

Agent-Based Computational Modelling



Contributions to Economics

www.springeronline.com/series/1262

Further volumes of this series can be found at our homepage.

Pablo Coto-Millán

General Equilibrium and Welfare

2002. ISBN 7908-1491-1

Wojciech W. Charemza/Krystyna Strzala (Eds.)

East European Transition and EU Enlargement

2002. ISBN 3-7908-1501-1

Natalja von Westernhagen

Systemic Transformation, Trade and Economic Growth

2002. ISBN 3-7908-1521-7

Josef Falkinger

A Theory of Employment in Firms

2002. ISBN 3-7908-1520-9

Engelbert Plassmann

Econometric Modelling of European Money Demand

2003. ISBN 3-7908-1522-5

Reginald Loyen/Erik Buyst/Greta Devos (Eds.)

Struggling for Leadership: Antwerp-Rotterdam Port Competition between 1870-2000

2003. ISBN 3-7908-1524-1

Pablo Coto-Millán

Utility and Production, 2nd Edition

2003. ISBN 3-7908-1423-7

Emilio Colombo/John Driffill (Eds.)

The Role of Financial Markets in the Transition Process

2003. ISBN 3-7908-0004-X

Guido S. Merzoni

Strategic Delegation in Firms and in the Trade Union

2003. ISBN 3-7908-1432-6

Jan B. Kuné

On Global Aging

2003. ISBN 3-7908-0030-9

Sugata Marjit, Rajat Acharyya

International Trade, Wage Inequality and the Developing Economy

2003. ISBN 3-7908-0031-7

Francesco C. Billari/Alexia Prskawetz (Eds.)

Agent-Based Computational Demography

2003. ISBN 3-7908-1550-0

Georg Bol/Gholamreza Nakhaeizadeh/

Svetlozar T. Rachev/Thomas Ridder/

Karl-Heinz Vollmer (Eds.)

Credit Risk

2003. ISBN 3-7908-0054-6

Christian Müller

Money Demand in Europe

2003. ISBN 3-7908-0064-3

Cristina Nardi Spiller

The Dynamics of the Price Structure and the Business Cycle

2003. ISBN 3-7908-0063-5

Michael Bräuninger

Public Debt and Endogenous Growth

2003. ISBN 3-7908-0056-1

Brigitte Preissl/Laura Solimene

The Dynamics of Clusters and Innovation

2003. ISBN 3-7908-0077-5

Markus Gangl

Unemployment Dynamics in the United States and West Germany

2003. ISBN 3-7908-1533-0

Pablo Coto-Millán (Ed.)

Essays on Microeconomics and Industrial Organisation, 2nd Edition

2004. ISBN 3-7908-0104-6

Wendelin Schnedler

The Value of Signals in Hidden Action Models

2004. ISBN 3-7908-0173-9

Carsten Schröder

Variable Income Equivalence Scales

2004. ISBN 3-7908-0183-6

Wilhelm J. Meester

Locational Preferences of Entrepreneurs

2004. ISBN 3-7908-0178-X

Russel Cooper/Gary Madden (Eds.)

Frontiers of Broadband, Electronic and Mobile Commerce

2004. ISBN 3-7908-0087-2

Sardar M. N. Islam

Empirical Finance

2004. ISBN 3-7908-1551-9

Jan-Egbert Sturm/Timo Wollmershäuser (Eds.)

Ifo Survey Data in Business Cycle and Monetary Policy Analysis

2005. ISBN 3-7908-0174-7

Bernard Michael Gilroy/Thomas Gries/

Willem A. Naudé (Eds.)

Multinational Enterprises, Foreign Direct Investment and Growth in Africa

2005. ISBN 3-7908-0276-X

Günter S. Heiduk/Kar-yiu Wong (Eds.)

WTO and World Trade

2005. ISBN 3-7908-1579-9

Emilio Colombo/Luca Stanca

Financial Market Imperfections and Corporate Decisions

2006. ISBN 3-7908-1581-0

Birgit Mattil

Pension Systems

2006. ISBN 3-7908-1675-2

Francesco C. Billari · Thomas Fent
Alexia Prskawetz · Jürgen Scheffran
(Editors)

Agent-Based Computational Modelling

Applications in Demography,
Social, Economic
and Environmental Sciences

With 95 Figures and 19 Tables

Physica-Verlag

A Springer Company

Series Editors

Werner A. Müller
Martina Bihn

Editors

Professor Dr. Francesco C. Billari
Università Bocconi & IGIER
Istituto di Metodi Quantitativi
Viale Isonzo 25
20135 Milano
Italy
francesco.billari@uni-bocconi.it

Dr. Jürgen Scheffran
University of Illinois, ACDIS
505 East Armory Ave.
Champaign, IL 61820
USA
scheffra@uiuc.edu

Dr. Thomas Fent
Univ. Doz. Dr. Alexia Prskawetz
Vienna Institute of Demography
Prinz Eugen-Straße 8–10
1040 Vienna
Austria
thomas.fent@oeaw.ac.at
alexia.fuernkranz-prskawetz@oeaw.ac.at

ISSN 1431-1933

ISBN-10 3-7908-1640-X Physica-Verlag Heidelberg New York

ISBN-13 978-3-7908-1640-2 Physica-Verlag 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 Physica-Verlag. Violations are liable for prosecution under the German Copyright Law.

Physica-Verlag is a part of Springer Science+Business Media

springer.com

© Physica-Verlag Heidelberg 2006
Printed in Germany

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.

Cover-Design: Erich Kirchner, Heidelberg

SPIN 11531616 88/3153-5 4 3 2 1 0 – Printed on acid-free and non-aging paper

To our children

Preface

This book is the outcome of a project that started with the organisation of the Topical Workshop on “Agent-Based Computational Modelling. An Instrument for Analysing Complex Adaptive Systems in Demography, Economics and Environment” at the Vienna Institute of Demography, December 4-6, 2003. The workshop brought together scholars from several disciplines, allowing both for serious scientific debate and for informal conversation over a cup coffee or during a visit to the wonderful museums of Vienna. One of the nicest features of Agent-Based Modelling is indeed the opportunity that scholars find a common language and discuss from their disciplinary perspective, in turn learning from other perspectives. Given the success of the meeting, we found it important to pursue the purpose of collecting these interdisciplinary contributions in a volume. In order to ensure the highest scientific standards for the book, we decided that all the contributions (with the sole exception of the introductory chapter) should have been accepted conditional on peer reviews. Generous help was provided by reviewers, some of whom were neither directly involved in the workshop nor in the book. All this would not have been possible without the funding provided by the Complex Systems Network of Excellence (*Exystence*) funded by the European Union, the Vienna Institute of Demography of the Austrian Academy of Sciences, Università Bocconi, and ARC Systems Research GmbH, and the help of the wonderful staff of the Vienna Institute of Demography (in particular, Ani Minassian and Belinda Aparicio Diaz). Agent-Based Modelling is important, interesting and also fun—we hope this book contributes to showing that.

Milano
Zurich
Vienna
Champaign

Francesco C. Billari
Thomas Fent
Alexia Prskawetz
Jürgen Scheffran

Contents

Agent-Based Computational Modelling: An Introduction
Francesco C. Billari, Thomas Fent, Alexia Prskawetz, Jürgen Scheffran . 1

Socio-Economics

Agent-Based Modelling – A Methodology for the Analysis of Qualitative Development Processes
Andreas Pyka, Thomas Grebel 17

On the Analysis of Asymmetric Directed Communication Structures in Electronic Election Markets
Markus Franke, Andreas Geyer-Schulz, Bettina Hoser 37

Population and Demography

The Role of Assortative Mating on Population Growth in Contemporary Developed Societies
Mike Murphy 61

An Agent-Based Simulation Model of Age-at-Marriage Norms
Belinda Aparicio Diaz, Thomas Fent 85

The Strength of Social Interactions and Obesity among Women
Mary A. Burke, Frank Heiland 117

Ecology and Environment

Agent-Based Models in Ecology: Patterns and Alternative Theories of Adaptive Behaviour
Volker Grimm, Steven F. Railsback 139

Agent-Based Modelling of Self-Organisation Processes to Support Adaptive Forest Management
Ernst Gebetsroither, Alexander Kaufmann, Ute Gigler, Andreas Resetarits 153

Vampire Bats & The Micro-Macro Link
Rosaria Conte, Mario Paolucci, Gennaro Di Tosto 173

General Aspects

How Are Physical and Social Spaces Related? – Cognitive Agents as the Necessary “Glue”
Bruce Edmonds 195

Agent Design for Agent-Based Modelling
Jim Doran 215

List of Contributors 225

Agent-Based Computational Modelling: An Introduction

Francesco C. Billari¹, Thomas Fent², Alexia Prskawetz³ and Jürgen Scheffran⁴

¹ Istituto di Metodi Quantitativi, Università Bocconi, Milano, Italy
`francesco.billari@unibocconi.it`

² Department Management, Technology, and Economics, ETH Zurich, Switzerland
`tfent@ethz.ch`

³ Vienna Institute of Demography, Austrian Academy of Sciences, Austria
`alexia.fuernkranz-prskawetz@oeaw.ac.at`

⁴ ACDIS, University of Illinois at Urbana-Champaign, USA
`scheffra@uiuc.edu`

Summary. Agent-based models (ABMs) are increasingly used in studying complex adaptive systems. Micro-level interactions between heterogeneous agents are at the heart of recent advances in modelling of problems in the social sciences, including economics, political science, sociology, geography and demography, and related disciplines such as ecology and environmental sciences. Scientific journals and societies related to ABMs have flourished. Some of the trends will be discussed, both in terms of the underlying principles and the fields of application, some of which are introduced in the contributions to this book.

1 Agent-Based Modelling: An Emerging Field in Complex Adaptive Systems

Since Thomas C. Schelling's pathbreaking early study on the emergence of racial segregation in cities [32], a whole new field of research on socio-economic systems has emerged, dubbed with a diversity of names, such as social simulation, artificial societies, individual-based modelling in ecology, agent-based computational economics (ACE), agent-based computational demography (ABCD). Accordingly, the literature on agent-based modelling in social sciences has flourished recently, particularly in economics⁵, political science⁶,

⁵ See i.e. the special issue on agent-based computational economics of the Journal of Economic Dynamics and Control [34], especially the introduction by Leigh Tesfatsion, as well as the website maintained by Tesfatsion <http://www.econ.iastate.edu/tesfatsi/ace.htm>.

⁶ See i.e. the review paper by Johnson [23].

and – to a lesser extent – sociology⁷. During the 1990s, this computational approach to the study of human behaviour developed through a vast quantity of literature. These include approaches that range from the so-called evolutionary computation (genetic algorithms and evolution of groups of rules) to the study of the social evolution of adaptive behaviours, of learning, of innovation, or of the possible social interactions connected to the theory of games.

Different to the approach of experimental economics and other fields of behavioural science that aim to understand why specific rules are applied by humans, agent-based computational models pre-suppose rules of behaviour and verify whether these micro-based rules can explain macroscopic regularities. The development in computational agent-based models has been made possible by the progress in information technology (in hardware as well as software agent technology), and by the presence of some problems that are unlikely to be resolved by simply linking behavioural theories and empirical observations through adequate statistical techniques. The crucial idea that is at the heart of these approaches is to use computing as an aid to the development of theories of human behaviour. The main emphasis is placed on the explanation rather than on the prediction of behaviour, and the model is based on individual agents.

As outlined in Axelrod ([1, p.4]), agent-based computational modelling may be compared to the principles of induction and deduction. “Whereas the purpose of induction is to find patterns in data and that of deduction is to find consequences of assumptions, the purpose of agent-based modelling is to aid intuition”. As with deduction, agent-based modelling starts with assumptions. However, unlike deduction, it does not prove theorems. The simulated data of agent-based models can be analysed inductively, even though the data are not from the real world as in case of induction.

2 From Rational Actors to Agent-Based Models

Established economic theory is based on the rational actor paradigm which assumes that individual actors know their preferences, often measured by a utility function, and the best possible decision, based on complete information about their environment and the supposed consequences. Decision theory deals with the ranking and selection of the options of actors, according to their preferences. Usually a single rational decision-maker maximizes utility (value) under given constraints, where a wide range of methods have been developed to search for and find the optimum. While rational actors may be adequate in environments with a few number of state and control variables, they have limits in complex and uncertain environments and with real human beings of bounded rationality and restrained computational capabilities.

One of the conditions that restrains rationality is the social environment itself, in particular the unpredictable behaviour of other agents. Game theory

⁷ See i.e. the review paper by Macy and Willer [26], or the review of Halpin [18].

is trying to extend rational decision-making to two and more players, each pursuing their own preferences and utilities in response to the expected or observed decisions of other players. Game theory becomes more difficult to handle when a large number of players interact in a dynamic environment. Dynamic game models describe the interaction between multiple players according to situation-dependent decision rules and reaction functions. In repeated games players can learn and adapt their behaviour to the strategies of other players, possibly leading to the evolution of cooperation. Evolutionary games analyse the selection among competing populations of game strategies according to their fitness in replication.

Recent years saw a transition from rational actor models to agent-based modelling, and from top-down macro decision-making to bottom-up micro-simulation. A common feature of ABMs is that individual agents act according to rules, where utility optimization is just one of many possible rules. Thanks to increasing computational capabilities, it became possible to analyse interactions between multiple agents, forming complex social patterns. Computers turned into laboratories of artificial societies ([12], [13]). Simulations have now the character of experiments in virtual worlds, often with demanding computational requirements.

In cellular automata models, agents behave like insects in virtual landscapes [41]. For a large number of homogenous agents, methods from statistical physics, non-linear dynamics and complexity science are applicable [17], describing self-organization or phase transitions when observed macroscopic properties emerge from the behaviour and interactions of the component agents. Approaches to collective phenomena have been transferred to interdisciplinary fields such as socio-physics and econo-physics, with applications ranging from moving crowds and traffic systems to urban, demographic and environmental planning ([22],[39],[33]).

Key challenges are to find a conceptual framework to structure the diverse field of ABMs, to calibrate the models with data and to integrate ABMs into real-world applications. The selection of strategies and decision rules in computer-based simulation models can be based on observation and include real-world actors and stakeholders, offering a wide field of experimental games for educational and research purposes as well as for decision support and policy advice. Special modelling-simulation environments or toolkits of various kinds are available for performing experiments, which abstract from the details and can be duplicated by other researchers.

3 Structure, Behaviour and Interaction of Agents

Agent-based models are usually based on a set of autonomous agents capable to interact with each other as well as with the environment according to rules of behaviour, which can be simple or complex, deterministic or stochastic, fixed or adaptive. An agent can be any organisational entity that is able to

act according to its own set of rules and objectives. All agents can be of the same type (homogenous) or each agent can be different from the other (heterogeneous).

One core question is related to the structure of agents: should agents be simple or should they be complex? Proponents of the simplicity of agents, such as Robert Axelrod [1], support the so-called KISS principle (keep it simple, stupid), and point out that the most interesting analytical results are obtained when simple micro-level dynamics produce complex patterns at the macro level. This approach is analogous to mathematical models where complex dynamics may arise from simple rules. Proponents of the complexity of agents base their views especially in economics, sociology and cognitive psychology, assuming that agents are possibly guided by a set of behavioural rules and objective functions which evolved as a result of interaction and learning in complex environments and shape the individual structure of each agent. Reality tends to be between simplicity and complexity, and agents should be kept as “simple as suitable”. Real agents seek to reduce complexity according to their needs and adjust to their social environment, sometimes leading to rather simple collective behaviour, despite the potential for individual complexity.

Agents can include many details matching reality, at different spatial and temporal scales. Depending on the agents’ number, their attributes and behavioural rules in their respective environments, ABM’s can be of great variety and complexity, making them hard to analyse or predict. Using sensors, agents can perceive their local neighbourhood and receive or send messages ([14]).

Cognitive agents may have cognitive capabilities “to perceive signals, react, act, making decisions, etc. according to a set of rules” ([9]). Their intended actions are shaped by what they think to know about the world (beliefs), based on experience and perception, and what they would like to achieve (desired goals), both represented by an internal model of the external environment. Agents can be autonomous and act independently of any controlling agency, or they can directly interact with or depend on other agents. In their environment agents need information to react and adapt to their observation and to respond to changes in the environment, and they can communicate with other agents via a language. Pursuing goals, agents need to be pro-active, and they can be rational by following a well-defined and logical set of decision rules to achieve these goals.

Adaptive agents have the capability to learn, i.e. rather than following a fixed stimulus-response pattern, they continuously adapt to changes in their environment according to their expectations and objectives. They evolve in a learning cycle of acting, evaluating the results of the actions dependent on the response of the environment and updating the objective or the actions. By acting an agent employs resources and directs them onto its environment, in order to achieve the objective. Evaluation compares the results of the actions and their impacts with the expectations and objectives. Searching tries to find

better routines for achieving the objective. Adaptive agents can change their objectives and routines.

A general framework for agent-based modelling can be characterized by the following elements (see the contribution by Gebetsroither et al. in this volume):

- Values, targets and objectives
- Resources or production factors
- Observation, expectation and update
- Rules, search routines and actions

These elements occur repeatedly in a cycle of action, evaluation and update. A more comprehensive analysis would consider the complete multi-step process of decision-making, interaction and management, including the following phases [31]:

1. Situational analysis and problem structuring
2. Option identification and scenario modelling
3. Concept development and criteria-based evaluation
4. Decision-making and negotiation
5. Planning and action
6. Monitoring and learning

The different phases are connected by processes such as evaluation, communication, capacity building, information, simulation, validation. Usually ABMs do not apply all phases of this cycle but only selected elements which are of particular relevance for a given problem.

4 From Micro to Macro: Modelling Population Processes from the Bottom-Up

Agent-based simulations are increasingly applied in the social sciences. Artificial computational environments serve in fact as small laboratories to simulate social behaviours and interaction among a large number of actors. This includes the study of the complex dynamics evolving from heterogenous populations. Populations are by definition aggregates of individuals, and as such they constitute entities at the aggregate or “macro” level, whereas individual lives contribute to numbers of events, person years and survivors, which are used in the statistical analysis of populations. Demography as such is concerned with the study of populations, and has been traditionally focusing on the macro side of population dynamics, on “macro-demography”. However, during the last decades of the Twentieth Century a “micro-demography” emerged with a specific emphasis on the unfolding of individual-level demographic trajectories and on the consequences of individual heterogeneity for the study of population dynamics.

Perhaps surprisingly, other disciplines than the one focusing on population per se have attempted at micro-founding the study of specific types of behaviour using some type of “methodological individualism” approach. In particular, we refer to ecology, sociology, and economics, disciplines that are in particular represented in this book.

In *ecology*, “individual-based modelling” (IBM), e.g. for the study of animal and plant populations, has emerged starting from the mid-1970s as a research program that has led to significant contributions (for a review see [15]). According to Grimm and Railsback [16], individual-based models in ecology fulfill, to a certain degree, four criteria: first, they explicitly consider individual-level development; second, they represent explicitly the dynamics of the resources an individual has access to; third, individuals are treated as discrete entities and models are built using the mathematics of discrete events rather than rates; fourth, they consider variation between individuals of the same age. Individual-based models in ecology are aimed at producing “patterns” that can be compared to patterns observed in reality. The sustainable use and management of natural resources is an important issue but difficult to model because it is characterized by complexity, a high degree of uncertainty, information deficits and asymmetries.

There are not many examples of agent-based models concerning the management of natural resources. A complete agent-based model would have to comprise both social and natural systems and respective agents, which is a challenging task.

In *sociology*, the approach proposed by James Coleman (see [8] Ch. 1) proposes to found social theory ultimately on the micro-level decisions of individuals. Coleman proposes to use a three-part schema for explaining macro-level phenomena, consisting of three types of relations: 1) the “macro-to-micro transition – that is, how the macro-level situation affects individuals; 2) “purposeful action of individuals” – that is, how individual choices are affected by micro-level factors; 3) the “micro-to-macro transition” – that is, how macro-level phenomena emerge from micro-level action and interaction.

Coleman’s conceptual framework is embedded in the notion of “social mechanism” as the key concept to explain behaviour in the social sciences, proposed by Hedström and Swedberg [21], who see the three types of relationships as 1) situational mechanisms, representing the case in which “The individual actor is exposed to a specific social situation, and this situation will affect him or her in a particular way”; 2) action formation mechanisms, representing “a specific combination of individual desires, beliefs, and action opportunities (that) generate a specific action”; 3) transformational mechanisms, specifying “how these individual actions are transformed into some kind of collective outcome, be it intended or unintended”. The framework is very similar to the one presented recently by Daniel Courgeau [11] in a review on the macro-micro link.

As we noticed before, the micro level is the natural point of departure in *economics*, also when pointing to the macro level as the important out-

come. While the first generation of economic simulation models was rather focused on stylized empirical phenomena, the emergence of agent-based modelling during the last 10 years has shifted the emphasis from macro simplicity to micro complexity of the socio-economic reality. As noted by van den Bergh and Gowdy [36, p. 65] “During the last quarter century, the microfoundations approach to macroeconomic theory has become dominant”. Mainstream economics, also known as “neoclassical” economics traditionally considers a “representative agent” who maximizes a potentially complex utility function subject to potentially complex budget constraints. This and other hypotheses lead to mathematically tractable models of macro-level outcomes. The new economics approach that applies the toolkit of neoclassical economics to demographic choices has been a key success of the work of Gary Becker (see e.g. [6]). This approach has now reached a level of maturity that can be attested from the literature on population economics (see e.g. [42]). That we ought to start from the micro level is also clearly stated by an economist who is particularly interested in population matters, Jere Behrman, who states that “For both good conditional predictions and good policy formation regarding most dimensions of population change and economic development, a perspective firmly grounded in understanding the micro determinants - at the level of individuals, households, farms, firms, and public sector providers of goods and services of population changes and of the interactions between population and development is essential” [7].

The attention on the policy relevance of research on population (including policy implications of results) is undoubtedly the main characteristic that comes to the surface when looking at research on population economics. Micro-based theories of behaviour are thus used to cast “conditional prediction” of reactions to a given policy, with these reactions affecting macro-level outcomes. Within economics, several scholars have objected to the neoclassical paradigm from various perspectives (see e.g. [7] for objections to critiques concerning population-development relationships). Of particular interest are the critiques on mainstream economics that concern the assumption that agents are homogeneous and the lack of explicit interaction between agents (see e.g. Kirman [24]). Kirman’s point is that even if individuals are all utility maximizers (an idea that has also been challenged by several scholars), the assumption that the behaviour of a group of heterogeneous and interacting agents can be mimicked by that of a single representative individual whose choices coincide with the aggregate choices of the group is unjustified and leads to misleading and often wrong conclusions.

To overcome this micro-macro “aggregation” problem, that is the transformational mechanism in Coleman’s scheme, some economists have proposed to build models that resemble that of IBM in ecology. Models in agent-based computational economics (ACE) explicitly allow the interaction between heterogeneous agents (see e.g. the review by Tesfatsion [34]).

5 Population Dynamics from the Bottom-Up: ABCD

We now document the emergence of the agent-based modelling approach in demography as a specific case-study.

Without the strong paradigm of the “representative agent” that underlies mainstream economics, demography has to solve aggregation problems taking into account that demographic choices are made by heterogeneous and interacting individuals, and that sometimes demographic choices are made by more than one individual (a couple, a household). For these reasons, and for the natural links to current micro-demography, computer simulation provides a way to transform micro into macro without having to impose unnecessary assumptions on the micro level (among those homogeneity, lack of interaction).

Agent-based computational demography (ABCD) has been shaped by a set of tools that models population processes, including their macro level dynamics, from the bottom up, that is by starting from assumptions at the micro level [4]. Agent-based computational demography includes also micro-simulation that has been used to derive macro-level outcomes from empirical models of micro-level demographic processes (i.e. event history models), but also formal models of demographic behaviour that describe micro-level decisions with macro-level outcomes.

It is interesting to notice that demography has for a long time been using simulation techniques, and microsimulation has become one of the principal techniques in this discipline, being a widely discussed and applied instrument in the study of family and kinship networks and family life cycle ([19]; [38]; [30]; [20]; [35]). Microsimulation has also been widely used in the study of human reproduction and fecundability ([29]; [27]), migratory movements [10] or whole populations [25], and its role has been discussed in the general context of longitudinal data analysis [40]. Evert van Imhoff and Wendy Post [37] provide a general overview of the topic. Microsimulation has been used to study and predict the evolution of a population using a model for individuals.

What does ABCD add to demographic microsimulation in helping to bridge the gap between micro-demography and macro-demography? The emphasis of demographic microsimulation has been on the macro-level impact of a certain set of parameters estimated at the micro-level from actual empirical data. There has been no particular emphasis on the theoretical side. Agent-based models do not necessarily include only parameters estimated from actual empirical data, but it may include parameters that are relevant for a specific theoretical meaning. In fact, microsimulation is to the event history analysis what macrosimulation (i.e. population projection based on aggregate-level quantities like in the cohort-component model) is to traditional, macro-level, formal demography. On the other hand, agent-based computational demography is the micro-based functional equivalent of mathematical demography.

Some of the reasons why ABCD helps bridging the macro-micro gap in demography are mentioned in this context (see [5] for a full discussion).

First, it is relatively easy to include feedback mechanisms and to integrate micro-based demographic behavioural theories (and results from individual-level statistical models of demographic behaviour such as event history models) with aggregate-level demographic outcomes. This ability to include feedback is possibly the most important gain of ABCD models. In such models, space and networks can be formalised as additional entities through which the agents will interact.

Second, compared to mathematical modelling, it is relatively easy to introduce heterogeneous agents that are not fully rational. Hence, the paradigm of the representative, fully rational agent that has and often still penetrates many economic and sociological applications can easily be relaxed in agent-based modelling.

Third, when building agent-based computational models, it is indispensable to adopt simple formulations of theoretical statements. Although agent-based modelling employs simulation, it does not aim to provide an accurate representation of a particular empirical application. Instead, the goal of agent-based modelling should be to enrich our understanding of fundamental processes that may appear in a variety of applications. This requires adhering to the KISS principle.

Fourth, using agent-based approaches, it is possible to construct models for which explicit analytical solutions do not exist, for example social interaction and generally non-linear models. Agent-based models are used to understand the functioning of the model and the precision of theories need not be limited to mathematical tractability. Simplifying assumptions can then be relaxed in the framework of an agent-based computational model. But as Axtell [2] notes, even when models could be solved analytically or numerically, agent-based modelling techniques may be applied since their output is mostly visual and therefore easier to communicate to people outside academia. In general, we can see formal modelling of population dynamics using differential equations and agent-based computational models as two ends of a continuum along the macro-micro dimension [28].

Finally, it is possible to conceive artificial societies that need not necessarily resemble present societies; such artificial societies can be seen as computational laboratories or may allow to reproduce past macro-events from the bottom-up.

6 Contributions of ABMs to Economic, Demographic and Ecological Analysis

The present book describes the methodology to set up agent-based models and to study emerging patterns in complex adaptive systems resulting from multi-agent interaction. It presents and combines different approaches, with applications in demography, socio-economic and environmental sciences.

6.1 Socio-Economics

Andreas Pyka and *Thomas Grebel* provide a basic instruction on how to model qualitative change using an agent-based modelling procedure. The reasons to focus on qualitative change are discussed, agent-based modelling is explained and finally an evolutionary economics model of entrepreneurial behaviour is given as an example. The conceptual framework for the analysis of entrepreneurial behaviour is composed of several building blocks (actors, actions, endowments, interaction, evaluation and decision processes), which are not separate and unrelated entities but represent the conceptual view on the issue, as a result of a systematization process. Actors are not modelled by a representative agent but by a population of heterogeneous agents. For any of two subpopulations (agents and firms) rules and routines are derived which govern the particular actions of the agents, the interaction and interrelation of the agents within and among the sub-populations. The nature of the actors and their heterogeneity is shaped by the endowment with resources and their individual routines, which are related to the satisficing behaviour and bounded rationality of the actors. Routinized behaviour causes some inertia and stability of the system. Some actors join networks with other actors and found a firm, others disentangle their networks or even go bankrupt. The basic conceptual building blocks are implemented in the actual model of entrepreneurial behaviour.

In their contribution, *Markus Franke*, *Andreas Geyer-Schulz* and *Bettina Hoser* analyse asymmetric directed communication structures in electronic election markets. They introduce a new general method of transforming asymmetric directed communication structures represented as complex adjacency matrices into Hermitian adjacency matrices which are linear self-adjoint operators in a Hilbert space. With this method no information is lost, no arbitrary decision on metrics is involved, and all eigenvalues are real and easily interpretable. The analysis of the resulting eigensystem helps in the detection of substructures and general patterns. The formal method is applied in the context of analysing market structure and behaviour based on market transaction data from the eigensystem. As an example, the results of a political stock exchange for the 2002 federal elections in Germany are analysed. Market efficiency is of special interest for detecting locally inefficient submarkets in energy markets.

6.2 Population and Demography

Mike Murphy discusses the role of assortative mating on population growth in contemporary developed societies. Assortative mating is a widespread feature of human behaviour, with a number of suggested benefits. The question of whether it contributes to population growth in contemporary societies is considered using the micro simulation program SOCSIM. Ways of parameterising heterogeneous fertility and nuptiality, and the relationship of such parameters

to those of both fathers and mothers are considered. One conclusion is that the effect of assortative mating in which the fertility backgrounds of spouses are positively correlated leads to higher population growth. A population with a higher long term rate of growth, no matter how small the advantage, will come to dominate numerically any population with a lower one and the overall population eventually becomes effectively homogeneous and consists only of the higher growth population. Further progress will require developments in theory, data, modelling and technology, but assortative mating remains one of the most persistent and enduring features of humans and other species.

Belinda Aparicio Diaz and *Thomas Fent* analyse an agent-based model designed to understand the dynamics of the intergenerational transmission of age-at-marriage norms. A norm in this context is an acceptable age interval to get married. It is assumed that this age-interval is defined at the individual level and the individuals' age-at-marriage norms are transmitted from parents to their children. The authors compare four different transmission mechanisms to investigate the long term persistence or disappearance of norms under different regimes of transmission. They investigate whether results also hold in a complex setup that takes into account heterogeneity with respect to age and sex as well as the timing of union formation and fertility. To create a more realistic model of evolving age norms, the characteristics of the agents are extended, and the age-at-marriage norms are split into two sex-specific age-at-marriage norms. The results provide information about how additional characteristics and new parameters can influence the evolution of age-at-marriage norms.

To explain the differences in obesity rates among women in the United States by education, *Mary A. Burke* and *Frank Heiland* model a social process in which body weight norms are determined endogenously in relation to the empirical weight distribution of the peer group. The dramatic growth in obesity rates in the United States since the early 1980's to close to 30% in 2000 has been widely publicised and raised attention to the problem of obesity. Obesity significantly elevates the risks of diabetes, heart disease, hypertension, and a number of cancers, and remains a prominent public health priority. The agent-based model embeds a biologically accurate representation of variation of metabolism which enables to describe a distribution of weights. Individuals are compared to others with the same level of educational attainment. The agents are biologically complex, boundedly rational individuals that interact within a social group. Using heterogeneous metabolism and differences in average energy expenditure, an entire population distribution of body weights is generated. Weight norms are defined as a function of aggregate behaviour, and deviation from the norm is costly. Consistent with the observed distribution of body weights among women in the U.S. population, the model predicts lower average weights and less dispersion of weight among more educated women. While previous models have made qualitative predictions of differential obesity rates across social groups, they have not captured the differences in the overall weight distributions that this model is able to reproduce. The model is

also used to investigate competing hypotheses based on behavioural or genetic differences across education groups.

6.3 Ecology and Environment

Volker Grimm and *Steven F. Railsback* specify agent-based models in ecology by discussing two modelling strategies that have proven particularly useful: pattern-oriented modelling (POM), and a theory for the adaptive behaviour of individuals. These two strategies are demonstrated with example models of schooling behaviour in fish, spatiotemporal dynamics in forests, and dispersal of brown bears. Schooling-like behaviour is based on simple assumptions on individual behaviour: individuals try to match the velocity of neighbouring individuals, and to stay close to neighbours which leads to the emergence of school-like aggregations. This demonstrates how simple behavioural rules and local interactions give rise to a collection of individuals which are more or less regularly spaced and move as one coherent entity. The question is discussed how to learn about how real fish behave by combining observed patterns, data, and an IBM. Specific properties of real fish schools are quantified, such as nearest neighbour distance and polarisation, i.e. the average angle of deviation between the mean direction of the entire school and the swimming direction of each fish.

Ernst Gebetsroither, Alexander Kaufmann, Ute Gigler and *Andreas Reserits* present a preliminary version of an agent-based model of self-organisation processes to support adaptive forest management. The modular approach consists of two separate, but interlinked submodels. While the forest submodel includes a very large number of comparatively simple agents, the socio-economic submodel comprises only a few complex agents defined by a fixed set of an objective and several routines, technologies and resources. The use of forest resources is determined by the interrelations between specific forest management methods and the specific demand for timber of industries producing wood-based goods. The timber market includes two types of agents which belong to the sectors “forestry” offering timber with a long-term planning horizon and “industry” producing wood-based goods with a short-time perspective. Their relation is characterised by imperfect competition, imperfect information, strategic behaviour and learning. Other potentially important agents are either not included in this model (e.g. tourists, hunters) or considered as exogenous forces (e.g. state authorities, communities, demand for wood-based products, competing sources of timber supply). The main question is how self-organisation processes on the timber market (demand for the forest resource “timber”) as well as in forest succession (available stock of timber) influence each other and which effects of adaptive management methods can be expected on the overall system’s behaviour. Running simulations with an empirically calibrated model (using forestry data and interviews of experts) allows to test specific forest management routines under controlled conditions and restrictions.

Rosaria Conte, Mario Paolucci and Gennaro Di Tosto use an evolutionary variant of the Micro-Macro Link (MML) theory in biological evolution to understand the emergence of altruism, applied to food sharing among vampire bats. Behaviour at the individual level generates higher level structures (bottom-up) which feed back to the lower level (top-down). Starting from ethological data a multi-agent model is used to analyse the key features of altruistic behaviour. Every agent in the simulation is designed to reproduce hunting and social activity of the common vampire bats. During night, the simulated animals hunt, during day they perform social activities (grooming and food-sharing). A high number of small groups (roosts) provide social barriers preventing altruists from being invaded by non-altruists (simple loop). When the ecological conditions vary (e.g., the number of individuals per group increases), altruism is at risk, and other properties at the individual level evolve in order to keep non-altruists from dominating, and to protect the whole group (complex loop). The two loops are illustrated by simulation experimenting on individual properties, allowing altruists to survive and neutralise non-altruists even under unfavourable demographic conditions.

6.4 General Aspects

To establish the potential importance of the interplay between social and physical spaces, *Bruce Edmonds* exhibits a couple of agent-based simulations which involve both physical and social spaces. The first of these is a more abstract model whose purpose is simply to show how the topology of the social space can have a direct influence upon spatial self-organisation, and the second is a more descriptive model which aims to show how a suitable agent-based model may inform observation of social phenomena by suggesting questions and issues that need to be investigated. Taking the physical and social embeddedness of actors seriously, their interactions in both of these “dimensions” need to be modeled. In his view, agent-based simulation seems to be the only tool presently available that can adequately model and explore the consequences of the interaction of social and physical space. It provides the “cognitive glue” inside the agents that connects physical and social spaces.

To build an agent-based computational model of a specific socio- environmental system, *Jim Doran* discusses designs to create the software agents. The currently available range of agent designs is considered, along with their limitations and inter-relationships. How to choose a design to meet the requirements of a particular modelling task is illustrated by reference to designing an informative agent-based model of a segmented, polycentric and integrated network (SPIN) organization. As an example, a social movement in the context of environmental activism is discussed, representing a segmentary, polycentric and integrated network composed of many diverse groups, which grow and die, divide and fuse, proliferate and contract. The adaptive structure of SPINs prevents effective suppression by authorities and opponents, an aspect that

is relevant for the stability and disruption of networks, in particular terrorist networks.

References

1. Axelrod, R. (1997) *The Complexity of Cooperation. Agent-Based Models of Competition and Collaboration*. Princeton University Press, Princeton.
2. Axtell, R. (2000) Why agents? On the varied motivations for agent computing in the social sciences. Washington, The Brookings Institution, Center on Social and Economic Dynamics, 22 p. (CSED Working Paper, n 17). Available online at: <http://www.brook.edu/dybdocroot/es/dynamics/papers/agents/agents.htm>.
3. Billari, F.C. (2000) Searching for mates using 'fast and frugal' heuristics : a demographic perspective. In: Gerard Ballot and Gerard Weisbuch (ed.), *Applications of simulation to social sciences*, p. 53-65. Oxford, Hermes Science Publications, 461 p.
4. Billari, F.C., Prskawetz, A. (ed.) (2003) *Agent-Based Computational Demography. Using Simulation to Improve our Understanding of Demographic Behaviour*, Heidelberg, Physica/Springer.
5. Billari, F.C., Ongaro, F. and Prskawetz, A. (2003) Introduction: agent-based computational demography. In: Billari and Prskawetz (2003), pp. 1-17.
6. Becker, G.S. (1981) *A treatise on the family*. Cambridge (Mass.), Harvard University Press (2nd edition 1991).
7. Behrman, J.R. (2001) Why micro matters. In: Nancy Birdsall, Allen C. Kelley and Steven W. Sinding (ed.), *Population matters. Demographic change, economic growth and poverty in the developing world*, p. 371-410. Oxford, Oxford University Press.
8. Coleman, J.S. (1990) *Foundations of Social Theory*. Cambridge (Massachusetts), Belknap Press.
9. Conte, R., Castelfranchi, C. (1995) *Cognitive and social action*. London, UCL Press limited.
10. Courgeau, D. (1995) Migration theories and behavioral models. *International Journal of Population Geography*, vol. 1, no.1, p. 19-27.
11. Courgeau, D. (2003) General introduction. In: Daniel Courgeau, *Methodology and epistemology of multilevel analysis. Approaches from different social sciences*, p. 1-23. Boston, Dordrecht and London, Kluwer Academic Publishers.
12. Epstein, J.M., Axtell, R. (1997) *Growing Artificial Societies*. Cambridge MIT Press.
13. Gaylord, R.J., D'Andria, L.J. (1998) *Simulating Society - A Mathematica Toolkit for Modeling Socioeconomic Behaviour*. Heidelberg, Springer/Telos.
14. Gilbert, N., Troitzsch, K.G. (2000) *Simulation for the Social Scientist*. Buckingham, PA, Open University Press.
15. Grimm, V. (1999) Ten years of individual-based modeling in ecology. What have we learned, and what could we learn in the future? *Ecological Modelling*, vol. 115, no. 2-3, p. 129-148.
16. Grimm, V. and Railsback, S.F. (2005) *Individual-based Modeling and Ecology*. Princeton, Princeton University Press.
17. Haken, H. (1977) *Synergetics*. Heidelberg, Springer.

18. Halpin, B. (1999) Simulation in Society. *American Behavioral Scientist* 42(10), p. 1488-1508.
19. Hammel, E.A., McDaniel, C.K. and Wachter, K.W. (1979) Demographic consequences of incest tabus: a microsimulation analysis, *Science*, vol. 205, no. 4410, p. 972-977.
20. Hammel, E.A. and Wachter, K.W. (1996) Evaluating the Slavonian Census of 1698. Part II : A microsimulation test and extension of the evidence. *European Journal of Population*, vol. 12, no. 4, p. 295-326.
21. Hedström, P. and Swedberg, R. (1999) *Social mechanisms. An analytical approach to social theory.* New York and Cambridge (UK), Cambridge University Press.
22. Helbing, D. (1995) *Quantitative Sociodynamics Stochastic Methods and Models of Social Interaction Processes.* Boston/London/Dordrecht: Kluwer.
23. Johnson, P.E. (1999) Simulation Modeling in Political Science. *American Behavioral Scientist* 42(10), p. 1509-1530.
24. Kirman, A.P. (1992) Whom or what does the representative individual represent?, *Journal of Economic Perspectives*, vol. 6, no. 2, p. 117-136.
25. Land, K.C. (1986) Method for national population forecasts: a review. *Journal of the American Statistical Association*, vol. 81, no. 347, p. 888-901.
26. Macy, M.W., Willer, R. (2002) From Factors to Actors: Computational Sociology and Agent-Based Modeling. *Annual Review of Sociology* 28, p. 143-166.
27. Nakazawa, M. and Ohtsuka, R. (1997) Analysis of completed parity using microsimulation modeling. *Mathematical Population Studies*, vol. 6, n 3, p. 173-186.
28. Rahmandad, H. and Sterman, J. (2004) Heterogeneity and network structure in the dynamics of diffusion: comparing agent-based and differential equation models. Cambridge (Massachusetts), Massachusetts Institute of Technology (MIT Sloan Working Paper 4512-04).
29. Ridley, J.C. and Sheps, M.C. (1996) An analytic simulation model of human reproduction with demographic and biological components, *Population Studies*, vol. 19, no. 3, p. 297-310.
30. Ruggles, S. (1993) Confessions of a microsimulator. *Historical methods*, vol. 26, no. 4, p. 161-169.
31. Scheffran, J. (2006) Tools in Stakeholder Assessment and Interaction. In: S. Stoll-Kleemann, M. Welp (Eds.), *Stakeholder dialogues in natural resources management and integrated assessments: Theory and practice* (forthcoming).
32. Schelling, T.C. (1978) *Micromotives and Macrobehavior* New York: Norton.
33. Schweitzer, F. (ed.) (1997) *Self-Organization of Complex Structures: From Individual to Collective Dynamics*, vol. II. London, Gordon and Breach.
34. Tesfatsion, L. (Eds.) (2001) Special Issue of Agent-Based Computational Economics. *Journal of Economic Dynamics & Control* 25, p. 281-654.
35. Tomassini, C. and Wolf, D. (2000) Shrinking kin networks in Italy due to sustained low fertility, *European Journal of Population*, vol. 16, no. 4, p. 353-372.
36. van den Bergh, J. C.J.M. and Gowdy, J.M. (2003) The microfoundations of macroeconomics: an evolutionary perspective. *Cambridge Journal of Economics*, vol. 27, no. 1, p. 65-84.
37. van Imhoff, E. and Post, W.J. (1998) Microsimulation methods for population projection. *Population: An English Selection*, vol. 10, no. 1, p. 971-938. (Special issue on "New Methodological Approaches in the Social Sciences").

38. Wachter, K. W. (1987) Microsimulation of household cycles. In: Bongaarts, J., Burch, T. K., Wachter, K.W. (Eds.): Family demography. Methods and their application, Clarendon Press, Oxford.
39. Weidlich, W. (2000) Sociodynamics - A Systematic Approach to Mathematical Modelling in the Social Sciences. Harwood Academic Publishers.
40. Wolf, D. (2001) The role of microsimulation in longitudinal data analysis. Canadian Studies in Population, vol. 28, n 2, p. 313-339.
41. Wolfram, S. (2002) A New Kind of Science. Champaign, Wolfram Research.
42. Zimmermann, K.F. and Vogler, M.(ed.) (2003) Family, Household and Work. Heidelberg, Springer.

Agent-Based Modelling – A Methodology for the Analysis of Qualitative Development Processes

Andreas Pyka¹ and Thomas Grebel²

¹ University of Augsburg, Germany
`andreas.pyka@wiwi.uni-augsburg.de`

² GREDEG-IDEFI, CNRS, Valbonne, France
`grebel@idefi.cnrs.fr`

Summary. In economics numerical approaches are increasingly used for the analysis of dynamic phenomena of economic development since almost 30 years. The first generation of simulation models was rather focused on stylized empirical phenomena. With the emergence of agent-based modelling the last 10 years, however, the trade-off between simplicity and abstracting in modelling, and taking into account the complexity of the socio-economic reality has been enhanced to a large extent. This paper serves as a basic instruction on how to model qualitative change using an agent-based modelling procedure. The necessity to focus on qualitative change is discussed, agent-based modelling is explained and finally an example is given to show the basic simplicity in modelling.

1 Introduction

The tremendous development of an easy access to computational power within the last 30 years has led to the widespread use of numerical approaches in almost all scientific disciplines. Nevertheless, while the engineering sciences focused on the applied use of simulation techniques from the very beginning, in the social sciences most of the early examples of numerical approaches were purely theoretical. There are two reasons for this. First, since the middle of the 20th century, starting with economics, equilibrium-oriented analytical techniques flourished and were developed to a highly sophisticated level. This led to the widely shared view that within the elegant and formal framework of linear analysis offered by neoclassical economics, the social sciences could reach a level of accuracy not previously thought to be possible. Including also important non-linearities to this framework, on the one hand was opening the discussion of important dynamic phenomena, on the other hand, however, this was already questioning the achievement of accuracy due to the problem of multiple equilibria and the difficulties of equilibrium selection (e.g. [6]).

Second, within the same period, new phenomena of structural change exerted a strong influence on the social and economic realms. Despite the mainstream neoclassical successes in shifting the social sciences to a more mathematical foundation, an increasing dissatisfaction with this approach emerged. For example, by the 1970s the benchmark of atomistic competition in neoclassical economics had already been replaced by the idea of monopolistic and oligopolistic structures under the heading of workable competition (e.g. [28]). A similar development emphasizing positive feedback effects and increasing returns to scale caused by innovation led to the attribute “new” in macroeconomic growth theory in the 1980s [26]. In addition to these stepwise renewals of mainstream methodology, an increasingly larger group claimed that the general toolbox of economic theory, emphasizing rational behaviour and equilibrium, is no longer suitable for the analysis of complex social and economic changes. In a speech at the International Conference on Complex Systems organized by the New England Complex Systems Institute in 2000, Kenneth Arrow stated that until the 1980s the “sea of truth” in economics lay in simplicity, whereas since then it has become recognized that “the sea of truth lies in complexity”. Adequate tools have therefore to include the heterogeneous composition of agents (e.g. [27]), the possibility of multilevel feedback effects (e.g. [4]) and a realistic representation of dynamic processes in historical time (e.g. [1]). These requirements are congruent with the possibilities offered by simulation approaches. It is not surprising that within economics the first numerical exercises were within evolutionary economics, where phenomena of qualitative change and development are at the front of the research programme. The first generation simulation models were highly stylized and did not focus on empirical phenomena. Instead, they were designed to analyse the logic of dynamic economic and social processes, exploring the possibilities of complex systems behaviour. However, since the end of the 1990s, more and more specific simulation models aiming at particular empirically observed phenomena have been developed focusing on the interaction of heterogeneous actors responsible for qualitative change and development processes. Modellers have had to wrestle with an unavoidable trade-off between the demands of a general theoretical approach and the descriptive accuracy required to model a particular phenomenon. A new class of simulation models has shown to be well adapted to this challenge, basically by shifting outwards this trade-off:³ So-called agent-based models are increasingly used for the modelling of socio-economic developments. Our Chap. deals with the changed requirements for modelling caused by the necessity to focus on qualitative developments which is generally highlighted within evolutionary economics and the possibilities given by agent-based models. The next Sect. is concerned with the importance of an analysis of qualitative development and it is shown that evolutionary economics is offering an adequate framework for this. Section 3 then focuses on agent-based-modelling as “the” tool that allows incorporat-

³ See e.g. [12].

ing endogenously caused development processes. Section 4 gives an illustrative example of an agent-based-model. Section 5 summarizes the whole story.

2 Qualitative Change in an Evolutionary Economics Perspective

When concerned with the examination of change and development processes within industrialized economies economists usually focus on the movement of certain variables they consider a good description of the basic effects of economic growth and development. In mainstream economics the phenomenon of economic development is e.g. empirically analysed on the macro-economic level as the improvement of total factor productivity in time which lowers prices and leads to the growth of incomes. Accordingly, most often the GDP per capita is used as an indicator describing economic development in a quantitative fashion. Although it is impressing to observe the growth of income in economies over a long time span, this indicator, due to its quantitative nature only, does not give any idea about the structural and qualitative dimensions underlying economic development. This becomes even more obvious on the sectoral level where the analysis is most often restricted to long-run equilibrium structure describing e.g. the number of firms in a particular industry without putting emphasis on those factors driving the emergence and maturation of industries. By restricting their analysis on the quantitative dimension, the economic mainstream implicitly confines itself to the analysis of a system characterized by a constant set of activities basically neglecting innovation processes.⁴

However, in less orthodox economic approaches it is argued, and it is indeed also one of Schumpeter's major contributions that economic development does also include prominently qualitative changes not only as an outcome but also as an essential ingredient which justifies us to speak of transformation processes going on. Qualitative change manifests itself basically via innovation of different categories of which technological innovation very likely is among the most important ones (others are social, legal, organizational changes). Qualitative change is the transformation of an economic system, characterized by a set of components and interactions into another system, with different components and different interrelationships (e.g. [27]). An analysis of qualitative change therefore necessarily has to include the actors, their activities and objects which are responsible for the ongoing economic development. An example for the significance of qualitative changes can be found in Fig. 1 which displays the emergence of new industries in the internet sector in the 1990s for Germany by showing the number of firm entries. What strikes immediately is

⁴ [10]: Economic growth can be described at the macro-economic level, but it can never be explained at that level. Economic growth results from the interaction of a variety of actors who create and use technology and demanding costumers.

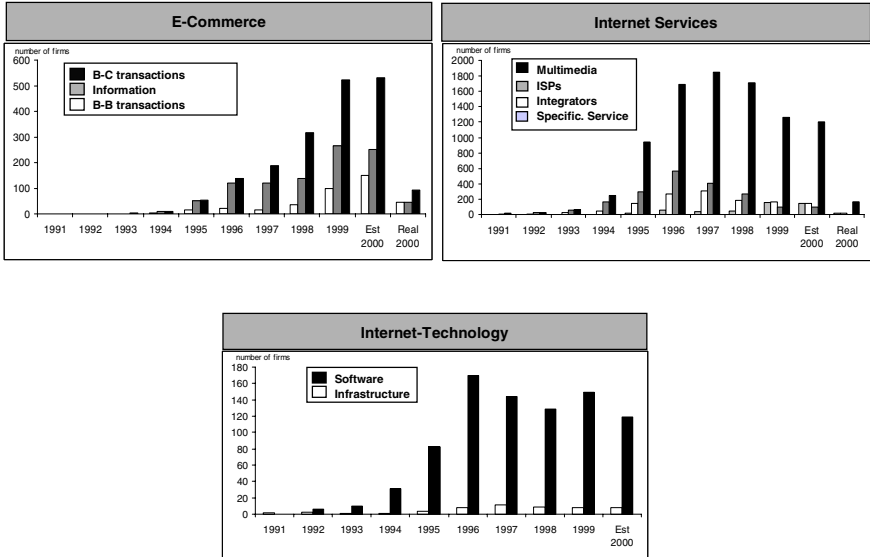


Fig. 1. Swarms of new firms in the internet industries (Germany)
 Source: e-startup.org database, survey of 332 venture capital firms, analyses of insolvency databases, Newsfeed and other public sources.

that anything but an equilibrated regular or proportional development is visible: Instead new firms appear in swarms, to use a notion coined by Schumpeter [29] and sometimes almost no activities occur. Of course there are many other variables which do also reflect the importance of the qualitative dimensions of economic development e.g. on a macro-economic level the changing composition of the employment structure (Fourastier Hypothesis), on a meso-level the regional specialization patterns or on a micro-economic level the obsolescence of old and the emergence of new knowledge like the biotechnology revolution in pharmaceuticals, to name a few. By its very nature, the transformation of an economic system is a multi-faceted phenomenon. Accordingly, it is misleading to focus only on quantitative changes of the economy when analysing the driving factors of the transformation of economic systems over time. To better understand the mechanisms and dynamics behind the observed developments one has to explicitly include the qualitative dimensions. To achieve this, economic analysis has to consider – besides the prevailing cost-orientation – an important knowledge- and learning-orientation.

The following paragraphs are concerned with the implications of this knowledge-orientation, which can also be considered as the heart of the matter of evolutionary economics.

Knowledge-Based Approach of Evolutionary Economics

It is beyond the scope of this contribution to discuss in detail the criticism brought forth by evolutionary economics with respect to assumptions underlying the mainstream economic reasoning. A major discussion can be found among others, in [7], [5] and [31]. For our purposes it is sufficient to mention three major points, evolutionary economists claim to be of outstanding importance in the discussion of economic development processes and which are incompatible with traditional economic approaches. These points are also constitutive for that strand of literature within evolutionary economics which is concerned with industry evolution and technological progress namely the Neo-Schumpeterian approach. Here, instead of the resource- or incentive-orientation of neoclassical industrial economics a knowledge-orientation focuses on the investigation of industries and innovation processes in particular.

- First of all, the Neo-Schumpeterian theory wants to explain how innovations emerge and diffuse over time. A specific feature of these processes is uncertainty, which cannot be treated adequately by drawing on stochastic distributions referring to the concept of risk. Therefore, the assumption of perfect rationality, underlying traditional models cannot be maintained, instead the concepts of bounded and procedural rationality are invoked. Consequently, actors in Neo-Schumpeterian models are characterized by incomplete knowledge bases and capabilities.
- Closely connected, the second point concerns the important role heterogeneity and variety plays. Due to the assumption of perfect rationality i.e. optimal decisions, in mainstream models homogeneous actors and technologies are analysed. E.g. every deviation from an optimal technology would lead to the exit of the respective firm applying by definition a sub-optimal technology. Heterogeneity as a source of learning and novelty is by and large neglected or treated as an only temporary deviation.
- Finally, the third point deals with the time dimension in which learning and the emergence of novelties take place. By their very nature, these processes are truly dynamic, meaning that they occur in historical time. The possibility of irreversibility, however, does not exist in the mainstream approaches, relying on linearity and equilibrium.

Thus, traditional economic theories, summarized under the heading of incentive-based approaches, with their focus on cost-based and rational decisions only, are excluding crucial aspects of actors' behaviours and interactions, which are influenced by a couple of factors lying by their very nature beyond the scope of these approaches. Although, of course, cost-benefit calculations (with respect to innovation itself a problematic activity) play an important role, the actors' behaviour is influenced additionally by several other factors as learning, individual and collective motivation, trust etc. It is the role of these factors the knowledge-based approach of evolutionary economics explicitly takes into account.

By switching from the incentive-based perspective to the knowledge-based perspective the Neo-Schumpeterian approaches have realized a decisive change in the analysis of the transformation of economic systems. In this light the introduction of novelties mutate from optimal cost-benefit considerations to collective experimental and problem solving processes [9]. The knowledge-base of the actors is no longer perfect, instead a gap between the competencies and difficulties which are to be mastered opens up ([16, C-D gap]). There are two reasons responsible for this C-D gap when it comes to innovation: on the one hand, technological uncertainty introduces errors and surprises. On the other hand, the very nature of knowledge avoids an unrestricted access. Knowledge in general, and new technological know-how in particular, are no longer considered as freely available, but as local (technology specific), tacit (firm specific), and complex (based on a variety of technology and scientific fields). To understand and use the respective know-how, specific competencies are necessary, which have to be built up in a cumulative process in the course of time. Following this, knowledge and the underlying learning processes are important sources for the observed heterogeneity among agents.

Challenges for Analysing Qualitative Change

From the discussion above we can identify two major challenges for an analysis of qualitative change:

The first challenge is that a theoretical framework adequately displaying our notion of qualitative change has to incorporate concepts that comply with the notion of development of evolutionary economics in the sense Nelson [25] discussed. Basically he refers to path-dependencies, dynamic returns and their interaction as constitutive ingredients for evolutionary processes in the socio-economic realm.

The second challenge is that we generally have to focus on both the micro- and meso-level of the economy as to our understanding the term qualitative change refers to a changing composition of components and interaction of and in the economic system. In doing so, we can identify some stylized facts that are considered of crucial importance when qualitative change in an economy is considered. The most obvious ones are:

First, an increasing importance of knowledge generation and diffusion activities is observed at least in those sectors of the economy that are considered to be the most dynamic and innovative ones. This coins the notion of a transformation of the economy into a knowledge-based economy. Second, this is accompanied by a continuously increasing specialisation and related to this an increasing variety of products and services coexisting simultaneously. Third, specialisation and differentiation goes hand in hand with an increasing importance of (market and non-market) interactions between the agents. Fourth, behind this increasing variety we observe innovation processes that at the same time improve efficiency of the production process and the quality of

the products. Fifth, this innovation process is driven by competition selecting between different technological alternatives. Finally, the environmental constraints can be considered as filter- and focusing devices in this selection process either supporting or suppressing the diffusion of new technologies.

Once the relevance of these facts for the transformation of an economy is accepted the research has to account for those developments adequately.

