. Moore. *Horizontalists and verticalists: The macroeconomics of credit money*. Cambridge University Press, New York, 1988.
- D.L. Turcotte. Self-organized criticality. *Reports on Progress in Physics*, 62:1377–1429, 1999.
- D.J. Watts and S.H. Strogatz. Collective dynamics of “small world” networks. *Nature*, 393:440–442, 1998.

Monetary Policy Experiments in an Artificial Multi-Market Economy with Reservation Wages

Marco Raberto, Andrea Teglio, and Silvano Cincotti

DIBE-CINEF, University of Genova, Italy
{cincotti, raberto, teglio}@dibe.unige.it

3.1 Introduction

The agent-based framework provides an useful computational facility for economics, where performing experiments on policy design issues in a realistic environment, characterized by non-clearing markets and bounded rational agents (see Tesfatsion and Judd, 2006, for a recent survey). Under this respect, this study addresses the issue of monetary policy design by investigating an appropriate rule for the central bank interest rate. Our work consists in pursuing a general equilibrium approach to the problem by considering a multi-market economy characterized by a goods, a labor and a credit market, where agents are price makers on the supply side and act according to sensible rules of thumb. A previous paper (Raberto et al., 2006) by the authors showed the absence of real effects of monetary policy in an agent-based model characterized by price-taking agents. However, if agents are price makers, prices may be set far away from their market clearing values, thereby allowing potential real effects of monetary policy.

The concept of price stickiness as a source of monetary non-neutrality is central in the new-Keynesian literature (Clarida et al., 1999; Mankiw and Romer, 1991; McCallum and Nelson, 2004), where models are usually characterized by a limited number of dynamic forward-looking equations, derived from a log-linear approximation of a general equilibrium model with optimizing, representative and homogenous agents, e.g., a representative consumer and a continuum of homogeneous firms. While recent developments regarded the introduction of learning within the usual new-Keynesian framework (see e.g. Casaccia et al., 2006; Evans and Honkapohja, 2003), this paper may offer a new contribution to the study of monetary policy from the perspective of the economics of

heterogenous and interacting agents. In particular, we study the effects of a nominal interest rate as the operational instrument of monetary policy, according to the current approach within monetary economics (see e.g. Walsh, 2003; Woodford, 2003) we investigate an interest rate rule which depends on the gap between the current output of the economy and the full-employment output (i.e., all households that apply for a job are hired). Being labor the only factor of production, when a full employment state is reached, the output can not be further increased, causing the price-setting productive sector to strongly increase prices if it faces a higher demand with respect to its productive capacity, thus generating an price inflation. This may give rise to instability and undermine the economy. In order to keep the inflation monitored and to guarantee stability, a monetary policy that keeps the output somewhat below the maximum potential output may be effective. It is worth noting that, in the optimizing sticky price model of the new-Keynesian literature (Clarida et al., 1999), a concept of output gap, defined as the deviation of output from its level under flexible prices, plays a central role both as a source of fluctuations in inflation (represented by the new-Keynesian Phillips curve), and as a policy target (e.g., the well-known Taylor's rule Taylor (1993)). It is worth noting that, irrespective of the different definition provided in our model, the output gap has a similar role here both as a determinant of inflation dynamics and as key policy variable. Furthermore, the maximum potential output is not fixed in the model but it is an endogenous variable, because the productive capacity of the firm is bounded by the households' labor supply, and this in turn depends on the current real wage. Indeed, the principal driver of the labor supply dynamics resides in the heterogeneity of the reservation wages, i.e., each household is characterized by a reservation real wage that indicates the wage that makes households indifferent between taking a job or remaining unemployed. These features gives rise to a very rich economic behavior which poses challenging issues to the monetary policy maker.

The paper is organized as follows. The model is outlined in Section 2. Computational experiments and results are discussed in Section 3. Section 4 provides some concluding remarks.

3.2 The model

The model is composed by a labor, a goods and a credit market. Households supply the labor force in the labor market and are organized in a trade union that sets the nominal wage. Each worker is characterized

by a reservation wage, i.e., a minimum real wage in order to apply for a job. A monopolistic firm hires workers to produce the scheduled quantity of output. The firm acts in the goods market as a price setter, and supplies the output according to a profit maximizing behavior. The aggregate demand is given by the sum of each household's demand, which is modeled according to a rule of thumb proposed by Deaton (1991a,b), based on the assumption that households, if liquidity constrained, save in order to smooth consumption over time. The individual consumption rule has been adapted here to our framework. The firm borrows money from the central bank in the credit market in order to pay wages, the bank sets an the interest rate according to the policy rule.

3.2.1 Households

Households take two key decisions in the model, determining the labor supply, according to their heterogeneous reservation wages, and how much to save or to consume in order to smooth consumption over time. Furthermore, a trade union sets the nominal wage w in order to increase the aggregate real labor income U , given by $(w/p)N$, henceforth workers' utility, where N is the number of workers (with $N \leq M$, M being the total number of households) and w/p is the real wage (being p the price level). The wage policy of the trade union is based on a backward looking behavior. If the correlation $\rho(dU, dw)$ between nominal wage variations dw and variations of workers' utility dU , computed in a backward time window T^U , is positive, i.e., nominal wage increments dw led in the past to an increase of workers' utility, the trade union raises the nominal wage. If the correlation is negative, the trade union keeps the nominal wage unchanged. In the former case, the wage bill is increased according to a fixed rate π^* set by the central bank, corresponding to a fixed planned rate of inflation. The trade union's decision rule can be summarized as:

$$w_t = \begin{cases} w_{t-1}(1 + \pi^*) & \text{if } \rho(dU, dw) \geq 0, \\ w_{t-1} & \text{if } \rho(dU, dw) < 0. \end{cases} \quad (3.1)$$

This wage indexation rule has been selected in accordance with the current practice in some European countries, e.g., Germany and Italy.

Reservation wages

Households' labor supply depends on the comparison between the current real wage and the reservation wage of each household. If the current real wage exceeds its reservation wage, then household i -th applies for

a job, if not, the household does not apply for a job unless its financial condition does not allow it to buy the essential goods for survival. Formally, the i -th household applies for a job according to the following rule,

$$\left\{ \begin{array}{l} (\frac{w_t^i}{p_t} \geq w_i^R) \cup (\frac{X_{t-1}^i}{p_t} \leq S^C) \longrightarrow \text{job application,} \\ (\frac{w_t^i}{p_t} < w_i^R) \cap (\frac{X_{t-1}^i}{p_t} > S^C) \longrightarrow \text{no job application,} \end{array} \right. \quad (3.2)$$

where w_i^R is the i -th household real reservation wage, X_{t-1}^i is its available cash at the end of the previous period and S^C , expressed in real terms, represents the indispensable quantity of goods to consume, that we call survival real cash, taken as homogeneous among agents and constant in the model. This is due to the fact that we consider S^C as a parameter that characterizes the whole population of households with a cultural attitude towards spending. S^C has therefore to be interpreted as a minimal arbitrary quantity of goods for a decent living rather than as a survival level *tout court*. Reservation wages w_i^R , defined as the wage that makes households indifferent between taking a job or remaining unemployed, are heterogeneous but constant. They represent a sort of social stratification for households that is kept constant along time. It is worth noting that in this model reservation wages are not, at least directly, a determinant of the actual wage, that is fixed by the trade union in order to increase workers' utility, but they have an essential part in determining the unemployment rate (see Hogan, 2004, for empirical evidence on these topics). Each household is endowed with a specific reservation wage according to a uniform distribution that varies from a minimum level $w_{\min}^R$ (generally set to zero) to a maximum level $w_{\max}^R$ that is used as a varying parameter for computational experiments.

Consumption rule

Household consumption choice is based on the theory of buffer-stock saving pioneered by Deaton (1991a,b), which states that households, if restricted in their ability to borrow to finance consumption, have a precautionary demand for saving in order to smooth consumption in case of bad draws of income, e.g., unemployment. The theory proposes accordingly a rule-of thumb as an approximation of the usual intertemporal maximization problem for the determination of the consumption path. The rule-of-thumb has been modified in order to take into account price inflation and is based on the comparison between the current income and past income stream realized in the last time window T^i . Let us define as X_{t-1}^i the quantity of cash at the i -th household disposal

before its consumption choice c_t^i at period t . The households's disposable income for consumption I_t^i is composed by the previous period wage, w_{t-1} , and the dividends from profits that the firm made in the previous period, i.e., $I_t^i = \delta_{t-1}^i w_{t-1} + m_{t-1}^i d_{t-1}$, where δ_{t-1}^i is equal to 0 or 1, depending on the employment status of the household at time $t-1$ and the integer m_{t-1}^i is the number of shares in the portfolio of household i at the end of previous period. Dividends d_{t-1} are given by $p_{t-1} \Pi_{t-1} / K$, where Π_{t-1} are the real profits realized by the firm in the previous period and K is the total number of shares of the monopolistic firm. The households' target is to maintain a stable rate of consumption, i.e., saving when income is high in order to accumulate cash for periods of low income. Deaton assumes that individuals consume cash as long as current nominal income I_t^i is less, in real terms, than the average past real income $\bar{I}_t^i$, while, if the income exceeds $\bar{I}_t^i$, households save a constant fraction $(1 - v)$ of the excess income. Thus, given the price p_t set by the firm in the current period, Deaton's decision rule can be formalized as:

$$c_t^i = \begin{cases} \min(\bar{I}_t^i, (I_t^i + X_{t-1}^i)/p_t) & \text{if } I_t^i/p_t \leq \bar{I}_t^i, \\ \bar{I}_t^i + v(I_t^i/p_t - \bar{I}_t^i) & \text{if } I_t^i/p_t > \bar{I}_t^i. \end{cases} \quad (3.3)$$

Aggregate goods demand Y_t^d is then given by $Y_t^d = \sum_i c_t^i$.

3.2.2 The monopolistic firm

The model includes a single monopolistic firm whose role is:

- to set the price and the quantity of the goods to be produced, according to a profit maximizing behavior,
- to hire workers, to produce and sell the goods,
- to distribute profits to households.

The firm produces an homogeneous perishable good according to a production function whose only input is labor:

$$Y_t = \zeta N_t^\alpha. \quad (3.4)$$

The parameters $\zeta > 0$ and $\alpha > 0$ are determined by the current technology and are kept constant in our computational experiments. The firm knows the nominal wage w_t that has been already set by the trade union, and acts as a price setter, facing the problem to decide the price p_t of the good and the quantity Y_t of goods to be produced. The firm also knows the labor supply N_t^s and has a perfect knowledge of the

demand elasticity. In order to set the price, the firm takes into consideration a set of hypothetical prices p_t^h , that lie in a neighborhood of the last market price p_{t-1} . The prices p_t^h are chosen inside a grid parameterized by $(1+j\epsilon)p_{t-1}$, with $j = -n, -n+1, \dots, n-1, n$, where ϵ represents the minimum relative variation of the price and $n\epsilon$ is the higher bound for variation. Consequently, the firm calculates the exact goods' demand relative to each price, i.e., $Y_t^d(p_t^h)$. Therefore, the firm computes, for each pair $(p_t^h, Y_t^d(p_t^h))$, the value of real profits, considering nominal costs given by:

$$\mathcal{C}_t = (1 + r_t^L)w_t N_t, \quad (3.5)$$

where r_t^L is the interest that has to be paid on the loan $w_t N_t$, and $N_t = (Y_t/\zeta)^{1/\alpha}$, with the constraint $N_t \leq N_t^s$. The price and quantity couple $(p, Y)_t$ is therefore chosen as the one that corresponds to the higher real profits, i.e.,

$$(p, Y)_t = \operatorname{argmax}_{(p, Y)_t} \Pi_t, \quad (3.6)$$

where

$$\Pi_t = Y_t - \mathcal{C}_t/p_t. \quad (3.7)$$

Finally, the firm distributes profits to households. Each household will receive dividends at the beginning of the next period, proportionally to the number of stocks it owns.

3.2.3 The central bank

The model incorporates a bank, which fulfills the functions of both a commercial bank and a central bank. The bank performs the following actions:

- to set an inflation target π^* ,
- to remunerate the household's cash account at a fixed rate r^D ,
- to provide credit to firms at a lending rate r_t^L ,
- to set r_t^L according to a monetary policy rule.

The rate on deposit r^D is set by the bank at the target level of inflation π^* , in order to let the money aggregate of the households grow at the same rate of inflation. A policy rule, that uses the nominal interest rate r^L as the operational instrument, has been designed. It is based on the control of the output gap, and sets the lending rate r_t^L as:

$$r_t^L = r^{L_{\min}} + \phi \exp \left(-\beta \frac{Y_t^p - Y_t}{Y_t^p} \right), \quad (3.8)$$

where β is a policy tuning parameter, ϕ represents the policy strength, varied to compare different monetary policies in the computational experiments, and Y_t^p is the potential output, given by the quantity of goods that would have been produced if all the available labor force had been employed, i.e., $Y_t^p = \zeta(N_t^s)^\alpha$. The effect of the rule on the economy is discussed in Section 3.3.

3.2.4 Rationing and accounting

As described in Section 3.2.2, the firm knows in advance the aggregate demand curve. Accordingly, it sets the price and chooses the goods supply in order to match the goods demand at that price. Therefore, the firm is never rationed in the goods market. The firms also knows in advance the labor supply. Consequently, while taking price and quantity decisions, it considers its possible rationing in the labor market. If there is not enough labor supply to produce the desired quantity, the firm hires all the households applying for a job and produces a quantity of goods lower than the demand level, implying that households will be rationed in the goods market. Furthermore, households are often rationed in the labor market. In both cases, households are rationed according to a random priority list, drawn from a uniform distribution.

After transactions in the goods and in the labor markets, households cash is reallocated for the next period, i.e., for the i -th agent:

$$X_t^i = X_{t-1}^i + \delta_{t-1}^i w_{t-1} + m_{t-1}^i d_{t-1} - p_t c_t^i + r^D X_{t-1}^i, \quad (3.9)$$

where r^D is the fixed rate on deposit of the bank, δ_{t-1}^i indicates the employment state of agent i at time $t - 1$ (i.e., $\delta_{t-1}^i = 1$ or $\delta_{t-1}^i = 0$ denote employment or unemployment state at $t - 1$ respectively), and $m_{t-1}^i d_{t-1}$ denotes the capital income due to m_{t-1}^i stocks paying dividend d_{t-1} . The term $p_t c_t^i$ takes into account the nominal expenses for consumption and $r^D X_{t-1}^i$ is the remuneration of the saving account.

3.3 Computational results and discussion

We present a study on the effects of using a nominal interest rate as the operational instrument of monetary policy. The interest rate r^L has an influence on the economy through the decision making of the firm, which borrows money to pay wages. Given the nominal wage set by the trade union, nominal labor costs incurred by the firm depend directly on the interest rate level, as shown by Eq. 3.5. As an example, a rise of




Fig. 3.1. The dynamics of price and output for three different values of the monetary policy strength parameter ϕ .

the interest rate at time t implies an increase of costs, and determines at the same time step an upward shift of the firm's supply curve in the (Y, p) plane. Due to the fact that the aggregate demand curve at time t is not yet affected by this interest rise, the goods market clears at a higher price and at a lower quantity. However, a rise of the interest rate, if negative in the short run, may have positive effects in the long run by keeping the economy below its full capacity.

The computational experiment presented here have been realized considering the following parameters values: $M = 1000$, $T^U = 20$, $\pi^* = 0.5\%$, $\zeta = 1$, $\alpha = 0.9$, $\delta = 0.1$, $n = 50$ (implying a maximum price variation of $\pm 5\%$), $r^D = 0.005$, $r_{\min}^L = 0.1$, and $w_{\min}^R = 0$. No stock trading is considered, and each household has been endowed with the same amount of stock holdings $m^i = K/M$ which is keep constant over time. Different values for the maximum reservation wage $w_{\max}^R$ and for the monetary policy strength parameter ϕ have been considered. Figure 1 shows the dynamics of price and output for three different values of the monetary policy strength parameter ϕ . Figure 1 points out the presence of price spikes and simultaneous falls of production; they occur when

the economy reaches its full capacity. In that case, the firm can not hire more workers to increase production and profits. Thus, obeying to its profits maximizing behavior, it has to rise the price. Consequently, consumers' demand is depressed and this reduces the market clearing output of the economy. Furthermore, the price rise reduces the real wage and therefore the labor supply, which is a binding production factor when the economy is running at its full capacity. The objective of the monetary policy rule, outlined in Eq. 3.8, is to prevent, by rising the interest rate, the firm from scheduling a production level that could not be sustained by the labor supply, i.e., to prevent the economy from reaching its full capacity. The policy strength parameter ϕ weights the importance of the output gap in the interest rate setting. The monetary policy experiments showed in Figure 1 point out that a monetary policy which takes into account the output gap, i.e., $\phi > 0$, gives rise to an higher inflation rate and lower output growth in the short run, but it is able to contain output negative fluctuations and to significantly reduces the volatility of prices and output in the long run. Furthermore, it is worth remarking that caution has to be payed in tightening the monetary policy, and a tradeoff is necessary between a lower inflation rate ($\phi = 0.5$) and an higher output growth rate ($\phi = 0.3$). Figure 3.2 shows how reservation wages affect the labor market dynamics and therefore the output level of the economy. Simulation shown in Figure 3.2 has been performed with the usual monetary policy rule with ϕ set at 0.1. The reservation wages are heterogeneous among households and uniformly distributed between zero and a maximum value $w_{\max}^R$. Raising $w_{\max}^R$, the reservation wages range becomes larger and there are more households that require a higher salary to work. Figure 3.2 allows one to observe that a higher lever of the maximum reservation wage $w_{\max}^R$ has a clear depressive effect on the labor supply, e.g., for $w_{\max}^R = 0.35$, there is a few number of households that are disposable to work (around 250). Thereby, this contraction on the labor supply has a depressive effect on the long-term level of production in the economy. An effective monetary policy should then take into account this feature in order to address the usual trade-off between long-run output level and inflation control. Figure 3 shows the mean value of the relative output gap, i.e., $(Y_p - Y)/Y_p$, relative to the last 500 time steps. Each bin represents an output gap value calculated for a couple $(w_{\max}^R, \phi)$ where $w_{\max}^R$ varies from 0.1 to 0.5 with a step interval of 0.025 and ϕ varies from 0.1 to 0.5 with a step interval of 0.1. As expected, the value of relative output gap diminishes with the increasing monetary policy strength. Furthermore, it is worth noting that the relative output gap




Fig. 3.2. The dynamics of labor supply and output for three different values of the maximum reservation wage $w_{\max}^R$.

shows a clear dependence on the maximum reservation wage. Indeed, according to our model, the maximum value of the reservation wage, setting the labor supply, determines the potential output of the economy, i.e., higher $w_{\max}^R$ correspond to lower Y_p . Besides, Figure 3 points out that also the relative value of the output gap decreases for rising values of $w_{\max}^R$. This results suggests that a milder monetary policy, i.e., lower values of ϕ , should be considered as more appropriate in the case of lower $w_{\max}^R$.

3.4 Concluding remarks

The model presented should contribute to the agent-based approach for monetary policy design along two main different perspectives. First, it provides a sensible micro-foundation to a monetary policy rule based on output gap control, with a complementary approach to the new-Keynesian Phillips curve literature. Under this respect, it shows that the optimizing behavior of a price setting monopolistic firm is sufficient




Fig. 3.3. Average relative output gap for different values of ϕ and $w_{\max}^R$.

to produce inflation when the economy reaches full employment. Second, the model shows how an effective monetary policy design should take into account the distributional property of reservation wages among agents, i.e., an individual heterogeneous feature which determine aggregate labor supply, and, thus, the potential output of the economy at full employment.

Acknowledgement. This work has been partially supported by the University of Genoa, and by the Italian Ministry of Education, University and Research (MIUR) under grants FIRB 2001 and COFIN 2004, and by the European Union under IST-FET STREP Project EURACE.

References

S. Casaccia, S. Cincotti, M. Raberto, and A. Tegli. Monetary policy subject to measurement errors of private sector adaptive expectations. In G. Setti and T. Ushio, editors, *Proceedings of the 2006 International Symposium on Nonlinear Theory and its Applications*,

- pages 335–338. Research Society of Nonlinear Theory and its Applications, 2006.
- R. Clarida, J. Gali, and M. Gertler. The science of monetary policy: A new keynesian perspective. *Journal of Economic Literature*, 37(4): 1661–1707, December 1999.
- A. Deaton. Savings and liquidity constraints. *Econometrica*, 59(5): 1221–1248, September 1991a.
- A. Deaton. Household saving in ldc's: credit markets, insurance and welfare. *The Scand. J. of Economics*, 94(2):253–273, 1991b.
- G.W. Evans and S. Honkapohja. Adaptive learning and monetary policy design. *Journal of money credit and banking*, 35(6):1045–1072, December 2003.
- V. Hohan. Wage aspirations and unemployment persistence. *Journal of Monetary Economics*, 51(8):1623–1643, November 2004.
- N.G. Mankiw and D. Romer, editors. *New Keynesian Economics*. MIT Press, 1991.
- B.T. McCallum and E. Nelson. Timeless perspective vs. discretionary monetary policy in forward looking models. Technical report, The Federal Reserve Bank of St. Louis, 2004.
- M. Raberto, A. Tegli, and S. Cincotti. A general equilibrium model of a production economy with asset markets. *Physica A*, 370(1):75–80, 2006.
- J.B. Taylor. Discretion versus policy rule in practice. In *Carnegie-Rochester Conference Series on Public Policy*, volume 39, pages 195–214, 1993.
- L. Tesfatsion and K. Judd. *Agent-Based Computational Economics*, volume 2 of *Handbook of Computational Economics*. North Holland, 2006.
- C.E. Walsh. *Monetary theory and policy*. The MIT Press, 2003.
- M. Woodford. *Interest & Prices*. Princeton University Press, 2003.

Market Mechanisms and Agents Behavior

Testing Double Auction as a Component Within a Generic Market Model Architecture

Julien Derveeuw¹, Bruno Beaufils¹, Philippe Mathieu¹, and Olivier Brandouy²

¹ Laboratoire d'Informatique Fondamentale de Lille, USTL, France
`{derveeuw,beaufils,mathieu}@lifl.fr`

² Lille Economie et Management, USTL, France
`olivier.brandouy@univ-lille1.fr`

4.1 Introduction

Artificial stock markets are models designed to capture essential properties of real stock markets in order to reproduce, analyze or understand market dynamics with computational experiments. Despite research advances in modern finance many questions remain unsolved: market dynamics exhibit, for instance, particular statistical properties, called *stylized facts*, which origins are not clear. As real markets are complex systems, it is really hard to study them directly because too many parameters stay out of control. Hence, multi-agents simulations of these markets seem to be a key for a better understanding of their properties.

Building such models implies to simplify reality as most as it can be in order to keep markets most representative and characteristic features. In the literature (see for example LeBaron et al. (1999), Cincotti et al. (2006) or Ghoulmie et al. (2005)) real markets structure complexity is often circumvented by the use of an equation weighting the balance between bids and offers as a price formation model. This simplification is in complete contradiction with the reality of stock markets where prices *emerge* from agents interactions through an order book which do not act as a central weighting entity but as a peer-to-peer meeting point used by agents to exchange stocks. However, such studies manage to reproduce realistic price series, which seems odd regarding market models used. We can then wonder if some of these models are more suited than others to capture market dynamics.

To answer this question, it seems that a comparison between these models needs to be realized in order to put them to the proof and investigate their robustness. Hence, we propose in this article a generic market model architecture based on four independent entities, each of which can be modeled in different ways. We show that existing models found in literature fit well in this architecture. We then propose an artificial stock market model which takes into account real markets characteristics: trading activity takes place *continuously* through an *asynchronous* mechanism. Agents interact through the market by posting *orders* in an *order book*, as it happens on real market places. We show that without making any strong assumption on agents behaviors, this model exhibits many statistical properties of real stock markets.

4.2 Quick review of different ASMs architectures

Since the first artificial stock market was developed in the early nineties at the Santa-Fe Institute Palmer et al. (1994), many market models have been developed. Though almost all of them aim to reproduce the same market properties (the so-called stylized facts) with the same multi-agents simulation methods, they all exhibit different properties: some are synchronous, while others are asynchronous. Some of them require agents to emit realistic orders (direction/price/quantity) while others only require a direction (buy/sell) to compute the new stock price. Without pretending to be completely exhaustive, we investigate in this section some of these models in order to identify the most represented microstructures and trading rules in artificial stock markets.

The Santa-Fe artificial stock market

Historically, the first model to be developed was the *Santa-Fe Artificial Stock Market*. This model is mainly characterized by the use of a macroscopic equation based on demand and supply law to compute the new traded stock price. Hence, agents take their decisions synchronously and emit their desires as a direction $a_{i,t}$ (buy $a_{i,t} = 1$ or sell $a_{i,t} = -1$) to the market, which calculates the imbalance between demand and supply ($I_t = \sum_i a_{i,t}$), to finally compute the price according to equation 4.1.

$$p_{t+1} = p_t(1 + \beta \times I_t) \quad (4.1)$$

Though this model may seem attractive due to its relative simplicity, its lack of realism regarding real market microstructure is obvious:

agents take their decisions synchronously without being able to reason about others beliefs; moreover, agents are not even aware of the quantity of stocks they will trade due to the clearing process used to realize exchanges between agents once the price is calculated.

The \$-game

To solve the question of market clearing, a possible solution is to add a market maker to the model, so agents are always satisfied with the quantity they want to trade. This feature was incorporated in the \$-game ASM Andersen and Sornette (2003). As the market maker provides liquidity to the market (e.g. he buys excess stocks and provides supplementary stocks when needed), his position needs to be covered to avoid bankruptcy. Hence, Andersen et al. use in their model a slightly modified version of the previous price calculation equation. Instead of only considering the current imbalance between demand and supply, they also take into account the global imbalance since the beginning of the simulation, which is the market maker current position. Using the same naming as above, the price update equation is then given by 4.2.

$$(\ln(p_t) - \ln(p_{t-1})) = \frac{I_t + \sum_{i=0}^{t-1} I_i}{\lambda} \quad (4.2)$$

Though this model correctly addresses the problem of stock liquidity and market clearing, it can't be considered as a realistic one: agents still interact synchronously with the market and only emit a desired quantity to trade, without having the ability to associate it with a desired price for the transaction.

The Genoa artificial stock market

To bring more realism to synchronous models, researchers from Genoa proposed a model called the *Genoa artificial stock market* in which agents are allowed to emit classical limit orders to the market (see Raberto et al. (2001), Cincotti et al. (2003) or Raberto et al. (2003)). In this model, agents still take their decisions synchronously, but as they associate a limit price to the desire they pass to the market, a different clearing mechanism needs to be used to ensure that agents do not buy or sell stocks for a different price than the limit they asked for. This is achieved by computing a clearing price, which is defined as the crossing of the demand quantity curve function of price and of the supply quantity curve function of price (see equations 4.4 and 4.3 for a definition of these two series).

$$f_{t+1}(p) = \sum_{u|p_u \geq p} q_u^b \quad (4.3)$$

$$g_{t+1}(p) = \sum_{v|p_v \geq p} q_v^s \quad (4.4)$$

Though this model is more realistic than the previous ones, it still lacks an essential feature of real markets microstructure: the asynchronism of transactions.

Toy model of an asynchronous double auction

In order to get a more realistic time handling process in artificial markets, some researchers proposed models in which transactions take place asynchronously. This is the case of the toy model proposed in Bak et al. (1996). In this model, there are only $\frac{N}{2}$ stocks on the market, where N is the number of agents. Agents do not have the right to own more than one share at a time. They can therefore be sellers if they own a share, or buyers if they own nothing.

At each time step, an agent is given speak randomly and has the possibility to emit a desire according to the pre-cited rules. This desire is a composed of a price and a direction. If this agent finds an other one who is willing to make the opposite transaction with a compatible price, they immediately exchange one share. If no counterparts are available, the agent's order is saved in a list until a counterpart is found.

Even if it is a toy model, this model is one of the first to take into account the asynchronism of exchanges on real market places. Agents act in a random order and a simplified order book is used to save agents desires. A criticism which can be made is that the market rules used (an agent can at most own one share) tend to make the market illiquid and prevent from testing realistic investment strategies.

We have seen in this section that many different market models are used to reproduce high frequency dynamics from real stock markets. Despite of their heterogeneity, they are used to reproduce the same three main stylized facts: the shape of the return distribution (which is fat-tailed and leptokurtic), the autocorrelation of absolute returns and clustered volatility. We can notice strong differences in the way agents express their desires, in the set of information they are able to get from the market, and in the way they are given speak by the market. This a major problem regarding our main goal, which is to be able to compare heterogeneous market models in similar experimental environment.

4.3 A generic market model architecture

In the previous section, we have presented some of the most representative market models microstructures and trading rules found in the literature. Their diversity is so great that it seems difficult to correctly identify which of these models parts are responsible for the statistical properties of computed price dynamics: are they due to the microstructure of the market ? to the way time is handled ? to the agent investment strategies ? In order to address these open questions, we expose in this part a generic model of market architecture which allows to unify these different models. We also show that this formalization allowed us to develop a concrete implementation of this generic architecture, which will make us able to compare artificial stock markets.

4.3.1 The abstract generic model




Fig. 4.1. General market model architecture.

If we look at how markets operate, we can decompose them in three parts: the *market*, which allows agents to exchange stocks, *agents*, who trade through this market, and the *external world*, which can for example influence agents with information. This situation is summed up in figure 4.1: agents communicate their desires to the market, being influenced by their peers or exogenous information. They can also be influenced by public information available from the market. If we make a parallel between this abstract model and multi-agents models of markets, we can see from the previous section that each of these three components can be modeled in different ways: the market can be an averaging equation or a complex microstructure; agents can be either cognitive, reactive or replaced by equations.

4.3.2 The concrete generic model

In order to experiment the influence of each of these modules on price dynamics, we need to be able to compose heterogeneous modules coming from the literature. For example, to investigate the influence of

market microstructure on prices, it seems interesting to study some of their different implementations for a given set of agents behaviors. Unfortunately, as we have seen in section 4.2, most of market models require agents to emit their desires in many different ways: there are sometimes expressed as a direction, sometimes as a quantity or even as limit orders. Hence, it seems obvious that to make our generic architecture practical, we need to propose some concrete details on its implementation.

Information

In our formalism (see figure 4.1), we showed that agents were able to use some information coming from the market in order to take a decision. As we saw in the first section, information required by agents or published by market models are heterogeneous: some market models only publish the last transaction price, while others make all of the agents current positions public. Hence, to be able to compose any agents model with any market model, it is necessary to define the maximum set of information needed by agents models and to define how all of these information can be approximated when they are not present in a given market model.

According to our literature review, agents use *at most* the following information from the market:

- the last transaction price, which is an information available on every market model
- other agents desires (which represent the order book in asynchronous models).
- current demand and supply disequilibrium, which is available in most synchronous models. In asynchronous model, it is easy to deduce this information from the current order book state by summing quantities available in both sides of the order book.

To be able to compose any market model with any agent model, we have to define a set of translators able to fill missing information from some market models if it is required by the agents. An example of such a translator (or wrapper) is described in table 4.1. Though we provide in our framework a full set of information translators which allow to translate any type of emitted information in any type of required information, the effect of these translators on experimental results still has to be investigated.

emitted by market	required by agent	translator description
(price, agents positions)	(price, disequilibrium)	translator sums up quantity associated to agents positions in order to compute the global demand/supply imbalance

Table 4.1. An example of information translator.

Agents desires

In figure 4.1, we identified that agents emit trading desires to the market, which are then interpreted according to market model trading rules. These desires, in artificial markets as well as on real ones, are defined by a composition of the three following characteristics: a direction, a price and a quantity. Obviously, the direction is the minimal requirement in order to get a valid desire (emitting a desire to a market without saying if one wants to buy or sell makes no sense). The two others desires properties (price and quantity) are optional according to the agent or market model. As we would like to compose any agents and markets models which emit or require different types of desires, we need to define a translation system to make this composition possible.

Assuming that a direction is the minimum required to express an economic desire and that the maximum is a direction, a price and a quantity (which was deduced from our intensive literature investigation), it is possible to propose a first set of translators (which are called wrappers in computer science) that are required to allow communication between any agent model and any market model. The effect of these wrappers on agents and market behaviors still has to be investigated. An example of such a translator is described in table 4.2.

emitted by agent	required by market	translator description
(d, , q)	(d, p, q)	interpret order as a market order, and fill the missing price with the best offer in opposite direction

Table 4.2. An example of desire wrapper.

Time handling

In addition to the differences between information required or emitted by the different modules of a market, time handling is managed in very

different ways regarding the market model used: some are synchronous while others are asynchronous. Moreover, each of this time handling philosophy can be implemented in several different ways. These differences are a major problem to solve while trying to compose heterogeneous modules: if an agent strategy is built to operate in an asynchronous context, is it possible to make its strategy make sense in a synchronous one ?

To address this problem, we have split time handling from the market module and separated it in what we call a *simulation engine*. This additional module is responsible for giving the ability to talk to the agents and for making the market treat agent desires when it is time to do so. For example, a synchronous simulation engine will give to all of the agents the ability to talk, and will then ask the market to compute the new stock price, whereas an asynchronous one will perhaps pick randomly an agent and then immediately ask the market to take his desire in account.

Global framework layout

Due to lack of space, we can not explain further all of the implementation details that are needed to allow free market modules composition. Figure 4.2 sums up the general layout of our simulation framework, which we detail step by step:




Fig. 4.2. Framework functioning.

- *step 1*: The simulation engine gives speak to the agent(s) who are allowed to speak at current time according to the time policy in use.

- *step 2*: Before taking a decision, agents are able to ask the market some information about its current state (best offers, current stock price, demand/supply imbalance, etc). As each market model can exhibit different public information, they need to be treated by a wrapper which traduce them so they can be used by any agent model. Agents can also ask external world about its current state if their decision making process requires such an exogenous information.
- *step 3*: Once agents have sufficient information to take their decisions, they can emit a desire to the market. As we have seen before, this desire can be expressed in many different ways, so it needs to be traduced by a wrapper to be understood by any market model. These desires are then stored in an accumulator, which is useful to keep track of agents desires, in particular if the simulation engine is synchronous.
- *step 4*: Each time the market receives an agent desire, it immediately informs the emitter about its validity. This is required as some market models require agents to meet specific conditions to be able to emit desires.
- *step 5*: Once the simulation engine has given speak to the agents allowed to do so, it notifies the market that it is time to take the agents desires into account. If the market is order book based, this means “insert new desires in the book”, whereas in equation-based models, this means “enter in a clearing phase and compute a new price”.
- *step 6*: The simulation engine finally gives the possibility to the world model to update itself.

Limitations

Even if our generic architecture is implemented and practical, it still has some limitations inherent to the major differences between models we try to compose one with another.

For example, some information translators need to be able to translate an information expressed as a single price in an information expressed as other agents positions. Even if other agents positions may be assimilated to the current stock price, impact of such translations on agent trading strategy have still to be investigated. The same observation can be made about the composition of agents designed to work in an asynchronous context with market models designed to work in a synchronous one.

Hence, our generic architecture still has to be improved and validated with intensive experiments, in order to make sure that translators do not bias simulations results. Even at this early stage, this generic architecture can however be merely considered as a formalism able to describe any artificial stock market model through their components.

4.3.3 An example of application: the market component as double auction

We have seen in section 4.2 that most of existing market models lack realism: some do not respect real markets asynchronism while others oversimplify the way agents emit desires to the market. In consequence, we choose to illustrate the use of our generic market simulation framework by implementing a simple asynchronous double auction model following our formalism. This model can be linked up to the one used in Raberto et al. (2005). We will detail in this section how each module is defined according to the formalism we presented before.

The market component

The market component is a classical order book similar to the one used on market places such as EURONEXT. This order book requires agents desires to be expressed as a *direction*, a *price* and a *quantity*, which defines an order. These orders are all *limit prices orders*, which means that the price associated to the order is the maximum (respectively minimum) price the agent is willing to buy (sell) stocks. When an order is received by the market, it is stored in the order book according to price and time priorities if it has no counter part. When a counterpart is found, a transaction occurs immediately and the price of this transaction is published.

The simulation engine component

In order book based markets, time handling does not follow the same logic as in equation-based ones: central quotation system does not aggregate agents decisions at particular time steps and market participants are free to talk when they want. Hence, we need to implement the simulation engine component as a process which asks agents to speak asynchronously and which asks the market to update its current state each time an agent has spoken.

Our choice is to give randomly an agent the opportunity to talk regardless to the fact he has already spoken or not. The major inconvenient of this method is that some agents can be out of the market

(have never the opportunity to speak) because of the random generator used in the scheduler. However, on real markets, some agents are very active (speak a lot) whereas others rarely interact with the market. For these reasons, this is the scheduling principle we choose.

The agent component

Following the works of Gode and Sunder (1993), our agents are designed as purely reactive ones (as simple as possible), which implies that we do not make any strong hypothesis about the agents reasoning capabilities, nor on the information set they use to take their decisions, as it is done in most of other studies. The choice of using simple agents behaviors in this article is hence deliberate: our goal, here, is not to design realistic agents but to validate our microstructure model separately from the two other components of the market architecture.

These agents can be assimilated to *zero intelligence traders* who post orders with a random direction, a random price for a random quantity of stocks. When an agent emits a new order, he stops emitting new ones until his order is fulfilled or until the order reached his *timeout*. This *timeout* is randomly assigned to each agent at the beginning of the simulation and stays constant over time. This mainly guarantees that an order with a price too far from the current limits of the order book won't stay in it for an endless time.

4.4 Experiments

We present in this section some experiments we have designed to test our generic framework. Only a part of the statistical tests we made are reproduced here due to lack of space. Full experimental tools and results used to produced data presented in this paper may be downloaded at <http://cisco.univ-lille1.fr/papers/ae2007>. This experiments are realized using the market model, agents and scheduler exposed in previous section. All of our experiments are run on 20 000 time steps with 100 agents.

First, we interested ourselves to the returns distribution as its shape is one of the major characteristic of real price dynamics. This distribution, on real markets, is leptokurtic and exhibits fat tails. Table 4.3 shows some statistical results: the excess kurtosis measured oscillates around 4.5 which is similar to what can be observed with real markets data (see right column for a comparison). To further illustrate this property, figure 4.3 shows one of our experimental returns distribution compared to a theoretical normal distribution.

Description	Value (experimental)	Value (real data)
Excess kurtosis	4.52	4.158
Aug. Dickey-Fuller	-20.47	-18.47
ARCH	100%	100%

Table 4.3. Statistical results obtained with our interaction-based model, compared to the one obtained on real data (BMW daily stock returns coming from DAX30).




Fig. 4.3. Experimental returns distribution compared to a theoretical normal distribution with same mean and variance.

Another major characteristic of returns is that they do not exhibit significant autocorrelation but that a short-range autocorrelation decaying over time exists when looking at their absolute value. Figure 4.4 presents the ACF plot for both returns and absolute returns. Comparing them to the ones obtained with real market data, we can see that returns properties similar to reality can be obtained with our interaction-based model. These properties are further verified by the use of the Augmented Dickey-Fuller test which tests for the null hypothesis “*The serie has a unit root*”. Table 4.3 shows its result on our

time series: the presence of a unit-root is rejected at a high confidence level as with real data (right column).




Fig. 4.4. ACF of returns and squared returns obtained in our experiments.

We have seen in this section that time series obtained with our model exhibit the same statistical properties as real data sets. This results improve the preliminary ones obtained by Raberto et al. (2005). This shows that our asynchronous and continuous auction model is able to reproduce most of markets characteristics without making any assumption on agents behaviors or on an external world model.

4.5 Conclusion

In this article, we introduced a generic architecture of artificial market models. This architecture is composed of four independent parts: a model for the external world, another for agents behaviors, one for the market structure and a last for time handling. We have shown that most of existing market models can fit in this architecture, so it can therefore be considered as a description formalism of artificial stock markets. Moreover, our generic architecture allows to compose existing market and agent models, which is a major benefit if one plans to compare market models between them: it is now possible to do such comparisons in identical environments (e.g. with the same agents) and to draw strong conclusions from these experiments, which was not the case before. However, some of the effects of our generic model still needs

to be investigated in order to make sure that translators do not bias simulation results.

We have also presented and tested an artificial stock market component based on an order book, which implies that quotation is *asynchronous* and *continuous* as on real markets. This is opposed to classical approaches, which aggregates agents decisions synchronously with an equation as a substitute for market interaction mechanism. First results show that it is possible to reproduce most of the *stylized facts* observable on real markets with a pure multi-agents model based on local interactions. This may confirm recent statements implying that most market features are due to the exchange process more than to agents behaviors.

We argue that such continuous and asynchronous models should be used in stock markets simulations. The order book model is so close to reality that no validation problems subsist about the mechanism used to make the agents exchange stocks. Moreover, developing agents behaviors is simplified: real traders investment strategies could be implemented “as is”, without having to modify their output to match the model requirements.

Concerning technical issues, we can notice that the order book does not require specific parameters: this ensures that no hazardous tweaking is necessary to make the market model work in a proper way. Moreover, our model is carefully designed with respect to multi-agents modeling paradigms: by adapting blackboard mechanism and well-known techniques of scheduling to the field of market simulation, we reduce the probability to get unwanted side effects due to technical issues in our simulations.

Now that we both have a realistic market model and a generic market architecture, we are going to be able to compare our model with other ones from the literature. By doing such intensive experiments, we hope to bring some more elements to the theories which impute most of the stylized facts to the market structure. We will also be able to test new investment strategies coming from classical economic literature such as the self referential agents proposed in Orlean (1999).

Acknowledgement. This work is supported by European funds of FEDER and the CPER TAC of Nord-Pas-de-Calais region.

References

- J. Andersen and D. Sornette. The \$-game. *European Physics Journal*, 31:141–145, 2003.
- P. Bak, M. Paczuski, and M. Shubik. Price variations in a stock market with many agents. *Physica A*, 246:430–453, 1996.
- S. Cincotti, M. Raberto, S.M. Focardi, and M. Marchesi. Who wins ? study of long-run trader survival in an artificial stock market. *Physica A*, 324:227–233, 2003.
- S. Cincotti, L. Ponta, and S. Pastore. Information-based multi-assets artificial stock market with heterogeneous agents. In *WEHIA06*, 2006.
- F. Ghoulmie, R. Cont, and J.P. Nadal. Heterogeneity and feedback in an agent-based market model. *Journal of Physics: Condensed Matter*, 17:1259–1268, 2005.
- D.K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101:119–137, 1993.
- B. LeBaron, W.B. Arthur, and R. Palmer. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, 23:1487–1516, 1999.
- A. Orlan. *Le pouvoir de la finance*. Odile Jacobs, 1999.
- R.G. Palmer, W.B. Arthur, J.H. Holland, B. LeBaron, and P. Tayler. Artificial economic life: A simple model of a stockmarket. *Physica D*, 75:264–274, 1994.
- M. Raberto, S. Cincotti, S.M. Focardi, and M. Marchesi. Agent-based simulation of a financial market. *Physica A*, 299:320–328, 2001.
- M. Raberto, S. Cincotti, S.M. Focardi, and M. Marchesi. Traders’ long-run wealth in an artificial financial market. *Computational Economics*, 22:255–272, 2003.
- M. Raberto, S. Cincotti, C. Dose, S.M. Focardi, and M. Marchesi. Price formation in an artificial market: limit order book versus matching of supply and demand. *Nonlinear Dynamics and Heterogenous Interacting Agents*, 2005.

A Conceptual Framework for the Evaluation of Agent-Based Trading and Technical Analysis

Olivier Brandouy¹ and Philippe Mathieu²

¹ Lille Economie et Management, USTL, France
olivier.brandouy@univ-lille1.fr

² Laboratoire d'Informatique Fondamentale de Lille, USTL, France
philippe.mathieu@lifl.fr

5.1 Introduction

The major part of research dedicated to technical analysis and active trading (*i.e.*, the management of financial portfolios using chartism or moving average indicators for instance) generally focuses on single “signals” giving the opportunity to buy or sell a financial commodity *frequently a well diversified portfolio* (see the extensive survey of Park and Irwin, 2004). In this context, it has been extensively argued that technical analysis is useless in order to outperform the market (Jensen and Benington, 1969). The reason for that is, assuming informational efficiency (Fama, 1970), all relevant piece of information is instantaneously aggregated in prices. Therefore, there is nothing to extract from previous quotations relevant for one willing to trade on this basis. Since information is, by definition, unpredictable, next price fluctuations will be driven by innovation and the price motion will fluctuate randomly as a result. Nevertheless, empirical investigations tackling this question of “technical trading” exhibit heterogeneous results. On the one hand, a large part of these researches shows that, once risk taken into account, no-one can seriously expect any rate of return over what can be earned with a simple *Buy and Hold* strategy (henceforth B&H). On the other hand, some intriguing results seem to attest that technical analysis is useful to a certain extent (Brock et al., 1992; Dempster and Jone, 2005; Detry and Gregoire, 2001). More generally speaking, this idea is trusted and shared by many practitioners.

We argue here that this confusion depicted by this heterogeneous set of results comes from ill-defined concepts, confusing measures and fuzzy evaluation procedures. We propose in this paper some elements to cor-

rect these imprecision and to elaborate a conceptual framework for technical analysis evaluation.

We consider these elements using an Agent-Based approach because we ultimately would like to investigate large sets of technical trading strategies, to encompass automatic trading issues and to generalize as much as possible our investigations. Thus, in this research, an *agent* is systematically an *artificial agent*, that is, a virtual entity endowed with *Artificial Intelligence*, mimicking a *real investor*, and able to deal with information, learning, and adaptation procedures.

Therefore, our propositions are a contribution to organize as rigorously as possible the large set of problems linked to the evaluation of automatic trading, technical analysis and related topics including those where *Artificial Intelligence* is used to investigate large sets of investment strategies.

This paper is organized as follows. In section 1, we discuss the basis upon which technical analysis is usually analyzed. We show why it must be distinguished between *signals*, *strategies* and *behaviors* although this distinction is seldom done in other researches. Section 2 focuses on the problematic link between technical indicators received by the traders and their ability to benefit from them when they try to implement them in “winning” strategies. Section 3 deals with the value added by increasing cognitive capabilities of the agents in plugging sets of technical signals rather than a single signal in their rationality. Section 4 enlarges the discussion to several strategies and addresses several questions around the design of tests for weak-form Market efficiency and automatic trading. It also serves as a conclusion.

5.2 Why confusing elements have lead to a controversy

If one considers the basic elements in most researches dealing with technical analysis or weak-form market efficiency, it is often the case that one specific “strategy” (or a limited set of strategies) is systematically replicated over various time windows, using real stock market data. Performance is computed comparing this active-investment strategy to a specific benchmark, like a simple B&H behavior. Some refinements concerning the statistical properties of the performance distribution is also usually proposed, such as Monte-Carlo simulations or Bootstrap Reality Checks (see for instance White, 2000).

However, no one can seriously sustain that these tests directly assess what a real technical trader would do. This practitioner would certainly mix a large number of “*receipts*” to strategize his behavior. His perfor-

mance is supposed to be grounded on various “signals”, “special skills” allowing him to have a correct diagnosis, and a professional “know-how” : this mix makes any evaluation complex because the origin of performance (or lack of performance) is not easily observable. To make this point clearer, let’s consider briefly figure 5.1. A large part of the evaluation complexity arises from the interaction between:

- elements constituting investors’ intelligence (and consequently, virtual agents’ artificial intelligence) : their cognition is a structured mix of information – extracted from market observation – and knowledge coming from the organization of these information *plus* the result of their past behavior,
- and external constraints : what kind of commodity are they allowed to trade? Are they subject to budget or credit constraints? Can they go *short* or not?




Fig. 5.1. Elements of complexity in performance evaluation

Every evaluation problem has to take into account these elements to be satisfying. The most difficult of them, and clearly the less treated in the literature, tackles the ability of technical traders to evolve and

to mix various elements to achieve good performance in the market. We will present this point at the end of the following example : to introduce the discussion, we propose first a basic situation, a very direct evaluation problem where trader's intelligence is limited since she applies systematically a strategy based on a "Mixed Moving Average 90-10 signal" (MMA_{90-10}). A "moving average" of range K (K being equal, in our example, to 10 and 90 since we mix these indicators) and a "mixed moving average (n, p)" are respectively proposed in the following expressions 5.1 and 5.2 :

$$MA_{K,t} = \frac{1}{K} \sum_{t-K+1}^t p_t \quad (5.1)$$

$$MMA_{n,p} = \{MA_{n,t}, MA_{p,t}\} \quad (5.2)$$

Chartists consider the situation in which the short-term moving average crosses the long term one from the bottom to the top as a "buy" signal (*resp.* from the top to the bottom as a "sell" signal). We use the daily closing value of the Dow-Jones from 26/05/1896 to 22/11/2005 (27424 quotation days) to generate a series of 478 signals (figure 5.2 shows a subset of this signals from 21/05/1996 till the end). We consider that it is always possible to trade a tracker based on this index. The allocation rule for the trader is simply maximum investment (that is, to buy as much trackers as possible or to sell them massively).

