

THE EFFECTS OF DUAL INHERITANCE MECHANISMS OF CULTURAL EVOLUTION
ON EMERGENT SOCIAL NETWORKS

by

Peter Revay
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
In Partial fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Computational Social Science

Committee:

_____	Dr. Claudio Cioffi-Revilla, Dissertation Director
_____	Dr. Andrew Crooks, Committee Member
_____	Dr. Sean Luke, Committee Member
_____	Dr. J. Daniel Rogers, Committee Member
_____	Dr. Jason Kinser, Acting Director, Department of Computational and Data Sciences
_____	Dr. Donna M. Fox, Associate Dean, Office of Student Affairs & Special Programs, College of Science
_____	Dr. Peggy Agouris, Dean, College of Science
Date: _____	Fall Semester 2017 George Mason University Fairfax, VA

The Effects of Dual Inheritance Mechanisms of Cultural Evolution on Emergent Social
Networks

A dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy at George Mason University

By

Peter Revay
Master of Science
University of Vermont, 2014
Bachelor of Science
Masaryk University, 2010

Director: Dr. Claudio Cioffi-Revilla, Professor
Department of Computational and Data Sciences

Fall Semester 2017
George Mason University
Fairfax, VA

Copyright © 2017 by Peter Revay
All Rights Reserved

Dedication

To my wife Shauna, who has shown me her profound love and unwavering support throughout all of these years.

Acknowledgments

First and foremost, I want to thank Professor Claudio Cioffi-Revilla for his guidance while writing this dissertation, advice on all things academia, as well as his generous financial support, which was made possible thanks to the Office of Naval Research and the National Science Foundation (NSF CDI grant no. IIS-1125171 and ONR Minerva grant no. N00014130054). I would also like to thank my parents, Anna and Dalibor, for their continuing support and encouragement during my entire academic career. Furthermore, I wish to thank my committee members, Professor Andrew Crooks, Professor Sean Luke and Dr. Daniel Rogers, for their helpful feedback during various stages of the dissertation process. More thanks go out to Martin Solich, Dr. Stephan Henn and the other participants of the Considering Cultural Complexity in ABMs workshop, Dr. Jessica Turnley and other participants of the 2016 Conference of the Computational Social Science Society of the Americas, Professor David Reitter along with other participants of the 2017 SBP-BRiMS conference, and Professor Bill Kennedy for useful reactions and criticisms; Professor Kevin Curtin and Dr. Alec Barker for their financial support during parts of my doctoral studies; Bob and Mary Simeone for their emotional and financial support; Karen Underwood and Dorothy Kondal for their administrative support; finally, Matt Oldham, Tom Briggs, and many other students in the CSS PhD program for advice and friendship.

Table of Contents

	Page
List of Tables	viii
List of Figures	ix
Abstract	xiv
1 Introduction	1
1.1 Motivation and Statement of the Problem	1
1.2 Background	2
1.2.1 Overview of Cultural Evolution Theories	2
1.2.2 Overview of Cultural Agent-Based Models	5
1.2.3 Overview of Evolutionary Computation in Agent-Based Models	7
1.3 Research Questions	8
1.3.1 Research Question 1	8
1.3.2 Research Question 2	9
1.3.3 Research Question 3	10
1.3.4 Research Question 4	11
1.4 Structure of the Dissertation	12
2 Survey of Evolutionary Computation Methods in Social Agent-Based Modeling	
Studies	14
2.1 Introduction	15
2.2 Overview of Evolutionary Algorithms: History, Purpose and Applications	17
2.3 A Simple Evolutionary Algorithm	20
2.4 Evolutionary Algorithms in Agent-Based Models	22
2.4.1 Evolution of Cooperation	24
2.4.2 Ecological Models	29
2.4.3 Evolution of Cultural Practice	32
2.4.4 Adaptation	33
2.4.5 Evolution of Sign Systems	35
2.4.6 Evolution of Bounded Rationality	36

2.5	Design of Evolutionary Algorithms and Practice in Agent-Based Modeling	
	Studies	37
2.5.1	Representation	37
2.5.2	Selection and Survival	40
2.5.3	Recombination	43
2.5.4	Mutation	46
2.5.5	Lamarckian Methods	47
2.5.6	Co-Evolution	48
2.6	Discussion	49
3	A Model of Co-Evolution of Signs and Cultural Traits	57
3.1	Introduction	58
3.2	Background	58
3.3	Overview	60
	3.3.1 Purpose	61
	3.3.2 Entities and Variables	61
	3.3.3 Process Overview and Scheduling	65
3.4	Design Concepts	70
3.5	Initialization and Inputs	75
3.6	Verification and Validation	76
3.7	Experimental Design	76
3.8	Results	80
3.9	Discussion	83
3.10	Conclusion	85
4	Modeling the Co-Evolution of Culture, Signs and Network Structure	93
4.1	Introduction	94
4.2	Background	95
4.3	Overview	96
	4.3.1 Purpose	97
	4.3.2 Entities and Variables	97
	4.3.3 Process Overview and Scheduling	99
4.4	Design Concepts	105
4.5	Initialization and Inputs	110
4.6	Verification and Validation	110
4.7	Experimental Design	111
4.8	Results	114
4.9	Discussion	119

5	Cultural Aspects of Political Polarization: Simulating U.S. Congress Co-Sponsorship Networks	123
5.1	Introduction	124
5.2	Background	126
5.3	Methods	128
5.3.1	Entities and Variables	128
5.3.2	Process Overview and Scheduling	130
5.3.3	Initialization and Inputs	135
5.4	Verification and Validation	136
5.5	Experimental Design	136
5.6	Results	144
5.7	Discussion	150
5.8	Conclusion	153
6	Cultures of Collaboration: Simulating Scientific Co-Authorship Networks	154
6.1	Introduction and Background	155
6.2	Methods	158
6.2.1	Entities and Variables	158
6.2.2	Process Overview and Scheduling	161
6.2.3	Initialization and Inputs	165
6.3	Verification and Validation	165
6.4	Experiment Design and Data Processing	166
6.5	Results	167
6.6	Discussion	170
6.7	Conclusion	173
7	Conclusion	174
7.1	Main Contributions	174
7.2	Discussion	174
7.3	Summary of Results	176
7.4	Limitations and Future Work	177
7.5	Conclusion	178
	Appendix A: Model Version Overview	179
	Appendix B: Mathematical Representation of the Basic Agent-Based Model	180
	Appendix C: List of Abbreviations	182
	References	183

List of Tables

Table	Page
2.1 Comparison of the implementation of evolutionary mechanism and concepts in different ABMs	54
2.2 (Continued) Comparison of the implementation of evolutionary mechanism and concepts in different ABMs	55
3.1 Agent variables	62
3.2 Model Parameters	64
3.3 Model parameter descriptions	64
3.4 Model parameter values	77
3.5 Network parameters	77
4.1 Agent variables	98
4.2 Model Parameters	99
4.3 Model parameter values	111
5.1 Agent variables	129
5.2 Model Parameters	130
5.3 Model parameter values	136
5.4 Co-sponsorship network statistics in the U.S. House of Representatives . . .	144
6.1 Agent variables	160
6.2 Model Parameters	161
6.3 Model parameter values	168
A.1 Overview of differences in implementation of selected model concepts in individual model versions	179

List of Figures

Figure		Page
2.1	Evolutionary algorithm cycle	22
2.2	An example of a binary phenotype representation (bottom) encoding for a genotype consisting of a list of real-valued parameters (top).	38
2.3	A general schema of survivor and parent selection.	40
2.4	One-point crossover recombination	44
2.5	Uniform crossover recombination	45
3.1	An outline of the logic of agent activity	66
3.2	Simulation statistics for complete networks as a function of the horizontal/vertical transmission ratio (a–e). Each point in parts a–e represents the average of the sum of the values for the given measure at the end each of the 100 generations. Part f shows the average share of the most common tag as a function of time when $H/V = 20$	86
3.3	Simulation statistics for random networks as a function of the horizontal/vertical transmission ratio and average degree for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)–(c) show tag skewness. Parts (d)–(f) show trait skewness. Parts (g)–(i) show fitness skewness. Parts (j)–(l) show percentage of rejected interactions.	87
3.4	Simulation statistics for random networks as a function of the horizontal/vertical transmission ratio and average degree for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)–(c) show tag entropy. Parts (d)–(f) show tag modularity. Parts (g)–(i) show trait modularity. . . .	88

3.5	Simulation statistics for small-world networks as a function of the edge rewiring probability and average degree for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show trait skewness. Parts (d)-(f) show tag skewness. Parts (g)-(i) show fitness skewness. $H/V = 10$	89
3.6	Simulation statistics for small-world networks as a function of the edge rewiring probability and average degree for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show tag entropy. Parts (d)-(f) show trait modularity. Parts (g)-(i) show tag modularity. $H/V = 10$	90
3.7	Simulation statistics for small-world networks as a function of the edge rewiring probability and transmission ratio for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show trait skewness. Parts (d)-(f) show tag skewness. Parts (g)-(i) show fitness skewness. $\hat{k} = 100$	91
3.8	Simulation statistics for small-world networks as a function of the edge rewiring probability and transmission for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show tag entropy. Parts (d)-(f) show trait modularity. Parts (g)-(i) show tag modularity. $\hat{k} = 100$	92
4.1	An outline of the logic of agent activity	101
4.2	Example of a base-level activation for some input over time. The spikes in the chart coincide with instances of processing the input, followed by gradual decay.	108
4.3	Model statistics for $\langle k \rangle = 8, c = 0.16$ and $H/V = 10$	116