Micro- and Meso-Perspective

Obviously this aim can only be accomplished by abandoning an aggregate perspective and instead focusing on a micro- or meso-level population approach [23]. This allows for examining diverse agents, their interaction and the knowledge-induced transformation of both. By doing this, modelling openly has to take into account the importance of micro-macro-micro feedback effects (e.g. [31]). In their decisions actors obviously consider macro (-economic) constraints, but they also exert a significant influence on the altering of these constraints [7]. The interrelated inspection of the meso- and the micro-level reflects the idea that analysis on the aggregated meso-level relies on description whereas the analysis of the micro-level focuses on explanation of the phenomena found on the meso-level [7].

Knowledge

Considering this will lead to a revision of standard economic models as analysis here follows reality closely. Traditional 'production functions' include labour, capital, materials and energy. Knowledge and technology are only external influences on production. However, recent analytical approaches have been developed allowing the explicit consideration of knowledge as well as learning of actors as a means of acquiring new knowledge. Improvements in the knowledge base are likely not only to increase the productive capacity of other production factors, leading to the introduction of new products, as a visible outcome of the transformation process, but also to alter the organizational processes of knowledge creation, namely the interrelationships between the actors. Thus, transformation relates to a result- and a process-dimension similar to the terminology elaborated in [17].

Consequently, it cannot be assumed that there exists a fixed set of activities and relationships in the social and economic sphere, especially when it comes to knowledge generation and learning. But this does by no means imply that no such set exists at all. It does exist, although, by its very nature it is evolving continuously. In this respect transformation does not only refer to the feedback processes, but it does also and with major relevance refer to the change of the set itself during the process. This is evolution, and evolution is the very reason for not using static equilibrium theories or dynamic models to analyse qualitative developments as they are based on the notion of reversibility. The notion of evolution demands that we resort to ideas of irreversibility and path-dependence.

3 A Modelling Approach Allowing for Qualitative Change: Agent-Based Modelling

An exploration of settings fulfilling the above requirements very likely needs numerical techniques, which are regarded as a major tool in evolutionary economics ([19],[2]). Although simulation analysis comes in various flavours most of them reflect Boulding's call that we need to develop 'mathematics which is suitable to social systems, which the sort of 18th-century mathematics which we use is not' [3]. An increasingly growing literature today now is concerned with the application of so-called agent-based models. This approach consists of a decentralized collection of agents acting autonomously in various contexts. The massively parallel and local interactions can give rise to path dependencies, dynamic returns and their interaction. In such an environment global phenomena such as the development and diffusion of technologies, the emergence of networks, herd-behaviour etc. which cause the transformation of the observed system can be modelled adequately. This modelling approach focuses on depicting the agents, their relationships and the processes governing the transformation. Very broadly, the application of an *agent based modelling approach* offers two major advantages with respect to the knowledge- and learning-orientation:

The first advantage of agent based models is their capability to show how collective phenomena came about and how the interaction of the autonomous and heterogeneous agents leads to the genesis of these phenomena. Furthermore agent-based modelling aims at the isolation of critical behaviour in order to identify agents that more than others drive the collective result of the system. It also endeavors to single out points in time where the system exhibits qualitative rather than sheer quantitative change [32]. In this light it becomes clear why agent-based modelling conforms with the principles of evolutionary economics ([20], [21]). It is 'the' modelling approach to be pursued in evolutionary settings.

The second advantage of agent-based modelling, which is complementary to the first one, is a more normative one. Agent-Based models are not only used to get a deeper understanding of the inherent forces that drive a system and influence the characteristics of a system. Agent based modellers use their models as computational laboratories to explore various institutional arrangements, various potential paths of development so as to assist and guide e.g. firms, policy makers etc. in their particular decision context.

Agent-Based modelling thus uses methods and insights from diverse disciplines such as evolutionary economics, cognitive science and computer science in its attempt to model the bottom-up emergence of phenomena and the top down influence of the collective phenomena on individual behaviour.

The recent developments in new techniques in particular the advent of powerful tools of computation such as evolutionary computation (for a summary of the use of evolutionary computation and genetic programming in par-

ticular see [8]) opens up the opportunity for economists to model economic systems on a more realistic i.e. more complex basis [32].

There is no entity, even though it may exist without the actors, which has no influence on the current state of the system or the development of the system. To illustrate this point, bits of information have no influence on the system as long as they are not put into the appropriate context by a capable individual, influencing its activities. No resource can change the system as long as it is not used for carrying out certain activities that change the nature and the structure of the system. Hence in the centre of the stage there is the actor and its activities.

In the following Sects. a typical example for an agent-based model is introduced in order to highlight the specialties of this methodology. In particular the model deals with the emergence of new firms which are considered the outcome of entrepreneurial decisions of individual agents pooled together in networks. As the focus of this Chap. lies on the methodology of agent-based modelling we cannot go into detail with respect to the economic implications of the model but refer instead to [14] and [15] where all the economic concepts and the formal description used in the model are described in detail.

4 An Illustrative Example: An Evolutionary Economics Model of Entrepreneurial Behaviour

4.1 The General Building Blocks

A conceptual framework for the analysis of entrepreneurial behaviour can be composed of several building blocks. In particular we consider *actors*, *action*, *endowments*, *interaction and evaluation and decision processes* as the decisive building blocks. However, the building blocks discussed here are not separate and unrelated entities. Rather they are the result of a systematization process. They represent our conceptual view on the issue, developed to clarify the analytical concepts, and to facilitate the implementation of the simulation model in the second step. In the following Sects. we sketch the building blocks.

Actors

We consider actors and their interactive decisions being the major driving force in the evolution. As such we regard them as the reason for the manifestation of qualitative developments going on in the system. They are the crucial components of the system. The model requires a multi-agent approach, which assumes that agents populating the model can be divided into various categories according to their initial endowments concerning the availability of capital, an entrepreneurial attitude as well as the respective technological competencies.

Accordingly, a central issue is the general design of the actors. Actors are represented as code that has the standard attributes of intelligent agents [33]:

- *autonomy*, which means that agents operate without other agents having direct control of their actions and internal states. This is a necessary condition for implementing heterogeneity.
- *social ability*, i.e. agents are able to interact with other agents not only in terms of competition but also in terms of cooperation. This includes the possibility to model agents that show various forms of interaction blended from competition and cooperation.
- *reactivity*, agents are able to perceive their environment and respond to it.
- finally, *proactivity* enables the agents to take the initiative. This means that they are not only adapting to changing circumstances, rather are they engaged in goal-directed behaviour.

The above points indicate that the actors in the simulation are able not only to adapt their behaviour to a given set of circumstances but they are also in a Neo-Schumpeterian sense able to learn from their own experience and to modify their behaviour creatively so as to change the circumstances themselves.

When modelling the features and characteristics of the artificial agents the above mentioned standard attributes have to be implemented. As the agents in our conceptual framework can be characterized by their actions, endowments, interactions and their evaluation and decision processes, these conceptual building blocks have to be designed such as to reflect these attributes.

Actions

The different actions (e.g. founding decisions, production decisions etc) performed by the actors enable us to classify certain groups of actors. Not only is it the actions that we use as a demarcation of different groups of actors - their endowment might be another criterion for differentiation - but actions is one of the most striking one and connected to the other features such as endowments, etc. that will be discussed below. Basically we distinguish between individual agents and firms as networks of agents.

Routines

The actors are not modelled by a representative agent but by a population of heterogeneous agents. For any of our two subpopulations (agents and firms) rules and routines can be derived which govern the particular actions of the agents, the interaction and the interrelation of the agents within and among the sub-populations. Actions and routines are conceptually closely related and the latter can be considered as realizations of actions.

Hence it is routines through which the actors manipulate reality. It is not only the endowment with resources that shapes the nature of the actors, it is their individual routines that make up a large part of the actors heterogeneity.

Nelson and Winter [24] relate routines to the satisficing behaviour and the bounded rationality of actors.⁵ Routinized behaviour causes some stickiness and some inertia of the system that results in some stability of the system - stability, at least to a certain degree.

Evaluation and Decision Processes

The discussion up to this point reveals that we have to cope with a heterogeneous set of actors. Some actors join networks with other actors and found a firm, other disentangle their networks or even go bankrupt with their previously founded firm. The question here is, how to unify the decision process of such a diverse set of actors while preserving the possibility for heterogeneity. After having introduced the basic conceptual building blocks in a rather abstract and general way the following Sects. deal with the actual model of entrepreneurial behaviour.

4.2 Modelling Entrepreneurial Behaviour

The starting point of the model is the micro-level. The driving force of an agent-based model is the agent. While an incentive-based model would rather focus on facts and phenomena, external to the actor, a knowledge-based view has to thoroughly investigate the agent. This also raises methodological issues. In principal, the former - that is orthodox methodology - uses the Newtonian mechanics which requires the concept of a *homo oeconomicus* as a necessary condition. The *homo oeconomicus* performs a robustly optimal behaviour. In case behaviour is deterministic, the usage of analytical tools (equilibrium analysis) becomes legitimate. In return, this methodology makes it difficult to discuss psychological and sociological aspects of agents. The *homo oeconomicus* has been deprived from any psychological and sociological qualities that indeed affect individual (economic) behaviour. As much as orthodox methodology asks for such a perfectly rational and therefore homogeneous agent within a supposedly deterministic world, the need to shed some light on the non-deterministic aspects - the heterogeneity in agents' behaviour - asks for an adequate methodology. Agent-Based modelling allows us to cope with the complexity emerging from the behaviour of heterogeneous actors.

In the following, a sketch of an agent-based model of entrepreneurial behaviour will be drawn.

Actors

Actors are boundedly rational. Their current individual state is the result of an ongoing path-dependent, cumulative and irreversible process. Their knowl-

⁵ An example of a routine applied by agents in the innovation process is: Invest x percent of the turnover of the previous period in R&D.

edge, their capabilities and their resources are the result of a congenitally determined learning and decision-making process. On these grounds, actors will make future decisions.

When we exemplarily investigate the emergence of entrepreneurial behaviour, a stereotypical agent may look as follows: the decision to become an entrepreneur might be driven by entrepreneurial traits ([29], [22]), by its knowledge and capabilities acquired by education and work experience ([18], [13]) and last not least, sufficient financial resources ([30], [11]). For simplicity, these three components, we call the entrepreneurial (*ec*), the capability (*cc*) and the financial (*fc*) component of our basic, bounded rational actor as shown in Fig. 2.




Fig. 2. An actor and its endowment

Using a Schumpeterian concept of the entrepreneur, innovative behaviour will be dependent on the actor's capacity to make use of a new technology. Besides the fact that the actor's cognitive capacity might prevent him/her from innovating on a new technology in the first place, the possibility of not receiving the knowledge about the new technology may never make an actor an entrepreneur either, though having the potentials. As a result, the diffusion of knowledge is constrained by individual factors as well as the fashion of social interaction, which is the means to pass on such knowledge. This knowledge diffusion process can be easily modelled with a cellular automaton using percolation theory [15].

For simplicity let us now consider only those agents who have received and understood the application of a new technology. Then, it is more probable that these agents might undertake entrepreneurial actions, even if this is not a necessary consequence.

Social Interaction

An entrepreneurial decision cannot be considered in isolation. The context of a social group plays an important role in an actor's decision-making process: either supporting or disapproving a decision such as starting a new business. Some actors might be interested in a new technology (e.g. the internet) and begin proactively to gather new information and knowledge about it. Thus, a dynamic social interaction process keeps the agents forming new networks and thus building and restructuring connections as depicted in Fig. 3.




Fig. 3. Actors forming a network

This social networking dynamic is an indeterministic, quasi-random process⁶. With a cellular automaton [15] such quasi-random behaviour can be implemented into a model: placing the actors on a lattice, giving them certain rules when and where to move, a self-organizing process evolves that makes actors of a kind (having a comparable set of endowments) happen to bump into each other.

At any time an actor evaluates his/her chances to found a firm successfully, and so does he/her evaluate the chances of other network members to start a business and finally, they may decide to establish a firm altogether. See Fig. 4.




Fig. 4. A firm founded by three actors

So far the micro-level has been substantiated with the specificities of the actors, their endowments, their routines behaviour and their social interaction. At any point in time, each network constitutes a potential firm.

⁶ Quasi-random means that such a search process is neither perfectly deterministic saying that the result of this process is always the optimal network, nor is the search process completely chaotic. Agents act goal-oriented but not perfectly rational.

Micro-Meso/Macro Feedback

Up to this point, the modelling procedure was strictly bottom-up. However, this is not the whole story. Though indeed entrepreneurial decisions are micro-level decisions with all the psychological and sociological aspects involved, the meso and macro level also have a major influence on such kind of decisions. The emergence of new industries (e.g. e-commerce, etc.) is an endogenous process. It is an ever-changing process principally driven by micro behaviour. Nevertheless, economic indicators such as the rate of entry and exit, market concentration and the stage of an industry's life cycle have a feedback effect from the meso/macro level onto the micro level. Hence, the agents create their common economic reality and at the same time are guided by the same economic reality.

As a consequence, economic data (entry, exit, etc.) have to be taken into account within the decision-making process of the agents. Whether an actual firm is actually established not only depends on individual factors, the self and the group evaluation process, but also it depends on the evaluation of economic opportunities of a new technology, i.e. meso and macro data.

A Heterogeneous Oligopoly

In order to implement this endogenous change, we have to use a module that produces this data given the agents' actions. This module is a simple heterogeneous oligopoly module [15], which produces the data required. Once firms are founded, they take part in market competition. Each firm faces an individual demand curve which depends on the firm's competitiveness relative to the remaining incumbent firms' competitiveness. Thereby, the competitiveness is determined by the firm's balance in endowments. For example, firms that have less in the capacity component have worse chances than others. Firms learn over time and improve their efficiency, i.e. there is a first-mover advantage. This way, the firm's competitiveness is an endogenous result of the quasi-random search process of the agents.

The Founding Threshold

The heterogeneous oligopoly only serves as a selection criterion to generate the necessary data which influences micro behaviour. The module is interchangeable. The continuously produced data is fed via the so-called founding threshold into the decision-making process of the agents. The founding threshold is perceived by the agents from the observation of the overall industry performance (e.g. number of entry and exits). Thus, the focus on micro behaviour is guaranteed and the model is kept parsimonious and simple. See Fig. 5 .




Fig. 5. Micro-macro-feedback effects

Results

An agents-based model as it is developed in [14] has been conceptually outlaid above. The formal part and the code can be found in the literature cited.

Now, with the model as it was stated above various scenarios can be run. The endowments of actors can be fleshed out with empirical data or, as it was done in the simulation run shown below, can just be pseudo-random numbers. So it is done with the routine behaviour of agents. Certain, specified rules make them interact with each other. Hence, the start-up decision is an economic, irreversible decision, contingent to psychological and sociological aspects.

As Fig. 6 shows firms are founded by agents driven by the positive data generated by market competition. The founding threshold thereby depicts the ups and downs of the agents common attitude towards the economic development, which agents adapt their behaviour to.

Figure 7 serves to illustrate the heterogeneity among firms. Each firm founded has its individual competitiveness relative to others. Not all of the firms are successful and survive the early phase of competition. Some become insolvent and may exit the market, whereas others survive and grow. Furthermore, it has to be emphasized that the heterogeneity of firms is no arbitrary assumption but the result of a decision-making process of bounded rational and therefore heterogeneous actors.

5 Conclusions

The Chap. deals with one of the most prominent challenges in social sciences today, namely the analysis of qualitative change. It is shown that evolutionary economics is offering an adequate framework for this, overcoming the severe




Fig. 6. A selection of simulation results

restrictions orthodox economic approaches are confronted with. Emphasizing the role of learning, true uncertainty, heterogeneity of agents and irreversibility, within evolutionary economics qualitative development becomes an endogenous process driven by agents and their interactions.

Agent-Based models allow for an explicit consideration of these characteristic features and therefore can be considered as “the” modelling tool for the analysis of qualitative development and transformation processes. After having worked out the basic features and requirements of agent-based models, their functioning is exemplarily shown by introducing an agent-based model of entrepreneurial behaviour. This particular model is also demonstrating a second crucial advantage of this modelling technique: agent-based models are offering a platform for inter- and trans-disciplinary research. In the model of entrepreneurial behaviour for example, several insights from psychology




Fig. 7. The performance of new firms

as well as the theory of social networks are embedded. In a way, agent based models can be considered a systemic approach, allowing the consideration and integration of different social “realities” which makes them an extremely valuable tool for the analysis of social processes which generally can be considered as multifaceted phenomena.

References

1. Arthur, W. B. (1988) Competing Technologies: an Overview. In: Dosi, G. et al. (eds.) *Technical Change and Economic Theory*. Pinter, London, pp. 590-607.
2. Aruka, Y. (2001) *Evolutionary Controversies in Economics: A New Transdisciplinary Approach*. Springer, Tokyo, Berlin and others.
3. Boulding, K. E. (1991) What is evolutionary economics? *Journal of Evolutionary Economics*, 1, 9-17.
4. Cantner, U., Pyka, A. (1998) Technological Evolution – an Analysis within the knowledge-based Approach. *Structural Change and Economic Dynamics*, Vol. 9, pp. 85-108.
5. Clark, N. and Juma, C. (1987) *Long Run Economics – An Evolutionary Approach to Economic Growth*. Pinter Publishers, London.
6. Delli Gatti M., Gallegatti, M., Kirman, A. (2000) *Interaction and Market Structure, Essays on Heterogeneity in Economics*. Series: *Lecture Notes in Economics and Mathematical Systems*, Vol. 484, Springer, Berlin.
7. Dopfer, K. (2001) Evolutionary economics: Framework for analysis. In: K. Dopfer (ed.) *Evolutionary Economics: Program and Scope*. Kluwer Academic Publishers, Boston, Dordrecht, London, Chap. 1, pp. 1-44.
8. Ebersberger, B. (2002) *Genetische Programmierung: Ein Instrument zur empirischen Fundierung ökonomischer Modelle*. DUV, Wiesbaden.

9. Eliasson, G. (1988) Schumpeterian Innovation, Market Structure and the Stability of Industrial Development. In: H. Hanusch (ed.) *Evolutionary Economics – Applications of Schumpeterian Ideas*. Cambridge, Cambridge University Press.
10. Eliasson, G. and Carlsson, B. (2004) Industrial Dynamics and Endogenous Growth. *Industry and Innovation*, Vol. 10, Number 4, December 2003, pp. 453-455.
11. Evans, E. S. and Jovanovic, B. (1989) An estimated model entrepreneurial choice under liquidity constraints. *Journal of Political Economy* (97), 808-827.
12. Gilbert, N. and Troitzsch, K. (1999) *Simulation for the Social Scientist*. Open University Press, Milton Keynes.
13. Goebel, P. (1990) *Erfolgreiche Jungunternehmer*. mvg – Moderne Verlagsgesellschaft, München.
14. Grebel, T., Pyka, A. and Hanusch, H. (2004) An Evolutionary Approach to the Theory of Entrepreneurship. *Industry and Innovation*, Vol. 10, Number 4, December 2003, pp. 493-514.
15. Grebel, T. (2004) *Entrepreneurship – A New Perspective*. Routledge, Cheltenham.
16. Heiner, R. A. (1983) The origin of predictable behaviour. *American Economic Review* 73, 560-595.
17. Hermann-Pillath, C. (2001) On the ontological foundation of evolutionary economics. In: K. Dopfer (ed.) *Evolutionary Economics: Program and Scope*, Kluwer Academic Publishers, Boston, Dordrecht, London, Chap. 1, pp. 89-139.
18. Hunsdiek, D. and May-Strobl, E. (1986) *Entwicklungslinien und Entwicklungsrisiken neugegründeter Unternehmen*. Poeschel, Stuttgart.
19. Kwasnicki, W. (1998) Simulation methodology in evolutionary economics. In: F. Schweitzer and G. Silverberg (eds) *Evolution and Selbstorganisation in der Ökonomie*, Vol. 9 *Selbstorganisation: Jahrbuch für Komplexität*, Duncker und Humblot, Berlin, pp. 161-186.
20. Lane, D. (1993) Artificial worlds and economics, part I, *Journal of Evolutionary Economics* 3, 89-107.
21. Lane, D. (1993) Artificial worlds and economics, part II, *Journal of Evolutionary Economics* 3, 177-197.
22. McClelland, D. C. (1961) *The Achieving Society*, Van Nostrand Reinhold, New York.
23. Metcalfe, J. S. (2001) Evolutionary approaches to population thinking and the problem of growth and development. In: K. Dopfer (ed.) *Evolutionary Economics: Program and Scope*. Kluwer Academic Publishers, Boston, Dordrecht, London, pp. 141-164.
24. Nelson, R. R. and Winter, S. G. (1982) *An Evolutionary Theory of Economic Change*. Cambridge University Press, Cambridge, MA.
25. Nelson, R. R. (2001) Evolutionary perspectives on economic growth. In: K. Dopfer (ed.) *Evolutionary Economics: Program and Scope*. Kluwer Academic Publishers, Boston, Dordrecht, London, Chap. 1, pp. 165-194.
26. Romer, P. M. (1990) Endogenous Technological Change, *Journal of Political Economy*, Vol.98, pp. 77-102.
27. Saviotti, P. P. (1996) *Technological Evolution, Variety and the Economy*. Edward Elgar, Aldershot.
28. Scherer, F.M., Ross, D. (1990) *Industrial Market Structure and Economic Performance*. Houghton and Mifflin Company, Boston.

29. Schumpeter, J. A. (1912) *Theorie der wirtschaftlichen Entwicklung*. Duncker Humblot, Berlin.
30. Schumpeter, J. A. (1939) *Business Cycles*. McGraw-Hill, New York.
31. Silverberg, G. (1988) Modelling economic dynamics and technical change: Mathematical approaches to self-organization and evolution. In: G. Dosi et al. (eds) *Technical Change and Economic Theory*. Pinter Publishers, London, New York.
32. Tesfatsion, L. (2001) Agent-based modelling of evolutionary economic systems. *IEEE Transactions on Evolutionary Computation* 5, 1-6.
33. Wooldridge, M. and Jennings, N.R. (1995) Intelligent agents: theory and practice. *Knowledge Engineering Review*, 10: 115-152.

On the Analysis of Asymmetric Directed Communication Structures in Electronic Election Markets

Markus Franke, Andreas Geyer-Schulz and Bettina Hoser

Information Services and Electronic Markets, University of Karlsruhe, Germany

Markus.Franke@em.uni-karlsruhe.de

Andreas.Geyer-Schulz@em.uni-karlsruhe.de

Bettina.Hoser@em.uni-karlsruhe.de

Summary. In this article we introduce a new general method of representing trading structures as complex adjacency matrices and transforming these into Hermitian adjacency matrices which are linear self-adjoint operators in a Hilbert space. The main advantages of the method are that no information is lost, no arbitrary decision on metrics is involved, and that all eigenvalues are real and, therefore, easily interpretable. The analysis of the resulting eigensystem helps in the detection of substructures and general patterns. While this approach is general, we apply the method in the context of analyzing market structure and behaviour based on the eigensystem of market transaction data and we demonstrate the method by analyzing the results of a political stock exchange for the 2002 federal elections in Germany.

1 Introduction

The analysis of directed asymmetric communication structures has its origins in military signal intelligence in the 1st and 2nd World War. With the spread of information and communication technology in almost every aspect of everyday life and with today's trends towards embedded, wearable, and networked devices we expect an ever increasing stream of communication and traffic data. Wellman's vision of the rise of network societies requires new tools for enhancing social capital, tools from social network analysis which analyze these data streams (see e.g. [39], [38]). H. A. Simons's early investigation of the behavioural aspects of organizations show the potential of analyzing organizational communication streams [32]. In information retrieval Google's page rank index exploits the "frozen" asymmetric communication between documents in the form of directed hyperlinks [26]. For a more detailed analysis see [20]. Another application area – this time on the level of the technical infrastructure – is the surveillance of distributed systems by means of analyzing

their communication streams. However, in this contribution we concentrate on the analysis of the communication streams in electronic double-auction markets.

This contribution is structured in two parts: In Sect. 2 we present the notation and the theoretical foundations for the eigensystem analysis of asymmetric directed communication streams and in Sect. 3 we apply the method to the analysis of market structure and behaviour in an experimental forecasting market for the 2002 federal elections of Germany which was organized by the department of Information Services and Electronic Markets.

2 Hilbert Space, Hermitian Matrices, and Asymmetric Communication Streams

To make our presentation self-contained we introduce in the following both notation and basic definitions that we need. By z we denote a complex number in algebraic form or in exponential form $z = a + ib = |z|e^{i\phi}$. $Re(z) = a$ denotes the real part of z , $Im(z) = b$ the imaginary part. The absolute value of z is $|z| = \sqrt{a^2 + b^2}$, and the phase $\phi = \arccos \frac{Re(z)}{|z|}$, $0 \leq \phi \leq \pi$, with i the imaginary unit ($i^2 = -1$). \bar{z} denotes the complex conjugate $\bar{z} = a - ib$ of z . The following rules serve as a reminder how to compute with complex numbers:

$$z_1 z_2 = |z_1| |z_2| e^{i(\phi_1 + \phi_2)} \quad (1)$$

$$z + \bar{z} = 2Re(z) \quad (2)$$

$$z = \bar{z} \quad \text{if and only if} \quad z \in \mathbb{R} \quad (3)$$

$$z\bar{z} = |z|^2 \quad (4)$$

2.1 Hilbert Space

We use the following notational conventions: Unless otherwise stated all numbers are complex ($\in \mathbb{C}$). A column vector is printed in bold face \mathbf{x} , its components are $x_j, j = 1 \dots n$. The vector space is defined by $\mathbf{V} = \mathbb{C}^n$. Matrices are denoted as capital letters A . a_{kl} represents the entry in the k -th row and the l -th column. Greek letters denote eigenvalues. λ_k represents the k -th eigenvalue. The complex conjugate transpose of a vector \mathbf{x} is defined as \mathbf{x}^* . The transpose of a vector \mathbf{x} is \mathbf{x}^t . The outer product of two vectors \mathbf{x} and \mathbf{y} is defined as:

$$\mathbf{xy}^* = \begin{pmatrix} x_1 \bar{y}_1 & \dots & x_1 \bar{y}_n \\ \dots & \dots & \dots \\ x_n \bar{y}_1 & \dots & x_n \bar{y}_n \end{pmatrix} \quad (5)$$

We represent the inner product of \mathbf{x} and \mathbf{y} which is a semilinear form on a given vector space \mathbf{V} as:

$$\langle \mathbf{x} | \mathbf{y} \rangle = \mathbf{x}^* \mathbf{y} = \sum_{k=1}^n \overline{x_k} y_k \tag{6}$$

For the vector space \mathbf{V} the following rules hold:

$$\langle \mathbf{x} | \mathbf{x} \rangle \geq 0 \quad \text{with } \langle \mathbf{x} | \mathbf{x} \rangle = 0 \text{ if and only if } \mathbf{x} = \mathbf{0} \tag{7}$$

$$\langle a\mathbf{x} | \mathbf{y} \rangle = \overline{a} \langle \mathbf{x} | \mathbf{y} \rangle; \quad \langle \mathbf{x} | a\mathbf{y} \rangle = a \langle \mathbf{x} | \mathbf{y} \rangle \tag{8}$$

$$\langle \mathbf{x} + \mathbf{y} | \mathbf{z} \rangle = \langle \mathbf{x} | \mathbf{z} \rangle + \langle \mathbf{y} | \mathbf{z} \rangle \tag{9}$$

$$\langle \mathbf{x} | \mathbf{y} \rangle = \overline{\langle \mathbf{y} | \mathbf{x} \rangle} \tag{10}$$

The norm, denoted by $\| \mathbf{x} \|$, is defined as follows:

$$\sqrt{\langle \mathbf{x} | \mathbf{x} \rangle} = \| \mathbf{x} \| \tag{11}$$

Note that the distance between two vectors \mathbf{x} and \mathbf{y} is defined by $\| \mathbf{x} - \mathbf{y} \|$. For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ the distance reduces to the familiar Euclidean norm $\sqrt{\sum_i^n (x_i - y_i)^2}$.

A Hilbert space is a complete normed inner product space as defined by Eqs.(6) - (11). We only consider separable Hilbert spaces in this contribution. In addition, for every $n \in \mathbb{N}$, a set of linearly independent elements $\mathbf{x}_1, \dots, \mathbf{x}_n$ exists such that the equation $a_1 \mathbf{x}_1 + \dots + a_n \mathbf{x}_n = \mathbf{0}$ holds only when $a_1 = \dots = a_n = 0$. (See e.g. M. H. Stone [37].)

The Cauchy-Schwartz inequality

$$|\langle \mathbf{x} | \mathbf{y} \rangle| \leq \sqrt{\langle \mathbf{x} | \mathbf{x} \rangle} \sqrt{\langle \mathbf{y} | \mathbf{y} \rangle} \tag{12}$$

(With equality if $\mathbf{x} = a\mathbf{y}$, $a \in \mathbb{C}$.)

is the basis to the Bessel inequality. For any vector \mathbf{h} and an orthonormal basis \mathbf{x}_j

$$\sum_{j=1}^n |\langle \mathbf{h} | \mathbf{x}_j \rangle|^2 \leq \| \mathbf{h} \|^2 \tag{13}$$

(Bessel inequality.)

For any vector h and a complete orthonormal basis \mathbf{x}_k the Bessel inequality becomes Parseval's equality:

$$\sum_{j=1}^n |\langle \mathbf{h} | \mathbf{x}_j \rangle|^2 = \| \mathbf{h} \|^2 \tag{14}$$

2.2 Hermitian Matrices

The adjoint space \mathbf{X}^* of \mathbf{X} is the set of all semilinear forms (Eq. (8)) (e.g. inner product) on vector space \mathbf{X} [19, p. 11]. A selfadjoint linear operator is called Hermitian. A matrix H is called Hermitian, if and only if

$$H^* = H \quad (15)$$

with H^* representing the conjugate complex transpose of H . This means that the matrix entries can be written $h_{lk} = \overline{h_{kl}}$. Hermitian matrices are also normal:

$$HH^* = H^*H \quad (16)$$

The eigenvalue equation

$$H\mathbf{x} = \lambda\mathbf{x} \quad (17)$$

of a complex Hermitian matrix H can be represented due to its complete orthonormal eigenvector system (which means of full rank) in the Fourier sum representation:

$$H = \sum_{k=1}^n \lambda_k P_k; \quad P_k = \mathbf{x}_k \mathbf{x}_k^* \quad (18)$$

with λ_k the k -th eigenvalue, \mathbf{x}_k the k -th eigenvector, and P_k the orthogonal projectors. Note that $\sum_{k=1}^n P_k = I$, $P_k^* = P_k$, $P_k^2 = P_k$. The set of all eigenvalues is called spectrum.

Any orthogonal basis can be chosen such that:

$$\langle \mathbf{x}_k | \mathbf{x}_l \rangle = \delta_{kl} \quad \text{follows from (15)} \quad (19)$$

with

$$\delta_{kl} = \begin{cases} 0 & \text{if } k \neq l, \\ 1 & \text{if } k = l \end{cases} \quad (20)$$

This also holds true for arbitrary rotation:

$$\langle e^{i\phi_k} \mathbf{x}_k | e^{i\phi_l} \mathbf{x}_l \rangle = e^{-i\phi_k} e^{i\phi_l} \langle \mathbf{x}_k | \mathbf{x}_l \rangle = e^{i(\phi_l - \phi_k)} \langle \mathbf{x}_k | \mathbf{x}_l \rangle = e^{i(\phi_l - \phi_k)} \delta_{kl} \quad (21)$$

Hermitian matrices thus have full rank and, therefore all eigenvalues are real:

$$\lambda_k \in \mathbb{R} \quad \forall k \quad (22)$$

because $\langle H\mathbf{x} | \mathbf{x} \rangle = \langle \lambda\mathbf{x} | \mathbf{x} \rangle = \overline{\lambda} \langle \mathbf{x} | \mathbf{x} \rangle$ and $\langle \mathbf{x} | H\mathbf{x} \rangle = \langle \mathbf{x} | \lambda\mathbf{x} \rangle = \lambda \langle \mathbf{x} | \mathbf{x} \rangle$ and $H^* = H$ imply $\langle H\mathbf{x} | \mathbf{x} \rangle = \langle \mathbf{x} | H\mathbf{x} \rangle$ and thus $\overline{\lambda} = \lambda$ which means $\lambda \in \mathbb{R}$ (see [23, p. 548], [19, p. 53]).

Since all eigenvalues of a Hermitian matrix are real (Eq.(22)) the interpretation of the eigenvalues does not pose the difficulty of interpreting complex eigenvalues of non-symmetric real matrices.

For a complex Hermitian matrix with $\text{tr}(H) = 0$ some eigenvalues have to be negative due to the fact that

$$\text{tr}(H) = \sum_{k=1}^n h_{kk} = \sum_{k=1}^n \lambda_k = 0 \quad (23)$$

As a special case consider a matrix B of order $l = n + m$ with

$$B = \begin{pmatrix} 0_{n \times n} & A \\ A^* & 0_{m \times m} \end{pmatrix} \quad (24)$$

with A representing a n by m matrix of rank r . In the special case of $n = 1$ and $m = l - 1$ this matrix represents a directed, weighted star graph. The spectrum of that system is described by the following non-zero eigenvalues:

$$\sigma(B) = \{+\lambda_1, -\lambda_1, \dots, +\lambda_r, -\lambda_r\} \quad (25)$$

(This follows from

$$\begin{pmatrix} 0 & A \\ A^* & 0 \end{pmatrix} \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix} = \lambda \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix} \Rightarrow \left\{ \begin{array}{l} A\mathbf{x}_2 = \lambda\mathbf{x}_1 \\ A^*\mathbf{x}_1 = \lambda\mathbf{x}_2 \end{array} \right\} \Rightarrow A^*A\mathbf{x}_2 = \lambda^2\mathbf{x}_2,$$

see [23, p. 555]). For the special case of a weighted star ($n = 1$), the non-zero eigenvalues are $\sigma(B) = \{+\lambda_1, -\lambda_1\}$.

Furthermore, the eigenvalues of any matrix A depend continuously upon its entries a_{ij} , because the zeroes of a polynomial depend continuously on the coefficients of the polynomial. This property is made precise by the following theorem (see [15, p. 539-540]):

Theorem 1. *Let $n \geq 1$ and $p(x) = \sum_{i=0}^n a_i x^i$ is a polynomial with $a_n \neq 0$. Then for every $\epsilon > 0$, a $\delta > 0$ exists, so that for any polynomial $q(x) = \sum_{i=0}^n b_i x^i$ with $b_n \neq 0$ and $\max_{0 \leq i \leq n} |a_i - b_i| < \delta$ we have*

$$\min_{i \in \tau} \max_{1 \leq j \leq n} |\lambda_j - \mu_{\tau(i)}| < \epsilon,$$

where λ and μ are the zeros of $p(x)$ and $q(x)$ in some order and the minimum is taken over all permutations τ of $1, \dots, n$.

In addition, perturbation bounds for invariant subspaces of a Hermitian matrix A and its Hermitian perturbation $A + E$ can be found in [18] or [10].

2.3 Coding Communication Streams as Hermitian Matrices

We consider communication streams defined as a weighted directed graph $G = (N, E)$ with N denoting the set of nodes of the graph, $E \in N \times N \times \mathbb{R}$ the set of edges weighted with the number of messages m from node k to node l and $k \neq l$ holds. A complex adjacency matrix H from such a graph G is constructed as follows:

1. We construct a square complex adjacency matrix A with n members by

$$a_{kl} = m + ip \quad (26)$$

with m the number of outbound messages from node k to node l , and p the number of inbound messages from node l to node k , and i representing the imaginary unit. As can be seen $a_{kl} = i\overline{a_{lk}}$.

2. We rotate A by multiplying A with $e^{-i\frac{\pi}{4}}$ (complex multiplication is equivalent to rotation, see Eq. (1)) in order to obtain a Hermitian matrix H :

$$H = A \cdot e^{-i\frac{\pi}{4}} \quad (27)$$

Proof:

- a) $a_{kl} = re^{i\phi}$
 - b) $a_{lk} = ire^{-i\phi} = i\overline{a_{kl}}$
 - c) $a_{kk} = 0$ because of the exclusion of self references
 - d) $a_{kl_r} = re^{i\phi}e^{i\psi} = re^{i(\phi+\psi)}$
 - e) $a_{lk_r} = ire^{-i\phi}e^{i\psi} = e^{i\frac{\pi}{2}}re^{i(\psi-\phi)} = re^{i(\frac{\pi}{2}+\psi-\phi)}$
 - f) For Hermitian matrices $a_{kl_r} = \overline{a_{lk_r}}$ must hold, therefore, $re^{i(\phi+\psi)} = re^{-i(\frac{\pi}{2}+\psi-\phi)}$.
 - g) This holds, if $\phi + \psi = -\frac{\pi}{2} - \psi + \phi$.
 - h) Solving for ψ leads to $2\psi = -\frac{\pi}{2}$. Thus $\psi = -\frac{\pi}{4}$. qed
3. Under this rotation the coordinate independent characteristics of the original communication patterns is kept [23, p.256], no information is lost.

Furthermore, since H is a Hermitian matrix (and Hermitian matrices are of full rank) and rotation is rank preserving, this implies that A is also of full rank.

Table 1 shows how for four different types of communication behaviour the characteristic is transferred after the rotation described above. Note that each type remains identifiable.

Table 1. Communication behaviour representation

communication behaviour	$a_{kl} = m + ip$	$h_{kl} = m_r + ip_r$
no self reference	$a_{kk} = 0$	$h_{kk} = 0$
$k \rightarrow l > l \rightarrow k$	$m > p$	$p_r > 0$
$l \rightarrow k > k \rightarrow l$	$m < p$	$p_r < 0$
$k \rightarrow l = l \rightarrow k$	$m = p$	$p_r = 0$

3 Analyzing Market Structure and Behaviour in an Experimental Forecasting Market

Recently researchers investigated and recommended virtual stock exchanges as marketing research tools comparable to opinion polls (e.g. [34], [35]). These

forecasting markets have been pioneered by Forsythe, Nelson, Neumann and Wright with the Iowa Election Market which successfully predicted the US presidential elections in 1988 and 1992 and outperformed the polls [9]. For a survey on German election markets in 1990, 1991, 1994, and 1998 as well as on other European election markets organized during the last decade see [3] who show that these markets in general did not as well in predicting the outcome of elections as their US counterparts. While a number of explanations have been put forward (marginal trader and judgement bias [9], political insiders [3], ...) we tend towards a different explanation, namely the possible breakdown of Vernon Smith's induced value theory when moving election markets from a university environment into the field [33].

We illustrate this hypotheses with anecdotal evidence from the political stock markets run during the Austrian presidential elections in 1998. In the campaign for this election Richard "Mörtl" Lugner ran as outsider against Thomas Klestil. Election markets were fashionable in Austria at this time, several newspapers (Kurier, Salzburger Nachrichten) and two universities (TU-Wien, WU-Wien) organized election markets – each of a different kind. The press election markets received high media coverage, whereas the smallest market, namely the market of the WU-Wien, was more of an internal market (and a software performance test) for an interdisciplinary seminar organized by a research group of the second author. Surprisingly, on the market of the Salzburger Nachrichten Richard Lugner was traded at approximately 20 percent of the votes triggering a large number of stories covering his life, his life-style and his election program in the press as well as numerous appearances at talk shows.

Practically overnight, Richard Lugner was Austria's political shooting star. Even more surprisingly, only the small, "inefficient", and unknown market of the WU-Wien performed better than the opinion polls. Rumours on money losses in the order of 20 000 Euro were reported to the second author and discussed within the Austrian election market group. However, due to lack of hard data this incident remained unpublished. Shortly after the election, mass media researchers estimated the market value of Richard Lugner for advertising at 1.5 million Euro. And from this anecdote we conclude that it may be rational to loose money in one market, if and only if we can offset these losses in another market. Clearly, running an election market which serves as an information channel for the election proper (another market with higher rewards) implies loss of experimental control, because it seems almost impossible to establish a proper incentive structure for this kind of game in non-isolated markets.

As an alternative to incentive compatibility several researchers (e.g. [9], [3], [34], and [35]) propose explanation models for the prediction error based on explanatory variables not observable in the market as e.g. self-reports and surveys of traders (e.g. demographics [3] and expectations of election results [25]) or on explanatory variables based on measures of market behaviour as e.g. the volatility of the market on the last trading day (e.g. [34]). For a critic of

the first approach we refer the reader to Michael Berlemann [1] who identified untestable hypotheses, institutional differences, failing incentive systems, and difficulties in classifying trading behaviour as the main problems. The choice of the definition of volatility as the coefficient of the variation of price changes $\frac{(\sqrt{\sum_{i=1}^n \frac{(d_i - \bar{d})^2}{n}})}{\bar{d}}$ with p_i the i -th price, $d_i = p_i - p_{i-1}$, and $\bar{d} = \frac{1}{n} \sum_{i=1}^n d_i$ in the second approach seems to be arbitrary. Other choices for volatility measures (e.g. the semi-variance) can be found in [27]. Further options are measures as the bid-ask spread or Roll's effective bid-ask spread [28] or Gomber's market impact as a measure of liquidity [12]. In addition, in an experimental setting all these measures are based only on a fraction of the available information in the market transaction data.

The rest of this section is structured as follows: First, we discuss different aspects of market and information efficiency and their connection to forecasting markets, namely that efficient markets have a high forecasting accuracy. Since market efficiency is essential for forecasting accuracy, researchers in forecasting markets try to induce market efficiency either by enforcing incentive compatibility or by explanation models on the prediction error. In the following we propose an alternative to achieve this goal: To analyze the eigensystem of market transactions and to compare the eigensystem of an actual market with the eigensystem of a monopoly/monopsony and with the eigensystem of an (efficient) competitive market. In order to do this, we first show that the transactions recorded in the accounting system of a market can be interpreted as a weighted asymmetric communication structure. Next, we present the eigensystem of a typical monopoly/monopsony and the eigensystem of an (efficient) competitive market. Finally, we analyze a stylized fact about the market structure and behaviour in an experimental forecasting market for the 2002 federal elections in Germany.

3.1 Market Efficiency and Forecasting Markets

Let us recall that the idea of market efficiency has a long tradition in economics with a first peak of interest in the comparison of socialist/communist centralized planning economies with decentralized market economies after the Russian revolution (see [17] for a survey). Hayek's main objections against the feasibility of centralized planning economies are based on the economics of information processing: First, the very large number of variables and equations for optimizing a centralized economy and the too large computational complexity of solving this problem ("a task which, with any means known at the present, could not be carried out in a lifetime" [13, p. 212]), and, second, the difficulty of placing all the relevant information at the disposal of a single central agency, because of the dispersal of information throughout the society [14]. Under the (natural) assumption that each economic unit has only information about itself, Hayek concluded that information efficiency requires a minimization of information transfers between the economic units as

well as a decentralized solution of the optimization problem in such a way that the information transfer between the economic units is restricted to the price vector (without the numeraire) and that each economic unit solves its own optimization problem (as a price taker). Hayek claimed that decentralized market-based economies are exactly such informationally decentralized procedures.

A formalization of Hayek's result due to an idea of K. Arrow can be found in [24]: It can be shown that in an economy with private property, with rational companies which know the prices and their own production technology and maximize their profits, with rational consumers which know the prices and their own preferences and which maximize their own utility, and with prices which balance supply and demand the allocation of resources is efficient and the incentive problem is solved because each party follows only its own interests. Markets are allocationally efficient if prices equate the (risk adjusted) marginal return of all parties. Note, however, that this formalization at the same time allows an analysis of various market failures, either caused by external effects or the interplay of information decentralization and incentive compatibility (see [17]). The last problem is fundamental and of high relevance for forecasting markets.

Efficient (capital) markets are more realistic than perfect markets. Perfect capital markets have the following properties [5, p. 331]:

1. No friction. In perfect markets no transaction costs and taxes exist (by definition operationally efficient). All assets are perfectly divisible. No regulations constrain the market.
2. Perfect competition in both product and securities markets. Producers supply goods and services at minimum average cost, in security markets all participants act as price takers.
3. Informationally efficient means information is costless and received simultaneously by all parties.
4. All parties are rational expected utility maximizers.

Compare these conditions with capital market efficiency: "... in an efficient capital market, prices fully and instantaneously reflect all available information" [5, p. 331]. Prices are accurate signals for capital allocation. And this remains true, even if we relax some of the perfect market conditions. We can still have efficient markets with taxes, brokerage fees, costly information, imperfect competition in product markets, etc.

The operationalization of this definition in the form of three different and testable versions of the efficient capital market hypotheses is due to Fama (see [6], [7], and [8]). The weak-form efficiency says that the information in past prices is not relevant in earning excess returns. The semi-strong form efficiency claims that no excess returns can be earned from any publicly available information. The strong-form efficiency stipulates that no excess returns can be earned using any information. Rubinstein [29] and Latham [21] have given an even stronger definition of capital market efficiency: A market is efficient

with regard to an information event, if that event causes no portfolio changes - that is no trade occurs.

Formally, in an efficient market (Fama [7])

$$f(\mathbf{p}_t | \Phi_{t-i}) = f_m(\mathbf{p}_t | \Phi_{t-1}^m), \quad (28)$$

where \mathbf{p}_t denotes the security price vector at time t , Φ_{t-1} is the available information set at t and Φ_{t-1}^m is the information set used by the market, $f_m(\mathbf{p}_t | \Phi_{t-1}^m)$ is the density function estimated by the market for \mathbf{p}_t , and $f(\mathbf{p}_t | \Phi_{t-i})$ is the true density function of \mathbf{p}_t . To make Eq. (28) testable a link between $f_m(\mathbf{p}_t | \Phi_{t-1}^m)$ and \mathbf{p}_{t-1} must be established. Usually, the assumption that the conditions of market equilibrium can be expressed in terms of expected returns serves this purpose:

$$p_{j,t-1} = \frac{E_m(\bar{p}_{j,t} | \Phi_{t-1}^m)}{1 + E_m(\bar{r}_{j,t} | \Phi_{t-1}^m)} \quad (29)$$

where $E_m(\bar{r}_{j,t} | \Phi_{t-1}^m)$ is the expected return of security j in market equilibrium implied by $f_m(\mathbf{p}_t | \Phi_{t-1}^m)$ and $E_m(\bar{p}_{j,t} | \Phi_{t-1}^m)$ is the expected value of the price of security j at time t of the market. $E(\bar{p}_{j,t} | \Phi_{t-1})$ denotes the true expected price of security j at time t implied by $f(\mathbf{p}_t | \Phi_{t-1})$ and $E(\bar{r}_{j,t} | \Phi_{t-1})$ the true expected return implied by $E(\bar{p}_{j,t} | \Phi_{t-1})$ and $f(\mathbf{p}_t | \Phi_{t-1})$. If Eq. (28) holds, we have $E(\bar{p}_{j,t} | \Phi_{t-1}) = E_m(\bar{p}_{j,t} | \Phi_{t-1}^m)$ and $E(\bar{r}_{j,t} | \Phi_{t-1}) = E_m(\bar{r}_{j,t} | \Phi_{t-1}^m)$. Thus in an efficient market, the true expected return equals its expected value in market equilibrium and the expected prices in market equilibrium equal the true prices - and, therefore, markets can be used for forecasting prices as in election markets. The forecasting quality of a market is directly related to market efficiency.

However, in capital markets tests on market efficiency always require the joint specification of a market equilibrium model (that is, how $E_m(\bar{r}_{j,t} | \Phi_{t-1}^m)$ is related to $f_m(\mathbf{p}_t | \Phi_{t-1}^m)$) but the conditions of market efficiency do not restrict the choice of a market equilibrium model. (For a survey on empirically testing the various asset pricing models, see [8]).

Fortunately, in a forecasting market we can dispense with the problem of jointly specifying and testing a market equilibrium model. At least ex-post (after the event at time T) the true prices are known and they can be compared with the market prices (by the Euclidean distance of \mathbf{p}_T (the observed price vector) and $\bar{\mathbf{p}}_t$ (the market equilibrium price vector)). However, in a forecasting market, the interesting question ex-ante is if $\bar{\mathbf{p}}_t$ is a good predictor of \mathbf{p}_T ? Again, this is related to market efficiency and to incentive compatibility, as the anecdote on Richard Lugner demonstrates. In financial markets behaviourists like Andrei Shleifer argue that market efficiency is caused by rational traders, by independent deviations from rationality, and by arbitrage [36]. Moreover, they argue that even a single condition is enough. For a forecasting market, that is, unfortunately, only true, if incentive compatibility is assured.

Nevertheless, we take in the following a behaviouristic position. To support this, we repeat here A. Samuelson’s argumentation that observed choice acts reveal consumer preferences [30]. In our own forecasting markets we can completely observe trading and bidding behaviour and we turn the questions about market efficiency and incentive compatibility into questions about the observed trading and bidding behaviour in a market accounting system. Thus, in our setting more variables of the trading process are observable and potentially taken into account. More specifically, in this contribution we concentrate in the following on identifying trading patterns consistent with perfect competition as opposed to trading patterns indicating a monopolistic/monopsonistic behaviour in the eigensystem of the market accounting system in an attempt to link monopolistic/monopsonistic behaviour with less efficient markets and larger forecasting errors. At the moment we can do this only qualitatively. We are well aware that this constitutes only a first small step towards a behavioural foundation for market efficiency.

3.2 Financial Accounting and Asymmetric Communication Structures

Next, consider the small example of financial accounting shown in Fig. 1 with eight transactions of the form trader k sells to trader l shares and receives an amount of money s . The T-accounts of all traders involved are shown with the 8 trades recorded according to the conventions of financial accounting. (For a short introduction to financial accounting see [11]). The journal entries of the records numbered (1) to (8) are shown next to the T-accounts. The graph in Fig. 2 shows the monetary flow between the traders. That is the edge is weighted by the money flow from trader k to trader l . In the example, the money flow on the edges of the graph in the observation period is generated by a single transaction. In general and in the rest of this paper, the edge weight is the sum of all monetary flows s from trader k to trader l . And we see that the monetary flow between accounts can be represented as an asymmetric directed communication structure. Note that the weighting function of the transaction stream can be chosen as required by the intended application.

	Trader 1	Trader 3	
(1) Record trader 1 trader 2, 9 Euro	(1) 9	(3) 4	
(2) Record trader 1 trader 3, 10 Euro	(2) 10	(3) 4 (2) 10	
(3) Record trader 3 trader 1, 4 Euro	(5) 8	(7) 5	
(4) Record trader 4 trader 1, 7 Euro	(6) 9	(8) 2	
(5) Record trader 1 trader 4, 8 Euro	(8) 2	(4) 7 (5) 8	
(6) Record trader 1 trader 5, 9 Euro	Trader 2	Trader 5	
(7) Record trader 5 trader 1, 5 Euro	(8) 2	(7) 5 (6) 9	
(8) Record trader 2 trader 1, 2 Euro	(1) 9		

Fig. 1. 5 Accounts and 8 records

3.3 Star Graphs and their Interpretation

As a basic construct consider a directed and weighted star graph with 5 members as in Fig. 2.




Fig. 2. Unperturbed star graph: A monopoly/monopsony or a market place with almost perfect competition?

The complex adjacency matrix belonging to this graph is:

$$\begin{pmatrix} 0 & 2 + 9i & 4 + 10i & 7 + 8i & 5 + 9i \\ 9 + 2i & 0 & 0 & 0 & 0 \\ 10 + 4i & 0 & 0 & 0 & 0 \\ 8 + 7i & 0 & 0 & 0 & 0 \\ 9 + 5i & 0 & 0 & 0 & 0 \end{pmatrix} \tag{30}$$

Which after rotation becomes the Hermitian matrix:

$$\begin{pmatrix} 0. & 7.7 + 4.9i & 9.8 + 4.2i & 10.5 + 0.7i & 9.8 + 2.8i \\ 7.7 - 4.9i & 0. & 0. & 0. & 0. \\ 9.8 - 4.2i & 0. & 0. & 0. & 0. \\ 10.5 - 0.7i & 0. & 0. & 0. & 0. \\ 9.8 - 2.8 & 0. & 0. & 0. & 0. \end{pmatrix} \tag{31}$$

As can be seen in Table 2 there are two non-zero eigenvalues of the same absolute value but different sign. As was shown in Eq. (25) this is indeed the characteristic of a star graph complex Hermitian adjacency matrix. As can also be seen the eigenvectors belonging to the two eigenvalues are the same in absolute values but differ by π in phase. Note that the trader with ID 1 is the center of the star graph and is indicated as such by the highest absolute value of the eigenvector component. This result and all other numerical results in the following are determined only up to an arbitrary rotation. For the rest of the paper this rotation is always chosen in such a way that the eigenvector element with the largest absolute value (in polar coordinates) is real (has a phase of 0). We have used Mathematica 4.0 to calculate our results [40]. E.g. for a perturbed star we refer the reader to [16].

Interpreting this graph as a transaction stream in the accounting system of a market requires knowledge about the market organisation. Two interpretations corresponding to two types of market organization, namely a market

with bilateral trades between parties and a market with a market mechanism as formal counterparty to all trades are discussed: The account with ID 1 belongs to a trader as in the example shown in Fig. 1 or the account with ID 1 belongs to a formal counterparty for all trades (“the market mechanism”). In the case study in Sect. 3.5 the market mechanism is that of a continuous double auction market.

In the first interpretation, we see that the trader with ID 1 is the only counterparty for all trades. Depending on the strength of the asymmetry of the flows in the graph the trader with ID 1 can be identified either as a monopolist (only seller, more inflows of money), a monopsonist (only buyer, more outflows of money), or as a market maker (liquidity provider, symmetric flow of money). The eigensystem of a completely balanced market maker has a real solution of the form characteristic for a star with all trader accounts having a phase shift of π (as theoretically expected). The eigensystems of a monopolist/monopsonist are conjugate complex and exhibit the typical star pattern which we expect from theory. A monopolist has an inbound star pattern, a monopsonist an outbound star pattern. In our randomly chosen example, the trader with ID 1 is slightly on the monopolistic side. We see this from the phase distribution in the first two eigenvectors in Table 2.

However, in the second interpretation the account with ID 1 belongs to the market mechanism. In this case market rules require that the market mechanism is a formal counterparty to all traders and so a star is naturally formed. The absolute values of the eigenvector elements for traders 2, 3, 4 and 5 are almost uniformly distributed, although with different phase. Under this interpretation this would indicate almost perfect competition of the traders 2, 3, 4, and 5. For a double auction mechanism, the account of the formal counterparty must always be balanced. And it is this second interpretation which we use in the example. This is a consequence of the bipartite graph-structure of a continuous double auction market which is shown in Fig. 6. Moreover, Eq. (25) defines the spectrum of a continuous double auction market with a market accounting system with the structure shown in Fig. 6.

Table 2. Eigensystem and its eigenvectors for MStar5 in Eq. (31) with $z = |z|e^{i\phi}$

λ_k		\mathbf{x}_{kl}				
1.	abs(z)	0.71	0.32	0.37	0.37	0.35
	arg(z)	0	-0.57	-0.40	-0.07	-0.28
-1.	abs(z)	0.71	0.32	0.37	0.37	0.35
	arg(z)	0	2.57	2.74	3.07	2.86

Suppose now, that trader 3 in the example has considerable market power. He was the counterparty for the trades of traders 2, 4, and 5 via the double auction mechanism with ID 1. The resulting graph is shown in Fig. 3.




Fig. 3. Monopolistic/monopsonistic behaviour in a continuous double auction: A star graph

Table 3. Eigensystem and its eigenvectors for graph in Fig. 3 with $z = |z|e^{i\phi}$

λ_k		\mathbf{x}_{kl}				
1.	abs(z)	0.71	0.19	0.61	0.22	0.21
	arg(z)	0	-0.57	0.29	-0.07	-0.28
-1.	abs(z)	0.71	0.19	0.61	0.22	0.21
	arg(z)	0	2.57	-2.85	3.07	2.86

Next, we compare the eigensystem in Table 2 with the eigensystem in Table 3. We see that both eigensystems exhibit the star structure. However, in Table 3 we see that trader 3 as sole counterparty to all trades has the second largest absolute value of the eigenvector with 0.61, this is considerably larger than 0.37 in Table 2. Since trader 3 has received a net money inflow of 12 (he has sold shares for 26 and bought shares for 14), he has a positive phase of 0.29 because imbalances in money flow show in the phase. The strongest eigenvector element corresponds – as expected to the market mechanism with account 1. By definition, a market mechanism must be balanced and has, therefore, a phase of 0. The other three traders retained their phase and have approximately the same absolute value of their eigenvector elements.