On the basis of the signals, the agent trades the DJ-tracker 477 times (the first signal being a "sell" signal). Do these chartist signals actually signal something useful for trading or not (**question 1**)? Would a portfolio, solely composed of trackers based on the Dow-Jones Industrial, have benefited from such a trading rule if one considers various performance indexes (**question 2**)? Especially, do these signals allow smart traders to elaborate strategies that outperform the market (**question 3**)? Is it possible to improve dramatically this Limited Intelligence Trader's performance in endowing her with higher cognitive skills (**question 4**)?

question 1:

The idea behind this question is the actual power of chartist signals to predict correctly, regularly and with a sufficient reliability, the next moves of one specific market. We can quantify this power with a very simple indicator called "Hit Rate":

$$HR_{MMA_{90-10}} = \frac{\text{correct signals}}{\text{total number of signals}} \quad (5.3)$$




Fig. 5.2. Moving Averages 10 and 90 over 10 years of DJI

Others definitions of such indicator can be found in the literature (Hellström and Holmström, 1998). The average score of our chartist signal, in terms of Hit Rate is here of 52%. This score may vary significantly over sub samples of time, and to some extent, it is hard to say that this 52% score is better than what a pure random rule would do. Nevertheless, we can still hypothesize that a subset of rules (whatever these rules are) in the infinite space of possible rules actually performs well.

question 2:

Assuming the MMA_{90-10} signal has been selected by a trader, would she be able to obtain a good performance implementing it in a basic strategy ³? Graphs 5.2 shows the evolution of a trader's portfolio composed of one tracker (at date 26/05/1896) and managing her portfolio with a basic strategy using MMA_{90-10} signals against a passive trader receiving at the same date the same tracker, and playing a B&H strategy. Rules for managing the portfolio are as follows: when a trader

³ in other terms, following what the signals suggest: to buy when the market is supposed to rise, to sell when it is supposed to decrease

decides to sell her portfolio, all the trackers she holds are sold. When she decides to buy, she invests all her cash in trackers (considering she will have to pay in both cases transaction costs at $x\%$). One can easily observe that when transactions costs are zero (graph 5.3), the basic MMA_{90-10} seems to perform well whereas it is a good road to ruin when transaction costs are non-zero, even if they are extremely low (graph 5.4).




Fig. 5.3. Without trans. costs




Fig. 5.4. With 0.5% trans. costs

question 3:

The previous graphical analysis is obviously not sufficient. When applying standard performance indexes, especially those including a risk/return analysis (a Sharpe ratios as instance), one can clearly see that all supposed advantages for a MMA_{90-10} vanish as soon as transaction costs are considered (see Table 5.2).

Transaction Costs		0%	0.5%
Buy&Hold	Mean return	2.0344 E-4	
	σ	0.01145	
	Sharpe Ratio	0.0177	
	Portfolio	10871.43	
Basic Strat.	Mean return	2.0964 E-4	1.2267 E-4
	σ	0.00724	0.0073
	Sharpe Ratio	0.02893	0.01679
	Portfolio	12852.61	1183.57

Table 5.1. Performance evaluation based on a MMA90-10

question 4:

Do increasing cognitive abilities for the agents lead to better results in terms of risk/return performance? In other terms, assuming that “*perceiving good signals*” necessarily leads to “*achieving a good strategy*” – and this assertion will be extensively discussed – can we design agents sufficiently smart to adapt their behavior to many signals and many external constraints to outperform the market? Would this kind of agents *prove* really any ability in this game? How should we design an evaluation framework taking these elements into account if we want to design automatic trading platforms and/or tests of market efficiency with agents duplicating as well as possible the behavior and the cognition of true technical traders? What kind of implications, both theoretical and practical, these considerations can highlight?

5.3 On the link between *good market signals* and the capacity for building up simple good strategies

In this section, empirical investigations use daily data from the Euronext Paris Stock Exchange between 1988 and 2005. The traded tracker is now based on the CAC40 index. Agents have access only to past values of this index. We first present some technical/theoretical arguments and propose a series of illustrations afterwards.

We first propose to distinguish two fundamental concepts that must be considered separately previously to be articulated. Technical trading is always based on “signals” indicating either that the market is about to increase or to decrease, and “strategies” based on these signals as well.

1. As evoked previously, a “signal” is generally grounded on the (controversial) idea that profitable persistence or inertia characterize the price motion in stock markets. One difficulty here is to detect which “signal” is actually able to reveal such persistency. We consider in this paper a large number of instances of signals; these instances are based on several generally accepted technical rules (moving averages, rectangle, triangle, RSI, momentum ...), each of them being modeled as a parametric function. These signals will be active or not, depending on the existence of “patterns” in prices provoking their activation. Once activated, the signal sends a recommendation to the trader expressed like: “*according to my own logic*”, “*the market should increase*” or “*should decrease*”.
2. A “strategy” is the way agents use these signals to build a trading behavior.

- a) Some agents will only observe one signal (some being endowed with multiple signals), and will follow it systematically (we call this behavior “Basic Strategy”).
- b) Others will be “contrarians” (i.e. will follow an “Inverse Strategy”)
- c) Others will choose sometimes to follow the signals, sometimes to ignore them. We call them “Lunatic” traders.
- d) ...

Extracting best candidates from a large soup of signals

In this section, a limited sample of results from a series of massive empirical investigations is reported. We select, among many thousands of chartist/technical signals, some of them exhibiting good “Hit Rates” (HR, see equation 5.3) and a minimum activity (that is, signals frequently activated and useful to manage a portfolio – at least one signal per week –). Table 5.3 shows a limited subset of this “good signals” (a “signature” is simply the name and the parameters used to compute this signal).

Num. of signals	with $HR \geq 50\%$	with min activity
110288	6640 (6.02%)	97 (0.08%)
Signature		
MMA-1-4 ; MMA-1-6 ; MMA-1-7		
Momentum-2-1 ; Momentum-5-0		
Variation-1-1-4 ; Variation-1-5-1		
Variation-1-7-1 ; Variation-1-8-1 ; Variation-1-9-1		

Table 5.2. Subsets of “good signals”

Executing these best candidates with a simple strategy

We show how we can use the signals selected in section 5.3 to design “pseudo-good” strategies.

An agent decides, at each time step and according to the set of information it accesses, to manage the portfolio, selling, buying or letting the number of held trackers unchanged. This set of information is as follows:

- $S_1, S_2, \dots, S_n$, the set of signals exploited by her strategy.
- $HR_1, HR_2, \dots, HR_n$, the corresponding set of Hit Rates.

One can notice here that we did not design a very complex set of information, including performance evaluation in terms of risks-returns, rate of activity, memory etc... This is obviously possible but leads to an increasing computing time and a huge amount of data to analyze.

We focus here on the simplest imaginable strategy: one signal, one Hit Rate, no evolution, and a strict application of what the signal suggests : if the market is identified as a rising market, "*Buy*", if identified as a decreasing one, "*Sell*". In all other circumstances, "*stay unchanged*". Table 5.3 presents the results for 10 strategies based on the signals in Table 5.3, for two transaction costs levels. It is illustrated that no strategy is able to outperform the market when transaction costs are fixed at 0.5%.

Signature	0% rate				0.5% rate			
	Mean (10^{-4})	σ (10^{-2})	Sharpe Ratio	Rank /97	Mean (10^{-4})	σ	Sharpe Ratio	Rank /97
B&H	3.0879	1.09	0.0283	—	3.0879	0.0109	0.0283	—
MMA-1-4	3.9428	0.74	0.0529 *	13	-0.010	0.00870	-0.1151	88
MMA-1-6	4.4715	0.74	0.0599 *	2	-6.006	0.00852	-0.07043	69
MMA-1-7	4.1877	0.74	0.0562 *	7	-5.6647	0.00845	-0.06701	67
Mom.-2-1	3.2710	0.81	0.0402 *	51	-9.3296	0.00915	-0.10194	83
Mom.-5-0	4.1035	0.74	0.0552 *	8	-6.3746	0.00835	-0.07630	70
Var.-1-1-4	3.9915	0.74	0.0535 *	12	-6.5312	0.0083	-0.0778	72
Var.-1-5-1	3.1547	1.07	0.0294 *	67	1.5685	0.01081	0.01450	23
Var.-1-7-1	3.0960	1.08	0.0285 *	70	2.8503	0.01087	0.02622	5
Var.-1-8-1	3.0403	1.08	0.0279	71	2.8839	0.01088	0.02650	4
Var.-1-9-1	2.9761	1.08	0.0273	73	2.9091	0.01088	0.0267	3

MMA: mixed moving average, Mom.: momentum, Var.: Variation

* stands for "actually outperform the Market"

Table 5.3. Performance evaluation of 10 strategies based on "good signals"

In figure 5.5 we show that a good signal (MMA 1-4) can lead to disastrous results when transaction-costs are non-zero, while it can be profitable when transaction costs are not paid. This is linked to the fact that a good Hit Rate can produce a lot of activity that will not be profitable because the costs for transacting exceed the benefits one can obtain with small upwards or downwards in prices.

In figures 5.6 and 5.7 and we have extended this analysis including the entire set of agents endowed with signals presenting a HR > 50% (6640 signals, see section 5.3). They are plotted in a risk/return space.




Fig. 5.5. Agents using MMA(1,4) signals

Agents under the market line (black plain line) underperform the B&H strategy.




Fig. 5.6. Without trans. costs




Fig. 5.7. With 0.5% trans. costs

It appears that once transaction costs are implemented, the number of agents being able to exploit their signals in order to “outperform” the market decreases extremely rapidly with limited increments for these costs. It is noticeable that the agents seemingly well-performing are not those endowed with best signals.

5.4 Do intelligent agents outperform ZIT?

In this section, we want to address the following question: do agents endowed with a set of signals of size N behave systematically better than agents endowed with a set of signals of size $N - i, i \in [1, N - 1]$? Do “smart” agents behave better than Zero Intelligence Traders (ZIT)? In other terms, does increasing cognitive skills, that is, the ability to detect potential opportunities to trade, actually lead to a better performance? This is a recurrent question in economics and finance that has provoked many intriguing results (see for instance Gode and Sunder, 1993; Greenwald and Stone, 2001). As stated previously, a first obstacle is the profitable implementation of good signals in the agent. One potential solution could be to allow the agents to select the signals upon which they trade on the basis of their individual Hit-Rate (or some indicator based on this measure).

Technical elements

The first step here consists in allowing each agent to let her rationality evolve along time. To a certain extent, we must consider agents endowed with learning capabilities or adaptive reasoning. This is a specific topic in Agent Based literature (see for example Weiss, 1996), which is not developed here. We just exhibit a limited treatment for this problem:

1. Agents are endowed with N signals (in the following examples $N \in [1, 11]$), previously selected on a large set of signals in order to ensure some (arbitrary) level of “effectiveness”⁴.
2. At each time-step, agents compute for each signal the corresponding Hit-Rate.
3. Every P time steps, agents observe which signal has performed well in terms of HR and select this predictor to trade over the next P time steps. In the following developments, and for the sake of simplicity, $P = 100$.

It is relatively easy to imagine various learning and adaptive procedure that may lead to better results, and it could be argued here that the results shown might be dramatically improved. This is presumably true, although this should be done with a correlative increased complexity of agents’ design, solution which has not been retained in this article.

⁴ We mix three indicators: individual Hit-Rate, number of emitted signals, balance between “buy” and “sell” signals.

Basic strategies based on sets of best signals

We present now one typical answer to an instance of the generic question proposed at the beginning of this section: *“On the basis of the 10 best signals proposed in Table 5.3, is it possible to create basic strategies using many signals (2, 3, ..., 10) in order to outperform the market?”*

It is particularly contra-intuitive to imagine that adding cognitive skills to the agents should lead to a decrease in performance. One should expect to observe a rise in performance for agents accessing a larger set of decision rules when evolving in the market. This is not actually the case.

To answer these questions we create a series of agents endowed with an increasing number of signals, from 1 to 10, $Agent_i$ being endowed with the i – *th* first signals in terms of Hit-Rate. We then investigate their relative performance when transaction costs are respectively fixed at a 0% rate and 0.5% using the adaptive procedure proposed in the technical discussion above. Figure 5.8 and 5.9 clearly show that increasing the number of signals in the agents do not systematically allow for obtaining a higher level of performance in terms of Sharpe Ratio. This




Fig. 5.8. With 0% trans. costs




Fig. 5.9. With 0.5% trans. costs

is obvious when transaction costs affect the agents’ global performance, but it is also generally true with no transaction costs. We have tested all possible values for P between 10 to 500 days, and obtained similar results. These considerations suggest that either the complexity of agents is not appropriate to increase their performance, either an other kind of rule should be implemented to select “good signals” (like their average profitability in terms of return, which is especially complex), or that the market being efficient, technical trading is definitively useless.

5.5 On the validity of technical trading arguments

A last point must be explicitly evoked now: technical trading, and more generally speaking the weak-form market efficiency have been studied using sophisticated statistical tests⁵ to verify if *simple* technical rules can convincingly outperform the market. Nevertheless, a research tackling the question of the relative performance for *complex* technical trading rules, including artificial intelligence agents, able to evolve in a wide decision-rules universe, has still to be done.

This would be the ultimate stage to obtain a strong test for market efficiency. As it has been shown in a previous communication Brandouy and Mathieu (2006), even if one explores an enormous number of signals individually “plugged” in artificial traders playing a “Basic” strategy, it seems to be impossible to obtain risk-adjusted rates of return in excess to a simple Buy and Hold strategy. This is an empirical evidence that strongly support the weak-form EMH.

The following illustrations suggest that if one does not accept to increase significantly the complexity of the agent-based architecture used in this kind of research, it will certainly not be possible to obtain strong evidence of an abnormal over-performance.

Four strategies and the Tale of Technical trading efficiency

In this last empirical investigation, we report results that clearly illustrate the previous discussion. We consider here four strategies using various sets of “good signals”. These four strategies are:

1. Basic strategy, that will serve as a benchmark.
2. Inverse Strategy
3. Deterministic Lunatic Strategy
4. Stochastic Lunatic Strategy

Firstly, we focus on agents endowed with multiple signals⁶ applying them on the daily closing price of the Dow-Jones (see section 5.2). These signals have been selected considering their Hit-Rate over a sub-sample of observations. Agents try to exploit these signals using various strategies, as proposed previously. Their relative performance are compared to a simple Buy and Hold behavior on the same sample. In this example there is no transaction costs.

⁵ Including risk/return measures, in-sample selection and out-of sample tests, data-snooping control procedures see Lo and MacKinlay (1990) for a technical point and Park and Irwin (2004) for a general survey about these topics.

⁶ RSI_{42-20} , RSI_{15-34} , $Momentum_{17-6}$, $Momentum_{13-10}$.

Strategy	Mean return	Standard deviation of returns	Sharpe Ratio
BH	$2.0152 \cdot 10^{-4}$	0.0113	0.0177
Basic	$1.6459 \cdot 10^{-4}$	0.0072	0.0227 *
Inverse	$1.9481 \cdot 10^{-4}$	0.007466	0.02608 *
Lunatic D.	$1.8826 \cdot 10^{-4}$	0.0080	0.0234 *
Lunatic S.	$1.0989 \cdot 10^{-4}$	0.008160	0.01346 *

* stands for “actually outperform the Market”

Table 5.4. Performance of 4 strategies based on “(pseudo)good signals”

Considering this simple illustration, one can see that the best strategy here consists in doing exactly *the opposite* of what the signals suggest (*i.e.* to follow an Inverse Strategy, see table 5.5) . One can also achieve a better Sharpe Ratio with the “Deterministic-Lunatic” strategy than with the “Basic” strategy. One has to keep in mind that this result does not prove any inefficiency in the market because it might well be due to data-snooping, because its stability and robustness has not been checked, and last but not least, because it has been obtained without transaction costs. It is proposed for the sake of illustration and we therefore do not argue that it *proves* any dominance in performance. We only highlight the fact that whatever the “strategy” we consider, one can achieve a similar result with any other kind of strategy (apart “Stochastic-Lunatic”, which basically is similar to a coin toss).

Some other amazing results

We now briefly propose some results of massive investigation on French data (see section 5.3) leading to similar conclusion.

Cheating is not playing: The following “strategy” is only given to fix some kind of boundaries. We call it the “cheating strategy”. It has been designed to allow the agents to know at date t what will happen at date $t + 1$. They can therefore directly benefit from this information to (easily) outperform the market. The result of this behavior (Sharpe Ratio = 0.46349) is presented in figure 5.10. Our best non-cheating agent using a single signal is only able to produce 14.35% of this performance.

Good performance on bad basis: It is perfectly possible to design good agents (obtaining a Sharpe Ratio over the B&H one). As instance, signals “Variation-2-7-14” and “MA-85” obtain very bad Hit-Rates.

When these signals are “plugged” in an agent playing a Basic strategy and switching from one to the other every 500 dates (with respect to their relative Hit-Rate at these dates) we obtain a very satisfying performance with a Sharpe Ratio of 0.0288, while B&H Sharpe Ratio is 0.0283.

Signals do not signal anything: As quoted previously, it is frequently better to do exactly the opposite of what the signals suggest: if one wants to trade using a signal “indic-7-18-5” with an “Inverse Strategy” one should obtain a Sharpe Ratio of 0.0666 while following this signal would lead to a Sharpe equal to -0.0292 with a “Basic Strategy”.

On the nature of the best strategies: Our set of signals is composed of 360.288 elements, 250.000 of them being “periodic signals”: they propose to go long after “n” days and to go short after “m” other days. They cannot really be called “technical” signals but they can catch some special patterns such as the so-called “Monday Effect”. Nevertheless, many of them can simply be analyzed as stochastic signals or zero-intelligence signals. Nevertheless, each of the 200 first agents ranked by Sharpe Ratio use these kind of signals. The best agent is therefore plugged with a “periodic signal 21-56” (obtaining a Sharpe Ratio equal to 0.0467). It is easy to find a similar agent using an “Inverse” strategy based on periodic signals, and behaving nearly as well as this pseudo-champion.

Thus, if one only scratches the surface of weak-form market efficiency, there is nothing to expect from technical trading. In other words, little evidence in terms of superior performance should arise from a cautious analysis of simple active trading rules. Nevertheless one cannot seriously affirm that these last tests completely answer the question.

This set of results as well of the elements we have discussed in this paper strongly suggest that:

1. Automatic trading based on technical analysis depends upon external factors such as leverage, transaction costs. There is an enormous variability in performance linked to these parameters.
2. It appears necessary to separate at least “signals” and “strategies”. Naïve increases in agents cognitive skills are also useless to achieve satisfactory levels of performance (once incorporating risk). A fine-tuning aiming to balance the complexity of agents’ capabilities and information resources is necessary.
3. To go deeper in this analysis would imply the definition of generic strategies describing learning procedures, adaptation and decision making processes.




Fig. 5.10. With 0% trans. costs

Therefore, from a conceptual point of view, a robust framework for the evaluation of Agent-Based trading and technical analysis should systematically answer each of these 3 points at least, which obviously constitute a first step before rigorous statistical examinations.

Acknowledgement. This work has received a grant from European Community – FEDER – and “Region Nord-Pas de Calais” – CPER TAC – and from ACI “Systèmes Complexes en SHS”

References

- O. Brandouy and P. Mathieu. Large-scale agent-based simulations and the efficient markets hypothesis. In C. Bruun, editor, *Advances in Artificial Economics*, volume 584 of *Lecture Notes in Economics and Mathematical Systems*, pages 15–27. Springer, 2006. 3-540-37247-4.
- W. Brock, J. Lakonishock, and B. LeBaron. Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, 47:1731–1764, 1992.

- M. Dempster and C. Jone. Can technical pattern trading be profitably automated ? 1. channels. Working paper, Center for Financial Research Working paper, uge Institute, University of Cambridge, 2005.
- P.J. Detry and P. Gregoire. Other evidences of the predictive power of technical analysis: the moving-average rules on european indices. In *Proceedings of the Lugano European Financial Management Association Conference*, 2001.
- E. Fama. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25:383–417, 1970.
- D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, February 1993.
- A.R. Greenwald and P. Stone. Autonomous bidding agents in the trading agent competition. *IEEE Internet Computing*, 5(2), 2001.
- T. Hellström and K. Holmström. Predicting the stock market. Technical Report IMA-TOM-1997-07, Center for Mathematical Modeling, Department of Mathematics and Physics, Märdalen University, Sweden, August 1998.
- M.C. Jensen and G.A. Benington. Random walks and technical theories: Some additional evidence. *Journal of Finance*, 25:469–482, 1969.
- A.W. Lo and A.C. MacKinlay. Data-snooping biases in tests of financial asset pricing models. *Review of Financial Studies*, 3(3):431–467, 1990.
- C.H. Park and S.H. Irwin. The profitability of technical analysis. *Ag-MAS project Research Report 2004-4*, October 2004.
- G. Weiss. Adaptation and learning in multi-agent systems: some remarks and a bibliography. In G. Weiss and S. Sen, editors, *Adaptation and Learning in Multi-Agent Systems*, volume 1042 of *Lecture Notes in Artificial Intelligence*, pages 1–21. Springer-Verlag, 1996.
- H. White. A reality check for data snooping. *Econometrica*, 68(5): 1097–1126, September 2000.

Which Market Protocols Facilitate Fair Trading?

Marco LiCalzi and Paolo Pellizzari

Dept. Applied Mathematics and SSAV, U. of Venice, Italy
[licalzi,paolop]@unive.it

6.1 Introduction

The evaluation of an exchange market is a multi-faceted problem. An important criterion is the ability to achieve allocative efficiency. Gode and Sunder (1993) shows that a continuous double auction for single-unit trades leads to an efficient allocation even when the traders exhibit “zero-intelligence”; in other words, market protocols are active contributors in the search for a better outcome. Under reasonable circumstances, most of the commonly used market protocols share the ability to help traders discover an efficient allocation.

As suggested in Hurwicz (1994), however, the attainment of allocative efficiency is only a necessary condition for the effectiveness of a trading protocol and one should take into account other dimensions. Assuming zero intelligence, LiCalzi and Pellizzari (2007) compares the performance of different market protocols with regard both to allocative efficiency and other criteria such as excess volume or price dispersion. Their study considers agents with decreasing marginal utility that can repeatedly make single-unit trades and examines four common protocols: batch auction, continuous double auction, nondiscretionary dealership, and a hybrid of these latter two. All protocols exhibit a remarkable capacity to achieve allocative efficiency. However, stark differences in performance emerge over the other dimensions. These differences persist even when the assumption of zero intelligence is removed; LiCalzi and Pellizzari (see 2006).

The general conclusion is that although common market protocols may be close substitutes in helping (even zero-intelligent) traders to attain efficiency, they behave quite differently in many respects. This paper expands this line of research moving from the evaluation of al-

locative effectiveness to the assessment of allocative fairness. See Fehr et al. (1993) for a different line of attack on this theme.

Any trading protocol that attains allocative efficiency has two effects. From a static point of view, it moves the traders from their initial endowment to a final (efficient) position where no further paretian improvements are possible and all gains from trade are realized. This ability to help traders discover and exploit all gains from trade pertains to the allocative effectiveness of a market protocol. From a dynamic point of view, on the other hand, the denouement of a trading session decides how these gains are distributed among the traders. The performance of a trading protocol in this respect pertains to its allocative fairness.

A protocol that is allocatively efficient never leaves unrealized gains from trade. A protocol that is allocatively fair makes sure that these gains are equitably distributed among the traders. While many definitions of equitability are possible, there is a general sense that each traders should be entitled to a share of the gains from trade that his being in the market creates. In this paper, we consider the same four common protocols that we have shown to be allocatively efficient (even under zero intelligence) and we ask the following question. Suppose that the market is populated with only two families of agents. Both families are using trading strategies that are individually rational, but the second family enjoys a potential trading advantage on the first one. Which of these market protocols is more effective in making sure that the first family of agents overall loses the least on his “fair share” of gains from trade?

The organization of the paper is the following. Section 6.2 describes the model tested in our computational experiments and formalizes our research question. Section 6.3 details the experimental design and provides detailed instructions for its replication. Section 6.4 reports on the results obtained and Section 6.5 offers our conclusions.

6.2 The model

We use the same setup as in LiCalzi and Pellizzari (2007), where a simple exchange economy admits a unique efficient allocation for the single good to be traded. Given that the market protocols attain allocative efficiency, this implies convergence to the same final allocation of the good and facilitates comparisons.

6.2.1 The environment

We consider an economy with n traders. There is cash and one good, which we call “stock”. Each trader i has an initial endowment of cash $c_i \geq 0$ and shares $s_i \geq 0$. Each trader i has CARA preferences over his final wealth, with a coefficient of risk tolerance $k_i > 0$. Therefore, trader i ’s excess demand function for stock (net of his endowment s_i) is the linear function

$$q_i(p) = \tau k_i(\mu - p) - s_i. \quad (6.1)$$

where μ is the mean and $\tau = 1/\sigma^2$ is the reciprocal of the variance (a.k.a. as the “precision”) of the distribution of the final value of the stock. Each trader knows μ and τ as well as his endowment and his demand function, but otherwise has no information on the other agents.

Let $K = \sum_i k_i$ be the sum of traders’ coefficients of risk tolerance, while $S = \sum_i s_i$ and $C = \sum_i c_i$ are the total stock and cash endowments. The unique efficient allocation of shares in this economy requires that trader i holds $s_i^* = (S/K)k_i$ shares of the stock. This is also achieved in the (unique) competitive equilibrium at price $p^* = \mu - S/(\tau K)$; see Wilson (1968). Clearly, the unique efficient allocation of shares is associated with a continuum of feasible allocations for cash; each of these determines a different apportionment of the gains from trade. Therefore, allocative efficiency corresponds to handing out stock in a unique way; allocative fairness has to do with how cash is redistributed during the trading that takes place before the efficient stock allocation is attained.

We emphasize that our setup is not meant to replicate the structure of a stock market; in particular, informational effects are ruled out. The underlying economy can be described as an exchange market for one good, where traders have strictly decreasing linear demands and heterogeneous preferences that are driven by a particularly simple parameterizations.

6.2.2 The market protocols

We compare the performances of four market protocols: a batch auction, a continuous double auction, a nondiscretionary dealership, and a hybrid of these last two. The first protocol is simultaneous, while the other three are sequential. The following features are common to all protocols. See LiCalzi and Pellizzari (2006, 2007) for a complete description of the protocols and details on their implementation.

A protocol is organized in trading sessions (or days). Agents participate in every trading session, but each of them can exchange at most one share per session. Reaching an efficient allocation requires multiple rounds of trading. If the protocol is sequential, the order in which agents place their orders is randomly chosen for each trading session. If the protocol is simultaneous, all orders are made known and processed simultaneously so the time of their submission is irrelevant. The books are completely cleared at the end of each trading session. Prices are ticked and, for convenience, the tick is set equal to 1; in other words, prices must be integers.

6.2.3 Behavioral assumptions and fair shares

The following behavioral assumptions hold for each trader. An agent is restricted to trade one unit at a time. Budget constraints must be satisfied. Given the demand function (6.1), trader i has decreasing marginal utility for additional units. If the current endowment of a trader is s , his valuation for the next unit to trade is

$$v_i(\pm 1) = \mu - \frac{s \pm 1}{\tau k_i} \quad (6.2)$$

where the $\pm$ sign depends on whether the attempted trade is a purchase or a sale. Hence, his reservation price depends on the side of the transaction he is entering and on his current endowment s_i . Moreover, his certainty equivalent for holding quantities c and s of cash and stock is

$$m_i(c, s) = c + \left(\mu - \frac{s}{2\tau k_i} \right) s \quad (6.3)$$

It is worth noting that the certainty equivalent m_i accounts for c at face value but evaluates s using an individual “price of risk” $\mu - [s/(2\tau k_i)]$.

The initial endowment (c_i^0, s_i^0) of a trader i provides him with a certainty equivalent $m_i^0 = m_i(c_i^0, s_i^0)$. We define his “fair share” m_i^* of gains from trade as the certainty equivalent he would attain under the fictitious protocol of Walrasian tâtonnement, where a centralized market maker iteratively elicit traders’ excess demand functions and keeps adjusting prices to equilibrate them *before* trade takes actually place. Under standard conditions, this protocol is a natural benchmark because it attains allocative efficiency in one giant step, while simultaneously minimizing both the volume of transactions and price dispersion. For later use please note that, under this protocol, a trader with *ex ante* knowledge of the equilibrium price p^* would attain exactly the

same final certainty equivalent and thus would not be able to increase his fair share.

Under the Walrasian protocol, a trader i ends up with cash $c_i^* = c_i^0 - p_i^*(s_i^* - s_i^0)$ and stock $s_i^* = (S/K)k_i$. After substitution, the certainty equivalent of his fair share is

$$\begin{aligned} m_i^* &= c_i^0 - p^*(s_i^* - s_i^0) + s_i^* \left(\mu - \frac{s_i^*}{2\tau k_i} \right) = \\ &= c_i^0 + \left(\mu - \frac{S}{2\tau K} \right) s_i^0 + \frac{S}{2\tau K} \left(\frac{S}{K} k_i - s_i^0 \right) \end{aligned}$$

which nicely decomposes into the sum of three terms. The first one is the initial cash endowment of trader i ; the second is the “value” of his initial stock endowment at the market price of risk; the third one is a positive correction term that is increasing in the difference between the efficient and the initial stock endowment for i . Since trading is voluntary, individual rationality implies that the difference between the fair share and the initial certainty equivalent for each i is positive:

$$m_i^* - m_i^0 = \frac{(Ks_i^0 - k_i S)^2}{2\tau k_i K^2} \geq 0$$

We expect that market protocols affect how much of their fair share different families of agents manage to obtain in the end. This requires to aggregate social welfare over groups of agents. We measure the social welfare of a group G by the sum of the certainty equivalents across the traders in G . Given the initial endowments (c_i^0, s_i^0) of each trader i , the (initial) social welfare of the entire traders’ population is $M^0 = \sum_i m_i(c_i^0, s_i^0)$. After reaching an efficient allocation, the social welfare increases to

$$M^* = \sum_i m_i^* = C + \left(\mu - \frac{S}{2\tau K} \right) S \quad (6.4)$$

which is the analog of Equation (6.3) at the market level. We slightly abuse notation here, because M^* is achieved by any efficient allocation including (but not limited to) the one induced by the Walrasian procedure. Looking at the left-hand side of Figure 6.1, efficient trading expands the pie from M^0 to M^* .

Consider now a strict subset G of traders. They start with an initial endowment that corresponds to a social welfare $M_G^0 = \sum_{i \in G} m_i^0$ for the group G . The fair share of this group is $M_G^* = \sum_{i \in G} m_i^* \geq M_G^0$. In the right-hand side of Figure 6.1, we represent M_G^0 as the circular sector from the inside circle and M_G^* as the union of M_G^0 and the annular sector



Fig. 6.1. Gains from trade and fairness.

topping it. In general, M_G^* expands but need not be proportional to M_G^0 . Suppose now that at the end of a trading protocol, the social welfare of a group G is $M_G^* \cup A_G$ so that the group G is extracting higher gains from trade than its fair share. Then we say that the protocol has been too favorable to the traders in G or, equivalently, that it has been unfair to the traders in the complementary set G^c . Hence, allocative fairness is about how the larger pie created by trading is redistributed among different groups of traders. Similarly to a zero-sum game, a trader gets more than his fair share by taking away a piece of someone's else fair share.

Our approach to study allocative fairness is to split the traders' population into two families and compare the ability of market protocols to prevent one group from exploiting the other one. For realism, we assume that all agents are individually rational: regardless of which family he belongs to, each agent accepts a trade only if this cannot decrease his current certainty equivalent. An agent who undertakes a sequence of trades over time increases (possibly, weakly) his own certainty equivalent in each transaction. This assumption, for instance, is consistent with zero-intelligence.

Our two families of interacting traders are chosen to emphasize differences in the ability to appropriate gains from trade. Notably, individual rationality alone cannot prevent a purchase from an inframarginal seller even if this reduces the potential gains from a specific trade. Put differently, individual rationality protects a buyer from making a personal loss on a trade but does not imply that he is trading with the "right" counterpart. This stronger guarantee requires knowledge of the equilibrium price p^* in order to spot and refuse inframarginal trades. We assume that some traders satisfy only individual rationality while others can do better because they know p^* as well.¹

The first group is formed by the truth-telling (from now on, TT) traders described in LiCalzi and Pellizzari (2007). At the start of a

¹ An alternative assumption is that only the second type of traders are able to compute or deduce p^* from the available information.

trading session, a TT trader chooses with equal probability on which side of the market (buy or sell) he attempts to trade one unit. Suppose he goes for a purchase; the case of a sale is analogous. Given his current endowment, the agent knows that his valuation for the next unit to buy is $v_i(+1)$ from Equation (6.2). In a batch auction, he truthfully bids $v_i(+1)$. In a sequential protocol, he checks first if the best current ask price is $p \leq v_i(+1)$; if so, he buys one unit at p . Otherwise, he places a bid equal to $v_i(+1)$. In other words, when no better deal is available, a TT buyer posts a bid equal to his current valuation for the next unit to buy and thus “truthfully” reveals his reservation value. Compared to zero-intelligence trading, a TT agent is less greedy because he posts the largest bid that is individually rational given his own valuation. When a TT agent buys one unit at a price p higher than the equilibrium price p^* , he increases his certainty equivalent but eats up a piece $(p - p^*)$ of his fair share.

The second group of agents consists of traders that know the correct equilibrium price; we call them price-informed (from now on, PI). This extreme assumption is a very parsimonious way to endow these agents with the ability to cut down on inframarginal trades and make sure that they never lose on their fair share. Given his current endowment, a PI agent knows that he should be a buyer if $v_i(+1) \geq p^*$ and a seller if $v_i(-1) \leq p^*$. Therefore, he never needs to guess which side he should take.

Suppose that the PI agent should be a buyer; the opposite case is analogous. In a batch auction, he simply bids p^* . In a sequential protocol, a PI trader must take action when he is called out and cannot wait for better terms. When it is his turn, he first looks for “sure deals” by checking whether the best current ask price is $p_a \leq p^*$ or the best bid price is $p_b \geq p^*$; if so, he buys or sell one unit, respectively. Otherwise, and limitedly to the two book-based protocols, a PI agent places a bid that improves the current best bid p_b by one tick and achieves time-price priority at a buying price never greater than p^* . In general, the trading strategy of a PI agent has three characteristics: first, he never fails to exploit opportunities for trading off the equilibrium price; second, he never trades at a price worse than p^* (and hence never loses on his fair share); third, conditional on these two constraints, he maximizes the probability of trading in the right direction. This last restriction is chosen to emphasize the ability of PI traders to take advantage of TT agents.

Depending on the protocols and the random sequence of trades, the attainment of full allocative efficiency may sometimes fail. For

instance, in the nondiscretionary dealership, the existence of a fixed bid-ask spread may prevent two or more TT agents from completing their few last trades. This may (albeit marginally) reduce the overall gains from trade and lower allocative efficiency, confusing our study of allocative fairness. To rule out this spurious effect, after all trading opportunities within the protocol are exhausted, we force agents to carry out all residual efficient trades at price p^* . We emphasize that this has the only purpose of actually realizing the full pie M^* so that we can concentrate on its redistribution; in particular, none of these final trades eats up on the fair share of a trader.

Let M_G^T be the final fair share of a group G when trading takes place using a trading protocol T . Given their information and trading strategies, only PI agents can “exploit” TT traders. Therefore, whenever allocative efficiency is attained, $M_G^T \geq M_G^*$ for $G = PI$ and any protocol T among the four we consider. We can thus test the ability of a trading protocol T to foster a fair allocation by comparing $M_G^T - M_G^*$ for $G = PI$.

Clearly, the ability of the PI group to exploit TT traders depends also on the proportion π of PI traders in the market. The more the exploiters, the harder becomes the competition for trades at prices different from p^* . Therefore, we study how allocative fairness is affected by the proportion π in $(0, 1)$. Endpoints of the interval are ruled out to avoid trivialities.

6.3 Experimental design

6.3.1 Identification

The global parameters are the number n of traders, the mean μ and the variance σ^2 of the realization value of the asset, the number t of trading sessions, and the number λ of PI traders. (The proportion of PI agents is $\pi = \lambda/n$.) Individually, a trader i is characterized by his coefficient k_i of risk tolerance and by his endowment of cash c_i and asset shares s_i . Finally, for protocols involving the dealer, we need to select her initial quotes and a (fixed) spread.

The exemplar for our simulations is similar to that one used in LiCalzi and Pellizzari (2006). The basic parametric configuration is reported in Table 6.1. The ratio $S/K = 2$ implies that the competitive equilibrium price is $p^* = \mu - \sigma^2(S/K) = 760$. The initial dealer’s quotes in the nondiscretionary dealership are a bid of 755 and an ask of 765, with a fixed bid-ask spread of 10. In the hybrid protocol, where

	Parameters	Initialization
Global	n	$= 1,000$
	μ	$= 1,000$
	σ^2	$= 120$
	t	$= 500$
	λ	$= \text{integer in } (0, n)$
Trader	k_i	$= \text{divisors of } \sigma^2 \text{ in } \{10, \dots, 40\}$
	c_i	$= 50,000$
	s_i	$= \text{permutation of } 2k_i$

Table 6.1. Exemplar for identification.

the dealer's presence restricts the ability of PI traders to steal better deals, the initial bid and ask prices of 745 and 775 exactly straddle the equilibrium price of 760, with a fixed spread of 30.

The robustness tests reported in Section 6.4.1 change one parameter at a time with respect to this exemplar. We have worked out simulations where the ratio S/K is 1 (or 3), making the equilibrium price higher (lower); where the dealer's fixed spread in the nondiscretionary dealership is 6 (or 30), making the market more (less) liquid; and where the fixed spread in the hybrid protocol takes different values between 4 and 300, making the dealer's presence more or less influential.

6.3.2 Simulations and data representation

A round of testing simulates traders' behavior in 4 different protocols for different values of λ . A typical cycle is run as follows. We fix an integer value of λ in the range $\{1, \dots, n-1\}$ and then we randomly choose different queues of traders for each trading session. These choices are kept fixed across the four protocols, so that each of them is tested using the same fraction of PI traders and the same orderings in each trading sessions. All other parameters are instantiated as per the exemplar in Table 6.1. The number of agents is $n = 1000$; we run 999 trials per cycle and test each value of λ from 1 to 999. At the end of each simulation, we compute and record all relevant statistics. The simulations are run using a package of routines written in Pascal. The statistical and graphical analysis of the data are made using R, an open-source environment for statistical computing available at <http://www.r-project.org/>.

We use two (normalized) measures to assess the allocative fairness of a protocol. Let M_G^T be the final share of the group G when trading takes place using a trading protocol T and M_G^* their fair share (using

the Walrasian protocol). As discussed above, only PI agents can “exploit” TT traders; hence, we fix $G = PI$ for the rest of the paper. The first measure is the absolute *excess gain* $(M_G^T - M_G^*)/M^0$ for the group G . The division by the size M^0 of the initial pie is a normalization introduced to make the index scale-free and allow direct comparisons; however, for simplicity, in the rest of the paper we write the absolute excess gain as $M_G^T - M_G^*$ and leave the normalization implicit. The second measure is the relative *excess gain* $(M_G^T - M_G^*)/M_G^*$. The absolute excess gain reports how much welfare PI traders collectively take away from TT traders with respect to the initial pie. The relative excess gain measures how much (on average) a PI trader is expected to improve his final welfare by trading within a given protocol.

A graphical representation of each set of data is obtained as follows. Given a protocol T , we plot the 999 data points produced in a simulation. We then fit a smoothing function generated by applying a Friedman smoother to all the data points associated with the same protocol; see Venables and Ripley (2002). Reading Figure 6.2 from left to right exemplifies this procedure for the case of a continuous double auction.




Fig. 6.2. Realizations (left) and a superimposed Friedman smoother (right).

6.4 Results

Figure 6.3 shows two representative pictures based on our exemplar. The figure on the left reports the (normalized) absolute excess gain

$M_G^T - M_G^*$ collectively achieved by the PI traders as a function of their cardinality λ for four protocols: batch auction, continuous double auction, nondiscretionary dealership, and the hybridization of these two latter protocols. Note that dividing λ by $n = 1000$ gives the proportion π of PI agents active in the market.




Fig. 6.3. Absolute (left) and relative (right) excess gain for PI traders.

The first comment is that the batch auction protects the TT traders much more effectively than any other protocol for both measures and for any number of PI traders. This is not surprising: the batch auction protocol requires simultaneous submission of trading orders and is therefore much more difficult for PI agents to exploit. By posting an order at p^* , each PI trader maximizes the probability of trading under the constraint of never losing on his fair share. Whenever the trading price issued in a session of the batch auction is different from p^* , he cuts away a piece of a TT trader's fair share. However, because the batch auction aggregates all the orders received in a trading session, it is very unlikely to issue a trading price different from p^* . We can thus shift our focus of interest to the three sequential protocols. For completeness, however, we report also the data relative to the batch auction.

The second comment is that in general the absolute excess gain for sequential protocols is a unimodal function of λ . Therefore, the collective ability of PI to exploit TT traders peaks at some intermediate value of λ . In this respect, there is a *natural ordering* of protocols from dealership to hybrid protocol to continuous double auction that appears

twice. First, the maximum excess gain for PI traders is increasing in the natural ordering of protocols. Simultaneously, the value of λ that maximizes the PI excess gain is decreasing. In other words, PI traders can achieve a greater excess gain in a continuous double auction, but their most effective proportion in such protocol is lower.

The result that the excess gain is increasing in the natural ordering is a direct consequence of the “protection” that the dealership provides. Because the dealer posts bid and ask prices that tend to straddle the correct p^* , the transaction price is never too different from this latter price; hence, no much fair share can be lost. The result that the λ 's maximizing total excess gain are decreasing in the natural ordering can be heuristically explained by the combination of two effects. Intuitively, PI traders are most effective to exploit TT traders when their proportion is neither too low (there must be enough exploiters around) neither too high (there must be enough people to exploit). But we can put a bit more flesh on this explanation.

Consider the continuous double auction. The overall fair share for the TT group that PI traders can appropriate is roughly proportional to $(1 - \pi)$. On the other hand, taken as a group, the TT traders can lose a piece of their fair share only when one of them trades with a PI agent at a price different from p^* . The probability of a PI agent being matched for trade with a TT agent is roughly proportional to $\pi(1 - \pi)$. Therefore, the excess gain appropriated by the PI group in the continuous double auction are approximately proportional to $\pi(1 - \pi)^2$ and the maximum should be attained around $\hat{\pi} = 1/3$, corresponding to $\lambda = n\hat{\pi} = 333$ in our exemplar. The actual value is somewhat lower because some matchings between PI and TT agents do not lead to any trade.

Consider now the dealership. The overall fair share that PI traders can appropriate is still roughly proportional to $(1 - \pi)$. Moreover, because they can only trade with an impersonal dealer, the probability of a trade involving a PI agent is roughly proportional to the fraction π . Therefore, the excess gain for the PI group is now approximately proportional to $\pi(1 - \pi)$ and the maximum should be attained around $\hat{\pi} = 1/2$, corresponding to $\lambda = 500$ in our exemplar. As before, the exact value of the maximizer is affected by microstructural considerations that this heuristic argument does not capture. Finally, the corresponding values for the hybrid protocol are a convex combination of those of the parent protocols.

The third comment is that there are no important differences among sequential protocols when π (or λ) is sufficiently large, because there

are too few TT traders to be exploited. The overall fair share to be appropriated is roughly proportional to $(1 - \pi)$ and for large π there is simply too little to be taken away by the PI group. Moreover, markets with a high proportion of PI traders tend to exhibit a similar degree of allocative fairness because a PI agent never loses on his fair share. Therefore, we restrict the following comparisons to $\pi \leq 40\%$, corresponding to $\lambda \leq 400$ in our exemplar. For any proportion $\pi \leq 40\%$, the ranking over sequential mechanisms concerning their ability to prevent PI agents from eroding TT traders' fair shares is clear-cut and follows the natural ordering.

The right-hand side of Figure 6.3 reports the relative extra gain $(M_S^T - M_S^*)/M_S^*$ collectively achieved by the PI traders as a function of their number λ for the three sequential protocols. Unsurprisingly, this shows that increasing the number of exploiters makes their "looting" less effective for each protocol. Moreover, the ranking is again clear-cut and follows the natural ordering. Finally, this effect is essentially unchanged in all the additional tests reported in the following section.

6.4.1 Tests of robustness

We have run some robustness tests by changing one parameter at a time in the exemplar. The first test looks at differences in the total endowment of stock, leading to a different equilibrium price p^* . The exemplar has a ratio $S/K = 2$ yielding $p^* = 1000 - 120(S/K) = 760$ and generates the left-hand side of Figure 6.3. We keep the same k_i for each trader i , but endow him with a different multiple of his original endowment s_i . This changes the ratio s_i/k_i and of course S/K as well. Figure 6.4 reports data when $S/K = 1$ (on the left) and $S/K = 3$ (on the right), corresponding respectively to a smaller and to a larger total endowment of stock. The equilibrium prices are now 880 and 640, respectively. We adjust the initial dealer's quotes accordingly, making sure that they always exactly straddle the equilibrium price.

Comparing the two figures from Figure 6.4 (as well as the right-hand side of Figure 6.3) shows that a larger stock endowment S increases the relative excess gain $(M_G^T - M_G^*)/M_G^*$ of PI for each protocol and each λ . We do not report the figures for the absolute excess gain to preserve space, but they exhibit a similar increasing effect. In fact, the following argument shows that, over our range of choices for S/K , an increasing relative excess gain implies an increasing absolute excess gain. The quantity M_G^* is roughly proportional to πM^* ; in turn, M^* is increasing in (S/K) as far as $\tau\mu \geq (S/K)$ — as seen by differentiating (6.4) with respect to S . As this inequality holds for our choices of S/K ,




Fig. 6.4. Different equilibrium prices: $S/K = 1$ (left) and $S/K = 3$ (right).

the denominator of the relative excess gain is increasing and the claim follows. Put differently, this shows that a higher stock endowment S brings about a roughly proportional increase in the absolute excess gain.

The increase in the relative excess gain exhibited by Figure 6.4 is a stronger property that is explained by a second perhaps less obvious effect. *Ceteris paribus*, a larger S increases the number of trades that need to be carried out in order to reach the allocative efficiency. Each of these trades is a potential opportunity for PI agents to exploit, making them more likely to extract excess gain from the TT agents. This second effect accounts for the increase in the relative excess gain.

The second test considers the effect of changing the dealer's fixed spread in the nondiscretionary dealership, while keeping his initial quotes centered around the equilibrium price. The exemplar has a fixed spread of 10. Figure 6.5 reports the absolute excess gain when the fixed spread is 6 (bottom), 10 (middle), or 30 (top). The lower the spread, the more influential is the dealer's ability to constrain prices within a narrow band that individually rational trading naturally tends to keep around the equilibrium price p^* . Forcing the transaction price to lie in a band, of course, protects TT agents from more serious mispricings and hence reduces the ability of PI traders to exploit them. Accordingly, we see in Figure 6.5 that the absolute excess gain is increasing in the dealer's fixed spread for any number λ of PI traders.

A third test checks the effect of changing the dealer's fixed spread in the hybrid protocol dealership where an agent has access both to the dealer's quotes and to a book fed with limit orders from other traders.




Fig. 6.5. Different fixed spreads in the dealership: 6, 10, 30 (bottom to top).

The exemplar has a fixed spread of 30. Figure 6.6 reports the absolute excess gain when this fixed spread takes five different values from 4 (bottom) to 300 (top). The absolute excess gain is increasing in the dealer's fixed spread for any number λ of PI traders. This effect and its explanation are analogous to the above. There is a second more interesting effect to note. In terms of the ability to control the absolute excess gain, the continuous double auction is the limit case of the hybrid protocol as the fixed spread goes to $+\infty$. When the dealer posts bid and ask that are too far apart, trading takes place only on the book. Accordingly, as we move from a low to a high spread, the excess gain curve morphs from the shape associated with a dealership to the shape associated with a continuous double auction; for instance, the peak increases and shifts leftward.

6.5 Conclusions

We have studied the performance of four market protocols with regard to their ability to equitably distribute the gains from trade among two groups of participants in an exchange economy. We assume Walrasian tâtonnement as benchmark and define the fair share that should accrue to a trader as the certainty equivalent he would attain under this procedure.

When necessary, the first group of traders bids or asks their reservation value; this makes sure that trading never decreases their own certainty equivalent but exposes them to a possible loss on their fair




Fig. 6.6. Different dealer's spreads in the hybrid protocol: 4, 10, 30, 100, 300 (bottom to top).

share. The second group of traders knows (or can compute) the equilibrium price p^* and uses this information to make sure that trading cannot reduce either their certainty equivalent or their own fair share.