4.4	Visualizations of the agent network in the Lamarckian bias (left) and unbiased (right) configurations after 5, 20 and 100 generations. The border color of a node determines its external marker, while the inner color determines its cultural trait variant. $\langle k \rangle = 8, c = 0.16, H/V = 10$	117
4.5	Model statistics as a function of maintenance cost and interaction/transmission ratio (H/V) for the unbiased (left), genetic bias (centre), and Lamarckian bias (right). Parts (a)-(c) show the average local clustering coefficients at the end of runs compared to regular lattices. Parts (d)-(f) show average path lengths at the end of runs compared to regular lattices. Parts (g)-(i) show the sum of tag modularity values at the end of each generation. Parts (j)-(l) show the sum of trait modularity values at the end of each generation. $\hat{k} = 32$.	118
4.6	Model statistics as a function of maintenance cost and average degree of the seed network for the unbiased (left), genetic bias (centre), and Lamarckian bias (right). Parts (a)-(c) show the average local clustering coefficients at the end of runs compared to regular lattices. Parts (d)-(f) show average path lengths at the end of runs compared to regular lattices. Parts (g)-(i) show the sum of tag modularity values at the end of each generation. Parts (j)-(l) show the sum of trait modularity values at the end of each generation. $H/V = 10$.	120
5.1	Illustration of the function of the tolerance attribute (T). In the above example the top case results in a failed interaction, while the bottom case results in a success.	131
5.2	An example of how the ideological position of legislators could be biased by party affiliation.	131
5.3	Expected relative payoffs on a random network as a function of tolerance and ideological positions	138
5.4	Examples of the the social force function F for different values of x	140
5.5	Cosponsorship networks in the 95th Congress. Part (a) shows the empirical networks. Remaining parts show networks simulated by different models under the fixed payoff scheme: (b) party-based ($RewiringLambda = 20, cost = 0.5$), (c) social ($w = -0.5, cost = 0.5$), (d) individual ($cost = 0.5$), (e) ideological ($RewiringLambda = 20, cost = 0.5$). In the simulated networks the different shades of red and blue reflect the agents' positions on the ideological spectrum.	142

5.6	Cosponsorship networks in the 108th Congress. Part (a) shows the empirical networks. Remaining parts show networks simulated by different models under the fixed payoff scheme: (b) party-based ($RewiringLambda = 20$, $cost = 0.5$), (c) social ($w = -0.5$, $cost = 0.5$), (d) individual ($cost = 0.5$), (e) ideological ($RewiringLambda = 20$, $cost = 0.5$). In the simulated networks the different shades of red and blue reflect the agents' positions on the ideological spectrum.	143
5.7	Violin plots of the distribution of measured network statistics for the four models across the entire search space.	145
5.8	Time series of the measured network statistics for the four models at a single point in the search space ($RewiringLambda = 20$, $MaintenanceCost = 0.5$, $SocialWeight = -0.5$, uniform payoff schedule).	147
5.9	Gradient fields of the measured network statistics for the party-based model under the uniform payoff scheme as a function of α and t	148
5.10	Gradient fields of the measured network statistics for the party-based model under the fixed payoff scheme as a function of α and t	149
5.11	Gradient fields of average tolerance and average pairwise distance for the party-based model under the uniform payoff scheme as a function of α and t	150
5.12	Gradient fields of average tolerance and average pairwise distance for the party-based model under the fixed payoff scheme as a function of α and t	151
5.13	Violin plots of ideological trait assortativity distributions across the entire search space over time in the party-based model.	152
6.1	Example of a base-level activation for some input over time. The spikes in the chart coincide with instances of processing the input, followed by gradual decay.	163
6.2	Examples of the the social force function F for different values of x	164
6.3	Observed statistics for the empirical co-authorship networks in given years.	169
6.4	Gradient fields of the RMSE in terms of (a) average degree, (b) clustering coefficient and (c) average path length relative to the Economics data for the biased (right) and the unbiased (left bar) model configurations as a function of social weight and maintenance cost.	169
6.5	Gradient fields of the RMSE in terms of (a) average degree, (b) clustering coefficient and (c) average path length relative to the AI data for the biased (right) and the unbiased (left bar) model configurations as a function of social weight and maintenance cost.	170

6.6	Time series of the (a) average degree, (b) clustering coefficient and (c) average path length for the Economics biased model runs (blue) and the Economics data (red).	171
6.7	Time series of the (a) average degree, (b) clustering coefficient and (c) average path length for the AI biased model runs (blue) and the AI data (red).	171

Abstract

THE EFFECTS OF DUAL INHERITANCE MECHANISMS OF CULTURAL EVOLUTION ON EMERGENT SOCIAL NETWORKS

Peter Revay, PhD

George Mason University, 2017

Dissertation Director: Dr. Claudio Cioffi-Revilla

Understanding cultural dynamics in human societies is the first step towards solving many complex social issues. In this dissertation I focus on the drivers of diffusion and adoption of cultural traits, such as values, beliefs, and behaviors. I adopt an evolutionary view of cultural dynamics. Particularly, I use concepts from dual inheritance theories of cultural evolution to develop and test an agent-based model capable of simulating the changing distributions of cultural traits in a large population of actors over the course of prolonged periods of time. Particularly, I pay close attention to the mechanisms of indirectly biased transmission of traits and guided variation, which are both hypothesized to be significant drivers of cultural dynamics. Indirectly biased transmission consists of the adoption of specific trait variants on the basis of possession of initially unrelated external markers. Guided variation is then individual adaptation driven by self-exploration.

I use various methods to explore the pathways of cultural evolution. Among them are agent-based modeling, evolutionary computation, complex network analysis and statistical data analysis. Furthermore, I make use of large publicly available datasets to validate my models. The first one of these is the database of bill co-sponsorship in the U.S. House of Representatives from 1973 to 2008. The other is a comprehensive dataset of scientific

co-authorship in various disciplines stretching back for over a century. The results show that cultural evolution models based on indirectly biased transmission and guided variation are suitable to explaining the dynamics of various complex social networks. Furthermore, I show that this type of cultural evolution leads to emergence of meaningful cultural signs by gradually associating previously independent external markers with specific cultural trait variants. Finally, I describe how the proposed model leads to plausible social network configurations.

Chapter 1: Introduction

In this chapter I first state the problem at hand and provide motivation for its study. I then provide some background on concepts that play an important role in addressing the problem. In the sections that follow, I give brief overviews of the individual research questions that I attempt to answer in my dissertation. I finish the chapter with comments on the structure of this dissertation.

1.1 Motivation and Statement of the Problem

Cultural practices structure human conduct in contexts ranging from modes of governance and natural resource management to conflict resolution, power relationships and religious practices. Due to their wide-ranging effects and observed persistence, such practices can serve as either seemingly insurmountable roadblocks or powerful catalysts for collaboration in different areas of life, such as economics, politics or international relations. If we could properly understand the life-cycle of cultural practices and the underlying mechanisms, we would be better equipped to solve complex issues in these areas of interest. For example, global climate change or terrorism are both phenomena, which are decidedly affected by culture (Tohme, 1992; Kluch and Vaux, 2016). Many scientific models of culture have been developed precisely for these purposes. However, existing models either lack explanatory power and remain static and descriptive in their nature, or their dynamics are fairly simple and schematic. Furthermore many models lack sufficient formalization and their replication and validation proves rather difficult. The goal of my doctoral dissertation is to contribute to the study of culture and its evolution by developing a formal, reproducible and general agent-based model of cultural evolution grounded in established theory, and extending the validity of the theory by comparing the model to empirical evidence of cultural evolution in

qualitatively different domains of human culture.

In particular, I am interested in how actors' imperfect knowledge of cultural attributes and the resulting reliance on unrelated external markers as signs of culture affects population-level cultural trajectories from an evolutionary perspective. This mechanism for cultural evolution has been proposed in dual inheritance theory (Boyd & Richerson, 1985), but to the best of my knowledge it has not been tested with the use of an agent-based model.

The work is positioned at the intersection of *cultural* and *social* realms of human action. Here, by the *social* realm I mean anything that pertains to relations between actors. As Weber (1978, p.4) puts it, any action “is social insofar as its subjective meaning takes account of the behavior of other and is thereby oriented in its course”. The *cultural* realm then refers to what Geertz (1973, p.89) describes as “a system of inherited conceptions expressed in symbolic forms by means of which men [*sic*] communicate, perpetuate, and develop their knowledge about and attitudes toward life”.

In the chapters that follow, I draw on theoretical foundations of cultural evolution to develop an agent-based model of culture. As the focus here will be on evolution and its trajectories, I implement an evolutionary algorithm within the agent-based model. For now, I therefore introduce the various themes that pervade this work.

1.2 Background

1.2.1 Overview of Cultural Evolution Theories

Many social theorists have proposed models of cultural dynamics over the years. Berger and Luckmann (1967) have developed a compelling model of social construction of reality and the role of culture-maintenance organizations. Bourdieu (1984) has attributed the stratification of society into classes based on the possession of different types of cultural and symbolic capital. Marxist theorists such as Gramsci (2000) or Williams (1978) have put forth comprehensive models on the spread and life-cycle of ideologies. However, these sociological theories suffer from insufficient formalization, and their empirical validation

often proves to be difficult.

The cultural debate has not been confined exclusively to sociology. In modern anthropology neo-evolutionary theories have discarded several core concepts of classical social evolutionism; doing away with notions of social progress and determinism. Neoevolutionism embraced stochasticity and individual decision-making among other driving forces of cultural emergence. Nevertheless, it remained firmly grounded in the Darwinian evolutionary framework (e.g. Sahlins, 1960). Similarly, in the biological sciences, evolutionary explanations of culture emerged under the heading of sociobiology. Wilson (1975) pioneered this approach, taking a strictly genetic view of cultural evolution. These ideas were rejected by many intellectual opponents, claiming them to be biologically deterministic, ethnocentric and a slippery slope leading to social darwinism and eugenics (Sahlins, 1976; Gould, 1981; Lewontin et al., 1984). Memetics, first introduced by Dawkins (1976), further extended sociobiological theories by positing that cultural units (or “memes”) are self-replicating entities akin to biological viruses.