The following idea of measuring competitiveness in a market is derived from Martin Shubik's analysis of symmetric oligopolies and his discussion of perfect competition and monopoly in a real world context [31]. In perfect competition, each trader has the same market power which is reflected in equal absolute values of the eigenvector elements. In order to compare different markets with respect to their market power, we introduce a uniform distribution \mathbf{u} with $u_i = \frac{1}{n}$, with n the number of active traders in the market. For each real market, we compute a normed vector \mathbf{y} with $y_i = \frac{|x_i|}{\sum_I |x_i|}$ with $i \in I$ where I denotes the index set of active traders in the market and x_i the eigenvector component. The market power in a market is defined by the norm

$$m_y = \|\mathbf{y} - \mathbf{u}\|, \quad (32)$$

where $m_y = 0$ indicates perfect competition. For the eigensystem in Table 2 the distance from the ideal perfect competition market is 0.03, for the eigen-

system in Table 3 the distance is 0.26, so that we conclude that the second example shows less competition than the first. We use this comparison in the case study. This approach works because of the continuity property of Hermitian matrices and their eigensystems presented in theorem 1.

3.4 Complete Graphs and Perfect Competition

Without a market mechanism, perfect competition always leads to a complete graph in the eigensystem. For example, let us consider a complete graph as presented in Fig. 4.




Fig. 4. A complete graph: An indication of almost perfect competition

The matrix representing this graph is Mat. (33).

$$\begin{pmatrix} 0 & 4 + 2i & 6 + 5i & 3 + 6i & 9 + 10i \\ 2 + 4i & 0 & 8 + 7i & 10 + 7i & 1 + 10i \\ 5 + 6i & 7 + 8i & 0 & 6 + 4i & 3 + 3i \\ 6 + 3i & 7 + 10i & 4 + 6i & 0 & 6 + i \\ 10 + 9i & 10 + i & 3 + 3i & 1 + 6i & 0 \end{pmatrix} \tag{33}$$

which after rotation becomes Hermitian:

$$\begin{pmatrix} 0. & 4.2 - 1.4i & 7.7 - 0.7i & 6.3 + 2.1i & 13.3 + 0.7i \\ 4.2 + 1.4i & 0. & 10.5 - 0.7i & 11.9 - 2.1i & 7.7 + 6.3i \\ 7.7 + 0.7i & 10.5 + 0.7i & 0. & 7. - 1.4i & 4.2 \\ 6.3 - 2.1i & 11.9 + 2.1i & 7. + 1.4i & 0. & 4.9 - 3.5i \\ 13.3 - 0.7i & 7.7 - 6.3i & 4.2 & 4.9 + 3.5i & 0. \end{pmatrix} \tag{34}$$

The eigensystem belonging to this graph has the form shown in Table 4. Under the interpretation of the transaction stream of the accounting system of a market with direct transactions between the traders we see from

Table 4. Eigensystem and its eigenvectors for MComplete5 in Eq. (34) with $z = |z|e^{i\phi}$

λ_k		\mathbf{x}_{kl}				
1.	abs(z)	0.448	0.486	0.425	0.427	0.447
	arg(z)	-0.13	0	-0.094	-0.148	-0.226
-0.6	abs(z)	0.379	0.583	0.256	0.384	0.550
	arg(z)	-0.422	0	-2.937	-2.845	2.253
-0.35	abs(z)	0.602	0.332	0.344	0.523	0.369
	arg(z)	0	-2.350	3.002	0.994	-2.456
-0.20	abs(z)	0.299	0.342	0.745	0.406	0.270
	arg(z)	-0.148	-2.80	0	2.929	2.770
0.15	abs(z)	0.451	0.443	0.283	0.481	0.537
	arg(z)	-0.198	2.577	3.035	-2.424	0

the graph in Fig. 4 that all traders trade with each other and the trading volume is of the same order of magnitude. Such a complete graph indicates a market with almost perfect competition. Why almost? A market with perfect competition requires a complete graph where all edges have the same weight. Because we have no self-references, Eq. (23) must hold and this implies that in a complete graph $\lambda_1 + \sum_{i=2}^r \lambda_i = 0$ where λ_1 is the largest eigenvalue > 0 and all other eigenvalues have the same value. In addition, we see that trader 2 is the center of a strongly perturbed star because the largest absolute value in the eigenvectors 1 and 2 corresponds to this trader. From the Table 4 we see that for a complete graph the absolute values of the eigenvalues λ_k drop considerably. The eigenvector \mathbf{x}_1 corresponding to the largest absolute eigenvalue λ_1 almost follows a uniform distribution.

With a double auction mechanism as formal counter-party, the graph shown in Fig. 4 can never be observed. What we observe, is the graph shown in Fig. 5. This graph is constructed by running all network flows over the new node 6 which represents the account of the market mechanism (an intermediary clearing account).

Table 5. Eigensystem and its eigenvectors for graph in Fig. 5 with $z = |z|e^{i\phi}$

λ_k		\mathbf{x}_{kl}					
1.	abs(z)	0.32	0.35	0.30	0.31	0.31	0.71
	arg(z)	0.02	0.14	0	-0.07	-0.12	0
-1.	abs(z)	0.32	0.35	0.30	0.31	0.31	0.71
	arg(z)	-3.12	-3.0	3.14	3.07	3.02	0

The eigensystem for the graph in Fig. 5 is shown in Table 5. The result is, of course, no surprise. We notice the star structure typical for markets, where all trades are handled via a market mechanism. The largest absolute value in the eigenvector is the account of the market mechanism, all other traders




Fig. 5. Perfect competition with a double auction mechanism: A star graph

have approximately the same market power with the distance from perfect competition 0.024. The example in this section shows that reorganization of the market accounting system leads to different eigensystems. However, the interpretation remains the same.

3.5 Case Study: Analysis of Transaction Data for a Political Stock Market for the 2002 Federal Elections in Germany

As a case study, we scrutinize a political stock market for the German federal elections in 2002. In Germany, the federal election system basically is a proportional election system in which citizens vote for parties. Seats in parliament are currently computed according to the method of Niemeyer which is described in [4]. In 2002 the following five major German parties, namely SPD (Social Democratic Party of Germany), CDU (Christian Democratic Union), Grüne (The Greens), FDP (Liberal Democratic Party), and PDS (Party of Democratic Socialism), were explicitly included in the market, whereas all other parties were aggregated under the label “Others”.

Participants in the political stock market for the German federal election market simply registered at the web-site of the market and received upon registration a fixed amount of game money. Leaving the market was possible at all times, no money transfers occurred. No monetary transaction costs occurred. Trading on the market was free, so no trading fees were collected.

In such a market, two kinds of submarkets can be distinguished: On the primary market, portfolios containing one of each share present in the market can be bought or sold at a fixed price. Thus only portfolios are traded on the primary market, single shares are not traded and, therefore, they are not priced on the primary market. Furthermore, for each party exists a secondary

market where the shares of this party can be traded in a continuous double auction. After the close of the market, the shares in the portfolio of each trader are “bought back” by the market operator at a price that reflects the election’s result. The only incentive in the market was that traders were simply ranked according to the final value of their account.

Other political stock markets (e.g. the Iowa political stock markets) have been run with monetary incentives. In such a setting each trader transfers a usually fixed sum of money to the market operator. This money transfer usually is converted into game money credited to the trader’s account at a fixed rate. After the closing of the market the market operator evaluates the portfolio of each trader’s shares, and converts trader’s account balance back into real money and the real money balance is transferred back to the trader.

The market in question was open from beginning of June 2001 till the closing of the polling stations on 11/22/2002. During this period, 118 traders submitted 2905 offers for the six shares (SPD, CDU, Grüne, FDP, PDS, and others), resulting in a total of 2144 transactions. There was no monetary incentive for participation in the market, instead the rank of the players was used for motivation purposes.

The political stock market implies a specific structure for the communication occurring in it. Since the transactions take place in an anonymous setting (the counterparty is not known to the trader), the visible communication partner for a participant is the respective market for a share or the portfolio market. Thus, as communication graph, we obtain a bipartite graph with the traders on one side and the markets on the other as shown in Fig. 6. As a consequence, the basic communication pattern is that of a star: For each trader, the communication matrix shows a star with him as center. Equally, for each share and the portfolio, there is a star pattern with the market as center.




Fig. 6. Star graphs as building blocks of the communication structure

Traditionally, market quality is seen as dependent on the presence of self selected traders, an incentive system (monetary or motivational) that includes

incentives for truth-telling, the experiences of the participants and their familiarity with the market institution, the representativity of the participants, and the number and market power of traders [9, 34, 35].

For the analysis of the German federal elections market 2002, we used a volume weighting for the construction of the communication matrix: In general, we interpret each offer request that a trader sends to a market as a communication between the trader and the market. Inversely, each transaction that is executed by the market is counted as a communication between market and trader. Each communication increases the weight on the respective edge by the number of shares involved times the limit for offers (trader-share) or by the number of shares times the price for transactions (share-trader).

In this experiment, the market was not as efficient in predicting the outcome of the election as opinion polls. Table 6 gives an overview of the forecasts along with their precision. RMSE is an abbreviation for Root Mean Square Error defined as $\sqrt{\sum_{p \in P} (\bar{x}_p - x_p)^2}$ with \bar{x}_p denoting the forecast for party p , x_p the observed final election result for party p and P the set of all parties. Allensbach, Emnid, Forsa, FGW (Forschungsgruppe Wahlen) and Infratest are German pollsters.

Table 6. Forecasts for the German federal elections 2002

	Final Result	Allensbach	Emnid	Forsa	FGW	Infratest	PSM
	09/22/02	09/20/02	09/14/02	09/20/02	09/13/02	09/13/02	09/22/02
SPD	38.5	37.5	39	39	40	38.5	40
CDU	38.5	37	37	37.5	37	36	36.5
Grüne	8.6	7.5	7	7	7	8	7
FDP	7.4	9.5	8	7.5	7.5	8.5	8.2
PDS	4	4.5	5	4.25	4.5	4.7	4.1
Others	3	4	4	4.75	4	4.3	4
RMSE		3.18	2.72	2.67	2.88	3.16	3.23

Table 7 shows two submarkets, namely for CDU and for Grüne. We show only the positive eigenvalues. The eigenvalue is an indication of the amount of trading volume explained by the corresponding subspace. Comparing the eigenvalues, the market for CDU with an eigenvalue of 9349.7 had a considerably larger trading volume than the market for Grüne with an eigenvalue of 353.34.

We see that the submarket for CDU is actually very strongly dominated by trader number 139. The submarket for Grüne also has one strong trader, namely trader number 107. Since an efficient market would require perfect competition, we would expect that the weights of the traders in the eigenvector approximately follow a uniform distribution. We see for both markets that this is not the case and that therefore the forecasting precision of both submarkets is expected not to be high. Since the deviation from a uniform distribution

Table 7. Eigenvectors for the CDU and Grüne submarkets

Submarket CDU Eigenvalue 9349.76			Submarket Grüne Eigenvalue 353.343		
Trader	$ z $	Phase ϕ	Trader	$ z $	Phase ϕ
CDU	0.707	0	Grüne	0.707	0
139	0.706	-0.780	107	0.460	3.137
113	0.040	0	161	0.282	π
102	0.011	-0.062	102	0.210	π
101	0.008	-0.010	168	0.184	π
158	0.007	0	195	0.184	2.934
168	0.005	0	101	0.132	π
201	0.005	-0.098	123	0.120	-2.974
...

(see Eq. 32) is higher for the submarket for CDU (0.83) than for Grüne (0.22), we expect a larger forecasting error in the submarket for CDU than in the submarket for Grüne. This is indeed the case, the RMSE forecasting error for CDU is 2.0 per cent, and 1.6 per cent for Grüne.

This result is not surprising. Since in our experimental market only motivational incentives were present and these are generally thought to be weaker than monetary incentives, incentive compatibility for traders is not guaranteed. The Iowa political stock market, however, has traditionally been run with monetary incentives, so an investigation of Iowa market data with the method presented in this paper would be desirable for a second study. Unfortunately, the data publicly available from the Iowa markets do not allow the construction of the monetary flow over the market accounting system. Further experiments with different types of incentive systems are planned.

In election markets identification of the power structure is an indirect indication of market inefficiency. However, in energy markets identification of temporary power structures may serve as an indication of load pockets and give hints on necessary infrastructure improvements.

4 Summary

This result shows the application potential of the analysis of eigensystems of complex Hermitian matrices for markets. For this purpose, we succeeded in showing how arbitrary accounting systems can be represented as asymmetric network flows and we transformed them into Hermitian adjacency matrices. In addition, we presented accounting examples in different market structures and we generally described the eigensystems of continuous double auction markets and we applied this to the analysis of an election market.

In order to link market efficiency, incentive compatibility and forecasting accuracy we concentrated on the competitiveness of a market as expressed as the distance from a perfectly competitive market in the eigensystem. For the

German Federal Elections 2002 forecasting market this link can be observed empirically.

However, the discovery and measuring of other stylized behaviour patterns in continuous double auction markets (e.g. arbitrage, fraud, irrationality, speculation, ...) remains to be done.

It is beyond the scope of this paper to develop statistical tests on eigensystems. For a discussion of problems related to tests on eigensystems as e.g. the Johansen test on cointegration we refer the reader e.g. to [22]. In addition, we defer the formal definition of the bidding and trading process of a forecasting market with a fixed-price primary portfolio market and continuous double auction mechanism for the secondary share markets as a dynamic system as well as its estimation to the future.

To validate and to elaborate this type of analysis, a lot more work and a systematic exposition of the interpretation of Hilbert space theory in the context of markets as well as the analysis of experiments is necessary and planned for the future.

Acknowledgement

We gratefully acknowledge the funding of the project “SESAM” by the Bundesministerium für Bildung und Forschung within the scope of the research initiative “Internetökonomie” and the very helpful and constructive comments of two anonymous reviewers.

References

1. Berlemann M. (2002) Comment on Gregor Brüggelambert and Arwed Crüger, Forrest Nelson and Reinhard Tietz, and Jörg Bochow, Peter Raupach and Mark Wahrenburg: What can we learn from experimental asset markets? In: Bolle and Lehmann-Waffenschmidt [2], pp. 251 – 257.
2. Bolle F., Lehmann-Waffenschmidt M. (eds.) (2002) *Surveys in Experimental Economics: Bargaining, Cooperation and Election Stock Markets*, Contribution to Economics, Heidelberg, Physica-Verlag.
3. Brüggelambert G., Crüger A. (2002) Election markets: Experiences from a complex experiment. In: Bolle and Lehmann-Waffenschmidt [2], pp. 167 – 191.
4. Bundeswahlleiter (2002) ABC der Bundestagswahl 2002. Technical report, Der Bundeswahlleiter, Wiesbaden.
5. Copeland T., Weston F. (1988) *Financial Theory and Corporate Policy*. Addison-Wesley, Reading, 3 edition.
6. Fama, E. F. (1970) Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2):383–417.
7. Fama E. F. (1976) Efficient capital markets: Reply. *Journal of Finance*, 31(1):143–145.

8. Fama E. F. (1991) Efficient capital markets: II. *Journal of Finance*, 46(5):1575–1617.
9. Forsythe R., Nelson F., Neumann G. R., Wright J. (1992) Anatomy of an experimental political stock market. *The American Economic Review*, 82(5):1142–1161.
10. Froilán J. M., Dopico M., Molera J. M. (2000) Weyl-type relative perturbation bounds for eigensystems of hermitian matrices. *Linear Algebra and its Applications*, 309(1):3–18.
11. Geyer-Schulz A. (1991) An Introduction to Financial Accounting with APL2. Technical report, ACM SIGAPL, New York, N.Y..
12. Gomber P., Schweikert U. (2002) Der Market Impact: Liquiditätsmaß im elektronischen Wertpapierhandel. *Die Bank*.
13. Hayek F. A. (1935) The present state of the debate. In: Hayek F. A., (ed.) *Collectivist Economic Planning*, pp. 201–243, London.
14. Hayek F. A. (1945) The use of knowledge in society. *American Economic Review*, 35:519–530.
15. Horn R., Johnson C. R. (1990) *Matrix Analysis*. Cambridge University Press, Cambridge.
16. Hoser B., Geyer-Schulz A. (2005) Eigenspectralanalysis of hermitian adjacency matrices for the analysis of group substructures. *Journal of Mathematical Sociology*. Accepted for publication.
17. Hurwicz L. (1973) The design of mechanisms for resource allocation. *American Economic Review*, 63(2):1–30.
18. Ipsen I. C. F. (2003) A note on unifying absolute and relative perturbation bounds. *Linear Algebra and its Applications*, 358:239 – 253.
19. Kato T. (1995) *Perturbation Theory for Linear Operators*. Springer, New York, 2 edition.
20. Kleinberg J. M. (1999) Authoritative sources in a hyperlinked environment. *JACM*, 46(5):604–632.
21. Latham M. (1985) Defining capital market efficiency. Technical report, Finance Working Paper 150, Institute for Business and Economic Research, University of California, Berkeley.
22. Maddala G. S., Kim I.-M. (2001) *Unit Roots, Cointegration, and Structural Change*. Cambridge University Press, Cambridge.
23. Meyer C. D. (2000) *Matrix Analysis and Applied Linear Algebra*. Society for Industrial and Applied Mathematics, Philadelphia.
24. Milgrom P., Roberts J. (1992) *Economics, Organisation and Management*. Prentice Hall, Upper Saddle River, 1 edition.
25. Nelson F., Tietz R. (2002) Expectations and rational actions in an experimental financial market. In Bolle and Lehmann-Waffenschmidt [2], pp. 193 – 227.
26. Page L., Brin S., Motwani R., Winograd T. (1998) The Page Rank Citation Ranking: Bringing Order to the Web. Technical report, Computer Science Department, Stanford University.
27. Pinches G. E., Kinney W. R. Jr. (1971) The measurement of the volatility of common stock prices. *Journal of Finance*, 26(1):119–125.
28. Roll R. (1984) A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance*, 39(4):1127–1139.
29. Rubinstein M. (1975) Securities market efficiency in an arrow-debreu economy. *American Economic Review*, 65(5):812–824.

30. Samuelson P. A. (1948) Consumption theory in terms of revealed preference. *Economica*, 15(60):243–253.
31. Shubik M. (1980) *Market Structure and Behavior*. Harvard University Press, Cambridge.
32. Simon H. A. (2000) *Administrative Behavior : A Study of Decision-Making Processes in Administrative Organizations*. Free Press, New York, 4 edition.
33. Smith V. L. (1976) Experimental economics: Induced value theory. *American Economic Review*, 66(2):274–279.
34. Spann M., Skiera B. (2003) Internet-based virtual stock markets for business forecasting. *Management Science*, 49(10):1310 – 1326.
35. Spann M., Skiera B. (2004) Einsatzmöglichkeiten virtueller Börsen in der Marktforschung. *ZfB – Zeitschrift für Betriebswirtschaft (Ergänzungsheft)*, (2):25–48.
36. Jaffe J., Ross S. A., Westerfield R. W. (2005) *Corporate Finance*. McGraw Hill, Boston, 7 edition.
37. Stone M. H. (1932) *Linear Transformations in Hilbert Space and their Applications to Analysis*, Vol. 15 of *American Mathematical Society, Colloquium Publications*. American Mathematical Society, New York.
38. Wasserman K., Faust S. (1994) *Social Network Analysis, Methods and Applications*. Cambridge University Press, Cambridge.
39. Wellman B. (2001) Computer networks as social networks. *Science*, 293:2031–2034.
40. Wolfram Research Inc. (1999) *Mathematica*. Urbana-Champaign, Illinois, version 4 edition.

The Role of Assortative Mating on Population Growth in Contemporary Developed Societies

Mike Murphy

Department of Social Policy, London School of Economics, United Kingdom
m.murphy@lse.ac.uk

Summary. Assortative mating is a widespread feature of human behaviour, which has a number of suggested benefits. The question of whether it contributes to population growth in contemporary societies is considered using the micro simulation program SOCSIM. Ways of parameterising heterogeneous fertility and nuptiality, and the relationship of such parameters to those of both fathers and mothers are considered. The magnitude of the correlation between numbers of sibs of the partners is similar to that of the correlation between number of sibs of the mother and the number of her own children. Models that can generate such degrees of similarity are discussed.

If continued for 250 years, populations with long-standing assortative mating, and with demographic parameter values that bound those found in practice would have fertility levels about 2% to 30% higher than those without assortative mating, and also earlier age at first marriage. Population size is between about 30% and 200% higher at the end of the period. I conclude that the effect of assortative mating in which the fertility backgrounds of spouses are positively correlated leads to higher population growth.

1 Introduction: Determinants of Fertility - Consequences and Possible Causes

A population with a higher long-term rate of growth, no matter how small the advantage, will come to dominate numerically any population with a lower one and the overall population eventually becomes effectively homogeneous and consists only of the higher growth population (e.g. [27, Sect.1.5]). A related result is found in Mendelian genetics: Fisher's fundamental theorem of natural selection (FTNS) [15] which states that any trait correlated with fitness should have a heritability¹ of zero. On the other hand, some empirical studies show that there is very large variability in fitness: for example, species such as the

¹ Heritability is usually measured as the proportion of total variance of a trait that is due to genetic factors, although other definitions exist [24]. Heritability can be zero either because the trait is not inherited, or because there is no genetic

sage grouse in which more attractive males have hugely greater reproductive success in mating, producing the so-called ‘lek paradox’ [8], [46].

The rate of natural growth is determined by levels of fertility and mortality, and intergenerational continuities in these will lead to growth rates being similar in different time periods. The mechanism that is responsible for similarities between generations may be genetic, environmental or an interaction of the two, but the relative contributions of nature and nurture to such correlations are largely irrelevant. If inherited factors exist that are correlated with reproductive fitness, then it is simply the overall value of the correlation that will have implications for population dynamics.

In contemporary developed societies, variations in mortality have little influence on reproductive success, therefore fertility is the major determinant, which, in turn, is heavily affected by partnership behaviour since fertility rates are much higher for married women. Studies of recent populations have indicated that there is now a correlation coefficient value of about 0.2 for fertility of parents and children averaged across a number of different studies [37]. More specific studies based on kinship behaviour genetic models and twin studies have found a relatively strong heritable component of fertility and related behaviours, including age at first birth, nuptiality and marital breakdown in recent periods [34], [11], [29], [49], [28].

There are a number of reasons for the superficially incompatible empirical and theoretical results on heritability of fertility. Explanations for persistent variability in fitness include a host-parasite ‘arms race’ [19] and the role of genetic mutation [30], [23]. However, Fisher’s fundamental theorem of natural selection assumes essentially long-term stability by a large population in an unchanged environment with Mendelian inheritance and non-assortative mating. The last of these conditions is the focus of this paper. Non-random mating is an almost universal feature of human experience, and spousal similarity in humans is observed on a wide range of socio-economic, psychometric and anthropometric traits (e.g. [20], [25], [56], [13], [50], [9], [10], [42], [55], [14], [26], [5]. In this paper I aim to quantify the magnitude of the effect of assortative mating on population dynamics in contemporary developed societies.

The key determinant of population change in modern societies is fertility since only mortality under about age 45 is relevant to population growth and mortality rates at these ages are so low that changes have little impact on overall growth (if the ‘grandmother effect’ — the beneficial effect that post-menopausal grandmothers have on the reproductive success of her children and the survival of her grandchildren even if she is unable to bear children herself [22], [51], [52], [31] — is ignored; I also ignore migration at this stage since I am concerned solely with natural increase). The phrase ‘inherited’ does not necessarily imply individual-level transmission, since inter-generational conti-

variability in the population. Clearly we inherit a propensity to reproduce [16].

nities are observed in cross-national comparisons, nor does it imply causation since it may reflect other factors: for example in cross-national comparisons, observed similarities between generations in levels of fertility or mortality would be likely to be attributed to factors such as level of development in the countries concerned. However, at the individual level, an inherited propensity to have children is found, at least if the number of children born acts as a proxy for this unobserved variable. Although information on fathers is less readily available than for mothers, the evidence is that fertility is also partially inherited from fathers, and that the magnitude of the effect is broadly similar [37], even though socialisation theory might suggest that mothers should have a more influential role since they typically play a greater role in bringing up children. Assuming that both parents play a role is certainly more reasonable than assuming transmission only occurs through the mother. (There are, however, two problems with including fathers: the first is that not all children have a known father, in this application, this is the case for all births outside of a married or cohabiting union when only the characteristics of the mother are available; the second is that the nominal father may not be the biological father.)

Table 1 shows the association between number of live-born children and their sibs of women aged 35 to 50 from the 1976 British Family Formation Survey [12], and also with the number of sibs of their current partner. The values of the correlation coefficients are very similar, and positively related to the couples' fertility outcomes (not all children will be those of the current partner, but the great majority will be so, since these data refer to childbearing around the 1960s). Table 1 also shows the association between numbers of sibs of the two partners for all women aged 16 to 50 in a partnership. The magnitude of this correlation between partners is about 0.2 in all age groups, essentially the same as that for parents and children. Thus the fact that partners (who are the parents of the subsequent generation) come from families with positively associated inherited fertility will also have an influence on the demographic characteristics of kinship distributions, and possibly on long-term population dynamics: the latter is the question that is addressed here.

Some degree of positive assortative mating would appear to be advantageous for fitness – if this was not the case, a population that did not practice it (or even had negative assortative mating) would have at least as great fitness, and assortative mating would not have been selected for (possibly some intermediate evolutionary stable solution (ESS) could arise). If mating was so non-assortative that it took place with another species, fitness would be zero. However, there are limits to the degree of acceptable assortative mating: many societies, 44%, have prohibitions on incest [57] although the reasons for this may have to do with strategies for resource accumulation as well as with restricting inbreeding, which leads to lower fitness (there is some aversion to mating with those children with whom one is brought up, the Wester-

Table 1. Correlation coefficients of number of sibs of female respondent and partner, and number of her live-born children, Family Formation Survey, 1976

Age-Group	Own & partner's sibs			Sibs & own children			Partner's sibs & own children		
	Correl-ation	Stand. error	Sample size	Correl-ation	Stand. error	Sample size	Correl-ation	Stand. error	Sample size
16-19	0.10	0.07	241	-	-	-	-	-	-
20-24	0.11	0.04	649	-	-	-	-	-	-
25-29	0.20	0.03	999	-	-	-	-	-	-
30-34	0.19	0.04	938	-	-	-	-	-	-
35-39	0.14	0.03	868	0.15	0.03	906	0.16	0.04	868
40-44	0.18	0.03	837	0.14	0.03	891	0.16	0.04	837
45-50	0.19	0.03	829	0.14	0.03	876	0.14	0.04	829

marck effect). As usual, social, biological and developmental factors are all relevant as determinants of observed patterns. The benefits of positive assortative mating include both social reasons (people are more likely to meet, and to feel comfortable with people who are similar to them), and also biological or evolutionary reasons, such as Genetic Similarity Theory in sociobiology, that relates spousal similarity to the existence of altruistic behaviours in humans, and Kin Selection Theory, whereby an individual will be more willing to engage in self-sacrificing behaviours if the benefits to close genetic relatives outweigh the costs [18]. All sexually reproducing species exhibit choice in selection of partners, especially females, for whom the penalties of a poor choice are greater, since they have to invest more in their offspring than do males in almost all cases [4], [9], [2, Chap.5], so there are plausible social and genetic mechanisms for the existence of assortative mating. However, mechanisms that led to assortative mating being selected for in the past, such as reduced mortality, might not be observed in contemporary populations with low mortality rates. However, if the mechanism was to increase fitness by increasing fertility over what it might otherwise have been, then such an effect might still be observable.

I discuss the role of assortative mating on population growth in Sect. 3. But before doing this, however, I consider the nature of inherited factors, how they are modelled, and how they are related to actual fertility in Sect. 2.

2 Microsimulation Modelling of Fertility and Nuptiality

Demographic micro simulation is the principal method used to elucidate kinship patterns in historical, contemporary and future populations ([53], [60], [62], [64], [59]). This analysis is based on the SOCSIM model developed by Professors Hammel, Wachter and colleagues at Berkeley over a number of decades [61], [21], [39], which has a number of features that make it attractive

for this application. It is a closed model so that alternative marriage strategies can be assessed. The model, written in the C language, is freely available and the code can be amended or extended by users, in this case to include assortative mating (the program includes both cohabitation and formal marriage, and fertility of cohabitantes is intermediate between non-partnered and married people in this application: the terms partnership, mating and nuptiality are therefore used interchangeably in this paper).

An initial population structure of that of England in 1850 taken from the 1851 England and Wales Census of Population evolves under the given monthly probabilities of fertility, mortality and nuptiality using empirically derived values for each series from 1850 to 2000. These rates are based on vital registration data following the establishment of the system in 1837. Fertility varies by age, marital status and parity; mortality and nuptiality vary by sex, age and marital status. Some of these baseline indicators of fertility and nuptiality over the simulation period are shown in Table 2. These baseline values represent the best estimate of how demographic parameters have varied in Britain since 1850; further details are given in [38].

Two populations with initial sizes of 40,000 people are included, one with and one without assortative mating; the way in which this is operationalised is discussed later. Fertility is below replacement-level from about 1970, and remains slightly below it, in line with current trends and expectations [43]. Life expectancy is about 83 for males and 89 for females in 2100. Total population size is nearly 200 thousand in 2000, before declining somewhat by 2100. The relatively large sizes means that it was unnecessary to replicate the runs: for example, mortality rates were the same in both populations, and the almost identical values for life expectancy show that the role of stochastic variability is trivial.

Since the model is closed [60], partners have to be found within the existing simulation population using an algorithm to ensure a realistic distribution of spousal age differences (there is also a prohibition on incest). Closed models are more complex than open models in which a partner is created when required, but as they do not come with any demographic background, it is impossible to investigate general kinship and other relationships, and crucially for this paper, the issue of assortative mating since only one partner in a couple will usually have such information.²

While an intergenerational fertility correlation of about 0.2 is typical of values found in practice, for genetic transmission between generations of inherited characteristics, *achieved* parental family size is not the most appropriate starting point. Sibship size is, at least in part, a socialization, environmental and life course variable since it affects an individual's childhood circumstances,

² This is not possible if a woman gives birth without an identified partner (in this analysis, any birth outside of a cohabiting or marital union), since such fathers are regarded as unknown by the program.

Table 2. Summary of population values

Period	Pop ⁿ Size		TFR		FM ^{male}		FM ^{female}	
	N-A	A	N-A	A	N-A	A	N-A	A
1850-1860	40000	40000	4.43	4.42	29.1	29.1	26.6	26.7
1860-1870	43574	43523	4.44	4.41	28.5	28.2	25.3	25.1
1870-1880	46793	46555	4.46	4.48	27.9	28.0	25.0	25.1
1880-1890	51016	51181	4.60	4.66	28.1	28.0	25.2	25.2
1890-1900	57094	57394	4.12	4.15	27.9	27.8	24.8	24.8
1900-1910	62525	62828	3.62	3.67	27.6	27.6	24.8	24.8
1910-1920	69158	69736	2.78	2.84	28.0	28.0	25.0	24.9
1920-1930	74365	74962	2.33	2.38	28.0	27.8	24.5	24.6
1930-1940	78834	79677	1.72	1.80	27.9	27.8	24.4	24.3
1940-1950	79920	81490	2.01	2.11	27.1	26.8	23.5	23.5
1950-1960	83083	85546	2.26	2.35	26.2	25.8	22.8	22.7
1960-1970	85596	88722	2.66	2.77	25.3	25.1	23.0	22.8
1970-1980	89694	94186	1.81	1.85	25.3	25.3	22.3	22.2
1980-1990	89032	94445	1.69	1.77	28.1	27.9	24.4	24.3
1990-2000	88450	95398	1.62	1.66	29.7	29.5	27.1	27.0
2000-2010	88346	95893	1.74	1.84	30.5	30.4	28.1	28.2
2010-2020	88076	96909	1.70	1.82	30.6	30.4	28.6	28.3
2020-2030	86500	96631	1.71	1.86	30.9	30.7	28.9	28.8
2030-2040	85378	96703	1.71	1.83	30.8	30.6	28.7	28.8
2040-2050	81442	93620	1.73	1.80	30.8	30.6	28.9	28.7
2050-2060	76300	89233	1.71	1.85	30.8	30.6	29.1	28.9
2060-2070	71574	84996	1.71	1.81	31.0	30.8	29.1	28.8
2070-2080	67534	81584	1.75	1.79	30.7	30.7	29.0	28.9
2080-2090	64235	78726	1.78	1.79	30.9	30.9	29.0	29.2
2090-2100	60949	75733	1.73	1.87	30.7	30.7	29.0	29.0
2100-	57839	73282						

Note: N-A non-assortative, A assortative;
 FM^{male} FM^{female} mean age at first marriage of males and females respectively.

including through mechanisms such as education or standard of living, and consequently it may affect the child's later behaviour including subsequent childbearing (see, for example, [54]). These standard social science mechanisms have often been assumed to be the ways in which intergenerational continuities in behaviour arise [58]. In fact, fertility cannot be inherited; rather it is the propensity to have a particular family size. In terms of a proximate determinants framework [6], this can be decomposed into the probability of a woman being in a sexual union, and the chance of a full-term pregnancy arising from that union. Both factors are known to have an inherited character, and therefore we consider models in which the propensity to form (and to dissolve) partnerships is also inherited. It is possible that actual sibship size may have a separate effect, in that people may choose a family size that is

similar to that of their own upbringing. However, even with identical average sizes in successive generations, the average family size a child is brought up in will be much larger than their own average number of children [47].

For investigating models of intergenerational transmission of fertility behaviour, there must be heterogeneous probability of giving birth. The monthly probability of conception or birth (depending on the context) to a woman at risk of the event is referred to as fecundability [17], [32]. I therefore consider models in which this underlying propensity to give birth (a) is heterogeneous; (b) may be intergenerationally transmitted; and (c) is correlated between partners. In SOCSIM, an individual woman's monthly probability of birth is given by the appropriate baseline distribution multiplied by a random variate, which is fixed at birth. Wachter [61] suggested that the heterogeneity distribution have a mean value of 1.0, so making the specification of models much easier, since the average fertility rates of the group of women are equal to their baseline input rates.³ Without such an assumption, the model is under-determined, since the same result can be obtained by multiplying the heterogeneity parameter by an arbitrary value, and dividing the baseline value by the same amount. Inclusion of such heterogeneity produces distributions of fertility that are close to the empirically observed baseline values, but with greater and more realistic variability in the fertility distribution than is the case for a homogeneous population, and, in particular, it permits the role of intergenerational transmission of heterogeneity to be assessed. In this analysis, periods of susceptibility to pregnancy are not identified, apart from a minimum gap of 12 months between births. I therefore call this fertility multiplier, which gives the *relative probability* of a woman in a particular age, marital status and parity category giving birth in a given month, 'naïve standardized fecundability' (NSF) - naïve because it subsumes all the other proximate determinants and standardized because the mean value is one [39]. As with the usual definition of fecundability, it obscures some aspects of the reproductive process but the present formulation has some advantages by permitting concentration on intergenerational continuities.

The form of the distribution of the lifetime fertility multiplier must be constrained, since the distribution is non-negative, unimodal and of mean one. In most analyses, fecundability is assumed to follow a skewed distribution such as the beta distribution, which is mathematically convenient for analysing heterogeneity [1], [63]. Although the evidence for a specific distributional form is lacking, the precise form is probably unimportant. As conventionally done, I choose a mathematically convenient form with the desired properties. For analysing intergenerational transmission, when there is a need to have a correlation in the values of naïve standardized fecundability between successive

³ This refers to rates within age, marital status and parity. Measures of overall fertility such as total fertility rate (TFR) will also depend on the proportion of time spent in these states. The paper will consider the relative contributions of these two components later.

generations, the beta distribution has the disadvantage that it does not easily produce correlated distributions⁴, therefore I use an approximate beta distribution with mean one and variance of 0.5 based on a transformed normal distribution with mean μ and standard deviation σ , $N(\mu, \sigma)$, to generate the naïve standardized fecundability distribution (NSF) as follows (Fig. 1):

$$NSF = 2.65 / \{1.0 + \exp(N(0.592, 0.934))\} \quad (1)$$




Fig. 1. Approximation to beta distribution

The underlying normal variate for person i , N_i , is generated to have a correlation with the corresponding value for his/her parents, N_m and N_f , for mother and father respectively.

$$N_i = \text{alphafert} \times (N_m + N_f) + \text{sqrt}(1. - 2.0 \times \text{alphafert}^2 \times 1.3) \times N_a \quad (2)$$

where N_a is an independently generated variate and *alphafert* is a constant chosen by the investigator to produce a given level of intergenerational correlation.⁵ Note that quantities such as naïve standardized fecundability are unobserved, rather than directly measured ones such as completed fertility or

⁴ In its original formulation in the SOCSIM model, the lifetime fertility multiplier was generated by a cubic approximation to the beta distribution with mean one, variance of 0.416 (coefficient of variation of 0.645), maximum of 2.4 and minimum of zero [61].

⁵ In the original version of SOCSIM, the basic model of fertility inheritance is that the value of naïve standardized fecundability of a daughter, f_d is given by $f_d = \alpha f_m + (1-\alpha)f$ with the notation as above, where f is generated as in footnote

mean age at marriage. The use of an underlying normal distribution means that correlated values can be readily obtained since it is straightforward to generate multivariate normal variates with a pre-determined level of correlation. In practice, the correlation between variates such as those of (1) based on correlated normal variates is very close to that of the underlying normal distribution, which makes it straightforward to specify the degree of correlation between successive generations. Thus the correlation in the transformed beta-like distribution is very close to the value *alphafert*, which is the inter-generational correlation of the untransformed values.

I also use the same approach as in (1) and (2) above to model intergenerationally transmitted heterogeneous partnership behaviour, since there are known to be intergenerational transmission of such behaviours. I call the heterogeneity multiplier applied to the baseline partnership rates (both marriage and cohabitation) naïve standardized nuptiality (NSN), and the analogous one applied to partnership breakdown rates, naïve standardized divorce (NSD). Since fertility rates differ substantially between partnership statuses, variations in overall fertility depend both on nuptiality and partnership-specific fertility rates, and variability in each of these two components contributes approximately equally to variability in total fertility (I do not include mortality heterogeneity in the model since the effect on population dynamics is trivial as noted above). In this application, I assume that the three heterogeneity parameters are independently distributed, but in principle they could be allowed to be correlated.

If there is no inherited component to fertility or nuptiality, the mean value of naïve standardized fecundability and nuptiality will remain one, and the observed population fertility rates will be equal to the baseline values at that time period (which clearly makes model specification easier). However, if there is inheritance, then on average more fecund women will have larger numbers of daughters who will themselves have higher than average values of naïve standardized fecundability and who will form a higher proportion of their generation, so that the mean value of naïve standardized fecundability will increase generation by generation (we assume that such correlations will be positive, as found empirically, but negative correlations are also theoretically possible). Thus the overall level of fertility will be higher than in the absence of such transmission because it is given by the original baseline value multiplied by the average population value of naïve standardized fecundability.

The existence of intergenerationally-transmitted naïve standardized nuptiality, will further tend to inflate the correlation of fertility between mothers

4. α can run from zero (no intergenerational transmission) to one (each daughter has exactly the same naïve standardized fecundability as her mother). However, if there is inheritance, the variance will decline. For example, if f_m and f are distributed independently from the same distribution, then if α is equal to 0.5 the variance f_d in the next generation will be only about 0.5 of that in the previous one. Therefore the variability (and the shape) of the distribution and hence its influence in subsequent generations will decline.

and daughters since early-marrying mothers will spend more of their fertile period in the higher-fertility married state, and thus have more children than average, who will themselves in turn have more children because of their higher propensity to form partnerships at young ages (although historical data suggest a relatively low intergenerational correlation for age at marriage, [33]). If the heterogeneous trait is uncorrelated with reproductive success ('fitness'), then the distribution will remain constant over time; an example of this is marital breakdown propensity (naïve standardized divorce), which is essentially independent of childbearing level (at least with the assumption that the multipliers for fertility, nuptiality and marital breakdown are independent).

The effect of transmitted heterogeneity will cumulate over generations (and there is empirical evidence of multi-generational transmission, [40], [36]). If the trait is correlated with fitness, then not only the mean value, but also the form of the distribution will change. Since those with high values will have greater reproductive success, over time this will lead to the distribution converging to the highest value *and* to a reduction in its variability if the distribution is bounded, which will be the case for quantities such as a monthly probability of birth. Such an outcome is implausible since such homogeneity in reproductive behaviour is not observed. Moreover, it is likely that childbearing performance is relative rather than absolute: for example, a family size of six children would be considered very high in contemporary European societies, whereas around 100 years ago, 60% of fertile women who married at ages 20-24 were found to have had eight or more children by the end of the childbearing ages in the 1911 Census of England and Wales [48, Table XVIII, p xlv]. Thus if 'high' and 'low' are related to the norms and values of the society in question, the values of the moments of *relative* distributions, such as naïve standardized fecundability and naïve standardized nuptiality, would also be expected to remain largely constant over time.

Since the baseline rates used in the model are equal to those actually observed in order to replicate the real situation, then the heterogeneity parameter *must* have mean values close to one and, as pointed out above, there is a strong case that the variability of the heterogeneity distributions would remain relatively constant. Clearly the variance in fitness is not collapsing, as would happen if values converged towards a fixed upper limit. Thus, since there is no reason to believe that the variance is changing, mechanisms are required to maintain approximately fixed levels over time of mean values and variability. In particular, if variability changes, it is impossible to separately identify changes due to assortative marriage from the effects of changing variability *per se*, and I want to concentrate on the former by constraining the latter. The coefficient of 1.3 in (2) above is an empirical value to adjust naïve standardized fecundability and naïve standardized nuptiality for assortative mating, so that their variances remain reasonably constant over time (this is unnecessary for naïve standardized divorce and no adjustment was done in this case). The justification for this value is as follows:

If ϵ_c is the child's, ϵ_m the mother's and ϵ_f the father's values, and ϵ a new independently generated value; the coefficients relating them are α_1 and β_1 and we want the variances to remain constant (for convenience, in all distributions, variance taken as 1.0 and the mean as 0.0).

Thus

$$\epsilon_c = \alpha_1(\epsilon_m + \epsilon_f) + \beta_1\epsilon$$

Calculating the variance therefore gives

$$1 = 2\alpha_1^2(1 + \rho_{mf}) + \beta_1^2$$

$$\rho_{cm} = \alpha_1(1 + \rho_{mf})$$

Where ρ_{mf} is the correlation between values for mothers and fathers, and ρ_{cm} is the correlation between values for mothers and children, if $\rho_{mf} = 0$, i.e. non-assortative mating, then

$$1 = 2\alpha_1^2 + \beta_1^2$$

$$\rho_{cm} = \alpha_1$$

if $\rho_{mf} = \rho_{cm}$ i.e. correlation between parents and parent-child are similar in magnitude (as in Table 1)

$$\rho = \alpha(1 + \rho)$$

$$\rho = \alpha/(1 - \rho)$$

$$\alpha = \rho/(1 + \rho)$$

$$\beta^2 = 1 - 2\alpha\rho$$

$\rho_{mf} = \gamma\rho_{cm}$ i.e. ratio of correlation between parents and parent-child is fixed

$$\rho_{cm} = \alpha_1(1 + \gamma\rho_{cm})$$

$$\rho_{cm} = \alpha_1/(1 - \alpha_1\gamma)$$

if $\beta^2 = 1 - 2\alpha_1^2$, then

$$\sigma_c^2 = \sigma_m^2(1 + 2\alpha_1^2\rho_{mf})$$

Now if we choose a plausible value for the correlations, e.g. $\rho_{mf} = .3$, $\rho_{cm} = .3$, $\alpha_1 = .3/1.3$

$$\beta^2 = 1 - (2 \times .3^2 \times 1.3)/1.3^2$$

so $\alpha_1 = .23$, $\beta_1 = .83$.

Thus an adjustment factor of the order of 1.3 is needed to maintain a fixed variance over time with correlations of the order of 0.3.

The second issue is that the higher reproductive success of those with higher values of naïve standardized fecundability or naïve standardized nuptiality will tend to increase the mean value of these distributions over time. In order to maintain a constant value for mean naïve standardized fecundability

and naïve standardized nuptiality, the innovation mean is reduced to counteract this tendency. In practice, this may be achieved by reducing the mean of the innovation process, N_a , for naïve standardized fecundability in (2) from zero by $mean_value(NSF) \times 3.0 \times \alpha_{fertility}$, and correspondingly for naïve standardized nuptiality.

Partners with Positive Correlations

So far, the extent to which people choose partners with characteristics similar to themselves has not been considered, and I now turn to assortative marriage. In the case of intergenerational transmission, a child’s characteristic can be relatively easily related to that of his/her parents: for example, if the parent has a trait distributed as $N(0,1)$, then the child’s value of

$$\alpha \times parental_value + \sqrt{(1 - \alpha^2) \times N(0,1)} \quad (3)$$

will have a correlation of α with the parent’s value and also be distributed as $N(0,1)$. However, mating is a process of mutual adjustment in which men and women attempt to maximise their chances of success [35]. The preferred way of finding a partner of the opposite sex in closed micro-simulation models is to form a ‘partnership market’, in which people seek a partner from the pool of eligible members using an algorithm to identify the two partners (e.g. [7]).

I therefore defined a function that calculated a score for each pair of potential partners as follows: when a person joins the partnership market, with probability determined by the observed partnership rates, the potential spouse with the highest score is selected provided the score is above a threshold value, or by random selection if there is more than one pair with the same maximum value. The assumption is that people have a preference for a particular type of partner who is more like themselves than the population at large; an alternative would be to assume that everyone wanted to partner with the most attractive person, and that person would chose the individual of the opposite sex with the highest score, i.e. every person would be ranked on an absolute scale, rather than one relative to a given potential partner (e.g. [35, pp. 198–215]). The former seems to be more in accord within people’s behaviours, although much work remains to be done in this area [3].

Clearly, age is a major factor that cannot be ignored; men tend to partner with women who are slightly younger than themselves on average, and the number of partnerships with large age differences is relatively small [44] — although the question of how the choice is made and the relative roles of men and women in the process remains a subject of lively debate, there exist clear preferences for a partner within specific age bounds ([45], [2]). The non-assortative model had only a simple age-preference scheme, whereby a score was attached to each potential match: this score was set at zero for couples in which the man was 12.5 years older than the woman, or 10 years younger. The score was set at a constant maximum value in the range where the man is

between 5 years older and 2.5 years younger than the woman. At intermediate ages, the score declines linearly to zero at the ends of the eligibility range. The highest non-zero score was used to select the partner, or if there was more than one potential partner with the same highest non-zero score, the partner was chosen randomly from these. It can be argued that including age preference represents choice to some extent, but not to include some such age preference would lead to unrealistic results, thus these assumptions produce plausible distributions of spousal age differences, but spouses are otherwise independent in their characteristics.

The function used in the assortative model includes age as above, but, in addition, a man and a woman with more similar fertility and nuptiality traits are more likely to form a partnership. A major objective of this study was to obtain a function of the spouses' scores on unobserved variables such as naïve standardized fecundability and naïve standardized nuptiality that were positively correlated. This was done by defining a score for the two potential partners that penalises differences as follows:

$$-0.1 \times \{(NSF_h - NSF_w)^2 + (NSN_h - NSN_w)^2 + (NSD_h - NSD_w)^2\} + N(0, \lambda) \quad (4)$$

(with h denoting potential husband and w denoting potential wife)

The larger the value of the constant λ in the second component, the less weight is given to similarity on the other variables of the potential partners, and consequently, the lower the correlation between spouses on these characteristics. The variance of the first component is about 0.01, and the two values chosen for λ were 1 ('Standard correlation'), and 0 ('High correlation'). The high correlation case is chosen as an extreme case, and all assortative results refer to the standard correlation case, apart from those of Fig. 4 and Table 4 presented later.

3 Results of Models under Alternative Assortative Mating Assumptions

I am attempting to identify the role of a particular factor, that of assortative mating, over and above the overall effect of inherited behaviour, which is known to be large, so it is necessary to be able to isolate the effect of assortative mating, which is likely to be small in comparison. I compare two models, both with inherited fertility and nuptiality. I set α in (2) to 0.4 for naïve standardized fecundability, nuptiality, and divorce over the simulation period — note that the intergenerational correlation between these underlying variables and corresponding observed quantities such as achieved family size, or average age at first marriage will be less than this, so the value of the intergenerational correlation coefficient in variables such as naïve standardized fecundability will be higher than the values shown in Table 1.

Macro-Level Effects

Mothers and fathers with a high (or low) propensity to form or dissolve partnerships will pass on this trait to their children. In comparing two populations with different patterns of inheritance, it is important that they are subject to the same demographic regime which is achieved by treating them as two groups within the same simulation run (there is an option in SOCSIM to do this) which would not be possible with independent runs. Exactly the same baseline demographic parameters and model assumptions are used in both groups, except that they differ only in that one excludes and one includes assortative partnership. Thus any differences observed between the two populations are due to deviations from the overall average values produced by the sub-group differences in marriage function. Table 3 and Fig. 2 show that the means, standard deviations, and overall distributions of naïve standardized fecundability, naïve standardized nuptiality and naïve standardized divorce for selected birth cohorts from 1850 to 2100 remain very similar over time after the adjustments of Sect. 2.

Table 2 and Fig. 3 show the main simulation results comparing the assortative and non-assortative cases. The values of mortality (shown as life expectancy) and divorce are essentially the same in the two cases, but the level of fertility is slightly higher, and the average age at marriage is slightly lower in the assortative case. Statistics referring to the average number of children born are based on people who reached at least age 45, since to include all individuals in the simulation, including those who do not even reach the age of reproduction, would distort the interpretation of the relationships shown here (especially important in the last period, with censoring at 2100 and the fact that birth to unknown fathers do not contribute to male fertility). However, the values of the three heterogeneity variables shown are not subject to this problem, and all ages are included. Looking first at the simplest case, that of naïve standardized divorce, the mean and standard deviation remain largely constant without any further adjustment (Tables 3(a) and 3(b), see also Fig. 2), because it is effectively independent of the other variables in the Table, i.e. it is uncorrelated with fitness. The standard deviations of naïve standardized fecundability and naïve standardized nuptiality also remain fairly constant over the period, showing that the adjustment process described earlier has produced the desired result, and so changes in measures such as population size or TFR are not artefacts of changing variability in the heterogeneity distributions as could be the case if such adjustments are not made. The mean values are close to one, but not exactly equal to it, but these differences do not affect the interpretation of differences between groups. In Table 3(c), the values of the correlations between these variables and number of children born to the same individuals are based on people who reached age 40 or over, so that they have effectively completed childbearing. Note that the magnitude of correlations of overall fertility with naïve standardized fecundability and

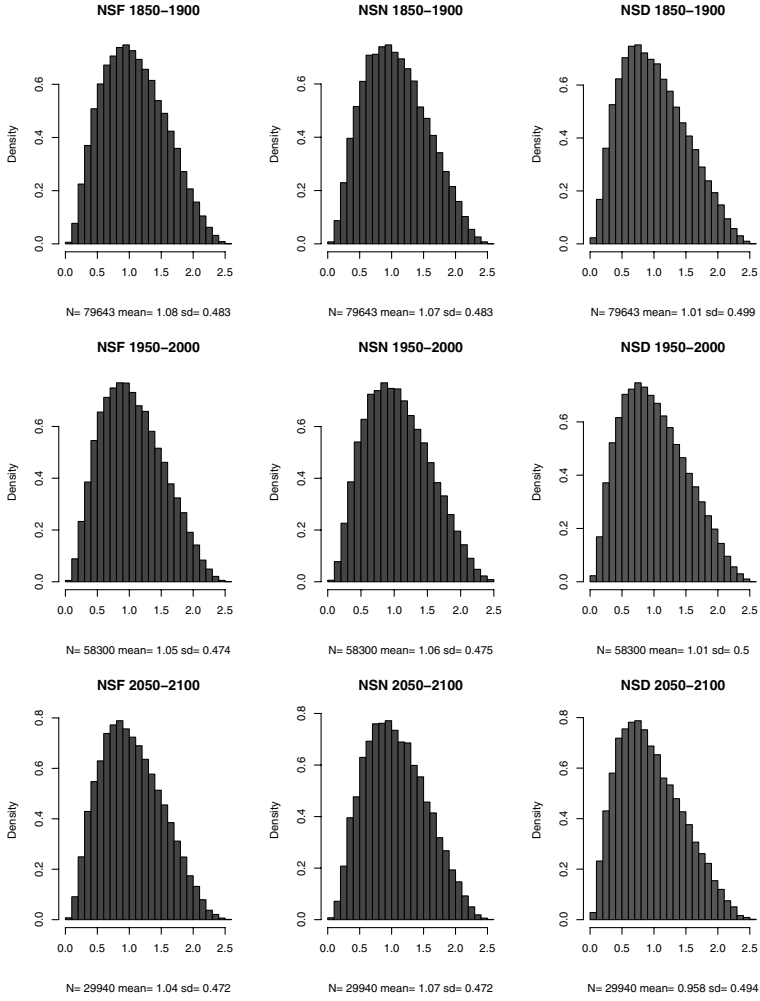


Fig. 2. Distributions of NSF, NSN & NSD by birth cohort

Table 3. Fertility and nuptiality parameters

(a) Mean values									
Year of birth	Children born		NSF		NSN		NSD		
	N-A	A	N-A	A	N-A	A	N-A	A	
1900-1950	1.561	1.631	1.082	1.116	1.079	1.098	1.012	0.988	
1950-2000	1.591	1.684	1.054	1.097	1.060	1.080	1.011	0.986	
2000-2050	1.609	1.684	1.038	1.080	1.072	1.085	0.984	0.969	
2050-2100	0.731	0.746	1.042	1.070	1.069	1.089	0.958	0.955	
(b) Standard deviations									
Year of birth	Children born		NSF		NSN		NSD		
	N-A	A	N-A	A	N-A	A	N-A	A	
1900-1950	1.656	1.692	0.478	0.494	0.477	0.487	0.502	0.499	
1950-2000	1.588	1.631	0.474	0.487	0.475	0.484	0.500	0.499	
1950-2000	1.617	1.653	0.472	0.483	0.476	0.484	0.498	0.495	
2000-2050	1.281	1.283	0.472	0.480	0.472	0.482	0.494	0.493	
(c) Correlation coefficients									
Year of birth	Children & NSF		Children & NSN		Children & NSD				
	N-A	A	N-A	A	N-A	A			
1900-1950	0.198	0.215	0.159	0.164	-0.006	-0.011			
1950-2000	0.196	0.227	0.230	0.233	-0.028	-0.035			
2000-2050	0.215	0.224	0.234	0.241	-0.039	-0.035			
2050-2100	0.118	0.128	0.149	0.145	-0.016	-0.021			

Note: N-A non-assortative; A assortative;
 NSF naïve standardized fecundability;
 NSN naïve standardized nuptiality;
 NSD naïve standardized divorce.

naïve standardized nuptiality are approximately equal, indicating that they have similar effects on achieved fertility.

The mean values for naïve standardized fecundability and naïve standardized nuptiality are about 4% and 2% higher in the assortative marriage group than in the non-assortative one (Table 3(a)). However, the continuation of relatively small differences in fertility with mortality being the same, will lead to the difference in overall population size increasing over time. Thus while the effect of assortative marriage is not substantial in a given period, but if continued over a long period it clearly has a substantial effect (Fig. 3). Because the initial population at 1850 was not assortatively mated, the effect of assortative mating takes time to build up as people form partnerships beyond 1850, and they subsequently have children, so it takes some time for the intergenerational patterns in population size to emerge. Thus while the difference in TFR around 1950 is only about 4% greater in the assortative case, by 2100, this had increased to 8%. By 1950, the assortative population size is only about 2% higher than the non-assortative one, but by 2050, it is 18% larger,

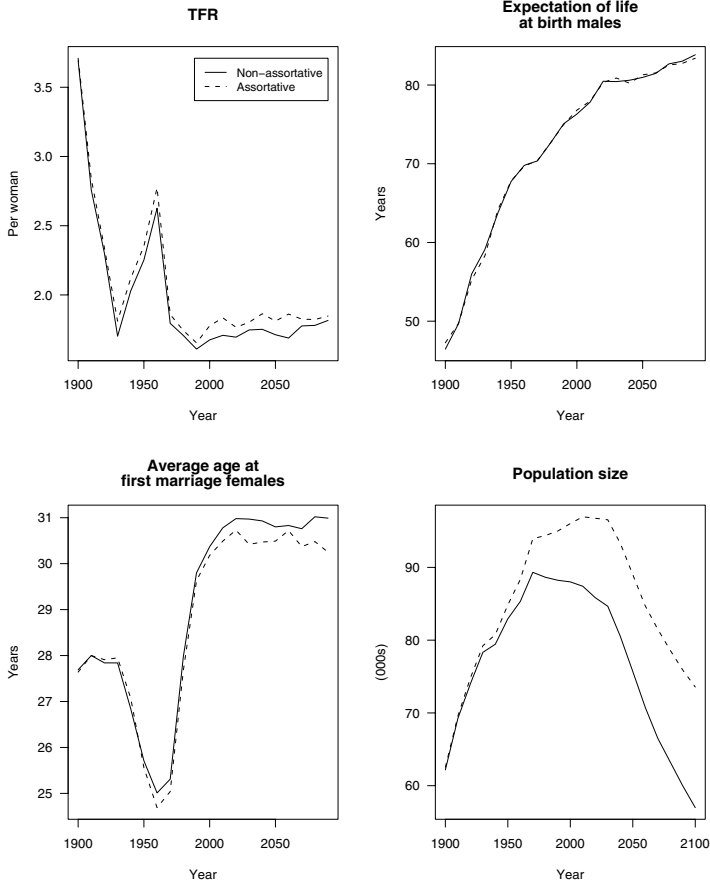


Fig. 3. Population parameters for non-assortative & assortative regimes standard correlation case

and 27% larger by 2100. These effects are much smaller than those associated with the absence or presence of inherited fertility where the population after 250 years was over four times higher in the latter case [41]. The slightly earlier age at first marriage in the assortative case also becomes apparent over time (although the great majority of people’s first partnership in these models in the 21st century is a cohabiting rather than a formal marriage, so this may understate differences in partnership behaviour).