We test the allocative fairness of protocols by running (computerized) experiments where these two families of traders interact with each other. We find that there is a clear-cut ranking of protocols with respect to allocative fairness, defined as their ability to prevent PI agents from eroding TT traders' fair shares. Going from best to worst, this ranking is: batch auction, nondiscretionary dealership, the hybridization of a dealership and a continuous double auction, and finally the pure continuous double auction. The same ranking holds when we replace the absolute excess gain for PI traders with their relative excess gain.

Acknowledgement. We thank two referees for their comments and acknowledge financial support from MIUR.

References

- E. Fehr, G. Kirchsteiger, and A. Riedl. Does fairness prevent market clearing? an experimental investigation. *The Quarterly Journal of Economics*, 108(2):437–459, May 1993.
- D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual ra-

- tionality. *Journal of Political Economy*, 101(1):119–137, February 1993.
- L. Hurwicz. Economic design, adjustment and mechanism. *Economic Design*, 1:1–14, 1994.
- M. LiCalzi and P. Pellizzari. The allocative effectiveness of market protocols under intelligent trading. In C. Bruun, editor, *Advances in Artificial Economics*, pages 17–29. Springer, 2006.
- M. LiCalzi and P. Pellizzari. Simple market protocols for efficient risk sharing. *Journal of Economic Dynamics and Control*, 2007. Forthcoming.
- W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S*. Springer, fourth edition, 2002.
- R. Wilson. The theory of syndicates. *Econometrica*, pages 119–132, 1968.

Market Dynamics and Efficiency

An Artificial Economics View of the Walrasian and Marshallian Stability

Marta Posada, Cesáreo Hernández, and Adolfo López-Paredes

INSISOC, University of Valladolid, Spain
{posada,cesareo,adolfo}@insisoc.org

7.1 Introduction

The experiments discussed below are an attempt to examine two concepts of instability which stem from two different models of market adjustment used in Economics: Walrasian (W) and Marshallian (M) instability. The M model views volume as adjusting in response to the difference between demand price and supply price at that volume. The W model views price as changing in response to excess demand at that price. Do the M and the W models have a firm foundation on micro-motives, or are they just macro abstractions that we could dispense of in Microeconomics?

If there is an awkward question one has to suffer when teaching economics, this is *who does the job of the so called market adjustment?* In contrast with the usual microeconomic models, these processes of adjustment do not represent any optimization of the economic agents' behavior, but just differential equations brought out from nowhere. As Nicholson (1997, chapter 19) states in his well known microeconomics book: "*in the last instance this speculation on the adjustment mechanism hardly makes any sense, because neither the Walrasian nor the Marshallian adjustment reflect the real behavior of the economic agents*". Being virtual market adjustment models taken from the mechanical analogy of vector fields, they can be compatible with different micro behaviors of the participant agents. Nevertheless it is interesting to further asses the relevance of both concepts of instability, and to which extent they emerge from individual agents' behavior. To this end we will use the empirical evidence obtained from the simulation with an Multi Agent Based Model (MABS).

MABS has become a popular tool to theorize about distributed but interacting phenomena as markets and economic activity are. Economics is a Social Science. Being social inherits complexity and being a science calls for experimentation. Historical records are not enough to test economic theories. Experimental Economics with humans has provided a replicable Lab that provides further empirical data to suggest new solutions to economic problems and to test standing models and theories. But human behavior in the experiments is not directly controllable and the question of what the agents' behavior is remains open. Artificial Economics, mainly MABS, has broadened the scope of Experimental Economics, allowing the experimenter to check alternative individual behavior. In this sense MABS is a "killer application" of Economic Theory.

This paper addresses four questions: Can both stable and unstable equilibria be observed in a Continuous Double Auction? If markets do exhibit instability, which of the two models, M/W, will lead to the right equilibria prediction? How robust are the results against alternative learning agents? In view of these results, what is the interest of both instability concepts for policy modeling and simulation?

7.2 Scope and related work

Some comments are convenient to assess the scope of the paper. There are two strands of researchers in Artificial Economics: people interested in the computational properties and policy issues of a given aggregated market and those interested in growing markets with desirable properties from agents with micromotives. Since in microeconomics we work with aggregated markets at a higher level of abstraction than in MABS, an important issue is often overlooked. The market has three dimensions (Smith, 1989): the institution (it is both the exchange rules and the way the contracts are closed, and the information network), the environment (agents' endowments and values, resources and knowledge) and the agents' behavior. These dimensions are frequently overlooked in Economics, leading to confusing terms such as market adjustment.

In the theory of general equilibrium it was undoubtedly the imposition of high standards of mathematical formalism that led to the almost exclusive concern with the existence and characterization of equilibria rather than the adjustment process that lead to them. In the nineties there was an upsurge in interest to determining the adjustment process and how the economic agents might learn their way into equilibria. But these works have been focused on the question of learning about

the parameters of the model as captured by some (Walrasian or not) equilibrium model. They have concentrated on proving the existence of such an equilibrium and perhaps some evolutionary mechanism that will lead to the fixed point equilibrium, under the forceful assumption of rational expectations.

Real markets are information gathering tools that will allow many reasonable and realistic heuristic learning-decision models from agents yet achieving market convergence and equilibrium. Assuming that learning is primary focused towards finding the true parameters, that may describe the macroscopic market behavior as a "virtual" entity, is just totally unnecessary and even unrealistic. Why should the agents have a "virtual correct" macro model in mind and try to capture the parameters of the model? Their search may be local and no super agent such as the Walrasian auctioneer is needed for a market clearing-price as the evidence from Experimental Economics has proved.

The advantage to move from human to artificial agents (from Experimental to Artificial Economics) is that in our experiment we have control over the agents' behavior (Lopez-Paredes et al., 2002).

An extreme case was the Zero-Intelligent agents model of Gode and Sunder (1993) and Sunder (2004). It is a simple artificial model that does a reasonable job of capturing the dynamics of the competitive market and assures convergence and efficiency without any tâtonnement process. Of course such a challenging fact has been questioned recently by Brewer et al. (2002); Cliff and Bruten (1997); Gjerstad and Dickhaut (1998); Posada (2006). The main conclusions are that the institutional design matters and so does the agents' intelligence. That perfect competition is compatible with strategic agent's behaviour.

Our paper is not about solving an M or a W set of general equilibrium system of equations, along the lines of Colander (1995) although we grow a model that, yes, at the same time computes the solution. We do not intend to inquire about the microfoundations of top-down models e.g. differences between Keynesian and Walrasian economics, although again, this will be an interesting application of MABS in a similar vein to this paper.

In this paper we replicate and generalize with artificial agents the results of Plott and George (1992) and Brewer et al. (2002) experiments of M/W instability with humans and a forward falling supply function. They found that in a continuous double auction the M stability model captures the observed phenomena whereas the W model does not. We wonder as well if their results depend on the particular human agents' behavior because nothing is said about this issue.

The choice of the environment to introduce instability is essential to isolate the two sources of instability: either from the agents' behavior or from the environment. In environments where the supply is positively sloped and the demand is negatively sloped (as in Figure 1.a), there is only one equilibrium and it is both W and M stable (MS/WS). The controversy arises when a priori there are several points that should be M (W) stable (unstable). For example, in Figure 1.b, if the supply is negatively sloped and if the demand cuts the supply from above (below), the equilibrium is W unstable (MU) and M stable (WS) and Marshallian unstable (MU).

The forward falling supply is not an exceptionally abnormal supply, although this issue is of a secondary relevance for our arguments. It is the case, for instance, in information technologies products where marginal costs are practically zero and externalities and learning are present.




Fig. 7.1. Environments: (a) standard, (b) forward falling supply

7.3 The model

The main features of our model are described using the triple (IxExA):

a) *The institution, I.* Under the continuous double auction (CDA) rules, any trader can send (or accept) an order at any time during the trading period. A new bid/ask has to improve previous pre-existing bid/ask. A trade occurs when a new ask is made that is less than a pre-existing bid, or when a new bid is made that is greater than a pre-existing ask. The trading is equal to that of the pre-existing bid/ask, whose acceptance is triggered automatically by the new entry.

b) *The environment, E.* We use a forward falling supply (Figure 1b). We have used the same valuations which were used in Plott and George (1992)'s experiments with human agents.

In the demand side, there are six buyers (each one with six units). There are two buyers of each type. Each buyer of a given type has identical reserve price given by [960, 600, 440, 350, 330, 0], [880, 640, 410, 390, 310, 0] and [800, 720, 410, 390, 290, 0], respectively. Buyers know their reserve price with certainty.

In the supply side, there are six sellers. Each seller is uncertain about his marginal costs because they depend on their own output and the output of all other sellers. The externality implies that, as market volume increases, the marginal cost decreases even though the individual seller's marginal cost increases with an increase in his own volume. Each seller has eight units to trade. There are two sellers of each type (a, b, c). In Table 1 we show the marginal costs of type-a sellers. Each seller of a given type has identical marginal cost. The marginal cost of the first unit when the volume of others is 0, is 820 (for type-b is 800 and for type-c is 780). As unit increases, the marginal cost value increases by 80. Note that as the volume of others increases, the marginal cost value decreases by 30 per unit. The exceptions are every fifth unit starting at 5 (at 3 for type-b and at 1 for type-c) at which point the increment is 80 as opposed to 30.

Table 7.1. Marginal costs of a type-a sellers

	Volume of others							
	0	1	2	3	4	5	6 ...	18
1st unit	820	790	760	730	700	620	590 ...	130
2nd unit	900	870	840	810	780	700	670 ...	210
3rd unit	980	950	920	890	860	780	750 ...	290
	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
8th unit	1380	1350	1320	1290	1260	1180	1150 ...	690

Table 7.2 lists the equilibria according to both Walrasian and Marshallian theories.

c) *Agent's behavior, A.* The first decision which is taken by each seller is to estimate the volume sold in the market by others. Sellers are uncertain about their marginal costs, because they depend on their output and the output of all the other sellers. Sellers form expectations $\hat{q}_{n+1}$ on the volume using their own past experience. They update their volume expectations according to the actual volume observed in the

Table 7.2. Equilibria according to both Walrasian and Marshallian theories

Price	Quantity	Marshall	Walras
500-540	12	stable	unstable
380-410	18	unstable	stable
140-180	30	stable	unstable

market (q_n). In particular, any seller uses the following simple updating rule:

$$\hat{q}_{n+1} = (1 - \lambda)\hat{q}_n + \lambda q_n.$$

The learning rate λ measures the responsiveness of sellers' volume estimates to new data (a memory weighting factor). If $\lambda = 1$ the agents believe that the quantity sold by others in the next period will be equal to the traded volumes observed in the current period, and they will under-estimate their marginal cost. If $\lambda = 0$ the agents do not use the information generated in the market to improve the estimation of the initial traded volume (myopic behaviour).

Buyers do not need to estimate their reserve prices because they know them with certainty.

In the next step, agents (sellers and buyers) face the following three decisions: How much should they bid or ask? When should they submit an order? When should they accept an outstanding order? To take these decisions each agent only knows his own valuations, which are private, and the information generated in the market.

How much should he offer? We try two alternative bidding strategies: ZI (Gode and Sunder [6]) and GD (Gjerstad and Dickhaut [4]). Each bidding strategy has different answers to this question. Each ZI agent chooses his order randomly between his private valuations and the best order outstanding in the market. Each GD agent chooses the order that maximizes his expected surplus, defined as the product of the gain from trade (price minus private valuation) and the probability $\prod_a$ for an order to be accepted:

$$\max_a \prod_a (price - MaC), \quad (7.1)$$

where GD agents estimate the probability $\prod_a$ learning to modify their beliefs using the history of the recent market activity. GD sellers calculate a belief value $q(a)$ for an order a using *AAG* (accepted asks greater than a), *BG* (bids greater than a) and *RAL* (rejected asks less than a). Interpolation is used for values at which no orders are registered.

$$q(a) = \frac{AAG(a) + BG(a)}{AAG(a) + BG(a) + RAL(a)}. \quad (7.2)$$

When should he submit an order? When an agent is active, he may submit an order (a new or a replacement an open order). The agents have a constant activation probability of 25%. Of course, orders must be also in agreement with spread reduction rule of the institution.

When should he accept an outstanding order? A seller accepts the current bid if his ask (submitted or not) is equal to or greater than the current bid. A buyer accepts the current ask if his bid (submitted or not) is equal to or less than the current ask.

7.4 Some results

Note that the agents have two basic learning tasks: learning about the externality as represented by the total amount to be traded, and learning how to bid/ask. How is the performance of humans, ZI and GD agents in these two tasks? Following we discuss the meaning of the results of prices and volume.

Price volatility. Figure 2 shows the price distributions of the transactions in ten trading periods (each one with 100 rounds). Here we can see that for homogeneous ZI populations, the transaction prices are more volatile than for homogeneous GD populations, where the price convergence is very clear. The learning rate plays an important role on the convergence to the Marshallian stable equilibrium. All transactions in GD populations are much closer to the Marshallian stable equilibrium (540-500) when their λ learning rate is 0.9 than when their λ learning rate is 0.1 (myopic behaviour).

Price convergence. Although the transaction prices in homogeneous ZI populations are widely distributed, there is a false convergence appearance. In Figure 3 we show the time series of price transactions. We observe that for ZI homogeneous populations, there is no price convergence to the equilibrium price.

Figure 4 shows the time series of transaction with a GD population. We have represented with a discontinue line the price range (540-500) which is Marshallian stable and Walrasian unstable equilibrium. Here we observe that the transaction prices remain slightly over 540 when $\lambda = 0.1$, with a positive bias, as it happens with the experiment with humans. The transaction prices are unbiased and they cluster around the theoretical equilibrium of (540-500) when $\lambda = 0.9$.




Fig. 7.2. Pricing distribution of both ZI and GD homogeneous populations under extreme learning rates ($\lambda = 0.1$ and $\lambda = 0.9$)

Do humans make mistakes when they estimate the volume sold by others as GD sellers with a low learning rate make? In CDA markets, efficiency is achieved even by ZI agents. This is a consequence of the robustness of the CDA institution against learning. It was not surprising to find in our experiment a market volume of 12 which is Marshallian stable and Walrasian unstable, in every period and for any population.

In GD populations the bias comes from a poor estimation of the externality effect, as measured by the traded volume. The ZI agents make their offers randomly, but they estimate well the traded volume. If we assume that the human agents are more intelligent than the ZI agents, we can conclude that they also achieve the right estimates of the traded volume. We may then conjecture that for humans the source of the bias towards the M equilibrium comes from the offering decisions. This example shows that Experimental Economics and MABS can be used for testing economic aggregated models. But MABS allows us also to calibrate to what extent the results from Experimental Economics




Fig. 7.3. Price dynamics for Homogeneous ZI populations




Fig. 7.4. Price dynamics for Homogeneous GD populations

are robust against the unknown behaviour of the participants in the experiment.

7.5 Conclusions

MABS is a "killer" application in Economics. We can validate and calibrate the results from Experimental Economics with humans that have enlightened Economics for the last two decades, since we can control for agents' behavior. If we go down to micro behavior we can grow a market model for some fixed set of $A \times I \times E$.

Grown markets, as the real ones, are information gathering tools that will allow many reasonable and realistic heuristic learning-decision models from agents yet achieving market convergence and equilibrium. The agents search may be local and no super agent such as the Walrasian auctioneer is needed for a market clearing-price as our experiment shows. This focus opens the way to a new strand of research into microfoundations of aggregated economic models.

Growing Agent-Based Models in Economics, we could dispense of some of the concepts used to describe microeconomic equilibrium since we can trace the full process towards equilibrium and can calibrate it for different arrangements of the $A \times I \times E$. Nevertheless aggregate models are, no doubt, useful for policy design. However, to assume that the behavior that is true for the agents is also true for the aggregated system, may be wrong and should always be checked.

Our simulation results show that the theoretical predictions of the Walrasian market adjustment are wrong and that the agents behavior is compatible with the Marshallian model for the chosen environment (E), confirming previous results in Experimental Economics. The results are not robust against alternative agents learning models. M and W equilibrium adjustment is an analogy taken from vector fields that is unnecessary if not confusing. We have the tools (MABS) to model a market fully, allowing individual and social learning, and richer equilibrium concepts.

In view of the results, the following comment by Axtell (2005) is very appropriate: *"In the end we advocate not the jettisoning of this useful abstraction (Walrasian equilibrium) but merely its circumspect use whenever focused on questions for which it has limited ability to adjudicate an appropriate answer: distributional issues and actual prices. a direct consequence of the results described above is to at least cast a pale on the utility of such analysis, if not vitiate them altogether."*

Acknowledgement. This work has received financial support from the Spanish MEC, 2005-05676, and from Junta de Castilla y León, VA029B06. We acknowledge the reviewers comments that helped us to clear the object and the scope of the paper.

References

- R. Axtell. The complexity of exchange. *Economic Journal*, 115(504): 193–210, June 2005.
- P. Brewer, M. Huang, B. Nelson, and C. R. Plott. On the behavioral foundations of the law of supply and demand: Human convergence and robot randomness. *Experimental Economics*, 5(3):179–208, December 2002.
- D. Cliff and J. Bruten. Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. Technical Report HP-1997-141, Hewlett Packard Laboratories, Bristol, England, 1997.
- D. Colander. Marshallian general equilibrium analysis. *Eastern Economic Journal*, 21(3):281–293, Summer 1995.
- S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.
- D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, February 1993.
- A. Lopez-Paredes, C. Hernandez-Iglesias, and J. Gutierrez Pajares. Towards a new experimental socio-economics: Complex behaviour in bargaining. *The Journal of Socio-Economics*, 31(4):423–429, 2002.
- W. Nicholson. *Microeconomic Theory*. Dryden Press, 1997.
- C. R. Plott and G. George. Marshallian vs. walrasian stability in an experimental market. *Economic Journal*, 102(412):437–460, May 1992.
- M. Posada. Strategic software agents in continuous double auction under dynamic environments. In Berlin Heidelberg New York, editor, *Lecture Notes in Computer Science*, volume 4224, pages 1223–1233. Springer-Verlag, 2006.
- V. L. Smith. Theory, experiment and economics. *Journal of Economic Perspectives*, 3(1):151–169, Winter 1989.
- S. Sunder. Markets as artifacts. In M. Augier and J. March, editors, *Models of a Man: essays in memory of Herbert A. Simon*. MIT Press, 2004.

The Performance of Option–Trading Software Agents: Initial Results

Omar Baqueiro, Wiebe Van der Hoek, and Peter McBurney

Department of Computer Science, University of Liverpool, UK {omar, wiebe, peter}@csc.liv.ac.uk

8.1 Introduction

The growth of e-commerce and the development of distributed processing systems have led to interest among computer scientists in methods for resource allocations across multiple participants (Chevalayre and Dunne, 2005). GRID systems, for example, allow multiple users access to some resource, such as computer processing power or use of an electron microscope (Foster and Kesselman, 1999).

If resources are limited, each agent in a GRID system or other online marketplace faces the possibility of not being able to obtain resources when needed. If resources are allocated according to a market mechanism (either with real-world money or with tokens), agents also face the possibility of not being able to afford to purchase resources, even when they are available. As computational resource allocation systems become increasingly common, participants will require agents able to reserve future resources on their behalf, and hedge against future risks.

Derivatives are financial products whose values depend on the value of some other asset, usually a physical product. Option derivatives provide traders with the right to purchase or the right to sell the underlying assets at agreed future times, under agreed conditions. In this way, traders attempt to hedge against falls in the price of the underlying asset or to gain from price rises, and so manage the risks associated with the uncertainty of asset prices.

Elsewhere, a multi-agent framework has been presented in which BDI-type agents could be vested with decision rules allowing them to trade some product (Espinosa et al., 2005). We use this framework to create agents with similar decision rules for trading of option derivatives, and then undertake a Monte Carlo simulation to compare the

marketplace performance of agents trading options with those which do not.

Our contribution comprises the results of the Monte Carlo simulation for which the paper concludes with a discussion of the work. It is important to stress that our focus throughout is not on the exchange mechanism by which agents trade options, or its properties; rather, our concern is with the relative benefits or disbenefits to agents undertaking Options trading.

8.2 Model description

We created a multi agent market framework based on the model of Palmer and Arthur (1994). In our model we consider *goods* instead of *stock*, the goods cannot be divided, we only consider one type of good or asset and the price of the asset is fixed from a external price series. In addition to the standard asset trading mechanism, our model provides means to exchange Option contracts among the agents. We make use of the basic properties of real financial Option contracts to define the Options that agents can trade. Price series of the underlying asset is set from an exogenous discrete time series and Option prices are calculated at each step using the Black and Scholes (1973) model for Option pricing.

8.2.1 The market

The market is composed of a set of agents $A = \{1, 2, 3, \dots, N\}$. A_i is composed of two subsets of agents, agents that can trade options and goods A_o and agents that can only trade goods (assets) A_g .

We consider discrete time points $t = \{0, 1, 2, 3, \dots, T\}$ and refer to a period of time as the t th period (or step t) $[t, t + 1]$. The market has also a risk free rate of return r .

At each t each agent i has a number of goods $g_i(t)$ and an amount of cash $c_i(t)$. The total number of goods in the model is fixed, being $\sum_i g_i(t) = G$ for all t . Each agent also has an Option portfolio $\mathcal{O}_i = \mathcal{O}_i^w \cup \mathcal{O}_i^h$ which is composed by the Options the agent holds ($\mathcal{O}^h$) and the ones it wrote ($\mathcal{O}^w$).

An Option α is defined as:

$$\alpha = \langle X^\alpha, t^\alpha, v^\alpha, \tau^\alpha \rangle \quad (8.1)$$

Where X^α is the exercise price of that Option (the price agreed to pay for each good); t^α is the expiration time; v^α is the volume (the quantity

of goods to trade with that Option) and τ^α is the type of Option (call or put). Each Option α has a corresponding premium price $p_\alpha(t)$. This is the price an agent will have to pay its counter-party to hold the Option.

The Options provided by the market are a set of standard *templates* for Option contracts that the agents can trade. Agents are only allowed to exchange Options that comply with the specified templates, this is similar to a real Option regulated market. The number of available Option templates is constant over all time steps.

Pricing mechanism

The asset price $p(t)$ will be provided to the model from an external time series. Option pricing is calculated each step using the Black-Scholes model for option pricing defined in Black and Scholes (1973)¹. Using this model, the price of an Option is calculated from the price of the good $p(t)$, the variance of the asset price (σ) and a predefined exercise price X^α . The exercise price of an Option is obtained by the following formula:

$$X^\alpha = p(t) \times (1 + k)$$

Where k is a uniformly distributed pseudo-random number within the range $[-SP_k, SP_k]$.

Market timeline

Each period of time starts when the market *publishes* the new price for the asset. After obtaining this price the *Option clearing* phase will run where the market will receive instructions from the agents to exercise any Option that expires at this time. The agents holding any expiring Option must either decide to exercise or lose the Option at this time. Any non exercised Option should be removed from agent's held Options set O_i^h . Any request to exercise an already expired Option will be ignored by the market. In the event of an Option being exercised, the agents will clear the Option, trading the corresponding asset immediately.

¹ It is worth noting that other option pricing mechanisms could have been used, in fact some experiments were also ran using the binomial option pricing model by Cox et al. (1979) without any relevant difference in the outcomes.

Afterwards, the market will publish the different Option templates to trade on that period and the *trading phase* will start where the agents will submit their offers to buy and sell assets or hold and write Options. Next, the market will try to match randomly the asset buy and sell offers and will also try to match the hold and write offers. Finally, the market will clear the matched offers by making the agents exchange the assets or confirming the matched Option contracts. A graphical representation of the time-line is shown in Figure 8.1.




Fig. 8.1. Timeline for one time step of the market

8.2.2 Trading Agents

An agent i is defined by the tuple:

$$\langle g_i, c_i, w_i, \mathcal{O}_i, \mathcal{S}_i, \mathcal{F}_i \rangle \quad (8.2)$$

At time t , the term $g_i(t)$ is the number of goods the agent owns at time t ; $c_i(t)$ denotes the quantity of cash the agent has. The term $w(t)$ denotes the *wealth* of the agent which is obtained by the equation:

$$w(t) = p(t) \times g(t) + c(t) \quad (8.3)$$

The agent also owns a set of Option contracts $\mathcal{O}$ which represent a contract to buy or sell one asset at a specific time. The set of Options

is composed by two subsets, O^α is the set containing the held Options and O^β contains the Options written. Specifically $\mathcal{O} = O^\alpha \cup O^\beta$.

The term $\mathcal{S}$ is the agent’s strategy (See 8.2.4) comprising an action or chain of actions to execute. The set of actions an agent can execute are listed in Table 8.1. Finally, $\mathcal{F}_i$ is the forecast strategy used by the agent.

Action	Description
$buy(g, t)$	Make an offer to buy an asset at time t .
$sell(g, t)$	Make an offer to sell an asset at time t
$hold(\alpha, t)$	Make an offer to hold Option α at time t
$write(\alpha, t)$	Make an offer to write an Option at time t

Table 8.1. Available actions for the agents at time t .

8.2.3 Forecasting and perceived risk

The forecasting process of the agent is comprised by two parts, firstly the agent obtains forecasted price of the asset for future time steps and secondly it uses these forecasts to obtain its perceived risk of executing the possible actions. At each time step agents calculate a forecasted price for future time steps. Agents obtain this price using a forecasting function. Although other types of time series forecasting formulae could be used, our model implements two forecasting mechanisms.

Simple moving average forecasting

The first forecasting mechanism is based on the Simple Moving Average (SMA). Prices at future times are obtained by first calculating the SMA for the interval $[t - n, t]$ as $p_{SMA}(t)$ and then the price at future time steps is obtained by extrapolating the price at current time using the formula:

$$p_i(t + m) = p(t) + m \times (p(t) - p_{SMA}(t)) \quad (8.4)$$

Where $p_i(t + m)$ is the agent’s forecasted price for time $t + m$ and $p(t)$ is the market price at time t .

α –perfect forecasting

Using the second forecasting mechanism called the α –perfect forecast the agent will obtain the future prices from the real time series with

some added random variability (noise). The forecasted price will be calculated as:

$$p_i(t + m) = p(t + m) \times (1 + r_\alpha) \quad (8.5)$$

Where $p_i(t + m)$ is the agent's forecasted price at $t + m$, $p(t + m)$ is the real market asset price at time $t + m$ and r_α is a uniformly distributed pseudo-random number within the range $[1 - \alpha, \alpha - 1]$, being α within the range of $[0, 1]$. Using this mechanism, it is possible for the agent to have complete knowledge of the future prices when α equals 0.

Perceived risk

Inspired by Holton (2004), we model risk as the probability that the agent loses wealth when it carries out a specific action. We assume prices are distributed Normally. Under this assumption each agent can calculate the probability of wealth loss $\rho(a)$ for each possible action a at each step in time t .

This is achieved by using the cumulative standard normal distribution to obtain the cumulative probability of the agent forecasted price being in the wrong direction, assuming that the distribution's mean is $p_i(t + m)$ (the price at the forecasted time step).

8.2.4 Trading strategies

There are two types of agents trading in the market, asset traders and Option traders; asset traders can only trade the underlying asset in the market whereas Option traders can trade assets and Option contracts.

Asset trading strategies

Asset traders trade in the market using one of two strategies: the Random trading strategy in which the agents select an action randomly and Speculator strategy in which agents select an action to buy or sell an asset according to their forecast of the price at the next step.

Option trading strategies

Option traders can trade using the *Minimize Risk* strategy in which agents create an action tree with all the possible combinations of actions for a specific number of time steps and select the path which yields the minimum combined risk. An agent that uses this strategy

will choose the sequence of actions from the action tree (a path) where the combination of the actions' risk loss factor ρ is the minimum from all possible combinations. Let a strategy S be defined by the sequence of actions $\langle a_1, a_2, \dots, a_n \rangle$ and also let $\rho(a_i)$ be the risk loss factor for doing some action, the combined risk loss $\rho_s(S)$ for such strategy is defined as:

$$\rho_s(S) = \prod_i^n \rho(a_i) \quad (8.6)$$

Option trading agents can also use the *Maximize wealth* strategy with which they select the next action after selecting the path which yields the maximum sum of wealth from an action tree. An agent that uses this strategy will choose the sequence of action from the action tree where the combination of each of the action's wealth difference is the maximum from all possible combinations. Let a strategy S be defined by the sequence of actions $\langle a_1, a_2, \dots, a_n \rangle$ and let $\Delta w(a_i)$ be the perceived wealth difference for doing an action (the wealth before executing the action subtracted from the wealth after executing the action), the combined wealth $\Delta w(S)$ is defined as:

$$\Delta w(S) = \sum_i^n \Delta w(a_i) \quad (8.7)$$

8.3 Experiments

Several experiments were run to compare the performance of agents under two different aspects. Firstly to test which of the strategies generated higher profits and secondly to compare the correlation between the agents' wealth and the price of the asset. Our hypothesis was that, the wealth of agents using Options would be lower than that of the ones trading only assets.

8.3.1 Environment setup

A simulation run for our model requires the specification of the parameters of Table 8.2 for the market setup and for the agents. The parameters are explained in Section 8.2 excepting O_s which is used to set the distance between the expiration time (t^α) of the Options generated at each step and $\sigma(0)$ which is the initial value for the standard deviation of the price series. These parameters were fixed for all the experiments.

Initial parameters for the market	
Parameter	Initial value
Simulation duration (T)	500
Number of available Option templates ($ O $)	3
Steps between available Options (O_s)	1
Strike Price multiplier (SP_k)	15
Risk free rate (r)	0.005
Initial price variance ($\sigma(0)$)	1
Initial parameters for agents	
Parameter	Initial value
Initial cash ($c_i(0)$)	1000
Initial goods ($g_i(0)$)	100

Table 8.2. Initial parameters for the experiments

For all the experiments we also populated the market with 4 sets of 20 agents. All agents within one set were initialized with the same parameters (including strategy and forecast function). Each set used one of the four defined strategies. All the experiment were done using each of the price series to be described.

Price series

To set the price of the underlying asset we used several price series in order to test the performance of the agents under different market conditions. We defined three categories for the price series: stock prices series, which were obtained from the closing prices of different stocks²; random prices, which are uniformly distributed pseudo-randomly generated series; and linear prices which are manually generated. Some statistical information for the price series is summarized in Table 8.3. The Dell, Microsoft, HP, and IBM price series were obtained from the stock prices of the corresponding companies; the RANDOM1 and RANDOM2 price are the pseudo-randomly generated; finally, the Increment price series was generated as a constantly increasing time series and the Decrement was generated as a constantly decreasing time series.

8.3.2 Experiments using SMA forecasting

For the experiments with the Simple Moving Average forecasting, all the agents were assigned this same forecasting function ($\mathcal{F}_i$) with the number of periods $t_{\mathcal{F}_i} = 10$. We ran 50 repetitions of each experiment

² Freely available online at <http://finance.yahoo.com/>

	N	Minimum	Maximum	Mean	Std. Deviation
DELL	500	16.15	30.63	25.78	2.44
HP	500	10.75	37.80	21.46	6.34
IBM	500	54.65	125.00	95.66	17.37
MICORSOFT	502	41.75	73.70	58.43	7.40
RANDOM1	500	1.20	200.00	103.18	58.30
RANDOM2	500	4.00	996.80	501.54	293.44
INCREMENT	500	10.00	510.00	260.00	144.77
DECREMENT	500	10.00	509.00	259.50	144.48

Table 8.3. Descriptive statistics of the used price series

and averaged the results. We also calculated the mean of the wealth for each set of agents to obtain the performance for each strategy.

Performance of strategies

We measured the performance of each strategy by obtaining the difference between the wealth of the agent at the last time step and the first time step, this resulted in the profits that each agent obtained for each simulation (see Table 8.4). In order to compare the profits among the

	Option Traders		Asset Traders		Mean
	MinRisk	MaxWealth	Speculator	Random	
DELL	210.890	-171.660	-74.060	34.830	960.000
IBM	-177.342	45.537	-9.782	141.587	-558.827
HP	-363.480	292.750	38.980	31.750	-1393.000
MICROSOFT	-3.892	-1.262	-39.422	44.578	817.003
RANDOM1	-259.283	265.357	-2384.593	2378.518	-11979.998
RANDOM2	2373.450	-946.150	-21119.950	19692.650	51073.250
INCREMENT	9922.925	-8430.975	8441.825	-9933.775	49899.975
DECREMENT	-3547.801	-2561.132	509.389	5599.545	-47310.719

Table 8.4. Average profit for each strategy with SMA forecasting.

agents we calculated the mean of the average profits (last column in Table 8.4) and subtracted it from the strategy profit. Figure 8.2 shows the relative profits for each strategy among the simulations; from this figure it can be seen that there is no clear advantage in the profits using any strategy.

Performance correlation with price

The second test we performed was an analysis correlation between the price series and the wealth of the agents. This test was conducted to




Fig. 8.2. Relative profit for each strategy with SMA forecasting

see whether the fluctuations on the price of the asset had less incidence in the wealth of an agent trading Options, the results³ on Table 8.5 suggest so, as the correlation between the wealth of the Option trading strategies is slightly less than of the asset trading strategies for three of the four stock market strategies.

	Option Traders		Asset Traders	
	MinRisk	MaxWealth	Speculator	Z.I.
DELL	0.994	0.974	0.998	0.999
HP	0.999	0.997	1.000	1.000
IBM	1.000	1.000	1.000	1.000
MICROSOFT	0.999	0.999	1.000	1.000
RANDOM1	0.999	1.000	0.970	0.985
RANDOM2	1.000	1.000	0.982	0.990
INCREMENT	1.000	1.000	1.000	1.000
DECREMENT	1.000	1.000	1.000	1.000

Table 8.5. Correlation between agents' wealth and price series with SMA.

³ All correlations were calculated as two tailed Pearson correlation significant to the 0.01 level.

8.3.3 Experiments using α -perfect forecasting

For the experiments with the α -Perfect forecasting function, all the agents were assigned this forecasting function ($\mathcal{F}_i$) with three different α values of 10, 20 and 40. We ran 50 repetitions of each experiment and averaged the results. We also calculated the mean of the wealth for each set of agents to obtain the performance for each strategy.

Strategies performance

As with the SMA experiments, the performance of each strategy was measured by obtaining the difference between the wealth of the agent at the last time step and the first time step, resulting in the profits that each agent obtained for each simulation. Figures 8.3 shows the resulting relative profit for each strategy with the different α values.




Fig. 8.3. Relative profit for each strategy with α -Perfect forecasting with different α values.

The wide difference in the performance of the Option trading strategies against the asset trading strategies suggests a clear advantage on the use of Options in the case of the α -Perfect forecasting.

Performance correlation with price

Finally, the correlation between the price series and the wealth of the agents was calculated for the α -Perfect experiments. The results of this are listed on Table 8.6, the lower correlation of the Option Trading strategies particularly appear to indicate that the use of Options decreases the influence of the price in the wealth of the agents trading them.

Correlation for $\alpha = 0.10$					Correlation for $\alpha = 0.20$				
	Option Traders		Asset Traders			Option Traders		Asset Traders	
	MinRisk	MaxWealth	Speculator	Random.		MinRisk	MaxWealth	Speculator	Random.
DELL	0.988	0.831	0.998	0.999	DELL	0.984	0.884	0.999	1.000
HP	0.999	0.993	1.000	1.000	HP	0.998	0.991	1.000	1.000
IBM	0.999	0.997	1.000	1.000	IBM	0.999	0.997	1.000	1.000
MICROSOFT	0.997	0.992	1.000	0.999	MICROSOFT	0.997	0.992	1.000	1.000
RANDOM1	0.999	0.718	0.841	1.000	RANDOM1	0.999	0.761	0.861	1.000
RANDOM2	0.999	0.718	0.841	1.000	RANDOM2	0.999	0.765	0.879	1.000
INCREMENT	1.000	1.000	1.000	1.000	INCREMENT	1.000	1.000	1.000	1.000
DECREMENT	1.000	1.000	1.000	1.000	DECREMENT	1.000	1.000	1.000	1.000

Correlation for $\alpha = 0.40$				
	Option Traders		Asset Traders	
	MinRisk	MaxWealth	Speculator	Random.
DELL	0.994	0.985	0.998	1.000
HP	0.997	0.992	1.000	1.000
IBM	1.000	0.999	1.000	1.000
MICROSOFT	0.998	0.998	1.000	1.000
RANDOM1	1.000	0.833	0.889	1.000
RANDOM2	0.999	0.833	0.904	1.000
INCREMENT	1.000	0.999	1.000	1.000
DECREMENT	1.000	1.000	1.000	1.000

Table 8.6. Correlation between agents’ wealth and price series with α -Perfect forecasting.

8.4 Conclusions

In this paper we demonstrated some of the results from the experiments performed in our proposed Option Market framework. The experiments so far show promising results. It is worth nothing that, although the differences in the results of the tests between Option traders and asset traders are low, we argue that the reason for this could be due to the simplicity of the market. Allowing the agents to trade more than one asset at each step in time and providing them with Options with a higher volume (more than one asset traded on each Option) might increase the differences among the agent’s performance. Also, it would be interesting to introduce the concept of *magnitude* of risk into the agents reasoning process.

Acknowledgement. This work has been partially funded by the Mexican Council of Science and Technology (Sponsorship No. 187564) and the Market Based Control project (EPSRC GR/T10664/01). The authors also wish to thank the anonymous reviewers for their very insightful comments and suggestions.

References

- F. Black and M.S. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–54, 1973.
- Y. Chevaleyre and P.E. Dunne. Multiagent resource allocation. *Knowledge Engineering Review*, 20(2):143–149, 2005.
- J.C. Cox, S.A. Ross, and M. Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263, 1979.
- B. Espinosa, O.W. van der Hoek, and P. McBurney. Designing agents for derivatives markets. a preliminary framework. In P. Gmytrasiewicz and S. Parsons, editors, *Game Theoretic and Decision Theoretic Agents*, pages 1–14, 2005.
- I. Foster and C. Kesselman. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan-Kaufmann, 1999.
- G.A. Holton. Defining risk. *Financial Analysts Journal*, 60:19–25, 2004.
- R.G. Palmer and W.B. Arthur. Artificial economic life: a simple model of a stockmarket. *Physica D*, 75:264–274, 1994.

Studies on the Impact of the Option Market on the Underlying Stock Market

Sabrina Ecca, Mario Locci, and Michele Marchesi

DIEE, University of Cagliari, Italy

{sabrina.ecca, mario.locci, michele}@diee.unica.it

9.1 Introduction

In the past thirty years, options have become an important financial instrument, and now they account for a substantial percentage of total trading activity. From a research perspective, a lot of research have been carried out about the theoretical computation of option prices, starting from the seminal works of Black and Scholes (1973) and Merton (1973). Several researchers also examined the issue of to which extent options interact with their underlying stocks, and in particular their possible effects on stock returns and volatility, and on the overall quality of the underlying security market.

Some studies claim that option trading may have a positive impact on the underlying asset market, reporting a decrease in volatility after the introduction of option trading. Among them, we may quote Nathan Associates (1974), perhaps the first to study the impact of listing options on the Chicago Board of Exchange. They reported that the introduction of options seemed to have helped stabilizing trading in the underlying stocks. Ross (1976) and Hakansson (1982) affirm that the options introduction improve incomplete asset markets by expanding the opportunity set facing investors, and reduce the volatility of the underlying stock. Kumar et al. (1998) claim that option listings have a beneficial impact on the stock market quality in terms of higher liquidity and greater pricing efficiency.

Other researchers affirm, on the other hand, that option trading causes an increase in volatility. because it favors large positions and increases the bid-ask spread. For instance, Wei et al. (1997) report an increase in volatility of options on OTC stocks in the USA.

A third opinion among researchers claims that option trading has no significant impact on price volatility of the underlying stock market. Among these, Bollen (1998) affirms that option introduction does not significantly affect stock return variance, while Kabir (2000) examined option listing in the Netherlands, studying the impact of option trading on the underlying market. He finds a significant decline in stock price, but no significant effect on the volatility.

In the past ten years, a significant new stream of research works introduced modeling and simulation using heterogeneous, boundedly-rational interacting agents as a new tool for studying financial markets (see LeBaron, 2006, for a recent survey). This new approach, while it is still debated and challenged, especially among classical economists relying on the “efficient market hypothesis”, is able to give new insight into how markets work. For instance, it is able to explain the so called “stylized facts”¹ shown by virtually every market price series, using endogenous mechanisms. Very many papers appeared proposing different models based on heterogeneous agents, and studying many different aspects of financial market trading. However, to our knowledge, no one has yet tried to study the effect of option trading using this approach.

This paper uses the heterogeneous agents, simulation approach to study the interaction between a stock option market and the underlying stock market. We analyze the effects of realistic option trading strategies on the stylized facts of financial time series, the long wealth distribution of traders and the price volatility. We consider three basic kinds of traders: traders who trade only in the stock market, traders who trade in the stock market, covering their positions using the option market, and a central Bank which issues option contracts in the option market, and trades in the stock market to cover these contracts, upon their expiration. There are four types of trading strategies in the stock market: random, fundamentalist, momentum and contrarian trading. Each trader consistently applies just one strategy, and cannot change it.

A given percentage of traders – spanning over all kinds of strategies – use options to cover their positions. Each time one of these traders places an order in the stock market, she also cover herself buying option contracts from the Bank, or uses a strategy based on the combination of call and put options, like a “Straddle”.

¹ The main “stylized facts” are: (i) unit root property of asset prices; (ii) power-law distribution of returns at weekly, daily and higher frequencies; (iii) volatility clustering of prices.

9.2 Method and model

In our model there is a market in which N agents trade a single stock, which pays no dividends, in exchange for cash, with no transaction costs. Each trader is modeled as an autonomous agent and is characterized by a wealth, constituted by the sum of her cash and stocks, valued at the current price. Traders' initial endowment – both in cash and stocks – is obtained by dividing agents into groups of 20 traders, and applying Zipf's law to each group, so that the difference in wealth among the richest and poorest traders at the start of the simulation is about twenty-fold.

The agents are divided into sub-populations that adopt different trading strategies. Besides the stock itself, which is traded in the stock market, there is an *European* option contract on the stock. A fixed percentage of traders is also enabled to buy and exercise options. We call them *option traders*.

Another, special type of trader is the *Bank*; only one Bank is present in the market. The Bank issues option contracts and, upon their expiration, guarantees their exercise.

At each simulation step, which roughly corresponds to a day of trading, each trader can place a buy or sell limit order to the stock market. This happens with a probability of 10%, so each trader is active on average every 10 time steps. The pricing mechanism of the stock market is based on the intersection of the demand-supply curve (Raberto et al., 2003).

At each time step, option traders may also buy from the Bank one or more European option contracts, in order to hedge their investment. These traders have a *long position* in the option market. Since we deal only with European options, their owners are allowed to exercise their rights only at the expiration date. The pricing of options is based on Black-Scholes formula.

Upon expiration, options can be classified as being *in the money* (ITM), *at the money* (ATM) and *out of the money* (OTM). A call (put) option is ITM if the strike price is less (greater) than the current market price of the stock, so it is profitable to exercise the option. On the other hand, OTM options are not exercised because they are not profitable, resulting in a net loss of the traders who bought them. A call or put option is ATM if the strike price is exactly equal to the current market price, making irrelevant to exercise or not the option. In practice ATM options are not exercised, and are equivalent to OTM ones.

Traders owning ITM options exercise them, asking the Bank to sell them, or to buy from them, the corresponding stocks at the strike price. If the total number of stocks sold to these traders is not equal to the total number of stocks bought from them, the Bank places on the stock market a market order to cover the imbalance.

The Bank sets all components of the option contract: strike price (which depends on the current price of the stock $p(t)$, expiration date, underlying quantity, premium (Hull, 2002). The computation of the price of the options is made using the formula first introduced by Black and Scholes (1973).

9.2.1 Trading strategies

Stock market traders play the market according to four different kinds of strategies, that roughly mimic traders' behaviors in real markets. These strategies are described in depth in Raberto et al. (2003), and are summarized below. Some strategies require a time window to compute some significant parameters. In this case, each trader has a specific time window whose length is an integer randomly extracted from a uniform distribution, in the interval $2 - 10$.

Random traders: Random traders are characterized by the simplest trading strategy, representing the “bulk” of traders who do not try to beat the market, but trade for exogenous reasons linked to their needs. They are traders with zero intelligence, issuing random orders. If a random trader decides to issue an order, this may be a buy or sell limit order with equal probability. The order amount is computed at random with uniform probability, but cannot exceed the trader's cash and stock availability. The limit price is set at random too, in such a way to increase the probability the order is satisfied.

Fundamentalist traders: Fundamentalist traders believe that stocks have a fundamental price due to factors external to the market, and that, in the long run, the price of the stock will revert to this fundamental price, p_f . Consequently, they sell stocks if the price $p(t)$ is higher than fundamental price and buy stocks in the opposite case. The fundamental price is the same for all fundamentalists and corresponds to the “equilibrium” price, when the total cash owned by all traders, C_{tot} is equal to the value of all the stocks owned by all traders, S_{tot} . The order amount is proportional to the distance between the current price and p_f . The limit price is equal to p_f , or

to the current price $p(t)$ plus or minus 20%, whichever is closer to the current price.

Momentum traders: Momentum traders speculate that, if prices are rising, they will keep rising, and if prices are falling, they will keep falling. Their orders are buy orders if the past price trend is positive, and sell orders if the trend is negative. The order amount is computed at random in the same way as random traders, while the limit price is set by extrapolating the price trend.

Contrarian traders: Contrarian traders speculate that, if the stock price is rising, it will stop rising and will decrease, so it is better to sell near the maximum, and vice-versa. So, their orders are sell orders if the past price trend is positive, and buy orders if the trend is negative. The order amounts are computed in the same way as random traders, while limit prices are set by reversing the trend, using as pivot the current price.

9.2.2 The Bank

The Bank is a special trader with infinite wealth, able to issue and sell call and put European options to other traders. The components of an option contract are:

Expiration date: it is fixed on the third Friday of the month. In our model, all months are nominally 20 working days long, thus the expiration dates are days 15, 35, 55, ..., $20k + 15$, We use realistic expiration dates, that depend whether the option is bought before or after the third Friday of the current month (see Hull, 2002). In the former case, the expiration month can be the current month, or the month whose index is equal to the current one, plus 1, 3 or 6. In the latter case, the expiration month can be the month whose index is equal to the current one, plus 1, 2, 3 or 6.

Premium: the premium to be paid for an option is computed using the Black and Scholes formula (Black and Scholes, 1973; Hull, 2002). This formula uses five parameters: the stock price $p(t)$ at the time the option is valued, the strike price X , the time to expiration ΔT , the price volatility, computed in a given time window whose length is in our case 50 time steps, and the short-term interest rate, which in our case is set to zero. The basic idea underlying Black and Scholes formula is that the prices of the stocks follow a random walk, implying that the underlying asset prices are lognormally distributed with a constant mean and standard deviation. In our

artificial stock market model, however, the price process is characterized by a strong mean reverting behavior toward a price, p_f , equal to the ratio between the total number of stocks and the total cash owned by traders (Raberto et al., 2001), due to the finiteness of resources of the traders. This leads to overpricing the options, causing steady losses to option traders. For this reason, the option premium computed using Black and Scholes formula is multiplied by a correction factor C that depends on time to expiration ΔT and typically varies between 0.75 (in the case $\Delta T = 120$) and 0.96 (in the case $\Delta T = 10$). These values have been empirically computed through many simulations. Using the correction factor $C(\Delta T)$, we were able to use an option premium that is fair with respect to our finite resources, mean reverting market model.

Strike price: it is the price X at which the option can be exercised. It depends on the current price $p(t)$ of the stock. We consider three different possible strike prices, given by eq. 9.1.

$$X \in \{p(t) - \delta, p(t), p(t) + \delta\} \quad (9.1)$$

The value of δ depends in turn on $p(t)$. In US Dollar-quoted markets, δ is given by the following formula (Hull, 2002):

$$\delta = \begin{cases} 1.5\$ & \text{if } p(t) \leq 25\$ \\ 3\$ & \text{if } 25\$ < p(t) \leq 200\$ \\ 6\$ & \text{if } p(t) > 200\$ \end{cases} \quad (9.2)$$

For instance, if $p(t) = 42.7\$$, then $\delta = 3$, and the possible three strike prices are $X = 39.7\$$, $X = 42.7\$$, or $X = 45.7\$$.

When the Bank sells an option contract, it earns the premium, updating its cash. On expiration dates, that is every 20 simulation steps, if option traders have expiring ITM options, they ask the Bank to honor the contracts, selling them the stocks at the strike price in the case of call options, and buying from them the stocks at the strike price in the case of put options. If required, the Bank places a buy or a sell limit order on the market, at the market stock price, plus or minus a proper percentage (set to 2% in our model), to cover its position and be able to satisfy all its obligations. The Bank has unlimited wealth. In practice, it starts with a cash and a number of stocks set to zero, but these values are unbounded, and can assume any value, even negative.

9.2.3 Option traders

Option traders are those traders who are allowed to trade both in the option market and in the underlying stock market. As regards the stock

market, they exhibit one of the four possible trading strategies described in section 9.2.1. When they trade in the option market, they can only buy options and possibly exercise them on their expiration dates. Option traders can buy option contracts from the Bank only if their residual cash is higher than the premium of the option contract.