Dual inheritance theories respond to this strain of research and hypothesize that culture is transmitted in human populations both vertically—from generation to generation—by forces that can be modeled with mechanisms similar to biological evolution, as well as horizontally—within generations—by social influence. Boyd and Richerson (1985, 2005) use genetic evolution as a partial metaphor for the evolution of culture, although they point out major differences. Most importantly, genetic evolution drivers such as selection for reproduction, gene recombination and mutation act on the genotype of individuals, while only indirectly affecting the phenotype. In Boyd & Richerson’s (1985) model of cultural evolution the phenotype is also affected directly by social forces of cultural transmission. Particularly important in this model are the concepts of biased transmission and guided variation. Biased transmission is based on the notion that evolved preferences for specific phenotypic traits can drive the selection of certain genetic (or in this case cultural) traits. Guided variation then serves the purpose of transmitting cultural traits within a single generation. The process of guided variation depends on an adaptive standard that determines

which cultural variants are learned.

Although dual inheritance theories have emerged in direct response and even opposition to earlier evolutionary perspectives of culture, they have often suffered from criticism on the same grounds as the former frameworks (Fracchia and Lewontin, 1999; Atran, 2001; Gabora, 2011). However, Henrich et al. (2008) claim that most of the critiques are rooted in a fundamental misunderstanding of the dual inheritance paradigm of cultural evolution. Unlike sociobiologists and their opponents who either categorically accept or reject the analogy between cultural and genetic evolution, dual inheritance theorists think in quantitative terms and recognize the analogy between culture and genes as loose, even if highly useful for the modeling of various empirical processes (ibid). Still others have deemed dual inheritance approaches as the most promising of the evolutionary frameworks (Laland and Brown, 2002); citing overly mathematical language, as well as the lack of empirical research on the topic as main obstacles for wider adoption.

In this work I will be mostly using the dual inheritance definition of culture due to Boyd and Richerson. If different definitions are discussed elsewhere, they will be stated explicitly.

Definition 1.1 (Boyd & Richerson, 2005, p.6). *Culture is information capable of affecting individuals' behavior that they acquire from other members of their species by teaching, imitation, and other forms of social transmission.*

Dual inheritance theory is an established framework that has been formalized in the form of mathematical models, and has been since applied to a number of specific cases of cultural evolution (e.g. Henrich & Boyd, 2001; Henrich et al., 2012). In this work I provide a stronger argument for the validity of dual inheritance theory by accounting for the complex networks of interactions present in human societies and elucidate the micro-to-macro link by representing and simulating the mechanisms of dual inheritance theory in an agent-based model.

1.2.2 Overview of Cultural Agent-Based Models

Within the field of computational social science, researchers have also contributed towards explanations of cultural evolution. One of the most influential ABMs of culture is due to Axelrod (1997). In this model the probability of interaction between agents located in a rectangular grid depended on their cultural closeness, and when they interacted they copied distinct traits from each other. Axelrod observed local convergence, but global polarization. Moreover, he showed that the number of distinct cultural regions was positively correlated to the number of cultural features, but negatively correlated with the number of possible traits. A number of agent-based modelers have since revised and extended Axelrod’s original model. Klemm et al. (2003) have added small random perturbations to the agents’ cultural genome and found that such “cultural noise” actually leads to less diversity in the long run. Parisi et al. (2003) have replaced the homophily mechanism with a social influence mechanism, where agents set their cultural traits based on the average trait in their neighborhood. Centola et al. (2007) have extended the homophily mechanism towards evolving networks. In their model, agents whose similarity falls below a threshold value sever their link and instead select another neighbor randomly. This led to the discovery of distinct intervals of the threshold value which resulted either in complete individuality, multiple cultural clusters, or global monocultures. Finally, Flache and Macy (2011) tested a version including both social influence and network homophily. This resulted in highly robust and persistent cultural diversity.

However, these models assume agents possess perfect knowledge of their neighbors’ cultural “genome” (in the sense of dual inheritance theory). This is a strong assumption and goes against the principle of bounded rationality (Simon, 1996). Other agent-based models of culture have attempted to circumvent this shortcoming. Bednar and Page (2007) and Bednar et al. (2010) created and analyzed models in which agents play several games with different payoff matrices. Each game represents a different cultural feature. Agents possess limited cognitive capacities and their rule sets for both games are forced to overlap to a certain degree. The authors found that this leads to the spreading of cultural memes through the

transfer of dominant strategies from one cultural feature to another. A different approach to represent agents' indirect inference of others' cultural traits is the use of observable "tags" (e.g. physical features, markings, clothes, etc.). Hales (2000) developed a model in which agents played one-shot prisoner dilemma games only with agents possessing the same tag. Each agent was either an absolute cooperator or defector. Over the course of many generations this led to homogenization of the tag-groups into clusters consisting solely of either cooperators or defectors. Hammond and Axelrod (2006) also made use of tags in their agent-based model of ethnocentrism. The simulation results implied that the emergence of cooperation was contingent upon the evolution of in-group favoritism. External markers also appear in Janssen's (2005) work on cooperation in reputation systems. However, it is the case in these studies, that the tags are "ready-made" signs, in that agents either recognize them as indicators of group membership (Hales, 2000; Hammond & Axelrod, 2006) or are able to learn a pre-existing relationship between the tag and another trait (Janssen, 2005).

Many models approach the evolution of culture *qua* evolution of cooperation. Often, the mechanics of cooperative (or selfish) behavior are modeled with the Prisoner's Dilemma game (Axelrod, 1986; Miller, 1996; Macy & Skvoretz, 1998; Bowles & Gintis, 2004), although others have used the ultimatum game (Shutters, 2009) or the stag hunt and snowdrift games as well (Santos et al., 2006). Other models have moved away from issues of direct reciprocity and instead consider indirect reciprocity among groups of actors. For example, Takahashi (2000) explores conditions under which generalized exchange is globally sustainable. Thus, instead of mutual cooperation, agents offer help unilaterally in hopes that a third party will return the favor at some point in the future. In a similar fashion, the question of delayed reciprocity was tackled by de Vos et al. (2001). The authors assume that agents periodically face natural hazards, and they must ask others for help, who in turn must hope that their favor will be returned by someone else in the future. The model tests several altruistic as well as selfish strategies and their sensitivity to harshness of environmental conditions.

Finally, some agent-based modeling studies have considered the effect of network structure on the diffusion of cultural practices such as cooperation. Santos et al. (2006) investigate the

nature of cultural diffusion when agents have the option to update both their game-playing strategies as well as their neighborhoods. Chiang (2013) studies the relationship between dynamics of cooperation on networks and nodal attributes, and finds cooperative strategies are robust when conditional upon node centrality.

1.2.3 Overview of Evolutionary Computation in Agent-Based Models

The use of evolutionary computation methods in agent-based models has a long tradition. Evolutionary algorithms (EAs) were originally designed as problem solvers. Given a space of potential solutions, a large number of search points, or “individuals”, is randomly initialized. By means of evaluating individuals in their environment, a measure of quality can be assigned to them. A selection process is then used to choose a subset of high-quality individuals. Descendants of these individuals are then generated by randomized processes representing recombination and mutation. The processes of evaluation, selection and reproduction then continue until a sufficient solution is found (Back, 2000). However, EAs need not be used only for optimization purposes. They have also been used to provide robust on-line adaptation, machine intelligence, as well as to capture the essence of evolutionary processes through simulation (Fogel, 2000). In fact, Holland (1975) has designed the genetic algorithm (GA) to be a simulation of evolution, not to solve problems (De Jong, 1995). While it is true that problem-solving applications have been at the forefront of evolutionary computation, there is ample room for more work to be done in modeling evolutionary systems (De Jong, 2005).

Several efforts have used evolutionary algorithms in agent-based model studies of cooperation (Macy & Skvoretz, 1998; Santos et al., 2006; Hammond & Axelrod, 2006; Hales, 2000). Other ABMs employing EAs have focused on ecological issues such as strategies for resource acquisition (Saam & Harrer, 1999), crop cultivation (Perez-Losada & Fort, 2011), or efficiency of resource extraction (Lake & Crema, 2012). The various evolutionary ABM studies also differ in the types of populations they study, from human actors, to mental models (Edmonds, 1997; Dosi et al., 1999). Considerable heterogeneity is also present in the specific implementations of individual mechanisms that make up the EAs. While some

ABMs only use vertical transmission (Bowles & Gintis, 2004; de Vos et al., 2001; Kohler et al., 2012) others include both vertical and horizontal transmission (Powell et al., 2009; Premo & Hublin, 2009; Lake & Crema, 2012). Similarly, some models approach reproduction from generational perspectives (Klos, 1999; Lake, 2001; Pepper & Smuts, 2000), but others implement a survival mechanism (Kohler et al., 2012; Perez-Losada & Fort, 2011; Premo & Hublin, 2009). Further differences in approaches are demonstrated in other aspects, such as choices of fitness functions, recombination, or mutation methods.

Ultimately, most of these models suffer from serious shortcomings. Issues with reproducibility, insufficient formalization and narrow application pervade the existing literature (see Chapter 2 for details). One of the goals of my work is to formalize the framework of cultural evolution in the form of an agent-based model, its implementation in computer code, and its detailed, comprehensive and disambiguous description. Moreover, the goal for the model is to be as general and robust as possible, with a wide range of applications.

1.3 Research Questions

1.3.1 Research Question 1

How do the cultural mechanisms of indirectly biased transmission and guided variation affect the distribution of cultural practices from an evolutionary perspective?

Boyd and Richerson (1985, 2005) have extensively developed dual-inheritance models of cultural evolution. Their models are couched in various contexts and rest on a whole range of hypothesized mechanisms. Among the most important mechanisms are those of *indirectly biased transmission* and *guided variation*.