Micro-Level Effects

While Table 3 shows the relationships between different variables for the same individual, Table 4 shows the effect of assortative mating for pairs of people at the end of the period of analysis, for those born between 2050 to 2100 in the principal kin groups of parent/child and partners (the two partners are those of each individual's last partnership, whether intact or not). It also shows the effect of increasing the strength of assortative mating by setting the value λ in (4) to 0 ('High correlation'). Table 4(a) confirms that the correlation coefficients for partners are trivial in the non-assortative case. The small positive values observed for this group for naïve standardized fecundability and naïve standardized nuptiality are probably due to the fact that those with higher values of naïve standardized nuptiality of both sexes will enter the marriage market more quickly on average than those with lower values, so the characteristics of the pool of eligibles will be slightly different from the overall population. However, there is no preference for selection within the pool as in the assortative case, where the magnitude of the correlation coefficients between partners is about 0.1 for naïve standardized fecundability and naïve standardized nuptiality, and rather smaller for naïve standardized divorce in the standard case. The degree of association between parents and children was chosen to be about 0.4 as discussed earlier, but the effect of assortative mating is to increase the correlation by about 10%, from about 0.37 to 0.41 for naïve standardized fecundability and naïve standardized nuptiality (Table 4(b)). Assortative mating has the effect of increasing the correlations between parents and children: this is to be expected since the child's values are positively correlated with both parents' values, and the parents' values are themselves correlated. Since higher correlations lead to higher population growth [41]), this accounts for the larger population size in the assortative case in Fig. 3.

However, analyses so far do not show how large or small are the effect of the strength of the correlation between the spouses on that between parents and children. The alternative high model in Table 4(a) produces a correlation coefficient of about 0.9 between partners, compared with about 0.1 in the standard case. The effect is to increase substantially the correlation between parents and children, from about 0.4 to about 0.7 in the three variables shown (the values remain unaltered in the non-assortative case). In turn, this leads to considerably higher population growth, which substantially increases the differential between the non-assortative and assortative populations, so that the size of the assortative group is three times larger at the end of the period than the non-assortative one, and the differentials in fertility and nuptiality are much larger (as before, mortality differentials are effectively zero as would be expected, Fig. 4). This second case is extreme, and serves to put an upper limit on the magnitude of such effects, but the standard case shows a modest effect, in that the correlations between naïve standardized fecundability and

Table 4. Correlation coefficients with standard and high assortative levels

(a) Women (born 2050-2100) and partners						
Level	NSF		NSN		NSD	
	N-A	A	N-A	A	N-A	A
Standard	0.014	0.099	0.019	0.078	0.007	0.039
High	0.001	0.902	0.033	0.902	0.039	0.912
(b) Children (born 2050-2100) and parents						
Level	NSF		NSN		NSD	
	N-A	A	N-A	A	N-A	A
Standard	0.372	0.408	0.377	0.406	0.388	0.400
High	0.379	0.669	0.381	0.679	0.397	0.721

Note: N-A non-assortative; A assortative;
 NSF naïve standardized fecundability;
 NSN naïve standardized nuptiality;
 NSD naïve standardized divorce.

naïve standardized nuptiality, and achieved family size is only about 0.02, so that the correlation between sibship sizes of partners is considerably smaller than the values shown in Table 1, and therefore the standard case results may be regarded as a lower bound on the magnitude of the effect of assortative mating on population growth.

4 Conclusions

There are three main points that arise from these findings. The first is that the effect of persistent assortative mating in conjunction with intergenerational transmission of fertility has a clear effect in making the population larger than would otherwise be the case. This is in line with theoretical expectations about the positive effect of assortative mating on fitness. The second is that assortative mating does have an effect that should not be ignored in demographic and genetic studies since part of the observed relationship in fertility of mothers and children is due to that via the father and their children, which is similar in magnitude [37] and therefore contributes to explanation, whereas the great majority of studies consider neither fathers nor the role of assortative mating. The third is that the whole area of partner selection remains a relatively under-developed area. Models that require interaction and adjustment between agents are more complex than ones that provide a series of linked sequential actions. Progress will require developments in theory, data, modelling and technology, but assortative mating remains one of the most persistent and enduring features not only of humans, but other species as well.

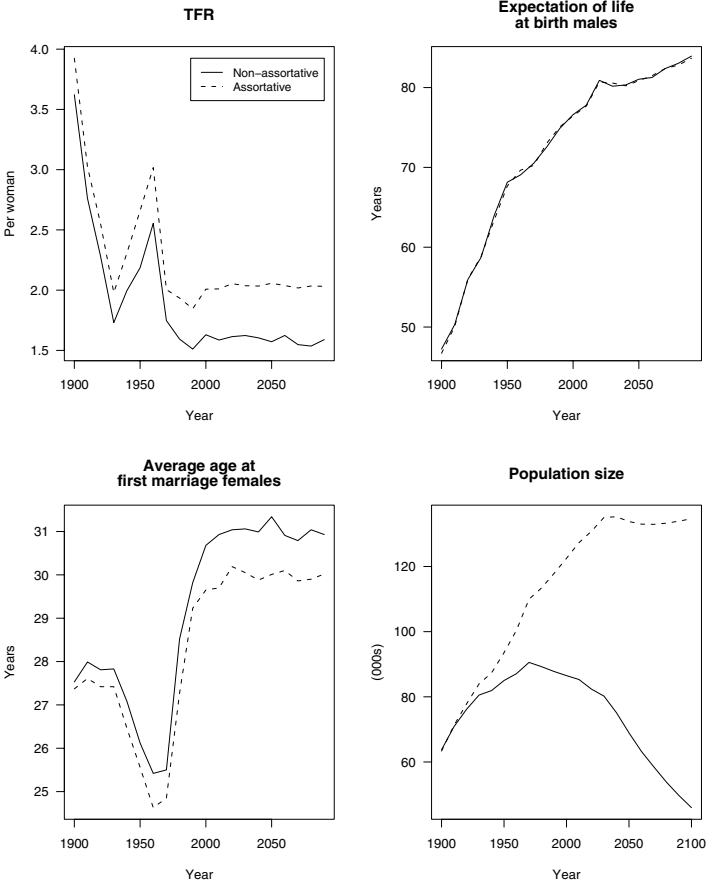


Fig. 4. Population parameters for non-assortative & assortative regimes high correlation case

Acknowledgements

Thanks are due to the Office for National Statistics for access to the 1976 Family Formation Survey.

References

1. Barrett, J.C. (1971) A Monte Carlo simulation of reproduction. In: W. Brass (Ed.), *Biological aspects of demography*. Symposia of the Society for the Study of Human Biology Vol.10. Taylor and Francis, London, Pp. 11-30.
2. Barrett, L., Dunbar, R., Lycett, J. (2002) *Human evolutionary psychology*. Princeton University Press, Princeton N.J..
3. Bateson, P. (ed) (1983) *Mate Selection*. Cambridge University Press, Cambridge.
4. Bateman, A.J. (1948) Intra-sexual selection in drosophila. *Heredity* 2, 349-368.
5. Blossfeld, H.-P., Timm, A. (Eds) (2003) *Who Marries Whom: Educational Systems as Marriage Markets in Modern Societies*. Kluwer Academic Publishers, Dordrecht.
6. Bongaarts, J., Potter, R.G. (1983) *Fertility, biology, and behavior: An analysis of the proximate determinants*. Academic Press, New York.
7. Bouffard, N., Easther, R., Johnson, T., Morrison, R.J., Vink, J. (2001) Matchmaker, Matchmaker, Make Me a Match. *Brazilian Electronic Journal of Economics* 4(2).
8. Boyce M.S. (1990) The Red Queen visits sage grouse leks. *American Zoologist*, 30, 263-270.
9. Buss, D.M. (1985) Human mate selection. *American Scientist* 73, 47-51.
10. Buss, D.M. (1999) *Evolutionary psychology: the new science of the mind*. Allyn and Bacon, Boston.
11. Dunne, M.P., Martin, N.G., Staham, D.J., Sltske, S.W., Dinwiddie, S.H., Bucholz, K.K., Madden, P.A., Heath, A.C. (1997) Genetic and environmental contributions to variance in age at first sexual intercourse. *Psychological Science*, 8, 211-216.
12. Dunnell, K. (1979) *Family Formation 1976*. H.M.S.O., London.
13. Epstein, E., Guttman, R. (1984) Mate selection in man: Evidence, theory, and outcome. *Social Biology* 31, 243-278.
14. Feng, D., Baker, L. (1994) Spouse similarity in attitudes, personality, and psychological well-being. *Behavioral Genetics* 24, 357-364.
15. Fisher, R.A. (1930) *The genetical theory of natural selection*. Clarendon Press, Oxford.
16. Foster, C. (2000) The limits to low fertility: a biosocial approach. *Population and Development Review*, 26, 209-234.
17. Gini, C. (1924) Premières recherches sur la fécondabilité de la femme. *Proceedings of Mathematical Congress*, Toronto, 889-92.
18. Hamilton, W.D. (1964) The genetical evolution of social behavior, I and II. *Journal of Theoretical Biology* 7, 1-52.
19. Hamilton, W.D., Zuk, M. (1982) Heritable true fitness and bright birds: a role for parasites? *Science* 218, 384-387.
20. Harrison, G.H., Gibson, J.B., Hiorns, R.W. (1976) Assortative mating for psychometric, personality and anthropometric variation in a group of Oxfordshire villages. *Journal of Biosocial Science*. 8, 145-153.
21. Hammel, E.A., Mason, C., Wachter, K.W. (1990) SOCSIM II, a sociodemographic microsimulation program, rev. 1.0, operating manual: Graduate Group in Demography Working Paper No. 29. Berkeley, Calif: University of California, Institute of International Studies, Program in Population Research.

22. Hawkes, K., O'Connell, J.F., Blurton Jones, N.G., Alvarez, H., Charnov, E.L. (1998) Grandmothering, menopause, and the evolution of human life histories. *Proceedings of the National Academy of Science, USA* 95, 1336-1339.
23. Houle, D. (1998) How should we explain variation in the genetic variabilities of traits? *Genetica* 102/103, 241-253.
24. Jacquard A. (1983) Heritability: one word, three concepts. *Biometrics* 39, 465-477.
25. Johnson, R.A. (1980) *Religious Assortative Marriage in the United States*. Academic Press, New York.
26. Kalmijn, M. (1994) Assortative Mating by Cultural and Economic Status. *American Journal of Sociology* 100, 422-452
27. Keyfitz, N. (2005) *Applied Mathematical Demography*. (3rd ed.). Springer-Verlag, New York.
28. Kirk, K.M., Blomberg, S.P., Duffy, D.L., Heath, A.C., Owens, I.P.F, Martin, N.G. (2001) Natural selection and quantitative genetics of life-history traits in Western women: a twin study. *Evolution*, 55, 423-435.
29. Kohler, H.-P., Rodgers, J.L., Christensen, K. (1999) Is fertility behaviour in our genes? Findings from a Danish twin study. *Population and Development Review*, 25, 253-288.
30. Kondrashov, A. (1988) Deleterious mutations and the evolution of sexual reproduction. *Nature* 336, 435-440.
31. Lahdenperä, M., Lummaa, V., Helle, S., Tremblay, M., Russell, A.F. (2004) Fitness benefits of prolonged post-reproductive lifespan in women. *Nature*, 428, 178-181.
32. Leridon, H. (1977) *Human Fertility: The Basic Components* (translated by J. F. Helzner). University of Chicago Press, Chicago.
33. Levine, D. (1982) 'For their own reasons': Individual marriage decisions and family life. *Journal of Family History*, 17, 255-264.
34. McGue, M., Lykken, D.T. (1992) Genetic influence on risk of divorce. *Psychological Science*, 3, 368-373.
35. Miller, G. (2001) *The Mating Mind: How Sexual Choice Shaped the Evolution of Human Nature*. Vintage, London.
36. Mueller, U. (2001) Is there a stabilizing selection around average fertility in modern human populations? *Population and Development Review*, 27(3), 469-498.
37. Murphy, M. (1999) Is the relationship between fertility of parents and children really weak? *Social Biology*, 46, 122-145.
38. Murphy, M. (2001) Family and kinship networks in the context of aging societies. Paper prepared for the Conference on Population Ageing in the Industrialized Countries: Challenges and Responses organised by the Committee on Population Age Structures and Public Policy of the International Union for the Scientific Study of Population (IUSSP) and the Nihon University Population Research Institute (NUPRI), Tokyo, Japan, 19-21 March 2001.
39. Murphy, M. (2003) Bringing behaviour back into micro-simulation: Feedback mechanisms in demographic models. In: F. Billari and A. Prskawetz (Eds) *Agent-Based Computational Demography: Using Simulation to Improve our Understanding of Demographic Behaviour*. Physica-Verlag, Heidelberg, Pp 159-174 .
40. Murphy, M., Wang, D. (2001) Family-level continuities in childbearing in low-fertility societies. *European Journal of Population*, 17, 75-96.

41. Murphy, M., Wang, D. (2002) The impact of intergenerationally-transmitted fertility and nuptiality on population dynamics in contemporary populations. In: J. Rodgers and H.-P. Kohler (Eds.) *Biodemography of Human Reproduction and Fertility*. Kluwer Academic Publishers, Boston, Pp. 209-228 .
42. Nagoshi, C.T., Johnson, R.C., Ahern, F.M. (1987) Phenotypic assortment vs social homogamy among Japanese and Chinese parents in the Hawaii Family Study of Cognition. *Behavioral Genetics* 17, 477-485.
43. National Statistics and Government Actuary's Department (2004) *National Population Projections 2002-based*. Series PP2 No. 24. T.S.O., London. (also available at http://www.gad.gov.uk/Publications/Demography_and_statistics.htm#ms)
44. Ní Bhrolcháin, M. (2000) Age difference asymmetry and a two-sex perspective, *European Journal of Population*, 8, 23-45
45. Ní Bhrolcháin, M. (2001) Flexibility in the marriage market. *Population: An English Selection* 13(2), 9-48.
46. Pomiankowski A., Moller, A. (1995) A resolution of the lek paradox. *Proceedings of the Royal Society (London) B*, 260, 21-29.
47. Preston S.H. (1976) Family sizes of children and family sizes of women. *Demography* 13, 105-114.
48. Registrar General (1923) *1911 Census of England and Wales*, Vol. XIII. *Fertility of marriages Part II*. H.M.S.O., London.
49. Rodgers, J.L., Doughty, D. (2000) Genetic and environmental influences on fertility expectations and outcomes using NLSY kinship data. In: J.L. Rodgers, D.C. Rowe, and W.B. Miller (Eds.) *Genetic Influences on Human Fertility and Sexuality*. Kluwer, Boston, Pp. 85-105.
50. Rushton, J.P., Russell, R.J.H., Wells, P.A. (1984) Genetic similarity theory: beyond kin selection. *Behavioral Genetics* 14, 179-193.
51. Sear, R., Mace, R., McGregor, I.A. (2000) Maternal grandmothers improve the nutritional status and survival of children in rural Gambia. *Proceedings of the Royal Society of London B* 267, 1641-1647.
52. Sear, R., Mace, R., McGregor, I.A. (2003) The effects of kin on female fertility in rural Gambia. *Evolution and Human Behaviour* 24, 25-42.
53. Smith J.E. (1987) Simulation of kin sets and kin counts. In: J. Bongaarts, T. Burch, K. Wachter (Eds.) *Family Demography: Methods and Their Application*. Clarendon Press, Oxford, Pp 249-266.
54. Sulloway, F. (1996) *Born to Rebel*. Little Brown and Co, London.
55. Tambs, K., Sundet, J.M., Berg, K. (1993) Correlations between identical twins and their spouses suggest social homogamy for intelligence in Norway. *Personal and Individual Differences* 14, 279-281.
56. Thiessen, D., Gregg, B. (1980) Human assortative mating and genetic equilibrium: an evolutionary perspective. *Ethology and Sociobiology* 1, 111-140.
57. Thornhill, N.W. (1991) An evolutionary analysis of rules regulating human inbreeding and marriage. *Behavioural and Brain Sciences* 14, 247-293.
58. Thornton, A. (1980) The influence of first generation fertility and economic status on second generation fertility. *Population and Environment*, 3, 51-72.
59. van Imhoff, E. Post, W. (1998) *Microsimulation Methods for Population Projection*. *Population: An English Selection* 10(1), 97-138.
60. Wachter, K.W. (1987) *Microsimulation of household cycles*. In: J. Bongaarts, T. Burch, K. Wachter (Eds.) *Family Demography: Methods and Their Application*. Clarendon Press, Oxford, Pp 215-227.

61. Wachter K.W. (n.d.) Socsim Technical Documentation, available at <http://www.demog.berkeley.edu/>
62. Wolf, D.A. (1994) The Elderly and Their Kin: Patterns of Availability and Access. In: L.G. Martin, S.H. Preston (Eds.) Demography of Aging. National Academy Press, Washington D.C., Pp. 146-194.
63. Wood, J.W., Weinstein, M. (1988) A model of age-specific fecundability. *Population Studies*, 42, 85-113.
64. Zhao, Z.W. (1996) The demographic transition in Victorian England and changes in English kinship networks. *Continuity and Change* 11(2), 243-272.

An Agent-Based Simulation Model of Age-at-Marriage Norms

Belinda Aparicio Diaz¹ and Thomas Fent^{2*}

¹ Vienna Institute of Demography, Austrian Academy of Sciences, Austria
`belinda.aparicio.diaz@oeaw.ac.at`

² ETH Zurich, Switzerland and Vienna Institute of Demography, Austrian Academy of Sciences, Austria
`Thomas.Fent@oeaw.ac.at`

Summary. In this Chap. we analyse an agent-based model designed to understand the dynamics of the intergenerational transmission of age-at-marriage norms. A norm in this context is an acceptable age interval to get married. We assume that this age-interval is defined at the individual level and the individuals' age-at-marriage norms are transmitted from parents to their children. We compare four different transmission mechanisms to investigate the long term persistence or disappearance of norms under different regimes of transmission. Our work is an extension of [4] that introduces a one-sex non-overlapping-generations version of an age-at-marriage model. Here we investigate whether their results also hold in a more complex setup. Therefore, we explicitly take into account heterogeneity with respect to age and sex. Moreover, we also include the timing of union formation and fertility into our model. To create a more realistic model of the evolution of age norms the characteristics of the agents are extended, some new parameters are added to the model and the age-at-marriage norms are split into two sex-specific age-at-marriage norms. A comparison of the results with those of the original model gives information about how additional characteristics and new parameters can influence the evolution of age-at-marriage norms.

1 Introduction

In this Chap. we present an agent-based model to simulate the evolution of age-at-marriage norms. While some individual decisions are mostly influenced by economic incentives and other decisions are mostly driven by social norms there can also be decisions depending on both. We postulate that individuals' choices regarding their age at first marriage is at least partially influenced

* We are grateful for comments and suggestions by Francesco Billari, Rosaria Conte, Johannes Fürnkranz, Alexia Prskawetz, Wojtek Przepiorka, and two anonymous referees.

by normative guidelines enforced by the society and put the emphasis of our model onto that aspect of age-at-marriage.

Social norms can only be effectively put into practice if there exist sanctions that punish deviant behaviour. In a densely connected society there exists a versatile bundle of sanctions to put norms into effect such as ostracism, physical retaliation, refusal of social approval, gossip etc. Diekmann and Voss [6] show that rational actors are able to enforce social norms with sanctions even in one-shot situations. However, the existence of a social network is a prerequisite for the successful implementation of norms.

Normative guidelines generally are a decision guidance whenever an individual has to decide about something important. Thus certain actions are influenced by social norms, namely social rules that state how individuals ought to behave in certain circumstances. Many papers address the presence of such social norms. For instance [11] shows that the presence of family-size norms can explain diverse experiences in income and population growth.

Although one might think that modernisation processes may have weakened traditional normative pressure, the effect of social norms may have been internalised in western societies rendering obsolete any need for external societal enforcement of social norms [9]. In post-industrial societies there seems to be a trend towards a diminishing normative regulation of schedules. A decreasing impact of social norms in the transition to adulthood may be related to a decreasing dependence from the traditional references to the family and to the church [5]. Nevertheless, according to [3] new types of norms may have substituted the old ones.

The impact of social norms on shaping individuals' lives has been addressed by [2]. They accomplish a theory-based empirical analysis of cross-sectional survey data on norms and sanctions concerning sexual life and marriage for young Italian university students. The survey shows the existence of lower and upper age limits on sexual debut and first marriage. Moreover, there exist perceived norms and sanctions connected for instance to the experience of some types of sexual behaviour. Social norms are supposed to be enforced by formal and informal sanctions. Their investigation exhibits strong evidence that sexual behaviour is subject to strong sanctions and that sexual behaviour is highly affected by social norms. Consequently, it seems reasonable to assume that there may be a corresponding normative pressure to adhere to the norm.

Our simulation model is motivated by [4] who introduce a more stylized model using a slightly simpler implementation of norms. The age-at-marriage norms serve as guidelines for individuals to take decisions about the right point in time to get married. Agent-based simulations are frequently applied in the social sciences, since they have proved to be a valuable tool to study the complex dynamics evolving from heterogenous populations. Here they are applied to observe the long-term persistence or dissolution of social norms and to investigate their evolution over time. Within the artificial environments which can be seen as small laboratories it is possible to simulate behaviours that are influenced by such social norms.

The remainder of this paper is organized as follows. Sect. 2 provides empirical evidence regarding the past development of age-at-marriage. Sect. 3 briefly summarizes the agent-based model of [4], who studied the evolution of age-at-marriage norms, their long term persistence or disappearance, the long term impact of the initial distribution of norms in a population, and the impact of random mutations. In Sect. 4 we describe our extended model, the details of the implementation are provided in Sect. 5, and Sect. 6 highlights the numerical results obtained. The concluding Sect. 7 draws a summary of the main results and elaborates on the following questions: Do the missing characteristics influence the results in an important way or does the simpler model serve as a description that is close enough to reality to detect the evolution of age-at-marriage norms? Does this extension provide a step forward to a closer approximation to the real world?

2 Empirics

Hajnal [8] describes two basic marriage patterns, a traditional or non-European pattern of early and universal marriage reflecting the typical behaviour in most of the developing regions and a European pattern of late marriage and high proportions of individuals who never get married characterizing Western European Societies.

Dixon [7] investigates timing and nuptiality in 57 countries according to censuses taken around 1960. In particular she looks at the proportions of men and women not being married at age 20–24 and at the corresponding proportions at age 40–44. First marriages after the age of 44 are not taken into consideration since they occur only rarely and are demographically of little impact. The data show that grooms are older than their brides in all societies. The main difference between the European and the traditional marriage pattern, however, is partially due to the fact that in societies where marriage occurs early, more people marry ultimately than in societies where marriage occurs relatively late. Dixon investigates in particular the *availability of mates*, determined by the sex ratio of persons of marriageable age and by the method of mate selection, the *feasibility of marriage*, determined by expectations with respect to financial and residential independence and the available resources, and the *desirability of marriage*, indicating the strength of the motivation to marry and depending on the available social and institutional alternatives to marriage and childbearing. When looking at the desirability, it is important not only to take the availability of alternatives into consideration but also whether these alternatives are considered rewarding. Dixon states that *the pressure toward marriage and the penalties of remaining single vary in kind and degree, and differ for men and women*. Thus, empirical investigation should also assess the penalties associated with marrying late or never and childlessness like social isolation, stigma, and the loss of economic and social opportunities. This qualifies the desirability of marriage to be an indicator

for the relevance of social norms on the decision whether and when to get married. The data investigated by Dixon reveal that *delayed marriage and celibacy are most highly correlated with indicators of the desirability of marriage, less so with feasibility, and least with availability*. Moreover, she arrives at the conclusions that *the degree of social isolation and stigma that bachelors and spinsters experience depends on the level of celibacy in each society and in those countries where romanticism provides a primary motivation for marriage, the unmarried person still experiences the discomfort of being visibly 'unwanted' in a society that idealizes personal attractiveness and individual happiness*.

Bhrolcháin [1] examines age preference data for measuring recent levels of partner availability in England, Wales, and the USA and for assessing time trends of partner supply in those countries. The data reveal that mean age differences in England and Wales do not exhibit a long-run secular trend driven by social and cultural change but rather fluctuate during the 20th century. During the same time period mean age differences have varied within a relatively narrow range in the USA, where a long-run but modest change resulted in a decline from an average of around four years around 1900 to 2.4 years in 1990.

Current empirical data regarding female first marriage in Europe reveal that there is no monotone trend in age-at-marriage over a long period but rather a turning point between two opposed trends. Figure 1 illustrates the mean age of women at first marriage for five European countries. It shows that from the birth cohorts in the 1930s to the birth cohorts in the 1940s the mean age at first marriage decreased by about one year and after that it started to increase to levels already higher than at the beginning of the observation. Moreover, the time series suggest that this increase has not finished yet.

Figure 2 shows the rate of first marriage per 1000 females in Italy by 5-year age-groups. In compliance with the previous graph it turns out that for younger women (< 20 and $20 - 24$) the rate of first marriage slightly increased between the 1960s and the 1970s while it declined afterwards. Furthermore, the rate of first marriage of the women aged 25 to 34 decreased between 1960 and 1975 and increased later on. This observation motivates that the lower age limit of first marriage underwent an increase from 1960 to the 1970s and a decrease later on. The picture is not so clearcut with respect to the upper age limit since the rate of first marriage simply fades away for higher ages. However, Settersten and Hägestad [12] investigated a survey of individuals belonging to different age groups in Chicago. Their analysis revealed that 82.3% of the respondents perceived an age deadline for marriage, i.e. an upper age after which it would not be appropriate to get married.




Fig. 1. Mean age of women at first marriage (below age 50), Source: Council of Europe, Demographic Year Book, 2004 Edition




Fig. 2. Sum, by five-year age-groups, of first marriage rates per 1000 females in Italy, Source: Council of Europe, Demographic Year Book, 2004 Edition

3 The One-Sex Model

To show the long term persistence of norms, a population of fixed size is generated, where several individuals are characterized by their age-at-marriage norm. This norm is an age interval, which describes the age at which an individual can marry. The agent's age at a certain time is not a relevant characterization in this setup and is therefore ignored. The absence of additional characterizations is one of the properties to be modified in the following Sect.

Starting with an initial generation with randomly generated age-at-marriage norms, the evolution of the experimental population is simulated. Within this model, the agents do not grow older — they are born, can marry each other, can reproduce if they find a partner, and they die — in the course of one generation. Thus, each time step represents one generation. Individuals can only marry another individual with compatible norms. The norms of two individuals are compatible if they overlap — i.e. if their intersection is nonempty. Other criteria, like age and sex do not matter. Only married individuals are allowed to reproduce, but the reproduction isn't connected with any other criteria like the current age or the duration of the marriage. The independence between age at marriage and the number of children a couple can have is another restriction of this model which will be relaxed in the following Sect. While searching for a partner each individual tries to find someone whose age interval overlaps with the own one. If an agent can't find an acceptable partner it remains single and isn't taken into consideration any longer. Otherwise both partners marry. Married individuals are removed from the list of marriage candidates.

The population within this model is stationary. This is achieved by assigning $\min(\lfloor s/c \rfloor, k)$ children to each married couple, where the parameter s means the size of the starting generation, c is the number of couples and $k \geq 0$ is a numerical parameter determining the minimal number of children

a couple could have. After that replenishment of the population is achieved by assigning further children to the couples until the original population size is attained. The case k equal to zero means there is no minimum number of children and therefore it is possible that some couples remain childless.

If a couple has children, their children inherit their age-at-marriage norms by means of a special transmission mechanism. Four different transmission mechanisms are applied,

- Intersection,
- Union,
- Random, and
- Uniform.

These mechanisms are also adopted in the extended model and are described in detail in the following Sect. Combinations of these mechanisms are also allowed. This is achieved by assigning to each individual one out of four mechanisms with the same probability. In this case children inherit both, the age norms and the transmission procedure, from their parents. In addition to the transmission of norms, two alternative forms of mutations are allowed. Thus, with a certain user-defined probability a child does not necessarily inherit the transmission mechanism or the age norm from its parents. In the former case the children are initialised randomly with the same method applied for the initial generation, in the latter case the lower and upper bounds of the child's age norm are calculated as the average of the lower and upper bounds of the parent generation. These two mutation techniques are not adopted to the new model in an attempt to keep the number of degrees of freedom at a tractable level.

4 The Extended Model

This Model is an extension of [4] which is described in brief in Sect. 3. It is designed to study the cultural evolution of age-at-marriage norms. The model is a system in which agents interact in a dynamic and evolving way. The agents search for a partner, marry, and reproduce. The existence of norms implies that marriage only takes place within a particular age interval. As these age norms prevent marriage outside of the personal age interval, they influence the demographic choices of individuals. The norms serve as a guideline for the timing of marriage and for choosing an acceptable partner. Because these norms restrict the individual life course choices, they are important for investigating the further deployment of the life course. The model was developed for simulating these dynamic age norms. In particular the long term persistence and the disappearance of age norms are examined. Further, the impact of the initial distribution of norms within the population is depicted. The long run persistence — i.e. the survival across several generations — of age norms can be investigated by means of agent-based modelling. Agent-based

models allow us to study the evolution of norms within setups determined by co-existence of norms.

The previously described model is now extended by adding the demographic structure characteristics age and sex. For this model a starting population of N agents is produced. The agents obtain a starting age between zero and the maximum age m , which is assigned randomly. The sex of the individuals is also chosen randomly. The sex ratio at birth resp. at initialization, srb , which means the ratio of male to female births, can be chosen arbitrarily.

Besides age and sex the individuals are characterized by two sex-specific age-at-marriage norms. The female age-at-marriage norm determines the acceptable age for a woman to marry. Therefore, a female individual recognizes her own marriage readiness by this interval whereas for a male individual it indicates the age his potential wife should have. The male age-at-marriage norm on the other hand determines the acceptable age for men to marry. Consequently, the male agents consider this norm to determine their own readiness while it determines the age range of an acceptable partner from the viewpoint of a female agent. It seems to be appealing to extend the length of the marriage intervals with age since the tolerance of an individual with respect to age differences may expand with increasing age. However, the age-norms in our model are not intended to determine the age differences within unions but to constitute a regulating mechanism with respect to age-at-marriage. Thus, expanding the marriage interval with age means to mess up two different processes. Therefore, we decided not to include this extension into our model. Moreover, in the simulation experiments this extension did not exhibit a significant impact on the results, it simply delayed some of the observed effects.

Each norm is represented by a lower acceptable age-at-marriage, l , and an upper acceptable age-at-marriage, u . The individual lower bound must be above the global minimum age l_a , which may be interpreted as a legal minimum age to get married. The upper bound of the individuals age-at-marriage intervals is only restricted by the age m at which agents are being removed from the simulation. Thus, the lower limit is situated between l_a (which is the minimum age for marriage) and m , and the according upper age limit is set between the lower limit and m .

The norms for the initial population are drawn from a random distribution. First a random number l_i^f satisfying $l_a \leq l_i^f \leq m$ is selected as the lower bound. Then a random number, u_i^f , which must be between this lower bound and m is chosen as the individual's upper bound. These two values describe the female age-at-marriage norm. After fixing the female norm the male age-at-marriage interval is determined according to the same procedure.³

³ This procedure does not generate a uniform distribution for both lower and upper bound. Only the lower bounds are uniformly distributed while the upper bounds are biased towards higher ages. However, Figs. 7 and 8 indicate that all possible norms occur in the initial population. As long as the population size is sufficiently

On the basis of this starting population, the evolution of this population is simulated, where each individual ages, may marry and get children. Since a simulation sequence now represents one year, the individuals grow one year older within one simulation step. The maximum age which an individual achieves is m years. As soon as an agent gets m years old, it is discarded from the model, since it has no influence on the dissemination of age-at-marriage norms due to age-specific fertility rates. For simplicity we neglect age-specific mortality rates and assume that all agents survive until the age m .

In every time step each individual who is in the marriageable age may search for an acceptable partner. An individual arrives at the marriageable age when its own age is situated within its appropriate sex-specific age interval. A potential partner is a marriageable single individual of the other sex, whose sex-specific age-at-marriage norms overlap with the agent's own norms. An unmarried female at marriageable age would search for any male single individual whose current age is within her male age-at-marriage interval and whose female and male age-at-marriage norm have nonempty intersections with her own female and male age-at-marriage norms (see Fig. 3).

Moreover, the chosen potential partner would only accept a partnership if her current age is also within his female age-at-marriage interval.




Fig. 3. Age intervals of a potential partner

Additionally, the potential partners must not have the same parents since marriage among siblings is prohibited in our simulation. This restriction is useful to avoid the persistence of a weak norm which is transmitted by only one couple. Given the population is sufficiently large the result of the simulation described so far would be that most individuals find an acceptable partner as

large this bias in the initial population will not have any significant impact on the results.

soon as they enter their individual age-at-marriage interval. Moreover, each agent would marry the very first partner encountered who is acceptable with respect to age-at-marriage norms. This is obviously not the way how dating and marriage happen in the real world. Besides the age of the potential partner there will also be criteria like sympathy, physical attraction, or social and economic status which play a role in mate choice of within human populations. These phenomena are discussed for instance by [13] and [15]. To overcome the problem of a marriage peak just after surpassing the lower age limit [13] and [15] applied a normally distributed courtship time in their model and [14] introduced a variation in the number of dates to be encountered during adolescence. Our solution is in line with these models. Here, every agent who finds an acceptable partner gets married with probability pm given by

$$pm = pm_0 + (1 - pm_0) \frac{a - l}{u - l}, \quad (1)$$

where $a - l$ is the number of years since the individual has reached the marriageable age and $u - l$ is the length of the personal age interval. The use of pm allows for an individual to marry as soon as she/he reaches the marriageable age with a certain probability but also to wait after entering the marriageable age-interval (e.g. someone who is sure of having found her/his partner would marry immediately whereas others might rather wait for a “better” partner). An individual who doesn’t find an acceptable partner or decides not to marry remains single for that period and continues to search for a partner in the next period if she/he is then still at marriageable age.

However, if female and male age norms of two marriageable individuals match, this couple may eventually get married. In that case these two agents are no longer potential partners for others. Each married couple can have children annually. In our simulation model the probability for a married woman at age a to give birth is

$$\frac{w(a)}{mw(a)} af(a) tfr, \quad (2)$$

where tfr is an adjustable parameter determining the period total fertility rate within the population. This parameter is multiplied by $af(a)$ to replicate empirically observed age-specific fertility patterns. Consequently, $af(a) tfr$ would represent the age-specific fertility of all female agents at age a within the population. Moreover, since children are assigned only among the married couples this age-specific fertility rate is multiplied by $w(a)/mw(a)$ where $w(a)$ is the total number of women at age a and $mw(a)$ is the number of married women at that age. Therefore, the fraction $w(a)/mw(a)$ must be greater or equal to one which causes the age-specific fertility of the married women to be greater or equal than the age-specific fertility of the whole female population. Each newborn inherits the conceptions concerning the marriage age of its parents due to a special transmission mechanism. In this model we apply the four transmission techniques introduced in [4].

Intersection: The child’s age norms $[l_c^f, u_c^f]$ and $[l_c^m, u_c^m]$ result from the intersection of its parent’s intervals, $l_c^f = \max(l_{p1}^f, l_{p2}^f)$, $u_c^f = \min(u_{p1}^f, u_{p2}^f)$, $l_c^m = \max(l_{p1}^m, l_{p2}^m)$, and $u_c^m = \min(u_{p1}^m, u_{p2}^m)$ (see Fig. 4).

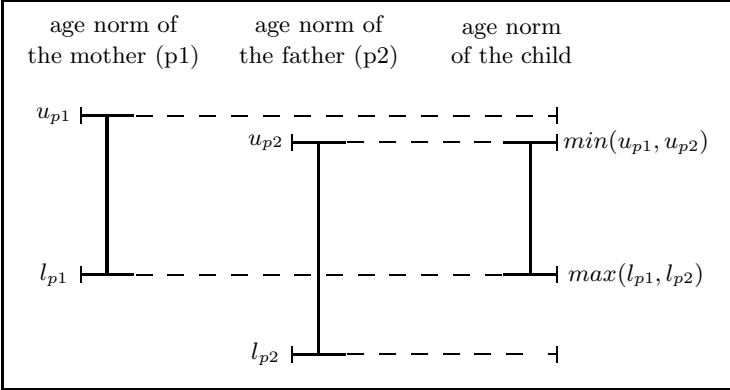


Fig. 4. Intersection of age intervals

Union: The age-at-marriage norms of the child are in each case the unions of the parents’ age intervals, $l_c^f = \min(l_{p1}^f, l_{p2}^f)$, $u_c^f = \max(u_{p1}^f, u_{p2}^f)$, $l_c^m = \min(l_{p1}^m, l_{p2}^m)$, and $u_c^m = \max(u_{p1}^m, u_{p2}^m)$ (see Fig. 5).

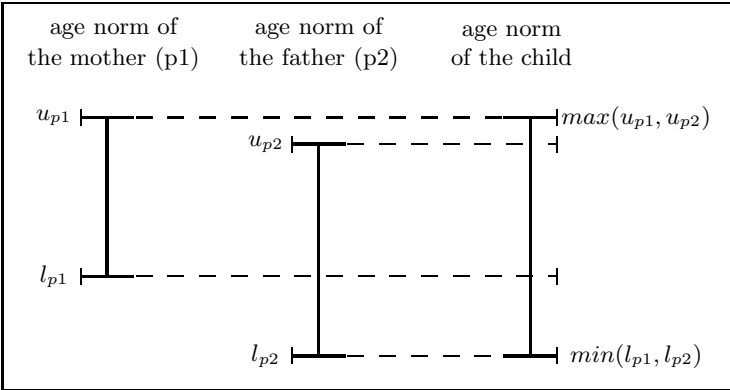


Fig. 5. Union of age intervals

Random: The boundaries of the female norm of the child are selected randomly from the respective boundaries of the parents’ female norms. Thus, the lower bound of the child may be either the lower bound of the mother or the lower bound of the father. The upper bound of the female norm as well as the lower and upper bound of the male norm are selected the same

way, $l_c^f = \text{random}(l_{p1}^f, l_{p2}^f)$, $u_c^f = \text{random}(u_{p1}^f, u_{p2}^f)$, $l_c^m = \text{random}(l_{p1}^m, l_{p2}^m)$, and $u_c^m = \text{random}(u_{p1}^m, u_{p2}^m)$ where $\text{random}(x, y)$ chooses either x or y with the same probability (see Fig. 6).




Fig. 6. One possibility for Random combiner. The upper limit is randomly chosen from (u_{p1}, u_{p2}) and the lower limit is chosen from (l_{p1}, l_{p2}) .

Uniform: The lower (upper) bound of one of the two norms of the child is a random number between the lower (upper) bound of the respective norm of the mother and the father⁴, $l_c^f = \text{uniform}(l_{p1}^f, l_{p2}^f)$, $u_c^f = \text{uniform}(u_{p1}^f, u_{p2}^f)$, $l_c^m = \text{uniform}(l_{p1}^m, l_{p2}^m)$, and $u_c^m = \text{uniform}(u_{p1}^m, u_{p2}^m)$ where $\text{uniform}(x, y)$ selects a number between x and y drawn from a uniform distribution (see Fig. 7). Similar mechanisms are used for instance by [10] and [16] to model opinion dynamics within an agent population. While [10] uses a weighted average of an agents current own opinion and the opinions of the other agents to get the agents opinion in the following period, in [16] only two agents communicate with each other and agree to a compromise by adjusting their own opinion slightly towards the opinion of the other agents. Here, the age-at-marriage norm takes over the role of an opinion and the norm of the child is a compromise of the parents norms. Unlike [10] and [16] the particular location of that compromise is not deterministic but results from a random process.

In Sect. 6 we will present results obtained in simulations with homogenous populations — i.e. populations of agents endowed with the same transmission mechanisms — as well as results gained from heterogenous populations. In the latter case the assignment of combiners to the original population is done randomly with each mechanism being chosen with equal probability. So if the user chooses two mechanisms, these are assigned to the individuals with a

⁴ The Uniform transmission is a random transmission with uniform distribution. This notation is consistent with [4].




Fig. 7. Uniform combiner

probability of 50 percent. If all four techniques are to be used, these four are divided among the individuals with a probability of 0.25 each. Individuals who are born during the simulation inherit the transmission technique of one parent, where the probability to inherit from the mother is just the same as to inherit from the father. The one mechanism that is inherited to the child is also the one that is used to compute the child's age norms from its parents' male and female norms.

The extensions described above allow us to demonstrate within the simulation that the long term persistence of norms depends not only on the transmission mechanism. Thus other female norms will persist than male norms since the age at marriage of a woman considerably influences the number of children she can give birth and consequently the possibility of passing on her age norms. This natural fact is simulated by the consideration of age specific fertility rates⁵ (see Table 1).

Table 1. Age specific fertility rates

	15-19	20-24	25-29	30-34	35-39	40-44	45-49
ASFR ⁵	14.8	69.6	98.1	65.6	25.2	5.2	0.4
$af(a)$ ⁶ (%)	1.06	4.97	7.01	4.69	1.80	0.37	0.03

⁵ Source: U.S Bureau of the Census, International Database, Table 028: Age specific fertility rates (in Austria in 2002)

⁶ The total fertility rate in Austria in 2002 was 1.399. That year a thousand women aged between 15 and 19 years on average gave birth to 14.8 (= 1.48%) children as indicated in the first line of the table. Equation (2) implies that we need a standardized age-specific fertility rate assuming a total fertility rate of 1. This standardized age-specific fertility rate is given in the second line of the table. For a women in the age group [15,19] we get $af(a) = 1.48/1.399 \approx 1.06$.

The implementation of the agents' ages also influences the evolution of norms since the assumption that all couples can get children with the same probability no matter at which age the couple has married does not comply with reality. Thus integration of the age causes a displacement of the lower age at marriage bounds downward.

5 Simulation Details

As already mentioned earlier there are two sex specific age-at-marriage norms. Above all, the introduction of an age-at-marriage minimum is of relevance, as in all countries exists a minimum age before which individuals are not allowed to get married. Although a maximum age for marriage isn't intended legitimately, we restrict the maximum age of the agents to m because the evolution of norms is not effected by agents above that age. Consequently, the upper limit of the agents' age-at-marriage intervals cannot exceed m . The separation of norms shows how age-at-marriage norms of women evolve differently than those of men. Furthermore sex is now a vital selection criteria since marriage between individuals of the same sex is not allowed. Another important extension is the specification of the agents' age. In contrast to the original model the individuals can't marry before they reach their personal marriageable age. This characteristic influences the marriage readiness as well as the reproduction. The probability of having children depends on the age of the female partner.

The inhabitants of our simulation enjoy the pleasure to live in a world in which nobody dies before the age of 60. However, then they are removed from the model. We dare to refrain from modelling mortality in a more accurate way because in highly developed countries the chance to survive until the age of 60 is very high⁷ and the dying after the age of 60 does not affect the evolution of age-at-marriage norms, which is the main subject of our study.

The four transmission mechanisms Intersection, Union, Random and Uniform have been retained unchanged but we did not include the two mutation mechanisms (a child does not inherit any information from its parents) into this model. We abandoned the implementation of a mutation operator introducing new age intervals completely randomly because a certain degree of randomness is already being provided by the Random and Uniform operator. Nevertheless, the randomness inserted by these two operators takes place on a well-regulated level.

⁷ For instance the period life table for Austria for the period 1990/92 indicates that the probability for females to survive until the age of 60 is 91.9% and for males it is 83.6%, (Source: Statistik Austria, Statistisches Jahrbuch 2004, p. 75)

5.1 Numerical Parameters

Some model parameters may be changed to show the effect of their values on the results. The values of other parameters are fixed and cannot be changed. This Sect. gives an overview of all parameters and their values used in the simulation.

- N initial population size, $N = 500 - 5000$. N agents described by randomly chosen characteristics are created. On the basis of this starting population, the evolution of this population is simulated.
- m maximum age an individual can achieve, $m = 60$. As soon as an agent becomes 60 years old it is removed from the model — it dies. Moreover, because agents are removed at age m , this parameter also takes over the role of a global upper bound for age at marriage.
- srb sex ratio at birth. To examine the effects of an imbalanced ratio between sexes, values from 0.5 to 2 are allowed. A srb of 1.05 means that 105 boys are born while 100 girls are born.
- l_a lower bound for age at marriage, $l_a = 15$ years.
- pm probability that an individual who has found an acceptable partner really marries.
- pm_0 probability of marriage in the first year after arriving at marriageable age.
- tfr total fertility rate. tfr can take values between 1.0 and 3.0.
- $af(a)$ age specific fertility rate of women at age a .
- l^f lower bound of female age at marriage, $l^f \in [15, 59]$.
- u^f upper bound of female age at marriage, $u^f \in [l^f, 59]$.
- l^m lower bound of male age at marriage, $l^m \in [15, 59]$.
- u^m upper bound of male age at marriage, $u^m \in [l^m, 59]$.

5.2 The Agents

Each agent is described by some characteristics determining its behaviour during the simulation. For the age-at-marriage model the following agent characteristics are defined:

Variables	Values	Description
index	0 -	identifier of the agent
age	0 - 59	indicates the agent's age
sex	male / female	shows its sex
married?	true / false	is set true when the agent has married an acceptable partner
mother	index	identifier of the agents mother – for the first generation the value of this variable is undefined.
father	index	identifier of the agents father
brosis ⁸	index	lists all agents who have the same mother and father. An agent is not allowed to marry one of these agents
partner	index	if an agent is married this variable shows its partner otherwise the value is nobody
pregnant?	true / false	for male agents and unmarried agents this value is always set false, for married women it is randomly assigned true based on the probability in equation (2)
female-lower-bound	15 - 59	lower bound of the agent's female age norm l^f
female-upper-bound	l^f - 59	upper bound of the agent's female age norm u^f
male-lower-bound	15 - 59	lower bound of the agent's male age norm l^m
male-upper-bound	l^m - 59	upper bound of the agent's male age norm u^m
of-marriageable-age?	true / false	is true as long as the agent's age is situated within its appropriate sex-specific age interval
transmission	intersection/ union/ random/ uniform	indicates the transmission mechanism used to inherit the age norms

⁸ abbreviation of brothers and sisters

6 Results

To investigate the effect of the four transmission mechanisms on the persistence or dissolution of age norms, the model was implemented in NetLogo⁹. Our simulations show that the transmission mechanism determines which norms survive and which disappear first. Moreover, we are interested in the impact of a combination of two or more transmission mechanisms on the persistence of age norms. To facilitate the comparison of the different transmission mechanisms, the values of some numerical parameters are kept constant. The following six simulations each are started with an initial population of 5000 agents, whose characteristics like age, sex and age norms are assigned randomly. Male and female agents are generated with the same probability ($srb = 1$). To avoid erratic fluctuations in the size of the agent population we set the total fertility rate tfr equal to two. Consequently, the female agents (about half of the population) give birth to two children on average. Finally the variable pm_0 is set equal to 35%.

Intersection combiner

If the child's norm is the intersection of the age intervals of its parents, its lower bound is the maximum of its parent's lower bounds and its upper bound is the minimum of its parent's upper bounds. Therefore, the child's age intervals are always shorter than, or at most, as long as the intervals of its parents. That implies that the mean length of the age at marriage interval decreases with time. Thus the average mean length converges to a very narrow age interval (see Fig. 8). In reality such a development may not persist in the long run. Nevertheless, in a period of increasing lower age limits and nearly constant upper age limits (see for instance the period from 1980 to 2000 in Fig. 2) a mechanism similar to this intersection combiner may be at work. Figures 1 and 2 suggest the interpretation that there is not one universal mechanism at work in the long run. Therefore, here and in the following we will not only look at the long run equilibria resulting from the transmission mechanisms but also investigate the intermediate dynamics.

The age-at-marriage norms of the initial population were chosen randomly. Hence, Figs. 9 and 10 show that the initial state contains practically every possible norm.

Already after 60 time steps the evolution of age norms exhibits a clear trend towards shorter age intervals. Those norms with the largest interval length become fewer (see Figs. 11 and 12).¹⁰ The female age-at-marriage

⁹ Within a short test of different simulation platforms we got the impression that NetLogo provides an easy to use programming environment, which enabled us to quickly implement the simulation model from the scratch. Further details can be found at <http://ccl.northwestern.edu/netlogo/>

¹⁰ In Figs. 11 and 12 there are more female norms remaining in the upper left corner than male norms which is just a random "accident" of that particular simulation run.




Fig. 8. Mean length of age norms — intersection combiner




Fig. 9. Female age-at-marriage norms in the initial population




Fig. 10. Male age-at-marriage norms in the initial population

norms with a lower bound of 50 years or above disappear within the first 60 years, until all agents of the first generation with random norms are removed from the model (see Fig. 11). This phenomenon holds for all simulations and can be explained easily. A female agent who marries at the age of 50 or above isn't able to have offsprings because the age specific fertility rate above the age of 50 is zero. Thus, no child can be born who inherits an age norm with a female lower bound above 50.

Because of a very low age specific fertility rate for 45 to 49 year old women, the age norms with lower bounds in this range disappear over the next few years. The persistence of male age norms doesn't show the above behaviour since we did not take into account male (age-specific) fertility rates. For both sexes the norms in the upper left corner of the diagram vanish gradually. These are the norms with the largest length which disappear because of the intersection combiner. Within the next years the female norms keep converging towards the lower half of the diagonal and the male norms converge towards the whole diagonal, but in the long run also the male norms converge towards the lower half of the diagonal, simply because of the fact that younger men




Fig. 11. Female age-at-marriage norms after 60 time steps




Fig. 12. Male age-at-marriage norms after 60 time steps

have more time to conceive more children to whom they can transmit their norms.

Although those norms along the diagonal have the shortest length of all, and therefore should remain, most of them also disappear. The dissolution of female norms with a higher lower age bound can be traced back to age specific fertility rates. Since in the simulation only married agents can give birth to children, those female agents who get married early have a longer time period for having children and are married during the time in their life with the highest age specific fertility rate. Consequently, these agents have higher chances to have children and pass their norms to the next generation. But there are some norms that died out to which this fact doesn't apply. In addition this phenomenon also occurs in Fig. 12 which shows the male age-at-marriage norms. The disappearance of these norms happens for some other reasons. Individuals who are characterised by such extremely short norms also have very little time to search for a partner. Especially an individual with an age norm at the diagonal is at marriageable age for only one year and his/her partner has to be at a specific age to be able to marry. This reduces the supply of potential partners enormously. Therefore, for those individuals the probability of remaining single is rather high due to the fact that even if there are enough individuals who are characterised by the same norm it is unlikely that they are also at a marriageable age at the same time. On this account also many of these short norms along the diagonal vanish. Finally there are just a handful of norms surviving which each account for a group of agents who are only allowed to marry among themselves (see Figs. 13 and 14). In [4] where the agent's age isn't included the norms converge toward the whole diagonal within a few generations.

Figure 15 illustrates the time development of the mean age at marriage. The dotted line represents one particular simulation run and the solid line shows the average over 10 simulations with the same set of numerical parameters. The mean age at marriage decreases to 18 years due to the fact that the norms surviving in the long run are clustered at the lower end of the diagonal.




Fig. 13. Female age-at-marriage norms after 250 time steps




Fig. 14. Male age-at-marriage norms after 250 time steps




Fig. 15. Mean age-at-marriage within 450 time steps — intersection combiner

Union combiner

Using the parents' union as the children's age norm causes the converse effect of the intersection combiner. Creating a new age norm by using the union combiner sets the lower bound to the minimum of the parents' lower bounds and its upper bound to the maximum of its parent's upper bound. Compared to the intersection, the interval lengths have to be longer than, or at least as long as the parents' intervals. Therefore, the mean length of age-at-marriage norms increases quickly until it reaches the maximum possible length of 44 years (Fig. 16).

Once again we start with an initial population of 5000 agents with random parameters. Compared to the previous simulation, those with the largest interval length do not become fewer, but those norms with the smallest interval length do become fewer. The norms along the diagonal are barely represented by now, whereas the norms amass at the upper left corner representing the norms with the highest possible interval lengths. After 100 time steps there is clear evidence that the norms converge toward the maximum (Figs. 17 and 18). Age norms with very short interval lengths completely disappear. Al-




Fig. 16. Mean length of age norms — union combiner

though the female and male age-norms converge toward the same point their evolution is slightly different. The female norms are soon dominated by the lowest possible lower bound 15, whereas the male norms are dominated by the highest possible upper bound 59 (Figs. 19 and 20). This artefact is caused by the fact that male fertility rates are neglected. After 350 years only the norm with a lower bound of 15 years and an upper bound of 59 has survived, all other norms have disappeared completely. Each individual is characterised by the two age norms with the maximum interval length. This implies that every individual of marriageable age easily finds an acceptable partner. Because of an annual probability to get married above 35% an agent remains single on average for only two years. Therefore the mean age-at-marriage becomes 17 years (Fig. 21).




Fig. 17. Female age-at-marriage norms after 100 time steps




Fig. 18. Male age-at-marriage norms after 100 time steps




Fig. 19. Female age-at-marriage norms after 250 time steps




Fig. 20. Male age-at-marriage norms after 250 time steps




Fig. 21. Mean age-at-marriage within 350 years — union combiner

Random combiner

After investigating the effect of children inheriting their norms as an intersection or as a union of their parents' norms, we demonstrate the consequences of a norm consisting of age bounds applied by chance. Using what we call the random combiner each new born inherits one of its parents' lower bounds with the same probability. The upper bound is chosen the same way. This assignment of bounds doesn't offer the appearance of new bounds but it allows for new combinations of already existing boundaries. Thus age bounds that already got lost during the evolution of norms cannot reappear. Compared to the intersection and union combiner the random combiner allows for more possibilities regarding the norms of the children but not as many as the uniform combiner described in the next Sect. Therefore, the random combiner can be seen as an intermediate mechanism bridging the gap between the two very deterministic combiners and the very undeterministic uniform combiner. This transmission mechanism does not predetermine the change of the mean length. The interval length may increase, decrease, or remain constant as well.

At the beginning the interval length increases which is due to the disappearance of some female norms during the first few years. During the following years there are short term increases as well as short term decreases which is due to the extinction of several bounds but the mean length always levels off at average values since the remaining lower bounds are combined with several upper bounds. Therefore nearly as many norms with a large interval (e.g. persisting male norm with largest length: $(15, 59) \Rightarrow \text{length} = 44$) as norms with a short interval (e.g. shortest remaining male norm with a lower bound of 15: $(15, 21) \Rightarrow \text{length} = 6$) remain. Within this simulation the mean length of the age intervals converges towards 25 years (Fig. 22). Since the random combiner possesses the ability to behave in the same way as the intersection or the union combiner but may also create norms by any other combination of the parents norms, the dynamics of the mean length of age-norms is also somewhere between the results obtained from the two extreme transmission mechanisms.




Fig. 22. Mean length of age-norms — random combiner

Already after a few simulation steps it can be seen that only the female norms in the left half of the diagram remain. It is obvious that norms consisting of a small lower age bound have a high chance to survive. But there are no upper bounds that are obviously superior to others. When looking at male norms there is also a clear trend to the left half of the diagram. However, the convergence happens much slower because the age-at-marriage of men does not have an immediate impact on the number of births. After 1050 time steps (more than twice the simulation time we used for intersection and union combiner) there is still no stable structure (Figs. 23 and 24). There are still numerous variations of possible age-at-marriage norms. Since the structure of norms within the population is not stable yet, the mean length of the age intervals can still change as well. Comparing Figs. 23 and 24 with the results

obtained from the union combiner we can conclude that the variety in lower age limits gets reduced in both cases but the random combiner sustains a higher variety of upper age bounds than the union combiner.