To be more specific, let us call $m_i(t)$ the cash owned by option trader i at simulation step t , and $s_i(t)$ the stocks owned by the same trader at the same step. Let us also suppose that, at step t , option trader i has p_i put option contracts not yet expired. These put options refer to quantity q_j^p , at a strike price of x_j^p , $j = 1, 2, \dots, p_i$ respectively. Conversely, let us suppose that at step t , option trader i has c_i call option contracts not yet expired. These call options refer to quantity q_j^c , at a strike price of x_j^c , $j = 1, 2, \dots, c_i$ respectively.

The total cash balance expected when all undersigned options are expired, m_B is estimated by eq. 9.3.

$$m_B = \sum_{j=1}^{p_i} q_j^p x_j^p - \sum_{j=1}^{c_i} q_j^c x_j^c \quad (9.3)$$

The total stock balance expected when all undersigned options are expired, s_B is estimated by eq. 9.4.

$$s_B = \sum_{j=1}^{c_i} q_j^c - \sum_{j=1}^{p_i} q_j^p \quad (9.4)$$

Note that in computing the balances we don't consider the options to be ITM or OTM with respect to the current price $p(t)$, but for the sake of simplicity we give all the options the chance to be ITM.

On the expiration date, if the option is ITM and if the trader holding it has enough money or stocks, she exercises it. Otherwise, she gets back the difference between the actual price and the strike price from the Bank.

If the option is OTM, the trader places on the stock market a buy limit order (if the option is a call), or a sell limit order (if the option is a put) at the current stock price for the underlying quantity. This quantity is reduced if the trader has no cash or stock enough to cover it completely.

Using options to cover a position

If options are used to cover a position, when option traders decide to place a buy or sell order, they also buy from the Bank a corresponding option to cover their position, provided they have cash enough to buy the option. The expiration date is always three months from the current month, so $\Delta T \approx 90$. If the order is a buy, they buy a put option, with a strike price X equal to the current price $p(t)$ minus δ as in eq. 9.2. If the order is a sell, they buy a call option with $X = p(t) + \delta$. In this way, option traders are guaranteed against losses exceeding δ , but have to pay the option premium, that is in any case subtracted from their cash.

Using straddles

If option traders use straddles, they simultaneously buy a put and call on the same underlying security, with the identical strike price and expiration month. The value of the strike price is the same of the current price of the underlying asset. So, both call and put options are ATM at the moment of purchase. Typically, the buyer of a straddle anticipates a substantial movement in the stock price, but is uncertain what direction it will be. Because the trader is betting on an large stock movement, the odds of losing are high. The buyer of a straddle risks only the amount of the premium. The maximum loss occurs when the price of the stock on the expiration date of the options is exactly equal to the strike price. In our model, in order to ease comparison with the case when option are bought to cover a position, option traders buy a straddle when they place an order on the stock market. The stock quantity of the straddle is the same of the stock market order, provided that the trader has cash enough to pay for the straddle premium.

9.3 Results and conclusion

In this section we describe the results of the computational experiments we performed. Each simulation was run with 5000 time steps and 400 agents. We varied the composition of the population performing various runs, and eventually decided to hold at 10% the percentage of fundamentalist, momentum and contrarian traders. Option traders can be 0%, 20% or 40%, equally divided in the four possible types. Random traders account for the remaining percentage.

The price volatility used in Black and Scholes formula is computed using a time window of 20 trading days. Also in the presence of option trading, our artificial stock market consistently exhibits realistic price series from a statistical point of view, showing the classical stylized facts, with fat tails of returns and volatility clustering.

During the simulation, when option traders decide to buy or sell stocks, they also buy options from the Bank. In doing this, these traders do not directly interfere with the stock market. The only indirect effect on the stock market is that they spend money to undersign options, so that in subsequent buy orders they can buy a smaller amount of stocks.

On expiration dates, on the contrary, the option traders interact with the underlying stock market. This may happen in three ways. The first is when the Bank needs to buy or to sell stocks to cover its cumulative position with respect to the owners of expiring ITM options. These stocks are bought or sold in the stock market, creating an unbalance.

The second way is when option traders have OTM options. In this case, they often buy or sell the stock, placing a buy or sell order of the amount of the option on the stock market.

The third, indirect way is that, by exercising the options, option traders change the composition of their portfolios, and this has an impact on their subsequent trading activity.

We divided the performed simulations in two main categories – option traders covering their positions, and option traders buying straddles. In principle, the effects of these two strategies could be quite different, because covering a position implies buying a single option contract at a time, while buying straddles is a more speculative strategy, and options are bought in pairs.

In both cases, we found that, despite the high percentage of option traders, the price series exhibit the “stylized facts” of real financial markets, and do not substantially differ from the case with no option trader.

9.3.1 Results when options cover a position

In the case option traders use options to cover their positions, as described in section 9.2.3, we performed many simulations, checking the behavior of trader wealth and of price volatility. Fig. 9.1 shows the wealth dynamics for a typical simulation.

Note that contrarian and fundamentalist traders, who use the “right” strategy for a mean-reverting price behavior, tend to increase their wealth, as already reported and discussed in Raberto et al. (2003).




Fig. 9.1. Dynamics of wealth of all kinds of traders for a typical simulation where option traders cover their positions using options. The total percentage of option traders is 40%.

On the other hand, the same kinds of traders using options to cover their positions tend to be much less profitable. This is because they spend money to buy options to cover positions that are unlikely to yield strong losses. The situation of momentum and random traders is completely different. These traders employ losing strategies, and in fact tend to lose wealth. Their option counterparts, however, tend to be much more profitable, because it is convenient to cover themselves with options, when the underlying strategy is bad.

9.3.2 Results when option traders use “straddles”

In another series of simulations, we considered a market where option traders use “straddles”, as defined in section 9.2.3. Fig. 9.2 shows the wealth dynamics for a typical simulation. Here the traders tends to gain less than in the other case, and only contrarian traders and fundamentalists – the latter both using and not using options – gain something. All other kinds of traders tend to lose money. All traders, but fundamentalists, who use straddles tend to lose money with respect to traders with the same strategy, not active in the option market. This is not unexpected, however, because in a limited resources, mean-reverting market, a strategy betting on high price variations, like the ‘straddle’, is unlikely to win.




Fig. 9.2. Dynamics of wealth of all kinds of traders for a typical simulation where option traders use “straddles”. The total percentage of option traders is 40%.

Table 9.1. Price volatility with and without option trading; each reported value refers to 20 simulation runs. The values are multiplied by 10^3 .

Strategy	Quantity	No option trader	20% option traders	40% option traders
Cover	mean	0.45	0.3	0.25
Cover	std. dev.	0.13	0.07	0.06
Straddle	mean	0.45	0.23	0.14
Straddle	st. dev.	0.13	0.07	0.03

In Table 9.1 we show how price volatility changes in the presence of option trading. In general, our simulations show a consistent, strong decrease in price volatility when options are traded. This despite the fact that once in every month the Bank places an order that might be very large, at a limit price able to cause strong price variations. When straddles are used, the number of traded options is doubled, and the volatility decreases even more. The presented figures refer to averaging volatility, computed every 50 time steps, on the whole simulation, and then averaging on 20 different simulation runs. Note that volatility does not show significant trends across a single simulation. These results seem to confirm the empirical findings that option trading stabilizes the

market and reduces the volatility (Hakansson, 1982; Nathan Associates, 1974; Ross, 1976).

Clearly, all the presented results are still preliminary, and more tests are needed to assess them. Future research directions we are working on include: (i) modeling dividends and interest rates; (ii) opening the market to external influences so that it is no longer mean-reverting, at least in the short run; (iii) studying other strategies using options; (iv) giving traders the ability to sell options.

Acknowledgement. This work was supported by RAMSES (Research on Agent-based Modeling and Simulation of Economic Systems) research project, sponsored by FIRB research fund of MIUR, grant number: RBAU01KZ7Z, by MIUR PRIN 2004 Project #2004131501, by ComplexMarkets E.U. STREP project #516446, under FP6-2003-NEST-PATH-1, and by EURACE E.U. STREP project #035086, under FP6-2005-IST-2005-2.3.4 -(xi) FET Proactive Initiative: Complex Systems.

References

- F. Black and M.S. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–54, 1973.
- N.P.B. Bollen. A note on the impact of options on stock return volatility. *Journal of Banking and Finance*, 22:1181–1191, 1998.
- N.H. Hakansson. Changes in the financial market: Welfare and price effects and the basic theorems of value conservation. *The Journal of Finance*, 37(4):977–1004, 1982.
- J.C. Hull. *Options, Futures, and Other Derivatives*. Prentice-Hall International, 5 edition, 2002.
- R. Kabir. The price and the volatility effects of stock option introductions: A reexamination. In I. Hasan and W.C. Hunter, editors, *Research in Banking and Finance*. Elsevier, 2000.
- R. Kumar, A. Sarin, and K. Shastri. The impact of options trading on the market quality of the underlying security: An empirical analysis. *The Journal of Finance*, 53(2), 1998.
- B. LeBaron. Agent-based computational finance. In K.L. Judd and L. Tesfatsion, editors, *Handbook of Computational Economics*. North-Holland, 2006.
- R. Merton. The theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4:141–183, 1973.
- Nathan Associates. Review of initial trading experience at the Chicago board options exchange, 1974.

- M. Raberto, S. Cincotti, S.M. Focardi, and M. Marchesi. Agent-based simulation of a financial market. *Physica A*, 299:320–328, 2001.
- M. Raberto, S. Cincotti, S.M. Focardi, and M. Marchesi. Traders' long-run wealth in an artificial financial market. *Computational Economics*, 22:255–272, 2003.
- S. Ross. Options and efficiency. *Quarterly Journal of Economics*, 90: 75–89, 1976.
- P. Wei, P.S. Poon, and S. Zee. The effect of option listing on bid-ask spreads, price volatility and trading activity of the underlying otc stocks. *Review of Quantitative Finance and Accounting*, pages 165–180, 1997.

On Rational Noise Trading and Market Impact

Florian Hauser

Innsbruck University School of Management, Department of Banking and Finance
florian.hauser@uibk.ac.at

10.1 Introduction

Since Black (1986) introduced noise as “expectations that need not follow rational rules”, noise traders are welcome in modelling financial markets as they provide liquidity and solve theoretical problems like the information-paradox formulated by Grossman and Stiglitz (1980). Unfortunately, those traders cannot expect to be honored for their contributions, or, as Black (1986) states, “if they expect to make profits from noise trading, they are incorrect”.

This paper presents a simulation model where traders can decide *endogenously* whether to rely on costly information or to act as a noise trader by adopting a random trading strategy. We show that, especially for traders with a low market impact, a random strategy may be the only rational choice as it promises higher returns than trading on incomplete information.

The paper is structured as follows: Sect. 10.2 presents the simulation model. To check the value of additional information in this model, Sect. 10.3 discusses an exogenously defined strategy allocation. Sects. 10.4 and 10.5 analyze equilibrium situations where every trader endogenously chooses his optimal strategy in terms of returns and market efficiency. Robustness checks are discussed in Sect. 10.6, and Sect. 10.7 concludes.

10.2 Simulation model

We simulate a one-period call market where one security is traded. The market consists of 1,023 computerized traders who can be classified according to their market impact. Category $T_{j \in \{1,2,\dots,10\}}$ consists of

2^{10-j} traders, each of them trading 2^{j-1} shares each period. With this categorization, the market impact of all traders in the market roughly follows Zipf's law (see Zipf, 1949). It results in one single trader in group T_{10} , transacting 512 shares per period, while each of the 512 traders in category T_1 trades only one security per period.

All traders are risk-neutral expected wealth maximizers and are free to act as random traders ($I = 0$) by deciding to buy or sell the security in period 0 with equal probability, or to use a fundamental trading strategy by adopting one of ten discrete information levels ($I \in \{1, 2, \dots, 10\}$). To keep our traders in the market even though some of them will suffer losses, it is assumed that all traders have an exogenous motivation to trade. However, they have to make a decision whether to base trades on a fundamental strategy or to behave as a noise trader.

To model the information system, we state that the intrinsic value V in period 1 is given by the product of 11 individual signals α_i , written as

$$V = \prod_{i=1}^{11} \exp(\alpha_i); \alpha \sim N(\mu = 0, \sigma = 0.05) . \quad (10.1)$$

Like in real markets we assume that what the poorly informed traders know should be known to the better informed traders as well. Therefore, a trader adopting information level I will receive the signal $\alpha_{i=I}$ as well as all signals $\alpha_{i < I}$. Assuming a risk-free interest rate of zero, this trader will predict the fair price in period 0 as

$$E(V) = \prod_{i=1}^I \exp(\alpha_i) + \varepsilon; \varepsilon \sim N(\mu = 0, \sigma = 0.00001) .^1 \quad (10.2)$$

As (10.1) describes a random process,² the expected relative prediction error $\varsigma = \ln(E(V)) - \ln(V)$ is normally distributed with a mean of 0 and a standard deviation of $0.05 \times \sqrt{(11 - I)}$ (see Table 10.1 for exact values). The cost of information levels is rising progressively with the quality of the information, given by

¹ Contrary to α , the error term ε is individually computed for every trader in each run. This creates a certain divergence of opinion among traders adopting the same information level.

² Note that the formula used to model the signals can also be used to model a random-walk process (geometric Brownian motion). Many dynamic models apply the same method by assuming that informed traders can predict the value of a security (following a random-walk process) in future periods.

$$C_I = (2^I - 1) \times c; \quad c = 0.02. \quad (10.3)$$

As can be seen in Table 10.1, adopting information level $I = 1$ is quite cheap. For a trader in class T_1 , the expenses will lower his absolute returns by 0.02. The cost function increases progressively, so adopting higher information levels will only be reasonable to traders in higher classes. As a trader in T_{10} trades 512 shares, adopting $I = 9$ will lower his absolute returns per share also only by 0.02.

Table 10.1. Relative prediction error ς and absolute costs C of information levels.

I	0	1	2	3	4	5	6	7	8	9	10
ς	0.69	0.16	0.15	0.14	0.13	0.12	0.11	0.10	0.09	0.07	0.05
C	0.00	0.02	0.06	0.14	0.30	0.62	1.26	2.54	5.10	10.22	20.46

After each trader has chosen his information level (we will explain how our traders learn to find their optimal strategies in Sect. 10.4), all informed traders place limit orders to buy the security if the market price $P < E(V)$ and to sell the security if $P > E(V)$. Random traders will decide to buy the security (placing a limit buy order for $P < 2$) or to sell for $P > 0.5$ with a 50% chance for both options. Orders are matched at the price allowing for the highest possible market volume. If this condition is met by a steady interval of prices, the market price is computed as the geometric mean of this interval. If the number of buyers and sellers differs, orders are only partially executed.³ As every trader class places limit orders for 512 shares, maximum market volume is 2,560 securities per period.

In period 1, all shares are liquidated at the intrinsic value V . According to this, relative gains/losses per share for buyers in class T_j (trading $s = 2^{j-1}$ shares) are calculated as

$$R_{T_j} = \ln\left(\frac{V \times s - C}{P \times s}\right). \quad (10.4)$$

As going short in the security implies that traders receive P in period 0 and pay V in period 1 to buy back the security, returns for sellers are given by

³ As a trader in T_{10} trades 512 securities with the same reservation price, the median of all reservation prices is likely to correspond to this trader's prediction. In this case, the number of buyers and sellers will differ.

$$R_{T_j} = \ln\left(\frac{P \times s - C}{V \times s}\right). \quad (10.5)$$

For all results presented in the following sections, settings with 1,000 simulation runs were computed. To derive stable results, all random numbers used in a setting were pre-calculated and recycled when the simulation had to be re-looped.⁴

10.3 Reference setting

To get first insights into the value of information in this model, a reference setting with exogenously defined strategies is analyzed. Suppose that all traders in class T_j choose information level $I = j - 1$, so traders in T_1 adopt a random strategy while all other traders rely on fundamental information. Average returns $\bar{R}$ for trader classes in 1,000 runs with and without costs (black bars and grey bars respectively) are shown in Figure 10.1.




Fig. 10.1. Average returns per trader class with/without costs (black/grey bars). The adopted information level I is plotted above the horizontal axis.

As can be seen, the random traders in T_1 receive an average return of zero although they cannot predict the fundamental value of the security

⁴ This is especially important for the calculation of equilibria, as an equilibrium that is stable with a certain set of random numbers might get unstable if random numbers are re-calculated.

at all. This is not surprising if we consider that all 512 traders in T_1 will randomly decide whether to go long or short. On average we find 50% of them on each side of the market, so their orders have hardly any impact on market prices. Hence, on average, we will also find 50% of them to have made the right decision.

If we look at returns before costs (grey bars), the market is a zero-sum game. We have seen that random traders will receive average returns of zero, as they make an independent, random decision. Looking at the results for fundamental traders, we find the well-informed trader classes $T_{[7;10]}$ are able to realize positive returns before costs while the less-informed traders in $T_{[2;6]}$ suffer losses. Remarkably, the least informed fundamental traders in T_2 are not the ones receiving the lowest returns. Before (and after) costs, this trader class performs clearly better than T_3 and T_4 .

The idea that additional information can worsen the investors' performance is already discussed in several studies (see e.g. Huber et al., 2007; Samuelson, 2004; Schredelseker, 2001). To explain this result with the given model, we have to consider two different effects. First, a higher information level will lower a trader's prediction error. However, as a second effect, additional information will also increase the chance of a trader to make joint mistakes with other trader classes relying on the same signals α_i . Traders in T_3 and T_4 suffer the most from this kind of herding behavior, as their price predictions are on the same time rather unprecise and highly correlated with predictions of other trader classes. For the least informed traders in T_2 , the second effect is less pronounced: as they receive only one signal, their price prediction is almost independent of that of the rest of the market.

10.4 Equilibrium

So far we have seen that a random strategy is likely to beat low-information strategies. This leaves us with the question why any trader should stick with a low-information strategy if he is better off as a random trader. Picking up this idea, we are interested in equilibrium situations, meaning strategy allocations where no trader has an incentive to change his strategy if all other traders stick with their strategies as well.

In a first step, we search for equilibria under the condition that all traders in one class choose the same strategy. To find equilibria, every trader class starts with an information level drawn from a discrete uniform distribution ($I_{T_j, \text{start}} \sim U(0, 10)$). For every single trader class,

the incentive to change strategy is calculated as $R_{T_j,\text{best}} - R_{T_j,\text{start}}$, with $R_{T_j,\text{best}}$ being the highest possible return for this class when all other classes stick with their original strategy $I_{T_j,\text{start}}$. Then, the trader class with the highest incentive to change adopts the best strategy and the simulation is re-looped. Equilibrium is reached when no trader class has any further incentive to change strategy.

The calculation of equilibrium is re-looped for several hundred times. For the setting presented here, we find only one stable equilibrium with a strategy allocation I_T shown in Table 10.2. One can see that in equilibrium, the 3 traders of T_9 and T_{10} are the only ones who process information. For all other trader classes it is rational to stick with a random strategy.

Table 10.2. Strategy allocation in equilibrium.

Class T_j	T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8	T_9	T_{10}
I_T	0	0	0	0	0	0	0	0	7	9

This method to find equilibria is then re-applied to several settings, having the same model parameters but different signals α_j and a new set of random numbers. Not in all settings do we find stable equilibria. If equilibria are found, strategy allocations turn out to be quite similar to the setting shown in Table 10.2. Most importantly, in all equilibria only traders in T_9 and T_{10} process information with the higher information level always being associated with T_{10} .

Average returns in equilibrium are shown in Fig. 10.2. As can be seen, only T_{10} is able to realize positive returns after costs while average returns before costs are positive for all 3 traders of T_9 and T_{10} . For all other traders, their random strategies now produce negative returns. Despite that, the equilibrium calculation shows that it is rational for them to stick with a random strategy as all other strategies will make their performance even worse. According to this, noise trading is no longer the result of irrational trading behavior, it may under certain circumstances be the only rational strategy for an investor.

For all trader classes adopting a random strategy, one can see that a higher market impact clearly worsens their performance. This is due to the fact that for a trader in T_8 , a random strategy implies that he has to decide whether to go long or short with all 128 shares. Hence, systematic behavior (e.g. all 4 traders in T_8 go long) is much more likely for this category than for T_1 , where 512 traders can decide independently whether to go long or short with one share. Because of this, market




Fig. 10.2. Average returns per trader class with/without costs (black/grey bars) for a strategy allocation as shown in Table 10.2. The adopted information level I is plotted above the horizontal axis.

impact in our model has a twofold effect on the strategy selection of traders: a higher market impact makes a fundamental strategy more profitable as the costs of information per share decline, and it makes a random strategy less profitable as the price impact increases.

Relating to capital markets, this questions if it is rational for private investors to buy index-driven products. As many other investors do the same, those investors choose to play a random strategy with a high market impact. Considering the results in our model, those products may perform worse than the market average, as also shown by Hanke and Schredelseker (2005).

As a second step, we allow for heterogenous choice of strategies in every trader class. Starting with the allocation in Table 10.2, we subsequently check for all trader classes $T_{j \in \{10, 9, \dots, 1\}}$ whether every single trader in T_j already adopts his optimal strategy. If not, the trader with the highest incentive to change in this class adopts his optimal strategy and we start again with T_{10} .

After a few loops, a strategy allocation is reached where not a single one of the 1,023 traders has any better option than the strategy he already uses. The exact strategy allocation as well as average returns after costs for $T_{[8;10]}$ are shown in Table 10.3, average returns for all trader classes in this situation are presented in Fig. 10.3.

As a first result, note that the strategy allocation does hardly change compared to the equilibrium with homogenous strategies. While one

Table 10.3. Strategy allocation and average returns in equilibrium with heterogeneous strategy choice in every trader class.

Class T_j	$j \leq 7$	T_8				T_9		T_{10}
Trader t	<i>all</i>	1	2	3	4	1	2	1
I_t	0	0	0	0	1	5	7	9
$\overline{R}$ in %		-1.1	-0.8	-0.8	-1.2	-1.8	-1.3	1.9

trader in T_8 leaves his random strategy to adopt $I = 1$, one of the two traders in T_9 now chooses a lower information level of $I = 5$. This leaves us with a situation where only 22.5% of all shares in the market are traded on information. Note that all fundamental traders in the market now use an individual strategy.




Fig. 10.3. Average returns per trader class with/without costs (black/grey bars) for a strategy allocation as shown in Table 10.3.

For the random traders, returns in the new equilibrium have increased markedly. As a higher share of stocks is traded on information now and the heterogeneity of information increased, the overall market impact of all random traders declined. Although in the new equilibrium still more than 75% of all shares are traded upon a random strategy, the losses of the random traders are quite moderate. Random traders in T_8 (the ones with the highest market impact) on average perform less than 1% worse than the market average, and they realize higher returns than the fundamental traders in T_8 and T_9 .

While returns for T_{10} did not change significantly, traders in T_9 are now the ones suffering the highest losses. Note that in the new equilibrium, they are the ones relying on the medium information levels. As shown in the reference setting (see Sect. 10.3), this strategy is fraught with problems: the information processed by T_9 is not good enough to compete with T_{10} . As the signals processed by both traders in T_9 are correlated, they will often make joint mistakes against T_{10} . The reason for them to lose against traders in $T_{[1;8]}$ is rooted in their market impact: contrary to the random traders, their correlated predictions will let them act as price-makers, driving market prices away from fundamental values and hence leading to higher negative returns. Nevertheless, the strategy chosen by traders in T_9 is their best available option: their market impact is too low to adopt very high information levels like $I = 9$, and it is too high to adopt a random- or low-information strategy.

Related to capital markets, the results may explain why Malkiel (2005) finds that actively managed investment funds on average perform worse than their reference indices. The relatively good (and costly) information processed by fund managers in combination with their market impact will work against them and can easily drive their returns below market average.

10.5 Market efficiency

While analyzing several different settings (see Sect. 10.6), we also find equilibria in heterogeneous strategies where all traders realize negative returns after costs. In those situations, the simulation shows a market that is truly efficient in the sense of Jensen (1978): with respect to the information system and the cost of information, no trader in this market is able to realize systematic, above-average gains. But, according to Fama (1970), an efficient market should reflect all available information. Concerning the low number of traders that process information in equilibrium (see Section 10.4), one might ask how good this market performs in predicting fundamental values. We therefore use raw returns, calculated as $r = \ln(V) - \ln(P)$, to measure the relative deviation of prices from intrinsic values.

In equilibrium, raw returns from all 1,000 runs have a mean of 0 and a standard deviation of 0.10, corresponding to the prediction error ς made by traders processing information level $I = 7$ (see Table 10.1 for prediction errors). Considering that less than a quarter of all shares in

the market are traded on information, prices reflect the intrinsic values quite well.

Next, we check what happens when we leave equilibrium. The left panel of Fig. 10.4 shows the standard deviation of returns, depending on the (homogenous) strategy used by T_1 . As can be seen, the price accuracy gets worse if traders in T_1 start to process information. Compared to the random strategy played in equilibrium, price accuracy only increases if traders adopt information levels $I \geq 8$.




Fig. 10.4. Standard deviation of returns when trader class T_1 (left panel) or a single trader in T_9 (right panel) varies his strategy.

The maximum standard deviation is reached at information level $I = 1$. As one trader in T_8 plays the same strategy, this information is already reflected in market prices. If further traders base trading decisions on the same information, they make joint mistakes and prices become biased. The same explanation holds for the peak at $I = 5$. As one trader in T_9 already processes this information, further traders adopting the same information level will make it more likely for all traders with this strategy to act as price-makers. Although the individual prediction error of T_1 rises when changing his strategy to $I = 3$ or $I = 4$, price fluctuations are lower in those cases. According to that, herding behavior does not only affect the performance of the herd, it is also likely to drive prices further away from fundamental values than trading on noise. This result is inline with the findings in Schredelseker (2001).

The right panel of Fig. 10.4 shows price fluctuations depending on the strategy adopted by the two traders in T_9 . Strategy changes of the trader playing $I = 7$ ($I = 5$) in equilibrium are plotted with black (grey) bars. The graph shows that the dependence of price accuracy and information level is transitive for both traders, with one exception

in each case: for the first trader, $I = 4$ causes a lower price fluctuation than $I = 5$, as the second trader already adopts $I = 5$. For the second trader, adopting $I = 6$ results in a lower price fluctuation than with $I = 7$, which is the equilibrium strategy of the first trader.

The graph also shows that the maximum number of noise traders is reached not only in terms of possible gains; if one trader in T_9 leaves his equilibrium strategy to become a noise trader, the average price fluctuation increases beyond the prediction error of $I = 1$.

Although this model states that in equilibrium, certain traders do have an incentive to process information, it cannot deliver a solution to the information-paradox as formulated by Grossman and Stiglitz (1980). By using the intrinsic value to calculate returns, we implicitly assume that prices converge to the fundamental value. But the model shows that one might need a large number of noise traders to drive prices away from fundamental values.

10.6 Robustness checks

We check the robustness of our results by varying model parameters in several ways: equilibria for homogenous strategies of trader classes are computed for linear transformations of the original cost function by changing c in (10.3). We find that the number of random traders in equilibrium does not change when the cost of information is increased. Traders in T_9 and T_{10} then adopt lower information levels, which makes it more difficult to find settings with stable equilibria. When decreasing costs beyond a certain level, we find that a higher number of traders process information in equilibrium. With very low costs, also traders in T_8 and T_9 can easily afford to adopt an information level of $I = 10$. In these situations, T_{10} cannot stand out from the crowd anymore, leading to several trader classes using the same (maximum) information level. According to this, the model seems to be well balanced with $c = 0.02$, as the highest information level adopted gets close but does not reach the maximum of $I = 10$. Equilibria for different values for c are shown in Table 10.4.

We also check equilibria for various linear cost functions, following the equation

$$C_I = I \times c . \quad (10.6)$$

As this makes medium- and low information levels more expensive, traders in T_9 then adopt lower information levels in equilibrium. How-

Table 10.4. Strategy allocation in equilibrium for different cost functions.

T	1	2	3	4	5	6	7	8	9	10
$I_{T,c=0.0002}$	0	0	0	0	0	0	2	10	10	10
$I_{T,c=0.002}$	0	0	0	0	0	0	0	5	10	10
$I_{T,c=0.02}$	0	0	0	0	0	0	0	0	7	9
$I_{T,c=0.2}$	0	0	0	0	0	0	0	0	3	5
$I_{T,c=2}$	0	0	0	0	0	0	0	0	1	2

ever, the proportion of fundamental traders and random traders in equilibrium does not change.

The information system is also varied in several ways. As shown in Table 10.1, the prediction error declines progressively with higher information levels. Using information systems where the prediction error is a linear or logarithmic function of the information level does not affect the main results.

The way random traders place their orders does affect equilibria. If one limits the price impact of random traders too strictly (e.g. by setting the rule that in case of more than 50% of all orders in the order book being buy-orders from random traders, market price is set to the best sell-order), random trading gets even more profitable. Due to this, the share of traders processing information in equilibrium declines below 20% and equilibria become less stable.

Several ways how traders' strategies converge to equilibrium are tried, e.g. by varying and randomizing the choice which traders change strategies first. All equilibria found share the same characteristics as presented in Sect. 10.4.

10.7 Conclusions

This paper presents a simulation model where traders with heterogeneous market impact can endogenously choose to act as random traders, or to adopt a fundamental strategy by processing costly information. In equilibrium, less than a quarter of all shares are traded on information. As the proportion of fundamental traders and random traders in equilibrium was found to be robust against several variations of the model, we suggest that for traders with a low or medium investment volume, a random strategy is likely to perform better than low-information strategies, and it should on average outperform actively as well as passively managed investment funds.

By analyzing market efficiency in equilibrium, we show that a small share of traders processing information is sufficient to keep market prices close to fundamental values. In equilibrium, prices reflect the fundamental values even better than in situations where additional traders process low- and medium-quality information. Referring to Black (1986), we should rather blame the low- and medium-informed fundamental traders for making prices behave like “a drunk, tending to wander farther and farther from his starting point”.

Acknowledgement. I thank Michael Hanke, Jürgen Huber, Sonja Huber, Michael Kirchler, Klaus Schredelseker and two anonymous referees for helpful suggestions.

References

- F. Black. Noise. *The Journal of Finance*, 41:529–543, 1986.
- E. F. Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25:761–777, 1970.
- S. J. Grossman and J. E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70:393–408, 1980.
- M. Hanke and K. Schredelseker. Index funds should be expected to underperform the index. Working Paper, University of Innsbruck, 2005.
- J. Huber, M. Kirchler, and M. Sutter. Is more information always better? experimental financial markets with cumulative information. *Journal of Economic Behavior and Organization*, forthcoming, 2007.
- M. Jensen. Some anomalous evidence regarding market efficiency. *Journal of Financial Economics*, 6:95–101, 1978.
- B. C. Malkiel. Reflections on the efficient market hypothesis: 30 years later. *The Financial Review*, 40:1–9, 2005.
- L. Samuelson. Modeling knowledge in economic analysis. *Journal of Economic Literature*, 42:367–403, 2004.
- K. Schredelseker. Is the usefulness approach useful? some reflections on the utility of public information. In *Contemporary Issues in Accounting Regulation*. Kluwer Academic Publishers, 2001.
- G.K. Zipf. *Human Behavior and the Principle of Least Effort*. Addison-Wesley (Reading MA), 1949.

Analysis of Economic and Social Networks

A Note on Symmetry in Job Contact Networks

Andrea Mario Lavezzi¹ and Nicola Meccheri²

¹ University of Palermo, Department of Studies on Politics, Law and Society, Italy.

`lavezzi@unipa.it`

² University of Pisa, Department of Economics, Italy.

`meccheri@ec.unipi.it`

11.1 Introduction

Since the seminal work of Granovetter (1995), the sociological literature highlighted the importance of social relationships, like friends, relatives and acquaintances, as sources of information on jobs in labor markets. Such importance is also confirmed by a number of empirical studies.³ More recently, economists have devoted considerable attention to this topic,⁴ so that the study of individual and aggregate economic outcomes produced by the presence of social relationships in labor markets is becoming a fruitful research area in economics.

An important issue in the studies on social networks refers to how the network structure matters, that is how network characteristics, such as topology and type of connections play a role in explaining the economic effects of the networks. For instance, the effects of networks symmetry have been often discussed qualitatively in the sociological literature (e.g. Granovetter (2005)). On the other hand, the quantitative effects that such network's property may produce on output and wage inequality have still not received the same attention.

In this paper we tackle this issue and study the effect of symmetry on workers' aggregate output and inequality.⁵ In particular, we adopt a version with heterogeneous jobs of the model by Calvo-Armengol and

³ See Montgomery (1991) for further discussion and references.

⁴ See Ioannides and Loury (2004) for a survey.

⁵ In the extended version of this paper (Lavezzi and Meccheri, 2005a) we also begin to study the role of other networks' properties on output and inequality, such as social exclusion and network density (see Lavezzi and Meccheri, 2005a, also for more details on the related literature).

Jackson (2007), in which exogenous social networks facilitate the transmission of information on job vacancies among workers.⁶ We find that: a) symmetric networks produce higher output and lower inequality than asymmetric networks and, b) the introduction of social links, having the function of “structural holes” (see Burt, 1992), has a larger positive effect on output and inequality if they are associated with symmetric networks.

The paper is organized as follows: Section 11.2 presents the theoretical model; Section 11.3 contains the results of the simulations; Section 11.4 concludes.

11.2 A Model of labor market with social networks

11.2.1 Production, wages and turnover

Time is discrete and indexed by $t = 0, 1, 2, \dots$. The economy is populated by homogenous, risk-neutral, infinitely-lived agents (workers) indexed by $i \in \{1, 2, \dots, N\}$. In each period a worker can be either unemployed or employed in a “good” or “bad” job. Thus, by indicating with s_{it} the employment status of worker i in period t , we have three possible agents’ states:

$$s_{it} = \begin{cases} g, & \text{employed in a good job} \\ b, & \text{employed in a bad job} \\ u, & \text{unemployed} \end{cases}$$

On the production side, we consider one-to-one employment relationships and assume a very simple form of a production function, in which productivity depends on the job offered by a firm to a worker. In particular, we denote by y_{it} the output of a firm employing worker i at time t or, in other words, the surplus generated by the match between a worker and a firm (output price is normalized to one).

We simply assume that output in a good job is higher than in a bad job, for instance because it is a hi-tech job. According to these assumptions, the parameter y^s , $s \in (g, b, u)$, indexing the productivity of a match, follows the rule:

$$y^g > y^b > 0 (= y^u).$$

⁶ Calvo-Armengol and Jackson (2004) provides some simulations on the role of networks’ topology for the simpler case of homogeneous workers and jobs.

Wages are a fraction of the match surplus, and are denoted by $w^s = \beta y^s$ with $\beta \in (0, 1)$.⁷ This produces an ordering of wages obtainable in a given match, which follows the ordering of outputs. Obviously, unemployed workers earn zero wages, and we normalize their reservation utility to zero.

The labor market is subject to the following turnover. Initially, all workers are unemployed. Every period (from $t = 0$ onwards) has two phases: at the beginning of the period each worker receives an offer of a job of type f , with $f \in \{g, b\}$, with arrival probability $a_f \in (0, 1)$.⁸ Parameter a_f captures all the information on vacancies which is not transmitted through the network, that is information from firms, agencies, newspapers, etc. When an agent receives an offer and she is already employed and not interested in the offer, in the sense that the offered job does not increase her wage, she passes the information to a friend/relative/acquaintance who is either unemployed or employed but receiving a lower wage than the one paid for the offered job. At the end of the period every employed worker loses the job with breakdown probability $d \in (0, 1)$.

11.2.2 Social links and job information transmission

social networks may be characterized by a graph G representing agents' links, where $G_{ij} = 1$ if i and j know each other, and $G_{ij} = 0$ indicates if they do not. It is assumed that $G_{ij} = G_{ji}$, meaning that the acquaintance relationship is reciprocal. Given the assumptions on wages and arrival probabilities, the probability of the joint event that agent i learns about a job and this job ends up in agent's j hands, is described by p_{ij} :

$$p_{ij}(s_{it}^\theta, f) = \begin{cases} a_b & \text{if } f = b \cup j = i \cup s_i = u \\ a_g & \text{if } f = g \cup j = i \cup (s_i = u \cap s_i = b) \\ a_b \frac{G_{ij}}{\sum_{k:s_k=u} G_{ik}} & \text{if } f = b \cup (s_i = b \cap s_i = g) \cup s_j = u \\ a_g \frac{G_{ij}}{\sum_{k:s_k \neq g} G_{ik}} & \text{if } f = g \cup s_i = g \cup (s_j = u \cap s_j = b) \\ 0 & \text{otherwise} \end{cases}$$

In the first two cases, worker i receives an offer with probability a_f , $f \in \{g, b\}$, and takes the offer for herself. This holds if she is either unemployed or employed in a bad job and receives an offer for a good job.

⁷ For instance β may represent the bargaining power of workers when wages are set by Nash bargaining, as is usual in search models. Clearly, profits are $(1 - \beta)y^s$.

⁸ That is, each agent can receive both an offer for a good and a bad job.

In the third case the worker i is employed and receives with probability a_b an offer for a bad job, that she passes only to an unemployed worker $j (\neq i)$. We assume that among all unemployed workers connected with i by a social link, i chooses j randomly. Hence, the probability that worker j receives the information by worker i is equal to $\frac{G_{ij}}{\sum_{k:s_k=u} G_{ik}}$. In the fourth case worker i receives with probability a_g an offer for a good job when she is already employed in a good job, thus she passes the offer, with probability $\frac{G_{ij}}{\sum_{k:s_k \neq g} G_{ik}}$, to a worker connected to her who is either unemployed or employed in a bad job. Clearly, $p_{ij} = 0$ in all remaining cases.

To sum up, a worker who receives an offer makes direct use of it if the new job opportunity increases her wage. Otherwise, she passes the information to someone who is connected to her. The choice of the worker to whom pass the information is “selective”, in the sense that the information is never passed to someone who does not need it,⁹ but it is random with respect to the subset of the connected workers who improve their condition (wage) exploiting such information (for example, a worker receiving a good job offer is indifferent to pass it to an unemployed contact or a contact employed in a bad job). Finally, we assume that a worker receiving both an offer for a bad and a good job when she does not need them, decides to transmit first the information about the bad job and then, possibly to the same agent, the information about the good job, and we exclude that each job information may be transmitted to more than one (connected) worker.¹⁰

Figure 11.1 shows the timing of the events for a generic period t (for convenience, the period has been represented as composed by four different consecutive sub-periods, with sub-periods $t.1$, $t.2$ and $t.4$ having negligible length).

⁹ For the sake of simplicity, we assume that a worker observe the state of her connections at the end of the previous period to make a decision on passing information. In other words, she cannot observe if her connections have already received an offer from someone else. If all of the worker’s acquaintances do not need the job information, then it is simply lost.

¹⁰ Calvo-Armengol and Jackson (2007) provide various extensions on the process of transmission of job information.




Fig. 11.1. Timing

11.3 On networks symmetry

In this section we present the results of our simulations.¹¹ Our aim is to assess how the structure of social networks affects the dynamics of output and wage inequality in the long run, as well as the correlation of workers' wages. We measure output by averaging over time the average output of the n workers in every period. Inequality is measured by the average Gini index over time.

As a preliminary general remark, it is important to point out that in this framework the network structure basically affects the possibility for the system to be in a state of maximum output (*SMO* henceforth), that is a state in which all agents are employed in the good job. Given our assumptions, *SMO* would be a steady state if the probability of losing the job was zero, as workers would be in the best possible position and would turn down any offer they received, directly or indirectly. In other words, without the exogenous breakdown probability, *SMO* would be an absorbing state for the system. In this respect, the network structure regulates the possibility to attain *SMO* and the speed at which the system recovers to it, after the occurrence of stochastic perturbations given by breakdowns of job relationships. Therefore, as we shall see, high average levels of output and low levels of inequality obtain when the system, driven by the network structure, reaches faster and persistently remains in *SMO* (note that in *SMO* inequality is clearly absent).

¹¹ Simulations are run for 500,000 periods and the parameters are: $a_b = 0.15$, $a_g = 0.10$, $d = 0.015$, $\beta = 0.4$, $y^g = 5$, $y^b = 1$. All simulations are programmed in R (<http://www.r-project.org/>). Codes are available upon request from the authors.

11.3.1 Symmetric vs. asymmetric networks

We begin by considering a simple case of symmetric vs. asymmetric networks. Technically speaking, symmetry implies that all agents are connected to the same number of other agents. Consider the network




Fig. 11.2. Networks G_A and G_B

structures in Figure 11.2, G_A and G_B .¹² Both networks have the same number of agents, n , and links, N ($N = n = 4$), and the same average number of links for each agent, that is $\mu = 2$.¹³ However, they have a different geometry: network G_B is obtained from G_A by simply rewiring one link. This introduces an asymmetry, as in network G_B agent 2 has three links and agent 3 has one link, while agents 1 and 4 maintain the same number of links. In other words, agents 1, 2 and 4 form a cluster of interconnected agents, from which agent 3 is partially excluded. In addition, there exists a difference in the number of links of the agents to whom every agent is connected. In network G_A any agent has two links with agents who have two links. Differently, in network G_B agents 1 and 4 have one link with an agent with two links (respectively agents 4 and 1), and one link with an agent with three links, agent 2. Agent 2 has two links with two agents, 1 and 4, who have two links, and one link with agent 3, who has one link. As we show in Table 11.1, this has consequences for both output and inequality.

We observe that, moving from G_A to G_B , output decreases and inequality increases. The emergence of a local cluster makes the network asymmetric, and affects both output and wage inequality. In particular, the decrease in output and the increase on inequality depend on the

¹² This case corresponds to Example 1 in Calvo-Armengol (2004) where, differently from here, workers and jobs are both homogeneous. In general, in our examples we will consider networks where not all possible links are formed as a simple way to consider the fact that link formation is costly. See Calvo-Armengol (2004) for a full treatment of network formation with costly links.

¹³ The simple formula to obtain μ , the average number of links per agent, is $2N/n$.

relative isolation of agent 3. Agent 3's average wage is sharply lower in network G_B . In this case the increase in the average wage of agent 2, due to an increase in the number of her connections, is not sufficient to counterbalance the decrease in the average wage of agent 3. Also notice that the variance of agent 2's wage is lower while the variance of agent 3's wage is higher in network G_B .

Table 11.1. Output, inequality and wages, networks G_A and G_B

Network	Output	Ineq.	Av. wages [1, 2, 3, 4]	Var. wages [1, 2, 3, 4]
G_A	4.818	0.034	1.927, 1.927, 1.928, 1.928	0.122, 0.123, 0.121, 0.120
G_B	4.802	0.038	1.924, 1.945, 1.889, 1.924	0.126, 0.091, 0.183, 0.127

Results are also different for agents 1 and 4 although the number of their connections is the same. In particular their average wage is lower and the variance is higher in network G_B . This can be explained by the fact that the number of links of their “connections” is different in network G_B , in particular they are both connected to agent 2 who has three links. This implies that their probability of receiving information on vacancies from agent 2 is lower in network G_B , as they have more “competitors” for information. This result is not so obvious since there could be also a positive effect deriving from a connection with an agent with many links, which should guarantee a more stable position in the state of employment and therefore have a higher propensity to transmit information on vacancies. We term the first effect as *competition effect*, and the second as *connection effect*, and note that the former dominates the latter in network G_B .

These results highlight the complexity of capturing the externalities produced by the structure of the network. In the present framework, the network exerts externalities on agents' utilities as it affects their job opportunities. However, to put these network externalities in closed form is not an easy task, as they derive from a network stochastic process.¹⁴ Our numerical results, however, clearly show that such externalities differ across individuals depending on their location in the network. Moreover, switching from a symmetric to an asymmetric structure, it appears that the negative externalities that derive seem to prevail on positive externalities, since in symmetric networks aggregate results improve.

¹⁴ In a different setting, strategic and static, Ballester et al. (2006) studies analytically the variance of network externalities.

worker	1	2	3	4	worker	1	2	3	4
1	1	0.031	0.026	0.025	1	1	0.038	0.014	0.048
2	0.031	1	0.027	0.026	2	0.038	1	0.022	0.038
3	0.026	0.027	1	0.020	3	0.014	0.022	1	0.010
4	0.025	0.026	0.020	1	4	0.048	0.038	0.010	1

Table 11.2. Correlation of workers' wages. Left panel, network G_A ; right panel, network G_B .

The creation of a local cluster also affects the distribution of wage correlations across each pair of agents in the network; from Table 11.2 we see that, as predictable, the values of the correlation of wages of the agents in the cluster (i.e. agents 1, 2 and 4) increase.¹⁵ In network G_B agent 3's correlations with any other agent decreases. Note that the correlation between agent 3 and 2's wages is lower, despite the fact that the two agents share a link as in network G_A . In network G_B , however, agent 2 has one extra link and, in practice, this weakens the connection between 2 and 3. Finally, the correlation between 3 and 4 is lower in network G_B because they are not directly connected.¹⁶

To sum up, the introduction of asymmetry in a network which preserves the same average number of links, causes output to decrease and inequality to increase.¹⁷ Next section takes a step further.

11.3.2 Symmetry and relational heterogeneity

In this section we go deeper on the role of social networks' symmetry in explaining economic outcomes. In fact, as remarked for instance by Ioannides and Loury (2004, p. 1064), there exist results related to social networks structure that may be explained by symmetry, while they have been often attributed to other network properties. Indeed, much of the early sociological research on the effects of job networks' properties

¹⁵ All these numerical results are in accordance with the analytics of Calvo-Armengol and Jackson (2007).

¹⁶ Other examples with larger networks, not reported here but available upon request, confirm these results.

¹⁷ These results are in line with those of Calvo-Armengol (2004) on welfare and unemployment. However, this depends on the fact that we introduced heterogeneity in jobs but preserved the homogeneity of workers. In fact, in the setting with heterogeneous jobs and workers of Lavezzi and Meccheri (2005b), we provide an example in which, when productivity differentials are sufficiently high, more symmetry can be associated to lower output. A full treatment of this aspect is beyond the scope of the present work.

mainly focused on *relational heterogeneity*, emphasizing that not all social relations (contacts) have the same role or strength in affecting employment outcomes. Here we aim to disentangle in our framework the effects of network symmetry on output and inequality with respect to some traditional concepts related to relational heterogeneity.

An important argument in the theory of social networks refers to the role of *structural holes*. As is well-known, Burt (1992) defines structural holes as the “gap” of non-redundant links: agents placed at structural holes of a network allow information to flow between otherwise unconnected groups of agents. The structural holes argument implies that networks with more non-redundant links (i.e. more agents placed at structural holes) can provide more information than network *of the same size*, but with more redundant links (see Ioannides and Loury, 2004, p. 1063). Thus, networks in which (structural holes) agents link otherwise unconnected groups should be characterized by more efficient outcomes, since information in such networks circulates more widely.¹⁸

Consider the different network structures in Figure 11.3, with $n = 8$, $N = 12$ and $\mu = 3$.




Fig. 11.3. Networks G_C , G_D , G_E : symmetry and structural holes

¹⁸ Note that the “structural hole effect” amplifies when information can be transmitted to indirect relationships (more than two-links away) by means of sequential passages, while in our simulations information may be transmitted only one time between direct contacts. However, in a long-run perspective, since the transmission of information can improve the state of one (connected) agent in a given period and this allows her to be more prone to transmit information to others in future periods, the same effect should apply.

In network G_C , there are two separated groups of four agents and, in each group, each agent is linked to each other. Clearly, this is a symmetric network, since all agents are connected to the same number of other agents (three, in this case). It is important to point out that, according to the theory of structural holes, some links in network G_C are, at least partially, redundant, since each pair of agents could be (indirectly) linked anyway via other agents in their group (e.g. agents 1 and 2 are linked via agent 7), and it would be more efficient to have some links to agents in the other group.

In network G_D some agents become structural holes: 1 and 8 for the first subgroup, 3 and 6 for the second. The two groups are linked through a bridge provided by structural holes. The network is not symmetric, since there are agents with a different number of links: i.e. agents 3 and 6 have now four links, while agents 2 and 7 have just two links. Finally, network G_E is a symmetric network with the same number of “structural holes agents” (1, 3, 5 and 7 in this case) of G_D .

Running simulations for these networks, we obtain the following aggregate results:

Table 11.3. Networks G_C , G_D , G_E : output and inequality

	Network Output	Inequality
G_C	4.863	0.027
G_D	4.862	0.027
G_E	4.867	0.026

We note that the introduction of structural holes and asymmetry in network G_D slightly reduces output and leaves inequality unchanged, while output is higher and inequality slightly lower in G_E . Also, even if the results are fairly close, this example seems to suggest that, in the aggregate, the positive effect on output which may derive from the introduction of “bridges” between different groups (as can be the case in a passage from G_C to G_E) could be counterbalanced if those bridges are created by rendering asymmetric the structure of the network and, consequently, the position of different agents.