Definition 1.2 (Boyd & Richerson, 1985, p.182). *Transmission is indirectly biased if naive individuals prefer some models over others based on an observable trait and use such preferences to determine the attractiveness of that model for other, hidden traits.*

Definition 1.3 (Boyd & Richerson, 1985, pp. 6, 82). *Guided variation is learning in the form of trial and error in which the trials are self-generated and explored by the individual*

itself, the results of which can be transmitted to other generations

Boyd and Richerson investigate via mathematical approaches what happens when individuals possess little information regarding the presence of the underlying trait variants in others and are therefore led to form preferences for alternative, observable markers. They posit, that this leads to the association of certain markers with specific trait variants and results in phenomena such as drift-away or runaway evolution.

While many ABMs have studied the topic of cultural evolution, most have done so from the perspective of the emergence of cooperation (Axelrod, 1997; Macy & Skvoretz, 1998; Santos et al., 2006). Such models generally assume that agents have perfect information regarding the cultural makeups of others. Certainly, other efforts have made use of *tags* (i.e. external markers) as proxies for cultural practices (Hales, 2000; Hammond & Axelrod, 2006; Janssen, 2005), but without exception the models assume a pre-existing association between the tags and the cultural trait variants. Finally, there have been studies focusing on the interplay between network structure and cultural evolution (Santos et al., 2006; Chiang, 2013). However, to the best of my knowledge, there has been no agent-based model that would simultaneously implement (a) the issue of imperfect knowledge regarding cultural behaviors, (b) the independent nature of tags and cultural trait variants, and (c) the co-evolution of network structure.

1.3.2 Research Question 2

How does the co-evolution of cultural traits, external markers and social network structure affect the cultural topology of agent populations?

As part of this objective, I expand on the previous research question by considering the social network structure in more detail. Specifically, I ask how the trajectories of cultural evolution change when social networks are assumed to be amenable to change by the actors themselves. I therefore once again begin with the same set of assumptions stemming from the dual inheritance approach of Boyd and Richerson (1985). However, in their work, considerations of the exact structure of social relations are largely absent. Elsewhere, Santos

et al. (2006) have built and studied an agent-based model of cultural evolution, giving the agents the opportunity to adjust not only their game-playing strategies but also their social ties. Specifically they tested the sensitivity of their model with respect to the ratio between game strategy adjustment and neighborhood adjustment. They have found that when social tie adjustment is much more common than strategy adjustment cooperation does not emerge. However, above a critical value of the ratio, cooperative regimes appeared. Although the study focuses on social network evolution, it uses Prisoner’s Dilemma games (as well as stag hunt and snowdrift games) to model cooperation dynamics, and does not touch on the issue of hidden traits and external markers. Another study by Efferson et al. (2008) considers both tags and unknown traits as well as in-group/out-group choices. The authors experimented on small groups of live subjects and found that when external markers align well with the possession of certain hidden behavioral traits the subjects display significant in-group favoritism.

To answer this research question I combine the approaches listed above (i.e. dual inheritance cultural theory and dynamic network models), and I extend my original cultural evolution model in order to allow agents to modify their neighborhoods throughout the course of the simulation. I then study how the resulting structure of the agent network and the distribution of cultural traits and tags within the network depends on the ways that agents adjust their social ties.

1.3.3 Research Question 3

Does the proposed co-evolutionary model faithfully reproduce the dynamics of bill co-sponsorship in the U.S. House of Representatives?

There is an increasing sense that American society is becoming politically polarized. A number of studies have attempted to validate this intuition in quantifiable terms. Abramowitz and Saunders (2007) have carried out survey-based study among U.S. voters only to confirm that measures of polarization are increasing. Fowler (2006a, 2006b) along with Zhang et al. (2008) have performed quantitative analyses of co-sponsorship networks, and report

increasing partisanship in both chambers of the U.S. Congress. Andris et al. (2015) have arrived at similar findings when considering roll-call voting patterns in both chambers. Alizadeh et al. (2016) have made a case for the emergence of extremist ideologies from a psychological perspective.

There are several hypothesized causes of political polarization, from the impacts of party system quantity and quality (Dalton, 2008), and the effect of media bias (Bernhardt, Krassa, & Polborn, 2008), to the “trickle-down” dynamics of partisanship among political elites (Baldassari & Gelman, 2008; Lachat, 2008).

Here, I posit that the increasing polarization among U.S. political elites also depends on the existing cultural structures within the larger legislative community and is constrained by the uncertainty regarding ideological positions as they pertain to political negotiations. Therefore, I modify the cultural evolution model and apply it to simulate dynamics of collaboration among legislators. To validate the model I use an extensive dataset that tracks bill co-sponsorship in the U.S. House of Representatives.

As with any sort of evolutionary process, the underlying principle here is the accumulation of small gradual changes over longer periods of time. Although this change occurs at scales different from human biological evolution, and transmission of cultural material is not as clearly delineated as genetic (sexual) transmission, the mapping of basic evolutionary principles is still possible. Ideas, such as ideological beliefs may be transmitted by means of social influence, learning, or socialization, from one legislator to another. Ideas may be combined with others and modified by chance (“mutated”). Selection of behavioral models also occurs by means of influence and reputation. For this reason I choose to approach co-sponsorship dynamics from the perspective of cultural evolution.

1.3.4 Research Question 4

Does the proposed co-evolutionary model faithfully reproduce the dynamics of co-authorship in different scientific fields?

Scientific co-authorship networks have been studied from sociological perspectives (Endersby et al. 1996; Moody, 2004) as well as using network science approaches (Barabasi et al. 2002; Newman, 2004). Barabasi et al. (2004) studied networks extracted from different datasets and reported on the scale-free structure of such co-authorship networks. Furthermore, they used the preferential attachment model to successfully reproduce these networks *in silico*. Newman (2004) has later confirmed the scale-free nature of co-authorship networks in another study. De Stefano et al. (2013) further note that such network analyses are sensitive to the choice of data as well as the specific network definitions (De Stefano et al., 2011).

To answer this question I modify my cultural evolution model to account for the nature of scientific communities, such as the growth in population size, the discovery of new traits and markers, and the social aspect of collaboration. I then validate the model against empirical data on two different scientific fields (Economics and Artificial Intelligence) from the Microsoft Academic Graph database (Sinha et al., 2015).

Once again, although the scales are different, the mapping of evolutionary concepts onto the domain of scientific collaboration is possible: I argue that scientific practices can spread, combine and mutate over time via means of influence, reputation and social selection of viable model individuals.

1.4 Structure of the Dissertation

This is a manuscript-based dissertation. Each of the chapters, excluding the Introduction and Conclusion chapters either has been submitted, or will be submitted for review in the near future. Chapter 2 has been accepted for publication in the Journal of Computational Social Science.

The structure of the dissertation can be compared to a tree: Chapter 2 represents the roots, Chapters 3 and 4 make up the trunk of the tree, while chapters 5 and 6 are the different

branches stemming from the trunk. In Chapter 2, I present an extensive survey of agent-based modeling literature that incorporates Evolutionary Computation methods. In Chapter 3, I develop and present a model of cultural evolution that assumes hidden cultural traits and the reliance of actors on external markers to ensure successful collaboration between actors. I investigate the resulting cultural landscapes stemming from these assumptions. Chapter 4 describes an extension of this model, one in which the social networks are dynamic as a result of the agents' ability to adjust their social neighborhoods. Next, I explore how the different constraints and approaches to neighborhood adjustment affect the resulting distributions of cultural traits and external markers. Chapters 5 and 6 represent application of the cultural evolution model to specific cases. Chapter 5 extends the model and applies it to networks of bill co-sponsorship in the U.S. House of Representatives in an attempt to explain political polarization and partisanship from a cultural perspective. In Chapter 6, I apply the model to simulating networks of scientific collaboration, modifying several of the evolutionary mechanisms to account for the specifics of the scientific community. Here, I use the Microsoft Academic Graph database to compare model results to empirical evidence. Finally, Chapter 7 provides a summary of the dissertation, where I revisit my research questions, show my research contributions, identify areas of future work, and provide a conclusion.

Chapter 2: Survey of Evolutionary Computation Methods in Social Agent-Based Modeling Studies

Abstract: Agent-based modeling is a well-established discipline today with a rich and vibrant research community. The field of Evolutionary Computation (EC) is also well recognized within the larger family of computational sciences. In the past decades many agent-based modeling studies of social systems have used EC methods to tackle various research questions. Despite the relative frequency of such efforts, no systematic review of the use of evolutionary computation in agent-based modeling has been put forth. Here, I review a number of prominent agent-based models of social systems that employ evolutionary algorithms as a method. I comment on some theoretical considerations, the state of current practice, and suggest some best practices for future work.

2.1 Introduction

The purpose of this chapter is to give an overview of the current state of evolutionary computation techniques in the field of agent-based modeling of social systems, and to confront the practices with the theory of evolutionary computation. The history of agent-based modeling is over four decades long today, and the concept has blossomed from experimental method to a full-fledged discipline with its own conferences, journals, and institutions. The literature on ABMs is both numerous and broad in scope. Agent-based simulation is used to explain phenomena in sociology (see Bianchi and Squazzoni, 2015 for a review of relevant studies), as well as in archaeology, geography, demographics, human biology, political science and economics (see Cioffi-Revilla, 2017; Cegielski and Rogers, 2016; Heppenstall et al. 2011; or Lake, 2014 for reviews). The field of agent-based modeling has established itself by publishing works in leading scientific journals (e.g. Macy and Skvoretz, 1998; Boyd et al., 2010; Deffuant et al., 2005 ; Powell et al., 2009), by organizing annual conferences such as the Social Simulation Conference, and by sustaining quality field-specific journals such as the Journal of Artificial Societies and Social Simulation. The very nature of agent-based simulations as powerful explanatory tools capable of modeling change and emergent phenomena made it a suitable method for tackling evolutionary questions. In fact at least a few dozen of “evolutionary ABMs” have been published in the past two decades. Most of these studies are at least loosely inspired by the Evolutionary Computation (EC) approach and employ its tools to a certain degree. EC itself is an established discipline dating back to the 1960s, situated within the broader field of Artificial Intelligence in the computer science community (Gilbert and Troitzsch, 2005, ch.10). The fundamental principle of the field lies in the harnessing of mechanisms of biological evolution as a powerful problem-solving tool. EC practitioners implement computational algorithms that mimic the processes of biological evolution to tackle various optimization problems as well as to study evolutionary dynamics. However, despite the significant histories of both fields, and the continuing intertwining and borrowing of concepts from both camps, there has been

very little explicit interaction and collaboration between agent-based modelers and EC researchers. This work aims to survey the rich ecosystem of evolutionary ABMs and to compare their approaches to established theory and practice in the domain of EC.