Fig. 23. Female age-at-marriage norms after 1050 time steps




Fig. 24. Male age-at-marriage norms after 1050 time steps

Since the norms do not converge toward an equilibrium, the average value for the age-at-marriage does not converge either but fluctuates between 19 and 23 years in the long run (Fig. 25). The reason for these values is the mean lower bound of female norms of 17 years and the mean lower bound of male norms of 22 years. Due to the mean interval length of 25 years it follows from (1) with $pm_0 = 0.35$, $a - l = 1$ (annual step), and $u - l = 25$ that the probability for individuals to marry increases by $(1 - 0.35) * 1/25 = 2.6$ percentage points per year. Because of this increase most agents who find an appropriate partner do not remain single for more than one year.¹¹




Fig. 25. Mean age-at-marriage within 1050 time steps — random combiner

¹¹ This does not mean that agents who find an appropriate partner but do not marry remain connected to that partner for future periods of the simulation.

Uniform combiner

In the following we will discuss the results obtained from simulation experiments based on the uniform transmission mechanism. Now the children may get any bound between the respective bounds of the parents. Therefore the mean length does not converge toward an extreme, but rather toward an intermediate value. In this case the value for the interval length is nine years as it can be seen in Fig. 26. During the first few decades, the mean length increases which is comparable to the increase of the intersections mean length in the beginning. Like in the previous simulations, some short female norms (those with a lower bound of 50 and above) disappear within the first 60 years, which causes the short increase of mean length at the beginning. From that moment the mean length decreases until it arrives at approximately 9 years. This shrinkage of interval lengths is due to the fact that children inherit bounds somewhere between the respective bounds of their parents, which results in a modest tendency towards shorter age intervals.




Fig. 26. Mean length of age norms — uniform combiner

The norm’s evolution can be anticipated soon. It can be seen that the first norms to disappear are those that have survived in two of the previous experiments (Figs. 27 and 28): The norms which prove to be the strongest in the experiment with the union transmission mechanism are those in the upper left corner. The norms that survived in the intersection experiment are those with the shortest length, which are those along the diagonal. These two groups of norms are the first to die out.

Norms that are nearby a maximum value or a minimum value vanish. These boundary values disappear because they are likely to be paired with a partner with a value that is further away from that bound. Consequently their children are likely to inherit a shorter age interval. Female norms with higher lower limits disappear like in the other tests. However, it takes more than 1000 time steps until only three female norms and tree male norms remain. All




Fig. 27. Female age-at-marriage norms after 150 time steps




Fig. 28. Male age-at-marriage norms after 150 time steps

remaining female norms already have the same upper bound namely 28 years. The female lower bound varies from 19 years to 21 years (Fig. 29). The male lower bound is for all agents 28 years (like the female upper bound). The male upper bounds still varies from 38 years to 40 years (Fig. 30). These bounds are explained by the fact that 21 (female lower bound) is the mean value for the lower bound weighted by the age specific fertility rate and 28 (female upper bound) is the weighted mean value between 21 and 59. The bounds of the male norms are weighted with the remaining time for conceiving children. The according weighted averages for the male lower and upper bound are thus 28 and 38 years, respectively.




Fig. 29. Female age-at-marriage norms after 1050 time steps




Fig. 30. Male age-at-marriage norms after 1050 time steps

Sooner or later there will be only one point left in each diagram representing the strongest norm. Women achieve their marriageable age between 19 years and 21 years while all men reach the marriageable age at 28. Thus the mean age at marriage among the whole population fluctuates between 23 and 25 years in the long run (Fig. 31).




Fig. 31. Mean age-at-marriage within 1050 time steps — uniform combiner

Intersection and union

Now we will investigate a mixed population containing agents with the intersection combiner as well as agents with the union combiner. These are those two transmission mechanisms which result in extreme age norms when they are used in a homogenous population of agents. The union combiner, which results in the age norm in the upper left corner having a maximum mean length of 44 years is combined with the intersection combiner that results in an age norm along the diagonal with short interval lengths. The aim of this experiment is to investigate how the dynamics differ within a heterogeneous population compared to the homogenous populations. Each initial agent is randomly assigned one transmission mechanism with the same probability (0.5 each). Newborn agents inherit the transmission mechanism from one of their parents.

At the beginning the mean length increases nonmonotonically. A little bit later it becomes monotonically increasing until it reaches the maximum possible value of 44 after only 200 years (see Fig. 32). The reason for the increasing length is that the union combiner (causing an increasing length) dominates the intersection combiner which causes a decreasing interval length. The union combiner is the stronger of the two because it allows for more acceptable partners for marriage, which results in a bigger number of couples who can hand down the union transmission mechanism. Only relatively few agents with the intersection combiner get married, and consequently fewer children with an inherited intersection transmission mechanism are born.

The union’s predominance against the intersection is very strong. After only 100 years already more than 80 percent of all individuals are characterised by the union transmission and after 450 years the weaker combiner does not occur anymore in the population (Fig. 33).

Therefore the only norm that remains until the end of the simulation is the one with the lower bound at age 15 and the upper bound at age 59, which is the same that survived in case of a homogenous population of agents applying




Fig. 32. Mean length of age norms — intersection and union combiner




Fig. 33. Proportion of intersection and union combiner

the union combiner. Since the union transmission mechanism dominates the intersection mechanism, the mean age at marriage levels off at 17 years, which was the mean age at marriage in the simulation applying the union combiner.

Combination of all four transmission mechanisms

In this Sect. we have a look at the evolution of norms applying a combination of all four transmission mechanisms — intersection, union, random and uniform. Considering the combination of the four transmission mechanisms the changes of the mean length are comparable to those evolving from the combination of the intersection and the union combiner (Fig. 34). During the first 60 years the graph is monotonically increasing since all four transmission mechanisms cause an increasing mean length due to the disappearance of all female norms with a lower bound above 50 (which all have a short mean length). The following decades show a mean interval length that does no longer increase monotonically. The intersection and the uniform combiner cause some decreases in the short term but as their joint proportion constitutes less than 30 percent after 150 years, the influence of these mechanisms is rather small. Therefore the curve soon is monotonically increasing again and reaches the maximum possible length of 44 years after only 350 years when already 65 percent of all agents inherit their age norms as the union of their parents' intervals (Figs. 34 and 35).

It takes more than 1000 years until the union combiner dominates all the other combiners and only the norm with the longest possible interval remains. Because of the persisted age norm (15, 59) individuals are allowed to marry when they are at the age of 15. Most agents stay single for two years and get married at the age of 17.

By influencing the evolution of norms, the choice of the transmission mechanism also influences the mean length of the age interval, the number of married couples and singles of marriageable age and the mean age at marriage.




Fig. 34. Mean length applying all four transmission mechanisms




Fig. 35. Proportion of the transmission mechanisms

Table 2 shows the respective values. To make the data comparable, the proportion of married couples of all simulations are taken after 300 years¹².

Table 2. Comparison of the transmission mechanisms

	intersection	union	random	uniform
proportion of married couples (%)	29.6	36.0	30.8	29.9
mean length	1	44	25	9
mean age at marriage	18	17	19–23	23–24

This comparison explains the results of the two investigated combinations of transmission mechanisms. Individuals who are characterised by the union combiner find considerably more potential partners than those who are characterised by the intersection combiner because of the union’s large interval length. Therefore, the union combiner dominates the simulations with heterogeneous populations. The random transmission mechanism, which has the second largest interval length and thus the second largest proportion of couples, does also persists in a combination of all transmission mechanisms. The execution of different simulations showed that after 1050 years 10% to 25% of the agents are characterised by the random combiner. After the intersection combiner has died out also the uniform combiner disappears because its marriage rate is not that high as well. Since the union transmission technique is the strongest one that dominates all other mechanisms both simulated combinations finally lead to the age norm (15, 59) with an age interval of 44 years and a mean age at marriage of 17.

¹² The proportion of married couples was measured after 300 years since these midterm results illustrate the development of the distribution of the combiners within the population. The mean length of the age interval and the mean age at marriage were taken at the end of the simulation since we are interested in the long-term equilibrium.

Parameter variations

So far we have only been looking at simulations based on a set of numerical parameters which are kept constant over the whole time horizon. The purpose of these simulations is to understand how the different transmission mechanisms work and what is their impact on the appearance or dissolution of norms within the agent population. However, the empirical data discussed in Chap. 2 give evidence that in real societies norms and values are not constant over time (see Figs. 1 and 2). Consequently, we want to investigate whether our model is capable to replicate the observed dynamics. From the previous simulations we conclude that the uniform combiner results in medium size interval lengths which is of course the kind of dynamic behaviour which is most appropriate to approximate real world dynamics. Therefore, we set up a simulation model based on an agent population which is homogenous in terms of the uniform transmission mechanism. At the beginning we fix the parameter pm_0 equal to 35% and simulate 100 time steps to arrive at stable age-at-marriage norms. After that we modify pm_0 in ten year time steps such that $pm_0 = 35, 65, 95, 100, 55, 25, 5, 0$ at $t = 100, 110, \dots, 170$. Fig. 36 reveals that — neglecting the fluctuations — in this setup the mean age at marriage decreases for some decades from around 24 to about 22 and later on increases to mean ages higher than at the beginning of the simulation. Looking at the age specific rate of marriage we see that the frequency of marriage among young agents increases between $t = 100$ and $t = 120$ but later on decreases to rather low levels (see Fig. 37). Thus, we can conclude that the time dynamics within the agent population are similar to those observed in the real world data discussed in Sect. 2. The pronounced fluctuability in the mean age at marriage in Fig. 36 is due to the rather low size of the agent population, $N = 2000$. Since real populations are much bigger, the curves in Fig. 1 are smoother.




Fig. 36. Mean age at marriage




Fig. 37. Rate of marriage

7 Concluding Remarks

In this Chap. we investigated the impact of the design of the transmission mechanism on appearance, shifts, and extinction of social norms within an agent population. Moreover, we looked at the impact of these social norms on the age-at-marriage and on the age specific rates of marriage. In particular we explored the effect of four different transmission mechanisms — intersection, union, random, and uniform — on the dynamic behaviour of the social norms.

The first simulation considered the evolution of norms within a homogeneous population of agents endowed with the intersection combiner resulting in a decreasing mean length of the age at marriage interval. The final interval length was 1 year. The age norms converged toward the diagonal and finally only a few age norms with a lower bound between 16 and 20 years and an upper bound of 18 to 20 years survived. Applying the union combiner caused an increasing mean length up to the maximum possible value of 44 years. Regarding age norms a convergence toward the upper left corner could be observed. The random combiner did not cause one isolated norm to survive but the variety of different bounds shrank. Some lower and some upper bounds vanished but the structure of norms within the population was still not stable after more than 1000 years. The interval length of the age norms varied. There was no clear increase or decrease. In case of the uniform combiner, the mean interval length leveled off to a narrow value. The number of age-at-marriage norms reduced until a single point (15,59) survived. But this process lasted considerably longer than it lasted using the intersection or the union combiner. In a heterogenous population of agents equipped with different transmission mechanisms the fraction of the population characterized by the union combiner increased until extinction of all other transmission mechanisms.

A simulation setup based on the uniform transmission alone combined with the variation of the parameter pm_0 determining the initial probability of getting married allowed us to approximate the time development of age-at-marriage and the age specific rate of marriage among birth cohorts observed in empirical data. It turned out that a temporary increase followed by a decrease of the initial probability to get married may be an explanation of the U-shaped curve indicating the mean age at first marriage in some European countries (see Fig. 1). Thus, such shifts in the initial probability of getting married may at least partially explain past trends in age-at-marriage. For instance, it is reasonable to assume that women born in the forties considered marriage and childbearing as major priorities in their lives while succeeding cohorts were more interested in getting a proper education and pursuing their professional career before marriage. Compared to the model investigated in [4] the extended model introduced in this Chap. proved to provide a significant step toward reality since the timing of union formation and childbearing is taken explicitly into account. Only with this explicit consideration of time and age it is possible to investigate the impact of social norms and param-

eter shifts on mean age at marriage and age specific marriage rates. Hence, the increased complexity of the model is needed to replicate the phenomena observed in empirical data.

Of course we are aware that social norms are not the only mechanism influencing individuals on their decision about getting married. There are several forces at work at the same time. For instance, the availability of appropriate mates, economic considerations, and attractiveness may influence the decision. However, the empirical studies summarized in Sect. 2 give evidence for the existence of such social norms and our simulation model shows clearly that the existence of social norms can generate a behaviour similar to empirical data. Taken together, these findings strongly support the assumption that age-at-marriage norms indeed have a major influence on the decision to get married.

References

1. Maire Ní Bhrolcháin (2004) Validating age preferences for marriage market analysis. S³RI Applications Working Paper A04/08, Division of Social Statistics, School of Social Sciences, University of Southampton.
2. Francesco C. Billari and Letizia Mencarini (2004) Norms and sanctions on sexual life and marriage of young adults. In: G. Dalia Zuana and C. Crisafulli, editors, *Sexual Behaviour of Italian Students*, pp. 361–379. Department of Statistics, University of Messina, Italy.
3. Francesco C. Billari and G.A. Micheli (2001) Social norms on agents demographic events. preliminaries for a (multi-)agent-based approach. In: Rosaria Conte and C. Dellarocas, editors, *Social Order in Multiagent Systems*. Kluwer, Dordrecht.
4. Francesco C. Billari, Alexia Prskawetz, and Johannes Fürnkranz (2003) On the cultural evolution of age-at-marriage norms. In: Francesco C. Billari and Alexia Prskawetz, editors, *Agent-Based Computational Demography: Using Simulation to Improve Our Understanding of Demographic Behaviour*, Contributions to Economics, pp. 139–157. Physica-Verlag.
5. M. Buchmann (1989) *The script of life in modern society. Entry into adulthood an a changing world*. University of Chicago Press, Chicago.
6. Andreas Diekmann and Thomas Voss (2003) Social norms and reciprocity. Paper presented on the session “Solidarity and social norms — models and mechanisms”, Sektion Modellbildung und Simulation, Leipzig, October 2002.
7. Ruth B. Dixon (1971) Explaining cross-cultural variations in age at marriage and proportions never marrying. *Population Studies*, 25(2):215–233.
8. John Hajnal (1953) Age at marriage and proportions marrying. *Population Studies*, 7:111–136.
9. J. Heckhausen (1999) *Developmental Regulation in Adulthood. Age-Normative and Sociostructural Constraints as Adaptive Challenges*. Cambridge University Press.
10. Rainer Hegselmann and Ulrich Krause (2002) Opinion dynamics and bounded confidence: Models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3).

11. Theodore Palivos (2001) Social norms, fertility and economic development. *Journal of Economic Dynamics & Control*, 25:1919–1934.
12. Richard A. Settersten and Gunhild O. Hägestad (1996) What’s the latest? cultural age deadlines for family transitions. *The Gerontologist*, 36(2):178–188.
13. Jorge Simao and Peter M. Todd (2003) Emergent patterns of mate choice in human populations. *Artificial Life*, 9:403–417.
14. Peter M. Todd and Francesco C. Billari (2003) Population-wide marriage patterns produced by individual mate-search heuristics. In: Francesco C. Billari and Alexia Fürnkranz, editors, *Agent-based computational demography*, pp. 117–137, Berlin, Springer-Verlag.
15. Peter M. Todd, Francesco C. Billari, and Jorge Simao (2005) Aggregate age-at-marriage patterns from individual mate search heuristics. *Demography*, 42.
16. Diemo Urbig (2003) Attitude dynamics with limited verbalisation capabilities. *Journal of Artificial Societies and Social Simulation*, 6(1).

The Strength of Social Interactions and Obesity among Women

Mary A. Burke¹ and Frank Heiland^{2*}

¹ Research Department, Federal Reserve Bank of Boston, U.S.A.

Mary.Burke@bos.frb.org

² Department of Economics, School of Computational Science and Center for Demography and Population Health, Florida State University, Tallahassee, U.S.A. fheiland@fsu.edu

Summary. In order to explain the differences in obesity rates among women in the United States by education, we model a social process in which body weight norms are determined endogenously in relation to the weight distribution of the peer group. The model features biologically grounded variation in metabolism, and enables us to describe a complete distribution of weights in equilibrium. We assume that individuals compare themselves to others with the same level of education, and that the importance of conforming to the group weight norm increases with education status. Consistent with observed body weights among women in the United States, the model predicts lower average weights and less dispersion of weight among more educated women.

1 Introduction

The dramatic growth in obesity rates in the United States since the early 1980's has been widely publicized and much fretted over in recent years. In the aggregate the percentage of obese adults in the U.S. increased from 13.6% in the 1970's to a rate of close to 30% in 2000. While official Centers for Disease Control (CDC) estimates of obesity-related mortality in the U.S. were recently adjusted downward [18], from a high estimate of 400,000 per year to about 65,000, obesity remains a prominent source of elevated disease risk, most notably for diabetes and heart disease, and has contributed significantly to increasing medical expenditures. The official definition of obesity employed by CDC and by the World Health Organization (WHO) is a body mass index (BMI) value of 30 or greater, where BMI is the ratio of weight, measured in

* We are grateful for the comments of two anonymous referees. An extended version of this article including the FORTRAN90 code used to solve the model can be found online at <http://mailer.fsu.edu/~fheiland/research.htm>.

kilograms, to squared height, measured in meters. BMI less than 18.5 is officially underweight, BMI values between 18.5 and 24.9 are considered healthy, and BMI between 25 and 29.9 is defined as overweight but not obese.

While the increases in obesity have cut across a wide swath of the U.S. population, obesity prevalence varies significantly across groups by education level. Among U.S. women observed over the 1990-2002 period, the median weight for 30-60 year old college graduates was 149 pounds, while the median for those in the same age group with just a high school diploma was 162 pounds.³ The difference in average weights between these groups is even larger, the respective means being 143 and 154 pounds as shown in Fig. 1.⁴ The box plots of the weight distributions also illustrate the fact that the dispersion of weight about the median among more educated women is substantially smaller than that for the less educated.⁵ For men, however, the differences by education are much less pronounced. For example, college educated men in the 30-60 age bracket weighed on average 187 pounds compared to 190 pounds for men with only a high school education. Median weights are almost identical (184 pounds) for the two groups (see Table 1). Additional analysis confirms that the education and gender patterns observed in the weight data also hold when comparing the distributions of BMI. In this paper we use an agent-based model of choice to explain the differences between the weight distributions of U.S. women by educational attainment. Based on the hypothesis that more educated women face stronger incentives to conform to a body weight norm than do other women, the model captures many of the distributional differences across education class. We also investigate competing explanations based on behavioral or genetic differences across the groups. While data limitations preclude a direct empirical test of the social interaction hypothesis, the agent-based simulation approach taken here serves as an alternative tool for assessing the explanatory power of the competing hypotheses.

³ The empirical findings presented in this paper are based on a sample of 30 to 60 year old Americans from the Behavioral Risk Factor Surveillance System (BRFSS), a survey administered by the Centers for Disease Control and Prevention. The BRFSS is a large random sample of the resident population 18 years and older in participating states of the U.S.. Self-reported information on actual weight, height and socio-economic and demographic characteristics is gathered annually between 1990 and 2002. We correct for potential bias of self-reported weights (see [39]) following the approach by Chou et al. [14] using data from the third wave of the National Health and Nutrition Examination Survey for the 30-60 year olds.

⁴ Kolmogorov-Smirnov tests easily reject equality of the weight distributions by education. Specifically, they provide strong statistical support that college educated women have the lowest weights while high school dropouts have the highest weights.

⁵ Both the size of the box (=difference between the 75th and the 25th percentile=interquartile range) and the distance between the upper and the lower adjacent values (=four times the length of the box) illustrate the dispersion of the data about the median.

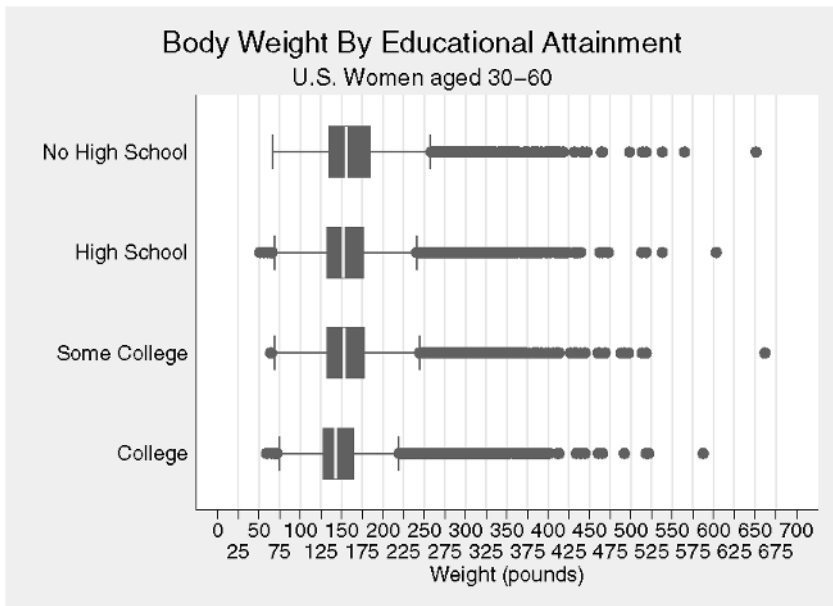


Fig. 1. Box plot of body weight distribution, 1990-2002, U.S. Women aged 30-60 (Source: authors’ computation from BRFSS 1990-2002 data)

Several recent papers in economics have sought to explain obesity among adults, including [14], [24], [16], [30], [22]. While these studies have offered some compelling explanations for the general increases in weights in the United States over time, they provide little insight into the relationship between education and weight, and, with the exception of [16], they focus on changes in average weight and ignore other distributional features. An earlier paper by the current authors [11] incorporates the influence of falling food prices into a model with social weight norms and genetic heterogeneity, and captures not just increases in average weights over time but also the growth in the upper tail of the distribution since the 1970’s. The current work employs the same basic choice framework as the earlier paper but focuses on cross-sectional differences by education rather than changes over time.

In our socio-economic choice model, individuals suffer disutility from deviating from an endogenously determined social weight norm. For a given value of the weight norm, M , the disutility is parameterized as $-J(W_i - M)^2$, where W_i is individual weight, and J captures the individual’s concern for conformity or social esteem. The greater is J the greater the disutility of a given deviation from a given norm. This formulation is standard in the social interactions literature, in which J (or some equivalent) is described as representing the “strength of social interactions” or the degree of concern for social esteem. Deviation costs may be self-imposed (a result of internalization of a norm), and/or imposed by others, as discussed below. We argue, based

on independent empirical evidence, that more educated women face stronger incentives to conform to social weight norms, in the context of competition for jobs, status, and marriage partners, than do less educated women. Given these stronger incentives, modelled as a higher value for J , we predict lower average weights among more educated women, a lower value for the endogenous weight norm, and a smaller variance in weights within this group, all of which are consistent with the empirical evidence.

In the model, each education class constitutes a self-referencing social group with a particular common value for the conformity parameter, J . We describe the qualitative effects of variation in J , and then simulate equilibrium weight distributions for different values of J , calibrating the model to American women 30-60 years of age. Preferences over food and non-food goods are identical across individuals, but metabolism is permitted to vary. Each agent draws a one-time metabolic shock that determines the relationship between her body weight and calories burned per day at every possible weight level; all shocks are drawn from the same distribution, whose parameters and characteristics are based on empirical studies of exogenous human metabolic variation. Equilibrium is group-specific, consisting of the stable distribution of weights and the corresponding equilibrium weight norm that together satisfy an internal consistency condition. Simulations based on this social interactions model are compared to simulated distributions under competing hypotheses in which individual characteristics vary systematically across groups and conformity may or may not be a concern. We find that, while several different hypotheses can explain variation in *average* weight across educational groups, our social interaction hypothesis does a superior job of explaining cross-group differences in additional distributional features such as variance and skewness.

The social interactions literature has long recognized the theoretical importance of the strength of social interactions, for example in determining the number and character of equilibria (among many others see [7]). And while in theory the parameter may be allowed to vary across individuals or across pairings of individuals depending on their characteristics, very little work has been done to measure J empirically or to identify the factors that affect its magnitude. Thus this investigation, in addition to seeking to explain variation in obesity rates across groups of women, also contributes to the study of the empirical determinants of social pressures.

The remainder of the paper is organized as follows. The remainder of Section 1 motivates the central assumptions of the model and presents competing hypotheses concerning obesity and education level. Section 2 describes the theoretical model and the comparative static effects of the strength of social interactions on equilibrium weights and the equilibrium weight norm. Section 3 presents the numerical simulations of the equilibrium weight distributions for women, illustrating the role of peer groups and of variation in the strength of social interactions by education. It also discusses the predictive power of alternative explanations that do not assume differences in the strength of social interaction. Section 4 concludes.

1.1 Sources of Non-Conformity Costs and their Relation to Education

The central hypothesis that drives the variation in weight distributions across education groups is that the degree of concern for conformity to a social weight standard (alternatively termed the strength of social interactions) increases with a woman's education level. While it is difficult to observe such concerns directly, there are a number of reasons to expect such a relationship. First, educational attainment might be construed as revealing (among other things) an individual's level of concern for social status. Given that educational attainment is highly correlated with measures of social and economic status such as employment and income, individuals who pursue higher education may do so because they have a relatively strong concern for social esteem. Concern for social esteem should influence other behaviors or choices that are socially visible, such as body weight and appearance in general [5].

Education might alternatively predict variation in the importance of conformity to body weight norms through its association with economic opportunity. Cawley [13] and Averett and Korenman [2], based on recent U.S. samples, show that overweight and obesity are associated with earnings penalties for women but not, interestingly, for men.^{6,7} Women with greater human capital are more likely to pursue professional careers and spend more time in the labor market, and hence penalties that stem from social interactions may be more salient for educated women. Consistent with greater costs of obesity for more educated women, Bhattacharya and Bundorf [6] find higher obesity-related wage penalties for women in high income occupations (the lower wages are in part to offset the greater healthcare costs of obese employees). A recent study by Carr and Friedman [12] on institutional and interpersonal discrimination finds that severely obese ($BMI > 35$) professionals are more likely than their nonprofessional counterparts to experience discrimination on the job. Also, Ross [35] finds that, among overweight women, the extent of depression attributable to negative self-perception is only significant for highly educated women.

Marriage market conditions may also vary substantially for women by educational attainment. While all women face large marriage market penalties for being overweight or obese [2], there are several reasons why such costs may be particularly large for better educated women. For example, women who invest in their education delay marriage and child-bearing relative to women that complete only high school or less. In a rational choice framework such delay represents a choice that is simultaneous with the choice of educational

⁶ A related literature finds that women in low prestige jobs tend to be more obese [28].

⁷ Employers may use BMI as a signal of the productivity of an employee. Using evidence from twin data [4] find that the correlation between excess weight and earnings disappears when earnings endowment, schooling, height and job experience are controlled for.

attainment. Given this choice the woman should feel pressure to stay attractive (relative to a particular social standard) for longer than women who marry earlier. In addition, since women have a preference to “marry up” in quality, and men exhibit a preference for women younger than themselves, the group of available and acceptable men shrinks with a woman’s education level and age, and competition for men may thus be more fierce among this group [34]. Within any given age cohort the more educated women are less likely to be married and, conditional on seeking marriage in the future, more likely to be concerned with achieving or maintaining a desirable body weight.

1.2 Competing Hypotheses

The recent economic literature on obesity generates potential alternative explanations for the relationship between education and body weight among women. Chou et al. [14], using pooled data from the BRFSS, find a negative association between educational attainment and obesity that remains after controlling for income, gender, race/ethnicity, and food prices. The authors do not explore the association between education and obesity further, but note that the effect may capture differences in the efficiency in household production, health awareness, discounting, or genetics (see [14], pp. 571, 579). Cutler et al. [16] suggest that the emergence of calorically dense convenience foods and falling food preparation costs are responsible for increased consumption of calories. The model predicts that those with the highest value for convenience foods would have gained the most weight as the price and time costs of convenience foods has fallen. But this framework implies, counter to evidence, that more educated women, with greater time opportunity costs, would have gained relatively more weight than low-opportunity cost women in recent years. Philipson and Posner [30] and Lakdawalla and Philipson [22] identify lower calorie prices and lower calorie expenditure due to labor-saving technologies as causes of rising obesity. The sign of the effect of education on weight implied by this explanation depends on whether more educated women expend fewer calories in their jobs and at home than less educated women. Brownson et al. [9, Table 3] provide evidence from a sample of U.S. women age 40 and older that more schooling is associated with more physical activity on the job but less physical activity at home. Based on this finding the labor-saving hypothesis does not readily explain the observed differences in average weights by education. Nonetheless the Philipson framework can achieve the predicted relationship between education and weight based on its assumption that closeness to ideal weight is a normal good (ideal weight is exogenous). In equilibrium higher income beyond a point leads to lower weight, and the association between income and education yields a similar prediction between education and weight.

Along these same lines a number of explanations based on preferences and information might be advanced. Intuitively we might simply assume that more educated individuals have greater awareness of nutrition, a better understand-

ing of the relationship between food intake, exercise, and weight, and greater awareness of the health risks associated with obesity. Such individuals should be more motivated and better equipped to maintain “healthy” weight levels as construed by medical standards. An observationally equivalent hypothesis would be that there are genetic differences that affect both educational attainment and either resting metabolism or food preferences. Another explanation may be that cultural composition varies by educational class, and with it cultural standards of ideal weight. For example, less educated women are more likely to be African-American or have a non-Western cultural background, and a number of studies have argued that African-Americans and some non-Western cultures exhibit a preference for higher female weight. Yet another factor is that women with higher education delay pregnancy and have fewer children on average, thus delaying or avoiding the (often persistent) weight gains associated with pregnancy. The latter explanation is problematic, however, given the endogeneity of fertility and pregnancy-related weight gain.

Income effects represent a prominent alternative explanation of the relationship between education and weight. More education leads to higher income, and it is often assumed that thinness, or closeness to a weight norm (as in [22]), is a normal good. While there is cross-sectional evidence of a modest negative relationship between income and BMI, the effects of education on weight remain significant even after controlling for income [14]. While Lakdawalla and Philipson [22] have noted a Kuznets-curve-type (inverted U-shape) relationship between income and weight, recent evidence indicates that the income-obesity gap has narrowed over time in the U.S. since the 1970’s [33]. In addition, unlike the social interactions model, a model that assumes closeness to an exogenous ideal weight is a normal good cannot explain observed changes in weight aspirations or income-weight relationships over time ([1], [37]).

The alternative explanations have the potential to explain the differences in average or median weight by education. However they do not readily generate predictions on higher moments of the weight distribution in relation to education. To tease out such predictions we simulate models of group-specific heterogeneity, in which weight norms may or may not enter preferences. In general the alternative models fail to explain the greater dispersion of weights among more educated women. Also, the alternative explanations should generally predict similar relationships between education and weight for women and men, counter to empirical evidence. Although it is beyond the scope of this article, we believe that our model has the potential to explain the gender-specific relationships using similar arguments about conformity incentives across groups.

1.3 Description of Weight Norms and Reference Groups

While the existence of weight standards may seem like an obvious social fact, the formal description of such standards is less obvious because there is no

scientific consensus on how they are formed. Norms of body weight and shape within a culture may derive from a number of sources, such as ideals promulgated in cultural images, the desire to be normal in relation to one's peers, and health standards promoted by physicians and public health officials. We do not provide a theory of how such sources might interact to determine weight aspirations. Instead we offer a model that aims to capture the net result of such influences in a contemporary Western context. The weight standard to which individuals aspire is defined as a fraction, less than one, of average weight in the reference population.⁸ This specification, in which people aim to be thinner than the average person in the reference population, combines two basic assumptions: (1) that in contemporary Western society thinness (up to a point) is prized, and (2) that individuals assess themselves in relation to their peers rather than against an absolute scale. Relative assessment within peer groups strikes us as appropriate to the context that motivates our model, namely that of positional competition among women for status, scarce jobs, and romantic partners [19].

As stated above, we define reference groups on the basis of educational attainment, such that the disutility of weight is assessed in relation to the weight standard that emerges from the reference group's weight distribution. This specification presumes that comparisons to peers with similar education levels are particularly salient. This will be true if individuals care only how they compare to others of the same educational class, despite interacting with people of diverse backgrounds, or if people tend to interact more frequently with others of similar education levels. Social stratification by education level is likely for a number of reasons: at the secondary level tracking by ability may lead to selection of friends on the basis of future educational attainment, leading to assortative friendships at later stages if the ties persist; individuals that attend college form long-lasting social ties with other college attendees; selection into jobs depends on education levels, leading to stratification in workplace as well as non-workplace interactions. More obviously, people may simply prefer to interact with others of similar educational backgrounds because they have more in common. The observed assortativity of U.S. marriage markets by education level constitutes evidence of such preferences ([29], [25]).

2 Theoretical Framework

2.1 Agent-Based Model

The model consists of boundedly rational individuals interacting within a social group. Each individual compares her own weight to the group's commonly-held norm or "reference" weight, and this comparison enters her optimization

⁸ Burke and Heiland [11] provide evidence from desired weight data from the BRFSS that women aged 30 to 60 between 1990 and 2002 desire a weight level that is about 15% below the average weight.

problem as described below. The reference weight itself is a function of the group's weight distribution. Equilibrium for the system is defined as a weight distribution and a norm that are mutually consistent. Each individual maximizes a myopic utility function over short-term food and non-food consumption taking the reference weight and prices into account. Food and non-food consumption are both goods, but deviation from the reference weight is a bad. A general expression of the one-period utility model is as follows:

$$U_{it}[F_t, C_t|W_{t-1}] = G_i[F_{it}, C_{it}] - J(W_{it}[F_{it}, W_{i,t-1}, \epsilon_i] - M_{t-1})^2. \quad (1)$$

F_t and C_t represent food and non-food consumption for period t , respectively. W_{t-1} is weight at the end of period $t - 1$, which is a product of past actions. Individual heterogeneity is captured by ϵ_i , which is a stationary shock to basal metabolism. G_i is the norm-independent or “private” component of utility: it is strictly increasing and strictly concave in C , and strictly concave but not necessarily monotonic in F . The term beginning with J gives the “social” component, which is the disutility of deviating from the reference weight (norm), M . The subscript on M indicates that agents observe the value of M at the end of period $t - 1$ and take this as fixed in the optimization; in particular they do not forecast the value of M that will emerge as a consequence of aggregate behavior in period t ; this specification facilitates solution of the model without altering the equilibrium outcome relative to a model with rational expectations of the norm.

The individual correctly anticipates her end-of-period weight as a function of food intake, and so takes into account the effect of current food consumption on the cost of deviating from the reference weight. This cost is symmetric—it is just as undesirable to be underweight relative to the norm as overweight—and is meant to capture a number of known types sanctions for non-conformity to weight norms. As mentioned above, overweight has been associated with significant wage penalties among women. Stigmatization of overweight (and underweight) individuals is well-documented [27], and may entail for example teasing, ostracism, discrimination in hiring, and fewer friendship and marriage opportunities. Peer pressure and contagion regarding eating behavior have also been observed, particularly among adolescent girls [15]. Ross [35] has identified depression as a consequence of overweight among women; she finds that educated women appear to suffer most acutely.

In addition to economic costs, social costs and mental health costs, extreme overweight and underweight entail significant physical health consequences (e.g., [26], [38]). Evidence from developing countries, where underweight is much more prevalent, indicates substantially elevated disease incidence among low weight (BMI below 20) individuals [17]. A model with deviation costs that depend on a mutable norm will capture these health costs only when the value of the norm lies within the medically recommended range. In the parameterizations we consider the emergent norms do in fact fall within this range, but in general the model does not constrain them to do so.

Our functional form and parametric assumptions guarantee convergence to a unique stable weight value for a given metabolic shock and a given weight norm (see [11]). This stable weight does not in general coincide with the stationary weight that optimizes the corresponding infinite horizon problem. The myopic specification may be taken to imply lack of self-control, although we do not explicitly model a time inconsistency problem. We believe the model is behaviorally plausible: individuals give some thought to the effects of calorie consumption on weight, but without full consideration of the lifetime implications.

For purposes of simulation and calibration we specify the maximization problem as follows:

$$\begin{aligned} \text{Max}_{\{F_t, C_t\}} U_{it}[F_t, C_t | W_{i,t-1}, \alpha, \delta, \beta, J, \gamma, \rho, \epsilon_i, M_{t-1}] = & \alpha F_{it} - \delta F_{it}^2 + \beta \log[C_{it} + 1] \\ & - J(W_{i,t-1} - (7/3500)(\gamma + (\rho + \epsilon_i)W_{i,t-1} - .00025W_{i,t-1}^2) + .9F_{it} - M_{t-1})^2, \\ & \text{s.t. } p_t F_t + C_t \leq Y \end{aligned} \quad (2)$$

Within the single period, calibrated to one week, the marginal utility of food, F , declines and eventually becomes negative. The expression inside the parentheses following J just amounts to the difference between end-of-period weight W_t , and M , as in equation (1).

Metabolism is quadratic in body weight, and the short-term relationship between food intake and weight (2) is:

$$W_t = W_{t-1} - (7/3500)(\gamma + (\rho + \epsilon_i)W_{t-1} - .00025W_{t-1}^2) + .9F_t. \quad (3)$$

Calories burned per day, not including those burned in digestion, is given in the above by the terms $\gamma + (\rho + \epsilon_i)W_{t-1} - .00025W_{t-1}^2$. The constant and the weight-linear terms are due to [36], and have been used in Cutler et al. and others. While the Schofield equations have become a de facto standard for predicting BMR, some aspects of the estimates have come into question by Horgan and Stubbs [20] and Pullicino [31]. In a perhaps biologically more accurate model, we allow a weight-BMR relationship that is quadratic in weight, such that BMR per unit of body weight declines in weight.

We multiply calories per day by 7/3500 to convert to pounds of body weight expended per week, based on the fact that burning 3500 calories implies loss of one pound of body weight. F_t is total food intake for the week, measured as calories divided by 3500, or the equivalent of the caloric intake in pounds of body weight accrued. Since digestion of a given amount of food requires on average 10% of calories consumed, we multiply by .9 to get food intake net of digestive metabolism [36]. Aside from the calories burned in digestion, we assume for simplicity that calorie expenditure is limited to the basal metabolic rate (BMR), or the calories needed only to sustain basic bodily functions such as lung and heart activity with the body at rest. The advantage of this assumption is that BMR is exogenous in body weight. Of course, variation in physical activity also contributes to variation in calorie

expenditure and therefore weight. By abstracting from endogenous physical activity in the model we assume that the number of actual calories burned is strongly correlated with BMR. We assume that individuals correctly perceive both their food intake and this metabolism function.⁹

As suggested by Leibel et al. [23] and Rand [32], we make the disturbances proportional to weight: the shock ϵ_i is normally and identically distributed with mean zero and standard deviation σ_ϵ . Because the shock is multiplied by weight the errors are heteroscedastic by weight class. Unlike the homoscedastic case, this specification implies an asymmetric equilibrium weight distribution, with a long upper tail mirroring the general shape of the weight distributions observed in the BRFSS data.

2.2 Definition of Equilibrium

Within a given education group, individuals are identical in all of the parameters of the utility function, α , β , ρ , γ , J , M , have identical incomes, and face the same price. The only explicit source of heterogeneity is the idiosyncratic metabolic shock, ϵ_i . The full equilibrium conditions can be expressed as follows:

$$\alpha - 2\delta F_i^S - 1.8J(W_i^S - M^S) = \lambda p, \quad (4)$$

$$F_i^S = (1.11)(7/3500)(\gamma + (\rho + \epsilon_i)W_i^S - .00025(W_i^S)^2), \quad (5)$$

$$M^S = \zeta \left(\frac{1}{N} \sum_i W_i^S \right), \quad (6)$$

$$\frac{\beta}{C_i^S + 1} = \lambda, \quad (7)$$

$$pF_i^S + C_i^S = Y_i. \quad (8)$$

The conditions apply to an interior equilibrium, in which stable food intake, F_i^S , stable weight, W_i^S , and stable non-food consumption, C_i^S are all strictly positive. M^S is the equilibrium weight norm, which according to equation (6) is some fraction, ζ , of the average stable weight that arises under this norm. Equation (4) gives the first-order condition on food consumption, where λ is the Lagrange multiplier. Equation (5) guarantees that per-period food intake maintains weight at the level W_i^S . Equations (7) and (8) are, respectively, the first order condition on non-food consumption and the budget constraint. The equilibrium norm depends on the relative price of food, the distribution of individual shocks, and the magnitude of J , because these determine the stable individual weights and consumption levels for any fixed M . The equilibrium norm (and therefore the weight distribution) also depends on ζ , which we will

⁹ There is evidence that people systematically underestimate their caloric intake [40], but we ignore this problem here.

set at .85 based on the evidence in [11]. Equilibrium depends on income levels and the remaining parameters as well, but we hold these fixed throughout the analysis.

2.3 Comparative Statics and the Strength of Social Interactions

Before generating quantitative experimental results, we analyze the qualitative effects of variation in J on equilibrium behavior. The non-conformity cost of a given deviation of individual weight from the group norm increases with J , and thus so does the incentive to conform. Holding the norm fixed, an increase in the value of J will reduce the absolute difference between an individual’s chosen stable weight and the weight norm: individuals initially below the weight norm gain weight (but remain below the norm); individuals initially above the norm lose weight but remain above; and an individual with stable weight equal to the initial norm will stay put. In the aggregate these weight changes alter average weight, and therefore the norm must be updated. This norm change in turn sets off additional weight adjustments, until a new equilibrium is reached. Thus we observe a “social multiplier” effect, as in [3], [8], [10], and [11], among others. The total effect of J on stable individual weight can be decomposed as follows:

$$\frac{dW_i^S}{dJ} = \frac{\partial W_i^S}{\partial J} + \frac{\partial W_i^S}{\partial M} \frac{dM^S}{dJ}, \tag{9}$$

where the expression $\frac{dM^S}{dJ}$ refers to the change in the equilibrium norm caused by the change in J . While we have already analyzed the sign of $\frac{\partial W_i^S}{\partial J}$, to get at the sign on $\frac{dW_i^S}{dJ}$ we need to sign the remaining terms in the expression. It is readily shown that $\frac{\partial W_i^S}{\partial M} \geq 0$ for all individuals. But the effect of J on the equilibrium norm, $\frac{dM^S}{dJ}$, is less obvious. The following decomposition helps to determine its sign:

$$\frac{dM^S}{dJ} = \frac{\zeta}{N} \sum_i \left(\frac{\partial W_i^S}{\partial J} + \frac{\partial W_i^S}{\partial M} \frac{dM^S}{dJ} \right) = \frac{\frac{\zeta}{N} \sum_i \frac{\partial W_i^S}{\partial J}}{1 - \frac{\zeta}{N} \sum_i \frac{\partial W_i^S}{\partial M}}. \tag{10}$$

The numerator in the last expression on the right represents the effect on average weight (times $\zeta = .85$) caused by the partial (norm constant) effects of J on individual weights. The denominator embeds the social multiplier effect: the effects in the numerator are amplified by the factor $1/(1 - m)$, where $m = \frac{\zeta}{N} \sum_i \frac{\partial W_i^S}{\partial M}$ is the multiplier itself, as in [3]. It can be shown that the multiplier m is strictly positive and strictly less than one for $J \geq 0$, such that $1/(1 - m) \geq 1$. Therefore the sign on $\frac{dM^S}{dJ}$ follows the sign of the numerator, but has greater magnitude.

The numerator represents the average (times $\zeta = .85$) of the partial effects of J on weight. Since some people gain and some lose weight the (scaled)

average effect depends on the number of gainers vs. losers and the magnitudes of gains and losses. However, within the confines of our functional forms and parameters, we always obtain $\frac{\zeta}{N} \sum_i \frac{\partial W_i^S}{\partial J} < 0$ and thus the sign of (10) is negative.

Turning back to the first expression, (9), representing the equilibrium adjustment of individual weight, we see that for overweight individuals the partial effect of J and the social multiplier effect reinforce each other, and so the value of $\frac{dW_i^S}{dJ}$ is unambiguously negative for initially overweight individuals. For underweight individuals the two effects are opposed, and the sign on $\frac{dW_i^S}{dJ}$ may go either way: for individuals very close to the norm initially, the magnitude of $\frac{\partial W_i^S}{\partial J}$ is very small, and is outweighed by the weight-reducing effect of the decline in M . Despite losing weight some of these individuals (the ones initially just below the norm) will wind up above the new, lower norm. However, there is a threshold value of the metabolic shock such that the individual exactly conforms to the norm. Individuals with shocks above this threshold gain weight on net when J increases, remain below the norm, and yet wind up closer to it than in the initial equilibrium. As J increases, the new perfect conformist will have a greater metabolic shock than the previous conformist.

3 Experiments

The simulation exercises use 50000 agents (50000 values from the initial metabolic shock distribution) to generate equilibrium weight distributions and norms under various specifications of the model. Equilibrium implies instantaneous adjustment of individuals to their stable weight for a given value of the norm, iterating until stable weights are consistent (myopically optimal) in relation to the emergent norm. As detailed below, we calibrate the model to women ages 30 to 60, setting a list of parameters to roughly match average weight for this group observed in the 1990-2002 BRFSS data, and the food budget share estimated by Huang [21]. The goal is to assess the power of the model to explain differences in the average, median, variance and dispersion of the weight distribution across education groups on the basis of variation in the incentive for conformity. To do this we examine the shape of the simulated equilibrium weight distribution at three different levels of J meant to roughly reproduce the different social environments faced by high school dropouts, high school graduates and college graduates. As explained above we are not aware of studies that provide quantitative evidence on the amount of conformity by education. To proceed with the numerical simulations we will select three values of J that result in predictions that approximately match the empirical distributions of body weight, holding all other parameters constant. We then compare our conformity costs model to competing explanations. The alternatives described in the introductory discussion are captured in two additional experiments. We first consider a model with no conformity effects, i.e.

Table 1. Summary of empirical and simulated weight distributions

Distribution	Mean	Std. Dev.	Min	Max	Median	95th ^a	99th ^b	Norm ^c
<i>Empirical Distribution</i>								
<i>I. U.S. Women Age 30-60^d</i>								
BRFSS-Data, 1990-2002	155.3	35.5	51	662	148	223	272	N.A.
High School Dropouts	162.2	38.7	66	651	154	236	288	N.A.
High School Graduates	157.2	36.2	51	603	151	225	277	N.A.
College Graduates	149.2	31.9	59	587	143	205	256	N.A.
<i>II. U.S. Men Age 30-60</i>								
BRFSS-Data, 1990-2002	188.6	36.3	49	738	184	251	304	N.A.
High School Dropouts	183.6	38.7	61	598	179	251	304	N.A.
High School Graduates	190.1	37.4	63	738	184	256	304	N.A.
College Graduates	186.8	33.7	59	636	183	249	293	N.A.
<i>Simulated Distribution - U.S. Women Age 30-60</i>								
<i>I. Endogenous Norms, Group-Specific Deviation Costs^e</i>								
High School Dropouts	161.5	39.6	85	736	153.9	232	294	137.3
High School Graduates	157.8	33.1	88	506	151.9	218	266	134.1
College Graduates	149.2	23.1	93	319	145.9	191	220	126.8
<i>II. Group-Specific Energy Expenditure^f</i>								
High School Dropouts	162.0	40.5	85	786	154.1	234	298	N.A.
High School Graduates	157.1	37.5	83	665	150.0	224	282	N.A.
College Graduates	148.3	32.6	81	520	142.3	207	256	N.A.
<i>III. Group-Specific Exogenous Norms</i>								
High School Dropouts	158.2	33.2	88	508	152.4	219	267	137.9
High School Graduates	157.7	33.1	88	506	151.9	218	266	133.6
College Graduates	157.0	32.9	87	503	151.1	217	264	126.8

Note: ^a95th Percentile. ^b99th Percentile. ^cWeight Norm, computed as 15% below the average weight. ^dSee Fig. 1. ^eSee Fig. 2. ^fSee Fig. 3.

$J = 0$, but with education-group-specific metabolism. This setup captures either genetic differences between educational groups, or behavioral differences based on predetermined preferences or information. Second we consider a model with conformity effects in which we exogenously vary the weight norm. This experiment illustrates the hypothesis that not differences in conformity costs but differences in representation of women with non-western body ideals across educational groups explain the observed distributional variation.

3.1 Calibration of the Model

The metabolism function is parameterized using the Schofield et al. [36] point estimates for the constant term and the coefficient on weight for women in the 30-60 year old group.¹⁰ In the first and third experiment we assume that all individuals regardless of educational groups face the same metabolic shock distribution. In this case, all ϵ are normally distributed with mean zero and

¹⁰ In particular, the values for the parameters are $\gamma = 844$ and $\rho = 8.13382/2.2$.

standard deviation of 0.75.¹¹ The quadratic coefficient of .00025 in expression (2) implies that metabolism for a 200 pound woman is 10 kcal less per day under the quadratic specification than under a (hypothetical) linear model of metabolism, implying a weight increase of 1.04 pounds annually holding caloric intake constant. For a 120 pound woman the corresponding annual weight gain would be about 0.4 pounds. In the second experiment we allow the metabolic shock distribution to differ across educational groups by shifting its mean.

The marginal utility of the first unit of food in a week is identical across individuals and exceeds the marginal utility of the first unit of non-food consumption by 20%.¹² Income and prices are chosen such that the average person is spending about 20% of her income on food purchases, which matches empirical measurements [21]. The price represents the price of 1 pound of body weight, or 3500 calories, which is about the amount burned in 1.5 days by a moderately active 140 pound woman. In nominal terms we use a price of \$36 and an income of \$600 per week (\$31, 200 per year).

3.2 Baseline Calibration: High School Graduates

Women whose highest degree is a high school diploma constitute the largest fraction (about 33%) in the U.S. population and hence serve as our baseline group. Using an intermediate level of $J = 0.001$ for the strength of social interactions ('Endogenous Norms, Group-Specific Deviation Costs'), we find that the model matches the empirical weight distribution very well. For example, the model predicts an average weight of 157.8 pounds, a standard deviation of 33.1, a 95th (99th) percentile weight of 218 (266) pounds and an implied obesity rate of 24.0%. The corresponding BRFSS values for 30-60 year old female high school graduates are 157.2 (mean weight), 36.2 (standard deviation), 225 pounds (95th percentile weight), 277 pounds (99th percentile weight) and 26.3% (obesity rate). Details can be found in Fig. 2 and Table 1. The social body weight norm predicted for high school graduates is 134.1 pounds ($=0.85*157.8$).

3.3 Main Hypothesis: Variation in Strength of Sanction

We now consider whether variation in the coefficient J , which captures the incentive for conformity, can predict the differences in the weight distributions that we observe across the three main groups of educational attainment. Using the high school graduates as the baseline, we chose $J = 0.0001$ for the high school dropouts. This level is one tenth of the one used for high school graduates, suggesting lower costs from deviating from the norm among dropouts.

¹¹ This parametrization of the metabolic shock distribution has an empirical basis: the distribution implies an average absolute deviation that is in line with the mean of the relevant (additive) Schofield residuals.

¹² The coefficients on the preference are $\alpha = 6.0$, $\delta = 0.9$, and $\beta = 5.0$.




Fig. 2. Box plot of predicted body weight distribution, endogenous norm with group-specific deviation costs, U.S. women aged 30-60

For college graduates we consider $J = 0.0035$, i.e. a 3.5-fold increase in the strength of social interaction compared to the baseline group.

As shown in Fig. 2 and Table 1, changes in the strength of social interactions can have visible effects on the weight distribution. The higher value of J for the more educated women implies lower average and median weights and less dispersion of the weight distribution. For the group that faces the greatest costs from not complying with the group standard (college educated women) the model predicts a mean of 149.2 pounds, a median of 145.9 pounds, and a standard deviation of 23.1 pounds. With the greater incentive to conform, we obtain a lower weight norm, 126.8 pounds ($=0.85 \times 149.2$), and a tighter distribution. In the other extreme, the group that faces the weakest non-conformity incentives (high school dropouts) is characterized by a greater average and median weight (161.5 and 153.9 pounds), greater standard deviation (39.6 pounds) and a longer right tail of the distribution (e.g. 95th percentile is at 232 pounds). These predictions, and the implied obesity rates for college educated women and high school dropouts of 13.0% and 27.9% respectively, compare well to the observed obesity rates of 18.4% and 33.9% in the BRFSS. The results also show, however, that the variation in the standard deviation of weight across groups predicted by our hypothesis is greater than the observed one.

3.4 Alternative Hypotheses

We now abstract from social norms, i.e. set J equal to 0, and assume that calorie expenditure differs systematically across educational groups to capture behavioral (e.g., exercising behavior not motivated by weight norms) or genetic differences between educational groups. Genetic differences could have emerged from evolutionary and cultural forces. If determinants that facilitate scholastic achievement (human capital accumulation) are correlated with metabolic differences, either behavioral or genetic, then the pattern in the weight distributions may be better explained by group-specific metabolic differences than by group-specific degrees of sanctioning (main hypothesis). An alternative source of systematic correlation could be greater health awareness among college educated women, leading to more frequent exercise and greater calorie burning on average. To capture these possibilities we assume that mean metabolism varies positively with educational attainment. We predict the weight distribution for high school graduates assuming normally distributed shocks with the same standard deviation as above, but with a mean of $+.10$. For the high school dropouts the distribution is the same as in the initial experiments, and for the college-educated women it is shifted to the right by two-thirds of a standard deviation (mean= $+.30$). The magnitude of the shift is chosen to roughly reproduce the averages in the empirical weight distributions.

The predictions based on this set of explanations ('Group-Specific Energy Expenditure') are shown in Fig. 3 and summarized in Table 1. As in the social sanction hypothesis, the model predicts that the group mean, median, and variance all decline with education level. Compared to the earlier model, this experiment does a better job capturing the variation in the standard deviation across the groups. However, as seen in Fig. 2 as compared with Fig. 1, the social sanctions model comes closer than the group-specific energy expenditure model to capturing the decrease in the inter-quartile range at higher education levels. Finally, we seek to explore the hypothesis that more educated women are more likely to adhere to western ideals of thinness as a result of cultural and ethnic stratification by education. We assume that $J = 0.001$ holds for all groups and exogenously set the value of the social norm at 15% below the group average, yielding norms of 137.9, 133.6, and 126.8 pounds for high school dropouts, high school graduates, and college graduates respectively. Predicted weights for high school dropouts and college graduates are more similar to the high school graduates than observed or predicted by the other hypotheses suggesting that the differences in social norms alone may not suffice to explain the observed differences in the data (see Table 1, 'Group-Specific Exogenous Norms').




Fig. 3. Box plot of predicted body weight distribution, group-specific unobserved heterogeneity, U.S. women aged 30-60

4 Conclusion

This paper presents an agent-based framework with biologically complex agents to explain the differences in the body weight and body mass distributions observed across women with different educational attainment in the U.S.. We show that the smaller average and median weights and smaller dispersion of weights that is observed for more educated women in the U.S. can be explained by two key social factors: stratification of individuals by education level, and greater sanctioning of deviation from the weight norm within the more educated groups.

We modify the framework to illustrate important alternative hypotheses including differences in average energy expenditure and exogenous differences in weight norms across educational groups. The alternatives can represent various sources of correlation between educational attainment and weight including behavioral, genetic, and cultural differences between women in the population. We show that these alternatives can also explain the observation of lower average weight and obesity rate for more educated women, but they do not capture the reduction in dispersion as well as our model does. The reason is that the alternative theories do not readily relate education to variability of weight within groups.

In the simulation of the cultural composition theory, we assumed differences only in the level of weight norms by education, not differences in the

variability of weight norms across groups. There is some evidence from the BRFSS that desired weight varies more among less educated women, consistent with greater dispersion of weight norms among the racially/ethnically more diverse lower educational strata. It is easy to see that allowing for variability of norms within the lower-educated group in the cultural composition experiment may predict variability in weights comparable to that which emerges under our main hypothesis. However, the observed differences in the variability of desired weights in the data appear to be small, and the effect of norm variability within groups on average weight would in general be ambiguous. Reproducing the stylized facts in this manner would require multiple tenuous assumptions, whereas the conformity costs model offers a more parsimonious explanation with further potential to explain changes in norms over time with changing economic incentives.