This appears more transparent if we look at Table 11.4, in which we report individual (average) wages of three workers (1, 2 and 3)¹⁹

¹⁹ The situation of the chosen workers changes differently when we move from network G_C to network G_D ; in this sense, they have been chosen as representing typical cases. Of course, the same qualitative results also hold for other workers in analogous situations.

in G_C , G_D and G_E . Table 11.4 also shows the wage correlation of two workers with no direct connections (1 and 6) in the different networks.

Table 11.4. Networks G_C and G_D : individual wages [1,2,3] and correlation [1;6]

Network	Av. wage [1]	Av. wage [2]	Av. wage [3]	Corr. wages [1;6]
G_C	1.944	1.944	1.944	0.000
G_D	1.947	1.929	1.959	0.010
G_E	1.947	1.947	1.948	0.012

Wages of agents 1 and 3 increase in network G_D . While the increase of agent 3's wage is largely due to the fact that she has now one extra link, the increase of agent 1's wage is related to a "structural hole" effect: given that the number of her connections is unchanged, now she is linked to the other group and can take advantage, directly or indirectly, from the presence of all workers in the economy. Agent 2 loses one link in G_D and her wage becomes lower than in G_C . In principle agent 2 could have benefited from the presence of a bridge connecting her group to the other, but this appears not sufficient to outweigh the negative effect of losing one link. Moreover, this negative effect appears so powerful that, although some agents become structural holes and have more links, aggregate output (which is proportional to wages) slightly decreases in network G_D .²⁰

In network G_E agents maintain the same number of links as in G_C , but some agents (1, 3, 5 and 7) become structural holes. The wage of agents 1 and 3 increases with respect to G_C , indicating that these agents benefit from a better circulation of information. The wage of agent 2, who is not a structural hole, increases as well with respect to G_C even if the number of links is the same, and is much higher than in G_D . The latter effect clearly depends on 2 having more links in G_E than in G_D .

Finally, note from Table 11.4's last column that the presence of bridges between the two groups of workers affects the structure of wages correlations. While wages of workers 1 and 6 are not correlated in network G_C , where they belong to two separated groups, the correlation becomes positive in networks G_D and G_E , even if those workers are not directly connected.

²⁰ This suggests the presence of decreasing returns from increases in the number of social links. This aspect is analyzed in more detail in Lavezzi and Meccheri (2005a).

Overall, the positive effects of bridges creation can be better appreciated in network G_E in which symmetry is preserved. In such a case, output increases and inequality slightly decreases with respect to other networks. This happens because the advantages of a wider circulation of information can be exploited at no costs for agents, in the sense that they maintain the same number of links.

11.4 Conclusions

In this paper we have provided an initial study of the effects of network symmetry on output and inequality. In particular, our results allow for a first set of considerations.

The relevance of symmetric social architectures, which appeared in our examples, points to the relevance of having an “egalitarian” society in which individuals are relatively similar in their degree of social interaction. Moreover, the importance of symmetry also appears in relation to the presence of structural holes, which effects may depend on their being related to symmetry or asymmetry of the network. More precisely, the importance of structural holes is in fact related to the possibility of connecting two or more otherwise disconnected groups of individuals by establishing a symmetric geometry.²¹

References

- C. Ballester, A. Calvó-Armengol, and Y. Zenou. Who’s who in networks. wanted: the key player. *Econometrica*, 74(5):1403–1471, 2006.
- R. S. Burt. *Structural Holes*. Harvard University Press, 1992.
- A. Calvo-Armengol. Job contact networks. *Journal of Economic Theory*, 115(1):191–206, March 2004.
- A. Calvo-Armengol and M. O. Jackson. Networks in labor markets: Wage and employment dynamics and inequality. *Journal of Economic Theory*, 127(1):27–46, January 2007.
- A. Calvo-Armengol and M. O. Jackson. The effects of social networks on employment and inequality. *American Economic Review*, 94(3): 426–454, June 2004.
- M. S. Granovetter. The impact of social structure on economic outcomes. *Journal of Economic Perspectives*, 19(1):33–50, Winter 2005.

²¹ In Lavezzi and Meccheri (2005a) we study symmetry with respect to the “strength of weak ties hypotheses” of Granovetter (1973).

- M. S. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78(6):1360–1380, 1973.
- M. S. Granovetter. *Getting a Job: a Study of Contacts and Careers*. University of Chicago Press, 1995.
- Y. M. Ioannides and L. D. Loury. Job information networks, neighborhood effects, and inequality. *Journal of Economic Literature*, 42(4): 1056–1093, December 2004.
- A. M. Lavezzi and N. Meccheri. Social networks in labor markets: The effects of symmetry, randomness and exclusion on output and inequality. *Computing in Economics and Finance 2005* 277, Society for Computational Economics, Nov 2005a.
- A. M. Lavezzi and N. Meccheri. Job contact networks, inequality and aggregate output. In N. Salvadori and R. Balducci, editors, *Innovation, Unemployment and Policy in the Theories of Growth and Distribution*, pages 145–167. Edward Elgar, 2005b.
- J. D. Montgomery. Social networks and labor-market outcomes: Toward an economic analysis. *American Economic Review*, 81(5):1407–1418, December 1991.

Innovation and Knowledge Spillovers in a Networked Industry

Jose I. Santos¹, Ricardo del Olmo¹, and Javier Pajares²

¹ University of Burgos, Spain
{jisantos,rdelolmo}@ubu.es

² University of Valladolid, Spain
pajares@eis.uva.es

12.1 Introduction

Knowledge and proximity are key concepts in the Geography of Innovation literature (Boschma, 2005). Innovating processes are uncertain because they often take place under unsure conditions and fierce business competitiveness. Geographical proximity can reduce this uncertainty since it potentially facilitates labor movement and knowledge interchange through personal contacts. Supporting this hypothesis, some scholars have highlighted the greater agglomeration of RD activities in technological industries where knowledge plays a significant economic role (Audretsch and Feldman, 1996a).

Our concept of *proximity* is not strictly geographical. There are other dimensions beyond spatial nearness, such as social proximity and cognitive proximity that could have diverse effects on innovation (Boschma, 2005). The social dimension can be modeled using a network of relations that shape agents' opportunities to interact, whereas the cognitive dimension can be captured using the absorptive capacity concept (Cohen and Levinthal, 1990): a firm's capacity to understand, learn and apply the knowledge generated outside itself.

The objective of this paper is to study the impact of network structure and knowledge proximity on the process of innovation and diffusion of knowledge through spillovers. We use agent-based modeling and computer simulation to implement and explore a formal model of an innovative industry, which could not be studied using other analytical approaches (Pajares et al., 2004).

12.2 The Model

In order to have a clearer vision of the proposed problem, we use a simplified abstraction of an innovative industry: a set of N innovating firms indexed in i and endowed with a scalar knowledge $K_i \in R$ which is assumed to affect their innovative effort.

Firms are organized in an undirected graph built following the well-known algorithm proposed by Watts and Strogatz (1998). We start with a one-dimensional ring of N vertices where each vertex links to its k closest neighbors; then, we randomly rewire every individual edge in the graph with probability p . This simple algorithm generates a family of networks that exhibit different values for their average path length and their clustering coefficient, depending on the value of the parameter p : from a fixed, highly clustered and ordered network where the average path length is large ($p = 0$), to the set of pure random networks ($p = 1$), which is characterized by low clustering coefficients and relatively short average path lengths.

Each firm can innovate and increase its knowledge. Following an evolutionary economic approach (Nelson and Winter, 1982) the innovating activity is modeled as a stochastic process. We define firm i 's probability of innovating at time t , $P_i^{in}(t)$, as an exponential function:

$$P_i^{in}(t) = p_{max}^{in} - (p_{max}^{in} - p_{min}^{in}) \exp(-\alpha \max(0, \overline{K}_{Vi} - \overline{K})) \quad (12.1)$$

This equation integrates two common concepts of the Innovation literature: the technological opportunities and the innovative effort. Innovating can be more difficult in some industries than in others (Klevorick et al., 1995); the possibilities to create significant technological novelties define the technological opportunities of an industry. The parameters p_{max}^{in} and p_{min}^{in} , which represent the maximum and minimum probability of innovating, determine the technological opportunity regime in the model.

Not every firm invests the same resources and time in innovative activities, and they are not equally efficient. In our model the intensity of the innovative effort varies from one firm to another according to the firm's relative advantage in knowledge over the industry, which is quantified as the difference between the average knowledge in the firm's neighborhood¹ $\overline{K}_{Vi}$ and the average knowledge in the industry $\overline{K}$. The equation (12.1) shows a positive effect of knowledge spillovers because

¹ It is the average knowledge of a firm and its immediate neighbors: vertices directly linked to it in the network.

knowledge interchange between a firm and its neighbors reinforces their innovative efforts, and therefore their probabilities of innovating. The effectiveness of this innovative effort depends on the parameter α that governs the growth rate of the probability function.

The assumption about the innovative effort responds to a simple cumulative causation: having comparatively greater knowledge gives firms an economic advantage that enhances their very process of innovation, and it therefore reinforces the probability of creating even greater knowledge. Thus, there is a positive feedback in the process of innovation: the more you innovate the more likely you are to innovate even more.

The increase in knowledge derived from a successful innovation is calculated using (12.2), where β controls the innovation jump.

$$K_i(t+1) = K_i(t) (1 + \beta) \quad (12.2)$$

Before the firm i has an opportunity to innovate, it can learn from its immediate neighbors j in the network, according to (12.3).

$$\Delta K_i(t+1) = AC_{ij}(t) \max(0, K_j(t) - K_i(t)) \quad (12.3)$$

The expression above describes a knowledge spillover from firm j to firm i as an interactive diffusion process modulated by the absorptive capacity AC_{ij} between the firms. The absorptive capacity AC_{ij} denotes the fraction of the (positive) difference in knowledge between the two firms ($K_j - K_i$) that firm i acquires. The value of AC_{ij} depends on the corresponding knowledge proximity ($K_i - K_j$), according to another exponential function (12.4).

$$AC_{ij}(t) = \frac{c_0}{1 + \exp(\gamma(K_j(t) - K_i(t) - d))} \quad (12.4)$$

The absorptive capacity AC_{ij} is close to the default value c_0 for short knowledge distances, gets the value $c_0/2$ for a knowledge distance equal to d and falls down to zero for larger distances. The parameter γ controls the sharpness of the fall, and the parameter d determines the distance where this fall happens. The last parameter is a measure of the firm's sensitivity to knowledge distance, the larger it is the more different knowledge a firm can learn. Fig. 12.1 depicts the functions (12.3) and (12.4) for different values of the parameter γ .




Fig. 12.1. On the left, the absorptive capacity function (12.4) for $c_0 = 1$, $d = 0.5$ and different values of γ ; on the right, the corresponding effective knowledge gain function (12.3).

Scheduling can be summarized as follows: each time-step one randomly chosen firm is given the opportunity to learn from its neighborhood and innovate, as described above. The model allows us to simulate complementary scenarios:

- Considering the technological opportunity regime: an industry with a few ($p_{max}^{in} \approx p_{min}^{in} \approx 0$) or many technological opportunities ($p_{max}^{in} \gg p_{min}^{in} \approx 0$).
- Looking at the knowledge spillovers process: a tacit knowledge regime with low values of the parameter d versus a codified knowledge regime with high values of d .
- Modifying the structure of the network with the parameter p we can compare innovation processes in highly regular networks of diffusion (low values of p) versus innovation processes in random networks of diffusion (values of p close to 1).

12.3 Simulations

Under the hypothesis of homogeneous firms (endowed with similar initial knowledge and capacity to learn through spillovers), we are interested in studying the impact of network structure (using parameter p) on innovation dynamics in an industry, using the rest of the model parameters to define different industry scenarios. In this section we

summarize the main simulation results that allow us to advance some conclusions of this work².

The industry dynamics are summarized here using the average knowledge (computed as the mean of the firms' knowledge). This industrial rate follows an increasing path without steady states, so we evaluate them close to the horizon of simulations, arbitrary fixed³.

In the proposed model the innovative advantage of any firm depends mainly on its innovative effort, measured by the difference between the average knowledge in the firm's neighborhood and the average knowledge in the industry. We can intuitively infer that clustered networks ($p \ll 1$) will get better rates because individual innovations diffuse to close neighbors, reinforcing their knowledge advantage above the rest faster than random networks ($p \gg 0$) where diffusion are like scattered showers. But inferences are not so obvious when we analyze the technological opportunities and the knowledge proximity effects on industry dynamics.

12.3.1 Innovation, networks and technological opportunities

We define a first scenario of an industry with few technological opportunities, which corresponds to the following parameterization: $p_{max}^{in} = 0.11$, $p_{min}^{in} = 0.1$ and $AC_{ij} = c_o = 0.1$. Fig. 12.2 shows the evolution of the average and the coefficient of variation of knowledge for a clustered network ($p = 0$) and a random network ($p = 1$). The graph depicts an expected phenomenon: knowledge dispersion drops due to knowledge diffusion, and the average knowledge grows due to firms' innovations. The simulation horizon of 300 time units is enough to evaluate the parameters' effects on the industry performance. We infer an interesting conclusion: when an industry is characterized by a few technological opportunities, modeled as a low probability of innovating similar for all agents, network structure does not play a critical role, although, unlike expected, a random network gets better results than a clustered network because the clustering effect is non significant in these cases.

² The model has been implemented in Repast (North et al., 2006) and replicated in Netlogo. The simulation model is composed of $N = 500$ firms endowed with a scalar knowledge, which is initialized from a Uniform distribution $U(0, 1)$. We set the value of $k = 6$, $\beta = 0.01$, $c_o = 0.1$ and $\gamma = 10$.

³ A simulation run is composed of 300 periods, and 50 replications are recorded for each one. The standard error is less than 1% for all average statistics shown in the graphs of this paper.




Fig. 12.2. The evolution of the average and the coefficient of variation of knowledge for a clustered network ($p = 0$) and a random network ($p = 1$) when there are a few technological opportunities in the industry ($p_{max}^{in} = 0.11$, $p_{min}^{in} = 0.1$). The industry evolution is very similar for all network structures, although a random networked industry always gets better results than a clustered industry.

The next scenario represents an industry characterized by more technological opportunities: $p_{max}^{in} = 0.3$, $p_{min}^{in} = 0.05$ and $AC_{ij} = c_o = 0.1$. Now the parameter α in (12.1), which modulates the firm's innovative effort, has an important meaning: the firm's ability to apply its innovative effort successfully. Note that we have translated the absorptive capacity concept (Cohen and Levinthal, 1990) into the corresponding absorptive capacity coefficient (12.4) and this skill, represented by α , which can be interpreted as the necessary firm's know-how to turn the knowledge advantage into more possibilities to innovate (12.1).

Fig. 12.3 shows the effect of network structure on the industry average knowledge for different values of the parameter α .

For low values of α we get similar results than the scenario with a few technological opportunities: most firms get the same probability of innovating (close to p_{min}^{in}) whatever their innovative effort, and randomness in network relationships improves the industry performance. For higher values of α clustered networks ($p \ll 1$) always get better industrial rates than random networks ($p \gg 0$). This result is in agreement with empirical studies of industrial clusters (Audretsch and Feldman, 1996b). According to our model, this phenomenon would be related with the existence of technological opportunities in an industry and firms with the ability to exploit them.




Fig. 12.3. In the graph above, the effect of network structure (p) on the average knowledge for different values of the parameter α when there are significant technological opportunities ($p_{max}^{in} = 0.3$, $p_{min}^{in} = 0.05$). In order to highlight the network structure effect, magnitudes are normalized dividing by the corresponding value at $p = 0$. For high values of α clustered networks ($p \ll 1$) always get better industrial rates than random networks ($p \gg 0$). Moreover, small world networks ($0 \ll p < 0.1$), characterized by high clustering and low average path, get better results than a pure clustered network ($p = 0$) for values of α high enough. In the graph bellow, the corresponding normalized network features: clustering C and average path length L .

Another interesting inference is that small world networks ($0 \ll p < 0.1$), characterized by high clustering and low average path, get better results than a pure clustered network ($p = 0$). When firms randomly rewired a few local links, they keep the cluster advantage and also benefit from learning from distant sources.

12.3.2 Innovation, networks and knowledge distance

We can imagine a scenario where the knowledge proximity sensitivity (d) is imposed to the industry, and firms can decide the randomness of their relations (p), thus determining the network structure⁴. This approach would be in line with the product life cycle theory, which

⁴ An interesting evolutionary agent-based model that studies in detail the life cycle approach to innovative industry dynamics is proposed in Pajares et al. (2003).

argues that knowledge is tacit in the first stages of the life cycle and more codified in the mature phases (Klepper, 1996).

Fig. 12.4 shows the effect of network structure, considering the knowledge proximity parameter d as an exogenous variable. Here, either there is no doubt that clustered networks ($p < 1$) get better industrial rates than random networks ($p > 0$) for every value of d . We see from (12.4) that low values of d , e.g. $d = 1$, limit the diffusion though knowledge spillovers and thus accentuating the positive effect of clustering in the innovation process. The results are very sensitive to the parameter d , for a bit higher value, e.g. $d = 2$, the industry evolution is similar to a scenario where the knowledge proximity between firms does not affect their absorptive capacity ($d = \infty$).




Fig. 12.4. In the graph above, the effect of network structure (p) on the average knowledge for different values of the parameter d when there are significant technological opportunities ($p_{max}^{in} = 0.3$, $p_{min}^{in} = 0.05$). Magnitudes are normalized dividing by the corresponding value at $p = 0$. The positive effect of clustering in the innovation process is accentuated when the knowledge proximity between firms affects significantly their absorptive capacity ($d = 1$). In the graph below, the corresponding normalized network features: clustering C and average path length L .

12.4 Conclusions

We have proposed a formal model of an innovative networked industry with knowledge spillovers. Innovation is modeled as a stochastic process where firms' probability of innovating depends on its innovative effort and technological opportunities in the industry. The intensity of the innovative effort varies from one firm to another according to the firm's relative advantage in knowledge over the industry, which is quantified as the difference between the average knowledge in the firm's neighborhood and the average knowledge in the industry.

Knowledge spillovers are modeled as a simple knowledge diffusion process restricted by the firms' absorptive capacity, which depends on the knowledge proximity between firms. Knowledge spillovers have two opposite effects on industry dynamics: a positive one since knowledge interchange between a firm and its neighbors reinforces their innovative efforts, and a negative one since the diffusion process limits knowledge appropriability cutting any initial innovative advantage.

We study the impact of network structure and knowledge proximity on industry dynamics. With these hypotheses if an industry is characterized by a few technological opportunities, the clustering effect is non significant and randomness in network relationships improves industry dynamics. However, when there are significant technological opportunities, clustered networks get better results than random networks. This result is in agreement with empirical studies of industrial clusters; according to our model, this phenomenon would be related with the existence of both technological opportunities in an industry and firms with the ability to exploit them. In most cases small world networks that exhibit high clustering and low average path get better results than a pure clustered network. Finally, when the knowledge proximity between firms affects their absorptive capacity and thus limiting the diffusion through knowledge spillovers, the positive effect of clustering on the innovation process is reinforced.

Acknowledgement. This work has been supported by the Spanish MICYT, research project DPI2004-06590 and SIGAME project DPI2005-05676.

References

- D. B. Audretsch and M. P. Feldman. R&d spillovers and the geography of innovation and production. *American Economic Review*, 86(3): 630–640, June 1996a.

- D. B. Audretsch and M. P. Feldman. Innovative clusters and the industry life cycle. *Review of Industrial Organization*, 11(2):253–273, 1996b.
- R. A. Boschma. Proximity and innovation: A critical assessment. *Regional Studies*, 39(1):61–74, January 2005.
- W. M. Cohen and D. A. Levinthal. Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly, Special Issue: Technology, Organizations and Innovation*, 35(1):128–152, March 1990.
- S. Klepper. Entry, exit, growth, and innovation over the product life cycle. *American Economic Review*, 86(3):562–583, June 1996.
- A. K. Klevorick, R. C. Levin, R. R. Nelson, and S. G. Winter. On the sources and significance of interindustry differences in technological opportunities. *Research Policy*, 24(2):185–205, March 1995.
- R. R. Nelson and S. G. Winter. *An Evolutionary Theory of Economical Change*. Harvard University Press, 1982.
- M. J. North, N. T. Collier, and J. R. Vos. Experiences creating three implementations of the repast agent modeling toolkit. *Transactions on Modeling and Computer Simulation*, 16(1):1–25, January 2006.
- J. Pajares, , A. Lopez-Paredes, and C. Hernandez-Iglesias. Industry as an organisation of agents: Innovation and r&d management. *Journal of Artificial Societies and Social Simulation*, 10, 2003. Available at <http://jasss.soc.surrey.ac.uk/6/2/7.html>.
- J. Pajares, C. Hernandez-Iglesias, and A. Lopez-Paredes. Modelling learning and r&d in innovative environments: a cognitive multi-agent approach. *Journal of Artificial Societies and Social Simulation*, 7(2), 2004. Available at <http://jasss.soc.surrey.ac.uk/7/2/7.html>.
- D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.

Heterogeneous Agents with Local Social Influence Networks: Path Dependence and Plurality of Equilibria in the ACE Noiseless Case

Denis Phan

GEMAS UMR, 8598 CNRS & University Paris IV
Sorbonne CREM UMR, 6211 CNRS & University of Rennes 1, France.
dphan@msh-paris.fr

13.1 Introduction

In this paper we explore numerically by means of ACE the impact of local social influence on binary choices. The basic model of binary choices with externality presented here (the “GNP model”) is based on Gordon et al. (2005); Nadal et al. (2005); Phan and Pajot (2006) (see Phan and Semeshenko (2007) for an introduction and a review of literature). GNP model has been generalized to a large class of distributions in Gordon et al. (2006). It allows to study the collective behavior of a population of interacting heterogeneous agents. Numerous papers in this field concern homogeneous agents with stochastic choices, in particular, among others: Brock and Durlauf (2001)—hereafter BD model. Our GNP class of models differs by the nature of the disorder. The former belongs to the classes of Random Utility Models (RUM): the utility is stochastic. The individual preferences have an identical deterministic part and the heterogeneity across agents comes from the random term of the RUM. In our *noiseless* GNP model, agents are heterogeneous with respect to their idiosyncratic preferences (IWA) which remain fixed and do not contain stochastic term. This model belongs to the class of the *Quenched Random Field Ising Models*, known in statistical physics.

The question of the local topologies of interactions has been recently examined by Ioannides (2006). In the following, we present equilibria results for models with a local regular network (cyclical, one and two dimensional with nearest neighbor). This work is an extension of the GNP

model presented previously to the case with local interactions. Therefore, the reader is assumed to be familiar with these references. Several important aspects of the analysis and simulation of the model which are discussed in this paper are not presented here but are mentioned briefly for the sake of completeness. Section 13.2 introduces the GNP model and shows how this framework is related to the population games, by summarizing previous contributions (Phan and Semeshenko, 2007). Section 13.3 presents and compares both probabilistic calculi for *infinite size* population and ACE based simulations for *finite size* populations in the case of a simple *regular local influence network* (lattice). Calculi in Section 13.3.1 and 13.3.2 are based upon a probabilistic method recently introduced by Shukla (2000) to calculate exactly the hysteresis path both starting from a homogeneous state (nobody adopts) and from any arbitrary initial state. The simulations were conducted using the multi-agent platform “Moduleco-Madkit” (Gutknecht and Ferber, 2000; Phan, 2004). A special attention is devoted to Sethna’s inner hysteresis (Sethna et al., 1993). For a given value of the external parameter (i.e. price), there is a multiplicity of equilibria, depending on the previous state of the system (*path-dependence*). Moreover, if this parameter returns back to the initial value, the system returns precisely to the same state from which it left. The inner loop illustrates the *return point memory effect*, in which the system remembers its former state.

13.2 GNP framework with local setting

13.2.1 Modelling the individual choice in a social context

We consider a set of N agents $i \in \Lambda_N \equiv \{1, 2, \dots, N\}$ with a classical linear willingness-to-adopt function. Each agent makes a simple binary choice, either to adopt ($\omega_i = 1$) or not ($\omega_i = 0$) (e.g., to buy or not one unit of a single good on a market, to adopt or not the social behavior, etc.). A rational agent chooses ω_i in order to maximize its surplus:

$$\max_{\omega_i \in \{0,1\}} [\omega_i V_i(\tilde{\omega}_{-i})], \quad (13.1a)$$

$$\text{where: } V_i(\tilde{\omega}_{-i}) = (H_i - C) + \frac{J_{ik}}{N_{\vartheta_i}} \sum_{k \in \vartheta_i} \tilde{\omega}_k, \quad (13.1b)$$

C is the cost of adoption, assumed to be the same for all agents, and H_i represents the idiosyncratic preference component.

The cost of adoption be subjective or objective - it may e.g. represent the price of one unit of a good. Each agent i is influenced by the (expected) choices $\tilde{\omega}_k$ of its neighbors $k \in \vartheta_i$ within a neighborhood $\vartheta_i \in \Lambda_N$ of size N_{ϑ_i} . Denoting J_{ik}/N_{ϑ_i} the corresponding weight, i.e. the marginal social influence on agent i from the decision of agent $k \in \vartheta_i$, the social influence is then a weighted sum of $\tilde{\omega}_k$ choices. When the weights are assumed to be positive, $J_{ik} > 0$, it is possible, according to Brock and Durlauf (2001), to identify this external effect as *strategic complementarities* in the agents' choices.

In the GNP model agents are heterogeneous with respect to their idiosyncratic preferences, which remain fixed and do not contain additively stochastic term. The *Idiosyncratic Willingness to Adopt* (IWA) of each agent is distributed according to the Probability Density Function (pdf) $f_y(y)$ of the auxiliary centered random variable Y , such as H is the average IWA of the population:

$$H_i = H + Y_i, \quad (13.2a)$$

$$\text{where: } \lim_{N \rightarrow \infty} \frac{1}{N} \sum_N Y_i = 0 \Rightarrow \lim_{N \rightarrow \infty} \frac{1}{N} \sum_N H_i = H. \quad (13.2b)$$

If Y_i remains fixed, the resulting distribution of agents over the network of relations is a *quenched random field*: the agents' choices are purely deterministic. As mentioned before, this contrasts with the random utility approach in the BD model. These two approaches may lead to different behaviors (Galam, 1997; Sethna et al., 1993, 2005). One advantage of the GNP model is that it does not constrain the distribution of the idiosyncratic willingness to adopt to be a priori logistic. Moreover, the qualitative feature of the results may be generalized to a large class of distributions (Gordon et al., 2006). We can assume hereafter without loss of generality that the idiosyncratic preferences are distributed according to a bounded, triangular pdf. This allows the analytical exact determination of the equilibrium properties in the case of complete connectivity (Phan and Semeshenko, 2007). In the following, we restrict to the case of regular nearest neighborhood, cyclical network of dimension one (circle, with $N_{\vartheta_i} = 2$) and two dimension (torus, with $N_{\vartheta_i} = 4$, von Neuman's neighborhood)). Moreover, for the sake of simplicity, we restrict to the case of *positive homogeneous influences*: $\forall i \in \Lambda_N, \forall k \in \vartheta_i : J_{ik} = J > 0$. For a given neighbor k the social influence is J/N_{ϑ_i} if the neighbor is an adopter ($\omega_k = 1$), and zero otherwise. Let η_i^e be i 's expected adoption rate within the neighborhood

$$\eta_i^e \equiv \eta_i^e(\tilde{\omega}_{-i}) \equiv \frac{1}{N_{\vartheta_i}} \sum_{k \in \vartheta_i} \tilde{\omega}_k. \quad (13.3)$$

With these assumptions the surplus of agent i if he adopted is: $H_i - C + J\eta_i^e$. The conditional probability of adoption, for a given η_i^e is:

$$P(\omega_i = 1 | \eta_i^e) = P(H_i > C - J\eta_i^e) \quad (13.4)$$

13.2.2 Interactions in the neighborhood as a population game

The interest of studying such idiosyncratic (exogenous) heterogeneity becomes clear if one reinterprets the GNP model within a game theoretic framework. Each agent i has only *two possible strategies*: to adopt ($\omega_i = 1$) or not adopt ($\omega_i = 0$). In the following, we assume agents have myopic expectations about the behavior of their neighbors: $\tilde{\omega}_{-i}(t) = \omega_{-i}(t-1) \equiv \omega_{-i}$: then $\eta_i^e(t) = \eta_i(t)$. The best response of an agent playing against its neighbors is formally equivalent to that of an agent playing against a *Neighborhood Representative Player* (NR) (Phan and Pajot, 2006; Phan and Semeshenko, 2007). *NR* player in turn plays a *mixed strategy* $\omega_{nr} = \eta_i \in [0, 1]$. In the present case of *finite neighborhood local interaction*, ω_{nr} takes its value in a discrete subset of $[0, 1]$. For example for $N_{\vartheta_i} = 2$, we have $\omega_{nr} = \eta_i \in \{0, 1/2, 1\}$ and for $N_{\vartheta_i} = 4$, we have $\omega_{nr} = \eta_i \in \{0, 1/4, 1/2, 3/4, 1\}$. The “*normal form*” payoff matrix $G1$ gives the total payoff for an agent i playing against this *fictitious* NR player. According to Monderer and Shapley (1996), the best-reply sets and dominance-orderings of the game $G1$ are unaffected if a constant term is added to a column (i.e. $C - H_i$). The coordination game matrix $G2$ in Table 13.1 (right) is said to be “best reply equivalent” to the matrix $G1$ of Table 13.1 (left). However, the values in $G2$ do not indicate the cumulated payoffs, contrary to the values in $G1$, but are a direct measure of the cost—the risk in the sense of Harsanyi and Selten (1998)—of a unilateral deviation from the coordinated solution ($\omega_i = \omega_{nr}$) in the case of the pure strategy framework.

Figure 13.1 (right) presents a (symmetric triangular) distribution and related best reply for a given cost C and a particular value of the IWA. If $C - J > H_i$, then *never adopt* ($\omega_i = 0$) is the strictly dominant strategy for all possible values of η_i (agents of type (0) in the light grey zone on the left). If $H_i > C$, *always adopt* ($\omega_i = 1$) is the strictly dominant strategy for all possible values of η_i^e (agents of type (1) in the dark grey zone on the right). If $C > H_i > C - J$ then the agent’s virtual

a-game $G1$	$\omega_{nr} = 0$	$\omega_{nr} = 1$
$\omega_i = 0$	0	0
$\omega_i = 1$	$H_i - C$	$H_i - C + J$

b-game $G2$	$\omega_{nr} = 0$	$\omega_{nr} = 1$
$\omega_i = 0$	$C - H_i$	0
$\omega_i = 1$	0	$H_i - C + J$

Table 13.1. Payoff matrix for an agent i (left) and best reply equivalent potential game (right). Player i in rows, fictitious NR Player—indexed nr —in columns.

surplus $V_i \equiv H_i - C + J\eta_i$ may be either positive or negative depending on the rate of adoption within the neighborhood η_i . These agents are *conditional adopters* and said to be of type (2). Within these agents, only those with $V_i > 0$ will adopt thanks to the social influence (hashed region). The relevant economic cases are the ones with (at least some) agents of type (2).

Figure 13.1 (left) exhibits a distribution of agents' type in the space $(J, H - C)$ for the symmetric triangular distribution on the interval $[-a, a]$. In the south-west light grey zone there are only agents of type (0), while in the north-dark-zone there are only agents of type (1). In the white zone there is a mixture of at least 2 types of agents, with necessarily some agents of type (2). If $H - C > a - J$, there is no agents of type (0). Conversely in the south zone, where $H - C < -a$, there is no agent of type (1). If both conditions hold then *all agents* are of type (2), corresponding to the hashed triangular zone in the east on Figure 13.2. This implies a sufficiently strength intensity of social effect, with respect to the dispersion of the preferences, that needs to be relatively moderate: $J > 2a$.




Fig. 13.1. Distribution of agents with respect to their type, on the pdf (right) and in the space $(J, H - C)$ (left) for the symmetric triangular distribution on the interval $[-a, a]$. Source: Phan and Semeshenko (2007).

In the case of global social influence (full connectivity) and bounded distribution of IWA, dominance-ordering analysis allows us to predict the issues of some classic configurations (i.e. symmetric Nash equilibrium), where all agents have the same structure of best reply. But this may be done for only some special cases. In general situations, one needs to use the methods from statistical mechanics. In the case of local neighborhood considered here, any simple result of that kind is available and probabilistic approach or numerical simulations are required. Within an approach of ACE as *complement* of traditional mathematical models, in the next section we compare results using the probabilistic approach for *infinite size* population with results from agent based simulation for *finite size* population.

13.3 Collective Behavior, hysteresis and local frozen domains with local externality: probabilistic approach for *infinite size* population and simulation approach for *finite size* population

As suggested before, the GNP model, as a socio-economic version of Quenched RFIM model, has some significantly different properties with respect to the BD model. First, when changing the external field (i.e. cost variation) an equilibrium depends on the previous equilibrium, but does not depend on the order in which the agents change their behaviour (i.e. adoption or not) through the process called avalanche. In other words, from the simulation point of view, both parallel and sequential updating lead the system to the same equilibrium. Second, the interesting property of Sethna's inner hysteresis phenomenon (Sethna et al., 1993) can be observed. This result from the *return point memory effect*: starting from an equilibrium state, if we change the cost by a given value and reverse the field by the same value, the system remembers its former state and returns exactly to the equilibrium point of departure. The corresponding trajectory is called "inner loop" "minor hysteresis" (Sethna et al., 2005). Finally, in the special case where we change the cost monotonically for a homogeneous state (everybody adopts or no-body adopts) the final equilibrium does not depend on the rate of variation in cost. A dramatic change from $C1$ to $C2$ or a succession of smaller monotonic changes from $C1$ to $C2$ leads to the same state. In this section, we experiment the effect of local social influence in discrete choice adoption process based on the GNP model by means of finite size population, agent-based simulation on the multi-agent platform "Moduleco-Madkit" (Gutknecht and Ferber, 2000; Phan, 2004).

Sections 13.3.1 and 13.3.2 compare analytical results with simulated outcome in the case of the cyclical one-dimensional nearest neighborhood network (circle). Section 3.2. is devoted to the calculus of the inner loop. Section 3.3. presents simulation outcome in the case of the cyclical two-dimensional regular network (von Neuman neighborhood on a torus).

13.3.1 Starting from a homogeneous state without adoption and going to the complete adoption and return: the larger hysteresis loop

Hysteresis within the Quenched RFIM is somewhat different in nature from the hysteresis used by economists that arise from a delayed response of a system (time lags) to a change in the external parameter (here cost). First accounts of such difference are (Amable et al., 1994) for the study of the wage-price spiral and zero-rot dynamics. Previous application of hysteresis in Quenched RFIM in socio-economic models are Galam (1997); Phan et al. (2004). In the case of a finite population, there are a very large number of equilibria and related thresholds between them. In this section, we use methodology and results from physics (Shukla, 2000) established for the ferromagnetic case ($J > 0$) for one dimensional, nearest neighborhood, cyclical and infinite size network. In that case the conditional probability of adoption (equation 13.4) can be expressed in a finite number of occurrences; see Figure 13.2) and relations (13.5)




Fig. 13.2. Agents' choices with respect to their IWA and neighborhood state (symmetric triangular distribution).

Probability of adoption for a given state of neighborhood $\eta_i \in \{0, 1/2, 1\}$ with $P_i \equiv P_i(C)$

$$P_0 = P(H_i > C) = P(\text{type } 1, \omega_i = 1 \text{ if } \eta = 0) \quad (13.5a)$$

$$P_1 = P(H_i > C - J/2) = P(\omega_i = 1 \text{ if } \eta = 1/2) \quad (13.5b)$$

$$P_2 = P(H_i > C - J) = P(\omega_i = 1 \text{ if } \eta = 1) \quad (13.5c)$$

$$\begin{aligned} P_1 - P_0(C) &= P(C > H_i > C - J/2) \\ &= P(\text{to be of type } 2 \text{ AND } \omega_i = 1 \text{ if } \eta = 1/2). \end{aligned} \quad (13.5d)$$

For a given cost C , the probability of adoption of an agent is:

$$P(\omega_i = 1|C) = P(\eta_i = 1) P_2 + P(\eta_i = 1/2) P_1 + P(\eta_i = 0) P_0 \quad (13.6)$$

where:

$$\eta(C) = P(\omega_i = 1|C) \quad (13.7a)$$

$$P(\eta_i = 1) = P(\omega_{i\pm 1} = 1|C, \omega_i = 0)^2 \quad (13.7b)$$

$$P(\eta_i = 1/2) = 2 P(\omega_{i\pm 1} = 1|C, \omega_i = 0) P(\omega_{i\pm 1} = 0|C, \omega_i = 0) \quad (13.7c)$$

$$P(\eta_i = 0) = P(\omega_{i\pm 1} = 0|C, \omega_i = 0)^2. \quad (13.7d)$$

Note that, $P(\omega_{i\pm 1} = 1|\omega_i = 0) \equiv P^*(C)$ can be calculated exactly in the infinite case. The probability that my neighbor adopts before me is equal to P_0 (the probability of type (1), then adopts even if no neighbor has adopted before). One must add again the probability for my neighbor to be of type (2) but to adopt as soon as the next neighbor has adopted, since this next agent ($\omega_{i\pm 2}$) is of type (1). The corresponding joint probability is equal to $[P_1 - P_0]P_0$. At the level 3, one must add again the probability for my neighbor and the next agent to be of type (2) but to adopt as soon as the next neighbor has adopted, given the probability that this next agent ($\omega_{i\pm 3}$) is of type (1). This joint probability is equal to $[P_1 - P_0]^2 P_0$, and so on. Summing over all cases:

$$\begin{aligned} P^*(C) &\equiv P(\omega_{i\pm 1} = 1|C, \omega_i = 0) \\ &= \lim_{m \rightarrow \infty} P_0 \sum_{k=0}^m [P_1 - P_0]^k = \frac{P_0}{1 - [P_1 - P_0]} \end{aligned} \quad (13.8a)$$

$$[1 - P^*(C)] \equiv P(\omega_{i\pm 1} = 0|C, \omega_i = 0) = \frac{1 - P_1}{1 - [P_1 - P_0]} \quad (13.8b)$$




Fig. 13.3. Theoretic (line) and simulated (dot) values for the main hysteresis with: $N = 1156$ agents, $N_\theta = 2$ (circle), $J = 4$, $H = 0$.

Using equations (13.6), (13.8a), and (13.8b) the global equilibrium rate of adoption in the population for a given cost C is equal to the probability of adoption of an agent taken at random within a symmetric triangular distribution of IWA:

$$\eta_+(C) = P(\omega_i = 1|C) \quad (13.9)$$

The upper half branch of the main hysteresis, for decreasing C from complete adoption to zero can be obtained by symmetry: $\eta_-(C) \sim -\eta_+(-C)$. Figure 13.3 provides a comparison between these theoretic values of the main hysteresis and the simulated ones, based on experiments with finite population (here 1156 agents).

13.3.2 The inner hysteresis loop: reversing the Cost from an arbitrary point on the exterior loop

In the limit of quasi-static driving (the change in prices remains constant within an avalanche), starting from a point on the upstream trajectory (grey) for $\eta = 40\%$ and $C = 1.25$ a backtracking increase in cost C induces a less than proportional decrease (avalanche) in the number of customers (black curve, upper inner loop) until $C = 2.49$ and $\eta = 30\%$. Then after reversing the cost changes at $C = 2.49$, as the cost decreases back to the initial value (grey curve, lower inner loop) the system returns precisely to the same state from which it left the outer loop ($C = 1.25$, $\eta = 40\%$). The inner loop can also go from a

branch of the main hysteresis loop to the other. For example, starting at $C = 1$ and $\eta = 68\%$, a backtracking increase in cost C induces a cross-trajectory between the upstream and the downstream branch of the main hysteresis loop. This cross-trajectory finishes at $C = 2.93$ and $\eta = 30\%$, when the equilibrium points are those of the main hysteresis. As established analytically, that confirms that there is a multiplicity of equilibria, depending on the previous state of the system (*path dependence*). Figure 13.4 (right panel) exhibits separated homogeneous domains (or cluster) in the network, due to the dominance of positive or negative effects of social influence as well as a particular distribution of heterogeneous IWA, enforced by both locality and finite size effect.




Fig. 13.4. Right panel, Sethna’s inner hysteresis $J = 4$, $N = 2$ (circle). Left panel, homogeneous domains (1D-clusters) within the network for $\eta = 40\%$ and $C = 1.27$.

As previously mentioned, it is possible to provide some hints to calculate the probability of adoption starting from an arbitrary point on the exterior loop. The method used here follows the lines of Shukla (2000). For the reversing formula and complete calculations in the case of the symmetric triangular pdf, see the long version of this work (to be presented at CEF 2007). This calculus is more difficult than in the previous case, because the choice of adoption depends now in a non trivial way on the rate of adoption in the neighborhood, which depends itself directly or indirectly on the state of the other agents over the network. The probability of adoption between the two branches of the external hysteresis is conditional to the cost C for which the backtrack starts. These analytical results fit correctly the numerical simulations in

the case of finite population experiments (see the long version for CEF 2007)). For a given cost C' , with backtracking at C , the probability of adoption of an agent is:

$$P(\omega_i = 1|C, C') = P(\omega_i = 1, C) - Q_2(C', C) + Q_1(C', C) - Q_0(C', C), \quad (13.10)$$

where:

$$\eta(C', C) = P(\omega_i = 1|C, C') \quad (13.11a)$$

$$Q_2(C', C) = P^*(C)^2 (P_2(C) - P_2(C')) \quad (13.11b)$$

$$Q_1(C', C) = 2 P^*(C) [Q_a(C', C) + Q_b(C', C)] (P_1(C) - P_1(C'))$$

$$Q_0(C', C) = [Q_a(C', C) + Q_b(C', C)]^2 (P_0(C) - P_0(C')) \quad (13.11c)$$

$$Q_a(C', C) = \frac{P^*(C) (1 - P_2(C)) + [1 - P^*(C)] (1 - P_1(C))}{1 - (P_1(C) - P_1(C'))} \quad (13.11d)$$

$$Q_b(C', C) = \frac{P^*(C) (P_2(C) - P_2(C'))}{1 - (P_1(C) - P_1(C'))} \quad (13.11e)$$

13.3.3 The two-dimensional von Neuman neighborhood network (Torus)

There is no analytical result at this time for the two-dimensional, cyclical network with von Neuman neighborhood (Torus). But the example on Figure 13.5 suggests that both Sethna's inner loop and homogeneous domains remain quite similar.

13.4 Conclusion

Using methods from statistical physics, we illustrated the stationary properties for particular cases of symmetric triangular distribution of IWA in the presence of local interactions. Simulation results allow us to observe numerous complex configurations on the adoption side, such as hysteresis and Sethnas inner-loop hysteresis. This complex social phenomenon depends significantly on the structure and parameters of the relevant network. Finally, the last section opens the question of finite size effects, also addressed by Glaeser and Scheinkman (2002); Krauth (2006) among others. The preliminary results in the case of a simple, regular network suggest new fields of investigation, as opposed to a standard focus on conditions of uniqueness of equilibrium, under a




Fig. 13.5. Right panel, Sethna's inner hysteresis $J = 4$, $N = 2$ (torus). Left panel, homogeneous domains (2D-clusters) within the network for $\eta = 40\%$ and $C = 1.41$.

moderate social influence assumption (Glaeser and Scheinkman, 2002). It would be interesting to compare more systematically the analytical predictions against the simulation results and to study the statistical properties of such a phenomenon for different values of J and different network's structure.

Acknowledgement. The author acknowledges M. B. Gordon, J.P. Nadal, and V. Semeshenko for their significant intellectual contribution to the development of methods used in this paper (see bibliography), to C. Deissenberg for valuable remarks on a previous version of this work and to Alexandra Frénoel for significant corrections on the original draft, D.P. is CNRS member.

References

- B. Amable, J. Henry, F. Lordon, and R. Topol. Strong hysteresis versus zero-root dynamics. *Economics Letters*, 44(1–2):43–47, 1994.
- W. A. Brock and S. N. Durlauf. Discrete choice with social interactions. *Review of Economic Studies*, 68(2):235–260, April 2001.
- S. Galam. Rational group decision making: A random field ising model at $t = 0$. *Physica A: Statistical and Theoretical Physics*, 238(1–4): 66–80, April 1997.
- E. L. Glaeser and J. A. Scheinkman. Non-market interactions. In Turnovsky. Eds. Dewatripont, Hansen., editor, *Advances in Economics and Econometrics: Theory and Applications*. Cambridge University Press, 2002.

- M. B. Gordon, J. P. Nadal, D. Phan, and J. Vannimenus. Seller's dilemma due to social interactions between customers. *Physica A: Statistical Mechanics and its Applications*, 356:628–640, October 2005.
- M. B. Gordon, J. P. Nadal, D. Phan, and V. Semeshenko. Discrete choices under social influence: Generic properties. In *Proceeding of WEHIA 2006*, 2006.
- O. Gutknecht and J. Ferber. Madkit: a generic multi-agent platform. In *AGENTS '00: Proceedings of the Fourth International Conference on Autonomous Agents*, pages 78–79, New York, NY, USA, 2000. ACM Press.
- J.C. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, 1998.
- Y. M. Ioannides. Topologies of social interactions. *Economic Theory*, 28(3):559–584, 2006.
- B. Krauth. Social interactions in small groups. *Canadian Journal of Economics*, 39(2):414–433, May 2006.
- D. Monderer and L.S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, May 1996.
- J. P. Nadal, D. Phan, M. B. Gordon, and J. Vannimenus. Multiple equilibria in a monopoly market with heterogeneous agents and externalities. *Quantitative Finance*, 5(6):557–568, December 2005.
- D. Phan. *From Agent-Based Computational Economics Toward Cognitive Economics*, chapter 22, pages 371–398. Springer, Berlin, 2004.
- D. Phan and S. Pajot. Complex behaviours in binary choice models with global or local social influence. In C. Bruun, editor, *Advances in Artificial Economics*, pages 203–220. Springer, 2006.
- D. Phan and V. Semeshenko. Equilibria in models of binary choice with heterogeneous agents and social influence. *European Journal of Economic and Social Systems*, 2007. revised version accepted.
- D. Phan, M.B. Gordon, and J.P. Nadal. *Social Interactions in Economic Theory: an Insight from Statistical Mechanics*, chapter 20, pages 335–358. Springer, Berlin, 2004.
- J. P. Sethna, K. Dahmen, S. Kartha, J. A. Krumhansl, B. W. Roberts, and J. D. Shore. Hysteresis and hierarchies: Dynamics of disorder-driven first-order phase transformations. *Physical Review Letters*, 70: 3347–3350, 1993.
- J. P. Sethna, K. A. Dahmen, and O. Perkovic. *Random-Field Ising Models of Hysteresis*, volume 2, chapter 2. Elsevier, Amsterdam, 2005.

- P. Shukla. Exact solution of return hysteresis loops in one dimensional random field ising model at zero temperature. *Physical Review E*, 62 (4):4725–4729, 2000.

Economy-Driven Shaping of Social Networks and Emerging Class Behaviors

Philippe Caillou, Frederic Dubut, and Michele Sebag

LRI, Université Paris Sud, France
{caillou,dubut,sebag}@lri.fr

14.1 Introduction

Agent-based Computational Economics (ACE) is a powerful framework for studying emergent complex systems resulting from the interactions of agents either mildly rational, or with incomplete information (Axelrod, 2004; Tesfatsion, 2002, 2006) or driven by the social network (Bala and Goyal, 2003; Carayol and Roux, 2004; Slikker and van den Nouweland, 2000).

This paper focuses on the interrelationship between social networks and economic activities. Compared to the state of the art, the main originality is that the social network dynamically evolves based on the rational decisions of agents: a loan granting activity is enabled by the network and the agents continuously re-shape the network to optimize their utility. Three classes of agents (rational agents, free riders and “investors”) are considered. The global welfare is investigated in relation with the agent diversity, examining the differential advantages/disadvantages of the agent classes depending on their distribution in the agent population. Lastly, the stability of e.g. the average interest rate is contrasted with the instability of the network structure.

After describing the state of the art and the problem tackled in this paper (section 14.2), we present a loan granting game played by a society of rational agents with a long-term utility function, conditioned by and shaping their social network (section 14.3). Section 14.4 reports on the simulation results; the main contributions of the approach are discussed together with perspectives for further studies in section 14.5.

14.2 State of the Art and Goal of the Study

Introduced by Epstein and Axtell (1996), Agent-based Computational Economics established two major results (the interested reader is referred to Tesfatsion (2002, 2006) for a comprehensive presentation). Firstly, the lack of a centralized walrasian auctioneer does not prevent a society of 0th-intelligence agents from converging towards an economical equilibrium when agents interact and exchange in a decentralized manner; secondly, this result does not hold any longer if the agents can die or evolve.

Meanwhile, after the pioneering Milgram (1967)'s experiment and many further studies (e.g. Watts and Strogatz, 1998), the structure of social networks is acknowledged a major factor of economics efficiency. A framework for analyzing social network economics was defined by Jackson and Wolinsky (1996), and exploited through either analytical approaches, or various simulation-based extensions (Bala and Goyal, 2003; Carayol and Roux, 2004; Slikker and van den Nouweland, 2000).