To carry out such a review, I must first define what I mean, when I say “evolution”. The term evolution has at least two distinct, commonly used meanings (along with a number of other, rarely used meanings). The first, colloquial meaning, refers to any kind of accumulation of change, or gradual directional change. The second, formal term, refers to biological evolution, which is defined as the change in the heritable characteristics of biological populations over successive generations (Hall & Hallgrimsson, 2008). As the first meaning is overly broad, and the second meaning too restrictive for the purposes of the study of social systems, I introduce a third definition, which I will use throughout this chapter. This definition is neither as relaxed as the colloquial use, nor as narrow as the biological definition. The definition used here is as below:

Definition 2.1 (Evolution). *Evolution is the cumulative change in intrinsic characteristics of individual entities over time, as a consequence of the suitability of those characteristics to changes external to the individuals.*

Note, that unlike in biological evolution the individuals need not be biological organisms and furthermore the traits need not necessarily be heritable (at least not in the strict genetic sense). On the other hand, this definition is unlike its colloquial counterpart in the sense that any change in individuals’ attributes must be a function of the suitability of the individuals’ characteristics to its current external environment¹. This allows us to study other crucial categories, such as cultural evolution. Thus, when I write about “Evolutionary ABMs” I mean agent-based models that simulate processes included under this definition of evolution. Furthermore, I only study agent-based models that either explicitly use Evolutionary Algorithms (EAs), or implicitly implement methods that closely resemble them, in order to study evolution as defined above.

¹Here by the *environment* I mean anything external to the individual that has the power to affect it and over which the individual has limited control.

In the section that follows I begin with an overview of the history and development of the field of EC. In section 2.3 I continue by outlining a simple prototype of an EA. Then, in section 2.4, I proceed by giving brief overviews of numerous examples of Evolutionary ABMs. I divide these into several categories, such as *evolution of cooperation*, *ecological models*, etc. Next, in section 2.5, I address each integral design component of an EA in its own subsection and contrast relevant theoretical considerations with actual implementations found in the surveyed studies. In section 2.6 I conclude by assessing the current state of evolutionary agent-based modeling literature.

2.2 Overview of Evolutionary Algorithms: History, Purpose and Applications

An EA is a stochastic, population-based metaheuristic that relies on random variation, selection, and other mechanisms with analogies in biological evolution (Back et al. 2000, p. xxv). The central idea of evolutionary algorithms is that of harnessing the power of biological evolution and translating it into an efficient computer-based optimization algorithm. EAs have found their place in agent-based modeling literature as well. For example Axelrod (1986) has used one in his study on social norms. In the model each agent possessed a pair of alleles tied to their propensity to cooperate with others. The agents' fitness was modulated in different ways based on the actions they took, whether selfish or altruistic. Only the fittest agents were then selected to seed the next generation of agents. In this way, Axelrod was able to track which alleles related to cooperation were able to survive and eventually dominate the population.

Evolutionary algorithms (EAs) are the subject of the field of Evolutionary Computation, which traces its history back to the 1960s. The notion that evolutionary processes are able to find local (and potentially also global) optima of solutions to problems defined by environmental constraints was familiar since the times of Charles Darwin. The proposition that such processes could be represented in the form of computational operations and that

their power could thus be harnessed by computers to solve various optimization problems was initially posed in the first half of the 20th century (De Jong, 2005, p. 24). It was not until the availability of inexpensive digital computers reached a sufficient threshold in the 1960s that these ideas began to be exploited in practice. Holland (1962) envisioned the role of evolutionary processes as crucial tools for the design of robust adaptive systems capable of navigating through complex dynamic environments. Rechenberg (1965) developed his idea of how evolutionary processes could be utilized to solve difficult parameter optimization problems into the family of tools now known as evolution strategies. Fogel et al. (1966) saw evolutionary computational techniques as a means of developing artificial intelligence.

From these different visions emerged the three techniques which enjoy the most interest in the EC community today: Evolutionary Strategies (ES), used mainly for real-valued parameter optimization, Genetic algorithms (GA), focusing on the genotype-phenotype mapping and the idea of genetic building-blocks, and finally Genetic Programming (GP), which is used for evolving formal symbolic statements, such as chunks of computer code, decision-making routines or mathematical equations, represented as nested tree-like structures. EC and the formal study of EAs is well established today. There are several recognized archival journals (such as *Evolutionary Computation* or *IEEE Transactions on Evolutionary Computation*), as well as a number of peer-reviewed conferences (such as the Genetic and Evolutionary Computation Conference or the Congress of Evolutionary Computation), and numerous high-quality monographs (e.g. Back, 1995; De Jong, 2005; Fogel, 1995; Holland, 1992; Koza, 1992; Mitchell, 1998).

Although EAs are used for a wide range of purposes, De Jong (2005, p.27) notes that perhaps the most natural application is as an optimizer, which is why significant effort was put into applying these techniques towards various types of optimization problems. From a systems analysis perspective we may identify three main components of a working system: its inputs, outputs, and its internal model which transforms the former into the latter. Armed with this interpretation we may define an *optimization* problem as one where the model of the system is known along with a series of outputs and it is left to identify the

inputs which best reproduce the outputs (Eiben & Smith, 2007, p. 9). Among some of the concrete applications of optimization problems one can count route planning (such as in the Traveling Salesman Problem) and scheduling, control (of robotic movement, etc.), design (e.g. aircraft parts) and classification (e.g. game-playing systems or image processing).

However, De Jong (2005, p.28) also cites other uses of EAs, such as models of cognition, the design of neural networks, but perhaps most importantly the potential of such techniques to provide significant insights into complex adaptive systems (CAS). He claims that even simple computational models can provide better understanding in contexts in which humans are known to make notoriously poor predictions (p. 231). In fact, there is a desire among parts of the EC community to capture the essence of evolution in a computer simulation and use the simulations to gain new insight into the physics of natural evolutionary processes (Bäck, Fogel, & Michalewicz, 2000, p. 2). From a systems point of view, this type of problem falls under the *modeling and system identification* heading. Eiben and Smith (2007) define system identification problems as those where both the inputs and outputs are known and a suitable model that translates between the two is to be found. A related type of problem is that of *simulation*; in this case the inputs and the model is known, and the goal is to identify the associated outputs. These kinds of problems are usually encountered in models of biological evolution, economic models and in models of social and cultural systems. In these cases evolutionary computing allows the researcher to design and perform experiments that fundamentally differ from the methods of classical biology, economics, or social science and offers possibilities that go far beyond the capabilities of studies based on such traditional approaches. There are, however, several caveats to this approach. One of these is the trade-off between biological (social, economic, etc.) fidelity and computational effectiveness. These pose conflicting design objectives in evolutionary algorithms (De Jong, 2005). Tightly related to this is the issue of how conclusions drawn in the computer simulations can be transferred to the biological (social, economic, etc.) medium (Eiben & Smith, 2007, p. 8). Finally, Eiben and Smith (2007, *ibid*) also note the lack of mutual awareness between computer scientists and other disciplines interested in modeling evolutionary processes as

the reason why such experimental studies are still scarce.

2.3 A Simple Evolutionary Algorithm

Although the particulars of different implementations of EAs vary, each consists of at least a subset of canonical components. These are:

- A **population** of individuals and their representation (see subsection 2.5.1)
- **Fitness** evaluation (2.5.2)
- **Parent selection** mechanism (2.5.2)
- **Survivor selection** mechanism (2.5.2)
- **Recombination** operator (2.5.3)
- **Mutation** operator (2.5.4)

A detailed discussion of each of these components along with an overview of its use by evolutionary agent-based modelers in practice will be given in section 2.5. For now, I will present only a brief outline. Each EA is first initialized with a **population** of random individuals (also “solutions” or “search points”). Most often these individuals are simply sets of parameters which pertain to the problem at hand. The set of all possible combinations of parameters is usually referred to as the **search space**. There are many options regarding the representation of the parameters in the computer. The most straight-forward way is the so-called **phenotypic approach**, where the **representation space** is identical with the search space. An example of this is when the parameters are all real-valued and they are also represented as real values in the computer program. This need not necessarily be the case, as is illustrated by the **genotypic approach** in which the search space is different from the representation space and a mapping between the two is necessary. The most common example of this is when integer- or real-valued parameters are represented as bit-strings. Although this intuitively seems as an unnecessary practice, it has certain advantages (as

well as shortcomings, which will be discussed later). Moreover, the genetic approach has its analogy in biological evolution.

Each of these individuals is then evaluated for its **fitness**. Here EAs can be differentiated into those with an **objective fitness function**, and those where the function is only **implicit**. In the first case the performance of each individual is solely the result of its position in the search space and can be objectively measured independent of the performance of other individuals. The latter case is usually found in co-evolutionary systems, that is, systems where performance of an individual can depend on the state of the remaining individuals in the population.