References

1. Anderson, J.L., Crawford, C.B., Nadeau, J., and T. Lindberg (1992) Was the Dutchess of Windsor right? A cross-cultural review of the socioecology of ideals of female body shape. *Ethology and Sociobiology*, 13, 197–227.
2. Averett, S., and S. Korenman (1996) The Economic Reality of the Beauty Myth. *Journal of Human Resources*, 31, 304–30.
3. Becker, G.S., and K.M. Murphy (2000) Social Economics: Market Behavior in a Social Environment. Harvard University Press, Cambridge, MA.
4. Behrman, J.R., and M.R. Rosenzweig (2001) The Returns to Increasing Body Weight. University of Pennsylvania Working Paper.
5. Bernheim, D.B. (1994) A Theory of Conformity. *Journal of Political Economy*, 102(5), 841–877.
6. Bhattacharya, J., and M.K. Bundorf (2005) The incidence of the healthcare costs of obesity. NBER Working Paper No. 11303.
7. Blume, L.E., and S. Durlauf (2004) The Interactions-Based Approach to Socioeconomic Behavior. In: Durlauf, S., and H.P. Young, eds., *Social Dynamics*, MIT Press, Cambridge, MA.
8. Brock, W., and S. Durlauf (2001) Discrete Choice with Social Interactions. *Review of Economic Studies*, 68 (2), 235–260.
9. Brownson, R.C., Eyster A.A., King A.C., Brown, D.R., Shyu, Y.-L., and J.F. Sallis (2000) Patterns and correlates of physical activity among U.S. women 40 years and older. *American Journal of Public Health*, 90(2), 264–70.
10. Burke, M.A., and K. Prasad (2003) Contracts with Social Multipliers. Florida State University working paper, http://mailer.fsu.edu/~mburke/Contracts_with_social_multipliers.pdf.
11. Burke, M.A., and F. Heiland (2005) Social Dynamics of Obesity. Brookings CSED Working Paper No. 40, http://www.brookings.edu/es/dynamics/papers/CSED_wp40.htm.
12. Carr, D., and M. Friedman (2005) Is Obesity Stigmatizing? Body Weight, Perceived Discrimination, and Psychological Well-Being in the United States. *Journal of Health and Social Behavior*, 46(3), 244–259.

13. Cawley, J. (2004) The Impact of Obesity on Wages. *Journal of Human Resources*, 39, 451–74.
14. Chou, S. Y., Grossman, M., and H. Saffer (2004) An economic analysis of adult obesity: results from the Behavioral Risk Factor Surveillance System. *Journal of Health Economics*, 23, 565–587.
15. Crandall, C. (1988) The social contagion of binge eating. *Journal of Personality and Social Psychology*, 55: 589599.
16. Cutler, D. M., Glaeser, E. L., and J. M. Shapiro (2003) Why have Americans Become More Obese? *Journal of Economic Perspectives*, 17, 93–118. Also NBER Working Paper No. 9446.
17. Ezzati, M., Lopez, A.D., Rodgers, A., Hoorn, S.V., Murray, C.J.L., and the Comparative Risk Assessment Collaborating Group (2002) Selected major risk factors and global and regional burden of disease. *The Lancet*, October 30.
18. Flegal, K.M., Graubard, B.I., Williamson, D.F., and M.H. Gail (2005) Excess Deaths Associated With Underweight, Overweight, and Obesity. *Journal of the American Medical Association*, 293:1861-1867.
19. Frank, R. (1999) *Luxury Fever: Why Money Fails to Satisfy in an Era of Excess*. New York: Free Press.
20. Horgan, G.W., and J. Stubbs (2003) Predicting basal metabolic rate in the obese is difficult. *European Journal of Clinical Nutrition*, 57, 335–340.
21. Huang, K.S. (1993) A Complete System of U.S. Demand for Food. USDA Economic Research Service, Technical Bulletin 1821.
22. Lakdawalla, D., and T.J. Philipson (2002) The growth of obesity and technological change: a theoretical and empirical examination. NBER Working Paper No. 8946.
23. Leibel, R.L., Rosenbaum, M., and J. Hirsch (1995) Changes in Energy Expenditure Resulting from Altered Body Weight. *The New England Journal of Medicine*, 332(10), 621–628.
24. Levy, A. (2002) Rational eating: can it lead to overweightness or underweightness? *Journal of Health Economics*, 21, 887–899.
25. Lewis, S.K., and V.K. Oppenheimer (2000) Educational assortative mating across marriage markets: non-Hispanic whites in the United States. *Demography*, 37(1), 29–40.
26. Must, A., Spadano, J., Coakley, E.H., Field, A.E., Colditz, G., and W.H. Dietz (1999) The disease burden associated with overweight and obesity. *Journal of the American Medical Association*, 282(16), 1523–1529.
27. Myers, A., and J.C. Rosen (1999) Obesity stigmatization and coping: relation to mental health symptoms, body image, and self-esteem. *International Journal of Obesity and Related Metabolic Disorders*, 23(3), 221–30.
28. Pagan, J. A., and A. Davila (1997) Obesity, Occupational Attainment, and Earnings. *Social Science Quarterly*, 78, 756–70.
29. Pencavel, J. (1998) Assortative Mating by Schooling and the Work Behavior of Wives and Husbands. *American Economic Review*, 88(2), Proceedings, 326–329.
30. Philipson, T.J., and R.A. Posner (1999) The long-run growth in obesity as a function of technological change. NBER Working Paper No. 7423.
31. Pullicino, E., Copperstone, C., Luzi, L., McNeill, G., and M. Elia (1996) Relationship between anthropometric indices of body fat distribution and basal energy metabolism in healthy Maltese women. *Acta Diabetology*, 33, 198–204.

32. Rand, W. M. (1982) The estimation of BMR. Rome: Joint FAO/WHO/UNU. Unpublished Report.
33. Robinson, J. and N. Maheshwari (2005) presentation at American Heart Association's 45th Annual Conference on Cardiovascular Disease Epidemiology and Prevention, Washington, DC.
34. Rose, E. (2004) Education, Hypergamy and the Success Gap. University of Washington unpublished working paper.
35. Ross, C.E. (1994) Overweight and Depression. *Journal of Health and Social Behavior*, 35(1), 63–79.
36. Schofield, W.N., Schofield, C., and W.P.T. James (1985) Predicting basal metabolic rate, new standards review of previous work. *Human Nutrition: Clinical Nutrition* 39C, Suppl. 1, 5–41.
37. Stearns, P.N. (1997) *Fat History: Bodies and Beauty in the Modern West* New York: New York University Press.
38. Troiano, R.P., Frongillo, E.A. Jr, Sobal, J., and D.A. Levitsky (1996) The relationship between body weight and mortality: a quantitative analysis of combined information from existing studies. *International Journal of Obesity and Related Metabolic Disorders*, 20(1), 63–75.
39. Villanueva, E.V. (2001) The validity of self-reported weight in U.S. adults: a population based cross-sectional study. *BMC Public Health*, 1-11
40. Wansink, B. (2004) Environmental factors that unknowingly increase food intake and consumption volume of unknowing consumers. *Annual Review of Nutrition*, 24, 455–479.

Agent-Based Models in Ecology: Patterns and Alternative Theories of Adaptive Behaviour

Volker Grimm¹ and Steven F. Railsback²

¹ UFZ Centre for Environmental Research Leipzig-Halle, Germany

`Volker.Grimm@ufz.de`

² Lang, Railsback & Associates, Arcata, California, USA `LRA@northcoast.com`

Summary. Ecologists have used agent-based models for a long time, but refer to them as “individual-based models” (IBMs). Common characteristics of IBMs are discrete representation of unique individuals; local interactions; use of adaptive, fitness-seeking behaviour; explicit representation of how individuals and their environment affect each other; and representation of full life cycles.

Ecology has contributed to agent-based modelling in general by showing how to use agent-based techniques to explain real systems. Ecologists have used IBMs to understand how dynamics of many real systems arise from traits of individuals and their environment. Two modelling strategies have proven particularly useful.

The first strategy is “pattern-oriented modelling” (POM). POM starts with identifying a variety of observed patterns, at different scales and at both individual and system levels, that characterize the system’s dynamics and mechanisms. These patterns, along with the problem being addressed and conceptual models of the system, provide the basis for designing and testing an IBM. A model’s variables and mechanisms are chosen because they are essential for reproducing these characteristic patterns. After an IBM is assembled, alternative versions (different theories for individual behaviour; different parameterizations) can be tested by how well they reproduce the patterns.

The second strategy is developing general and reusable theory for the adaptive behaviour of individuals. A “theory” is a model of some specific individual behaviour from which system-level dynamics emerge. Theory can be developed by hypothesizing alternative models for the behaviour, then using the IBM to see which alternative best reproduces a variety of patterns that characterize the system dynamics of interest. Empirical observations are used to develop both theories and the patterns used to test and falsify them.

These two strategies are demonstrated with example models of schooling behaviour in fish, spatiotemporal dynamics in forests, and dispersal of brown bears.

1 Introduction

Agent-based modelling has a long tradition in ecology but under a different label: individual-based modelling [13] [5] [8]. Due to the different histories of

agent-based [3] and individual-based modelling [11], they differ in definition and rationale, but in recent years the two approaches are merging [22]. In ecology, the original focus was on three facts that are ignored in analytical ecological models: individuals are discrete entities, they usually interact locally, and they are not all the same. However, some authors in ecology used the term “individual-based” very broadly, for example for models with individuals that are discrete but identical and lacking an explicit life cycle. Therefore Uchmański and Grimm [32] proposed a definition of “individual-based models” (IBMs) which includes discreteness, uniqueness and full life cycles of individuals. Railsback [22] identified an additional characteristic of IBMs: individuals are autonomous and show adaptive behaviour, i.e., they adapt their behaviour (growth, feeding, habitat selection, mate choice, etc.) to the state of themselves and their environment, to seek higher fitness. Fitness is either modelled explicitly [26] or implicitly by assuming that under certain conditions a certain behaviour improves fitness.

Regarding adaptive behaviour, most individual-based ecological models are still far behind agent-based models. Adaptive behaviour is either ignored, as in most plant IBMs, or behaviour is imposed, or adaptation, decision making, and fitness considerations are only implicitly assumed. Ecology has a lot to learn from agent-based modelling not only regarding adaptation but also regarding software platforms and concepts for representing individuals, their decisions, and their interactions [3].

On the other hand, ecological IBMs usually address natural systems, not engineered, designed, or artificial systems, so considerable know-how has accumulated in ecology regarding the design, analysis, verification, and validation of models describing real systems. “Getting results” [23], i.e. achieving understanding, testable predictions, management support, etc., from an IBM or ABM is not at all trivial. The step from just interesting demonstrations to insights into how real systems work and real agents behave requires appropriate strategies; many ABM projects, especially by beginners in modelling, seem stuck at the demonstration level.

Here we summarize two linked strategies which were developed in ecology [11] and are likely to be useful in agent-based modelling in general: using multiple patterns to design, test, validate, and parameterize ABMs; and contrasting alternative theories of behavioural traits. These strategies are explained in the following two Sects. and then example applications are presented.

2 Pattern-Oriented Modelling

One of the main problems in developing agent-based models is to find the appropriate level of detail: how many state variable should we use to describe agents? How variable should the model world be in space and time, and how much information should we assume an agent knows and uses to make decisions? If the model is too simple we might neglect essential mechanism of

the real system, which limits the model's potential to provide understanding and testable predictions. If the model is too complex, model analysis will be cumbersome and we might get bogged down in too much detail.

Pattern-oriented modelling [7] [10] [23] [9] [34] [11] is to use the indicators of internal organisation provided by the real systems themselves: patterns. A pattern is anything above random variation and thus indicates some kind of internal organisation. We should therefore choose a model's detail so it (in principle) allows the observed patterns to emerge in the model. This is the principal way natural sciences proceed: trying to explain observed patterns by developing models, or theories, which reproduce the patterns.

For example, early theories of the structure of the hydrogen atom were pointing in the right direction (a separation of nucleus and electron) but were incapable of reproducing many observed patterns. Patterns in atomic spectra, the discrete spectral lines, forced physicists to assume structures and mechanisms (i.e., quantum mechanics) which reproduced these discrete spectra. Similarly, patterns indicating internal organisation were the key that revealed the structure of DNA [33].

For developing agent-based models we need to ask from the very beginning: what characteristic patterns can we observe and how can we make sure that these patterns can also emerge in the model? If there is a spatial pattern, the model should be spatially explicit; if there are patterns in age or size structures, age and size should be variables describing individuals; if we know characteristic dispersal patterns, dispersal should be an explicit model process; if we have data about how the system responds to specific disturbances, e.g. a drought, then precipitation and response of individuals to different water availabilities should be in the model; if we know that for a certain plant community total biomass per spatial unit remains within certain limits, we should use biomass as an additional variable of individuals; if we know that at high densities individuals are more aggressive, we should think about how we model interaction and aggression; etc. Thus, observed patterns to a large degree can help us decide which structures (variables) and processes of a real system we include in the model.

Of course, the other main determinants of model structure beside patterns are the question we want to answer with the model and our current hypotheses about how the system works, i.e. our conceptual model of the system [6]. However, by including variables and processes that are motivated by observed patterns, our model automatically has a level of detail linked to the system's internal organisation. In this way, the model becomes testable at different hierarchical levels. We can test alternative assumptions about the model entities and their behaviour by how well they reproduce observed patterns. Of particular importance in IBMs are the models, or theories, for adaptive behaviour.

3 Developing Theory of Adaptive Behaviour

Individuals or agents adapt their behaviour to their current state and environment. In ecology, we assume that adaptive behaviour is a product of natural selection: individuals have behavioural traits that improve their fitness, i.e., increase their expectation of passing their genes to future generations. However, it is not easy to formulate realistic models of these traits, i.e. the set of decision rules used by the individuals: what do individuals know, what do they assume, and how do they predict the consequences of their decisions? Even when we think organisms do not actually identify and predict the consequences of alternative decisions, evolution has often provided organisms with behaviours that can be modeled as if they do (e.g., [26]).

To develop useful, general theory of adaptive behaviour, we propose using the scientific method of strong inference [19] by contrasting alternative theories [11]. We refer to “theories” instead of “models” to emphasize that the aim is to find general and tested theories which can be used to model the same traits of other species or in other environments. In the theory development cycle we propose, we first define the trait and ecological context of interest, for example habitat selection of European lynx during dispersal [30]; then we formulate alternative theories of the trait, including “null theories” that contain no adaptive ability; next we identify patterns at both the individual and system level (see previous Sect.); and finally we implement the alternative theories in an IBM and test how well they reproduce the patterns. Thus, we use the full IBM as a kind of “virtual laboratory” to test alternative theories of a certain adaptive behaviour. Usually, the steps of the cycle have to be repeated several times because first insights gained from the model lead to new theories, additional patterns, or modification of the entire IBM.

Testing alternative theories has several benefits: we are forced to be explicit about the theories and the way we test them; we can demonstrate how significant the specific formulation of a certain trait is for explaining observed patterns; in particular, we can demonstrate whether null theories lead to unrealistic results; and we can refine theories by identifying additional patterns.

The advantages of testing theories of adaptive traits become obvious when we compare two IBMs addressing the same phenomenon, schooling in fish. The famous “Boids” (=bird-oids) model of Reynolds [29] reproduces schooling-like behaviour. It is based on simple assumptions or, in our terminology, theories for individual behaviour: individuals try to avoid collision, to match the velocity of neighbouring individuals, and to stay close to neighbours. This simple theory indeed leads to the emergence of school-like aggregations, as seen in the numerous implementations of Boids available on the Internet (Fig. 1). However, Boids is only a demonstration of how simple behavioural rules and local interaction give rise to a collection of individuals which are more or less regularly spaced and move as one coherent entity. It does not allow us to infer how real fish or birds behave, and we do not know which elements of the boids behaviour are essential to reproduce schooling behaviour.

The fish school of Huth and Wissel [15] addressed a different question than Boids: how can we use observed patterns, data, and an IBM to learn about how real fish behave? The model is very similar in structure to Boids but the entire modelling strategy is different.

First, specific currencies were used to quantify properties of real fish schools: polarization p and nearest neighbour distance NND . Polarization is the average angle of deviation between the swimming direction of each fish and the mean direction of the entire school; p is 0° if all fish swim in the same direction and p approaches 90° if all fish swim in random directions. Values of p observed in real fish schools are $10\text{--}20^\circ$. The NND is the average distance between a fish and its nearest neighbor. In real fish schools, NND is typically about 1-2 times the average body length of a fish. The observed values of these two currencies are the patterns towards which the entire IBM is oriented.




Fig. 1. Fish school model, based on the Boids model of Reynolds [29] and developed and implemented by H. Hildenbrandt (figure courtesy of H. Hildenbrandt).

Second, alternative theories were formulated for how fish decide where and how fast to swim in the next time increment. For example, it seems reasonable to assume that fish average the influence of their nearest neighbours; this

assumption is also made in Boids [29]. But how many neighbours should be taken into account: a certain number or all neighbours within a given distance? It is also conceivable that fish first decide on one single neighbour fish, for example the closest one or the one swimming both close *and* in front of the fish, and then adjust their swimming to this single fish. Huth and Wissel [15] [14] formulated 11 alternative theories of swimming behaviour. Nine of these were based on an averaging decision, two on a priority decision. It turned out that priority decisions failed to reproduce realistic polarization values (Fig. 2).

What we also see in Fig. 2 is that looking at only one pattern may not be sufficient to identify the better theory: looking at *NND* alone might suggest that both priority and averaging models produce similar results, but in fact priority theories produce schools which are only as compact, but not as polarized as real schools. Unfortunately often in ecological theory, the focus is on only one pattern, for example cycles in abundance or levels of primary production, so the resulting models may still be quite unrealistic in structure and mechanism. If we use multiple patterns at different hierarchical levels and test alternative theories against these patterns, we are more likely to end up with models which are “structurally realistic” [34] and “mechanistically rich” [6].

Realism and richness in structure and mechanism also allow models to be validated with independent or secondary patterns which were not used to develop, parameterize, or verify the model. After testing their model, Huth and Wissel [15] [14] searched the empirical literature for additional patterns. Some were more qualitative, for example regarding the shape of fish schools; others were more quantitative, for example the *NND* of the first, second, and third nearest neighbour. In almost all empirical patterns, Huth and Wissel found a good, at least qualitative match between model and data. (The Huth-Wissel model is discussed in more detail in [4] and [11].)

4 Examples

The fish school example from the previous Sect. is from ethology, not ecology. Here we first present two IBM projects from ecology that are pattern-oriented or use alternative theories of adaptive behaviour. More than 30 such example IBMs are discussed in [11]. The third example IBM is also pattern-oriented, but its focus is on how patterns can be used to indirectly determine model parameters that are unknown or not directly accessible. For detailed descriptions of the three models, see the original publications and corresponding information on the internet.

4.1 Mid-European Natural Beech Forests

Forest IBMs have been widely used since the 1970s. The JABOWA forest model [2] [31] [1] is probably the most successful IBM because it gave rise

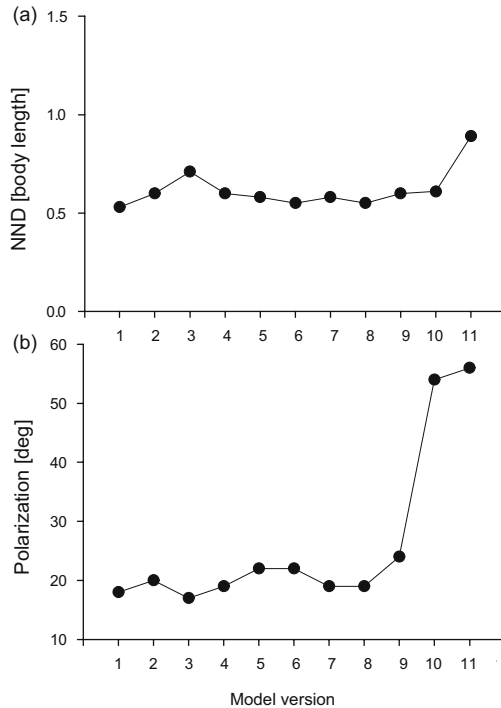


Fig. 2. Comparison of 11 alternative theories of the schooling behaviour of fish. Theories 10 and 11 are based on selecting a single neighbour fish for orientation, whereas theories 1-9 average the influence of neighbours within a certain range. The theories are tested for the two currencies, or patterns: nearest neighbour distance (*NND*, in body lengths) and polarization (in degrees). (After [14].)

to a large family of so-called “gap models” [16]. However, early gap models were not designed to address spatial patterns in forests, but temporal patterns (succession); and more recent gap models are spatially explicit, quite complex, and not easy to implement or parameterize. Therefore, Wissel [37] developed a very simple cellular automaton model to explain the mosaic structure of natural beech forests [28]. However, foresters were not convinced by this model, mainly because it was too poor in structure and mechanism. There was no way to analyse the model forest on different hierarchical levels to identify and test secondary predictions.

Neuert [17] developed a new beech forest model, BEFORE, which took into account not only the mosaic pattern but also patterns in the vertical structure. The different developmental stages are characterized by different vertical structures. For example, the “optimal stage” is characterized by a closed canopy layer and almost no understory. Moreover, the conceptual model of

many foresters about beech forests includes vertical structure and the generation and closing of gaps in the canopy as factors driving forest dynamics.

Following the strategy of pattern-oriented modelling, Neuert [17] included the vertical spatial dimension in BEFORE so that vertical structures could also emerge (Fig. 3). The behaviour of individual trees is described by empirical rules because foresters know quite well how growth and mortality depend on the local environment of a tree. Likewise, empirical information is available to define rules for the interaction of individuals in neighbouring spatial units.




Fig. 3. Vertical and horizontal structure of the beech forest model BEFORE [17] for an area of 10×10 grid cells. Within a grid cell, which has the area of a large canopy tree, four different vertical layers or size classes of trees are distinguished and characterized by their cover, ranging from 0% (white) to 100% (black) (A). For comparing the model results to horizontal patterns (mosaic structure) of real forests, patterns in the vertical structure are assigned to three different developmental stages (B). The resulting mosaic pattern is smoothed by a moving average (C). (After [20].)

BEFORE not only reproduces the mosaic and vertical patterns, it is also so rich in structure and mechanism that it could produce independent predictions of forest characteristics not considered at all during model development and testing [18] [20]. These predictions were about the age structure of the canopy and the spatial distribution of very old and large trees. All these predictions were in good agreement with observations, considerably increasing the model's credibility.

The use of multiple patterns to design the model obviously led to a model which was “structurally realistic”. This realism had the additional benefit of allowing the addition of rules to track woody debris, a process not originally considered in the model. Again, the amount and spatial distribution of woody debris in the model forest were in good agreement with observations of natural forest and old forest reserves [21].

This example shows that the usual complaint about the complexity of IBMs and ABMs is not necessarily true: if multiple patterns are used to design and verify a model, its complexity may no longer be only a burden making model analysis and understanding difficult, but can become a rich source for increasing the model's credibility and gaining understanding. True, the time frame for developing and testing models like BEFORE is much longer than for simple analytical models or for mere demonstrations, but once the modelling project is in the "exploitation phase" it can be used to rapidly and productively address a wide range of applied and basic questions.

4.2 Habitat Selection in Stream Trout

In rivers, dams and water diversions affect the flow regime and thus fish populations. Fish adapt to changes in flow by moving to different habitat, so to predict how fish populations react to new flow regimes we need to know how fish select their habitat. Habitat suitability (or "preference") models are widely used to predict habitat selection, but are limited because they are based on empirical relationships between observed distributions of fish and habitat characteristics. These empirical relationships are valid only for the conditions under which observations were made and may not be valid for altered conditions [27].

Therefore, IBMs of stream fish have been developed to model habitat selection and the consequent response of populations to flow alteration (e.g. [36]). These IBMs attempt to capture the important processes determining survival, growth, and reproduction of individual fish, and how these processes are affected by river flow. For example, mortality risks and growth of trout depend on habitat variables (depth, velocity, turbidity, etc.) and fish size; moreover, competition among trout resembles a size-based dominance hierarchy. However, existing foraging theory (e.g., that habitat is selected to maximize growth or minimize the ratio of growth to risk) cannot explain the ability of trout to make good tradeoffs between growth and risk in selecting habitat under a wide range of conditions. Therefore Railsback et al. [26] developed a new theory, which is based on the assumption that fish select habitat to maximize the most basic element of fitness: the probability of surviving over a future period. This new "state-based, predictive" theory was tested by demonstrating that it could reproduce, in a trout IBM, a wide range of habitat selection patterns observed in real trout populations [24] (Table 1).

The resulting structural realism of the trout IBM made it a powerful tool for addressing many basic and applied questions [25] [27]. The trout IBM reproduced system-level patterns observed in real trout including self-thinning relationships, periods of high density-dependent mortality among juveniles, density-dependence in juvenile size, and effects of habitat complexity on population age structure. In a management application, the trout IBM was used to predict the population-level consequences of stream turbidity [12]. Turbidity (cloudiness of the water) reduces both food intake and predation risk. The

Table 1. Pattern-oriented test of three alternative habitat selection theories for a trout IBM. Only the state-based, predictive theory caused all six observed patterns to be reproduced in the IBM. (After: [24] [11].)

Observed pattern	Maximize growth	Maximize survival	State-based, predictive
Hierarchical feeding	+		+
Response to high flow	+	+	+
Response to inter-specific competition	+		+
Response to predatory fish		+	+
Seasonal velocity preference			+
Response to reduced food availability			+

population-level consequences of these two offsetting individual-level effects would be very difficult to evaluate empirically, but was easily predicted using the IBM: over a wide range of parameter values, the negative effects of turbidity on growth (and, consequently, reproduction) outweighed the positive effects on risk.

4.3 Patterns for Parameterization: Spread of Brown Bears into the Alps

A major problem of agent-based modelling of real systems is parameterization. In virtually all cases, many parameters are uncertain or even unknown. Consequently, model results are uncertain and predictions and insights from the model become questionable. For example, in ecology many questions regarding population dynamics must include spatial heterogeneity of the landscape. However, it has been asked whether spatially explicit population models (SEPMs), which are often individual-based, are made useless by parameter uncertainty and error propagation.

Wiegand et al. [35] argue that this is not necessarily so. If a model is structurally realistic it captures key structures and processes of the real system (which by itself reduces the importance of parameter values). Then, observed patterns can be used to reduce parameter uncertainty. Wiegand et al. demonstrate this indirect parameter estimation (in other disciplines referred to as “inverse modelling”) with their SEPM of brown bears (*Ursus arctos*) spreading from Slovenia into the Alps. First, they determine parameter values from sparse data from the Alps and other regions (Scandinavia, Spain, USA) and

varied the remaining unknown parameters over wide ranges. A global sensitivity analysis of this uncalibrated parameter set revealed high uncertainty in model output (coefficient of variation of about 0.8).

Then, two data sets were used to identify five patterns. These patterns are (1) fluctuations in the abundance of females with cubs observed over 10 years, (2) the low density of females in a certain transition area in the Alps, and the density of bears observed in (3) central Austria, (4) the Carnic Alps, and (5) the Karawanken. For these five patterns, quantitative criteria were defined to obtain confidence intervals for the agreement between observed and simulated patterns.

The indirect parameter estimation started with 557 random parameter sets covering the entire range of parameter uncertainty. Then the five observed patterns were used as filters: parameter sets which failed to reproduce a pattern were discarded. The five filters were used alone and in combination (Table 2). The best combination of filter patterns reduced the number of feasible parameter sets to 10. The global sensitivity of this remaining data set was reduced to a coefficient of about 0.2.

Table 2. Filters defined through observed patterns in the distribution and abundance of brown bears spreading from Slovenia into the Alps (after [35]).

Filter description	Filters	Number of model parameterizations in agreement with observed pattern
No filter	0	557
Density of females in transition area	1	506
Bear observation in central Austria	2	138
Bear observation in the Carnic Alps	3	154
Bear observation in the Karawanken	4	180
Census time series of females with cubs	5	12
	2+3+4	13
	5+1	10
	2+3+4+1	11

It is noteworthy that both the census time series (filter 5 in Table 2) and the three spatial filters (2+3+4 together) narrowed down parameter uncertainty to a similar degree. This is reassuring because if a model really captures the essence of a system then different patterns at the population level should be redundant because they reflect the same underlying processes.

5 Discussion

Agent-based models in ecology, usually referred to as individual-based models (IBMs), have generally differed in two main aspects from agent-based models in other disciplines. First, adaptive behaviour and decision-making has not been a main issue (although this certainly will change in the future [22] [3]). Second, in ecology IBMs are more often oriented towards understanding and managing real systems, not hypothetical or “what-if” situations. This emphasis on real systems led to the development of a strategy—pattern-oriented modelling (POM)—to design IBMs so that they are structurally realistic and can be verified and validated.

We present POM and how it leads to IBMs which can be used as virtual laboratories to test alternative theories. However, there is a chicken-or-egg problem: we implement theories in an IBM to test them, but suitable IBMs must be based on tested theories—models of adaptive behaviour which are proven capable of reproducing observed patterns. A flawed IBM could lead to wrong conclusions about theories tested in it. One solution to this problem is that not more than one theory at a time should be tested in an IBM. Moreover, we cannot model all aspects of individual behaviour with the same resolution. The behaviours that are key to the questions we are addressing with a model should be modelled as adaptive traits: individuals decide what to do next depending on their current state and the state of their environment, to improve their expected fitness. Other traits can simply be imposed by the modeller. For example, in the brown bear model [35] habitat selection during each time step of dispersal was modelled as an adaptive process: the decision where to move next was based on the attractiveness of habitat in the neighbourhood. But mortality per time step was imposed—simply assumed to be constant. Finally, we cannot overstate the importance of testing a model and its theories by how many characteristic patterns it can reproduce at a variety of scales and at both individual and system levels. Much higher confidence in an IBM is justified when the model can be shown to reproduce a wide diversity of such patterns.

Pattern-oriented modelling is not really new. Many experienced modellers use observed patterns to design and test their models. What is new is our attempt to formalise the use of multiple patterns and the contrasting of alternative theories for behavioural traits as an explicit strategy for developing IBMs and ABMs. We hope that this strategy will contribute to make individual-based and agent-based modelling more efficient and coherent [11].

References

1. D.B. Botkin (1993) *Forest dynamics: an ecological model*. Oxford University Press, Oxford, New York.
2. D.B. Botkin, J.F. Janak, and J.R. Wallis (1972) Some ecological consequences of a computer model of forest growth. *Journal of Ecology*, 60:849–873.

3. F. Bousquet and C. Le Page (2004) Multi-agent simulations and ecosystem management: a review. *Ecological Modelling*, 176:313–332.
4. S. Camazine, J.-L. Deneubourg, N.R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau (2001) Self-organization in biological systems. Princeton Studies in Complexity. Princeton University Press, Princeton, New Jersey.
5. D.L. DeAngelis and L.J. Gross (1992) Individual-based models and approaches in ecology: populations, communities and ecosystems. In: D.L. DeAngelis and L.J. Gross, editors, *Individual-based models and approaches in ecology*, pp. 523–525. Chapman and Hall, New York.
6. D.L. DeAngelis and W.M. Mooij (2003) In praise of mechanistically-rich models. In: C.D. Canham, J.J. Cole, and W.K. Lauenroth, editors, *Models in ecosystem science*, pp. 63–82. Princeton University Press, Princeton, New Jersey.
7. V. Grimm (1994) Mathematical models and understanding in ecology. *Ecological Modelling*, 75/76:641–651.
8. V. Grimm (1999) Ten years of individual-based modelling in ecology: What have we learned, and what could we learn in the future? *Ecological Modelling*, 115:129–148.
9. V. Grimm and U. Berger (2003) Seeing the forest for the trees, and vice versa: pattern-oriented ecological modelling. In: L. Seuront and P.G. Strutton, editors, *Handbook of scaling methods in aquatic ecology: measurement, analysis, simulation*, pp. 411–428. CRC Press, Boca Raton.
10. V. Grimm, K. Frank, F. Jeltsch, R. Brandl, J. Uchmański, and C. Wissel (1996) Pattern-oriented modelling in population ecology. *Science of the Total Environment*, 183:151–166.
11. V. Grimm and S.F. Railsback (2005) Individual-based modeling and ecology. Princeton University Press, Princeton, N.J.
12. B.C. Harvey and S.F. Railsback. Elevated turbidity reduces abundance and biomass of stream trout in an individual-based model. in prep.
13. M. Huston, D. DeAngelis, and W. Post (1988) New computer models unify ecological theory. *BioScience*, 38:682–691.
14. A. Huth (1992) Ein Simulationsmodell zur Erklärung der kooperativen Bewegung von polarisierten Fischeschwärmen. Phd, Universität Marburg.
15. A. Huth and C. Wissel (1992) The simulation of the movement of fish schools. *Journal of Theoretical Biology*, 156:365–385.
16. J. Liu and P.S. Ashton (1995) Individual-based simulation models for forest succession and management. *Forest Ecology and Management*, 73:157–175.
17. C. Neuert (1999) Die Dynamik räumlicher Strukturen in naturnahen Buchenwäldern Mitteleuropas. Phd, Universität Marburg.
18. C. Neuert, C. Rademacher, V. Grundmann, C. Wissel, and V. Grimm (2001) Struktur und Dynamik von Buchenurwäldern: Ergebnisse des regelbasierten Modells BEFORE. *Naturschutz und Landschaftsplanung*, 33:173–183.
19. J.R. Platt (1964) Strong inference. *Science*, 146(3642):347–352.
20. C. Rademacher, C. Neuert, V. Grundmann, C. Wissel, and V. Grimm (2004) Reconstructing spatiotemporal dynamics of central european beech forests: the rule-based model BEFORE. *Forest Ecology and Management*, (194: 349-368).
21. C. Rademacher and S. Winter (2003) Totholz im Buchen-Urwald: generische Vorhersagen des Simulationsmodelles BEFORE-CWD zur Menge, räumlichen Verteilung und Verfügbarkeit. *Forstwissenschaftliches Centralblatt*, 122:337–357.

22. S.F. Railsback (2001) Concepts from complex adaptive systems as a framework for individual-based modelling. *Ecological Modelling*, 139:47–62.
23. S.F. Railsback (2001) Getting “results”: the pattern-oriented approach to analyzing natural systems with individual-based models. *Natural Resource Modeling*, 14:465–474.
24. S.F. Railsback and B.C. Harvey (2002) Analysis of habitat selection rules using an individual-based model. *Ecology*, 83:1817–1830.
25. S.F. Railsback, B.C. Harvey, R.H. Lamberson, D.E. Lee, N.J. Claasen, and S. Yoshihara (2002) Population-level analysis and validation of an individual-based cuthroat trout model. *Natural Resource Modeling*, 14:465–474.
26. S.F. Railsback, R.H. Lamberson, B.C. Harvey, and W.E. Duffy (1999) Movement rules for individual-based models of stream fish. *Ecological Modelling*, 123:73–89.
27. S.F. Railsback, H.B. Stauffer, and B.C. Harvey (2003) What can habitat preference models tell us? tests using a virtual trout population. *Ecological Applications*, 13:1580–1594.
28. H. Remmert (1991) The mosaic-cycle concept of ecosystems - an overview. In: H. Remmert, editor, *The mosaic-cycle concept of ecosystems (Ecological Studies 85)*, pp. 1–21. Springer, Berlin Heidelberg New York.
29. C.W. Reynolds (1987) Flocks, herds, and schools: a distributed behavioral model. *Computer Graphics*, 21:25–36.
30. S. Schadt, F. Knauer, P. Kaczensky, E. Revilla, T. Wiegand, and L. Trepl (2002) Rule-based assessment of suitable habitat and patch connectivity for the eurasian lynx. *Ecological Applications*, 12:1469–1483.
31. H.H. Shugart (1984) A theory of forest dynamics: the ecological implications of forest succession models. Springer-Verlag, New York.
32. J. Uchmański and V. Grimm (1996) Individual-based modelling in ecology: what makes the difference? *Trends in Ecology and Evolution*, 11:437–441.
33. J. Watson (1968) The double helix: a personal account of the discovery of the structure of DNA. Atheneum, New York.
34. T. Wiegand, F. Jeltsch, I. Hanski, and V. Grimm (2003) Using pattern-oriented modeling for revealing hidden information: a key for reconciling ecological theory and conservation practice. *Oikos*, 100:209–222.
35. T. Wiegand, E. Revilla, and F. Knauer (2004) Dealing with uncertainty in spatially explicit population models. *Biodiversity and Conservation*, 13:53–78.
36. W. Van Winkle, H.I. Jager, S.F. Railsback, B.D. Holcomb, T.K. Studley, and J.E. Baldrige (1998) Individual-based model of sympatric populations of brown and rainbow trout for instream flow assessment: model description and calibration. *Ecological Modelling*, 110:175–207.
37. C. Wissel (1992) Modelling the mosaic-cycle of a Middle European beech forest. *Ecological Modelling*, 63:29–43.

Agent-Based Modelling of Self-Organisation Processes to Support Adaptive Forest Management

Ernst Gebetsroither, Alexander Kaufmann, Ute Gigler, and Andreas Resetarits

ARC systems research GmbH, Seibersdorf, Austria
Ernst.Gebetsroither@arcs.ac.at
Alexander.Kaufmann@arcs.ac.at
Ute.Gigler@arcs.ac.at
Andreas.Resetarits@arcs.ac.at

Summary. Managing the numerous and interrelated processes between man and nature in order to use renewable resources in a sustainable way is confronted with conflicting objectives, external effects, complex interdependencies, uncertainty and other features that make it nearly impossible to come to unambiguous optimal decisions. Self-organisation in socioeconomic and ecological systems - the process of structuring a system by the elements of the system themselves without hierarchical or external control - is often the reason for ambiguity and uncertainty. Adaptive management is an approach to deal with these challenges. This natural resource management method is permanently monitoring both socioeconomic and ecological systems in order to be able to react rapidly on any development pushing the systems into an undesired direction. Understanding and simulating the underlying self-organisation processes helps to make the adaptive management of renewable resources both more effective and more efficient. In this chapter we present a simple model of self-organisation concerning the use of forest resources. It consists of two submodels: The submodel of the socioeconomic system comprises firms producing wood-based goods who buy and forestry companies who sell timber. The ecological system is represented by a forest succession model. After a brief description of both submodels, some preliminary results of simulating forest succession by using NetLogo are presented.

1 Introduction

The sustainable use of renewable resources is both a very important issue and a very difficult problem to tackle. It will remain a crucial task in the foreseeable future in part due to economic growth in rapidly developing countries like China which accounts for a large part of the worlds population. The sustainable use and management of natural resources is also a very challenging task

because it encompasses socioeconomic as well as ecosystem processes and it is characterized by complexity, a high degree of uncertainty, and information deficits and asymmetries [10].

Managing forest resources makes no exception. Foresters and forest owners have to meet multiple and often conflicting objectives such as maintaining a functioning ecosystem while simultaneously managing for high yields in order to be profitable. They also have to take into account current needs of the ecosystem, potential users of forest products and at the same time project into the future and attempt to predict the market situation decades later. Traditional short-rotation, even-aged, intensive management approaches often prove insufficient to sustain the ecosystem long-term. More sustainable forest management methods that build on natural regeneration and aim at maintaining biodiversity and minimizing impacts on the environment appear to hold more promise [6]. The major challenge for forest managers remains how to best handle an inherently complex, dynamic, non-linear and hard to predict forest ecosystem that is embedded in a socioeconomic framework characterized by frequent changes.

Adaptive management is a natural resource management approach specifically designed to deal with the issues described above by keeping the management process flexible, feedback-driven, and continually adaptive [8, 10]. Central elements of the approach include the use of experiments, comprehensive monitoring and involving all relevant stakeholders in the management process [21]. Even though there is often less attention paid to socioeconomic processes than to ecosystem functions, they too are viewed to be crucial elements of the adaptive management instrument [14, 10, 15, 9].

Self-organising processes in the biosphere and the anthroposphere are a major reason why sustainable management of renewable resources is a challenging problem. Self-organisation means that macro patterns emerge only or primarily by interactions between individual agents and not by external regulation. As a consequence, the outcome depends on initial conditions, development is path-dependent and irreversible, often characterized by bifurcations. Non-linearity can lead to rapid cumulative processes where tiny changes cause huge effects (the so-called “butterfly-effect”). If one or more of these features are present in a decision problem, as is usually the case in natural resources management, it is impossible to derive clear-cut optimal solutions. It is only possible to calculate several scenarios and to assess the range of potential outcomes. One way to deal with self-organisation is to build agent-based models and to run simulations. This is the method we have chosen because it captures the importance of feedback loops between decisions and actions of individual agents and helps us understand how self-organisation leads to emergent patterns. In our context such patterns might concern emerging stable or cyclical timber market situations, the establishment of cooperative structures comprising stakeholders with very different aims or the adaptation of trees to changing environmental conditions.

One of the weaknesses of the current practice of adaptive management is the lack of a thorough consideration and analysis of self-organisation processes, especially in socioeconomic systems. As a consequence, there is insufficient knowledge about the critical processes and the type of information that needs to be collected. In non-linear processes such as forest management, it is useful to follow a so-called “piecemeal engineering” model [16], which is a cautious small step-by-small step process. In order to do this successfully, regular monitoring of critical variables is crucial. We want to contribute to the adaptive management model by gaining an in-depth understanding of self-organisation processes in socioeconomic and ecological systems. For this purpose we have chosen agent-based modelling. This method shall provide, in addition to field surveys and expert knowledge, information for identifying key parameters and processes underlying the dynamic behaviour of the system to be managed - in our case the use of forest resources and forest management. Finding these key parameters and processes is very important, because explicitly monitoring them can help to recognize ongoing changes earlier and management can be adapted in time accordingly.

We present a preliminary, still simple version of a self-organisation model. The model consists of two separate, but interlinked submodels, one for socioeconomic, the other for ecological self-organisation. The modular approach has been chosen because of different modelling requirements: In the socioeconomic submodel there are only few, but complex agents whereas in the forest submodel there is a very large number of agents, but they are comparatively simple. Both submodels are described in detail in the next two Chaps. It is still work in progress, therefore the final Chap. will provide an outlook of how the model will be developed further.

Our main question is therefore: How do self-organisation processes on the timber market (determining demand for the forest resource ‘timber’) as well as in forest succession (determining the available stock of timber) mutually influence each other and which effects of certain adaptive management methods on the overall system’s behaviour can be expected?

Our paper is organized as follows: Chapter 2 and 3 present the design of the two submodels concerning the socioeconomic system and the forest submodel. In these Chaps. we describe the agents of the submodels and their interactions as well as the interface between the socioeconomic and the forest submodels. In Chap. 4 we discuss our first steps of programming which refer exclusively to the forest submodel. We conclude the paper with Chap. 5 giving a summary of the experience so far and an outlook of the next steps in our work.

2 Socioeconomic Subsystem

Agent-based modelling has gained importance in recent years in the field of economics and has been applied to understand a wide range of social and eco-

conomic problems or to forecast the effects of certain processes in socioeconomic systems. A famous early work concerned the emergence of racial segregation in cities [18]. Other issues were the emergence of cooperation [2] or the mutual influence of expectations in certain markets like the stock market [1]. It has even been argued that agent-based modelling could be used for simulating whole artificial economies [3]. These and many similar approaches have contributed to a common research agenda known as “agent-based computational economics” [20]. Unfortunately there are not many examples of agent-based modelling concerning the management of natural resources (e.g.[11]), where a complete agent-based model would have to comprise both social and natural systems and respective agents. Most often the focus is more on ecosystems or other biological systems neglecting socioeconomic processes. In our model we intend to avoid this one-sidedness.

2.1 Modelling the Use of Forest Resources

The aim of the socioeconomic submodel is to analyse how the use of forest resources is determined by the interrelations between specific forest management methods - some of them adaptive, some not - and the specific demand for timber of industries producing wood-based goods. The relations between forestry and industry form a system which is characterised by imperfect competition, imperfect information, strategic behaviour and learning. In order to analyse this self-organising system we have designed a multi-agent model focusing on a few key variables and relations only. Running simulations with an empirically calibrated model (using forestry data and interviews of experts) allows to test specific forest management routines under controlled conditions and restrictions. The “response” of the forest submodel is also modelled and will be described in Chap. 3.

In our model we focus on selling and buying of timber as the main kind of human influence on forests, at least in countries like Austria. We have designed a timber market with two types of agents which belong to the sectors ‘forestry’ offering timber and ‘industry’ producing wood-based goods. Other potentially important agents are either not included in this model (e.g. tourists, hunters) or considered as exogenous forces (e.g. state authorities, communities, demand for wood-based products, competing sources of timber supply).

In general, a socioeconomic agent is defined by a specific set of an objective and several routines and resources. Any organisational entity that is able to act independently (i.e. according to its own set of objectives and routines) is considered an agent. The agents in the socioeconomic submodel do not behave according to a constant set of stimulus-response rules but have the capability to learn. They continuously adapt to changes in their environment and they do this in a strategic way trying to achieve their objectives. Thus, agents evolve following a continuous cycle of formulating expectations of how to achieve their objectives, acting, evaluating the results of the actions (i.e. the response of the environment) and updating the objective or a routine in the

case of having missed the objective (see Fig. 1). Agents are interacting with their environment, including other socioeconomic agents, in three basic ways: acting, evaluating and searching. By acting the agent employs resources and directs them onto its environment (in order to achieve its objective). By evaluating it compares the results of actions or any impact from its environment with the objective. By searching it tries to find better routines for achieving its objective or a more suitable objective.

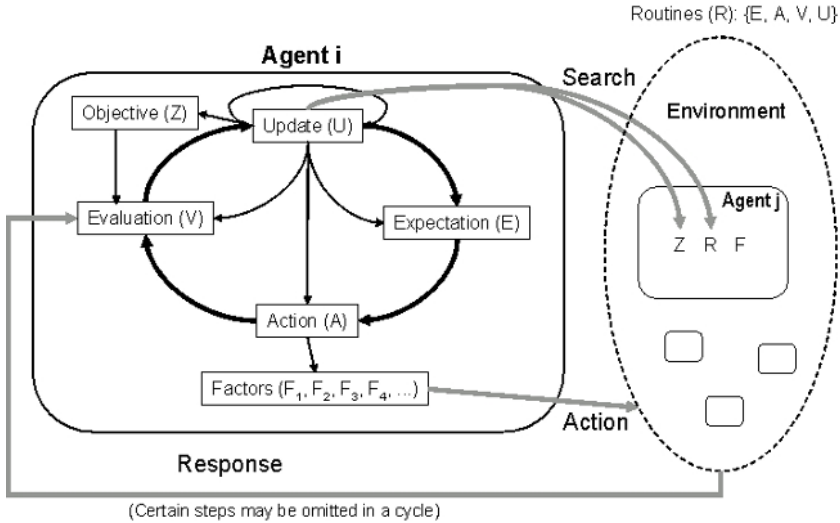


Fig. 1. Basic structure of a socioeconomic agent

Any agent is defined by a specific set encompassing objective, routines and resources (production factors) at a specific point of time t . At any time step this set can but need not be changed. An unchanged set of objectives and routines may be due either to perfect compliance with the objective or to a longer period during which one or more routines or the objective is decided to be fixed. The temporal perspective can be defined in the objective as well as in any of the routines. If the resources (production factors) remain the same, the agent has been inactive and has not been affected by its environment.

In the following, the elements of an agent are described in more detail:

- Objective (Z): Any agent has a single objective. While the objective functions of the agents are fixed (specific for each agent class), the target values can vary.
- Routines (R): This category comprises four basic types: expectation, action, evaluation and update.

- Expectation (E): An agent formulates at least one expectation about the effects of his actions on the environment, on reactions of the environment and on its development.
- Action (A): Agents dispose of at least one action routine which is required to achieve the objective.
- Evaluation (V): Any impact of the environment on the agent, be it the reaction on a specific action or any other impact, has to be evaluated by the agent whether the stated objective is still achieved or not. In the negative case, the update routine has to be invoked.
- Update (U): In the case of a missed objective, the agent searches for alternatives regarding the target value or one or several routines. Which elements of the (Z, R)-set are going to be adjusted has to be specified by the agent depending on the degree of failure.
- Resources or production factors (F): All resources an agent can dispose of in order to execute actions are called factors. In the model three basic types of factors are considered: fixed capital (production technology), variable capital (timber and wood-based goods) and money.

The fact that agents change their objectives and routines justifies to call them adaptive and the resulting changes in the population evolution. Innovation, however, is not considered in our model. There is a pool of routines and technologies available to the agents which is fixed throughout the simulation.

The temporal scale of the socioeconomic agents differs between industry and forestry. The former have a shorter time perspective. Their long-term planning covers typically 10 years. The latter consider longer periods of two or more generations (several decades). The model shall be able to simulate a small regional timber market comprising of a few dozen manufacturers of wood-based products as well as a similar or slightly smaller number of forest owners.

2.2 Timber Supply: Forestry Agents

Our model contains small forest owners (with forests of less than 200 ha) as well as large forestry companies. They differ with respect to their objectives, routines and available resources which are described in the following. Objectives and routines may vary between agents with regard to the functional form, more often, however, they will differ with regard to parameters only.

Objective ($Z_{forestry}$): Each forestry firm or owner defines a targeted profit rate for a certain period of time t . The profit rate equals revenue minus costs and is related to the capital employed (see equation 1). Price of timber (P_T) times volume sold (T_{sold}) gives the revenue. Costs comprise costs of capital, i.e. depreciation (C_F) of machinery (Tech), costs of extracting (C_T) the needed volume of timber ($T_{removed}$) and, if reforestation is done, costs of planting (C_R) new trees ($T_{reforested}$). Capital comprises machinery (Tech) and timber on stock (T_{stock}).

$$Z = \frac{P_T \times T_{sold} - (C_F \times Tech + C_T \times T_{removed} + C_R \times T_{reforested})}{Tech + P_T \times T_{stock}} \quad (1)$$

Exogenous cost parameters: C_F , C_T , C_R

Agents may have different target rates and time frames which represent different motives of the agents. If the agent is predominantly market oriented, the value of the profit rate will be high and the time frame short. If the agent also stresses nature conservation, the profit rate will be lower and the time frame longer.

Expectation ($E_{forestry}$): Forestry agents need expectations regarding the market demand for and price of timber as well as the future extractable volume of timber. As far as market expectation is concerned, a range of models may be applied. Some agents may use very simple forecasting models based on one or several previous periods while others may try to formulate more sophisticated models (e.g. looking for cyclical patterns). The second expectation concerns the growth of trees and the available stock of trees at any point in time. The expected volume of certain species and quality depends on the specific forestry method applied. The method also specifies if and how often the forecast will be adjusted based on monitoring of actual tree succession. Simplifying real forestry to two extreme approaches, two methods can be distinguished: The first method, we call “traditional” management, remains unchanged as long as there is always at least as much timber volume removable as expected. The second method, adaptive management, is regularly adjusted according to monitoring results of selected variables. Agents differ with regard to these expectations.

Action ($A_{forestry}$): Forestry agents execute two basic types of actions: managing forests and selling timber. Forest management consists of removal (harvesting) and reforestation. According to the expected market demand the respective volume of trees of a certain species and quality is removed. Quality is determined by diameter at breast height (DBH) and a factor derived by comparing actual growth with optimal growth (i.e. the growth under optimal environmental conditions). Areas which have been logged can be reforested with selected species, but they can also be left to regenerate naturally. There is a range of possible combinations of the forest management methods clear cutting, plantation and selective harvesting (“Plenterwirtschaft”). Each management method is characterized by specific fixed and variable costs.

Evaluation (V) concerns the comparison of the targeted and realised profit rates. The updating routine (U) consists of a coherent set of rules that specifies for any negative difference between stated goal and achieved result whether the target or a specific routine has to be modified and how they shall be adjusted. Agents’ update rules may differ with regard to thresholds and way of adjustment. For instance, a minor missing of the targeted profit ratio may lead to the adjustment of the market expectation only, while a bigger failure may require a new profit target. Furthermore, agents can apply varying search

spaces, ranging from global search (covering the whole pool of available routines) to local search (restricted to the routines used by agents with whom relations have been established).

Resources ($F_{forestry}$) of forestry firms comprise the production factors technology (machinery and other fixed capital necessary for timber harvesting) and money as well as the timber produced in period t and the timber that is on stock from former periods.

Forestry agents are interrelated with the forest succession model (see Chap. 3) via two activities: The first one is harvesting, the removal of trees which yield the volume and quality of timber expected to be sold. The second one, reforestation of certain tree species, is not necessary. Depending on the management method applied it is also possible to rely on natural rejuvenation. In our model forest management methods differ with respect to three routines:

- Expected removable volume of timber (expected growth of stock): Traditional management relies on total stock growth models whereas adaptive management tries to consider the forest succession process (see Chap. 3) in assessing the timber volume that will be removable in the future. There are also differences regarding the time horizon.
- Removal and reforestation: Adaptive management relies more on natural rejuvenation than traditional management where reforestation, planting trees according to estimated future demand, is more important.
- Adjustment of management: In traditional management trees are removed up to the expected demand (reduced by the timber on stock), restricted by the removable volume which is assessed on the basis of the stock growth model. The management is only changed if the available volume that can be removed is less than expected. In the case of adaptive management there is a monitoring programme watching actual growth of trees and key processes on a regular basis. The monitoring is used for the adjustment of management already long before a negative development becomes apparent.

Of course, real forest management cannot be reduced to two competing practices only. Our model allows the testing of a broader variety of management methods by combining a range of different routines.

2.3 Timber Demand: Wood-Based Industries

The present model comprises four industries, the three most important timber-consuming industries which are furniture, paper and pulp and construction, as well as bioenergy, a still small but growing sector, which is likely to increase in importance as a more sustainable form of energy production. It is not necessary to present the objectives, routines and resources of each industry separately, because all agents apply the same basic rules. What makes

industries distinct is that parameters vary more between firms in different industries than between firms within the same industry.

Objective ($Z_{industry}$): The objective has a similar form like in forestry. Each firm defines a targeted profit rate for a certain period of time t . Again, the price of the goods (P_W) times the volume sold (W_{sold}) gives the revenue. Except for the costs of capital, the costs of industry agents comprise slightly different categories: manufacturing costs (C_M) for the volume of produced goods (W_{prod}), costs of buying timber (P_T) and costs of transporting (C_T) between the seller and the buyer (D_{ij} being the distance between them) of the volume bought (T_{bought}). Capital comprises machinery (Tech) and goods on stock (W_{stock}).

$$Z = \frac{P_W \times W_{sold} - (C_F \times Tech + C_M \times W_{prod} + (P_T + C_T \times D_{ij}) \times T_{bought})}{Tech + P_W \times W_{stock}} \quad (2)$$

Exogenous cost parameters: C_F , C_T , C_M

Expectation ($E_{industry}$): Firms have to formulate expectations regarding demand for their wood-based products as well as their prices. As in forestry, several forecasting models can be applied (see Chap. 2.2.).

Action ($A_{industry}$): Firms perform three types of actions on a regular basis: they buy timber, produce wood-based goods and finally sell them. The production function depends on the technology used but will usually be characterised by a decreasing marginal product. The volume produced at any time t depends on the market expectation. The cost function consists of fixed and variable costs. The latter comprises costs of timber and transport from the saw mill to the firms plant.

Evaluation (V) is the same as in forestry. Update (U) rules, too, are based on the same logic but concern the specific routines of industry agents.

Resources of industry ($F_{industry}$) comprise machinery and other fixed capital necessary for manufacturing wood-based products, the stock of goods and money. It is assumed that all timber bought is processed within one period so there is no stock of timber at the end of a period.

The timber market is not a closed system. There are several exogenous factors and processes which are given as parameters and restrictions. This applies to the final demand of wood-based products, the competing supply of timber (from forests outside the model region), exports of timber (demand outside the model region), the pool of available technologies in forestry and industry and the cost of transportation between two sites. Furthermore there are important restrictions by public regulation. State authorities and communities are interested in maintaining or improving public functions of forests like protection against avalanches in alpine regions, preventing erosion, protecting the watershed or offering recreational areas. Instruments of regulation are detailed forest management standards (e.g. a ban on clear-cutting, restricting the maximum area to be clear-cut, rules concerning reforestation) and restrictions regarding land use (zoning).

2.4 Interactions between Agents: Timber Market

In our model there are only market relations. The individual decisions and actions of forestry firms determine supply of timber, the decisions and actions of firms producing wood-based goods determine demand for timber, and supply and demand are matched on the timber market.