While both domains of ACE and social networks are clearly related, to our best knowledge little attention has been devoted to the interdependent evolution of social and economical activities, considering the social network as both a result and an enabling support of the economical activity. In such a unified perspective, the stress is put on the complex system emerging through the interaction of social and economical activities. Along this line, this paper investigates the complex system made of a population of agents engaged in a loan granting activity, where the activity is simultaneously conditioned by, and shaping, the social network. Basically, every agent is endowed with an individual utility function parameterized after its fixed preference toward immediate rewards; it accordingly decides between borrowing or lending money from/to its neighbors at every time step. While the network thus governs the instant rational optimization problem faced by the agents, agents can decide to create/delete links and thereby modify the network. This setting contrasts with former studies (Bala and Goyal, 2003; Carayol and Roux, 2004; Jackson and Wolinsky, 1996; Slikker and van den Nouweland, 2000) modeling the social network as the end of the socio-economic game, that is, where the network only supports the exchange of information and agents are assessed based on their position in the network.

Furthermore, agents will not reveal their preference – as opposed to e.g. Epstein and Axtell (1996) where the exchange price is based on the preferences of both agents. The fact that agents do not reveal their preference is relevant to the study of socio-economic games in two

respects; firstly it is more realistic from a non-cooperative game perspective; secondly, the incompleteness of information might adversely affect the convergence of the game.

Finally, the study examines the impact of the agent models and strategies on the global welfare in a long term perspective. This contrasts with e.g. Jackson and Wolinsky (1996) focussing on the immediate network efficiency, and discarding the long term impact of current decisions.

14.3 Overview

This section presents the agent model, the interaction setting and the observed variables of the system. Due to space limitations, the reader is referred to Caillou et al. (2007) for details.

14.3.1 Agent model

The agent utility function models the intertemporal choice of the consumer after the standard economic theory (Fisher, 1930). Formally, agent A_i maximizes the sum over all time steps of its weighted instant utilities. The utility weight at time t , set to p_i^t ($0 < p_i < 1$), reflects the agent preference toward the present (parameter p_i). The instant utility reflects the current consumption level $C_{i,t}$, with a diminishing marginal utility modeled through parameter b_i ($0 < b_i < 1$), standing for the fact that the agent satisfaction is sublinear with its consumption level (Menger, 1871). Letting M_i denote the lifelength of agent A_i , it comes:

$$U_i = \sum_{t=0}^{M_i} (p_i^t C_{i,t}^{b_i}) \quad (14.1)$$

The instant neighborhood of agent A_i , noted $V_{i,t}$ involves all agents A_j such that link (i, j) belongs to the social network at time t . Additional agent parameters comprise:

- **Salary** R_i : A_i receives a fixed salary R_i at the beginning of each time step, and uses it to grant or pay back loans, to buy links, or for consumption.
- **Sociability factor** s_i ($0 < s_i < 1$): A_i creates a new link (i, j) (where j is uniformly chosen) with a probability s_i at each time step; in case A_i is isolated, a new link is automatically created.

- **Strategy S_i :** The social network comes at a cost, i.e. every link (i, j) must be paid by agents A_i or A_j or both. Three social strategies (classes of agents) are defined:

Optimizers accept to pay for a link iff it was profitable during the last five time steps (if the utility increase due to this link offsets the link cost). This strategy, referred to as rational strategy, deletes all links which are not sufficiently useful.

Free Riders never pay for a link. While the free rider minimizes its social cost (the link cost), it does not optimize its neighborhood which might adversely affect its utility (see below).

Investors always accept to pay for a link. On the one hand this strategy gives the agent every means to optimize its economic activities, and possibly maintain beneficial relations with isolated agents; on the other hand, it suffers the cost of possibly many useless links.

Under mild assumptions (Caillou et al., 2007), agent A_i can compute its threshold interest rate r_i (lower bound for grant activities and upper bound for loan activities). Note that this rate needs be updated after every elementary transaction as it depends on the agent current and expected capital.

14.3.2 Interaction protocol

Every agent lives a sequence of epochs, where each epoch involves four phases: i) salary and loans payback, ii) negotiation, iii) consumption, iv) social activity (link creation/deletion).

During the first phase, agent A_i receives its salary R_i , reimburses the money borrowed (plus interests) and is reimbursed for the money lent (plus interests). The negotiation phase involves a variable number of transactions. At each step, A_i determines the best possible borrowing and lending rate; it maintains its estimation $r_{i,j}$ of the interest rate for a transaction (borrow or grant) with every agent A_j in its neighborhood, and proposes the best possible transaction for one currency unit. Depending on whether the transaction is accepted, $r_{i,j}$ is updated (Alg. 1). Agent A_j accepts a borrow transaction if the proposed rate is lower than i) its limit rate r_j and ii) its last borrow rates during this negotiation phase (similar conditions hold for lend transactions).

The transactions proceed until no more transactions are realized.

```

BestRate  $r^*=0$ ;
foreach  $A_j \in V_i$  such that  $r_{ij} > r_i$  do
    Propose Loan(rate= $r_{ij}$ );
    if accepted then
        if  $r_{ij} > r^*$  then  $r^* = r_{ij}$ ;
        Increase( $r_{ij}$ )
    else
        Decrease( $r_{ij}$ )
    end
end
if  $r^* > 0$  then Lend one currency unit at rate  $r^*$ 

```

Algorithm 1: Lending transactions (borrowing transactions proceed likewise)

During the consumption phase, the agent computes its optimal fraction of consumption (see Caillou et al., 2007, and scores the corresponding utility) .

During the social phase, each agent decides whether it maintains its links depending on its strategy and whether the link has been profitable in the last five epochs. Link (i, j) is either maintained by agents A_i and/or A_j , or deleted. Independently, A_i creates a new link (i, j) with probability s_i (its sociability factor), where j is uniformly randomly selected. If A_i has no neighbor, a link (i, j) is automatically created.

After M_i epochs, agent A_i dies. It is then replaced by a new agent (reinitializing all agent parameters) *with same neighborhood*.

14.3.3 Fitness and Global Welfare

The socio-economical system will be assessed from the global welfare of the agents. As the agent utilities cannot be directly compared (parameters p_i and b_i depend on the agent), they are normalized w.r.t. the canonical consumer-only alternative strategy. Each A_i , would it have adopted the consumer-only strategy, would get utility The consumer-only agent, spending its whole salary in each time step, gets utility:

$$U_i^* = \sum_{t=0}^{M_i} (p_i^t R_i^{b_i}) = R_i^{b_i} \frac{1 - p_i^{M_i+1}}{1 - p_i}$$

Accordingly, the normalized fitness of A_i is defined as:

$$F_i = \left(\frac{U_i}{U_i^*} \right)^{\frac{1}{b_i}} - 1$$

Note that if A_i had spent a fixed fraction α , $0 \leq \alpha \leq 1$ of its salary in each time step (without engaging in any borrowing or lending transactions), it would score a normalized fitness $\alpha - 1$. In brief, agent A_i benefits from the social network iff its fitness F_i is positive.

The efficiency of the socio-economical system is thus measured from the average normalized fitness of the individuals, and its standard deviation. Further, each class (optimizers, free-riders and investors) will also be assessed from the average normalized fitness of the individuals belonging to this class.

14.4 Results

After the description of the experimental setting, this section reports on the impact of the network and agent dynamics on the global efficiency of the system.

14.4.1 Experimental settings

The socio-economical game is implemented and simulated within the Moduleco framework (Phan, 2004). The initial structure of the social network is a ring, where each agent is connected to its two neighbors. Agents are initialized by independently drawing their parameters using Gaussian or uniform laws as follows.

- Time preference $p_i \sim \mathcal{N}(0.8, 0.075)$
- Utility factor $b_i \sim \mathcal{N}(0.5, 0.1)$,
- Sociability factor $s_i \sim \mathcal{N}(0.05, 0.05)$,
- Salary $R_i \sim \mathcal{N}(20, 5)$,
- Life expectancy $M_i \sim U(20, 100)$,

The link cost is set to .2 in the remainder of the paper. Complementary experiments with varying values of the link cost are reported in Caillou et al. (2007). Experiments were conducted with a population size ranging from 25 to 100, with similar results. All reported results are averaged over 25 independent experiments conducted with 25 agents over 1000 epochs. The global (respectively, class) fitness is computed by averaging the normalized fitness of agents (resp. belonging to the class) that died before the 1000th. epoch.

14.4.2 Complete and costless information

The classical economic theory relies on the assumption of a complete and costless information, e.g. gathered and disseminated by the “wal-rasian auctioneer”, enforcing the convergence of the interest rate toward the equilibrium rate. As formally shown in Caillou et al. (2007), the equilibrium rate can be analytically derived from the agent utility functions.

The first experiment, as a sanity check, thus considers the fully connected social network and compares the empirical interest rate toward the equilibrium rate. As expected, the average interest rate rapidly converges toward the equilibrium value τ ($\tau = .24$ in the experimental setting, Fig. 14.1). The standard deviation ($< .0025$ after 5 epochs in the fully connected case) is explained from the experimental noise, discrete loan amount and limited number of agents.

Interestingly, randomly removing edges in the social network only delays the convergence toward the equilibrium rate, although the standard deviation of the interest rate significantly increases for social networks with low density. Fig. 14.1 displays the standard deviation vs the percentage of edges in the social network after 20 epochs.




Fig. 14.1. Interest rate and standard deviation within a complete network

14.4.3 Rational and immortal agents

The second experiment focuses on rational agents (optimizer strategy) with infinite lifelength. Despite the fact that the social network can evolve with the rational agent decisions, the empirical interest rate still




Fig. 14.2. Impact of the network connectivity on the standard deviation of interest rate after 20 epochs

converges toward the equilibrium rate. Still, the convergence is slower than in the previous case, and the standard deviation remains high after 1000 epochs (Fig. 14.3).




Fig. 14.3. Rational and Immortal Agents: Standard Deviation of the Interest Rate

Most surprisingly, while the interest rate reaches the equilibrium, it does so with a continuously changing social network; no edge in the network appears to last more than a few epochs, as agents endlessly

optimize their neighborhood. Indeed, in either competitive or monopolistic situations, there always exists some profitable link creation or deletion. The canonical case of a 3-agent network, depicted in Fig. 14.4, involves three possible configurations (being reminded that a link is created automatically when an agent is isolated), all of which are unstable. In the n -agent case, instability is increased by cascading effects, the creation/deletion of a link leading to further link deletions or creations.




Fig. 14.4. 3-agent network configurations; arrows indicate the lender-borrower pair

- 1 In the triangle (clique) case, there are two possible situations:
 - If A can lend the desired amount to C with a rate higher than the B limit rate r_b , A is not interested in maintaining link AB, which will thus be deleted.
 - If the C limit rate r_c decreases (because of the loans contracted by C) and becomes lower than r_b , A will grant loans to B and C (with a rate lower than r_b). Therefore B will be unable to lend money to C, since C will refuse to borrow money with rate higher than r_b . Thus the BC link becomes useless and will be deleted.
- 2 In the line case, B borrows from A at rate $\tau_{AB} < r_B$. C borrows from B at rate $\tau_{BC} > r_B$. When A or C will create the link AC, it will be stable because A will accept to grant loans to C at a rate $\tau_{BC} - \epsilon$ which will be higher (and thus more profitable) than τ_{AB} . We are back to case 1.
- 3 In the star case (case 3a), agent A is the only one lending money. In this monopolistic situation, A will progressively increase the loan granting rate, until the BC link becomes profitable and thus stable when it will be created. We are back to case 1. Same analysis holds for case 3b (C will decrease its borrow rate).

14.4.4 Mixed populations and global welfare

Let us consider the mixed population cases. Each possible distribution of the strategies (or classes) in the population is represented as a point

in the 2D plane, where the x (resp. y) coordinate stands for the proportion of Free Riders (resp. Investors) in the population (Fig. 14.5). Point (x, y) is associated with the global population welfare or fitness. Considering the three pure strategies (optimizers only, $(0,0)$; investors only $(0,1)$; free-riders only $(1,0)$), the Optimizer strategy is by far the best one; this fact is explained as the free-rider-only population generates a sparse random network, while the investor-only population generates a clique (and pays the price for it).




Fig. 14.5. Average Fitness vs Strategy Distribution in the Population. Point (x, y) correspond to the distribution of $x\%$ free-riders, $y\%$ investors and $1 - x - y$ % optimizers. Each fitness level differs by .5 from the neighbor fitness levels.

Most interestingly, in the case of mortal agents, mixed populations outperform optimizer-only populations; e.g. the uniform distribution ($1/3$ optimizers, $1/3$ free-riders and $1/3$ investors) gets an average fitness significantly higher than the optimizer-only population (complementary experiments show that this result also holds when the link cost is significantly higher or lower, see Caillou et al. (2007)). This fact is explained as the useless links paid for by investors, are actually very useful to quickly reorganize the network when an agent dies.

14.4.5 Class behaviors

While agent diversity is beneficial on average, it remains to examine whether “class behaviors” appear in the population. Two issues will be specifically investigated. Firstly, do the average class fitness depend

on its representativity, ie, its percentage of the population (intra-class effects); secondly, do the class representativity affect the average fitness of the other classes (inter-class effects).




Fig. 14.6. Strategy average fitness; each level represents an increase/decrease of 1 for the strategy average fitness

The free-rider class does present a class behavior (Fig. 14.6): their fitness is excellent when they are a minority, and it decreases rapidly as the free-rider representativity increases. The investor class also displays a class behavior: their fitness is good when the representativity of optimizers is sufficiently high, as the optimizers share the costs of the links. In summary, the average fitness of investors and free-riders depends on the class distribution; quite the opposite, the optimizer fitness does not.

We finally examine the impact on the global and class welfare, of the arrival of a new agent depending on its class (Fig. 14.7).




Fig. 14.7. Adding a new agent: impacts on the fitness of the other classes

The arrival of an investor is globally beneficial to other classes. and even more so when there a few optimizers. Indeed, investors fund the

infrastructure used by the other agents; their impact is greater when the network is poor (when there are few optimizers).

The arrival of a free rider has a negative impact on the rest of the economy, except when there are many investors. In this case, the free rider decreases the graph connectivity level (and cost) and allows a faster reorganization of the network.

Optimizers, which are not influenced by economic class structure, do not influence it either. Their impact is mostly neutral, though it might be positive in the extreme cases where there are no investors or no free riders.

14.5 Conclusion and perspectives

The socio-economic game presented in this paper, based on autonomous and diversified agents, leads to two lessons. The first one concerns the fact that the convergence of economic macro-variables such as the interest rate is compatible with the instability of the social network supporting the economic activities. Secondly, the benefits of the population diversity have been empirically demonstrated and interpreted in terms of the emerging class behaviors. While investors and free-riders display class behaviors (their fitness depends on the population structure, they have an impact on the welfare of the other classes), the optimizer class seems to be almost unaffected by the population structure, and exerts little influence in return.

Further research perspectives are concerned with the dynamics of the class structure, examining how the best fit agents can influence the distribution, preferences and strategies of the new-born agents.

Acknowledgement. We warmly thank D. Phan for making available his Moduleco Framework. This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778.

References

- R. Axelrod. Advancing the art of simulation in the social sciences. *Advances in Complex Systems*, 7(1):77–92, 2004.
- Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Japanese Journal for Management Information Systems*, 12(3), 2003.

- P. Caillou, F. Dubut, and M. Sebag. Loans in social networks: when sub-optimality increases global welfare. Technical report, LRI, 2007.
- N. Carayol and P. Roux. Behavioral foundations and equilibrium notions for social network formation processes. *Advances in Complex Systems*, 7(1):77–92, 2004.
- J. Epstein and R. Axtell. *Growing Artificial Societies. Social Science From the Bottom Up*. The MIT Press, 1996.
- I. Fisher. *The Theory of Interest*. Macmillan, 1930.
- M. O. Jackson and A. Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, October 1996.
- C. Menger. *Principles of Economics*. Dingwall and BF Hoselitz, 1871.
- S. Milgram. The small-world problem. *Psychology today*, 2:60–67, 1967.
- D. Phan. From agent-based computational economics towards cognitive economics. In P. Bourguine and J.P. Nadal, editors, *Cognitive Economics*, Handbook of Computational Economics, pages 371–398. Springer Verlag, 2004.
- M. Slikker and A. van den Nouweland. Network formation models with costs for establishing links. *Review of Economic Design*, 5(3): 333–362, 2000.
- L. Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2002.
- L.S. Tesfatsion. A constructive approach to economic theory. In L. Kenneth, J. Tesfatsion, and L. Tesfatsion, editors, *Handbook of Computational Economics*, volume 2 of *Handbooks in Economic Series*. North-Holland, 2006.
- D.J. Watts and S.H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, 1998.

Group Effect, Productivity and Segregation Optimality

Raúl Conejeros¹ and Miguel Vargas²

¹ Department of Biochemical Engineering
Catholic University of Valparaíso, Chile
rconejer@ucv.cl

² Diego Portales University School of Business, Chile
miguel.vargas@udp.cl

15.1 Introduction

In literature on residential segregation (**RS** hereafter) there are two important results about the social optimality of this phenomenon. Firstly, the maximum level of **RS** can constitute a social optimum if one part of the population generates negative externalities on the remaining one. The population suffering the negative externalities can be called a prejudiced population. Under these circumstances traditional bid-rent models with externalities and general equilibrium models have showed the optimality of **RS**³. On the other hand, when the individuals' preferences are to live in balanced neighborhoods, high levels of **RS** diminish the aggregated utility, consequently, full integration is the social optimum. The well known Schelling's model has stressed this issue, showing how a population, in an artificial world, can evolve to a segregated society although individuals want to live in a perfectly balanced neighborhood, reaching a bad, but the only stable, equilibrium. If the prejudiced population coexist with a population preferring balanced neighborhood, the literature has proposed the payment of compensating transfers by the prejudiced population to the non-prejudiced ones to accept the exclusion (Anas, 2002).

These results rely upon the fact of these models considering just the individuals' utility depending on the neighborhood's characteristics and segregated individuals suffering **RS**'s negatives consequences

³ See for instance Anas (2002); Ando (1981); Bailey (1959); Kanemoto (1980); Papageorgiou (1978); Rose-Ackerman (1975); Schnare (1976); Yinger (1976).

just on their own⁴. Notwithstanding, here it is argued that the negative consequences of **RS** would affect the utility of the non-segregated population too. This is a quite reasonable statement if the kind of **RS** consequences are studied. As a matter of fact, literature has pointed out that **RS** can have impact upon the the level of joblessness, out-of-wedlock births, level of criminality, low educational achievement, income inequality and poverty traps, amongst others. **RS** can produce these effects mainly by two mechanisms, namely, lowering the consumption of public good, for instance poor neighborhoods have, in general, bad quality schools⁵; and peer-effects. As an example of the latter, there is fair evidence telling us that the school performance does not depend just on the individual capabilities, but also upon the desire of the rest of the classmates for having a good school performance⁶.

If one part of society is being affected by these sort of difficulties, it is almost sure that all the population is going to be affected too. For example, if **RS** generates spots of criminality, this criminality will reach, also, the non-segregated population. As another example can be considered the fact that if **RS** has a negative impact on educational achievement, ghettos of low-skilled laborers can emerge, a process that can be reinforced by itself. Consequently, the society will loss productivity, and therefore, all individuals' level of consumption will be lower, diminishing the welfare of every single individual, being segregated or not. If that is the case, the prejudiced population is going to face a trade-off between the desire of living just amongst peers and the lower level of consumption that **RS** produces.

Benabou (1993) is the first attempt to formalize the link among location decision, education investment and productivity. According to this article, segregation diminishes some communities' chances to reach good labor skills and therefore the labor force quality will be lower, affecting poor and rich households welfare. Education is assumed to be a local public good affected by community composition. Consequently, segregation arises because rich households want to live in communities with a high level of investment in education, bidding out poor families. Segregation efficiency depends on the form of the function cost of poor households education.

⁴ See for instance Charles et al. (2004); Clapp and Ross (2004); Dawkins et al. (2005); LaVeist (2003); Massey (2001); Wilson and Hammer (2001); Yinger (2001).

⁵ For instance, Berglas (1976) introduces skills to the Tiebout's model.

⁶ Arnot and Rowse (1987) compute the school optimal composition for different forms of peer group effects. De Bartolome (1990) studies the inefficiency in the communities composition due to peer groups effects on education.

Nevertheless, with the important insights that Benabou (1993) has provided, there are some aspect of urban structure that standard classical models cannot treat properly. Meen and Meen (2003) have provided a detailed discussion on these characteristic, explaining why a model considering them, must be used to model urban systems. Specifically, the characteristic stressed by Meen and Meen (2003) are self-organization, heterogeneity between agents, multiple equilibria, unstable equilibria, with the system being out-of-equilibrium for long periods of time (therefore dynamics are more interesting than long-run equilibria). Cellular automata, or agent-based models, have been largely used in literature to deal with system of this sort. These models are attractive for this purpose because they consider explicitly agents heterogeneity and their interactions, which can be more complex than in traditional game theory models. Because of these interactions complexity and self-organization can emerge from agent-based models, where individual action generates ordered pattern of behavior. Besides, multiple equilibria and out-of-equilibrium situations can be studied as an equilibrium formation process: a dynamic process with random events. Hence, for instance, it is possible to study the probability of a particular equilibriums emergence.

Because of the elements explained in the above paragraph, a cellular automaton has been developed here to investigate upon optimality properties of segregation, formalizing the link between location decision and productivity, focusing on peer-groups effects. This is the first time that a model featuring these characteristics is implemented to study the segregation phenomenon.

In section 15.2 the theoretical model is developed, explaining its main characteristics. In section 15.3 through the use of simulations the equilibria and social optima properties are studied and the effects of some public policies. In section 15.4 conclusions and final remarks are given.

15.2 The model

The model developed here is an extension of Schelling (1971) and its first mathematical formalization due to Zhang (2004). Therefore the first step, it is to define an artificial society made up of an advantaged prejudiced population and a non-prejudiced disadvantaged population. Each agent belongs just to a one particular population group. The proportion of these two populations are given by π_j with $j \in \{0, 1\}$, indexing the population's types. If $j = 1$ the agent is a disadvantaged

one and with $j = 0$ a non-disadvantaged one. Each one of this society members, or agents, is allocated in the vertex of a $N \times N$ lattice graph with a periodic boundary condition or, what it is the same under this feature, embedded on a torus.

Utility. Each agent j 's utility U is made up of two parts: a deterministic one u_i and a stochastic term ϵ . This stochastic term reflects the assumption of bounded rationality, because agents can make mistakes, but also guarantee agents' heterogeneity. The deterministic part depends on how many like-type neighbors he has in the local neighborhood and in his level of consumption. The stochastic part is assumed being independent across agents and locations. As Zhang (2004) points out, the latter it is because agents value different characteristics, and different locations have different idiosyncratic traits. The utility function u is assumed additively separable in two components: a location term ℓ and a consumption term c . The location term differs depending on the agent's type.

$$\ell_j = \begin{cases} A_j \left(\frac{x}{n_j} \right) & \text{if } x \leq n_j \\ (2A_j - B_j) + (B_j - A_j) \frac{x}{n_j} & \text{otherwise} \end{cases} \quad (15.1)$$

where $A > B > 0$ are parameters guaranteeing a linear kinked shape, increasing on the left side of n_j and decreasing on the right, being n_j the peak as it is showed in Figure 15.1. Therefore, n_j is the number of like-type neighbors in the local neighborhood that maximize ℓ_j . The total number of the local neighborhoods inhabitants can be represented by $Z_j n_j$, where $Z_j \in \mathbb{N}$. x is the actual number of like-type neighbors in the local neighborhood.

The consumption term for both agent's type is just the normalized level of consumption. Therefore, the deterministic term of the utility function is given by:

$$u_j = \beta \ell_j + \alpha c \quad (15.2)$$

being β and α positives parameters indicating the importance of location and consumption for agents. The utility function for any agent i of type j is:

$$U_{ij} = u_j + \epsilon \quad (15.3)$$

Finally, the aggregated utility is:




Fig. 15.1. Location Utility

$$\mathbf{U} = \sum_1^{N \times N} U_i \quad (15.4)$$

Production. Each agent is endowed in every period of time with one unit of “productivity”. In every period of time the agent can have just one unit of “productivity”. An important assumption it is made here is that the agent’s productivity can be affected by the local neighborhood’s characteristics (group effects). In particular, it is assumed that if in a local neighborhood the majority of inhabitants (more that the 50%) belong to the disadvantaged population, then the productivity of that type of agent decreases. Hence, the productivity for a non-disadvantaged agent is always 1, and for a disadvantaged one is:

$$p_1 = \begin{cases} \frac{x - Z_1 n_1}{\tau - Z_1 n_1} & \text{if } x > \tau \\ 1 & \text{otherwise} \end{cases} \quad (15.5)$$

where τ is the threshold value of like-types neighbors that triggers the productivity diminishing process as is depicted in Figure 15.2. The aggregated production is:

$$\mathbf{P} = \sum_1^{N \times N} p_i \quad (15.6)$$

the level of consumption of every agent is $\frac{\mathbf{P}}{N}$. This means that all the agents, prejudiced and non-prejudiced, are going to be affected negatively by segregation.




Fig. 15.2. Disadvantaged Agents Productivity Function

Log-linear behavioral rule. A Log-linear behavior is observed when agents change their location based upon their own personal interests. In each period two agents coming from different neighborhoods are selected and in order to perform a change, every agent will bid in an sort of auction for a better location, consequently the focus is on the sum of two chosen agents' utility. If the sum of switching is bigger than the sum of not switching, then agents will swap their locations. If the switch situation is called S and the opposite NS , and, for following exposition, the chosen agents are agent 1 and agent 2, and they do not swap locations, then:

$$U_1(\cdot|NS) + \epsilon_1 + U_2(\cdot|NS) + \epsilon_2 = U_1(\cdot|NS) + U_2(\cdot|NS) + \epsilon_1 + \epsilon_2 = V^{NS} + \eta$$

but if they do, then:

$$U_1(\cdot|S) + \epsilon_1 + U_2(\cdot|S) + \epsilon_2 = U_1(\cdot|S) + U_2(\cdot|S) + \epsilon_1 + \epsilon_2 = V^S + \varepsilon$$

Agents will change location if and only if $V^{NS} + \eta < V^S + \varepsilon$. It is assumed that η and ε are independent, and that they follow an identical extreme value distribution. Then, based on McFadden (1973), and following Zhang (2004), a log-linear switch rule can be settled as follow:

$$\begin{aligned}
 Pr(S) &= Pr(V^{NS} + \eta < V^S + \varepsilon) = Pr(\eta < V^S - V^{NS} + \varepsilon) = \\
 &= \int_{-\infty}^{+\infty} F(V^S - V^{NS} + \varepsilon) f(\varepsilon) d\varepsilon = \int_{-\infty}^{+\infty} e^{-e^{-V^S + V^{NS} - \varepsilon}} \\
 &e^{-\varepsilon - e^{-\varepsilon}} d\varepsilon = \int_{-\infty}^{+\infty} \exp \left[-\varepsilon - e^{-\varepsilon} \left(\frac{e^{V^{NS}} + e^{V^S}}{e^{V^{NS}}} \right) \right] d\varepsilon = \\
 &= \int_{-\infty}^{+\infty} \exp [-\varepsilon - e^{-\varepsilon} e^{\Phi}] d\varepsilon = \\
 &= e^{(-\Phi)} \int_{-\infty}^{+\infty} \exp [-(\varepsilon - \Phi) - e^{-\varepsilon - \Phi}] d(\varepsilon - \Phi) = \\
 &= \exp(-\Phi) \int_{-\infty}^{+\infty} f(\varepsilon - \Phi) d(\varepsilon - \Phi) = \\
 &= \exp(-\Phi) \cdot 1 = \left(\frac{e^{V^{NS}} + e^{V^S}}{e^{V^{NS}}} \right).
 \end{aligned}$$

Being $\Phi = \ln \left(\frac{e^{V^{NS}} + e^{V^S}}{e^{V^{NS}}} \right)$, hence,

$$Pr(S) = \left(\frac{e^{V^{NS}} + e^{V^S}}{e^{V^{NS}}} \right)$$

This behavioral rule depends just upon the deterministic utilities, therefore, it is possible to work avoiding the stochastic utilities, which are unobservable.

A segregation measure. In order to measure the level of segregation the index of dissimilarity is used. This index, due to Duncan and Duncan (1955), can be interpreted as the percentage of a group's population that would have to change residence for each local neighborhood to have the same percentage of that group as the metropolitan area overall. The index ranges from 0.0 (complete integration) to 1.0 (complete segregation), and is given by the following formula:

$$D = \sum_{i=1}^m \left[\frac{g_i |k_i - K|}{2GK(1 - K)} \right] \quad (15.7)$$

where m is the total number of local neighborhoods, g_i is the i local neighborhood's total population, G is the total population, k_i is the group of interest percentage in the local neighborhood i , and K is the group of interest total percentage in the city.

15.3 Simulations and well-being analysis

For simulations purposes an artificial society made up by 100 agents is considered. Hence, $N=10$. Both type of agents are equally distributed across the population, therefore $\pi_0=\pi_1=0.5$. The local neighborhoods used are Moore neighborhoods, that, for every agent, include the eight adjacent agents as neighbors. The reasons underlying the latter is because, as Zhang (2004) shows, using this kind of neighborhood the model converges faster, besides, the final outcome is independent of the neighborhood type. It is assumed too, that non-disadvantaged agents are prejudiced against the disadvantaged ones, therefore the former prefer just like-type neighbors, meanwhile disadvantaged ones prefer balanced neighbourhoods⁷. Consequently, $n_1 = 4$ and $4 < n_0 \leq 8$. As one of the main aims of the present research is the search for the conditions that define a social optimum linked to a positive level of segregation, but not absolute segregation, α and β values must be relatively balanced between each other, otherwise either integration or complete segregation will be the social optimum. Intuition tells us in a straightforward fashion, that agents must have balanced preferences between location and consumption. For instance, considering the case where agents have higher preferences on neighborhoods characteristics than on consumption, i.e. $\alpha \rightarrow 0$, with a prejudiced population, the social optimum will imply complete social exclusion. On the other hand, if neighborhoods characteristics are not important, i.e. $\beta \rightarrow 0$, then complete integration will be the social optimum. In order to facilitate calculations and normalize utility functions is picked $A_0 = 1$, $B_0 = \frac{n_0}{n_0 + x}$, $A_1 = 1$ and $B_1 = 0.6$.

The variable of interest is the optimal level of segregation OS and the stable equilibrium level of segregation ES . In the first one is the

⁷ The other two kind of neighborhoods used in the agent-based literature are Von Neumann and $r(2)$, where the former considers the four surrounding agents as neighbors and the latter covers 12 agents inside a circle with radius 2.

level of segregation associated to the maximum level of aggregated utility, and in the second one is the level of segregation that the system reaches by itself. The basic structure of the algorithm used to obtain the *ES* has been already explained in the previous section 15.2. The way to obtain the *OS* deserves some further explanation. As the natural evolution of the problem does not settle at the optimal social utility, a stochastic search method has been used, namely the Threshold Acceptance (TA) algorithm, in order to find the best utility neighboring distribution. This algorithm is related to the general class of Simulated Annealing (SA) type algorithms, enabling the stochastic search for effective solutions to highly combinatorial optimization problems but with a much easier implementation. The algorithm is shown in Figure 15.3.

1. Get an initial system configuration S and an initial threshold $\Delta_{new} = \Delta_{old}$
2. While outer loop stop criterion not satisfied do:
 - a) While inner loop stop criterion not satisfied do:
 - Select a trial solution S'
 - If $C(S') \leq C(S) + \Delta_{new}$, let $S = S'$
 End of inner loop
 - b) If S has changed reduce threshold:
 - $\Delta_{old} = \Delta_{new}$,
 - $\Delta_{new} = \alpha \cdot \Delta_{old}$.
 End of outer loop
3. Report best solution found

Fig. 15.3. The threshold acceptance algorithm

As this is an stochastic algorithm, every time that a maximum it has been searched for in the inner loop, it has been run 100 times, and the maximum value amongst these 100 “maxima” has been picked up.

Agents’ relevant elements, in order to choose a location, are the value of the location utility ℓ and the value of consumption utility c . The parameters values that have been chosen are: $\frac{\alpha}{\beta} = \frac{3}{4}$, $n_0 \in \{5, 6, 7, 8\}$ and $\tau \in \{4, 5, 6, 7, 8\}$. The results of these simulations are shown in Table 15.1.

The most striking fact arising from simulations is that the system always converges to the lowest level of segregation after a random perturbation (see Figure 15.4). This means that full integration comes to

Table 15.1. Optimal and equilibrium segregation

$n_0=8$	$\tau=4$	$\tau=5$	$\tau=6$	$\tau=7$	$\tau=8$	$n_0=7$	$\tau=4$	$\tau=5$	$\tau=6$	$\tau=7$	$\tau=8$
<i>OS</i>	0.00	0.12	0.36	0.48	0.96	<i>OS</i>	0.00	0.08	0.32	0.36	0.84
<i>ES</i>	0.04	0.04	0.04	0.04	0.04	<i>ES</i>	0.04	0.04	0.04	0.04	0.04
<i>ES-OS</i>	0.04	-0.08	-0.32	-0.44	-0.92	<i>ES-OS</i>	0.04	-0.04	-0.28	-0.32	-0.80
$n_0=6$	$\tau=4$	$\tau=5$	$\tau=6$	$\tau=7$	$\tau=8$	$n_0=5$	$\tau=4$	$\tau=5$	$\tau=6$	$\tau=7$	$\tau=8$
<i>OS</i>	0.00	0.08	0.12	0.36	0.84	<i>OS</i>	0.00	0.04	0.04	0.20	0.80
<i>ES</i>	0.04	0.04	0.04	0.04	0.04	<i>ES</i>	0.04	0.04	0.04	0.04	0.04
<i>ES-OS</i>	0.04	-0.04	-0.08	-0.28	-0.80	<i>ES-OS</i>	0.04	0.00	0.00	-0.16	-0.76

be a stable equilibrium. This finding is quite different from previous literature results, where the main fact that has been pointed out is exactly the opposite: the system always converging to the highest levels of segregation. The element that makes the difference is the agents taking into account of the segregation’s negative effect on their own level of consumption.




Fig. 15.4. Evolution of segregation

As the equilibrium segregation is the lowest possible, the difference between optimal and equilibrium segregation is almost all the time negative. This difference gets lower when the level of advantaged agents’ prejudice diminishes, because optimal segregation also diminishes (see Figure 15.5).




Fig. 15.5. Equilibrium v/s optimal segregation

An interesting thing is that if the relative consumption importance gets lower, the system could converge to higher levels of segregation greater than 0, by means of adjusting the consumption relative importance parameter α . An interesting question is if, following this process, the system can reach a stable equilibrium where equilibrium and optimal segregation are equal. After some trials, it was possible to find a combination of parameters' values where this situation is fulfilled. The specific parameters' values used to find this particular equilibrium were the following: $\frac{\alpha}{\beta} = \frac{1}{5}$, $\tau = 4$ and $n_0 = 8$. Under this setting the system converges to an equilibrium level of segregation of 0.64, the same value of the optimal segregation (see Figure 15.6). Therefore, optimal segregation is not necessarily an out-of-equilibrium situation or a corner solution.

15.4 Concluding remarks

The first interesting finding of the present research is that positive levels of segregation, lower than complete segregation, can be a social optimum. This is quite different to the previous literature, where just either full integration or full segregation can be a social optimum.




Fig. 15.6. Evolution to the optimum level of segregation

The optimal segregation value depends on the extend of prejudices and negative impact of group effects on productivity. For instance, with highly prejudiced agents and irrelevant groups effects, the social optimum will be full segregation. In the opposite case, the social optimum will be full integration. However, with more balanced parameters values the optimal segregation will be lying between 0 and 1 but not reaching the extremes.

The most striking result is full integration being a stable equilibrium. The reason that has made the difference is the agents taking into account of the segregation's negative effect on their own level of consumption at the moment of taking their location decisions. Also, it has been possible to find a combination of parameters' values where the system converges to an optimum level of segregation greater than 0 but less than 1 too. Therefore, a society can reach by itself an optimal level of segregation different from the less realistic cases of full integration and full segregation.

All these findings can have interesting policy implications. First of all, segregation must not be seen as a bad situation a priori. What is going to be the optimal level of segregation is something depending upon of how prejudiced are prejudiced agents, the extend of groups effects impact and the relative importance amongst local neighborhood characteristics and consumption on agents' utility. Consequently, it can be argued that there is not such a thing like a unique optimal level of

segregation or an absolute segregation target. If the aim is to improve welfare, policymakers must have a clear picture of individual preferences, group effects and other elements, that have not been treated in this research but it is worthy to mention them, as provision of local and non local public goods, before to implement any policy with the objective of reducing segregation. Besides, it is clear that two cities, or regions, have not the same characteristics, hence, every different city have different level of optimal segregation, and, therefore needs different policies.

References

- A. Anas. Prejudice, exclusion, and compensating transfers: the economic of ethnic segregation. *Journal of Urban Economics*, 52:409–432, 2002.
- A. Ando. *Development of a unified theory of urban land use*. PhD thesis, University of Pennsylvania, 1981.
- R. Arnot and R. Rowse. Peer group effects and educational attainment. *Journal of Public Economics*, 32:287–305, 1987.
- M. Bailey. Note on the economics of residential zoning and urban renewal. *Land Economics*, 35:288–292, 1959.
- R. Benabou. The workings of a city: Location, education, and production. *Quarterly Journal of Economics*, 108(3):619–652, 1993.
- E. Berglas. Distribution of tastes and skills and the provision of local public goods. *Journal of Public Economics*, 6:409–423, 1976.
- C. Charles, G. Dinwiddie, and D. Massey. The continuing consequences of segregation: Family stress and college academic performance. *Social Science Quarterly*, 85(5):1353–1373, 2004.
- J. Clapp and S. Ross. Schools and housing markets: An examination of school segregation and performance in connecticut. *Economic Journal*, 114(499):425–440, 2004.
- J. Dawkins, Q. Shen, and T. Sanchez. Race, space, and unemployment duration. *Journal of Urban Economics*, 58:91–113, 2005.
- C. De Bartolome. Equilibrium and inefficiency in a community model with peer group effects. *Journal of Political Economy*, 48:110–133, 1990.
- O. Duncan and B. Duncan. A methodological analysis of segregation indices. *American Sociological Review*, 20:210–217, 1955.
- Y. Kanemoto. *Theories of Urban Externalities*. North Holland, 1980.

- T. LaVeist. Racial segregation and longevity among african americans: An individual-level analysis. *Health Services Research*, 38:1719–1733, 2003.
- D. Massey. Segregation and violent crime in urban america. In *Problem of the century: Racial stratification in the United States*. Russell Sage Foundation, 2001.
- D. McFadden. Conditional logit analysis of qualitative choice behavior. In *Frontiers of Econometrics*. Academic Press, New York, 1973.
- D. Meen and G. Meen. Social behaviour as a basis for modelling the urban housing market: A review. *Urban Studies*, 40:917–35, 2003.
- J. Papageorgiou. Spatial externalities. i: theory. *Annals of the Association of American Geographers*, 68:465–476, 1978.
- S. Rose-Ackerman. Racism and urban structure. *Journal of Urban Economics*, 7:233–250, 1975.
- T. Schelling. Dynamics models of segregation. *Journal of Mathematical Sociology*, 1:143–186, 1971.
- A. Schnare. Racial and ethnic price differentials in an urban housing market. *Urban Studies*, 13:107–120, 1976.
- F. Wilson and R. Hammer. Ethnic residential segregation and its consequences. In *Urban inequality: Evidence from four cities*. Russell Sage Foundation, 2001.
- J. Yinger. Racial prejudice and racial residential segregation in an urban model. *Journal of Urban Economics*, 3:383–396, 1976.
- J. Yinger. Housing discrimination and residential segregation as causes of poverty. In *Understanding poverty*. Harvard University Press, 2001.
- J. Zhang. Residential segregation in an all-integrationist world. *Journal of Mathematical Behavior and Organization*, 54(4):533–550, 2004.

The Grass is Always Greener on the Other Side of the Fence: The Effect of Misperceived Signalling in a Network Formation Process

Simone Giansante¹, Alan Kirman², Sheri Markose¹, and Paolo Pin³

¹ CCFEA, University of Essex, UK

{sgians,scher}@essex.ac.uk

² GREQAM, Université d'Aix Marseille, France

kirman@ehess.univ-mrs.fr

³ Economics Department, University of Venice, and ICTP, Trieste, Italy

pin@unive.it

16.1 Introduction

Social and economic networks are becoming increasingly popular in the last ten years, because of both the application of game theory to the network formation processes⁴, and the study of stochastic processes that fit the statistical properties of real world social networks.⁵ In the very recent years there have also been attempts to combine the contribution of these two streams of research, trying to find strategic models whose equilibria resemble the empirical data.⁶ A well known source of debate in the game theoretical approach is the incompatibility between stability and efficiency: in most of the models Nash equilibria are actually not the network architectures that maximize the overall sum of utilities, as surveyed in Jackson (2003). On the other hand the econophysics approach is not interested in the utility of single nodes but has other measures of efficiency, which are essentially the probabilities of the network to maintain certain properties after random deletion of links or nodes.

⁴ The seminal paper is by Jackson and Wolinsky (1996), see Jackson (2006) for a survey of this literature.

⁵ The starting point of this second stream of research can be considered Albert and Barabási (1999), see Newman (2003) for a survey. Let us refer to this second scientific contribution as the econophysics approach.

⁶ As an example see Jackson and Rogers (2007).

We will consider a similar trade-off between stability and efficiency, in a game theoretical network formation model. The nature of our model is however such that both the efficient networks, and the most likely outcomes in real world applications, are different from how classical social networks look like. The social networks that are usually brought as examples are those of human informal relations (as friendships). They are strongly connected networks with short minimal paths between any two of their nodes. There are however different environments and related theoretical models where the outcomes are likely to be segregated clusters. Reasonable applications where likely outcomes are actually segregated clusters could be those where the nodes tend to mutually control each others. Think for example as an R&D setup where firms cooperate on secret projects, trying to keep low the risk of industrial espionage. Consider moreover informal contracts of mutual insurance (such as those in rural villages analyzed by Bramoullé and Kranton (2007) where reciprocal control is necessary to avoid moral hazard issues. A third example could be the market of perishable goods described in Weisbuch et al. (2000). Also Jackson and Wolinsky (1996), in their original coauthor-model, imagine a set of researchers that have a utility from working with other colleagues that is increasing in the number of coauthors, but decreasing in the coauthors of their coauthors. Moreover there are some costs for maintaining links. For low enough costs both the efficient outcome and the equilibria of the co-author model are networks segregated in fully connected clusters.

The present paper considers a particular case of the model proposed by Kirman et al. (2007), where agents try to maximize the total number of reciprocal links in their neighborhood. The nodes have a maximum number of links they can send to others. Directed networks are considered so that the notion of Nash equilibrium can be straightforward applied. This model has a Nash equilibrium network architecture which is also the efficient one, that is the case where agents cluster in isolated but complete (i.e. fully connected) subnetworks. However there are also other, less efficient, equilibria, so that the problem is a classical one of coordination among players. We use computer driven simulations, where agents are faced with the possibility to change some of their connections if better profits are present. It comes out that the coordination toward the efficient equilibrium is not always the case, especially if the players have many links to possibly cast, or if they are heterogeneous in their maximum possible number of links. We have found out that, amplifying the signal from a change in strategies, so that it may happen that agents change because they think to improve

their situation but this does not happen, the probability for the final outcome to be efficient, and in general the expected cumulative utility, increase.

Next Section formalize the model and the algorithm used in the simulations, while Section 16.3 shows the results. The main result may seem not surprising to those who know some theoretical physics or computer science: every heuristic optimization algorithm work with experimentation to avoid local minima. We will conclude and analyze this comparison in Section 16.4, arguing that it is not completely exact.

16.2 The model

The game we are considering is a one shot network formation game between N agents. The network resulting from any strategy-profile is a directed irreflexive unweighed one, in the spirit of Bala and Goyal (2000). Each agent i can casts up to l_i links (this number l_i will be the only specification of the agents' type) to other agents.⁷ This is the only action we have considered, so that the strategy set for agent i represents all the subset of the remaining $N - 1$ agents, with cardinality up to l_i . Let us identify the action of every agent with one such subset, and call it R_i . The strategy profile of the game will determine an undirected network of N nodes (the agents) and up to $\sum_i^N l_i$ links (the maximum possible number of links). The structure of the network determines the payoff for each agent. This payoff is a measure of how well interconnected is the neighborhood of degree one of each agent. Agent i counts all the links between all the nodes in $R_i \cup \{i\}$ (the nodes in her out-degree neighborhood and herself), in both directions, call this number π_i , i 's payoff is then exactly π_i .⁸

Let us start our analysis from the homogeneous case where all the l_i are equal to a certain value l . To simplify the analysis let us also assume that N is a multiple of $l + 1$. Given these hypothesis (that we will maintain in our first group of simulations), it becomes easy to analyze the welfare of the system. The most efficient network structure is the one in which the nodes are clustered in segregated fully connected sub-networks of $l + 1$ nodes each. It is easy to check that this is the only possible architecture that guarantees the maximum profit to every node: every cluster is for everyone of its members exactly the $R_i \cup \{i\}$

⁷ The limit in the number of links can have many economic explanations, all of them reducible to a story of limited resources and costly link formations.

⁸ This game is the "level 2 neighborhood" case of the general model considered in Kirman et al. (2007).

neighborhood, in which a maximum of $(l + 1) \cdot l$ links is present. Since the efficient network is feasible and inside it no agent can improve her payoff, it is straightforward to check that the efficient network is also a Nash equilibrium. Unfortunately this is not the only Nash equilibrium, even when l is low (so that coordination does not seem hard), as shown in Figure 16.1, where l is set to 2, while $N = 6$.




Fig. 16.1. Comparison between the efficient Nash equilibrium (left) and a possible non-efficient one. Here $N = 6$ and $l_i = 2 \forall i \in \{1, \dots, 6\}$.

The right hand side equilibrium in Figure 16.1 gives an equal utility to all the agents that is only $\frac{2}{3}$ of that received in the optimal left hand side case (4 links in the neighborhood instead of 6) because every node's neighbors are not connected between each other, and changing neighbors would only make things worse. The suggestion we get from this comparison is that, without the possibility of coordination between the agents, suboptimal outcomes cannot be avoided.

As intuition would suggest, the coordination problems will increase as soon as we introduce heterogeneity in the players of our game. Suppose for example that our N agents are divided in two sub-population of which one has a given value l' for the maximum number of allowed links, while the other has another value l'' . In order to simplify things let us assume that the N' agents whose value is l' are such that N' is a multiple of $l' + 1$, and the same for N'' and l'' (clearly $N' + N'' = N$). With similar consideration as the ones for the simpler case, it comes out that the only efficient network is the one in which agents cluster in fully connected sub-networks of cardinality $l' + 1$ (the N' agents of the first type) and $l'' + 1$ (the remaining N'' agents). Since this configuration maximizes the payoff of every player and no agent can improve her own profit, the efficient network is also a Nash equilibrium.

As in the previous case, however, the efficient equilibrium is not the only one. Consider $N = 30$ and $l_i = 2 \forall i \in \{1, \dots, 15\}$, while $l_j = 4 \forall j \in \{16, \dots, 30\}$. The efficient Nash equilibrium would be: 8 segregated clusters (3 quintuplets and 5 triplets). This is however not the only one. Examples of non-efficient Nash equilibria can be easily set up for this example as has been done in the homogeneous case

16.2.1 The adaptive mechanism

We describe here a very simple algorithm that would allow us to span the possible network configurations and search for the equilibria of the game described above. Imagine a discrete time process where at every time step we have a network configuration. The initial configuration is just a random network with the constraint that every node i casts exactly l_i links to other nodes (from the way payoff are computed we can exclude a priori that a rational agent would cast less links than allowed). At every time step every node: (i) considers all her links and compute which is the worst one in terms of marginal payoff (in case it is more than one uniform probabilities are applied); (ii) considers a random node to whom she is not already linked and compute the marginal payoff from changing her worst connection for this new one; (iii) if this change is profitable to her she changes her link for the better target.

It is not given that, if this algorithm is blocked in a stall network for all the possible deviations, then this network is a Nash equilibrium, since a deviation in strategies could happen also changing more than one single link.⁹ We know however what the efficient Nash equilibrium is. Thus the aggregate wealth of every configuration can always easily be computed, and also the ratio between the aggregate payoff (obtained when the algorithm stops) and the optimal achievable one.