Once the individuals are evaluated a **selection** mechanism is used to identify a subset of the fittest individuals (usually called the “**parents**”) who will reproduce and create additional offspring. This is done via the recombination operator, which dictates how the genetic material of two or more parents is combined in a deterministic or stochastic manner to create a new individual. The mutation operator is then responsible for making small, stochastic changes to the offspring’s genetic material.

Once the offspring have been created the **survival** mechanism determines which individuals proceed to the next time step. The simplest format is the so-called generational or age-based EA in which all the individuals from the previous generation are removed from the population and only the offspring are allowed to proceed. The complementary approach to this are steady-state or non-overlapping systems in which only a part of the previous generation is removed and replaced by offspring at any given step. The survival mechanism in this case can be a function of either age or fitness. In most EAs the selection, reproduction and survival phases are implemented in a way that keep the population size constant, although formats with variable population sizes are possible, and certainly more faithful analogues of biological and social systems.

This completes a single cycle of the EA. A simple illustration of the EA scheme is given in figure 2.1. The individual stages are then repeated many times until some stopping criterion is met. Although there is a number of nuances and possible additions to this scheme this

gives an introductory sketch of a basic EA design.




Figure 2.1: Evolutionary algorithm cycle

2.4 Evolutionary Algorithms in Agent-Based Models

Agent-based simulation is a powerful methodological tool because it is (1) inherently dynamic, and thus capable of providing a lucid account of change; (2) able to represent heterogeneity, (3) path-dependent and (4) focused on non-teleological, bottom-up explanations (Epstein & Axtell, 1996). These qualities make it uniquely positioned to answer questions regarding evolutionary processes. For this reason, the history of evolutionary ABMs is almost as long as that of agent-based modeling itself. Only a decade after Schelling's (1971) ground-breaking

study on patterns of segregation, Axelrod and Hamilton (1981) analyzed the dynamics of iterative Prisoners' Dilemma (PD) tournaments from an evolutionary perspective. Although their work did not explicitly employ EAs to study what the authors called the “evolutionary robustness and viability” of different strategies, it was perhaps the sign of an increasing interest of tackling the problem of evolution in artificial agent societies.

While evolutionary ABMs have been published sporadically throughout the 1980s and the early 1990s, the frequency of such research efforts has increased greatly since the late 1990s and throughout the 2000s. The scholars who contribute to the emerging niche come from a variety of institutions across North America, Europe, and Asia, as well as from a wide range of disciplines. ABMs utilizing evolutionary algorithms have been designed to answer questions in sociology (Santos et al., 2006), economics (Takahashi, 2000), anthropology (Pepper & Smuts, 2000), archaeology (Kohler et al., 2012), ecology (Perez-Losada & Fort, 2011), etc., and many of them touch on aspects of psychology, cognitive sciences and neuroscience as well. As discussed in section 2.2 most EA applications approach evolution as an *optimization* process. This is not the case in evolutionary ABMs. In fact, all of the ABMs surveyed in the following sections study evolution from the *simulation* and *modeling* perspective, and this will be reflected in the analysis of the employed practices.

In the following parts of this chapter I present what constitutes a representative cross-section with respect to discipline, research topic and methodology. The surveyed studies are summarized in Table 2.1. I begin by outlining the research categories under which evolutionary ABMs most commonly fall. I identified six common themes in such models, which may overlap to a certain degree:

1. Evolution of cooperation
2. Ecological models
3. Evolution of cultural practice
4. Evolution of bounded rationality
5. Evolution of sign systems

6. Adaptation

Finally, there are many ABMs that either include the term “evolution” or “evolutionary” in their descriptions or seem to be evolutionary at first glance, but do not exactly fit the proposed definition. Examples of this include Axelrod’s (1996) study of cultural evolution or the Kohler et al.’s (2000) original Village simulation. The former was not included because there is no selective pressure or concept of fitness, while the latter, despite modeling death and reproduction, does not allow for any change in the inherent characteristics of its agents.

2.4.1 Evolution of Cooperation

At first glance the existence of cooperative or altruistic behaviors, which lead the individual to temporarily defer its own well-being for the well-being of others, poses a logical paradox. Why would any individual act contrary to its self-interest if acting out of their own free will? These questions have led philosophers to reason that individuals must be compelled to follow the laws of “covenants” (Hobbes, 2013 [1651]) or social contracts (Rousseau, 1968 [1762]) by the fear of punishment from some greater force. The already tentative position of cooperation was certainly not improved by Darwin’s theory of evolution, expressed in competitive terms of survival of the fittest individuals. In the second half of the 20th century, game theory emerged as a viable approach to formalizing the multi-objective nature of social behavior. Especially the formulation of Prisoner’s Dilemma and its subsequent analysis helped elucidate the elementary dynamics of self-interest and altruism (see Axelrod, 1984 or Poundstone, 1992 for explanations of the dilemma and some theoretical considerations). For this reason it became a popular tool of researchers hoping to demonstrate conditions under which cooperation can emerge.

Axelrod (1986) used a modified n -person version of the Prisoner’s Dilemma to explore what conditions lead to the development of norms in societies. In this model agents take turns during which they have an opportunity to cooperate or defect. Defections give a boost to the agent, but they hurt the other players. There is also a possibility of each defection being seen by other agents. If agents see a defection they can either choose to let it go, or they can

punish the defector. Such a punishment costs the defector as well as the agent who executes the punishment. Axelrod then allowed for heterogeneity in strategies by giving each agent two parameters which he termed “boldness” and “vengefulness”. While boldness controlled the agent’s willingness to defect given a certain probability of being seen, vengefulness determined its probability of punishing a defection of other agents. In simulations of this model Axelrod failed to observe any sort of cooperative norm. This prompted him to introduce so-called “metanorms” in which agents could punish not only defectors but also those who see a defection and do not punish it. Under these new conditions Axelrod was able to observe the emergence of a cooperative norm. It should be noted however, that later efforts to replicate the model failed to reproduce Axelrod’s results (Galan & Izquierdo, 2005). This was mostly ascribed to the small number of runs, the small population sizes and the arbitrary choice of parameters.

Macy and Skvoretz (1998) created an evolutionary ABM to study how cooperation emerges in a society where repeated interaction is uncertain. Once again, they achieve this by having the agents play the Prisoner’s Dilemma. The agents are modeled as either unconditional cooperators or unconditional defectors. Furthermore, they can signal certain external markers and possess rules on how to evaluate others (whether based on markers, fitness, projection of ego etc.). These rules can be based on the possession of certain markers or agents’ fitness. They position these agents into neighborhoods. Agents have a higher probability of interacting with neighbors rather than strangers. The agents also have the option to exit the game (i.e. refuse to interact, or distrust the partner). Under these assumptions, the authors allow the agents to evolve and track the model’s sensitivity to the cost of exit, the neighborhood size, and the propensity to favor neighbors over strangers. They show that cooperation can emerge (first locally in a single vanguard neighborhood and then globally by spreading to other neighborhoods) when exiting is fairly cheap and neighborhoods are small, and there is a large share of within-neighborhood interaction. The authors note that this is perhaps how Protestantism could have succeeded in the U.S.: churches were small and tightly knit and Protestantism thus became a marker of trust

spreading to other communities.

Klos (1999) extended an earlier evolutionary ABM by Miller (1998). In the original model agents play one of 16 different iterated Prisoner’s Dilemma strategies which replicate, recombine and mutate over the course of generations. In the extended version the agents are placed on a toroidal grid and they compete against their neighbors. This changes the selection mechanism, which becomes local: agents only imitate strategies of their neighbor. Similarly as in Miller (1998), a cooperative regime emerges. However, Klos introduces a second tweak by which fitness becomes subjective: the agents only know the standing of those neighbors who they have played in the current generation. With this modification implemented, cooperation ceases to emerge in the model.

Another adjustment of the classic PD design was explored by Takahashi (2000), who explores conditions under which generalized exchange (i.e. indirect reciprocity) is globally sustainable. Thus, instead of mutual cooperation, agents offer help unilaterally in hopes that a third party will return the favor at some point in the future. In the model, agents employ a particular strategy according to which they give another randomly chosen agent resources, while holding the receiving individual to a certain standard of giving to others. If the standard is not met, the giver ceases to donate to that agent, and chooses another random individual who meets the standard. The agents are heterogeneous in that they have different standards (thresholds) of giving and different shares of resources that they give out. Subsequently, the assumption of perfect information possessed by the agents is relaxed by placing them on a grid, where the agents only know how charitable their Moore neighbors are. Even under this assumption generalized exchange emerges. It should be noted that the central assumption in the form of the general strategy employed requires a certain pre-developed sense of fairness on the part of the agents.

The question of the evolution of direct delayed reciprocity was tackled by de Vos et al. (2001). In their model agents can become distressed with an exogenous probability (representing “harshness of conditions”) each round. If agents become distressed they will seek help from others; if they do not receive it in time they die. However, if an agent

helps another one, this results in an increased risk (controlled by a cost parameter) of the helper becoming distraught in the next round. Agents possess one of three strategies. First, defectors ask different agents for help every time, while never reciprocating it. Second, committal agents only ask for help those they have helped in the past; they always answer to requests for help, as long as those requesting it have helped the agent at some point in the past or have not received help from the agent before. The third strategy is a book-balancing one in which agents ask for help those who have received help from them and not reciprocated yet, and provide help only to those who are not indebted to them. The agents remember all interactions and their outcomes, but do not know anything about other interactions (i.e. they do not exchange information). The authors then simulate the evolution of mixed strategy populations under different conditions of harshness, population size, and cost of helping. They conclude that commitment strategies are able to invade small populations in harsh conditions as soon as two committal mutants appear. The results showed that the commitment strategies fared better than strict book-balancing agents, even when costs of helping are high.