- Supply side: Agents in forestry formulate expectations regarding market price of and demand for timber of a certain quality. They calculate the volume of timber considering existing stock, and they calculate the minimum price of timber acceptable to achieve the profit target. Then they produce the volume of timber of this quality and offer the total volume (stock and production) on the market at the expected price.
- Demand side: Firms of the wood-based industries formulate expectations concerning market price of and demand for their wood-based products. They calculate the quantity of goods considering goods on stock and the maximum price of timber acceptable to achieve the profit target. Then they look for timber of the quality required for production at the expected price. The quality is at a minimum determined by tree species and size, or the diameter at breast height (DBH) in technical terms.
- Timber trading and market clearing: The timber market is primarily determined by selling and buying between the agents in the model (forestry and manufacturing firms). But there are also exogenous sources of demand and supply, i.e. exports and imports of timber to and from outside the model region. At first, all individual supply volumes and all individual demand volumes of the agents in the model are summed up, differentiated by quality. Then the imported volume of timber is added to and the exported volume is deducted from the total supply. Individual selling and buying offers where both prices are equal are immediately matched. Finally, those offers of timber which cannot be matched are removed from the market and added to the forestry firms stock of timber.
- Wood-based products trading and market clearing: It is assumed that all firms face one market demand curve for each industry which gives the maximum volume that can be sold at the chosen price. The shape of the demand function for wood-based products is given exogenously. Except for the exogenous nature of demand, the process of determining market price, matching supply and demand offers and adding unacceptable offers to the manufacturers stock is the same as in the case of the timber market.

3 Forest Subsystem – Forest Succession Model

Natural ecosystems and forests are constantly undergoing changes and responding to changes; they are open, in flux and are affected by a series of often stochastic factors. Managing non-linear, highly variable and complex

systems is challenging and requires a better understanding of ecosystems. Therefore scientific analysis and monitoring of key parameters are necessary, in order to ensure a sustainable management process [17]. An agent-based model can help analyse self-organisation processes during forest ecosystem succession which may help explain certain ecosystem processes and interactions. This helps to monitor the key parameters within a forest ecosystem, which, as a consequence, contributes to its adaptive management.

3.1 General Characteristics of the Agents in the Succession Model

The agents (individuals) of the forest model are different trees with different main characteristics. These agents are spatially immobile, but their influence due to seed dispersal and competition with other trees goes beyond their local position. During forest succession the competitive advantage changes between species. Many pioneer species which have advantages at the beginning of settlement or resettlement of an area (e.g. a gap within a forest ecosystem after a windthrow or a clear cut) cannot compete with other species in later succession stages. Therefore, the tree stock at a certain point of time often does not represent the biodiversity present during former stages of the succession process. For example, if we look at Austria's forest inventory, many species do not seem necessary to be modelled because of their marginal occurrences. However, if we want to simulate self-organisation processes, they are also important. In the first model the tree species birch, beech, fir, spruce and pine are included, but the number of species can be increased in future versions to cover the diversity in the succession process in different regions.

3.2 Self-Organisation in a Forest Ecosystem

A definition for self-organisation in biological systems is given in Camazine, Deneubourg et al. [5]: "Self-organisation is a process in which a pattern at the global level of a system emerges solely from numerous interactions among the lower-level components of the system." Within a forest ecosystem, trees can be seen as lower-level components which act mainly local while their interaction shapes global patterns. One main problem of analysing changes in forest ecosystems is that visible changes of patterns often take a very long time. A computer-based simulation model therefore can provide a tool to observe these trends earlier and help to find more sustainable management decisions. Hence, one major goal is to capture the main events and mechanisms that determine the temporal and spatial dynamics within a forest succession under given environmental circumstances. As we focus on the self-organisation processes, i.e. the spatial and temporal interaction and competition between trees, we found it appropriate to use a so-called forest gap model instead of the growth-yield models, more often used in forest management [13]. Our approach is based on the forest model JABOWA III from Botkin et al. [4], as several studies have shown [19, 7], is an efficient and widely used model.

In the first version of the model we focus only on those features of JABOWA III which are important in our modelling context. In the model, the trees are characterised by diameter (diameter at breast height, DBH), height and species. Every species is defined by a set of parameters and the individuals compete against each other to get more light, water and nitrogen. For each agent or tree a general growth equation is calculated (equation 3).

$$\Delta GGF = f(Dmax, Hmax, b_2, b_3, D, G) \times f(environment) \quad (3)$$

b_2, b_3 ...Parameters defining tree form which can be estimated with Dmax and Hmax [4]

D...Diameter at breast height (DBH)

G...Growth parameter which can be estimated according to maximum age, Dmax, Hmax

With this equation the maximum growth potential (the maximum changes per year, ΔGGF , if $f(environment)=1$) for each tree species is calculated as a function of the species-specific parameters maximum diameter at breast height (Dmax), maximum height (Hmax) and maximum age, the actual diameter of the tree, reduced by local environmental responses ($f(environment)$, a factor between 0 and 1). Environmental responses are correlated with available light and site conditions (equation 4).

$$f(environment) = f(light) \times Qi \quad (4)$$

Site conditions included in the present model are the “wilt” factor WiFi, an index of drought conditions a tree can withstand, and the index of tree response to nitrogen content of the soil (NFi). For the next version it is also planned to include the general temperature response function TFi and the “soil wetness” factor WeFi, an index of the amount of water saturation of the soil a tree can withstand (equation 5).

$$Qi = WiFi \times NFi \times TFi \times WeFi \quad (5)$$

In a natural forest ecosystem dead trees are common and an important factor, whereas they are often missing in a managed forest. This is important because some seeds chiefly germinate on dead trees. Within the model, mortality is simulated according to JABOWA III in two different ways. First, there is an inherent risk of death for any tree independent of the competition with other trees, e.g. caused by windthrow or, very important as it relates to the socioeconomic subsystem, harvesting. Second, there is competition-induced death. Trees that grow poorly over a certain period of time (e.g. ten years) have a higher probability to die than well growing trees. These are the responses of an existing tree to the environment. The natural reproduction of trees, i.e. regeneration, is also very important. A tree has to reach a species-specific age for seed production and produces characteristic seeds. To simplify

the model, we distinguish only between tolerant, intermediate and intolerant seeds for the tree parameters shade, nitrogen- and water-scarcity. Contrary to the forest model developed by Botkin et al. [4] (JABOWA III), we do not assume that there are always enough trees to produce seeds independent of the management practice and how many old trees in the area are able to produce seeds. We think that is closer to reality and as other investigations have shown, simultaneous events like increased seed production and increased open areas can strongly influence the dynamics [22] of the system. In the first version of the model we assume the concentric dispersal of seeds from the producing tree, but the model can be improved by correlating seed dispersion with the main wind direction and dynamic. Seeds of shade-tolerant species are able to germinate under the producing tree as opposed to shade-intolerant seeds which need an open area to germinate and grow. Self-organisation within the forest succession model emerges due to direct local competition of the trees for light and space (seed dispersal). A taller tree reduces available light for the smaller trees in its neighbourhood. Neighbouring trees with almost the same height do not reduce the amount of available light. Direct competition for nutrients and water is not implemented in the present version of the model as one tree does not influence the general availability of nitrogen and water of a neighbouring tree. Environmental influences (rainfall, snow, calamities etc.) can be summarised as external natural influences which are often difficult or even impossible to predict. They can have dramatic influences on the system depending on its current state.

The self-organisation mechanism within the forest succession model is visualised in Fig. 2. The positive feedback between “Ability to compete with other trees”, “Resources” and “General growth function” characterises the self-organisation process (Fig. 2). A tree with a higher ability to compete can absorb more resources (light, nitrogen, water) and hence grows faster which, in turn, increases its absorption capacity. To avoid an overshooting reaction, negative feedbacks, such as the inability of seeds of a tree which grows faster under full light conditions to grow under the producing tree itself, balance the system over time. Changes at the macro level (forest succession) are driven by the positive feedback mechanism at the micro level (induced by the interaction of individual trees).

3.3 Interface with the Socioeconomic Submodel

The interrelation with the socioeconomic subsystem simulates a specific human influence on the forest succession process, the extraction of timber by forestry. As shown in Fig. 2, the interface to the socioeconomic model encompasses two types of actions - *removal (harvesting)* and *reforestation*.

Removal/harvesting: Timber for different purposes like furniture, construction and paper production requires different tree species with different qualities (age, shape or DBH). Demand for different species and qualities of

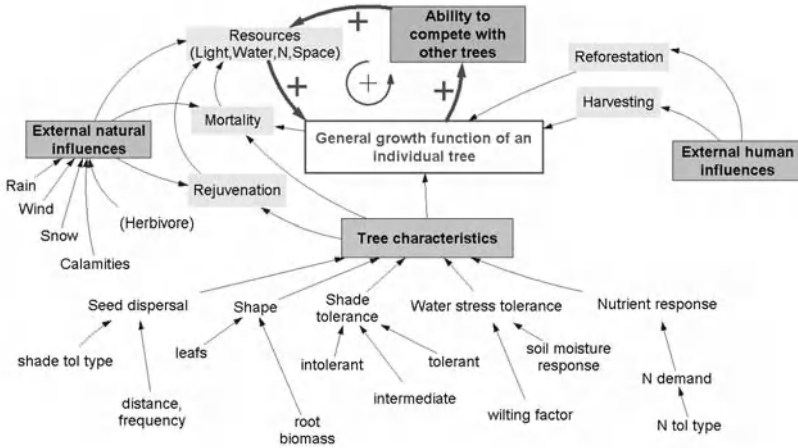


Fig. 2. General forest submodel scheme

timber is determined in the socioeconomic subsystem which, in turn, determines which trees at what location and by which harvesting method will be removed. The currently implemented harvesting methods are selective harvesting, clear-cutting, thinning and cleaning. Selective harvesting means that individual trees are removed depending on their DBH. They can be selected randomly within the whole forest area until the wood demand is fulfilled or they can be taken from a restricted site. Clear-cutting means that all trees within a certain area independent of their age, DBH or growing behaviour are removed. The newly opened site (gap) may be reforested with a chosen tree species or natural regeneration may be allowed to take place. Thinning means that all individual trees less than a certain DBH are removed. Thinning can be restricted to a small patch or encompass the entire forest. This harvesting method is mainly used to improve growing conditions by increasing available light. Cleaning means that all trees or a certain amount of trees are removed which exhibit poor growth over a period of time. These trees have poor wood quality and a higher possibility for natural mortality compared to trees that grow well. Removing them enhances growing conditions for all trees and normally improves germination of seedlings.

Reforestation: The socioeconomic model determines whether and which tree species should be planted. Reforestation has no direct influence on the general growth function (equation 3) of existing trees, but affects them indirectly, because competition for resources by newly planted trees affects existing trees.

4 Implementation of a First Version of the Forest Succession Model

We have used NetLogo to implement our first model-version and to test our approach. NetLogo is a freely available, multi-agent, programmable modelling environment from the Center for Connected Learning and Computer-Based Modelling of the Northwestern University (<http://ccl.northwestern.edu/>). One of the main advantages of NetLogo is that it is easy to use and therefore very valuable to test new model approaches. A disadvantage is the poor simulation speed compared to other agent based simulation platforms like RePast (a JAVA-based agent simulation toolkit). Since we need a large amount of agents (trees), this is a major argument for developing our next version with RePast. Figure 3 shows the graphical user interface developed with NetLogo. Number of trees, harvesting method, simulation duration, nitrogen and water availability, harvesting interval and reforestation behaviour are the main input parameters of the model which can be chosen by the user. The simulated area is 36 ha in size with an individual patch size of 5x5 meter (25 m²) which equals about 14,600 patches in total. Each patch has a certain amount of nitrogen and water available. This availability is initialised within the setup process and can range from 'no' to 'medium' and 'high' to reflect different site conditions. The available light on the ground of each patch is derived from the trees growing on the patch and in the neighbourhood. This amount of light is very important for germination and the growing potential of different seedlings.

Outputs include the visualisation of tree species and size within the simulation area, diagrams of the number of trees of the selected species, percentages of trees with high quality (the quality is calculated compared to the growing behaviour under optimal environmental conditions), the harvesting potential of trees with a DBH greater than 60 cm and the removed (harvested) amount of trees.

4.1 Calibration, Verification and Validation

Calibration (including parameterisation) of a model is a crucial step in the development of a model. One reason to use a relatively simple model approach is to find valuable data for calibration. Main inputs for the forest self-organisation model are tree characteristics, site conditions and information about stochastic events such as windthrow, seed dispersal or tree mortality. Within this version of the model, there is no feedback loop between the state of the forest and the occurrence of calamities (e.g. bark-beetle). The quality and validity of models has to pass several checks [12]. Verification means testing the model with regard to the underlying mathematical and computational components. Sensitivity tests play an important role in this context.

Validation comprises structural and outcome validation. In a model which is focused more towards shedding light on theoretical questions – which applies

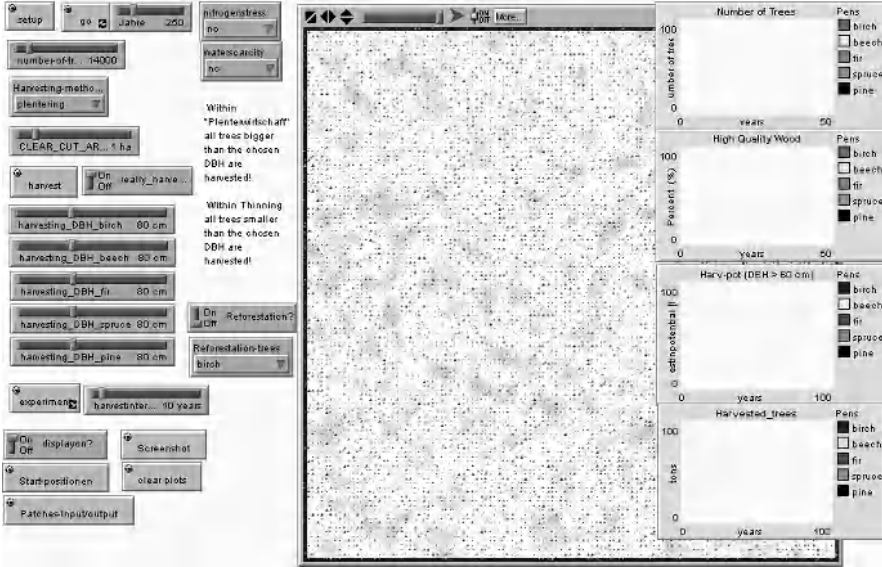


Fig. 3. General overview of the first model-version developed with NetLogo

to our model – verification and structural validation is more important [12]. The validation of the dynamics of the forest subsystem to prove the realistic behaviour of our model is a three step procedure:

- First, the potential growth curve of each tree species is simulated independently from competition with other trees and validated with experimental data.
- Second, the different environmental response functions of each tree species is checked independently and also validated with experimental data.
- Third, the competition between different tree species is analysed and validated with experimental data and expert knowledge.

The validation of the model developed with Netlogo has shown reliable behaviour as expected by experts. A more intensive validation, however, will be necessary for the RePast model, including more site-specific data from Austrian forests. The first simulation results have shown that some parameter settings are very sensitive which makes it necessary to calibrate and verify them very carefully.

4.2 Analysis of the Forest Succession Model

Agent based models such as the one introduced have extensive lists of parameters which can be combined to an almost infinite number of different parameter settings and simulation runs. We made more than 1,000 runs with

different parameter settings to test and analyse the behaviour. In the following we discuss some interesting outcomes of this analysis.

As is often the case in complex systems there are sensitive parameter ranges which makes the model verification more difficult [12]. Within one experiment all trees with a DBH of less than 15 cm were harvested at a certain interval. Figure 4 shows two different simulations with only a slight change in parameter settings namely the harvesting interval was increased from 15 (left) to 20 years (right).

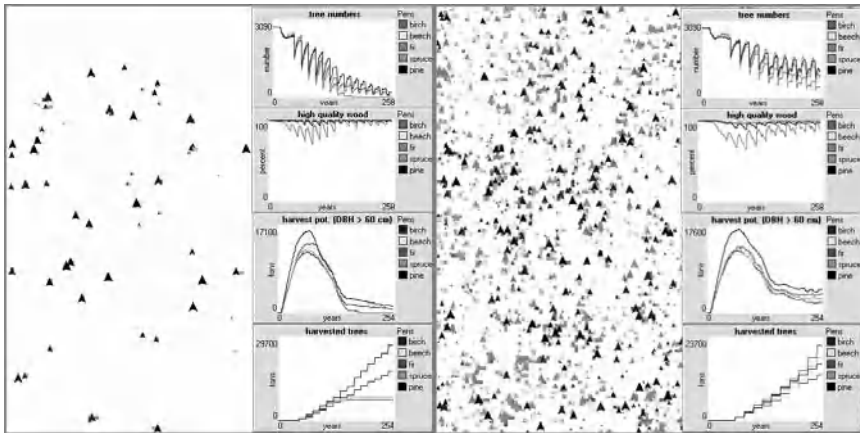


Fig. 4. Experiment: Thinning trees less than 15 cm DBH with an interval of 15 years left, 20 years right

Figure 4 (left) shows the development over the simulated 250 years with the result that almost all trees are extinct. In Fig. 4 (right) this is changed dramatically; there, a more or less sustainable behaviour can be observed. Main causes for this behaviour are individual tree characteristics, age at first seed production and number of seeds and growing performance. If trees grow slowly within 15 years they are not able to reach more than a DBH of 15 cm and are harvested. It is important to note that this result is independent of the number of trees present at the beginning. If the interval is changed from 20-25 years, we do not observe the same behaviour. The next Fig. shows an experiment where we have used the selective harvesting method for all trees larger than a certain DBH. The change we made between the runs was to increase the DBH from 70 (left) to 80 cm (right).

Changing the DBH from 70 cm to 80 cm changes growing conditions in particular for birch and pine which have similar assumed growing and reproduction behaviour, different from the other tree species. It further shows, that although more trees have been harvested in the first experiment (DBH 70 cm) at the beginning, at the end much more trees could be harvested with a DBH of more than 80 cm. The simple experiments displayed in figures 4 and 5 pro-




Fig. 5. Experiment: Selective harvesting of trees with more than 70 cm DHB left and 80 cm DHB right with an interval of 10 years.

vide an impression as to why simulation is valuable to learn more about system behaviour. A main element of adaptive management is long-term monitoring. In the experiments discussed above we found that the harvesting interval is one of the key parameters. The influence of this parameter on the number of trees and harvesting potential (trees > 60 cm DBH) indicates the long-term behaviour of the system. Without adaptive management, the trees would be extinct as in Figs. 4 and 5 (left).

5 Conclusions and Outlook

The succession model, presented in this paper, is able to show basic behaviour of self-organisation, e.g. competition of trees for light and space. The currently used agents within the forest model vary only with regard to few parameters, nevertheless we can observe interesting dynamics. For example, we did not include realistic seed dispersal rules, and although it can be assumed that these are critical parameters for determining succession, we have found characteristic succession patterns like early dominance of pioneer species followed by characteristic species of later stages. Self-organisation is obviously driven even by small changes due to the positive feedback between growth function, competitiveness and resources.

At present, the other, the socioeconomic submodel is being programmed. It will be the basis for simulation and analysis of the self-organisation processes on the timber market and the resulting use of forest resources. We already have started to program both submodels in RePast in order to ensure their smooth connection.

Before that, however, both submodels are being verified independently, because the technical requirements differ significantly between them. Espe-

cially, the large amount of agents in the forest model (individual trees) which is necessary when modelling a small regions timber market might require distributed simulation techniques.

Finally, the overall model, connecting the socioeconomic and ecological submodels, will be validated primarily by discussing them with experts in forest management and, to a lesser extent, by empirical data.

After having validated the model it will be possible to analyse simulation results regarding the most critical self-organisation processes. Different forest management practices will be tested on their effects on the socioeconomic and ecological subsystems. It will be particularly interesting to compare adaptive management with “traditional” management practices.

References

1. Arthur WB, Holland JH, LeBaron B, Palmer R, Tayler P (1997) Asset pricing under endogenous expectations in an artificial stock market. In: Arthur W.B., Durlauf S.N., Lane D.A. (eds) *The economy as an evolving complex system II*. Perseus Books.
2. Axelrod R (1984) *The evolution of cooperation*. Basic Books.
3. Batten DF (2000) *Discovering artificial economics. How agents learn and economies evolve*. Westview Press.
4. Botkin DB (1993) *Forest Dynamics, An Ecological Model*, Oxford University Press Inc, ISBN 0-19-506555-7:309, New York.
5. Camazine S, Deneubourg JL, Franks NR, Sneyd J, Theraulaz G, Bonabeau E (2001) *Self-organisation in Biological Systems, Princeton Studies in Complexity*, Princeton University Press, ISBN 0-691-01211-3:538, Princeton.
6. Cote MA, Kneeshaw D, Bouthillier L, Messier C (2001) Increasing partnerships between scientists and forest managers: lessons from an on-going interdisciplinary project in Quebec, *The Forestry Chronicle*, v.77, 1:85–89.
7. Green DG (1989) Simulated effects of fire, dispersal and spatial pattern on competition within forest mosaics. *Vegetatio* 82(2):139–153.
8. Grumbine RE (1994) What is Ecosystem Management?, *Conservation Biology*, 8(1):27–38.
9. Haney A, Power RL (1996) Adaptive Management for Sound Ecosystem Management, *Environmental Management*, 20(6):879–886.
10. Holling CS (eds) (1978) *Adaptive Environmental Assessment and Management*, John Wiley and Sons, London.
11. Janssen MA, Walker BH, Langridge J, Abel N (2000) An adaptive agent model for analysing co-evolution of management and policies in a complex rangeland system. *Ecological Modelling* 131:249–268.
12. Janssen MA (eds), Manson ST, M (2002) *Complexity and ecosystem management*. Edward Elgar Publishing, ISBN 1-84376-061-4, Cheltenham:63–74.
13. Liu J., Ashton, P, S (1995) Individual-based simulation models for forest succession and management. *Forest Ecology and Management*, 73:157–175.
14. McLain RJ, Lee RG (1996) Adaptive Management, Promises and Pitfalls, *Environmental Management*, 20(4):437–448.

15. Peterson G (1997) Uncertainty, Climate Change, and Adaptive Management, *Conservation Ecology*,1(2).
16. Popper K (1966) (5th edition): *The open society and its enemies*. First Princeton Paperback Printing.
17. Schaffer LL (1997) Lectureship in Forest Science, November 3. <http://www.forestry.ubc.ca/schaffer/botkin.html> (16.11.2004).
18. Schelling TS (1969) Models of segregation. *American Economic Review* 59/2, 488-493.
19. Shugart HH, West DC (1981) Long-term dynamics of forest ecosystems. *American Scientist* 69(6):647–652.
20. Tesfatsion L (2001) Introduction to the special issue on agent-based computational economics. *Journal of Economic Dynamics and Control* 25: 281–293.
21. Walters CJ, Holling CS (1990) Large-scale management experiments and learning by doing, *Ecology*, v. 71, 6:2060–2068.
22. Wiegand T, Milton SJ, Wissel C (1995) A Simulation Model for Shrub Ecosystem in Semiarid Karoo, South Africa, *Ecology*, Volume 76, Issue 7:2205–2221.

Vampire Bats & The Micro-Macro Link

Rosaria Conte, Mario Paolucci, and Gennaro Di Tosto

Institute of Cognitive Science and Technology, CNR, Rome, Italy

`rosaria.conte@istc.cnr.it`,

`mario.paolucci@istc.cnr.it`,

`gennaro.ditosto@istc.cnr.it`

Summary. An evolutionary variant of the Micro-Macro Link (MML) theory is proposed. According to the MML theory, behaviour at the individual level generates higher level structures (bottom-up process), which feed back to the lower level (top-down), reinforcing the producing behaviour either directly or indirectly, i.e. acting on the mind of the agent. Focusing on the evolution of pro-social behaviour, we explore these two types of link by means of simulation experiments. The experiments were aimed at studying the role of groups in the performances of populations of agents exchanging helping behaviour (simple loop) and the role of cognitive mediators (i.e. goals) and social norms in the spreading of altruism, described as a purposeful action (complex loop).

1 Introduction

As shown by a huge amount of experimental literature (see [1], for a review), humans are more likely to cooperate than expected by rationality theory. Why?

Answers to this question are essentially found in sociobiology, under the umbrella of either kin selection or reciprocal altruism.

Kin selection theorists expect altruism to be found among relatives, in a measure that is a direct function of the degree of relatedness between donors and recipients. On the other hand, reciprocal altruism accounts for the evolutionary advantage of altruism in terms of inclusive fitness.

However, a strong objection to this fascinating biological theory — which became very popular in the seventies and still is dominant among evolutionary scientists — comes from analytical models of social dynamics: according to the so called haystack models [2], altruists go extinct when randomly matched with non-altruists. Otherwise stated, unless altruists and non-altruists are assorted in such a way that the former are more likely to meet with one another than they are to meet with non-altruists, altruists are invaded by non-altruists, and the whole group is bound to die out soon.

On one hand, then, the survival and reproduction of groups seems largely to depend on the survival of altruists: groups of cooperators have a much longer life than groups of non-cooperators.

On the other hand, how did altruistic individuals and groups evolve at all, if they are so vulnerable to non-altruistic ones? In this paper, we propose an answer to this question in terms of an evolutionary variant of the Micro-Macro Link (MML) theory. The MML occurs when behaviour at the individual level generates higher level structures or groups (bottom-up process), which feedback to the lower level (top-down), either:

- *directly*, reinforcing the altruistic behaviour (**simple loop**);
- *indirectly*, acting on the minds of altruists (**complex loop**).

In our view, both loops are involved in different stages of the evolution of altruism. Inclusive fitness is not sufficient to explain the evolution of altruism, since group selection is also needed. As we will endeavour to demonstrate by means of artificial data, drawn from simulations of a classic ethological example of altruism, food-sharing among vampire bats, a high number of small groups (roosts) allow for a strong variance in the assortment of altruists and non-altruists over the whole population. Hence, roosts *de facto* provide social barriers preventing altruists from being invaded by non-altruists (simple loop).

However, this is not a robust effect, since as soon as the ecological conditions vary (e.g., the number of individuals *per* group increases), altruism is at risk. Other properties at the individual level must evolve in order to keep non-altruists from dominating, and to protect the whole group (complex loop).

We illustrate the two loops by means of simulation. In the first part of the paper, we will describe a simple loop: our simulation model of vampires' food-sharing will be applied to investigate the effects of roosts on the evolution of this behaviour. In the second part of the paper, we will investigate a complex loop, by experimenting on the individual properties allowing altruists to survive and neutralise non-altruists even under unfavourable demographic conditions, i.e. when the probability of meeting a non-altruist is higher than the other way around.

2 The Model

In nature, examples of altruism abound [3]. Among mammals, the most famous example of pro-social behaviour is *blood-sharing* in vampire bats [4]. The species studied by Wilkinsons lives in Central America, in small groups (a few dozen individuals) inhabiting the cavity of trees. We will call this basic unit group a roost¹. Their diet consists of ingesting each day an amount of fresh

¹ To be more precise, real vampire bats move in subgroups around several cavities, creating a fluid and territorial group system. Roosts contain usually only one

blood, which they suck from herbivores. However, each night about 7% of the adults find no prey to parasitize. In these occasions, they can survive thanks to luckier fellows regurgitating for them a portion of the food ingested. Wilkinson, who have studied this species in its natural settings, actually stated that such behaviour “depends equally and independently on degree of relatedness and an index of opportunity for reciprocation” [4].

Starting from these ethological data we have drawn the model for a multi-agent system to analyse the key features of altruistic behaviour and to explore the micro-macro link. We are not interested in the ethological debate around the life of vampire bats *per se*. Our goal is to model behaviour at an abstract level, taking inspiration from the case of vampire bats, which provides good evidence for building the simulation model in a non-arbitrary way.

Every agent in the simulation is then designed to reproduce the hunting and social activity of the common vampire bats. Each simulation time step includes one daily and one nightly stage. During the daily stage, the simulated animals perform social activities (*grooming* and *food-sharing*). In the night, they hunt. In our model, hunt is defined as an ecological parameter; in accordance with real-world data [5], its default value is set to 93%. In substance, each night 93% of the population will find food to survive until the next hunt. The remaining 7% will begin (or continue) to starve, unless they receive help from some fellow (under the form of regurgitation). Vampire bats do not accumulate resources: hunt is performed only for short-term food consumption. In addition, although the average lifetime of these animals lasts around 14 years, starvation and death are a constant menace to them, since each good hunt gives them no more than 60 hours autonomy. As a consequence, for a bat in isolation, two failures in a row are fatal. These are the harsh conditions characterizing the life of vampires, which face infrequent — in the simulation, about 1.65 episodes of double unsuccessful hunt per animal per year — but lethal food scarcity. The only way to prevent starvation and death is receiving help from fellows, which is what these animals appear to do in nature.

In the daily stage, the following actions can be performed:

Groom The condition for this action is that two agents are sorted out from the same roost. Grooming allows for help requests.

Ask for help The condition for the application of this action is that the requestor be starving. The request will be addressed to one agent in the same grooming network, if any. Otherwise, other in-roosts may be addressed as well. The effect of a request will be either donation or denial. In the first case, the requestor will ingest some blood and gain some hours of autonomy. In the second, it is bound to die.

Donate The condition for applying this action is that recipient is starving. The effect is that donor’s lifetime is reduced and the recipient’s is in-

alpha male, plus several other males and females in a rigid hierarchy, but we will not model this level of detail in our simulation.

creased. Donating, in accordance with physiological data, is a non-zero sum interaction: the receiver gains more than the donor loses.

Deny help The condition is that agent received a request for help. The effect is the requestor's death.

As to daily activities, in nature the rationale of grooming is at least twofold: animals familiarize thanks to and during it, and check their respective physical shape. Since satiation causes body volume inflation, a lucky hunter may grow to even twice as much as its normal size, as can be easily detected by its grooming partners. Likewise, a starving bat is also likely to be recognized. Bluff would immediately be found out in such extreme life conditions.

Each day, pairs are formed by each animal choosing one partner from the roost population. As in the real world, in our model grooming has the effect of increasing the probability of food-sharing among in-roosts: a starving bat will turn to grooming partners for help, and will avoid death if any of them is found to be full (having had a good hunt). Because of the bat's metabolism, the donor will lose much less than is gained by the recipient. In the simulation, the agent loses an amount of energy allowing them to survive for six hours; this amounts to losing a chance to ask for help during the last day, given two failures in a row following donation. In this set of experiments, we limited the number of partners per day to one.

The numbers obtained merge with what is known by ethological observations in presence of mutual help (the yearly rate of death for adults is about 24%) and with the results of a simulation carried on by Wilkinson [5] in absence of help. As said above, help is rare but critical: roosts in which all individuals deny help reduce their population by 82% in a year.

The case under study presents two peculiarities. Any kind of wealth accumulation is exceedingly difficult. Energy coming from a meal is dissipated after two nights, so that there can be no such thing as a wealthy individual. The lucky hunter of today has the same chances as everybody else to starve tomorrow. Moreover, direct retaliation is simply impossible in the present setting. The victim of cheating² dies on the spot; asking for help is the last resort, and given our restriction of one helping partner per night, a cheater is a dangerous killer that is really difficult to find out.

In the simulations, only starving animals are allowed to ask for help, and will be helped by their addressees if these are both altruists and satiated. No bluff is allowed — if an agent is not starving it will not ask for help. Agents have no memory of past interaction and cannot calculate the probability of reciprocation. No explicit mechanism for punishment of cheaters is implemented. In such conditions, how can reciprocity emerge as a mere “objective” effect, implying neither computation nor deliberation on the side of altruists?

² Following the convention used in the literature on cooperation a cheater is a defector, i.e. anybody who does not cooperate or, more specifically, who denies help.

In order to answer this question and to explore the effect of groups on the evolution of altruism, we run simulations with mixed populations (altruists and cheaters in variable combination) initially distributed over a given number of roosts. During the simulation, roosts can grow or collapse, depending upon the survival and reproduction rates of their members, which in turn exclusively depends on social attitudes (whether altruistic or not). Ecological conditions are kept equal for all roosts. At given times, roosts give rise to new roosts if the number of young individuals reaches a given threshold. This was meant as an operational simplification of the notion of group selection and reproduction.

3 Simple Loop: Groups and the Evolution of Altruism

This study was aimed to test different interpretations of the evolution of altruism by means of simulation. The reference example in the real world is the vampire bats' food-sharing. As previously recalled, this species offers clear evidence of the advantages of altruism on life expectancies. Wilkinson's simulation findings, also reproduced by our simulation, show that the probability of survival per year is around 76% of the population when food-sharing is activated, as opposed to a bare 18% in the opposite condition (no food-sharing). How to interpret these findings?

In line with the sociobiological theory of reciprocal altruism [6], one could say that vampires survive to such a greater extent when sharing food because of reciprocity [7], which adds to the individual fitness of donors much as altruism adds to the fitness of recipients. In other words, help-giving acts as a sort of investment, although non-deliberate nor acknowledged, on the part of the altruist, which accumulates credits to be refunded by means of reciprocity. Since vampires do not accumulate food, donors that are reciprocated later on in their lives will survive longer than if they had performed no altruistic act. Whereas the initial donation caused a mere reduction of the time interval before starvation, the following reciprocation prevents immediate death!

However, it is unclear whether and to what extent vampires take measures against cheaters. Wilkinson's findings refer to the comparison between an all-cooperators condition *vs* an all-defectors condition. What happens in intermediate conditions? Which is the minimal share of altruists for obtaining an increase of the survival rate with regard to the all-defectors condition? Moreover, does the increasing of survival rate effectively correspond to an increase of donors' fitness, or is it redistributed over the entire population? And if so, are individual donors always refunded or do they sustain a share of the costs of redistribution?

The latter question is crucial since if donors are not always reciprocated in person or along their future generations, there is reason to question the reciprocal altruism interpretation, and to look for another concurrent explanation. One good candidate is the group-selection theory, i.e. the theory that considers aggregates of non-kin individuals as units of biological selection and




Fig. 1. Evolutionary advantages of food sharing among common vampire bats. When food sharing is not allowed the annual rate of mortality increase from 24% to 82%.

evolution. In such a perspective, a given trait, like altruism is accounted for in terms of its contribution to the fitness of the group rather than to the individual (genetic or reproductive) fitness of its members. Consequently, vampires' food-sharing may be seen as a habit that evolved thanks to its positive effects on the fitness of roosts taken as wholes, rather than on the individual fitness of donors. In short, the reciprocal altruism theory proves adequate if donors are almost always reciprocated in their lives or in the lives of their offspring. In such a case, the altruistic gene spreads because the genes of donors survive and replicate through generations.

Instead, the concurrent group-selection theory proves right if (i) survival rate increases in altruistic roosts although donors are poorly reciprocated both in their lives and in their offspring's, provided (ii) the altruistic roosts

as wholes are fitter than their non-altruistic competitors. How to measure a roost's fitness? We turned to roost formation as one possible solution. In this sense, the higher the number of roosts that get formed from an original one, the fitter the latter. This is so, provided the rationale for roost formation is reproduction: the higher the number of reproduction of a parent roost, the higher the number of child roosts it will give rise to.

3.1 Simulating the Effects of Groups

By means of RePast simulation platform³, we constructed a pilot experiment that mimics the vampires' behaviours in what we perceive as its essential traits, introducing two different algorithms, *altruistic* (food-sharing) and *cheating* (no food-sharing). The former mimics the behaviour of lucky hunters that give away an extra amount of the blood ingested, if any, to the benefit of starving fellows asking for help. The latter reproduces the behaviour of selfish animals, denying help to their unlucky fellows, which starve to death.

Cheaters never give help when asked, even if they are full; unlike altruists, they sustain no costs. Since no retaliation mechanism is modeled, a first expectation could be that cheaters prosper, thereby reducing the efficiency of the system as a whole. This is indeed what happens in the short run; but when we start to consider longer simulations, the scenario changes dramatically.

Moreover, apart from the effects of cheating, we were also interested in a measure allowing to discriminate between group selection and inclusive fitness. To this purpose, we keep track of the lineage of the agents from the beginning of the simulation. Reproduction by cloning allows for clear tracks; if the mortality rate of one lineage is equal-lower than the average in the same roost but this produced a significantly higher number of children roosts than happens in the control condition, then group-selection seems an adequate interpretation of vampires' altruism. If instead donors' lineages show a mortality rate significantly lower than the average in the same roost, and the number of children roosts is not significantly higher than in the control condition, reciprocal altruism provides a more adequate interpretation.

Simulations have been run for a number of cycles corresponding to forty years, which includes about four generations of vampires. In Fig. 2 and 3, we present typical examples of what happens during the run for different initial shares of cheaters.

Results clearly show the reproductive advantage of the food-sharing condition. Selfish vampires are bound to go extinct in a few generations, leading also their roosts to collapse. They play a destructive role, by gradually reducing the reproductive capacity of their roosts until global extinction. However, when altruists far exceed cheaters, or some demographic catastrophe (triggered by cheaters themselves) leads to earlier extinction of cheaters, the reproduction of altruists takes off again and the number of roosts grows in proportion.

³ <http://repast.sourceforge.net>




Fig. 2. A typical experiment. Total population, number of roosts, and number of cheaters (multiplied by 10) are shown. Simulation starts with 300 agents in 10 roosts, for 8000 days, 10% of cheaters at startup.

This happens after a critical period during which cheaters go extinct, and the global fitness of the whole population is almost on the verge of collapse. After cheaters are totally extinguished, the population starts growing rapidly and indefinitely.

This observation leads to appraise the role of roosts. In fact, if the whole population were sharing one roost (see Fig. 4 for one example), the presence of cheaters would lead it to certain extinction. With a single roost, most simulations converge to zero after some period, with or without later resurgence; in any case, the presence of cheaters increases until they cause a catastrophic lowering of the population, after which they start to increase again until extinction. No reciprocity could emerge in a world in which cheaters are allowed to repeatedly exploit others, incurring neither retaliation nor isolation.




Fig. 3. Total population, number of roosts, and number of cheaters (multiplied by 10) are shown. Simulation starts with 300 agents in 10 roosts, for 8000 days, 40% of cheaters at startup.

Under these extreme conditions, in which help exchange is vital, after having exploited their altruistic in-roosts to death, cheaters find no way to face adversity, and are soon bound to share the same fate. If a few altruists happen to survive the extinction of cheaters, they will soon repopulate the roost and produce new ones. If no-one survives cheaters, which is the most likely event since they survive longer than their good fellows, the roost will collapse.

On the contrary, the phenomenon of roosting radically modifies the situation. Due to the presence of cheaters, most of the roosts disappear. However, if at a certain point any roost without cheaters will appear, it will grow and repopulate the world. This is what happens in Fig. 2, after a demographic catastrophe.




Fig. 4. Single roost with 300 agents, no new roost formation, 10% cheaters, 20000 days.

As for the question if the evolutionary advantage of food sharing is better explained by inclusive fitness or group selection, we tried several measures to discriminate between them. A clear signal in favour of group selection would be given if we were to find that helpers, that is, agents that helped others often during their lifetime, had less offspring than other altruists. The problem here is subtle, since agents that have had more chances for reproduction, i.e., have lived longer, also automatically have had more chances for helping others. For this reason, the correlation we found between help and fitness — defined as both the number of descendants and the number of living descendants — is obviously positive, but inadequate to support either explanation. Trying to pinpoint the effect, we selected a subset of the initial agent population, characterised as both long-lived and altruists. In the simulation, agents reaching an age of 10 years are automatically removed from the population.

However, by tracking their descendants, we can look for a correlation between the number of helps given and the size of the donor's offspring.

We run a set of 100 simulations, with 10 roosts of 30 agents each. Cheaters' percentage was set to 10%; simulations stopped at 5000 day, and only values for agents surviving from the beginning of the experiment to maximum possible age were extracted. For these agents, we kept track of the number of helps given and of the number of agents in the whole lineage up to the end of the simulation. We found only a very low factor (0.07) of positive correlation between these measures. In substance, there is not a linear relation between the number of help given in life and the total number of offspring in the simulation for the subset we choose to analyse.

3.2 Preliminary Discussion

The findings presented above point into two distinct directions.

On one hand, altruistic roosts survive longer and reproduce far more than mixed roosts, which include cheaters. Indeed, cheaters tend to cause the extinction of their roosts, including themselves. This result is not particularly original, since it essentially confirms what was found by Wilkinson through his simulation study of vampire bats' food-sharing. Although distribution of cheaters over the roost population makes a difference in terms of the reproductive success of the whole roost, even one single cheater may be enough to lead the roost to extinction. This is so for two reasons: first, cheaters usually survive longer than their fellow altruists; second, in a competitive environment, any roost that happens to be cheater-free will generate more offspring, both in terms of roosts and in terms of individuals. In this sense, the roost is a critical unit of selection, whose efficiency in finding out and eliminating cheaters is amazing. The relevance of roosts could be lessened by a different kind of cheater, not included in this study: a roaming cheater, which distributes the cost of its presence over the roosts it has access to. However, it must be noted that in nature, individuals requesting entrance to a new roost find it very hard to be accepted by the new roost inhabitants; an entrance fee could have co-evolved in order to protect the altruistic mechanism.

The second direction pointed to by our findings concerns the role of roost in the evolution of altruism. In this respect, grouping, or better, roosting seems to matter. If it is the case that altruistic groups survive longer than non-altruistic ones, the reverse is also true: altruists have better chances to survive if they roost together. In a one-roost world, inhabited by cheaters and altruists, the chances of survival of altruists — and of the entire population — would be close to zero. Conversely, in a multi-roost world, where altruists happen to co-exist in variable distribution with cheaters, an inter-roost competition for reproduction occurs. Since cheaters lead to the extinction of their roosts, only altruistic ones will survive enough as to reproduce, and these will soon populate the world. In this perspective the group-selection argument seems

to receive support. Groups act as units of selection and reproduction, much alike individual organisms.

However, this finding *per se* does not say much about the internal rationale of altruism. If it supports group-selection, it does not disclaim the concurrent sociobiological theory of reciprocal altruism. Indeed, precisely because no rule for reciprocity is explicitly represented in our simulation model, the only way for altruists to survive is to roost together, waiting, so to speak, for cheaters to go extinct. In this sense, and rather tautologically, reciprocity emerges only when cheating disappears.

Up to now, our findings also indicate that actual donors do not reach a significantly higher rate of survival and reproduction than the rest of the population; indeed, the correlation factor is too low to come to a conclusion on this matter. No negative correlation definitely disclaiming reciprocal altruism has been found. The global increase of fitness of the roost population is not obtained at the expenses of one share of it (the actual donors). On the other side, the final generations do not necessarily include the lineage of the actual donors. This is possibly due to the simplicity of the algorithm, which allows for no specific rule of reciprocity. On the other hand, it corresponds to the simple rationale of reciprocal altruism, for which agents neither aim nor calculate the probability of reciprocity, which should be an emergent effect of the altruism fitness. However, if actual donors are not reciprocated, their fitness decreases to the benefit of the global fitness. But if this is the case, as appears to be in our findings, inclusive fitness is not the reason accounting for the spread of altruism. Indeed, group selection theory gains the ground that is lost by the reciprocal altruism theory. Grouping matters and helps altruists to survive and reproduce even in presence of cheaters. Under the shelter of their roosts, animals helping each other will have better chances to reproduce, although some of them, finding themselves in roosts with high cheaters content, will pay dearly for such behaviour.

4 Complex Loop: Spreading of Altruistic Behaviour

As said in the introduction, the loop between an emergent macro-social level and its feedback on the micro-social level may be more or less complex. In the former part of this paper, we have seen a simple loop, in which an emergent macro-effect (groups' fitness) of the agent behaviour (food-sharing) directly acts upon the micro-level preventing donors from being exploited, thereby increasing altruists' fitness.

So far, we have shown that, at least in a possible (artificial) world, prosocial behaviour can emerge and stabilize, despite a minority of cheaters, thanks to a social circuit that we described as a *simple loop*. If in nature things actually worked out in such a way, although plausible, is far from having being ascertained.

However, these dynamics are much too simple. As already shown, the roost dimension matters: the larger the roosts, the lesser their capacity to protect from cheating. Hence, the simple loop proves rather fragile, as altruists can only put up with a fairly small minority of cheaters. As shown by analytical modelling, large and isolated roosts with randomly assorted altruists and non-altruists have none or poor chance of survival.

Should we then forget about a more robust form of altruism? Certainly, the simple loop appears to be sufficient⁴ for the evolution of food-sharing and the survival of vampire bats. This species has probably reached a demographic/ecological equilibrium with a stable distribution of the population over a myriad of small and fairly segregated roosts, in which newcomers are not let before providing long-lasting and robust evidence of an altruistic attitude. But what about other species, e.g. humans, which, for most part of their evolutionary history, have been characterised by a consistent demographic increase and the attitude to form ever enlarging settlements?

Let us assume that collective dilemmas were frequent in the evolutionary stage of our species, an assumption that is far from unrealistic [8, 9] given the necessity of defenceless individuals to live in groups, on one hand, and the scarcity of resources usually threatening the cohesion of these groups, on the other [10]. How could such a species have evolved at all, if a handful of cheaters was enough to invade and desegregate a settlement?

Here is where a more complex dynamics between the micro- and the macro-social levels seems to be strongly needed. In our view, the bottom-up process, from the micro-social to the macro-social level, coincides no more with the emergence of specific social structures (e.g., a myriad of small groups), but with the appearance of an immaterial social structure, a prescription that agents are expected to observe, namely the norm of reciprocity. By this means, altruism is supposed to spread over newer portions of the population, allowing the group to resist the inevitable offensive of cheaters. The emergent norm, acting on the individuals, modifies their behaviours. How can this be done in a robust way, i.e. in such a way that individual donors can put up with and reduce the costs, or the risks, of bad encounters?

Many good evolutionary scientists have been playing around with this question. However, the typical answer which this usually receives — in terms of familiarity, repeated interaction, or past experience — does not allow us to proceed further than we did with our simple loop: small stable groups allow agents to recognise and possibly avoid one another. But what about larger open groups? How can norm-abiders reduce the costs of norm obedience, if they cannot escape transgressors? Which mental construct is required for the norm to be applied in a robust way, thereby increasing the probability that the norm itself will survive?

These questions show that the complex loop includes two steps:

⁴ Sufficient but not necessary: for that we should wait for conclusive field evidence.

- an *emergent* process (the way up of the circuit): how does the norm come into existence?
- an *immergent* process (the way down): which mental constructs are required for the norm to be executed and stabilised?

Here we set to answer the latter question, while postponing to further work the study of the former.

4.1 Reciprocity: Artificial Vampire Bats with Credit Network

Reciprocity is considered as the main factor in the evolution of cooperation. Of the two forms of reciprocity, direct and indirect, however, the former was above argued to be inapplicable to the present context. How about indirect reciprocity, then?

The abundant literature on this issue [11, 12, 13] points to the effect of punishment, possibly based on reputation, on the emergence of indirect reciprocity. According to these studies altruism spreads via reciprocation. Is this applicable to our scenario? Let us see.

We simulated a population of artificial bats, with variable percentage of cheaters (bats that never donate food). In Fig. 5, the effect of a variable number of initial cheaters over global survival rate is shown. The left chart shows the number of living agents at the end of simulation, for 20 different values of the initial cheater/altruist ratio; for each value we run 200 simulations, for a time span of about 30 years. On the right, we show the success rate for the same set of simulations, defined as the percentage of the 200 runs in which at least an agent survives to the end.




Fig. 5. Left: boxplot, number of living agents at simulation end by cheating rate. 200 runs for each value of cheating percentage. Right: successful runs (number of simulation where population does not extinguish), by cheating rate.

Apparently, even a small number of cheaters leads the population to collapse. The population starts to decline as soon as the number of cheaters grows over a bare 10%; no roost reached the end of the simulation when the percentage of cheaters reaches or goes beyond 40%. Sooner or later, depending on the initial number of cheaters, a mixed population is bound to go extinct. It is easy to see why: in isolation, cheating is a self-defeating strategy. Cheaters reproduce more than altruists, as long as they find any to exploit. As altruists decline, the simulation returns to a no food sharing condition. Consequently, cheaters go extinct in a couple of years. Which ingredient is then needed for altruism to spread despite cheaters?

Optionally, we provided agents with the capacity for individual recognition. Every time a donation occurs, an agent adds to its memory the identity of the recipient. We maintained this information through a network, the nodes of the net being the agents. At any donation the credit network is updated. In fact, either a previous donor is refunded — in which case its credit is extinguished and one link removed — or a new credit is formed and a new link is activated between current donors and their recipients. Whenever donors are reciprocated, their corresponding credits are cancelled. In our implementation, to avoid reiterated initialization of the network, we decided to pass it on to one's offspring, which then inherit parents' features and credits. Consequently, a given credit can be extinguished during the donor's life or after its death to the benefit of its offspring. (Obviously, the more the credits passed on to future generations, the higher the probability of survival of one's offspring.) If the credit network is activated, it gets investigated any time a request of help is received and, if the postulant is found to be a debtor with a number of unreciprocated donations over a given threshold, help is denied. This action is seen as a mechanism to limit the exploitation of the altruistic behaviour by selfish agents, who — since they never reciprocate — will progressively lose the opportunities they receive from in-roosts.

When one thinks of it, the grooming network creates a familiarity as well as a reciprocity basin: giving help allows animals to achieve credits, which will be extinguished if and when help is returned. A lucky hunt may last the short space of one night, and a fat guy may soon shrink in starvation. Hence, it will be urged to go out for grooming in the hope to meet with a luckier (and fatter) debtor. In less metaphorical terms, the grooming network facilitates re-encounters and therefore the extinguishing of credits.

This credit network is checked any time a request of help is received. In a more restrictive condition, only if no credit link is already active with the requestor, the agent will give help. Otherwise, help is denied. We found that this condition is too restrictive, and does not allow for population survival. In a less restrictive condition, shown in this paper, help is denied only when the same requestor asks for help more than two times consecutively.

Agents search for potential donors within the grooming network. Only one trial is available. If help is denied, the postulant is bound to die.




Fig. 6. Artificial vampire bats with and without credit network. Successful runs (number of simulation where population does not extinguish), by cheating rate.

As shown in Fig. 6, the credit network does not improve things. Indeed, to investigate own credits when deciding whether or not to donate leads to a comparable number of cheaters in the population, with the same reduction of the global survival rate.

Apparently, these findings are counterintuitive: turning off credits before donating again appears as a quite rational strategy. How come it is non-influential rather than helpful? Looking at what effectively happens within the simulations is rather instructive, indeed. The credit network has one main effect: it allows for lesser donations. Agents more often deny help. But this punishment is unfair, as it penalizes not only cheaters but also unlucky altruists. Moreover, it is rather extreme: requestors will die out and those, among them, which cheat will have no time to learn. Learning to reciprocate requires time, but in such harsh life conditions, time is not available.

Should we conclude that learning is unfeasible? Or, perhaps, should we turn to a different form of learning leading to milder, rather than tougher, criteria for donations, including unconditioned altruism? Is it possible to learn such a form of altruism? Is it compatible with an autonomous agent? To these questions we turned our attention in the successive study, which will be described below.

4.2 Dynamic Goal-Directed Agents

In this Sect. we start to model a different type of agent. Until now all the agents were designed only to reproduce the activities of a common vampire bat in its natural environment. Let us call these first agents **Simple** agents. Every Simple agent is endowed with a set of routines that (s)he will use in order to obtain food from the environment and help from other agents. From now on, we will use the name **Smart** agents for agents endowed with altruistic motivations of variable intensity.

We aim at enabling autonomous agents (i.e., agents endowed with and acting to achieve their own goals) to become altruists. Rather than incompatible with altruism, autonomy augments agents' adaptability thanks to goal-dynamics. The social environment — as well as the physical one — deeply affects autonomous agents. Autonomous agents endowed with goal dynamics respond to events that affect them by modifying their goals, harbouring from the least to the most benevolent strategies towards their fellows.

Dynamic goal-directed action is an essential aspect of cognition [14, 15]. In a cognitive architecture, a goal is a highly dynamic mental construct, which, thanks to beliefs, may be generated, abandoned, worked out, suspended, interrupted, achieved, compromised, etc. as an effect of its varying intensity. In the simulation model at the present stage of development, goal-dynamic is only partially implemented. Autonomous agents change their behaviour according to their goals, and change their goals on the base of their beliefs. However, beliefs are not explicitly modelled in our system and goals vary as to their motivational force, rather than in their representational content. Future extensions of this work will investigate qualitative goal-dynamics.

Goal-directed agents are an extension of the simple ones. They have access to the same set of actions: repertoire of actions (give blood, deny help) is kept constant. However, they differ in the decision-making mechanism with respect to decisions about donations. Every time a request for blood is received two opposite goals are activated:

- a Normative Goal (*NG*), which prescribes to answer the requests received;
- a Survival Goal (*SG*), according to which agents try to keep constantly high their blood-autonomy.

The conflict between these two opposite goals is resolved toward the goal with the higher value, which in turn will activate the appropriate action. So, the goal values are a measure of the altruistic motivation of each smart agent.

In our case five plausible cases are derived. In Table 1, the outputs of this motivational interplay as agents actions are given, together with suggested named for the five strategies; action are characterised by the amount of autonomy sufficient to activate donation, in hours.

We then endowed agents with rules for modifying the values of their goals, consequently changing their strategy while interacting with each others. For

Table 1. Behaviour associated to each strategy. The goal dynamic (applied to the Normative Goal, NG) leads agents to pass from one strategy to another, either more or less altruistic.

STRATEGY	NG	ACTIONS
Cheaters	< -2	Always deny help
Prudents	-1	If you have 48 h of autonomy donate at a cost for you of 6 h. Otherwise deny help.
Fair	0	If you have 48 h of autonomy donate for 12 h. If you have 24 h of autonomy donate for 6 h. Otherwise deny help.
Generous	$1 - 2$	If you have 48 h of autonomy donate for 24 h. If you have 24 h of autonomy donate for 12 h. Otherwise deny help.
Martyrs	> 3	Even with 18 h of autonomy left donate at a cost for you of 6 h.

simplicity, we implemented only those affecting the altruistic goal (NG, keeping constant the survival one). In particular two heuristics have been explored:

Action-Based Learning The value of altruistic goals goes up or down as an effect of one's and others' actions. If one receives help, the force of the altruistic motivation increases, whereas it decreases if one gives help. This heuristic is apparently fair, but in fact is biased toward altruism. Everybody's altruistic motivations can intensify since everybody receives help. But only altruistic agents can see their motivations decreasing, because only such agents can give help. Therefore, we turned to more symmetric rules.

Credit-Based Learning After a given time interval (the average time in which at least one unsuccessful hunt can occur per agent, that is, about two hundred days), own credits are investigated. If help has been received within that period, the goal goes up; in the opposite case, it goes down. These heuristics should be regarded as complementary effects of an in-built norm of reciprocity [16]. In the present study we will take these heuristics for granted, postponing to follow-up simulations the important evolutionary question as to the minimal mental requirements for agents to form these rules.

Noticeably, here, autonomy is other than weak rationality. Unlike rational ones, goal-directed agents do not necessarily care about costs (see [17]). Whereas a prudent algorithm act as if agents take their utility into account (performing an altruistic action at the minimum cost), the goal-directed is

not. However, both are autonomous, in the sense that both are endowed with internal criteria for decision-making.

4.3 Findings

Simulations have been initialized either by setting agents' goals to equal values (one strategy at the onset), or to all values (all strategies at the onset), or finally to extreme values, cheaters and altruists (either martyr or fair or prudent). Results show that different strategies — corresponding to different patterns of relationships among agents' goals (see Fig. 7, 8) — emerge and their difference increases over time.




Fig. 7. Strategy differentiation with action-based learning (cheaters *C*, prudent *P*, fair *F*, generous *G*, martyrs *M*). Averages of population divided by strategy. From left to right, initial populations of all Prudents, all Fair, all Martyrs.

What is more interesting is that the dominant strategy appears to be the most altruistic. Whatever the initial strategy, agents learn to be more altruistic than anything else, and the population prospers.

The results of the different models explored are shown to be non-trivial. Indeed, prudent agents perform equally when they are modeled either as rigid and as dynamic systems: compare the results for 20% cheaters, given in Table 2, with the results obtained above; prudent rigid agents have a success rate of 0,67 without credit network and of 0,655 with credit network. The opposite is true for unconditioned altruists. However, in general, the more altruistic strategies are always dominant and lead the population to an exponential growth, which is controlled in our simulations by means of a carrying capacity set to 200.




Fig. 8. Strategy differentiation with credit-based learning (cheaters *C*, prudent *P*, fair *F*, generous *G*, martyrs *M*). Averages of population divided by strategy, initial populations of all Martyrs.

As expected, the comparison between the two heuristics for learning shows that credit-based learning is more symmetric than the other, which is instead slightly biased in favour of altruism. As a consequence, it is no surprise that altruistic strategies are less dominant with this than with the other rule. Still, in both cases, the majority of agents learn to exhibit an unconditioned form of altruism.