Formally, at every time step, every node i computes the marginal payoff from deleting every one (but only that single one) of her links, call this value $\Delta_\lambda \pi_i$ for every link $\lambda \in \{1, \dots, l_i\}$. She picks a link that minimizes this payoff: $\underline{\lambda} \in \arg \min \{\Delta_\lambda \pi_i\}$. Node i is then assigned randomly (with uniform probabilities) a node $j \notin R_i$ (j is not in i 's neighborhood) and she is informed on what would be her marginal payoff from severing link $\underline{\lambda}$ and connecting to j with a link $\lambda_{\rightarrow j}$. She computes the difference

$$D_{\underline{\lambda} \rightarrow j} = \Delta_{\lambda_{\rightarrow j}} \pi_i - \Delta_{\underline{\lambda}} \pi_i ,$$

⁹ The reverse is however true, the algorithm would surely stop in a Nash equilibrium.

if it is positive she makes the change.

The object of our empirical analysis is then: what would happen if the perception of the marginal payoff from the new connection is positively biased? In the formulas it would be as if we had a parameter $\alpha \geq 1$ (the value of bias) that linearly distort the external information. The perceived difference would become

$$D_{\Delta \rightarrow j}^{\alpha} = (\alpha \cdot \Delta_{\lambda \rightarrow j} \pi_i) - \Delta_{\lambda} \pi_i .$$

If this new formula is positive the agent will change her connection, but it is not guaranteed that also the real difference $D_{\Delta \rightarrow j} \leq D_{\Delta \rightarrow j}^{\alpha}$ is non-negative.

16.3 Results

We apply the adaptive mechanism described in previous Section to computer-based simulations, for different values of $\alpha \geq 1$. α is the positive bias in the perception of the payoff from new links, its value is exogenously fixed in every simulation and is the same for all the agents in the system.

The measurement that we will make for those simulations are two. First of all we will check if the algorithm eventually stops to an absorbing network configuration.¹⁰ As discussed above a stall network is however not guaranteed to be a Nash equilibrium. This holds even for $\alpha > 1$.¹¹ It is plausible to expect from the simulations that, as the value of α increases from 1, the probability that the algorithm will not stop in some fixed configuration becomes higher and higher. There will moreover be a value of α for which the system will never stop because, no matter the shape of the network and the resulting payoffs, there will always be an agent willing to change one of her links. Since we are running several simulations for every starting population and value of α , we keep track of the percentage of those simulations that happen to stop in some network configuration. We call this first measure the *Stability* of a simulation.

As clarified in Section 16.2 we know what is the optimal achievable payoff for every agent i , as soon as we know her link-capacity l_i . This

¹⁰ Technically we check if the system does not change for a long enough number of time-steps.

¹¹ For $\alpha > 1$ neither the reverse implication can be made: if the system is in a Nash equilibrium, it can still change configuration because some agents may get wrong signals and adopt a non-profitable change of a link.

maximal payoff is the one she gets in a fully connected cluster of similar agents and is the number of links in such a cluster, namely $l_i \cdot (l_i - 1)$. We choose the starting population so that there is surely an efficient Nash equilibrium where all the agents reach their optimal payoff. The second measure we take will identify how far the agents' final payoff are from the ideal efficient case. We measure the final utility of each agent and consider the ratio with the optimal one, then we express as a percentage the average value of this ratio across all the population.¹² This second measure will be called the *Wealth* of a simulation. We have then two indices, Stability and Wealth (both expressed in percentage), for every starting population and bias α .

The first group of simulation has been made with a population of 60 homogeneous agents, each with a link capacity $l_i = 2$. This is a 10 times larger group than the one considered in the example of Figure 16.1, but the efficient Nash equilibrium is still one in segregated triplets. Figure 16.2 shows the result when α is ranging from 1 to 2.5. For values of $\alpha > 1.5$ the stability of the system is not sure anymore. Full wealth optimality is never reached on average (the Wealth measure is always below 90%), but a simple second order polynomial interpolation of the results suggest that an expected maximum Wealth is reached between $\alpha = 1.5$ and $\alpha = 2$. The result is surprising for two reasons. First of all the higher expected wealth is not reached for $\alpha = 1$, but for a higher value. Secondly this optimum seems to stand where the algorithm is not surely stable anymore. The effect of bias undermines stability after a certain α^U but, up to a certain point $\alpha^* > \alpha^U$, it improves the expected payoff of the agents. As we will see the qualitative outcome of this first group of simulations will hold also in the following ones.

Figure 16.3 shows the result of simulations when the population of 60 players is still homogeneous, but now the link capacity of the agents is $l_i = 4$. The efficient Nash equilibrium is a network divided in 12 segregated quintuplets. As expected now the agents' coordination problem gets more tricky: the value α^U over which instability may arise becomes smaller, while the value α^* where there is an optimum expected wealth seems to increase. As a consequence we get the same kind of results (i.e. $1 < \alpha^U < \alpha^*$) as in the previous simpler model.

We have ran simulations also for the case of non-homogeneous populations, mixing between the two previous cases. We have considered the three quartile distributions with 60 agents: 15 agents with a link capacity of $l_i = 2$ and 45 with a link capacity of $l_j = 4$; the 30 – 30

¹² When the algorithm is not stable and there is not a final configuration, we measure the average wealth of the agents on a sufficiently large final time window.




Fig. 16.2. Simulations on a homogeneous population of 60 agents, with $l_i = 2$. α ranges from 1 to 2.5, with a grid of 0.1. For every value of α in this grid 20 simulations are ran. Small squares show Wealth in every simulation; big squares are for the average Wealth, which is interpolated with a second order polynomial; triangles indicate stability.




Fig. 16.3. Simulations on a homogeneous population of 60 agents, with $l_i = 4$. α ranges from 1 to 2.5, with a grid of 0.1. For every value of α in this grid 20 simulations are ran. Small circles show Wealth in every simulation; big circles are for the average Wealth, which is interpolated with a second order polynomial; triangles indicate stability.

case and the 45 – 15 case. All of them have efficient Nash equilibria where all the two-linked agents ($l_i = 2$) cluster in triplets, while the four-linked ones ($l_j = 4$) segregate in quintuplets. Figure 16.4 (page 231) shows the results. As when increasing link capacity with homogeneous populations, an even higher complexity reduces the likelihood of




Fig. 16.4. Simulations on a heterogeneous population of 60 agents: 15, 30 and 45 respectively with $l_i = 2$, the remaining ones with $l_j = 4$. α ranges from 1 to 2.5, with a grid of 0.1. For every value of α in this grid 20 simulations are ran. Small squares show Wealth of the 2-links agents in every simulation; big squares are for their average Wealth, which is interpolated with a second order polynomial; circles are used for the 4-links agents; triangles indicate stability.

coordination, that is less stability and less expected wealth. In all of the three cases Stability suffer a substantial reduction above $\alpha = 2$,¹³ but

¹³ We observe what physicians call a phase transition.

already at $\alpha = 1.1$ stability is not guaranteed (i.e. $\text{Stability} < 100\%$), so that we can locate α^U slightly above 1. The points α_2^* and α_4^* , where the maxima of expected wealth are located, are now shifted to the right, even if the agents with only two links seem to suffer less from the cohabitation, at least when stability is kept safely above 0.

From simulations we have the *unexpected* result that a value of $\alpha > 1$, i.e. an encouragement for trying new connections, leads instead to a more probable and faster type-biased segregation, that is the efficient equilibrium. The idea is that if agents receive exaggerated incentives linking, i.e. the false perception of better conditions than the actual ones, in the long run they do not facilitate connected networks but instead drive themselves with higher probability to segregation.

16.4 Conclusion

We have a game theoretical network-formation model with multiplicity of equilibria, of which only a very few are the efficient ones. We are facing a typical coordination problem. We briefly discussed in the Introduction which economic applications may have segregated equilibria as the optimal ones.

The result could have two possible implications, a regulatory one and a behavioral one. Think for example to openness and mobility in the job market. Incentives to mobility, even if positively biased in a misleading way, so that people find worse conditions than expected, may lead in the long run to more efficient allocations of the resources. The fact that human beings seem to overestimate the conditions of their similar in different neighborhoods, and the profit they could get elsewhere, could be an inborn attitude to implement experimentation.

The model and the results proposed here have clear analogies, at one side, with other segregation models, at the other side, with the heuristic optimization techniques that allow non-optimal steps in order to exclude local minima (sub-optimal but stable network configuration in our case). Our setup of heterogeneous agents that cluster in closed homogeneous subgroups seems very close to the Schelling (1971) Segregation Model (SSM), even if differences should be pointed out. Here the topology of the problem is unbounded, every agent can cast links (up to her maximum threshold) to every other agent and there is no fixed grid. The geometrical constraint is much stronger in that model. In SSM, moreover, a single agent's deviation is a jump to a completely different neighborhood, while our agents can change only one neighbor at time. In this way also the effect on the other agents is reduced

here in comparison to SSM. Another difference between our heterogeneous setup and SSM is that for us the distinction across the two sub-populations is not in preferences, but in the spread of their neighborhood. We do not reach segregation because of different utilities but because of different resources.

We think that the achieved result (that is: biased perception increases average wealth) are a counter-intuitive novelty, at least in an economic setting of individual profit maximization. They may however not seem so if we consider how our adaptive algorithm has clear analogies with the optimizing simulated annealing algorithm at non-zero temperature (See Kirkpatrick et al., 1983). Since the latter works well in finding optima also our process should increase the average wealth. There are however two considerations that should be made. The first one is that the simulated annealing works well at a stable temperature only if the global optimum is distant (i.e. different) enough from the local ones. This is not something that we have for granted in our model, where the value of bias α is kept fixed along any single simulation. The second, more important, point is that in heuristic optimization, when evaluating a new move, marginal global utility is considered and not the marginal utility of a single agent. The heuristic optimization would work exactly in a cooperative game, but not in our non-cooperative network formation process. In short: when an agent of ours decides to change a link she does not care about what would happen to her neighborhood, neither the old nor the new one. So, there are at least two reasons for which our results are in principle not predictable by the optimizing stochastic processes considered by theoretical physics and theoretical computer science.

References

- R. Albert and A. L. Barabási. Emerging of scaling in random networks. *Science*, 286:509–512, 1999.
- V. Bala and S. Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1230, September 2000.
- Y. Bramoullé and R. Kranton. Risk-sharing across communities. Technical report, Mimeo, January 2007.
- M. O. Jackson and B. W. Rogers. Meeting strangers and friends of friends: How random are social networks? *American Economic Review*, 2007. Forthcoming.

- M. O. Jackson and A. Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, October 1996.
- M.O. Jackson. The stability and efficiency of economic and social networks. In M. R. Sertel and S. Koray, editors, *Advances in Economic Design*, pages 319–362. Springer, 2003.
- M.O. Jackson. The economics of social networks. In Blundell R., Newey W., and Persson T., editors, *Advances in Economics and Econometrics, Theory and Applications*, volume 1, chapter 1. Cambridge University Press, 2006.
- S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science, Number 4598, 13 May 1983*, 220(4598): 671–680, May 1983.
- A. Kirman, S. Markose, S. Giasante, and P. Pin. Marginal contribution, reciprocity and equity in segregated groups: Bounded rationality and self-organization in social networks. *Journal of Economic Dynamics and Control*, 2007. Forthcoming.
- M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.
- C. Weisbuch, A. Kirman, and D. Herreiner. Market organization and trading relationships. *Economic Journal*, 110:411–436, 2000.

Methodological Issues and Applications

Market Selection of Competent Venture Capitalists

David Mas

ERMES-CNRS, Université Panthéon-Assas Paris II, France
David.Mas@u-paris2.fr

17.1 Introduction

Venture capital is actually considered a very efficient mean for financing innovation. Kortum and Lerner (2000) estimate that a dollar invested in venture capital is three times more effective in stimulating patenting than a dollar invested in traditional R&D. Understanding venture capital is therefore a central matter for designing innovation policies. Venture capital has a very peculiar functioning. Venture capitalists finance young firms whose only activity is to develop radical innovations (start-ups). These new firms have no access to the banking system because they are too risky and have no collateral. But their future is also too uncertain to allow them to enter the financial market. Venture capitalists are assuming this uncertainty because they are looking for high-risk/high-reward investments. Venture capitalist are not wealthy individuals risking their own money (business angels), but fund managers. This means they provide institutional investors (pension funds, insurance companies, investment bank) with the possibility to invest in an asset, the venture capital fund, whose risk is manageable with traditional financial methods, like portfolio diversification. Venture capital is in fact turning the uncertainty of investing in radical innovations into a simple, though high, risk.

I propose a model that explicitly describe this essential feature of the venture capital market. This model is intended to understand what are the conditions for the existence and for an efficient functioning of a venture capital market. Following Carlsson and Eliasson (2003), I consider that venture capitalists are competent actors whose function is to select and finance the most promising start-ups. This means that venture capitalists reduce uncertainty in choosing the right start-ups to

invest in. To formalize this idea, I propose a model of venture capital as a market of heterogeneous interacting agents. On the one hand there are start-up projects of different qualities that need financing, and on the other hand there are venture capitalists of different levels of competence trying to detect and finance the best start-up. The evaluation of start-up quality by venture capitalist is imperfect, and its accuracy depends on venture capitalist's competence. Despite the simplicity of the model, the interaction of heterogeneous agents in a stochastic environment makes it tractable only with a simple distribution of agents' qualities and some restrictive hypotheses. The use of simulation allows to overcome this limitation.

In a previous paper Mas and Vignes (2006) show under which conditions the competence of venture capitalists allow for an efficient screening of start-ups. In this paper I investigate the market selection of venture capitalists and show how an efficient venture capital industry can emerge from an initial random population of heterogeneous venture capitalists. I also identify and study the role of a particular institution of venture capital, the limited partnership, in venture capitalists selection. I show i) that the accuracy of the selection as well as the risk taken by institutional investors increases with the size of the limited partnership, i.e. with the number of investments a venture capitalist can make before having to raise an other fund, ii) that the size of the limited partnership has almost no influence on the final distribution of competence, iii) that the optimal choice for the size of the limited partnership can be determined by the computation of selection costs taking type I and type II errors into account.

17.2 Related literature

Kaplan and Stromberg (2001) distinguish three main roles for venture capitalists, which are screening, contracting and monitoring. The contracting role has been extensively studied both from a theoretical point of view and from an empirical point of view (Kaplan and Stromberg, 2000). The optimal contract approach has successfully proposed rationales for some stylized facts of venture capital, like the control right allocation in venture capital contracts (Hellmann, 1998) and the staging of venture capital investment (Gompers, 1995). The monitoring role of venture capitalists is often considered as the specific added value provided by venture capitalists. The early empirical studies (Gorman and Sahlman, 1989; Sapienza, 1992) emphasize the time spent by venture capitalists interacting with the firms of their portfolio. The later ones

(Hellmann and Puri, 2000, 2002; Lerner, 1995) show evidences of the active involvement of venture capitalists in the management of start-ups and the positive impact of this involvement on the start-up success. But the screening role of venture capitalists has received much less attention. I consider though, like Carlsson and Eliasson (2003), that the screening of start-ups is the main role of venture capitalists. The model presented in this paper places this role at the core of the venture capital market.

My work can be related to two other models. Chan (1983) proposes a model of the venture capital market in which information about the quality of investment is costly. These information costs justify the existence of venture capitalists. In this paper I propose an alternative for the role of venture capitalists : instead of costly information, I consider that information about start-up quality is only (imperfectly) available to competent venture capitalists. Amit et al. (1999) propose a model of venture capital that takes all agency risk (moral hazard and asymmetry of information) into account. They also propose a screening competence, called 'due diligence', that would allow to predict the quality of a given start-up project (and thus to minimize the asymmetry of information) : but they have not developed the model based on this hypothesis. In this paper I adopt a very similar hypothesis for the screening competence of venture capitalists, I build a model based on that idea and I simulate its functioning.

17.3 The model

The model contains two types of heterogeneous agents, the start-ups and the venture capitalists. I first present each agent and then the dynamics of the model. I only provide the essential equations, see Mas and Vignes (2006) for a more detailed presentation.

17.3.1 The start-ups

Start-up projects are heterogeneous in quality. The quality of a start-up project i determines her probability of success p_i .

$$p_i = q_i \bar{p} \quad (17.1)$$

I assume that the mean quality of the start-up projects $\bar{q}$ is one, so that $\bar{p}$ is the mean probability of success of the start-ups projects. Investing in a start-up is risky, it only generates profit in case of success.

For an initial investment I in a start-up i the expected profit is given by the following equation.

$$E(\pi_i) = (p_i g - 1) I \quad (17.2)$$

Equations (17.1) and (17.2) simply state that the higher the quality, the higher the probability of success and the higher the expected profit.

17.3.2 Venture capitalists

Competence is defined as the ability of a venture capitalist to evaluate the quality of a start-up project. For a venture capitalist of competence c_j , the evaluation of the start-up project quality q_i is :

$$\tilde{q}_{ij} = c_j q_i + (1 - c_j) u_{ij} \quad (17.3)$$

Here u_{ij} is a random noise with the same distribution as q_i . The accuracy of the evaluation varies from perfect information, when $c_j = 1$, to pure noise, when $c_j = 0$. The venture capitalists screen the start-up projects in sequence. Each venture capitalist evaluates the available start-up projects and picks the best one according to his own evaluation.

17.3.3 Dynamics

Start-ups only live for one period. They are created as projects and screened by venture capitalists. If they are selected, they may succeed and generate profits. Venture capitalists stay on the market as long as they have enough capital to invest. The failed venture capitalists are replaced by new ones. The number of start-up projects S and of venture capitalists V is constant in time.

Each venture capitalist raises a fund LP (Limited Partnership) that he invests totally, always financing one start-up per period. The eventual profits are accumulated by the limited partners. Once the fund is totally invested the limited partners may decide to give the venture capitalist a new fund to manage or not, according to his performance.

At each period n , we repeat the following steps:

1. new start-up projects are created.
2. venture capitalists screen the start-ups
3. the financed start-ups generate profits or losses.
4. venture capitalists accumulate the profits or losses realized by their start-up.
5. venture capitalists whose fund has been totally invested raise a new fund or are replaced by new ones.

17.4 Formal analysis

The aim of this section is to find the best estimator of the performance of venture capitalists and to study its influence on the selection process.

17.4.1 Best selection criterion

The competence of a venture capitalist determines in a non linear way the expected quality of the start-ups he finances, and hence their expected probability of success. This probability is also affected by the market conditions and structure (Mas and Vignes (2006)). Thus I can define for each venture capitalist j the expected probability of success p_j of his start-ups s_j as a function of his competence c and of the market conditions M . With given market conditions M , p_j only depends on c_j , and $\frac{\partial p_j}{\partial c_j} > 0$.

$$p_j = E(q_{s_j}) \bar{p} = f(c_j, M) \quad (17.4)$$

The outcome of an investment for the venture capitalist j is a Bernoulli trial of probability p_j . Thus the number of successful investments out of n follows a binomial law $B(n, p_j)$. From these considerations it follows that the average profits of all the investments made by a venture capitalists converge towards a value that depends linearly on p_j .

$$\hat{\pi}_j = \frac{1}{n} \sum_t \pi_{j,t} \xrightarrow{n \rightarrow \infty} (p_j g - 1) I \quad (17.5)$$

$$E(\hat{\pi}_j) = \frac{1}{n} (E(B(n, p_j))g - n) I = (p_j g - 1) I \quad (17.6)$$

$$var(\hat{\pi}_j) = \frac{1}{n^2} var(B(n, p_j)) g^2 I^2 = \frac{1}{n} p_j (1 - p_j) g^2 I^2 \quad (17.7)$$

The expected profits of a venture capitalist depend through p_j on his competence c_j . After any number of investments made by a venture capitalist, the average of his realized profits is the best estimator of his performance. Therefore the best selection criterion for venture capitalists is $\hat{\pi}_j > 0$.




Fig. 17.1. Probability of selection of a venture capitalist at the end of his first limited partnership as a function of his expected success rate p_j for various number of investments per fund $LP = k g$ with $g = 10$.

17.4.2 Influence of the size of the limited partnership

For the institutional investor who gives the venture capitalists a fund to manage, the finite time horizon of the limited partnership provides both a powerful incentive to the venture capitalist to do his best and an opportunity to monitor his performance. The institutional investor can choose the size of the fund LP which corresponds to a given number of investments after which he can choose to give another fund to the venture capitalist or to give his chance to a new one. What is the impact of this choice on the selection of venture capitalists?

Let us assume for simplicity that the size of the limited partnership is a multiple of the return in case of success $LP = k g$. From the previous analysis I can compute the probability of meeting the selection criterion at the end of the first fund.

$$P(S|p_j, LP) = P(B(k g, p_j) \geq k g - k g \mid I > 0) \quad (17.8)$$

$$= P(B(k g, p_j) \geq k g) \quad (17.9)$$

$$= 1 - F_{B(k g, p_j)}(k g) \quad (17.10)$$

For a given size of the limited partnership $LP = k g$ the probability of selection can be expressed as the value at $k g$ of the complementary cumulative distribution function (ccdf) of a binomial law of parameters $k g$ and p_j . Figure 17.1 shows the evolution of the selection function for various sizes of funds. As k increases it converges towards an Heaviside's

function $h(p_y - \frac{1}{g})$. The bigger the size of the limited partnership, the more accurate is the selection. But a bigger size of the fund corresponds also to a bigger risk for the institutional investor. Thus the choice of the size of the limited partnership is a balance between the accuracy of the selection and the risk taken by the institutional investor.

17.4.3 Global selection function

The outcome of the complete selection process is not tractable. In order to study it, I will simulate the model. After the end of the run it is possible to retrieve the global selection function using Bayes' rules. Let S be the event that the venture capitalist has been selected. Then the global selection function is :

$$P(S|c) = \frac{P(c|S) P(S)}{P(c)} \quad (17.11)$$

Where $P(c|S)$ is the final distribution of competence, $P(S)$ the ratio of the number of selected venture capitalists over the number of the venture capitalists that once entered the market and $P(c)$ is the initial distribution of competence.

17.5 Simulations results

I simulate the model with four different settings, with either a gaussian or an exponential distribution for the quality of start-ups and with a uniform or triangular distribution for the competence of venture capitalists. I then vary the size of the limited partnership. Following the formal analysis, I choose $LP = k g$, with $g = 10$. In each case I ran one hundred simulations of five thousand periods, with a population of one hundred venture capitalists that screen five hundred start-ups projects at each period.

The parameters of the model have been chosen such that the minimum level of competence required to achieve positive profits is $c_G = 0.5$ for the gaussian distribution of quality, and $c_E = 0.6$ for the exponential distribution. The essential difference between the two distributions of quality is the spread of the function $p_j = f(c_j, M)$. With the exponential distribution the difference between more competent investors and lesser ones is bigger. It should leads to a more discriminative selection. The difference between the distributions of venture capitalists' competence, uniform and triangular, is the proportion of very competent venture capitalists. With the triangular distribution sufficiently competent venture capitalists are rarer, making the selection more challenging.




Fig. 17.2. Distributions of competence after 5000 periods

17.5.1 Final distribution of competence

Figure 17.2 shows the resulting distribution of competence after five thousand time periods. These distributions are averaged across the hundred simulation run with each setting. The plain line represents the initial distribution of competence.

The first result is that the selection is efficient. For each setting, even with few very competent venture capitalists in the initial distribution, the final population in the market is essentially composed of sufficiently competent venture capitalists.

The second result is that the final distributions of competence are almost the same for all sizes of the limited partnership. This is surprising because the formal analysis shows that bigger sizes should lead to a more accurate selection. This means that even if each selection step is different, the result of the global selection process does not depends of the size of the limited partnership.

17.5.2 Market selection function

Figure 17.3 shows the global selection function for each process, computed using Equation (17.11). The plain line represents the perfect




Fig. 17.3. Selection functions after 5000 periods


selection function. The selection functions are approaching the perfect selection function, getting closer to it as the size of the limited partnership increases.

The selection functions are clearly better with the exponential distribution of quality, which corresponds to a more discriminative spread of the performance of venture capitalists. The market has its own selection capability, determined by the market conditions which are essentially the distribution of agents' characteristics. The choice of the size of the limited partnership magnifies this selection capability.


This analysis confirms that even if the size of the limited partnership does not change the final distribution of competence, it affects the accuracy of the selection.

17.5.3 Optimal size of the limited partnership


The optimal size for the limited partnership LP corresponds to the optimal balance between selection accuracy and risk. In order to determine it, I compute the total selection costs paid to obtain the (identical) final distribution of competence with the following formula :




(a) Q: gaussian - failure and losses



(b) Q: gaussian - selection costs



(c) Q: exp. - failures and losses



(d) Q: exp. - selection costs

Fig. 17.4. Selections costs as a function of the size of the limited partnership $LP = k g$, $g = 10$

$$C_{selection} = E(losses) * failures + c_{search} * failures \quad (17.12)$$

The first term is the total losses incurred by the institutional investors before finding a sufficiently competent venture capitalists. It corresponds to the costs associated with type I errors (financing a bad venture capitalist). The second term is the total search costs incurred when the institutional investors have to look for a new venture capitalists to replace a failed one. It corresponds to the costs associated with type II errors (missing a good venture capitalists).

Figure 17.4 shows on the left the evolution of the number of failures and of the expected losses in case of failure with the size of the limited partnership. In both cases, gaussian and exponential distribution of start-up qualities¹, expected losses increases linearly with LP , while the number of failures falls as a negative power of LP . As already stated, increasing LP increases both the risk and the accuracy of the selection.

¹ For concision the results are only showed for the uniform distribution of competence. They do not change with the triangular distribution of competence.

The graphs on the right of Figure 17.4 allow to determine the optimal size for the limited partnership for different values of the unitary search cost c . In both cases the institutional investor will prefer the smallest size in the absence of search costs. If they don't pay for type II errors, it is rational for them to minimize the number of type I errors only. As the search costs increase, the optimal size for LP also increases. In the gaussian case, the selections costs with no search costs are very close. Thus the optimal size increases very rapidly with c . In the exponential case, on the contrary the initial difference is much bigger, and only slowly compensated by increasing search costs. With the same value of unitary search cost $c = 1$, the institutional investors will prefer the size $LP = 20$ in the exponential case when they already prefer the size $LP = 100$ in the gaussian case.

17.6 Conclusion

The analysis of this agent based model provides a characterization of the selection process of competent venture capitalists with a noisy selection criterion. It shows that the market selection based on venture capitalists' past performance is efficient. It also shows that the accuracy of the selection increases with the size of the limited partnership, i.e. with the number of observations used for the evaluation. The final distribution of venture capitalists' competence, however, is independent of the size of the limited partnership.

The choice of the size of the limited partnership for the institutional investor is a balance between the accuracy of the selection and the risk he incurs. The optimal size can be determined by the computation of selection costs. It depends on the relative costs of type I and type II errors.

Acknowledgement. This research was supported by CO3 NEST 012410 EC grant.

References

- Raphael Amit, James Brander, and Christoph Zott. Venture-capital financing of entrepreneurship: Theory, empirical evidence and a research agenda, 1999.
- Bo Carlsson and Gunnar Eliasson. Industrial dynamics and endogenous growth. *Industry and Innovation*, 10(4), 2003.

- Yuk-Shee Chan. On the positive role of financial intermediation in allocation of venture capital in a market with imperfect information. *The Journal of Finance*, 37(5), 1983.
- Paul A. Gompers. Optimal investment, monitoring and the staging of venture-capital. *The Journal of Finance*, 50(5), 1995.
- Michael Gorman and William Sahlman. What do venture capitalists do ? *Journal of Business Venturing*, 1989.
- Thomas Hellmann. The allocation of control rights in venture capital contracts. *The Rand Journal of Economics*, 29(1), 1998.
- Thomas Hellmann and Manju Puri. The interaction between product market and financing strategy : The role of venture capital. *The Review of Financial Studies*, 13(4), 2000.
- Thomas Hellmann and Manju Puri. Venture capital and the professionalization of start-up firms: Empirical evidence. *The Journal of Finance*, 57(1), 2002.
- Steven Kaplan and Per Stromberg. Venture capitalists as principals: Contracting, screening, and monitoring. *The American Economic Review*, 91(2), 2001.
- Steven Niel Kaplan and Per Johan Stromberg. Financial contracting theory meets the real world: An empirical analysis of venture capital contracts. Working paper, University of Chicago, Graduate School of Business, 2000.
- S. Kortum and J. Lerner. Assessing the contribution of venture capital to innovation. *The Rand Journal of Economics*, 31(4), 2000.
- Joshua Lerner. Venture capitalists and the oversight of private firms. *The Journal of Finance*, 50(1), 1995.
- David Mas and Annick Vignes. Why do we need venture capitalists ? the role of competent intermediaries in the creation process. Working paper, ERMES-CNRS, 2006.
- Harry Sapienza. When do venture capitalists add value ? *Journal of Business Venturing*, 7, 1992.

A Binary Particle Swarm Optimization Algorithm for a Double Auction Market

Calogero Vetro and Domenico Tegolo

Department of Mathematics and Applications, University of Palermo, Italy
{vetro,tegolo}@math.unipa.it

18.1 Introduction

In this paper, we shall show the design of a multi-unit double auction (MDA) market. It should be enough robust, flexible and sufficiently efficient in facilitating exchanges. In a MDA market, sellers and buyers submit respectively asks and bids. A trade is made if a buyers bid exceeds a sellers ask. A sellers ask may match several buyers bids and a buyers bid may satisfy several sellers asks. The trading rule of a market defines the organization, information exchange process, trading procedure and clearance rules of the market. The mechanism is announced before the opening of the market so that every agent knows how the market will operate in advance. These autonomous agents pursue their own interests maximizing their own utilities. Therefore, we can view our market as a multi-agent system where the market mechanism defines the structure and rules of the environment in which agents will play the market game. An efficient market maximizes the total profit obtained by all participating agents (Fudenberg and Tirole, 1991). However, voluminous game theory literature focuses on auction markets. Satterthwaite and Williams (1989) were among the early researchers studying double auction markets. They designed a single-unit double auction (SDA) market where they eliminated the strategic behavior on the sellers side. McAfee (1992) allowed strategic behavior on both sides of a SDA market and required a market maker to balance the budget. Gjerstad and Dickhaut (1998) allowed agents to use simple rules to form beliefs about their opponents offers and showed that the market price converged to competitive equilibrium quickly. Das et al. (2001) carried over a series of experiments where humans and software agents competed with each other. Sandholm and Suri (2001) showed

that if a double auction market allows agents to submit discriminatory bids, the problem of clearing the market faced by the market maker is NP-Complete.

In this paper, a binary particle swarm optimization algorithm has been proposed to solve a quadratic assignment problem to achieve an optimal solution to a double auction market.

In section 2 it will be proposed the mathematical model, section 3 will introduce the *PSO* methodology as proposed by Kennedy and Eberhart, a special binary version of *PSO* will be shown in section 4, and finally sections 5 and 6 will respectively give results and conclusions.

18.2 Mathematical model

One ideal way of organizing an efficient MDA market is to let the buyers/sellers submit bids/asks about how many items they want to purchase/sell and at what reservation prices. Based on this information, the market maker solves an optimization problem to determine how many units each agent should purchase/sell and at what price to maximize the total profit of the market. In a MDA market with m buyers and n sellers, each buyer i wants to purchase Z_i unit items and each seller j has Y_j unit items to sell. We assume both Z_i and Y_j are known to every agent. The reservation prices, which are private, for buyer i and seller j are b_i and s_j . We assume the reservation price for each agent is static. Let q_{ij} denote the quantity buyer i buys from seller j . If all information is public, the maximum total market value can be obtained by solving the following linear programming problem:

$$\max \sum_{i=1}^m \sum_{j=1}^n q_{ij}(b_i - s_j), \quad (18.1)$$

where

$$\sum_{i=1}^m q_{ij} \leq Y_j \quad \forall j \quad (18.2)$$

and

$$\sum_{j=1}^n q_{ij} \leq Z_i \quad \forall i \quad (18.3)$$

with

$$q_{ij} \geq 0 \quad \text{for every } i, j. \quad (18.4)$$

Constraints (2) and (3) state that a seller sells no more than what he possesses and a buyer will not buy more than he needs. It is interesting to note that the trading price does not show up in the problem. In fact, if buyer i buys quantity q_{ij} from seller j at price p_{ij} , then the market value this transaction implements is the sum of buyer i utility plus seller j utility, which is $q_{ij}(b_i - s_j)$ independent of the trading price p_{ij} . However, it is clear that the trading price will affect each agents utility.

18.3 Particle swarm optimization

Particle Swarm Optimization (PSO) is a parallel population-based computation technique developed by Kennedy and Eberhart (1995). Their biological inspiration is based on the organisms behavior such as flocking of birds and schooling of fishes. In these groups the movement of the whole swarm is based on his own knowledge and on a leader, the one with the best performance. From this, one can learn that PSO shares many common points with genetic algorithm, in fact, they start with a randomly generated population, evaluate the population with a fitness function, update the population, search for the optimum with random techniques and do not guarantee optimality. PSO's major difference from genetic algorithm is that PSO uses the physical movements of the individuals in the swarm as a flexible mechanism to combine global search and local search avoiding local optima, whereas genetic algorithm uses genetic operators like mutation and crossover. Each individual of the swarm has a position in the solution hyperspace and a velocity, that is changed, at each step, to update individual position. Each particle knows its position and the value of the fitness function for that position. Besides each particle keeps track of its coordinates in the problem space, which are associated with the best fitness value it has achieved so far. Every particle knows also the best position among all of the particles and its fitness value. The update of the particle position is the result of a compromise among three alternatives: following its current pattern of exploration; going back towards its best previous position; going towards the overall best position. The updating processes are accomplished according to the following equations:

$$x_{ij}(k+1) = x_{ij}(k) + v_{ij}(k) \quad (18.5)$$

$$v_{ij}(k) = F \cdot [w \cdot v_{ij}(k-1) + v_{ijlocalbest}(k-1) + v_{ijglobalbest}(k-1)] \quad (18.6)$$

$$v_{ijlocalbest}(k) = c_1 \cdot r \cdot (x_{ijlocalbest}(k) - x_{ij}(k)) \quad (18.7)$$

$$v_{ijglobest}(k) = c_2 \cdot r \cdot (x_{ijglobest}(k) - x_{ij}(k)). \quad (18.8)$$

At each step, equation (6) calculates a new velocity for each particle in the swarm based on its velocity at previous step, the best position it has been achieved (locbest) and the best position (globest) the population has been achieved. Then, using the resultant velocity value the position of each particle is updated by equation (5). About the coefficients in equation (6), F is a constrictor factor to insure convergence, the use of an inertia weight w provides to improve performance in a number of applications, while the (0,1)-random constants r and the coefficients c_1 and c_2 represent the weighting of the stochastic acceleration terms that pull each particle towards xlocbest and xglobest locations. The PSO algorithm is stopped when the best particle position of the entire swarm cannot be improved further after a sufficiently large number of iterations. PSO algorithms were proposed, in many research fields, to solve continuous optimization problems. Now, it is very interesting the possibility to solve discrete optimization problems. We do this adopting the quantum discrete algorithm (Yang et al., 2004). The basic idea is that, in quantum theory, the minimum unit carrying information is a bit, that can be in any superposition of state 0 and 1. let R the swarm size and S the particle's length, we define the following quantum particle vector $V = [V_1, V_2, \dots, V_R]$ with $V_i = [v_{i1}, v_{i2}, \dots, v_{iS}]$ with $0 \leq v_{ij} \leq 1$ to represent the probability of the bit j of the particle i to be 0. Denoting with $X = [X_1, X_2, \dots, X_R]$ with $X_i = [x_{i1}, x_{i2}, \dots, x_{iS}]$ where $x_{ij} = \{0, 1\}$ a position vector associated with V , to update the position we made use of the following rule 1:

"for each v generate a random number ρ in $[0,1]$ and a $\text{sigmoid}(v)$ using equation (12), then if $\rho < \text{sigmoid}(v)$ then $x = 1$ else $x = 0$ ".

The equations for updating processes are modified as follows:

$$v_{ij}(k+1) = w \cdot v_{ij}(k) + c_1 \cdot v_{ijlocbest}(k) + c_2 \cdot v_{ijglobest}(k) \quad (18.9)$$

$$v_{ijglobest}(k) = \alpha \cdot x_{ijglobest}(k) + \beta \cdot (1 - x_{ijglobest}(k)) \quad (18.10)$$

$$v_{ijlocbest}(k) = \alpha \cdot x_{ijlocbest}(k) + \beta \cdot (1 - x_{ijlocbest}(k)) \quad (18.11)$$

$$\text{sigmoid}(v_{ij}) = (1 + \exp(-v_{ij}))^{-1} \quad (18.12)$$

where $\alpha + \beta = 1, 0 < \alpha, \beta < 1, 0 < w, c_1, c_2 \leq 2$.

18.4 Binary particle swarm optimization for double auction market

To maximize the collective utility using a Binary PSO algorithm (BPSO), we solved the following quadratic programming problem:

$$\max \sum_{i=1}^m \sum_{j=1}^n q_{ij} x_{ij} (b_i - s_j) \quad (18.13)$$

Where

$$\sum_{i=1}^m q_{ij} x_{ij} \leq Y_j \quad \forall j \quad (18.14)$$

$$\sum_{j=1}^n q_{ij} x_{ij} \leq Z_i \quad \forall i \quad (18.15)$$

with

$$x_{ij} = \{0, 1\}, \text{ and } q_{ij} \geq 0 \text{ for every } i, j. \quad (18.16)$$

In the objective function we have two sets of variables x_{ij} and q_{ij} . We introduced a binary variable x_{ij} that is equal to 1 if there is a transaction between buyer i and seller j independently by the number of units, q_{ij} , each agent should purchase/sell. Then we can solve our problem as an assignment problem. It is possible summarize the algorithm as follows:

1. Random initialization of the particles (candidate solutions);
2. Evaluation of their fitness using equation (13);
3. Updating of particles velocity and position using equations (9)-(12) and rule 1 until a maximal number of iterations.

18.5 Results

Validation of an analytical method through a series of experiments demonstrates that the method is suitable for its intended purpose. The algorithm was implemented on a Pentium Centrino Duo in Matlab 6.5. To realize a robust algorithm, it was very important to obtain proper values for the control parameters according to problem characteristics. To do this, the algorithm was run from random initial solutions under many different parameter settings. After many experiments, the parameters were set as in table 1.

Table 18.1. control parameters for BPSO algorithm

parameter value	
α	0.7
β	0.3
w	1
c_1	2
c_2	2
F	0.729

Problems of different size were solved to test the efficiency of the proposed algorithm and to evaluate its performance. Each instance was run 20 times from different initial solutions. Our result (average of 20 running) is displayed in Fig.1 and proves as, after about 20 iterations, the process reaches stability. Then we proposed the information entropy for measuring the similarity convergence among the particle vectors. We calculated the conditional probability that value 1 happens at the bit j given the total number of bits that take value 1 in the entire swarm as follows:

$$prob_j = \frac{\sum_{i=1}^R x_{ij}}{\sum_{i=1}^R \sum_{h=1}^S x_{ih}}. \quad (18.17)$$

Then the particle vector entropy can be defined as:

$$E = - \sum_{j=1}^S (prob_j \cdot \log_2(prob_j)). \quad (18.18)$$

If the particle vectors are highly similar to one another, the values of those non-zero $prob_j$ would be high, resulting in less entropy value. Fig.2 depicts the particle vector entropy versus the number of iterations. This is to testify that the swarm evolves to the same optimization goal and the best solution is not obtained by chance due to a lucky particle. Obviously, this is because each particle can adjust its direction and velocity depending on the results of its neighbors. The results are quite promising and show that the algorithm is applicable to nonlinear problems.

18.6 Conclusions

In this paper we introduce PSO mechanism into double auction market with discriminatory bids and asks. The algorithm doesn't need complex




Fig. 18.1. Convergence graph for a 25 x 35 buyers-sellers problem




Fig. 18.2. The particle vector entropy versus the number of iterations

encoding and decoding processes and also for this is efficient in running time. Many problems are left for future research. In particular, it's necessary further theoretical analysis for getting a better BPSO convergency and to investigate how uncertainty in transactions can influence the solution's stability. The presented example are merely illustrative. However, the modeling results seem to be very promising for analysis and planning especially in markets that use fixing procedures. Future work can be extended to self-adaptive algorithms for dynamic environments. Moreover, PSO method can be defined as a evolutionary technique and it encourages studies about social interaction among peoples. Some features applications of PSO could focus the own attention on a study concerning special methodologies for prey-predator problems or for finding the best path to reach a target in robotic field (DiGesù et al., 2006).

Acknowledgement. This paper has been partially supported by 60% project and PI2S2 project by COMETA Consortium. Moreover special thanks have to be given to reviewers for the useful comments.

References

- R. Das, J. Hanson, J. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 1169–1176, 2001.
- V. DiGesù, B. Lenzitti, G. Lo Bosco, and Tegolo D. Comparison of different cooperation strategies in the prey-predator problem. In *Proceedings of IEEE International Workshop on Computer Architecture for Machine Perception and Sensing*, pages 105–111, 2006.
- D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1991.
- S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22(1):1–29, 1998.
- J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceeding of the 1995 IEEE International Conference on Neural Network*, pages 1942–1948, 1995.
- R. McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450, 1992.
- T. Sandholm and S. Suri. Market clearability. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 1145–1151, 2001.

- M. Satterthwaite and S. Williams. The rate of convergence to efficiency in the buyer bid double auction as the market becomes large. *The Review of Economic Studies*, 56(4):477–498, 1989.
- S.W. Yang, M. Wang, and L.C. Jiao. A quantum particle swarm optimization. In *Proceeding of the 2004 IEEE Congress on Evolutionary Computation*, pages 320–324, 2004.

Better-Reply Strategies with Bounded Recall

Andriy Zapechelnyuk

Center for Rationality, the Hebrew University, Israel
andriy@vms.huji.ac.il

19.1 Introduction

In every (discrete) period of time a decision maker (for short, an *agent*) makes a decision and, simultaneously, Nature selects a state of the world. The agent receives a payoff which depends on both his action and the state. Nature's behavior is ex-ante unknown to the agent, it may be as simple as an i.i.d. environment or as sophisticated as a strategic play of a rational player. The agent's objective is to select a sequence of decisions which guarantees to him the long-run average payoff as large as the best-reply payoff against Nature's empirical distribution of play, *no matter what Nature does*. A behavior rule of the agent which fulfills this objective is called *universally consistent*¹: the rule is "consistent" if it is optimized against the empirical play of Nature; the word "universally" refers to its applicability to *any* behavior of Nature.

A range of problems can be described within this framework. One example, known as the *on-line decision problem*, deals with predicting a sequence of states of Nature, where at every period t the agent makes a prediction based on information known before t . The classical problem of predicting the sequence of 0's and 1's with "few" mistakes has been a subject of study in statistics, computer science and game theory for more than 40 years. In a more general problem, an agent's goal is to predict a sequence of states of Nature at least as well as the best expert from a given pool of experts²(see Cesa-Bianchi et al., 1997; Freund and Schapire, 1996; Littlestone and Warmuth, 1994; Vovk, 1998). Another

¹ The term "universal consistency" is due to Fudenberg and Levine (1995).

² By an "expert" we understand a given deterministic on-line prediction algorithm. Thus, "to do as well as the best expert" means to make predictions, on average, as close to the true sequence of states as the best of the given prediction algorithms.

example is *no-regret learning* in game-theory. A regret³ of an agent for action a is his average gain had he played constant action a instead of his actual past play; the agent's goal is to play a sequence of actions so that he has "no regrets" (e.g., Cesa-Bianchi and Lugosi, 2003; Foster and Vohra, 1999; Fudenberg and Levine, 1995; Hannan, 1957; Hart and Mas-Colell, 2000, 2001).

Action a is called a *better reply* to Nature's empirical play if the agent could have improved upon his average past play had he played action a instead of what he actually played in the past. In this paper, we assume that in every period the agent plays a *better reply* to Nature's past play. The better-reply play is a natural adaptive behavior of an unsophisticated, myopic, non-Bayesian decision maker. The class of better-reply strategies encompasses a big variety of behavior rules, such as fictitious play and smooth fictitious play.⁴ Hart and Mas-Colell (2000)'s "no-regret" strategy of playing an action with probability proportional to the regret for that action; the logistic (or exponential-weighted) algorithms used in both game theory and computer science (see Cesa-Bianchi et al., 1997; Freund and Schapire, 1996; Littlestone and Warmuth, 1994; Vovk, 1998); the polynomial (l_p -norm) "no-regret" strategies and potential-based strategies of Hart and Mas-Colell (2001) (see also Cesa-Bianchi and Lugosi, 2003).

The agent is said to have m -recall if he is capable of remembering the play of m last periods, and the empirical frequency of Nature's play to which the agent "better-replies" is the simple average across the time interval not exceeding the last m periods. A special case of agent with perfect recall ($m = \infty$) is well studied in the literature, and universally consistent better-reply strategies of an agent with perfect recall are well known (see Cesa-Bianchi and Lugosi, 2003; Foster and Vohra, 1999; Hannan, 1957; Hart and Mas-Colell, 2000, 2001).

The question that we pose in this paper is whether there are better-reply strategies for an agent with *bounded recall* ($m < \infty$) which are (nearly) universally consistent if the agent has sufficiently large length of recall. We show that an agent with long enough recall can approach the best reply to any i.i.d. environment. However, by a simple example we demonstrate that an agent cannot optimize his average play against general (non-i.i.d.) environment, no matter how long (yet, bounded)

³ This paper deals with the simplest notion of regret known as *external* (or *unconditional*) regret (see, e.g., Foster and Vohra, 1999).

⁴ In the original (Fudenberg and Levine, 1995)'s definition, the smooth fictitious play is *not* a better-reply strategy; however, certain versions of it, such as the l_p -norm strategy with large p and the exponential-adjustment strategy with small η , are better-reply strategies.

recall he has and no matter what better-reply strategy he employs. Formally, we say that a family of better-reply strategies with bounded recall is *asymptotically universally consistent* if for every $\varepsilon > 0$ and every sufficiently large $m = m(\varepsilon)$ an agent with recall length m has an ε -universally consistent strategy in this family. We prove the following statement.

There is no family of bounded-recall better-reply strategies which is asymptotically universally consistent.

The statement is proven by a counterexample. We construct a game where if Nature plays a certain form of the fictitious play, then, regardless of what better-reply strategy the agent uses, for every agent's recall length m the limit play forms a cycle. The average payoff of the agent along the cycle is bounded away from the best-reply payoff by a uniform bound for all m . Intuitively, the reason for a cyclical behavior is that in every period t the agent's learns a new observation, a pair (a_t, ω_t) , and forgets another observation, (a_{t-m}, ω_{t-m}) . An addition of the new observation shifts, in expectation, the agent's average payoff (across the last m periods) in a "better" direction, however, the loss of (a_{t-m}, ω_{t-m}) shifts it in an arbitrary direction. Since the magnitude of the two effects is the same, $1/m$, it may lead to a cyclical behavior of the play. Note that with unbounded recall, $m = \infty$, the second effect does not exist, i.e., the agent does not forget anything, and, consequently, a cyclical behavior is not possible.

A closely related work of Lehrer and Solan (2003) assumes bounded recall of agents and studies a certain form of a better-reply behavior. Lehrer and Solan describe an ε -universally consistent strategy, where the agent periodically "wipes out" his memory. Comparison of this work with our results brought into our paper an interesting insight that "better memory multiplies regrets": an agent can achieve a better average payoff by not using, or deliberately forgetting some information about the past (see Section 19.6 for further discussion).

19.2 Preliminaries

In every discrete period of time $t = 1, 2, \dots$ a decision maker (or an agent) chooses an action, a_t , from a finite set A of actions, and Nature chooses a state, ω_t , from a finite set Ω of states. Let $u : A \times \Omega \rightarrow \mathbb{R}$ be the agent's payoff function; $u(a_t, \omega_t)$ is the agent's payoff at period t . Denote by $h_t := ((a_1, \omega_1), \dots, (a_t, \omega_t))$ the history of play up to t . Let

$H_t = (A \times \Omega)^t$ be the set of histories of length t and let $H = \bigcup_{t=1}^{\infty} H_t$ be the set of all histories.

Let $p : H \rightarrow \Delta(A)$ and $q : H \rightarrow \Delta(\Omega)$ be behavior rules of the agent and Nature, respectively. For every period t , we will denote by $p_{t+1} := p(h_t)$ the next-period mixed action of the agent and by $q_{t+1} := q(h_t)$ the next-period distribution of states of Nature. A pair (p, q) and an initial history h_{t_0} induce a probability measure over H_t for all $t > t_0$.

We assume that the agent does not know q , that is, he plays against an unknown environment. We consider better-reply behavior rules, according to which the agent plays actions which are “better” than his actual past play against the observed empirical behavior of Nature. Formally, for every $a \in A$ and every period t define $R_t^m(a) \in \mathbb{R}_+$ as the average gain of the agent had he played a over the last m periods instead of his actual past play. Namely, let⁵

$$R_t^m(a) = \left[\frac{1}{m} \sum_{k=t-m+1}^t (u(a, \omega_k) - u(a_k, \omega_k)) \right]^+ \quad \text{for all } t \geq m$$

and

$$R_t^m(a) = \left[\frac{1}{t} \sum_{k=1}^t (u(a, \omega_k) - u(a_k, \omega_k)) \right]^+ \quad \text{for all } t < m.$$

We will refer to $R_t^m(a)$ as the agent’s *regret for action a* .