Hales (2000) explored another dimension of the evolution of cooperation. Unsatisfied with kin selection, group selection or reciprocity as explanations for emergence of cooperation, he set out to explore how biased interaction affects the onset of wide-spread cooperation in societies. In the model agents, who are either unconditional defectors or cooperators, possess one from a wide range of tags. Agents then play one-shot PD games in pairs. Agents pick an agent at random and if they have the same tag they play the game. In the other case they keep picking agents until they find one with the same tag. Hales shows that over many generations sustained cooperation emerges when the range of possible tags is high enough, i.e. the possibilities far outnumber the agent population.

Hammond and Axelrod (2006) used this tag-based approach to test the hypothesis that social dilemmas such as the PD game promote in-group favoritism. In their model agents possess three traits: an observable tag, an in-group PD strategy, and an out-group PD strategy (either unconditional defection or cooperation). They then play one-shot PD games

with Von Neumann neighbors on a toroidal grid and reproduce locally. The authors show that this leads to the evolution of ethnocentric strategies, i.e. in-group favoritism where agents cooperate within their group but defect outside of their groups. External markers also appear in Janssen’s (2005) work on cooperation in reputation systems. The motivation of the research was to explain how reputation systems with voluntary feedback (such as eBay user ratings) effectively work. In the ABM agents play PD games repeatedly with an option to exit. Agents also possess reputation, which is modified based on feedback provided stochastically by their partners in the PD. Agents leave the simulation if their reputation or payoffs become too low and are replaced by new agents possessing random strategies. In an extension of the model, agents also signal a number of different symbols, which can eventually become recognized by others through a learning process as signs standing for specific strategies. This allows the agents to be strategic about playing and exiting the PD. The results showed that in the absence of signaling symbols cooperation emerges only above a certain threshold of propensity for giving negative feedback. With signaling this threshold disappears and cooperation emerges in all cases.

Several articles have also focused on the co-evolution of cooperation and network structure. Santos et al. (2006) gave agents in their model the option to update their game-playing strategies as well as their neighborhood. The sensitivity of the model was then tested with respect to the ratio between the frequencies of strategy and network updates and the overall connectivity of the network. The authors showed that cooperation emerges as the network update becomes more frequent relative to strategy update. Moreover, the threshold ratio necessary for cooperation to evolve increases with the network’s connectivity, i.e. highly connected networks require more agile adaptability.

Chiang (2013) also focused on the interplay of cooperation and network structure. The question is how is cooperation able to spread, if it is assumed to be conditional on certain nodal attributes (such as degree or betweenness centrality). Agents’ strategies were defined as ranges of nodal attributes, and they only cooperated with agents whose attribute values fall within that range. A wide variety of populations with mixed strategies were simulated

and allowed to evolve. Chiang observed that under such assumptions the most robust strategies are those that cooperate with either distinctly lower valued or higher valued nodes.

Evolutionary ABMs inspired by anthropological motivations have also explored the issue of cooperation. Bowles and Gintis (2004), for example, simulated a hunter-gatherer society divided into distinct groups with migration between them. Members of the groups have an option to contribute to resource acquisition efforts at a certain cost. The acquired resources are then distributed evenly regardless of contribution. However, free-riders can be punished as in Axelrod's norm model. Assuming agents possess imperfect information about the system, the authors show that cooperative mutants are able to survive and replicate even in populations initially composed entirely of opportunistic free-riders.

Other anthropological ABMs also touch on the evolution of cooperation, such as Pepper and Smuts' (2000) research on alarm signals and resource sharing and their dependence on kin selection, or Kohler et al.'s (2012) model of hierarchical structure formation in public goods games.

2.4.2 Ecological Models

The models listed in the previous subsection focused mostly on the evolution of individual behaviors and their dependence on the nature of social interactions between the individuals. However, societies do not evolve in a vacuum separated from the natural world. The surrounding environment constrains the actions of individual agents and affects their results as well. Thus the relationship between the population and its milieu shapes the co-evolution of the ecosystem as a whole. A number of articles on evolution in ABMs reflect this reality and employ an ecological approach to modeling.

One of the first such models was introduced and analyzed in a paper by Conte and Castelfranchi (1998). They were interested in the evolutionary dynamics of normative strategies for resource gathering. Agents move around a toroidal grid where food is distributed randomly. This food can be marked as their property and eventually eaten to increase their strength. Agents possessed one of several strategies which dictate conditions under which

they attack others for food. The authors then analyzed the resulting strength inequality within the population as a function of the initial make-up of the populations in terms of strategies. Saam and Harrer (1999) subsequently added an evolutionary mechanism to the model whereby the agents were able reproduce by pooling their strength between offspring.

Pepper and Smuts (2000) explored how ecological variation contributes to the emergence of within-group cooperation using a similar simulation environment in which agents expend energy to locate food. The authors modeled food restraint and alarm calling as two qualitatively different dimensions of cooperation. Alarm calling was useful in notifying nearby agents of predators, thus reducing their chance of being killed while simultaneously increasing the caller's chances of being preyed upon. Both forms of cooperation emerged from mixed populations in patchy environments, which forced agents to concentrate in small cohesive groups, as relocation proved costly. However the two cooperative behaviors differed in their dependence on kin selection. Unlike food restraint, alarm calling was only able to spread when newly-born offspring were placed near their parents.

Another use of the same artificial ecosystem was motivated by the question of the onset of cultural learning in early hominids and its dependence on the variance of resource availability (Lake, 2001). In Lake's model agents possessed different propensities towards individual and cultural learning. When agents learned individually they recalled (possibly out-dated) information regarding resource availability only at cells that they have personally visited. When agents learned culturally they were able to communicate their experiences with others. However, there was an inherent trade-off as better knowledge by the group leads to more over-exploitation of the slowly regenerating resources. Lake's simulations show an evolutionary advantage of cultural learning in environments where resource patchiness is low.

Lake and Crema (2012) revisited this topic and studied the link between adaptation and resource exploitation in more depth. The authors simulated a population of agents mining resources from a common pool. The agents were equipped with different cultural traits, which determine how much of the resource they can extract. Crucially, the extraction payoffs

showed some degree of variance independent of the traits. Agents reproduced selectively based on their payoffs, inherited traits from parents and innovated these traits by imitating other successful agents. The simulations showed that higher rates of innovation are only beneficial when payoff variance is low, otherwise the resource pool is over-exploited and adaptation dwindles. Moreover, the authors were able to identify a “sweet spot” such that population-wide adaptation is faster when the innovation rate is neither too low, nor too high.

Kohler et al. (2012) studied how a public goods game in a landscape with variable resources contributes to the evolution of social hierarchies. They specifically focused on modeling pre-hispanic Pueblo societies in the U.S. Southwest. In the simulation agents made a choice between contributing to the public good and defecting, while receiving a fair share either way. Once again free-riding can be punished at a cost. Kohler et al. then introduced two types of agents—hierarchical and non-hierarchical—who occupied the same landscape. Hierarchical agents willingly elected leaders who tax the followers and who were entirely responsible for punishing free-riders while bearing the costs. The results showed that over time there is a large increase in the ratio of agents living in hierarchical communities. These communities evolved to be significantly larger than those of the non-hierarchical type, which is in line with the archaeological record.

Migratory phenomena are another example of ecological factors in evolutionary dynamics. Premo and Hublin (2009) built an ABM to test the hypothesis that culturally mediated migration was crucial to the loss of genetic diversity in hominid populations. Individuals in the model were stochastically forced to migrate, however they would only move to cells that were culturally similar enough. Each individual then possessed a series of genes and cultural traits which could be either selective or selectively neutral. The simulations showed that the extent of genetic diversity loss is sensitive to increasing the cultural similarity threshold. The authors also found that larger innovation rates and proportions of culturally neutral traits result in a lower similarity threshold necessary to trigger loss of diversity.

The effect of population density and its structure in space was tackled by Powell et al.

(2009) when they agentized Henrich’s (2004) mathematical model of cumulative cultural evolution during the Upper Pleistocene transition. Agents were defined by their skill level in an abstract cultural domain and placed in sub-populations, which were linked by stochastic migration events. The skill value was modified through unbiased vertical transmission from parents to offspring and by biased horizontal transmission which acts by copying the best individual in the agent’s current sub-population. In both cases transmission was imperfect and “noisy”. As in Henrich’s (2004) original model, the authors concluded that for any given level of transmission noise there exists a critical population size above which cultural innovations start spreading, and that smaller communities can only innovate through migration from and into denser sub-populations.

2.4.3 Evolution of Cultural Practice

Cultural practices have a strong grasp on the behaviors of individual humans, even entire societies, yet the origins of some of them remain largely unknown. Consider for example religion and ritual, which have been the subject of fierce sociological and anthropological debates for over a century now. On one end, under the Durkheimian interpretation, religion is considered a crucial adaptation of societies which fosters a feeling of solidarity among its members by resolving cognitive dissonance, satisfying emotional needs, or infusing action with meaning (Kertzer, 1989). Other explanations view rituals as an unintended consequence of adaptations serving entirely different functions, such as regulating the impact of environmental conditions (Rappaport, 1967). Gould and Lewontin (1979) compared the latter phenomenon to the work of mosaic artists who fill in the empty spaces in the spandrels below the arches of a cathedral’s dome, itself a design aspect necessary from the perspective of structural stability. In fact, Gould and Lewontin dismissed *post hoc* explanations of specific traits as meaningful evolutionary adaptation for being “just-so stories”; offering enticing, but ultimately untestable narratives. However, the explanatory and revealing nature of agent-based simulations proves, yet again, especially fitting for the purpose of testing hypotheses regarding the emergence of cultural phenomena, be they adaptive or

serendipitous in nature.