Obviously, the more altruistic the dominant strategy, the more the population grows, which confirms the well-known law that altruistic populations do better than non-altruistic ones. The question addressed in this paper is whether unconditioned altruism can be learned and whether it can be learned by autonomous intelligent agents. Our findings seem to provide a tentatively positive answer to this question.

In short, altruism emerges and spreads in populations of dynamic goal-based systems, at least in populations with a relatively small number of agents where everybody can change either for better and for worse. Under these conditions, agents learn to be altruist more than the other way around.

Table 2. Success rate per heuristic and initial population (20% cheaters). Initial populations of Prudents (*P*), All five strategies in equal proportion (*A*), and Martyrs (*M*).

HEURISTIC	P	A	M
Action-based	0,03	0,36	0,98
Credit-based	0,312	0,0	0,14

5 Final Discussion

In this paper, we have been trying to show two types of MML processes: what we have called simple loop and complex loop.

In the former, from a large number of small social structures (roosts), those with a strong numeric prevalence of altruists are likely to get reproduced and to consist of altruistic offspring. Roosts in turn protect their members, providing the social barriers that prevent non-altruists from invading the roosts and leading inhabitants to an early extinction. The simple loop was then found to be good enough for the evolution of altruism in a species demographically and socially stable.

Under different demographic and social conditions, instead, a more complex loop seems to be needed. In it, the function played by the social structures in the simple loop (protect altruists) is now accomplished by an enforcing artefact, namely a norm of reciprocity. This is supposed to mediate the process and allow the altruists to put up with and neutralise the non-altruists. The question is whether the agent typology involved in the simple loop is still sufficient or other properties and a more complex architecture of the agent must be implemented in order for altruists to observe the norm without being exploited by cheaters up to extinction. Our simulation findings seem to show that a dynamic form of intelligence, i.e. the capacity to learn from the consequences of others' actions on oneself, is needed for a complex loop of this sort to be realised. With dynamic intelligent agents, diversified and robust forms of altruism emerge, which can put up even with a majority of cheaters. It may be of some interest to observe that the most extreme altruistic forms tend to get reinforced more than the milder ones, as they appear to be more contagious than the latter.

There is a couple of questions that this paper did not address, and which should be investigated in future works. First, how does a norm of reciprocity come out? If we do not consider it only as a mental module, but also as a social prescription, how did it emerge? To be noticed, this question should not be confused with the question addressed (and to a large extent, answered) by the theorists of conventions *à la* Lewis, according to which norms are behavioural conformities. On the contrary, we ask how a social prescription, not a social regularity, comes out. So far, to our knowledge, nobody did answer such a question.

The second, and related, question concerns the dynamics of agents' goals. In this paper, we have been investigating the effects of a quantitative dynamics: agents' goals change value as a function of a number of contingent beliefs. However, we did not show how and why agents acquire goals anew, what is a crucial step in the study of normative compliance: where comes from the goal to observe the norms of reciprocity?

We will address both questions in our future work.

References

1. Ostrom, E. (2000) Social capital: A fad or fundamental concept? In: Dasgupta, P., Seragilden, I., eds.: *Social Capital: A Multifaceted Perspective*. World Bank, Washington DC.
2. Maynard-Smith, J. (1964) Group selection and kin selection. *Nature* **201**:1145–1147.
3. Brembs, B. (1996) Chaos, cheating and cooperation: potential solutions to the prisoner's dilemma. *Oikos* **76**:14–24.
4. Wilkinson, G.S. (1984) Reciprocal food sharing in the vampire bat. *Nature* **308**:181–184.
5. Wilkinson, G.S. (1990) Food sharing in vampire bats. *Scientific American* **2**:64–70.
6. Trivers, R. (1971) The evolution of reciprocal altruism. *Quarterly Review of Biology* **46**:35–56.
7. Dawkins, R. (1976) *The Selfish Gene*. Oxford University Press.
8. Cosmides, L., Tooby, J. (1992) Cognitive adaptation for social exchange. In: Barkow, J., Cosmides, L., Tooby, J., eds.: *The Adapted Mind*. Oxford University Press.
9. Ostrom, E. (2003) Altruistically inclined? the behavioral sciences, evolutionary theory, and the origins of reciprocity by alexander j. field. *Journal of Interdisciplinary History* **34**:63–64.
10. Wright, R. (1994) *The Moral Animal. Why We Are the Way We Are: the New Science of Evolutionary Psychology*. Pantheon Books.
11. Bowles, S., Gintis, H. (2001) The evolution of strong reciprocity. Santa Fe Working Paper.
12. Boyd, R., Richerson, P.J. (1992) Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethology and Sociobiology* **13**:171–195.
13. Nowak, M.A., Sigmund, K. (1998) Evolution of indirect reciprocity by image scoring. *Nature* **393**:573–577.
14. Castelfranchi, C. (1997) Principles of limited autonomy. In: Tuomela, R., Holstrom-Hintikka, G., eds.: *Contemporary Action Theory*. Kluwer.
15. Conte, R. (2000) Memes through (social) minds. In: Auger, R., ed.: *Darwinizing Culture: the Status of Memetics as a Science*. Oxford University Press.
16. Cosmides, L., Tooby, J. (1994) Better than rational: Evolutionary psychology and the invisible hand. *American Economic Review* **84**:327–332.
17. Conte, R., Pedone, R. (1998) Finding the best partner: The part-net system. In: Gilbert, N., Sichman, J.S., Conte, R., eds.: *Multi-agent systems and agent-based simulation*. Springer, Berlin 156–168.

How Are Physical and Social Spaces Related? – Cognitive Agents as the Necessary “Glue”

Bruce Edmonds

Centre for Policy Modelling, Manchester Metropolitan University, United Kingdom
bruce@edmonds.name

Summary. The paper argues that in many (if not most) cases, explicitly representing aspects of both physical and social space will be necessary in order to capture the outcomes of observed social processes (including those of spatial distribution). The connection between social and physical spaces for an actor will, almost inevitably involve some aspect of cognition. Thus, unless there is evidence to the contrary it is unsafe to try and represent such social distribution without representing key aspects of cognition linking social and spatial topologies. This argument is demonstrated by two counter-examples: an abstract simulation extending Schelling’s cellular automata model of racial segregation to include the social communication of fear; and a more descriptive simulation of social influence and domestic water consumption. Both models are sufficiently credible that one could not rule similar processes as occurring in reality, but in both the social and physical spaces that the agents are embedded in is critical to the global outcomes.

1 Introduction

Even social entities such as: humans, animals, households, firms etc. exist in physical space. The way these entities are distributed in that space is frequently important to us. Some of this distribution can be clearly attributed to economic and environmental factors that are essentially independent of the social interaction between these entities. However it is overwhelmingly likely that some of the distribution is due to the interaction between these entities¹ – i.e. the spatial organisation of the collection of such entities is (at least partially) self-organised via processes of social interaction.

It is difficult to study such self-organisation using purely statistical techniques. Statistical techniques are more suited to dealing with aggregate tendencies where the deviation from these tendencies in particular cases can be considered as essentially random. Thus not only can the detail of the spatial

¹ That is, the desired aspects of their spatial distribution can not be modelled without interaction between the entities that are distributed.

organisation be lost in the process of aggregation but in many social cases the deviations are not random. Furthermore statistical models, in practice, require fairly drastic assumptions in order to make them amenable to such techniques.

Mathematical models (e.g. those expressed as differential or difference equations) have the potential to capture the self-organisation, but only by disaggregating the model into many separate sets of equations for each entity (or place). This, except in a few special cases where strong assumptions hold, makes any analytic solution impossible. Thus if one tries to apply such techniques to study self-organised distribution one usually ends up by numerically simulating the results, rather than exploiting the analytic nature of the formalism.

It is for these reasons that the study of such self-organisational processes has been advanced primarily through the use of individual-based computational simulations. These are simulations where there are a number of individual entities in the simulation which are named and tracked in the process of the computation. It is now well established that considerable complexity and self-organisation can result in such models even where the properties and behaviour of the individuals in the models are fairly simple. Many of these models situate their component individuals within physical space, so that one can literally see the resulting spatial patterns that result from their interaction [1].

Some of these individual-based models seek to capture aspects of communicative interaction between actors. That is, the interaction between the modelled entities goes beyond simple cause and effect via their environment (as in market mechanisms, or the extraction of common resources) but tries to include the content or effects of meaningful communication between the actors. Another way of saying this is that the actors are socially embedded [9, 5]. That is to say that the particular network of social relations is important to the behaviour of the individual – or, to put it another way, a model which “assumes away” these relations will distort important aspects of the phenomena. Examples of this might include the spread of new land uses among a community of farmers or a request for households to use less water. In such models it is often the case that influence or communication occurs between individuals who are spatial neighbours – that is to say that physical space is used as a ‘proxy’ for social space. In such models communication or influence between individuals is either limited to local neighbourhoods or is totally global.

However, in the modern world humans have developed many media and devices that, in effect, allow communication at a distance². For example, farmer may drive many miles to their favourite pub to swap farming tips rather than converse with their immediate neighbours. Thus the network represented by

² Of course, when one looks at the detail of their operation they always involve some action, for example in the propagation of electromagnetic waves, but given the limited granularity of social models these are effectively non-local.

the communication patterns of the actors may be distinct from the spatial pattern. Recently there have been some models which seek to explore the effects of other communicative topologies. There has been particular focus on “small world” topologies, on the grounds that such topologies have properties that are found among the communicative webs of humans, in particular the structure of hyperlinks on the Internet. However such models are (so far) divorced from any reference to physical space, and focus on the organisation and interactions that can occur purely within the communicative web.

There have been very few models which explicitly include actions and effects within a physical space as well as communication and action within a social space. This paper argues that such models will be necessary if we are to understand how and why human entities organise themselves in physical space. A consequence of such an approach will involve a move away from relatively simple individual-based simulations towards more complex agent-based simulations due to the necessary encapsulation of the agents who act in space and communicate with peers. Thus some sort of cognitive agency will be necessary to connect the communication with the action of the individuals - communication only makes sense when put in the context of the actions of the individuals concerned as well as the actions of others that impinge on them and sensible action can only result when communications have been internalised. This parallels Carley’s call for social network models to be agentified [2].

Thus this paper argues that such cognitive agency will be unavoidable in adequate models of the spatial distribution of human-related actors and, further, that the spaces within which action and communication occur will have to be, at least somewhat, distinct. Thus the burden of proof is upon those modellers who omit such aspects - that the “short-cut” of abstracting away the cognition of individuals in models will not substantially bias the results away from that which is observed.

In other words this paper is explicitly questioning the extent to which models which fail to take into account these social facts will be successful at capturing self-organised social phenomena. I am suggesting that abstract “physics-type” models will not further our understanding of such phenomena and in particular the spatial distribution of many such phenomena.

To establish the potential importance of the interplay between social and physical spaces, and to illustrate the approach I am suggesting, I exhibit a couple of agent-based simulations which involve both physical and social spaces. The first of these is a more abstract model whose purpose is simply to show how the topology of the social space can have a direct influence upon spatial self-organisation, and the second is a more descriptive model which aims to show how a suitable agent-based model may inform observation of social phenomena by suggesting questions and issues that need to be investigated.

These models are counter-examples - it is sufficient for my point that they show that it is credible that similar processes are occurring in the phenomena we seek to model and, hence, that one can not rule out the necessity of representing both physical and social networks as well as the cognitive processes

which link these without good reason. This throws the burdon of proof back to those who abstract some of these away without knowing whether it is safe to do so. This is part of a more fundermental disagreement about the nature of modelling social phenomena and how would should approach this.

2 Example 1 – The Schelling Model of Racial Segregation Extended with a Friendship Network and Fear

2.1 The Schelling Model of Racial Segregation

To illustrate the interplay of social and physical spaces, I go back to Schelling’s pioneering model of racial segregation [12]. This was a simple model composed of black and white counters on a 2D grid. These counters are randomly distributed on the board to start with (there must be some empty squares left). There is a single important parameter, c , which is the ratio of counters of its own colour among the counters in its immediate neighbourhood (see the first diagram in 1 below) below which the counter will seek to move. Each generation of this game, each counter is considered and if the ratio of same coloured counters in its neighbourhood is less than c then it randomly selects an empty square next to it (if there is one) and moves there.

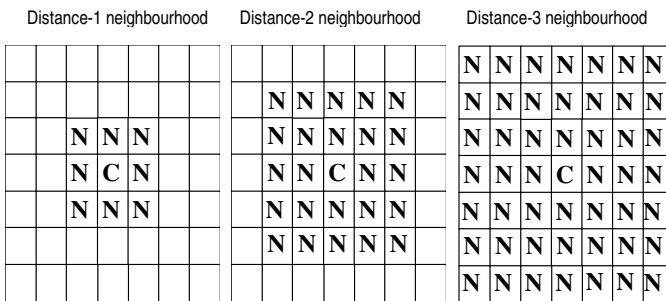


Fig. 1. Neighbourhoods of distances 1, 2 and 3 respectively in the 2D Schelling model

The point of the model was that self-organised segregation of colours resulted for surprisingly low levels of c , due to the movement around the 'edges' of segregated clumps. The interpretation is that, even if people are satisfied with their location if only 40% of their neighbours are the same colour as them, then racial segregation can still result – it does not take high levels of intolerance to cause such segregation. Figure 2 shows three stages of a typical run of this model for $c = 0.5$ and a neighbourhood of distance 1.




Fig. 2. A typical run of the Schelling model at iterations: 0, 4, and 40 ($c = 0.5$)

Figure 3 is a graph of the typical outcomes of this model in terms of the critical parameter, c , and the final average of the proportion of each counter’s neighbours who are of the same colour – this is an indication of the extent of the segregation that occurs. One can see that segregation gradually increases from low levels of c until at $c = 0.35$ a significant level of self-organised segregation is the result. The maximum is somewhere between $c = 0.6$ and $c = 0.65$. Above this level the segregation drops sharply off. This is due to the impossibility of all counters finding positions with 75% like counters and so counters at the edges of segregated clumps are continually randomly relocating destroying any clumping. As discussed in [7] this is not a fault of the model since its purpose was to show, as simply as possible, how segregation could occur at low levels of intolerance.




Fig. 3. The variation of resulting segregation with levels of intolerance (distance-1 neighbourhoods)

It is interesting to note that, even in Schelling’s model the social topology (in this case the neighbours each counter considers in the decision to move or not) can have an effect. Figure 4 shows the corresponding graph to Fig. 3 for runs of the Schelling model with a neighbourhood of distance 3. Since each counter has many more neighbours (in the later case 56 of them) it is

more likely that one is satisfied with a random mix at the beginning for low values of c . In other words, it is much less likely that a counter will find itself attached to a monolithic clump of the other colour at the beginning and so will not ever move. This “flattening” of segregation for low levels of c (i.e. $c < 0.35$) depends upon the random initialisation of the model. If one started from an already segregated pattern then increasing the size of neighbourhoods would have a less significant effect.




Fig. 4. The variation of resulting segregation with levels of intolerance (distance-3 neighbourhoods)

2.2 The Extended Model

I have extended the basic Schelling model by adding an explicit “social structure”. The purpose of this extended model is simply to demonstrate that social networks can interact with the physical space in ways that significantly affect the outcomes, in particular the spatial distribution that results. This is done in a model whose processes might credibly mirror observed social processes. The model is fairly simple – its structure is not justified by detailed psychological or social studies but rather it is only a demonstrator of processes that could not be ruled out as occurring in observed phenomena. It is thus a counter-example – it demonstrates that if, in this example, one abstracted away the social network one would have come up with significantly wrong outcomes. Hence that such abstraction is dangerous, unless there are good reasons to believe otherwise.

In this extended model the social structure is in the form of a friendship network. This is a directed graph between all counters to indicate who they consider are their friends. The topology of this network is randomly determined at the start according to three parameters: the number of friends, the local bias and the racial bias. The number of friends is how many friends each counter is allocated. The local bias controls how many of a counter’s friends come from the local neighbourhood – a value of 1 means all its friends come from its initial neighbourhood, and a value of 0 means that all counters are

equally likely to be a friend wherever they are. The racial bias controls the extent to which a counter's friends have its own colour – a value of 1 means that all its friends have the same colour as itself and a value of 0 that it is unbiased with respect to colour and friendship. In this model this friendship structure is then fixed for the duration of the run. This network has several functions: firstly, influence only occurs from a counter to a friend, secondly, if it has sufficient friends in its neighbourhood a counter is unlikely to seek to move and, thirdly, (depending on the movement strategy set for the run) if a counter has decided to move it may seek to move nearer to its friends (even if this move is not local).

The motivation for moving is different from the original model as well. Instead of being driven by intolerance, the idea is that it is driven by fear. Each counter has a fear level. This is increased by incidents (which occur at completely at random) happening in their location and by fear being transmitted from friend to friend. There are two critical levels of fear: when the fear reaches the first level the counter (randomly with a probability each time) transmits a percentage of its fear to a friend, this transmitted fear is added to the friends fear. Thus fear is not conserved but naturally increases and feeds on itself. When fear reaches the second critical level the counter seeks to move to be closer to friends (or away from non-friends). It only moves if there is a location with more friends than its present location. When it does so its fear decreases. Fear also naturally decays each iteration to represent a sort of memory effect. This is not a very realistic modelling of fear, since fear is usually fear of something, but can be cumulative and is communicable. The incidents that cause fear occur completely randomly over all locations with a low probability. The other reason for moving is simply that a counter has no friends in its neighbourhood.

Thus in this model the influence of counter colour upon counter movement is indirect: counter colour influences the social structure (depending on the local and racial biases) and the social structure influences relocation. Thus I separate its position in space and the social driving force behind relocation.

The dynamics of this model are roughly as follows. There is some movement at the beginning as counters seek locations with some friends but initially levels of fear are zero. Slowly fear builds up as incidents occur (depending on the rate of forgetting compared to the rate of incidence occurrence). Fear suddenly increases in small “avalanches” as it reaches the first critical level in many counters – this is because fear is suddenly transmitted over the social network causing others to pass the critical level etc. When fear reaches the second critical level then counters start moving towards where its friends are concentrated (or away from non-friends depending on the move strategy that is globally set). This process continues and eventually settles down as counters only move if they can go to where there are more friends than its current location.

Table 1 shows some of the key parameters of the simulation for the range that I will talk about. I have chosen this range of values, because it is relevant

to my point and seems to be the critical region of change in this model. The three parameters I will vary are those that affect the topology of the social network: the number of friends each counter has; the bias towards (initially) local friends; and the racial bias. Each of these parameters are tried with each of 5 values, giving 125 runs in total. In each of the graphs below each line represents the average value over 25 different runs.

Table 1. Parameter settings

Parameter	Range of Values
Number of Cells Up	20
Number of Cells Across	20
Number of Black Counters	150
Number of White Counters	150
Neighbourhood Distance	2
Local Bias	0.45-0.55
Racial Bias	0.75-1.0
Num Friends	2-10

The Effect of Network Connectedness

The greatest effect results from the number of friends each counter has, i.e. the connectedness of the social network.




Fig. 5. The average fear level for runs with different numbers of friends

Figure 5 shows the average fear levels for runs with different number of friends (25 runs each). The fear levels increase exponentially due to the fact that counters (probabilistically) communicate a proportion of their fear – an interpretation of this could be that if they are more fearful they communicate more fear. The amount of fear increases with the number of friends each

counter is allocated up to the value of 8 but this decreases for the runs with 10 friends. There seems to be two different competing effects: the more connected the social network the more fear is communicated and multiplied but also a high number of friends allows counters more opportunities to move and hence to dissipate a chunk of their fear.

Figure 6 shows the average number of communications and Fig. 7 the average number of movements. It is clear that the more friends one has the more chances there are of moving to a location with a higher number of friends (who themselves might well move etc.) so that the simulation takes a lot longer to “settle down” to a situation where no one can move somewhere better. It is notable that in Fig. 6 the number of communications is initially greatest for the runs with 10 friends but then drops down below those runs with 6 or 8 friends.




Fig. 6. The average of the number of communications for runs with different numbers of friends




Fig. 7. The average number of movements for runs with different numbers of friends

The greater possibility of movement for those runs with more friends means that the counters “sort themselves out” more so that friends are closer – this is shown in Fig. 8.




Fig. 8. The average of the average number of nearby friends for runs with different numbers of friends

Since in all these runs the racial biases run from 0.75 (each counter has 25% of other coloured friends than would the case if they were selected randomly) to 1 (no friends of the other colour), an outcome where friends are clustered means that counters of like colours will also be clustered, but to a much lesser extent. This set of outcomes is illustrated by Fig. 9.




Fig. 9. Average of the average proportion of nearby like coloured counters for runs with different numbers of friends

The Effect of Local Clustering

The local bias has the next most significant effect. Figure 10 shows the average of the average fear levels for different levels of local bias. What appears to be happening is that in the short run having a bias towards (initially) local friends means that the communication of fear is locally amplified and so increases locally, but in the long run self-reinforcement over a substantial part of the network will overtake this.

In this case the average number of movements and the average number of nearby friends is not much effected by the local bias, but the average propor-




Fig. 10. The average fear levels for runs with different local biases




Fig. 11. Average number of communications for runs with different local biases

tion of nearby counters of like colour is a bit higher for a local bias of 0.55 (Fig.12).




Fig. 12. The average of the average proportion of nearby counters of like colour for runs with different local biases

The Effect of Racial Bias (Disconnectedness of Sub-Networks)

The racial bias has the least dramatic effect in terms of outcomes.




Fig. 13. The average levels of average fear for runs with different racial biases

Figure 13 shows the average fear levels for runs with different racial biases, the greatest fear level resulting from the runs with just under a complete racial bias. The communication rates show that initially there are more communications in the runs with least racial bias but after a time the position reverses with a higher racial bias resulting in more communication.




Fig. 14. Average number of communications for runs with different racial biases




Fig. 15. The average number of nearby friends for runs with different racial biases

Figure 15 shows that the higher the racial bias the greater the average of nearby friends' results. Figure 16 shows that with a racial bias of 1 a much greater colour segregation results compared to even a slightly reduced racial bias.




Fig. 16. The average proportion of nearby same coloured counters for runs with different racial biases

These results clearly show how the structure of the social network relative can substantially effect the results both in quantitative terms as well as the kinds of dynamics that unfold in this model. Given that it is credible that a similar process could occur in corresponding social phenomena, one can conclude that, in such cases, abstracting away the social network would be unsafe.

3 Example 2 – Social Influence and Domestic Water Demand

The second example has a more descriptive flavour. It seeks to see how the quality of variation in domestic water demand in localities may be explained by mutual influence. It was developed as part of the FIRMA³ and CC:DEW⁴ projects for a more detailed description see [4]. The initial model was written by Scott Moss and then developed by Olivier Barthelemy. A fuller description of this model can be found in [8], but this does not include the comparison described below. Its role in this paper is to show that the structure of the social network is important to the results in a more detailed and descriptive model, one which has a closer relation with observation than the version of the Schelling model described above. It illustrates how the particular network structure is important as part of the model.

³ <http://firma.cfpm.org>

⁴ <http://www.sei.se/oxford/ccdew>

The core of this model is a set of agents, each representing a household, which are randomly situated on a 2D grid. Each of these households is allocated a set of water-using devices (such as toilets, washing machines etc.) in a similar distribution to those in the mid-Thames region of the UK. At the beginning of each month each household sets the frequency the appliance is used (and in some cases the volume at each use, depending on the appliance). Households are influenced as to their usage of appliances by several sources: their neighbours and particularly the neighbour most similar to themselves (for publicly observable appliances); the policy agent (either the local water company or a government representative); what they themselves did in the past; and occasionally the new kinds appliances that are available (in this case power showers, or water-saving washing machines). The individual household’s demands are summed to give the aggregate demand. Each month the ground water saturation is calculated based on weather data (which is either historical data or artificial data), if this is less than a critical amount for more than one consecutive month, this triggers the policy agent to suggest a lower usage of water. If a period of drought continues it progressively suggests using less and less water. The households are biased as to the extent that they attend to the influence of neighbours or the policy agent – the proportion of these biases are set by the simulator. The structure of the model is illustrated in Fig. 17.




Fig. 17. The structure of the water demand model

The neighbours in this model are those in the n spaces in the squares orthogonally adjacent to a household. The default value of this distance, n , was 4. The purpose of this neighbourhood shape was to produce a more complex set of neighbour relations than would be produced using a simple distance-related one as in the Schelling model (Fig. 1), but still retain the importance of the influence of neighbours.

To give an idea of the social topology that results from this neighbourhood I have shown the “most similar” neighbour influence pattern at a point in a typical run of the model in Fig. 19. Due to the fact that every neighbour has a unique neighbour who is most influential to it, the topology of this social

				N				
				N				
				N				
				N				
N	N	N	N	C	N	N	N	N
				N				
				N				
				N				
				N				

Fig. 18. The neighbourhood pattern for the households

network consists of a few pairs of mutually most influential neighbours and a tree of influence spreading out from these. The extent of the influence that is transmitted over any particular path of this network will depend upon the extent each node in the path is biased towards being influenced by neighbours.




Fig. 19. The most influential neighbour relation at a point in time during a typical run of the water demand model

Households are also (to a lesser extent) influenced by all its neighbours in its neighbourhood (the one shown in Fig. 18 above). Figure 20 illustrates all the effective neighbour relations between the households for the same instance. Note that the edges of this are not wrapped around into a torus in the examples described, so the households at the edges and corners have fewer neighbours than those in the middle. The reason for the chosen neighbourhood pattern is that the resulting patterns (as in Fig. 19 above and Fig. 1 above) seem to us a reasonable mix of locality and complexity. We have no good empirical basis for this, it just seems intuitively right to us and we could not find any evidence as to what the structure might be. I return to this issue in the discussion below.




Fig. 20. The totality of neighbour relations in the same case as Fig. 19 (each node connects to those 4 up, down, left and right along the indicated lines)

In each run the households are distributed and initialised randomly, whilst the overall distribution of the ownership and usages of appliances by the households and the biases of the households is approximately the same. In each run the same weather data is used, so droughts will occur at the same time and hence the policy agent will issue the same advice. Also in each run the new innovations (e.g. power showers) are introduced at the same date. Figure 21 shows the aggregate water demand for 10 runs of the model with the same settings (normalised so that 1973 is scaled to 100 for ease of comparison).

To illustrate the difference between the outcomes when the social network is active and when it is disrupted, I ran the simulation again with the same settings, social structures etc. except that whenever a household looks to other households, instead of perceiving that household's (public) patterns of water use, the patterns of a randomly selected household is substituted. Thus the neighbour-to-neighbour influence is randomly re-routed each time it occurs in the second case. This is designed to be the minimal possible disruption of the social network, for it does not effect the number of neighbours of each household, their cognition, or the external factors – these all remain the same. Figure 22 shows 12 runs of this version of the model, this can be compared to Fig. 21 where influence transmission is normal.

The qualitative difference between the two runs is reasonably clear. In the first set (Fig. 21) there is an almost uniform reaction to droughts (i.e. the advice of the policy maker), with almost all runs showing a reduction in water demand during these periods, whilst in the second (Fig. 22) the reaction to such periods is not a general reduction in demand but rather a period of increased volatility in demand. Secondly, the first set (Fig. 21) shows a much greater stability than the second (Fig. 22), which exhibits short term “oscillations” in demand.

What seems to be occurring is that, in the first experiment, small neighbourhoods mutually influence each other so as to adopt a set of usage patterns,




Fig. 21. The aggregate water demand for 10 runs of the model (55% Neighbour biased, historical scenario, historical innovation dates, dashed lines indicate major droughts, solid lines indicate introduction of new kinds of appliances)




Fig. 22. The aggregate water demand for 12 runs of the model where the neighbour relation is broken by randomisation

an this acts to dampen down any avalanches of influence that may otherwise occur, because such neighbourhoods are more difficult to influence from surrounding areas due to their internal self-reinforcement of usage patterns. In the disrupted case households “flip” between different patterns as the incoming perceived patterns vary, resulting in a lot of “noise” – this noise acts to drown out the relatively “faint” suggestions of the policy agent. A “mean field” style social influence model, where each household perceives the average of its neighbour’s water use patterns would be even smoother and more con-

sistent than the original model. Thus the social influence in the undisrupted model can be seen as somewhere “between” that of the randomised and the “mean-field” abstract models, a situation which allows for localised, mutually reinforcing patterns of behaviour to compete with and influence other localised patterns of water use.

Elsewhere Barthelemy [2] showed that the output of this model was qualitatively and significantly effected by different changes to the topology of the social network, including a change in the density of agent, whether the space has edges or not (i.e. whether the grid is toroidally wrapped around to itself or not); and the size of the agent neighbourhoods.

In any case it is clear that the social network does have a significant effect upon the resulting behaviour of the simulation, and that the structure of that network is important in producing the results. In other work on this model it is clear that changing the distribution of biases on the household (so that more or less are biased towards being influenced by their neighbours or the policy agent) also changes the qualitative nature of the aggregate water demand patterns that result.

4 Discussion

One of the main reasons for making more descriptive simulation models is that they suggest what to look for in the target phenomena – they inform good observation. Here a natural question that arises as a result of investigations into this model is what are the external influences upon households are – do they look to their immediate neighbours for cues to what is socially acceptable or does such influence spread mainly through local institutions such as the school, the pub or the place of work? However, as far as I can tell, not much is known about this. This points to an obvious “gap” in the field research – social networks do look at the structure of who talks to who, but does not relate this (as far as I can tell) to physical location – geographers do look at where people are located in space but do not generally investigate any social structure that is not based upon physical locality. There are some simple studies which start to touch upon this relation, (e.g.[13, 15]), but these only start to touch upon a single aspect of the relation of social and physical space corresponding to the local bias parameter in the modified Schelling Model above.

Social network theory is a long established field that studies the properties of social networks from both theoretical (e.g. [14]) and empirical approaches. However it tends to focus overmuch on the network as its abstraction of choice, largely leaving out the cognition of the agent [2]. Thus social network theory complements that of cellular automata which is an abstraction of physical action (e.g. [11]). Combining the two would lead to a richer and more complete model of many situations, however the interaction of physical and social space occurs primarily through the cognition of the agent. Thus to combine these

two spaces one needs a modelling tool that also allows the representation of this cognition, i.e. an agent-based model.

5 Conclusion

If we are to take the physical and social embeddedness of actors seriously we need to model their interactions in both of these “dimensions” – assuming these away to very abstract models will lead to different, and possibly very misleading, results. Agent-based simulation seems to be the only tool presently available that can adequately model and explore the consequences of the interaction of social and physical space. It provides the “cognitive glue” inside the agents that connects physical and social spaces. Statistical and mathematical tools are not well suited to this task which is perhaps why, up to now, models have been very abstract in nature. However, this situation has now changed – we now have an appropriate modelling tool, namely agent-based simulation. Thus, for the first time, it is no longer necessary to “simplify away” all of the real contingencies of social interaction but start to capture these in descriptive social simulations. Such models will then allow a more informed determination of when and how abstraction can safely be done – until then, we may find all sorts of interesting properties of networks and structures, but we will have no evidence as to whether they are relevant other than their intuitive appeal. It is not a question of agent-based simulation being “second best” to analytic models, for such complex phenomena it is the more appropriate, the better, tool.

Thus I am saying more than just that either: that there are aspects that are not covered in simpler models; or that an approach that starting with a model that is descriptively adequate to the available evidence is likely to be more productive than one that tries to start simply, i.e. a KIDS rather than a KISS approach as I advocate elsewhere [6]; but that models such as those above provide evidence that assuming- away the structure of the social network (that is separate from the physical topology) is unsafe. Thus the burden of proof is upon those that make this kind of simplifying assumption to show that such assumptions are, in fact, justified. So far I have not seen any such evidence, and so one must conclude that any such work is likely to be inadequate to capture the essence of many occurring socially self-organised processes.

Acknowledgements

Thanks to the many people with whom I have discussed these issues, including: Scott Moss, David Hales, Juliette Rouchier, Edmund Chattoe, Tom Downing, Nick Gotts, Olivier Barthelemy, Guiamme Deffuant and Fredrerick Amblard. Thanks also to Francesco Billari, Thomas Fent, Alexia Prskawetz, Jürgen

Scheffran and Ani Gragossian, the organisers and secretary of the Topical Workshop on Agent-Based Computational Modelling, who invited me.

References

1. Axtel, R. and Epstein, J. M. (1996) *Artificial Societies - social science from the bottom up*. MIT Press.
2. Barthélémy, O. (2003) The impact of the model structure in social simulations. 1st International Conference of the European Social Simulation Association (ESSA 2003), Gronigen, the Netherlands, September 2003. (<http://cfpm.org/cpmrep121.html>).
3. Carley, K. (2003) *Dynamic Network Theory*. In: Breiger, R., Carley, K. and Pattison, P. (eds.) *Dynamic Social Network Modelling and Analysis: Workshop Summary and Papers*. Washington: The National Academies Press, 133–145.
4. Downing, T.E, Butterfield, R.E., Edmonds, B., Knox, J.W., Moss, S., Piper, B.S. and Weatherhead, E.K. (and the CCDeW project team) (2003) *Climate Change and the Demand for Water*, Research Report, Stockholm Environment Institute Oxford Office, Oxford. (<http://www.sei.se/oxford/ccdew/>).
5. Edmonds, B. (1999) Capturing Social Embeddedness: a Constructivist Approach. *Adaptive Behaviour* 7:323–348.
6. Edmonds, B. and Moss, S. (2005) From KISS to KIDS an “anti-simplistic” modelling approach. In: P. Davidsson et al. (Eds.): *Multi Agent Based Simulation 2004*. Springer, *Lecture Notes in Artificial Intelligence*, 3415:130-144.
7. Edmonds, B. and Hales, D. (2005) Computational Simulation as Theoretical Experiment, *Journal of Mathematical Sociology* 29:209–232.
8. Edmonds, B. Barthélemy, O. and Moss, S. (2002) Domestic Water Demand and Social Influence - an agent-based modelling approach, CPM Report 02-103, MMU, 2002 (<http://cfpm.org/cpmrep103.html>).
9. Granovetter, M. (1985) Economic-Action and Social-Structure - The Problem of Embeddedness. *American Journal Of Sociology* 91:481-510.
10. Moss, S. and Edmonds, B. (2005) Sociology and Simulation: - Statistical and Qualitative Cross-Validation. *American Journal of Sociology*, 110(4) 1095-1131.
11. Polhill, J.G. , Gotts, N.M. and Law, A.N.R. (2001) Imitative Versus Non-Imitative Strategies in a Land Use Simulation. *Cybernetics and Systems* 32:285-307.
12. Schelling, T. (1969) Models of Segregation. *American Economic Review* 59:488-493.
13. Sudman, S. (1988) Experiments in Measuring Neighbor and Relative Social Networks. *Social Networks* 10:93-108.
14. Watts, D. J. (1999) *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press.
15. Wellman, B. (1996) Are personal communities local? A Dumptarian reconsideration. *Social Networks* 18:347-354.

Agent Design for Agent-Based Modelling

Jim Doran

Department of Computer Science, University of Essex, Colchester, United Kingdom
doraj@essex.ac.uk

Summary. To build an agent-based computational model of a specific socio-environmental system requires that answers be found to several important questions including: what social actors and contextual entities are to be modelled as software agents? With what mental functions must these software agents be equipped? What standard design or designs should be used to create the software agents? This paper concentrates upon the last of these questions. The currently available range of agent designs is considered, along with their limitations and inter-relationships. How best to choose a design to meet the requirements of a particular modelling task is discussed and illustrated by reference to the task of designing an informative agent-based model of a segmented, polycentric and integrated network (“SPIN”) organisation of the type analysed by Gerlach in the context of environmental activism.

1 Introduction

This paper is about how best to design, not merely to implement, an agent-based model on a computer of a real-world socio-environmental scenario. Although much has been written on this topic in recent years (e.g. [4], [5], [13, Chap.8], [1], [15], [16]), the various questions that arise have not been fully answered. Indeed, I fear that much agent-based model design continues to be arbitrary and ill conceived. Thus although some of what is to be said may seem commonplace, it is important to revisit the issues.

The start point is that given a real-world scenario (for example, a busy supermarket), and some questions about it to which we seek answers (for example, how best to reduce average evacuation time in case of fire), we must design a model capable of providing those answers:

SCENARIO + QUESTIONS → MODEL DESIGN

Major decisions to be made on the way to a suitable model design are: (a) whether to use a specific or a non-specific model, (b) the level of abstraction of the model, (c) the cognitive content of the agents, and (d) the choice of

agent architectures. Notice that all of these decisions must be made whether explicitly or not and whether arbitrarily or in a considered and informed way.

2 Specific and Generic Models

A model may refer to a specific situation, for example a particular supermarket at a particular location. Or it may refer to the essential features of a type of situation, for example the essential features of any supermarket at any location. This distinction is an important one. Specific models require much detailed observation to justify their particular and detailed structure and potentially yield specific insights and predictions. By contrast, generic models require less detailed observation and more preliminary interpretation and recognition of just what are the essentials of a type of situation. They potentially yield general insights. Specific models are naturally used to predict the likely outcomes of specific actions (including no action). Generic models are used to discover the properties that a real-world scenario has as a result of its structure and dynamics. A current and persuasive example of the use of generic models to explore social phenomena is the work of Hemelrijk and colleagues [14] who are investigating how certain types of animal social behaviour, especially those found in some ape communities, may emerge from relatively simple elements of behaviour. Their work is important not merely because it illustrates a certain type of agent-based modelling, but because it could significantly change our understanding of how the social behaviour in question arises.

3 Which Agents?

To set up an agent-based model, whether specific or generic, one must decide what are to be the agents within it. These will probably correspond in some fashion to actors in the real world. Indeed it is natural to envisage a two-staged process. First to decide what actors are to be “recognised” in the real world scenario in question, and then to decide how the actors recognised are to be mapped into agents within the model. It is clear that the decisions made must depend both on the reality of the scenario and on the objectives of the modelling study.

The problem of recognising actors in the world is far from trivial once cultural pre-conceptions are discarded¹. For example, how many actors are there in a University? Is the Department of Mathematics an actor? Are overseas students collectively an actor? Is the pleasant cedar just outside the Vice-Chancellor’s window that from time to time calms down the VC, an actor?

¹ See the analysis and discussion in [6], [7] with associated computational experiments and results.

Are the fly on the windowpane and the expert system on the VC's personal computer actors?

Given that a set of actors has somehow been systematically recognised in or "read into" the real-world scenario, according to current cultural norms or otherwise, the next question is how this set of actors is to be transformed into a set of agents within the model. A one-to-one mapping from actors to agents is likely to seem unnecessary or even impossible. Some notion is needed of what actors are relevant to the objectives of the modelling study and this involves the model as a whole.

4 What Level of Model Abstraction/Aggregation?

Deciding what are to be the agents in a model is closely connected with deciding the level of abstraction or aggregation of the model, that is, just how much detail is to be included. The following principles seem applicable to deciding the level of abstraction/aggregation:

Principle: The model must be sufficiently detailed that it can address the questions to be answered. For example, if the question posed concerns how many members of a group will support its leader in a certain eventuality, modeling the group solely as a single collective agent is inappropriate. However, too much detail is both unnecessary and likely to be computationally overwhelming. Hence:

Principle: The model should be as abstract/aggregated as possible subject to the requirement:

Each of a *completeset* of inter-variable relationships and processes in the model must

Either be reliably set from empirical observation (even if the relationship or process is non-linear)

Or be feasibly subjected to experimental variation (and in due course be so examined)

Or demonstrably have no significant impact (in which case it can probably be discarded).

Principle: Assumptions based on pre-conceptions are to be avoided.

It should be clear that meeting these requirements is rarely easy, wherein lies much of the difficulty in agent-based modelling. Perhaps the central issue is just what potential properties of the model can reliably and completely be observed in the real-world scenario. This is a practical matter. Note that generic models are somewhat less dependent on reliable observation than are specific models.

5 What Agent Cognition and what Agent Architecture?

The choice of cognition and structure to be designed into the agents in the model is an aspect of the foregoing considerations. We have to decide what cognitively oriented computations the agents are to perform in the model. This is so whether all agents perform the same computations or different agents perform different computations. Furthermore, we must decide what information (knowledge, belief) each agent is initially to possess. Again different agents may well possess different initial information.

It must further be decided what actual agent architectures are to be used. In this context an agent architecture is a structural and process design for the agent. The major differences between agent types lie in their architectural concepts rather than in implementation software. There is a significant range of recognised software agent architectures available for use in agent-based models ([17], [18]). They include architectures based upon:

1. sets of variables with associated condition-action rules (incl. “fuzzy” rules)
2. artificial neural networks
3. behaviours and subsumption
4. predictive planner and internal model
5. logic systems e.g. BDI
6. hybrid and/or multi-layer

Architectures with associated condition-action rules (1) can be very simple - no more than a couple of variables and a handful of (possibly fuzzy) rules - or can be highly complex and support advanced cognitive processes. Artificial neural networks (2) take as their building blocks simple computational abstractions of real brain neurones usually arranged in layers. These architectures are typically used to implement reactive agent behaviour with some kind of associated learning. Architectures based upon behaviours with subsumption relationships (3) between them were originally devised for robot control, where effective low-level behaviours are essential. In their original form they avoided the use of internal models or planning. Architectures that do use internal models and predictive planning (4) tend to be more “symbolic” in flavour and the more effective examples are based on the long tradition of AI work concerned with automatic planning systems. Formal logic based architectures (5), of which the best known is perhaps the BDI (“beliefs, desires intentions”) architecture, are derived from structures and processes in fragments of mathematical logic, but necessarily take a computational form in practice sometimes akin to (1) and (4) above. Finally, hybrid multi-layer architectures usually have a lower level that is reactive (an architecture of type (1), say), then a higher layer which uses modelling and planning (type (4), say) and often come with a further even higher layer with a more social function, perhaps involving communication and cooperative planning with other similar agents in a shared collective environment.

These architectures are not fully standardised. They vary not merely in their effective content, but in the degree to which they are well defined and available off the shelf. They can be mapped along at least three distinct dimensions: (a) by the type of their basic elements (e.g. rules, artificial neurones, symbols), (b) by the range of “cognitive” functions that are provided, and (c) by the definiteness with which they are specified. Most of these architectures are “open” in the sense that with due programming effort any cognitive function can be provided within them to at least some degree. In particular, all of these agent architectures can support adaptive agents.

Although their users often discuss particular types of architecture as if they were quite distinct from (and superior to) alternatives, in fact these architectures are not wholly independent. Furthermore there are some commonly encountered erroneous beliefs concerning them. For example, it is often assumed that control architectures typically deployed in robot contexts are quite different from and irrelevant to software modelling and vice versa; that “rule-based” architectures cannot involve learning and are necessarily limited in their human social significance because people are not “rule bound”; and that there is a major difference between “symbolic” and “neural network” architectures. In fact all of these beliefs are questionable if not downright mistaken. They arise from a too superficial and conventional view of the architectures in question. Further points that merit reiteration are that usually anything can be programmed in any agent architecture or software system (see next Sect.), and that “logic-based” software rarely gains any real power or reliability from its association with mathematical logic.

6 Software Platforms

Software platforms designed to support agent-based modelling typically provide some specific support for agent design and implementation. Platforms such as SWARM, CORMAS, SDML and RePast (see [12], [13, Chap.8]) make it easier to design and implement models and agents of certain types but do not necessarily guide the user to the right choices. Also, there are many software multi-agent support platforms that, although not designed with social modelling in mind, could easily be used for that purpose.

7 Choosing the Agent Architectures

Given previous choices of agents and of their required cognitive and information storage capabilities, how can we best to choose an agent architecture to meet the requirements of a particular modelling task? First we must consider why we do not design an agent architecture (or architectures) suitable for each particular model. The answer is that it is no easy matter to design and implement such an architecture. To attempt to do so is likely either to result

in a trivial architecture, or to set up a major and unnecessary design task. To ignore architectures already to hand is surely a mistake unless they can be demonstrated to be unsuitable.

It is sometimes argued that in a model which is to address social phenomena, since the agents must behave socially, we should design specifically social agent architectures, no doubt with a major emphasis on support for inter-agent communication. This argument is often supported by an explicit or implicit claim that the agent architectures listed earlier are not social. But the fact is that, in software, any distinctive “social cognition” that an agent may have will inevitably be built upon its non-social cognition, so that the semi-standard architectures remain relevant.

To sum up, we should aim to choose from amongst the semi-standard architectures, possibly embedded in a multi-agent software platform, by reference to the principles stated earlier. This is, of course, easier said than done!

8 An Example: Gerlach’s SPIN Organisations

By a SPIN, Gerlach [10] means a social movement that is a segmentary, polycentric and integrated network. He illustrates and discusses SPINs by reference to the environmental movement in the USA over the last four decades [11], and its component groups such as *Friends of the Earth* and the *Earth Liberation Front*, and, by contrast, to the *Wise Use* movement that sought to counter environmental activism.

Gerlach (*loc cit*, p. 289) focuses attention on the segmentary nature of SPINs. They are “composed of many diverse groups, which grow and die, divide and fuse, proliferate and contract”. They are also polycentric. They have “multiple, often temporary, and sometimes competing leaders or centers of influences”. And finally they are networked, “forming a loose, reticulate, integrated network with multiple linkages through travelers, overlapping membership, joint activities, common reading matter, and shared ideals and opponents”.

Gerlach argues that SPINs are adaptive and offers seven reasons why this is so (Gerlach, *loc cit*, pp. 302-6):

1. “... prevents effective suppression by the authorities and the opposition.”
2. “Factionalism and schism aid the penetration of the movement into a variety of social niches.”
3. “Multiplicity of groups permits division of labor and adaptation to circumstances.”
4. “... contributes to system reliability.”
5. “Competition between groups leads to escalation of effort.”
6. “facilitates trial-and-error learning through selective disavowal and emulation.”
7. “...promotes striving, innovation, and entrepreneurial experimentation in generating and implementing social change.”

9 Using Agent-Based Modeling to Verify SPIN Adaptive Functionality

To use agent-based modelling to test these claims for the adaptive functionality of SPINs it is inappropriate to model a specific instance of a SPIN (see discussion earlier). Rather we should create a computational version of a typical SPIN. This is made easier because Gerlach's discussion, although directed to the environmental movement and counter-movement in the USA, is largely pitched in general terms. Therefore to proceed we must define SPIN and traditional non-SPIN organisations, and then a range of strategies suitable to be used to attack them. Following Gerlach, we would then expect to demonstrate experimentally that the SPIN organisations fare better in most circumstances².

But what does testing Gerlach's claims for SPIN functionality require by way of agent structure and process? Gerlach's discussion implies that the SPIN actors deploy a comprehensive range of high-level cognitive functions. The issue is therefore whether we need to incorporate some or all of them in the agents of a SPIN model. Consider, however, a network's response to a "hostile" process that successively deletes actors, network members, at random - (see (1) above). It is difficult to see how the effect on the network of a deletion of a single actor in the network can be reliably predicted other than by empirical observation and subsequent generalisation. To try reliably to incorporate agent decision-making and cognitive processing within the model so that the "right" network responses will emerge is to aim at a level of complexity that is computationally and observationally overwhelming, ultimately leading, one can predict, to model and agent structures that are essentially arbitrary and unverifiable. This is so even if (implausibly) rationality is assumed in the model as a way of achieving predictability.

The conclusion is therefore a somewhat negative one: it is that no significant internal structure can safely be associated with the agents in the model. We reach this conclusion for this particular application because the actors in the real-world scenario (the members of the environmental movement) are not performing systematically observable role driven actions, or systematically observable routine learning, that can be reliably replicated within a software agent architecture. As discussed earlier, without reliable observation at a particular level of abstraction and aggregation, a reliable model cannot be grounded at that level.

10 SPINs and Terrorist Networks

Gerlach's first reason why SPINs are adaptive is that their structure "prevents effective suppression by the authorities and the opposition." His discussion of

² If our framework of consideration were more mathematical, we might reasonably hope to prove this conjecture as a theorem.

this point foreshadows ongoing work by Carley and her colleagues [2] that uses agent-based modelling to address the effectiveness of strategies for disrupting networks, in particular for destabilizing terrorist networks. Carley and colleagues, unlike Gerlach, are concerned with fully covert organisations, but this does not prevent there being close similarities in their work. Carley’s investigation has already provided interesting and important preliminary results notably that different strategies are required to destabilize cellular and distributed organisations from those that are effective on hierarchical organisations. Furthermore, different destabilisation strategies impact differently on different measures of the performance of the target organisation.

11 Conclusion

There exist a range of semi-standard software agent designs. But these are not always well understood by those creating agent-based models, and are rarely explicitly used by them. It is not clear how an agent architecture should be chosen on a particular occasion. At best a few broad principles can be stated, as we have done earlier. This lack of guidelines is surely important. Bonabeau’s [1] recent comments that “ABM is a mindset more than a technology” (p. 7280) and that model building remains “an art more than a science” (p.7287) confirm the difficulty of our situation.

Is it a limitation that we cannot create software agents to build into our agent-based models that display human levels of intelligence and consciousness³? It is tempting to answer “not at all”, on the grounds that models can and should be much simpler than that which they model. Yet human society is surely the emergent product of human intelligence and consciousness. It would be surprising if these fundamental human characteristics could be entirely ignored in models of human society.

References

1. Bonabeau, E. Agent-based modeling: methods and techniques for simulating human systems. PNAS, May 14th 2002 Vol. 99 suppl 3, 7280-7287.
2. Carley, K. M., Reminga, J., and Kamneva, N. (2003) Destabilizing Terrorist Networks NAACSOS conference proceedings, Pittsburgh, PA.
3. Doran, J. E.(1989) Distributed Artificial Intelligence and the Modelling of Socio-Cultural Systems. In: Intelligent Systems in a Human Context (eds. L A Murray and J. T. E. Richardson) Oxford Science Publications.
4. Doran, J. E. From Computer Simulation to Artificial Societies. Transactions SCS, 14(2), 69-77, June 1997 [Special Issue: Multi-Agent Systems and Simulation].

³ Consideration of consciousness is timely because there is much ongoing work. See, for example, the ESF Exploratory Workshop on “Models of Consciousness” held in Birmingham, UK, Sept 1-3rd 2003.

5. Doran, J. E. (2001) Can Agent-Based Modelling REALLY be Useful? In: Cooperative Agents: Applications in the Social Sciences (Eds. N. J Saam and B. Schmidt) Kluwer Academic Publishers pp 57-81.
6. Doran, J.E. Agents and MAS in STaMs. In: Foundations and Applications of Multi-Agent Systems: UKMAS Workshop 1996-2000, Selected Papers (Eds. M d'Inverno, M Luck, M Fisher, C Preist), Springer Verlag, LNCS 2403, July 2002, pp. 131-151.
7. Doran, J. E. (2002) Detecting Agents: the McCarthy and Hayes Martian Problem, Working Notes of UKMAS 2002 (The Fifth UK Workshop on Multi-Agent Systems) held at the Foresight Centre, University of Liverpool, December 18-19th.
8. Doran, J. E. and Gilbert, N. (1994) Simulating Societies: an Introduction. In: Simulating Societies (N. Gilbert and J. E. Doran eds.) UCL Press.
9. ESSA (2003) First Congress of the European Social Simulation Association, Groningen, September.
10. Gerlach, L. P. (2001) The Structure of Social Movements: Environmental Activism and its Opponents. In: Networks and Netwars: the Future of Terror, Crime and Militancy. (Eds. John Arquilla and David Ronfeldt) Santa Monica, CA: Rand.
11. Gerlach, L. P. and Hine, V. H. (1970) People, Power, Change: Movements of Social Transformation. Indianapolis: Bobbs-Merrill.
12. Gilbert, N. and Bankes, S. Platforms and methods for agent-based modeling PNAS, May 14th 2002 Vol. 99 supplement 3, 7197-7198.
13. Gilbert, N. and Troitzsch, K. (1999) Simulation for the Social Scientist UCL Press: London.
14. Hemelrijk, C. (2003) Social phenomena emerging by self-organisation in a competitive, virtual world ('DomWorld') First Congress of the European Social Simulation Association (ESSA), Groningen, September.
15. Moss, S. (2001) Messy Systems - The Target for Multi Agent Based Simulation. In: Multi-Agent Based Simulation (S. Moss and P Davidson eds.) Springer LNAI 1979.
16. Moss S. (2003) Presidential Address to the First Congress of the European Social Simulation Association (ESSA), Groningen, September.
17. Weiss G ed. (1999) Multiagent Systems: a Modern Approach to Distributed Artificial Intelligence. MIT Press.
18. Wooldridge M. (2002) An Introduction to MultiAgent Systems. John Wiley & Sons, Ltd.

List of Contributors

Belinda Aparicio Diaz

Vienna Institute of Demography
Austrian Academy of Sciences
Prinz-Eugen-Straße 8, 2nd floor
A-1040 Vienna, Austria
belinda.aparicio.diaz@oeaw.ac.at

Francesco C. Billari

IMQ and IGIER
Università Bocconi
Via Salasco, 5
I-20135 Milano, Italy
francesco.billari@unibocconi.it

Mary A. Burke

Research Department
Federal Reserve Bank of Boston
Box 55882, Boston
MA 02210, U.S.A.
mary.burke@bos.frb.org

Rosaria Conte

Institute of Cognitive Science
and Technology, CNR
Via San Martino della Battaglia, 44
I-00185 Rome, Italy
rosaria.conte@istc.cnr.it

Gennaro Di Tosto

Institute of Cognitive Science
and Technology, CNR
Via San Martino della Battaglia, 44
I-00185 Rome, Italy
gennaro.ditosto@istc.cnr.it

Jim Doran

Department of Computer Science
University of Essex
Wivenhoe Park
Colchester, C04 3SQ, UK
doraj@essex.ac.uk

Bruce Edmonds

Centre for Policy Modelling
Manchester Metropolitan University
Business School
Aytoun Building, Aytoun Street
Manchester, M1 3GH, UK
btuce@edmonds.name

Thomas Fent

ETH Zurich
Kreuplatz 5
CH-8032 Zurich, Switzerland
and
Vienna Institute of Demography
Austrian Academy of Sciences
Prinz-Eugen-Straße 8, 2nd floor
A-1040 Vienna, Austria
thomas.fent@oeaw.ac.at

Markus Franke

Information Services and
Electronic Markets
Universität Karlsruhe (TH)
Kaiserstraße 12
D-76128 Karlsruhe, Germany
markus.franke@em.uni-karlsruhe.de

Ernst Gebetsroither

ARC systems research GmbH
A-2444 Seibersdorf, Austria
ernst.gebetsroither@arcs.ac.at

Andreas Geyer-Schulz

Institute for Information
Engineering and Management
Universität Karlsruhe (TH)
Kaiserstraße 12
D-76128 Karlsruhe, Germany
andreas.geyer-schulz@em.uni-karlsruhe.de

Ute Gigler

ARC systems research GmbH
A-2444 Seibersdorf, Austria
ute.gigler@arcs.ac.at

Thomas Grebel

GREDEG-IDEFI, CNRS
250 Rue Albert Einstein
F-06560 Valbonne, France
grebel@idefi.cnrs.fr

Volker Grimm

UFZ Centre for
Environmental Research
P.O.Box 500136
D-04301 Leipzig, Germany
volker.grimm@ufz.de

Frank Heiland

Department of Economics
School of Computational Science
and Center for Demography
and Population Health
Florida State University
Tallahassee, FL 32301, U.S.A
fheiland@fsu.edu

Bettina Hoser

Information Services and
Electronic Markets
Universität Karlsruhe (TH)
D-76128 Karlsruhe, Germany
bettina.hoser@em.uni-karlsruhe.de

Alexander Kaufmann

ARC systems research GmbH
A-2444 Seibersdorf, Austria
alexander.kaufmann@arcs.ac.at

Mike Murphy

Department of Social Policy
London School of Economics
Houghton Street
London WC2A 2AE, UK
m.murphy@lse.ac.uk

Mario Paolucci

Institute of Cognitive Science
and Technology, CNR
Via San Martino della Battaglia, 44
I-00185 Rome, Italy
mario.paolucci@istc.cnr.it

Alexia Prskawetz

Vienna Institute of Demography
Austrian Academy of Sciences
Prinz-Eugen-Straße 8, 2nd floor
A-1040 Vienna, Austria
alexia.fuernkranz-prskawetz@oeaw.ac.at

Andreas Pyka

University of Augsburg
Universitätsstraße 16
D-86159 Augsburg, Germany
andreas.pyka@wiwi.uni-augsburg.de

Steven F. Railsback

Lang, Railsback &
Associates
250 California Ave.
Arcata, CA 95521, USA
LRA@northcoast.com

Andreas Resetarits

ARC systems research GmbH
A-2444 Seibersdorf, Austria
andreas.resetarits@arcs.ac.at

Jürgen Scheffran

ACDIS, University of Illinois
505 East Armory Ave.
Champaign, IL 61820, USA
scheffra@uiuc.edu