The parameter $m \in \{1, 2, \dots\} \cup \{\infty\}$ is the agent’s length of recall. An agent with a specified m is said to have *m -recall*. We shall distinguish the cases of *perfect recall* ($m = \infty$) and *bounded recall* ($m < \infty$).

Consider an agent with m -recall. Action a is called a better reply to Nature’s empirical play if the agent could have improved upon his average past play had he played action a instead of what he actually played in the last m periods.

Definition 1. Action $a \in A$ is a better-reply action if $R_t^m(a) > 0$.

A behavior rule is called a better-reply rule if the agent plays only better-reply actions, as long as there are such.

Definition 2. Behavior rule p is a better-reply rule if for every period t , whenever $\max_{a \in A} R_t^m(a) > 0$,

$$R_t^m(a) = 0 \quad \Rightarrow \quad p_{t+1}(a) = 0, \quad a \in A.$$

⁵ We write $[x]^+$ for the positive part of a scalar x , i.e., $[x]^+ = \max\{0, x\}$.

The focus of our study is how well better-reply rules perform against an unknown, possibly, hostile environment. To assess performance of a behavior rule, we use Fudenberg and Levine (1995)'s criterion of ε -universal consistency defined below.

An agent's behavior rule p is said to be *consistent with q* if the agent's long-run average payoff is at least as large as the best-reply payoff to the average empirical play of Nature which plays q .

Definition 3. Let $\varepsilon > 0$. A behavior rule p of the agent with m -recall is ε -consistent with q if for every initial history h_{t_0} there exists T such that for every⁶ $t \geq T$

$$\Pr_{(p,q,h_{t_0})} \left[\max_{a \in A} R_t^\infty(a) < \varepsilon \right] > 1 - \varepsilon.$$

A behavior rule p is consistent with q if it is ε -consistent with q for every $\varepsilon > 0$.

Let $\mathcal{Q}$ be the class of all behavior rules. An agent's behavior rule p is said to be *universally consistent* if it is consistent with *any* behavior of Nature.

Definition 4. A behavior rule p of the agent with m -recall is $(\varepsilon-)$ universally consistent if it is $(\varepsilon-)$ consistent with q for every $q \in \mathcal{Q}$.

19.3 Perfect recall and prior results

Suppose that the agent has perfect recall ($m = \infty$). This case has been extensively studied in the literature, starting from Hannan (1957), who proved the following theorem.⁷

Theorem 19.3.1 (Hannan, 1957) *There exists a better-reply rule which is universally consistent.*

Hart and Mas-Colell (2000) showed that the following rule is universally consistent:

$$p_{t+1}(a) := \begin{cases} \frac{R_t^\infty(a)}{\sum_{a' \in A} R_t^\infty(a')}, & \text{if } \sum_{a' \in A} R_t^\infty(a') > 0, \\ \text{arbitrary,} & \text{otherwise.} \end{cases} \quad (19.1)$$

⁶ $\Pr_{(p,q,h)}[E]$ denotes the probability of event E induced by strategies p and q , and initial history h .

⁷ The statements of theorems of Hannan (1957); Hart and Mas-Colell (2000) presented in this section are sufficient for this paper, though the authors obtained stronger results.

According to this rule, the agent assigns probability on action a proportional to his regret for a ; if there are no regrets, his play is arbitrary. This result is based on Blackwell (1956)'s Approachability Theorem. We shall refer to p in (19.1) as the *Blackwell strategy*.

The above result has been extended by Hart and Mas-Colell (2001) as follows. A behavior rule p is called a (*stationary*) *regret-based rule* if for every period t the agent's next-period behavior depends only on the current regret vector. That is, for every history h_t , the next-period mixed action of the agent is a function of $R_t^\infty = (R_t^\infty(a))_{a \in A}$ only: $p_{t+1} = \sigma(R_t^\infty)$. Hart and Mas-Colell proved that among better-reply rules, all “well-behaved” stationary regret-based rules are universally consistent.

Theorem 19.3.2 (Hart and Mas-Colell, 2001) *Suppose that a better-reply rule p satisfies the following:*

- (i) p is a stationary regret-based rule given for every t by $p_{t+1} = \sigma(R_t^\infty)$; and
- (ii) *There exists a continuously differential potential $P : \mathbb{R}_+^{|A|} \rightarrow \mathbb{R}_+$ such that $\sigma(x)$ is positively proportional to $\nabla P(x)$ for every $x \in \mathbb{R}_+^{|A|}$, $x \neq 0$.*

Then p is universally consistent.

The class of universally consistent behavior rules (or “no regret” strategies) which satisfy conditions of Theorem 19.3.2 includes the logistic (or exponential adjustment) strategy given for every t and every $a \in A$ by

$$p_{t+1}(a) = \frac{\exp(\eta R_t^m(a))}{\sum_{b \in A} \exp(\eta R_t^m(b))},$$

$\eta > 0$, used by Cesa-Bianchi et al. (1997); Freund and Schapire (1996); Littlestone and Warmuth (1994); Vovk (1998) and others; the smooth fictitious play⁸; the polynomial (l_p -norm) strategies and other strategies based on a separable potential (Cesa-Bianchi and Lugosi, 2003; Hart and Mas-Colell, 2001).

19.4 Bounded recall and i.i.d. environment

The previous section shows that the universal consistency can be achieved for agents with perfect recall. Considering the perfect recall

⁸ See footnote 4.

as the limit of m -recall as $m \rightarrow \infty$, one may wonder whether the universal consistency can be approached by bounded-recall agents with sufficiently large m .

We start with a result that establishes existence of better-reply rules which are consistent with any i.i.d. environment. Nature's behavior rule q is called an i.i.d. rule if $q_t = q_{t'}$ for all t, t' , independently of the history. Let $\mathcal{Q}_{i.i.d.} \subset \mathcal{Q}$ be the set of all i.i.d. behavior rules. An agent's behavior rule p is said to be *i.i.d. consistent* if it is consistent with any i.i.d. behavior of Nature.

Definition 5. *A behavior rule p of the agent with m -recall is $(\varepsilon-)$ i.i.d. consistent if it is $(\varepsilon-)$ consistent with q for every $q \in \mathcal{Q}_{i.i.d.}$.*

Denote by $\mathcal{P}^m$ the class of all better-reply rules for an agent with m -recall, $m \in \mathbb{N}$. Consider an indexed family of better-reply rules $\mathbf{p} = (p^1, p^2, \dots)$, where $p^m \in \mathcal{P}^m$, $m \in \mathbb{N}$.

Definition 6. *A family $\mathbf{p}$ is asymptotically i.i.d consistent if for every $\varepsilon > 0$ there exists m such that for every $m' \geq m$ rule $p^{m'}$ is ε -i.i.d. consistent.*

Theorem 19.4.1 *There exists a family $\mathbf{p}$ of better-reply rules which is asymptotically i.i.d. consistent.*

Proof Let $q^* \in \Delta(\Omega)$ and suppose that $q_t = q^*$ for all t . Denote by $\bar{q}_t^m$ the empirical distribution of Nature's play over the last m periods,

$$\bar{q}_t^m(\omega) = \frac{1}{m} |\{k \in \{t - m + 1, \dots, t\} : \omega_k = \omega\}|, \quad \omega \in \Omega.$$

Suppose that the agent plays the fictitious play with m -recall. Namely, the agent's next-period play, p_{t+1}^m , assigns probability 1 on an action in $\arg \max_{a \in A} u(a, \bar{q}_t^m)$, ties are resolved arbitrarily. Thus, the agent plays in every period a best reply to the average realization of m i.i.d. random variables with mean q^* . Since $\max_{a \in A} u(a, x)$ is continuous in x for $x \in \Delta(\Omega)$, the Law of Large Numbers implies that in every period the agent obtains an expected payoff which is ε_m -close to the best reply payoff to q^* with probability at least $1 - \varepsilon_m$, with $\varepsilon_m \rightarrow 0$ as $m \rightarrow \infty$. $\square$

19.5 A negative result

In this section we demonstrate that an agent with bounded recall cannot guarantee his play to be ε -optimized against the empirical play of

Nature, no matter how large recall length he has and no matter what better-reply rule he uses.

Definition 7. Family $\mathbf{p} = (p^1, p^2, \dots)$ of better-reply rules is asymptotically universally consistent if for every $\varepsilon > 0$ there exists m such that for every $m' \geq m$ rule $p^{m'}$ is ε -universally consistent.

Theorem 19.5.1 *There is no family of better-reply rules which is asymptotically universally consistent.*

The theorem is proven by a counterexample.

	L	M	R
U	1,0	0,1	0, $\frac{3}{4}$
D	0,1	1,0	0, $\frac{3}{4}$

Fig. 19.1.

Consider a repeated game Γ with the stage game given by Fig. 19.1, where the row player is the agent and the column player is Nature. For every m denote by p^m and q^m be the behavior rules of the agent and Nature, respectively. We shall show that for every $m_0 \in \mathbb{N}$ there exists $m \geq m_0$ such that the following holds.

Suppose that the agent with recall length m and Nature play game Γ . Then for every agent's better-reply rule p^m there exist behavior rule q^m of Nature, initial history h_{t_0} and period T such that for all $t \geq T$

$$\Pr_{(p^m, q^m, h_{t_0})} \left[\max_{a \in \{U, D\}} R_t^\infty(a) \geq \frac{1}{32} \right] \geq \frac{1}{32}.$$

Let $M = \{4j + 2 | j = 2, 3, \dots\}$. For every $m \in M$, let p^m be an arbitrary better-reply rule, and let q^m be the fictitious play with m -recall. Namely, denote by u_N the payoff function of Nature as given by Fig. 19.1, and denote by $\bar{p}_t$ the empirical distribution of the agent's play over the last m periods,

$$\bar{p}_t(a) = \frac{1}{m} |k \in \{t - m + 1, \dots, t\} : a_k = a|, \quad a \in A.$$

Then q_{t+1}^m assigns probability 1 to a state in $\arg \max_{\omega \in \{L, M, R\}} u_N(\bar{p}_t, \omega)$ (ties are resolved arbitrarily). Let P^m be the Markov chain with state space $H^m := (A \times \Omega)^m$ induced by p^m and q^m and an initial state h_{t_0} . A history of the last m periods, $h_t^m \in H^m$ will be called, for short, *history at t* . Denote by $H_C^m \subset H^m$ the set of states generated along the




Fig. 19.2. Closed cycle of Markov chain P^m .

following cycle (Fig. 19.2). The cycle has four phases. In two phases labeled (U,R) and (D,R), the play is deterministic, and the duration of each phase is exactly $m/2$ periods. In the two other phases, the play may randomize between two profiles (one written above the other), and the duration of each phase is $m/2$ or $m/2 + 1$ periods. First, we show that this cycle is closed in P^m , i.e., $h_t^m \in H_C^m$ implies $h_{t'}^m \in H_C^m$ for every $t' > t$.

Lemma 19.5.2 *For every $m \in M$, the set H_C^m is closed in P^m .*

The proof is in the Appendix.

Next, we show that the average regrets generated by this cycle are bounded away from zero by a uniform bound for all m .

Lemma 19.5.3 *For every $m \in M$, if $h_{t_0} \in H_C^m$, then there exists period T such that for all $t \geq T$*

$$\Pr_{(p^m, q^m, h_{t_0})} \left[\max_{a \in \{U, D\}} R_t^\infty(a) \geq \frac{1}{32} \right] \geq \frac{1}{32}.$$

The proof is in the Appendix. Lemmata 19.5.2 and 19.5.3 entail the statement of Theorem 19.5.1.

Remark 1 In the proof of Theorem 19.5.1, Nature plays the fictitious play with m -recall, which is a better-reply strategy for every m . Consequently, an agent with bounded recall cannot guarantee a nearly optimized behavior even if Nature's behavior is constrained to be in the class of better-reply strategies.

Remark 2 The result can be strengthened as follows. Suppose that whenever an agent has no regrets, then he plays a fully mixed action, i.e.,

$$\max_{a' \in A} R_t^m(a') = 0 \Rightarrow p_{t+1}^m(a) > 0 \text{ for all } a \in A. \quad (19.2)$$

The next lemma shows that if in game Γ the agent plays a better-reply strategy p^m which satisfies (19.2) and Nature plays the fictitious play with m -recall, then the Markov chain P^m converges to the cycle H_C^m regardless of an initial history. Thus the above negative result is not an isolated phenomenon, it is not peculiar to a small set of initial histories.

Lemma 19.5.4 *For every $m \in M$, if p^m satisfies (19.2), then for every initial history h_{t_0} the process P^m converges to H_C^m with probability 1.*

The proof is in the Appendix.

To see that the statement of Lemma 19.5.4 does not hold if p^m fails to satisfy (19.2), consider again game Γ with the agent playing a better-reply strategy p^m and Nature playing the fictitious play with m -recall, q^m . In addition, suppose that whenever $\max_{a' \in A} R_t^m(a') = 0$, $p_{t+1}^m(U) = 1$ if t is odd and 0 if t is even. Let t be even and let h_t consist of alternating (UR) and (DR). Clearly, $R_t^m(U) = R_t^m(D) = 0$, and Nature's best reply is R, thus, $q_{t+1}(R) = 1$. The following play is deterministic, alternating between (UR) and (DR) forever.

19.6 Concluding remarks

We conclude the paper with a few remarks.

1. Why does the better-reply play of an agent with bounded recall fail to exhibit a (nearly) optimized behavior (against Nature's empirical play)?

For every $a \in A$ denote by $v_t(a)$ the one-period regret for action a ,

$$v_t(a) = u(a, \omega_t) - u(a_t, \omega_t),$$

and let $v_t = (v_t(a))_{a \in A}$. Since $R_{t-1}^m = \frac{1}{m} \sum_{k=t-m}^{t-1} v_k$, we can consider how the regret vector changes from period $t-1$ to period t :

$$R_t^m = R_{t-1}^m + \frac{1}{m} v_t - \frac{1}{m} v_{t-m}.$$

Since the play at period t is a better reply to the empirical play over time interval $t-m, \dots, t-1$, the term $\frac{1}{m} v_t(a)$ shifts the regret vector, in expectation, towards zero, however, the term $-\frac{1}{m} v_{t-m}$ shifts the regret vector in an arbitrary direction. A carefully constructed example, as in Section 19.5, causes the regret vector to display a cyclical behavior.

2. The following model was introduced by Lehrer and Solan (2003). Suppose that the agent has bounded recall m . Divide the time into blocks of size m : the first block contains periods $1, \dots, m$, the second

block contains periods $m + 1, \dots, 2m$, etc. Let $n(t)$ be the first period of the current block,⁹ $n(t) = m \lceil t/m \rceil + 1$. The agent's regret for action $a \in A$ is defined by

$$\hat{R}_t^m(a) = \frac{1}{t - n(t) + 1} \sum_{\tau=n(t)}^t (u(a, \omega_\tau) - u(a_\tau, \omega_\tau)). \quad (19.3)$$

That is, $\hat{R}_t^m(a)$ is the agent's average increase in payoff had he played a constantly instead of his actual past play *within* in the current block. Let p^* be the Blackwell strategy (19.1) with better replies computed relative to (19.3). Clearly, this strategy can be implemented by the agent with m -recall. However, the agent behaves as if he remembers only the history of the current block, and at the beginning of a new block he “wipes out” the content of his memory. Notice that the induced probability distribution over histories within every block is identical between blocks and equal to the probability distribution over histories within first m periods in the model with a perfect-recall agent. The Blackwell (1956)'s Approachability Theorem (which is behind the result of Hart and Mas-Colell (2000) on the universal consistency of p^*) gives the rate of convergence of $1/\sqrt{t}$, hence, within each block the agent can approach $1/\sqrt{m}$ -best reply to the empirical distribution of Nature's play.

This result is a surprising contrast to the counterexample in Section 19.5. It shows that *an agent can achieve better average payoff by not using, or deliberately forgetting some information about the past*. Indeed, according to the example presented in Section 19.5, if the agent uses full information that he remembers, the play may eventually enter the cycle with far-from-optimal behavior, no matter with what initial history he starts.

3. Hart and Mas-Colell (2001) used a slightly different notion of better reply. Consider an agent with perfect recall and define for every period t and every $a \in A$

$$D_t^m(a) = \frac{1}{t} \sum_{k=1}^t (u(a, \omega_k) - u(a_k, \omega_k)).$$

Note that $R_t^m(a) = [D_t^m(a)]^+$. Action a is a *strict* better reply (to the empirical distribution of Nature's play) if $D_t^m(a) > 0$ and it is a *weak* better reply if $D_t^m(a) \geq 0$. According to Hart and Mas-Colell, behavior rule p is a better-reply rule if whenever there exist actions which are weak better replies, only such actions are played; formally, whenever $\max_{a \in A} D_t^m(a) \geq 0$,

⁹ $\lceil x \rceil$ denotes a number x rounded up to the nearest integer.

$$D_t^m(a) < 0 \Rightarrow p_{t+1}(a) = 0, \quad a \in A.$$

The definition of a better-reply rule used in this paper is the same as Hart and Mas-Colell's, except that the word “weak” is replaced by “strict”; formally, whenever $\max_{a \in A} D_t^m(a) > 0$,

$$D_t^m(a) \leq 0 \Rightarrow p_{t+1}(a) = 0, \quad a \in A.$$

These notions are very close, and one does not imply the other. To the best of our knowledge, all specific better-reply rules mentioned in the literature satisfy both notions of better reply. It can be verified that our results remain intact with either notion.

Appendix

A-1 Proof of Lemma 19.5.2.

Let $k = \frac{m-2}{4}$. Denote by z_t the empirical distribution of play, that is, for every $(a, \omega) \in A \times \Omega$, $z_t(a, \omega)$ is the frequency of (a, ω) in the history at t ,

$$z_t(a, \omega) := \frac{1}{m} |\{\tau \in \{t - m + 1, \dots, t\} : (a_\tau, \omega_\tau) = (a, \omega)\}|.$$

Let ζ_t be is the frequency of play of U in the last m periods, $\zeta_t = z_t(\text{U}, \text{L}) + z_t(\text{U}, \text{M}) + z_t(\text{U}, \text{R})$.

Fact 1. For every period t ,

$$\omega_{t+1} = \begin{cases} \text{L, if } \zeta_t < \frac{1}{4}, \\ \text{M, if } \zeta_t > \frac{3}{4}, \\ \text{R, if } \frac{1}{4} < \zeta_t < \frac{3}{4}. \end{cases}$$

Proof. Note that

$$\begin{aligned} u_N(\bar{p}_t, \text{L}) &= z_t(\text{D}, \text{L}) + z_t(\text{D}, \text{M}) + z_t(\text{D}, \text{R}) = 1 - \zeta_t, \\ u_N(\bar{p}_t, \text{M}) &= z_t(\text{U}, \text{L}) + z_t(\text{U}, \text{M}) + z_t(\text{U}, \text{R}) = \zeta_t, \\ u_N(\bar{p}_t, \text{R}) &= \frac{3}{4}. \end{aligned}$$

Since Nature plays fictitious play, at $t + 1$ it selects

$$\omega_{t+1} \in \arg \max_{\omega \in \{\text{L}, \text{M}, \text{R}\}} u_N(\bar{p}_t, \omega).$$

Note that ties never occur, since $m \in M$ and ζ_t is a multiple of $\frac{1}{m}$, thus $\zeta_t \neq \frac{1}{4}$ or $\frac{3}{4}$. $\square$ **Fact 2.** Suppose that $h_t^m \in H_C^m$ such that t is




Fig. 19.3. Three forms of the (U,M)/(D,M) phase.

the last period of the (D,R) phase, and suppose that the (U,M)/(D,M) phase preceding the (D,R) phase has form (a), (b) or (c), as shown in Fig. 19.3. Then the play for the next $m/2$, $m/2 + 1$, or $m/2 + 2$ periods constitute the full cycle as shown in Fig. 19.2, where phases (D,L)/(U,L) and (U,M)/(D,M) have forms¹⁰ (a), (b) or (c).

Proof. Suppose that h_t^m contains $m/2$ (D,R)'s, preceded by the (U,M)/(D,M) phase in form (a), (b), or (c). We shall show that the play in the next $m/2$ or $m/2 + 1$ periods constitute phase (D,L)/(U,L) in form (a), (b) or (c), followed by $m/2$ (U,R)'s. Once this is established, by considering the last period of phase (U,R) and repeating the arguments, we obtain Fact 2.

Case 1. Phase (U,M)/(D,M) preceding phase (D,R) has form (a) or (b). Note that whether the (U,M)/(D,M) phase has form (a) or (b), h_t^m is the same, since it contains only $2k + 1 \equiv m/2$ last periods of the (U,M)/(D,M) phase. Let t be the last period of the (D,R) phase. We have $\zeta_t = \frac{k}{m} < \frac{1}{4}$, thus by Fact 1, $\omega_{t+1} = L$. Also,

$$R_t(U) = z_t(D,L) - z_t(D,M) = -z_t(D,M) = -\frac{k+1}{m},$$

$$R_t(D) = z_t(U,M) - z_t(U,L) = z_t(U,M) = \frac{k}{m},$$

hence $a_{t+1} = D$. Further, in every period $t+j$, $j = 1, \dots, k$, $(a_{t+j}, \omega_{t+j}) = (D,L)$ is played and $(a_{t+j-m}, \omega_{t+j-m}) = (U,M)$ disappears from the history. At period $t+k$ we have

¹⁰ The forms of the (D,L)/(U,L) phase are symmetric to those of (U,M)/(D,M), obtained by replacement of (U,M) by (D,L) and (D,M) by (U,L).

$$R_{t+k}(U) = z_{t+k}(D,L) - z_{t+k}(D,M) = \frac{k}{m} - \frac{k+1}{m} = -\frac{1}{m},$$

$$R_{t+k}(D) = z_{t+k}(U,M) - z_{t+k}(U,L) = 0 - 0 = 0.$$

There are no regrets, and therefore both (U,L) and (D,L) may occur at $t+k+1$. Suppose that (D,L) occurs. Since $(a_{t+k-m}, \omega_{t+k-m}) = (D,M)$, it will disappear from the history at $t+k+1$, so, we have

$$R_{t+k+1}(U) = \frac{k+1}{m} - \frac{k}{m} = \frac{1}{m},$$

$$R_{t+k+1}(D) = 0 - 0 = 0,$$

and (U,L) occurs in periods $k+2, \dots, 2k+2$, until we reach $\zeta_{t+2k+2} = \frac{k+1}{m} > 1/4$. Thus, the phase (D,L)/(U,L) has $k+1$ (D,L)'s, then $k+1$ (U,L)'s, i.e., it takes form (b). If instead at $t+k+1$ action profile (U,L) occurs, then

$$R_{t+k+1}(U) = \frac{k}{m} - \frac{k}{m} = 0,$$

$$R_{t+k+1}(D) = 0 - \frac{1}{m} = -\frac{1}{m},$$

and, again, there are no regrets and both (U,L) and (D,L) may occur at $t+1$. If (U,L) occurs, then

$$R_{t+k+2}(U) = \frac{k}{m} - \frac{k-1}{m} = \frac{1}{m},$$

$$R_{t+k+1}(D) = 0 - \frac{2}{m} = -\frac{2}{m},$$

and (U,L) occurs in periods $k+3, \dots, 2k+1$, until we reach $\zeta_{t+2k+1} = \frac{k+1}{m} > 1/4$. Thus, the phase (D,L)/(U,L) has k (D,L)'s, then $k+1$ (U,L)'s, i.e., it takes form (a). Finally, if at $t+k+2$ (D,L) occurs, then

$$R_{t+k+1}(U) = \frac{k+1}{m} - \frac{k-1}{m} = \frac{2}{m},$$

$$R_{t+k+1}(D) = 0 - \frac{1}{m} = -\frac{1}{m},$$

and (U,L) occurs in periods $k+3, \dots, 2k+2$, until we reach $\zeta_{t+2k+2} = \frac{k+1}{m} > 1/4$. Thus, the phase (D,L)/(U,L) has k (D,L)'s, then single (U,L), then single (D,L), and then k (U,L)'s, i.e., it takes form (c).

Case 2. Phase (U,M)/(D,M) preceding phase (D,R) has form (c). Then, similarly to Case 1, we have $\zeta_t = \frac{k}{m} < \frac{1}{4}$, and (D,L) is deterministically played $k+1$ times, until

$$R_{t+k+1}(U) = z_{t+k+1}(D,L) - z_{t+k+1}(D,M) = \frac{k+1}{m} - \frac{k}{m} = \frac{1}{m},$$

$$R_{t+k+1}(D) = z_{t+k+1}(U,M) - z_{t+k+1}(U,L) = 0 - 0 = 0.$$

After that, (U,L) is played in periods $k+2, \dots, 2k+2$, until we reach $\zeta_{t+2k+2} = \frac{k+1}{m} > 1/4$. Thus, the phase (D,L)/(U,L) has $k+1$ (D,L)'s and then $k+1$ (U,L)'s, i.e., it takes form (b).

Let $t_1 = t + 2k + 1$ if the phase (D,L)/(U,L) had form (a) and $t_1 = t + 2k + 2$ if (b) or (c). Notice that at the end of the phase (D,L)/(U,L) we have $z_{t_1}(U,M) = z_{t_1}(D,M) = 0$, hence

$$R_{t_1}(U) = z_{t_1}(D,L) - z_{t_1}(D,M) > 0,$$

$$R_{t_1}(D) = z_{t_1}(U,M) - z_{t_1}(U,L) < 0,$$

Thus, (U,R) is played for the next $m/2 = 2k+1$ periods, until we reach $\zeta_{t_1+m/2} = \frac{3k+2}{m} > 3/4$, and phase (U,M)/(D,M) begins. $\square$

A-2 Proof of Lemma 19.5.3.

By Lemma 19.5.2, $h_{t_0} \in H_C^m$ implies $h_t^m \in H_C^m$ for all $t > t_0$. Let $h_t^m \in H_C^m$ such that t is the period at the end of the (D,R) phase. Since the history at t contains only (U,M)/(D,M) and (D,R) phases, we have $z_t(D,L) = z_t(U,L) = 0$. Also, since at the end of the (D,R) phase the number of U in the history is $\frac{m+2}{4}$, it implies that $z_t(U,M) = \frac{1}{4} + \frac{1}{2m}$. Therefore,

$$R_t(D) = z_t(U,M) - z_t(U,L) = z_t(U,M) = \frac{1}{4} + \frac{1}{2m} \equiv C$$

For every period τ , $|R_\tau(D) - R_{\tau+1}(D)| \leq \frac{2}{m}$, therefore, in periods $t-j$ and $t+j$ the regret for D must be at least $R_t(D) - 2j/m$. Since the duration of every cycle is at most $2m+2$, the average regret for D during the cycle is at least

$$\frac{1}{2m+2} \left(C + 2 \left[\left(C - \frac{2}{m} \right) + \left(C - \frac{4}{m} \right) + \dots + \left(C - \frac{2(m/4-2)}{m} \right) \right] \right) \geq \frac{1}{2m} \left(\frac{m}{2} C - \frac{2}{m} \frac{m^2-4}{32} \right) \geq \frac{1}{32}.$$

Let γ^m be the limit frequency of periods where at least one of the regrets exceeds ε ,

$$\gamma^m = \lim_{t \rightarrow \infty} \frac{1}{t} \left| \tau \in \{1, \dots, t\} : \max_{a \in \{U,D\}} R_\tau^m(a) \geq \varepsilon \right|.$$

Clearly, $\gamma^m > \varepsilon$ implies that for all large enough t

$$\Pr_{(p^m, q^m, h_{t_0})} [\max_{a \in \{U, D\}} R_t^\infty(a) \geq \varepsilon] \geq \varepsilon.$$

Combining (19.4) with the fact that γ^m is at least as large as the average regret for D during the cycle, we obtain $\gamma^m \geq 1/32$. $\square$

A-3 Proof of Lemma 19.5.4.

We shall prove that, regardless of the initial history, some event $H_E^m \subset H^m$ occurs infinitely often, and whenever it occurs, the process reaches the cycle, H_C^m , within at most $2m$ periods with strictly positive probability. It follows that the process reaches the cycle with probability 1 from any initial history.

Fact 3. Regardless of an initial state, L and M occur infinitely often.

Proof. Suppose that M never occurs from some time t on. Then at any t

$$\begin{aligned} R_t(U) &= z_t(D, L) - z_t(D, M) = z_t(D, L) \geq 0, \\ R_t(D) &= z_t(U, M) - z_t(U, L) = -z_t(U, L) \leq 0. \end{aligned}$$

Case 1. $z_t(D, L) > 0$. Suppose that L occurred last time at $t - j$, $0 \leq j \leq m - 1$. After that U must be played with probability 1 in every period $j' = t - j + 1, \dots$, until frequency of U increases above $\frac{3}{4}$ and, by Fact 1, Nature begins playing M. Contradiction.

Case 2. $z_t(D, L) = 0$. That is, the agent has no regrets, his play is defined arbitrarily. By assumption (19.2), $p_{t+1}^m(U) > 0$, and thus there is a positive probability that U occurs sufficiently many times that the frequency of U increases above $\frac{3}{4}$ and M is played. Contradiction.

The proof that L occurs infinitely often is analogous. $\square$

Fact 4. If $\omega_t = L$ and $\omega_{t+j} = M$, then $j > \frac{m}{2}$. Symmetrically, if $\omega_t = M$ and $\omega_{t+j} = L$, then $j > \frac{m}{2}$.

Proof. Suppose that $\omega_t = L$, then by Fact 1, $\zeta_{t-1} < \frac{1}{4}$. Clearly, it requires $j > \frac{m}{2}$ periods to reach ζ_{t+j-1} greater than $\frac{3}{4}$, which is required to have $\omega_{t+j} = M$. The second part of the fact is proved analogously. $\square$

Fact 5. Regardless of an initial state, the event $\{\omega_t = L \text{ and there are no more L in } h_t^m\}$ occurs infinitely often.

Proof. By Fact 3, both L and M occur infinitely often. By Fact 4, the minimal interval of occurrence of L and M is $\frac{m}{2}$, hence if L occurs first time after M, previous occurrence of L is at least $m + 1$ periods ago. $\square$

Fact 6. Suppose that $\omega_t = L$ and there are no more L in the history. Then after $j < m$ periods we obtain $\frac{1}{4} < \zeta_{t+j} < \frac{1}{4} + \frac{1}{m}$, and with strictly positive probability $R_{t+j}(U) > 0$ and $R_{t+j}(D) \leq 0$.

Proof. We have

$$\begin{aligned} R_t(U) &= z_t(D, L) - z_t(D, M), \\ R_t(D) &= z_t(U, M) - z_t(U, L). \end{aligned}$$

By Fact 1, $\omega_t = L$ implies $\zeta_{t-1} < \frac{1}{4}$, that is, U occurs at most k times in the history at $t-1$, thus $z_t(U, M) \leq z_{t-1}(U, M) \leq \frac{k}{m}$.

Case 1. $R_t(D) > 0$ and $R_t(U) > 0$ Then both (D,L) and (U,L) may be played. Since history at $t-1$ does not contain L, regardless of what disappears from the history, we have $R_t(U)$ nondecreasing and $R_t(D)$ nonincreasing. Thus, with positive probability, both (D,L) and (U,L) are played for j periods, until we obtain $\frac{1}{4} < \zeta_{t+j} < \frac{1}{4} + \frac{1}{m}$, $R_{t+j}(U) > 0$ and $R_{t+j}(D) \leq 0$. Note that $j < \frac{3}{4}m + 1$, since by Fact 4 the interval between the last occurrence of M and the first occurrence of L is at least $m/2$, thus after period $t + m/2$ there are no M in the history, $R_{t+m/2}(U) > 0$, $R_{t+m/2}(D) < 0$, and (U,L) is played at most $k+1 = \frac{m+2}{4}$ times until the frequency of U becomes above $1/4$.

Case 2. $R_t(D) > 0$, $R_t(U) \leq 0$. Then (D,L) is played for the next $j' = (z_t(D, L) - z_t(D, M)) \cdot m + 1$ periods. At period $t + j'$ we have $R_{t+j'}(D) > 0$ and $R_{t+j'}(U) > 0$, and proceed similarly to Case 1.

Case 3. $R_t(D) \leq 0$, $R_t(U) \leq 0$. That is, the agent has no regrets, his play is defined arbitrarily. By assumption, $p_{t+1}(D) > 0$, hence there is a positive probability that (D,L) occurs for $j' = z_t(D, M) \cdot m$ periods which will yield $R_{t+j'}(U) > 0$, Case 2.

Case 4. $R_t(D) \leq 0$, $R_t(U) > 0$. Then (U,L) is played for $j = 1$ or 2 periods (depending whether $(a_t, \omega_t) = (D, L)$ or (U, L)), and we have $\frac{1}{4} < \zeta_{t+j} < \frac{1}{4} + \frac{1}{m}$, $R_{t+j}(U) = R_t(U) > 0$ and $R_{t+j}(D) < R_t(D) \leq 0$. $\square$

Using Fact 6, we can now analyze the dynamics of the process. Suppose that $\frac{1}{4} < \zeta_t < \frac{1}{4} + \frac{1}{m}$, $R_t(U) > 0$, $R_t(D) \leq 0$. Then

I. (U,R) is played in the next $j_{UR} \geq \frac{m}{2}$ periods, and we obtain $\frac{3}{4} < \zeta_{t+j_{UR}} < \frac{3}{4} + \frac{1}{m}$. Since by now M has disappeared from the history, the regrets are

$$\begin{aligned} R_{t+j_{UR}}(U) &\geq z_t(D, L) > 0, \\ R_{t+j_{UR}}(D) &\leq -z_t(U, L) \leq 0. \end{aligned}$$

II. (U,M) is played for the next $j_{UM} = k+1$ periods. Since $j_{UR} + j_{UM} \geq \frac{m}{2} + k + 1 = 3k + 1$, it implies that $z_{t+j_{UR}+j_{UM}}(U, L) \leq k$, and

$$\begin{aligned} R_{t+j_{UR}+j_{UM}}(D) &= z_{t+j_{UR}+j_{UM}}(U, M) - z_{t+j_{UR}+j_{UM}}(U, L) \\ &\geq \frac{k+1}{m} - \frac{k}{m} = \frac{1}{m} > 0. \end{aligned}$$

III. With positive probability, (D,M) is played for the next $j_{DM} = k+1$ periods, and, since by now L is not in the history, we have

$$\begin{aligned}\zeta_{t+j_{UR}+j_{UM}+j_{DM}} &= 1 - \frac{j_{DM}}{m} = \frac{3k+1}{m} < \frac{3}{4}, \\ R_{t+j_{UR}+j_{UM}+j_{DM}}(U) &= -z_{t+j_{UR}+j_{UM}+j_{DM}}(D,M) < 0, \\ R_{t+j_{UR}+j_{UM}+j_{DM}}(D) &= z_{t+j_{UR}+j_{UM}+j_{DM}}(U,M) > 0.\end{aligned}$$

Notice that at period $t+j_{UR}+j_{UM}+j_{DM}$ the last m periods correspond to phases (U,R) and (U,M)/(D,M) of the cycle (the latter is in form (b)). $\square$

Acknowledgement. I thank Dean Foster, Sergiu Hart, Eilon Solan, Tymofiy Mylovanov, Peyton Young, and participants of the seminars at the Hebrew University and Tel Aviv University for helpful comments and suggestions. I gratefully acknowledge the financial support from Lady Davis and Golda Meir Fellowship Funds, the Hebrew University.

References

- D. Blackwell. An analog of the minmax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6:1–8, 1956.
- N. Cesa-Bianchi and G. Lugosi. Potential-based algorithms in on-line prediction and game theory. *Machine Learning*, 51:239–261, 2003.
- N. Cesa-Bianchi, Y. Freund, D. Helmbold, D. Haussler, R. Shapire, and M. Warmuth. How to use expert advice. *Journal of the ACM*, 44:427–485, 1997.
- Dean Foster and Rakesh Vohra. Regret in the online decision problem. *Games and Economic Behavior*, 29:7–35, 1999.
- Yoav Freund and Robert Schapire. Game theory, on-line prediction and boosting. In *Proceedings of the Ninth Annual Conference on Computational Learning Theory*, pages 325–332, 1996.
- Drew Fudenberg and David Levine. Universal consistency and cautious fictitious play. *Journal of Economic Dynamics and Control*, 19:1065–1089, 1995.
- J Hannan. Approximation to Bayes risk in repeated play. In M. Dresher, A. W. Tucker, and P. Wolfe, editors, *Contributions to the Theory of Games, Vol. III*, Annals of Mathematics Studies 39, pages 97–139. Princeton University Press, 1957.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68:1127–1150, 2000.

- Sergiu Hart and Andreu Mas-Colell. A general class of adaptive procedures. *Journal of Economic Theory*, 98:26–54, 2001.
- Ehud Lehrer and Eilon Solan. No regret with bounded computational capacity. The Center for Mathematical Studies in Economics and Management Science, Northwestern University. Discussion Paper 1373, 2003.
- N. Littlestone and M. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.
- V. Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56:153–173, 1998.

Lecture Notes in Economics and Mathematical Systems

For information about Vols. 1–512
please contact your bookseller or Springer-Verlag

- Vol. 513: K. Marti, Stochastic Optimization Techniques. VIII, 364 pages. 2002.
- Vol. 514: S. Wang, Y. Xia, Portfolio and Asset Pricing. XII, 200 pages. 2002.
- Vol. 515: G. Heisig, Planning Stability in Material Requirements Planning System. XII, 264 pages. 2002.
- Vol. 516: B. Schmid, Pricing Credit Linked Financial Instruments. X, 246 pages. 2002.
- Vol. 517: H. I. Meinhardt, Cooperative Decision Making in Common Pool Situations. VIII, 205 pages. 2002.
- Vol. 518: S. Napel, Bilateral Bargaining. VIII, 188 pages. 2002.
- Vol. 519: A. Klose, G. Speranza, L. N. Van Wassenhove (Eds.), Quantitative Approaches to Distribution Logistics and Supply Chain Management. XIII, 421 pages. 2002.
- Vol. 520: B. Glaser, Efficiency versus Sustainability in Dynamic Decision Making. IX, 252 pages. 2002.
- Vol. 521: R. Cowan, N. Jonard (Eds.), Heterogenous Agents, Interactions and Economic Performance. XIV, 339 pages. 2003.
- Vol. 522: C. Neff, Corporate Finance, Innovation, and Strategic Competition. IX, 218 pages. 2003.
- Vol. 523: W.-B. Zhang, A Theory of Interregional Dynamics. XI, 231 pages. 2003.
- Vol. 524: M. Frölich, Programme Evaluation and Treatment Choise. VIII, 191 pages. 2003.
- Vol. 525: S. Spinler, Capacity Reservation for Capital-Intensive Technologies. XVI, 139 pages. 2003.
- Vol. 526: C. F. Daganzo, A Theory of Supply Chains. VIII, 123 pages. 2003.
- Vol. 527: C. E. Metz, Information Dissemination in Currency Crises. XI, 231 pages. 2003.
- Vol. 528: R. Stolletz, Performance Analysis and Optimization of Inbound Call Centers. X, 219 pages. 2003.
- Vol. 529: W. Krabs, S. W. Pickl, Analysis, Controllability and Optimization of Time-Discrete Systems and Dynamical Games. XII, 187 pages. 2003.
- Vol. 530: R. Wapler, Unemployment, Market Structure and Growth. XXVII, 207 pages. 2003.
- Vol. 531: M. Gallegati, A. Kirman, M. Marsili (Eds.), The Complex Dynamics of Economic Interaction. XV, 402 pages. 2004.
- Vol. 532: K. Marti, Y. Ermoliev, G. Pflug (Eds.), Dynamic Stochastic Optimization. VIII, 336 pages. 2004.
- Vol. 533: G. Dudek, Collaborative Planning in Supply Chains. X, 234 pages. 2004.
- Vol. 534: M. Runkel, Environmental and Resource Policy for Consumer Durables. X, 197 pages. 2004.
- Vol. 535: X. Gandibleux, M. Sevaux, K. Sörensen, V. T'kindt (Eds.), Metaheuristics for Multiobjective Optimisation. IX, 249 pages. 2004.
- Vol. 536: R. Brüggemann, Model Reduction Methods for Vector Autoregressive Processes. X, 218 pages. 2004.
- Vol. 537: A. Esser, Pricing in (In)Complete Markets. XI, 122 pages. 2004.
- Vol. 538: S. Kokot, The Econometrics of Sequential Trade Models. XI, 193 pages. 2004.
- Vol. 539: N. Hautsch, Modelling Irregularly Spaced Financial Data. XII, 291 pages. 2004.
- Vol. 540: H. Kraft, Optimal Portfolios with Stochastic Interest Rates and Defaultable Assets. X, 173 pages. 2004.
- Vol. 541: G.-y. Chen, X. Huang, X. Yang, Vector Optimization. X, 306 pages. 2005.
- Vol. 542: J. Lingens, Union Wage Bargaining and Economic Growth. XIII, 199 pages. 2004.
- Vol. 543: C. Benkert, Default Risk in Bond and Credit Derivatives Markets. IX, 135 pages. 2004.
- Vol. 544: B. Fleischmann, A. Klose, Distribution Logistics. X, 284 pages. 2004.
- Vol. 545: R. Hafner, Stochastic Implied Volatility. XI, 229 pages. 2004.
- Vol. 546: D. Quadt, Lot-Sizing and Scheduling for Flexible Flow Lines. XVIII, 227 pages. 2004.
- Vol. 547: M. Wildi, Signal Extraction. XI, 279 pages. 2005.
- Vol. 548: D. Kuhn, Generalized Bounds for Convex Multistage Stochastic Programs. XI, 190 pages. 2005.
- Vol. 549: G. N. Krieg, Kanban-Controlled Manufacturing Systems. IX, 236 pages. 2005.
- Vol. 550: T. Lux, S. Reitz, E. Samanidou, Nonlinear Dynamics and Heterogeneous Interacting Agents. XIII, 327 pages. 2005.
- Vol. 551: J. Leskow, M. Puchet Anyul, L. F. Punzo, New Tools of Economic Dynamics. XIX, 392 pages. 2005.
- Vol. 552: C. Suerie, Time Continuity in Discrete Time Models. XVIII, 229 pages. 2005.
- Vol. 553: B. Mönch, Strategic Trading in Illiquid Markets. XIII, 116 pages. 2005.
- Vol. 554: R. Foellmi, Consumption Structure and Macroeconomics. IX, 152 pages. 2005.
- Vol. 555: J. Wenzelburger, Learning in Economic Systems with Expectations Feedback (planned) 2005.
- Vol. 556: R. Branzei, D. Dimitrov, S. Tijs, Models in Cooperative Game Theory. VIII, 135 pages. 2005.

- Vol. 557: S. Barbaro, Equity and Efficiency Considerations of Public Higher Education. XII, 128 pages, 2005.
- Vol. 558: M. Faliva, M. G. Zoia, Topics in Dynamic Model Analysis. X, 144 pages, 2005.
- Vol. 559: M. Schulmerich, Real Options Valuation. XVI, 357 pages, 2005.
- Vol. 560: A. von Schemde, Index and Stability in Bimatrix Games. X, 151 pages, 2005.
- Vol. 561: H. Bobzin, Principles of Network Economics. XX, 390 pages, 2006.
- Vol. 562: T. Langenberg, Standardization and Expectations. IX, 132 pages, 2006.
- Vol. 563: A. Seeger (Ed.), Recent Advances in Optimization. XI, 455 pages, 2006.
- Vol. 564: P. Mathieu, B. Beaufls, O. Brandouy (Eds.), Artificial Economics. XIII, 237 pages, 2005.
- Vol. 565: W. Lemke, Term Structure Modeling and Estimation in a State Space Framework. IX, 224 pages, 2006.
- Vol. 566: M. Genser, A Structural Framework for the Pricing of Corporate Securities. XIX, 176 pages, 2006.
- Vol. 567: A. Namatame, T. Kaizouji, Y. Aruga (Eds.), The Complex Networks of Economic Interactions. XI, 343 pages, 2006.
- Vol. 568: M. Caliendo, Microeconometric Evaluation of Labour Market Policies. XVII, 258 pages, 2006.
- Vol. 569: L. Neubecker, Strategic Competition in Oligopolies with Fluctuating Demand. IX, 233 pages, 2006.
- Vol. 570: J. Woo, The Political Economy of Fiscal Policy. X, 169 pages, 2006.
- Vol. 571: T. Herwig, Market-Conform Valuation of Options. VIII, 104 pages, 2006.
- Vol. 572: M. F. Jäkel, Pensionomics. XII, 316 pages, 2006.
- Vol. 573: J. Emami Namini, International Trade and Multinational Activity. X, 159 pages, 2006.
- Vol. 574: R. Kleber, Dynamic Inventory Management in Reverse Logistics. XII, 181 pages, 2006.
- Vol. 575: R. Hellermann, Capacity Options for Revenue Management. XV, 199 pages, 2006.
- Vol. 576: J. Zajac, Economics Dynamics, Information and Equilibrium. X, 284 pages, 2006.
- Vol. 577: K. Rudolph, Bargaining Power Effects in Financial Contracting. XVIII, 330 pages, 2006.
- Vol. 578: J. Kühn, Optimal Risk-Return Trade-Offs of Commercial Banks. IX, 149 pages, 2006.
- Vol. 579: D. Sondermann, Introduction to Stochastic Calculus for Finance. X, 136 pages, 2006.
- Vol. 580: S. Seifert, Posted Price Offers in Internet Auction Markets. IX, 186 pages, 2006.
- Vol. 581: K. Marti; Y. Ermoliev; M. Makowski; G. Pflug (Eds.), Coping with Uncertainty. XIII, 330 pages, 2006.
- Vol. 582: J. Andritzky, Sovereign Default Risks Valuation: Implications of Debt Crises and Bond Restructurings. VIII, 251 pages, 2006.
- Vol. 583: I.V. Konnov, D.T. Luc, A.M. Rubinov[†] (Eds.), Generalized Convexity and Related Topics. IX, 469 pages, 2006.
- Vol. 584: C. Bruun, Advances in Artificial Economics: The Economy as a Complex Dynamic System. XVI, 296 pages, 2006.
- Vol. 585: R. Pope, J. Leitner, U. Leopold-Wildburger, The Knowledge Ahead Approach to Risk. XVI, 218 pages, 2007 (planned).
- Vol. 586: B. Lebreton, Strategic Closed-Loop Supply Chain Management. X, 150 pages, 2007 (planned).
- Vol. 587: P. N. Baecker, Real Options and Intellectual Property: Capital Budgeting Under Imperfect Patent Protection. X, 276 pages, 2007.
- Vol. 588: D. Grundel, R. Murphey, P. Panos, O. Prokopyev (Eds.), Cooperative Systems: Control and Optimization. IX, 401 pages, 2007.
- Vol. 589: M. Schwind, Dynamic Pricing and Automated Resource Allocation for Information Services: Reinforcement Learning and Combinatorial Auctions. XII, 293 pages, 2007.
- Vol. 590: S. H. Oda, Developments on Experimental Economics: New Approaches to Solving Real-World Problems. XVI, 262 pages, 2007.
- Vol. 591: M. Lehmann-Waffenschmidt, Economic Evolution and Equilibrium: Bridging the Gap. VIII, 272 pages, 2007.
- Vol. 592: A. C.-L. Chian, Complex Systems Approach to Economic Dynamics. X, 95 pages, 2007.
- Vol. 593: J. Rubart, The Employment Effects of Technological Change: Heterogenous Labor, Wage Inequality and Unemployment. XII, 209 pages, 2007.
- Vol. 594: R. Hübner, Strategic Supply Chain Management in Process Industries: An Application to Specialty Chemicals Production Network Design. XII, 243 pages, 2007.
- Vol. 595: H. Gimpel, Preferences in Negotiations: The Attachment Effect. XIV, 268 pages, 2007.
- Vol. 596: M. Müller-Bungart, Revenue Management with Flexible Products: Models and Methods for the Broadcasting Industry. XXI, 297 pages, 2007.
- Vol. 597: C. Barz, Risk-Averse Capacity Control in Revenue Management. XIV, 163 pages, 2007.
- Vol. 598: A. Ule, Partner Choice and Cooperation in Networks: Theory and Experimental Evidence. Approx. 200 pages, 2007.
- Vol. 599: A. Consiglio, Artificial Markets Modeling: Methods and Applications. XV, 277 pages, 2007.
- Vol. 600: M. Hickman, P. Mirchandani, S. Voss (Eds.): Computer-Aided Scheduling of Public Transport. Approx. 424 pages, 2007.
- Vol. 601: D. Radulescu, CGE Models and Capital Income Tax Reforms: The Case of a Dual Income Tax for Germany. XVI, 168 pages, 2007.
- Vol. 602: N. Ehrentreich, Agent-Based Modeling: The Santa Fe Institute Artificial Stock Market Model Revisited. XVI, 225 pages, 2007.