Kachel et al. (2010) have used an evolutionary ABM for this reason, when testing the “grandmother hypothesis” (G. C. Williams, 1957). The hypothesis states that the lengthening of the post-menopausal period and thus the increasing longevity is the result of the inclusive fitness benefits of grandmothering. Each agent was defined by its longevity, fertility and reproductive age length, while the environment is defined by the exogenous parameters of maximum weaning age of offspring and the age of maturation. The critical assumption was that mothers who are weaning cannot have offspring. Two scenarios were tested: one in which grandmothers took over child-rearing duties, which allowed mothers to become fertile again earlier, and another in which grandmothering increased the chances of the offspring’s survival to maturity. Results showed that grandmothering behavior does not lead to an increase in longevity in either of the scenarios. However, in scenarios where grandmothering behavior had a direct effect on decreasing the weaning age, selection did favor shorter reproductive periods in women.

Perez-Losada and Fort (2011) set out to explain the loss of cultural diversity in farming during the neolithic in Central Europe through simulation methods. They modeled a landscape which is occupied by farming households who must share finite resources and take part in stochastic migration, as determined by the archaeological record. Each farm was defined by its strategy to cultivate any combination of eight possible crops. The households’ cropping practices are inherited by descendant households. From the simulated scenarios the authors observe that loss of diversity occurs in the most recently occupied cells and propagates “backwards”. The authors concluded that this is because cropping practices on the moving front of the migration have a better chance to catch on as a result of a lesser degree of competition.

2.4.4 Adaptation

Other evolutionary modeling efforts focus on explaining adaptation of cognitive agents to environmental constraints strictly during their lifetimes. This view of evolutionary processes

in which individual phenotypes are plastic and amenable to change after birth is known as Baldwinian or Lamarckian evolution. While in the Baldwinian approach changes in the phenotype are not retained in subsequent generations, the Lamarckian position holds that phenotypical adaptations are reflected in the genotype and passed onto offspring. While these paradigms are contentious, and particularly Lamarckian inheritance has been largely disproved as a mechanism of biological evolution, both have their place in general evolutionary theory. In fact, the existence of cultural transmission suggests that if cultural genotypes are constituted of socially constructed representations such as language, rather than genes, Lamarckian inheritance may now dominate genetic inheritance in human societies (Moravec, 1987). Even if the role of Lamarckian processes in nature and society is disputed, nothing precludes researchers from using evolutionary algorithms to represent such mechanisms. On the contrary, computer simulations devised in this way might even speak to the validity of such evolutionary models. For example, even though the Klos (1999) study uses a genetic algorithm complete with inheritance of strategies, its intention was never to model evolution over the course of many generations. Rather, the approach was to understand the algorithm as a model for the adaptation of individuals through abandoning and acquiring different social strategies over the course of their lifetime.

A good example of an evolutionary simulation model employing Lamarckian adaptation is Xue's et al. (2011) effort to properly identify conditions under which short-term fitness-enhancing innovations are advantageous in the long term. The authors argued that the long-standing intuition that the ability to predict consequences of one's own actions is beneficial to fitness has never been thoroughly tested. To carry out such a test, they designed an agent-based model in which individuals possessed a single phenotypical trait representing their adaptation to the environment, i.e. a measure of how close they are to the optimal fitness value. Then at each time-step innovations became available and agents had to decide whether to adopt them or not. Each agent had a probability of determining correctly whether the innovation will be beneficial or not. The environment was fully determined by the optimal fitness as a function of time. The function possessed a strong central tendency

throughout the first phase of the simulation, after which the value precipitated dramatically. Agents that were closer to the optimum had a higher chance of survival, and their offspring inherited the innovations adapted throughout their lifetimes. Several different types of environments were tested, distinguished by the rate at which the fitness value precipitates during reversal events. The results showed that the value of innovation significantly decreases in scenarios with faster environmental changes.

2.4.5 Evolution of Sign Systems

A crucial part of human cultural dynamics is the evolution of meaning: the forming of more or less arbitrary yet salient relationships between mental representations and physical phenomena. Complex sign systems such as languages convey references to physical objects or abstract concepts and their relations to one another through various means, such as the modulation of sound waves in the vocal chords, or the imprinting of various geometrical shapes on suitable materials. However, the assignment of specific signs to their representations is arbitrary. It is only a matter of convention that members of the feline family are in English referred to as cats and canines are referred to as dogs. Had it been the other way around, it would pose no impediment to the function of language whatsoever. The formal study of meaning in sign systems dates back to early 20th century, with the emergence of semiotics and the efforts of Saussure (1916) to deconstruct meaning into its constituent parts. Since then significant effort was put into the analytical study of sign systems, while considerations of mechanisms responsible for the emergence of meaning were given less attention. Agent-based modeling certainly offers a way to test hypotheses regarding the evolution of sign systems.

A handful of the models described in previous subsections make use of external markers (or tags) which can be observed by other agents. Through repeated interactions and their subsequent evaluation or classification, agents come to associate these markers with certain patterns of behavior (should any be present). Thus, the markers can stand in as signs for a behavior that cannot be *a priori* derived by any other means. In Hales' (2000) model of

conditional cooperation agents possessed a high degree of homophily, interacting only with those individuals who shared the same tag. The simulations showed that the proliferation of cooperation within tag groups was only possible when the discriminatory ability of the tags was minute enough, that is, when the tag alphabet was large enough to represent a wide range of small communities. Similarly Hammond and Axelrod (2006) showed that when agents are able to employ different strategies when encountering agents who possess different markers, the evolutionary dynamics of kin selection lead to the association of cooperation with sameness and defection with difference. Finally, Janssen (2005) showed that agent populations equipped with simple neural networks coupled with fitness-based selection were able to learn existing correlations between external markers and cooperative behaviors.

2.4.6 Evolution of Bounded Rationality

A special niche of evolutionary modeling is concerned with the evolution of cognitive models of reality in the minds of individual actors. While classical economists assumed that agents always act rationally to maximize their utility, modern approaches attempt to relax these often unrealistic assumptions. Simon (1996) pioneered the concept of procedural, or bounded rationality which posits that individuals form estimates of unknown utility functions, and that these estimates are periodically updated to reflect feedback to most recent actions and (imperfect) observations of a changing environment. Tree-based genotype representations and genetic programming algorithms are particularly well-suited for simulating the selection of such boundedly-rational decision-making processes.

Edmonds (1997) used this approach to simulate cognitive frameworks of agents. In his case a mental model is composed of a subset of possible elementary operations, variables, constants, and the relationships between them which lead to different outcomes based on the inputs. Agents would continuously evaluate the performance of their current models, create new ones by recombining older instances, and select the best ones to guide their next decision. Edmonds tested his agents in two contexts. The first was the management of a simple investment portfolio composed of two goods under an unknown utility function.

Agents had to estimate the function using their own models. By employing different utility functions, Edmonds was able to explore the limits of such a boundedly-rational cognitive structure. As a subsequent test Edmonds applied the same architecture to the El Farol crowding problem (Arthur, 1994). In this version the bar patrons continuously evolved a population of strategies for signaling their intentions and deciding on their true intention of whether or not to go to the bar (they only wanted to go, if it was not too crowded). The agents' models effectively co-evolved because outputs of one model can be used as inputs for another. This interaction ultimately lead to the emergence of heterogeneous strategies which the original El Farol study lacked.

Dosi et al. (1999) also argued that decision-making must follow some form of procedural rationality by pointing to the existence of a class of problems that cannot be algorithmically solved in optimal fashion. Just as Edmonds (1997), they used genetic programming to test the viability of their hypothesis. Once again, agents possessed a set of strategies defined by recursive tree-like structures which help them learn an unknown supply-demand curve. Agents had access to variables such as prices, quantities and market shares from a pre-determined number of previous steps. The authors showed that in monopoly cases a single agent is able to learn the optimal strategy fairly quickly. In cases with several agents interesting behaviors such as tit-for-tat appeared. However, the authors note that in cases of more complicated strategies their semantic interpretation becomes difficult. This points to the question of the realism of such a model and the degree of isomorphism to human cognitive architectures.

2.5 Design of Evolutionary Algorithms and Practice in Agent-Based Modeling Studies

2.5.1 Representation

Usually the first important choice a modeler must make when designing an evolutionary algorithm is the choice of how to represent the individuals, or the the candidate solutions, that




Figure 2.2: An example of a binary phenotype representation (bottom) encoding for a genotype consisting of a list of real-valued parameters (top).

make up the population of the EA. As most problems tackled by EAs are simply parameter estimation problems, this question is reduced to the representation of the parameter values themselves. Depending on the nature of the parameters it is perhaps natural to represent them directly as what they are: Boolean variables, integers or real values. This is what De Jong (2005) calls the *phenotypic approach*. In this case, individuals correspond directly to points in the solution space and the parameter values—the phenotypes—are the basic heritable units.

A different approach is to separate the representation space and the solution space and relate the two with some type of mapping. This approach, which De Jong (2005) dubs the *genotypic approach*, consists of encoding the parameter values using a universal code. In the case of computers it is natural to choose binary code (as shown in figure 2.2). This type of representation is essentially analogous to the concept of genetic inheritance, in that individual traits (points in the solutions space) are made of a combination of multiple genes (points in a multi-dimensional representation space). Evolution then operates on single bits, or individual genes rather than the actual phenotype. There has been considerable debate regarding the utility of binary representation. Holland (1975) and Goldberg (1989) argue that it is desirable from a theoretical point of view that genetic algorithms act on distinct

building-blocks or schemata that represent salient substructures of the larger structures being evolved. De Jong (2005) offers a different perspective, claiming that neither approach can be globally declared to be better and that each strategy has its own strengths and weaknesses. Certainly one weakness is the existence of hamming cliffs in binary representations. This refers to the fact that a single small change in the representation space (i.e. a single bit-flip) may produce a disproportionately large change in the solution space, depending on the salience of the bit being flipped. Effective strategies, such as Gray coding, have been proposed and analyzed to deal with this issue (Mattias & Whitley, 1994). However, even hamming cliffs aside, others have claimed it to be a mistake to encode real-valued parameters in binary (Eiben & Smith, 2007, p. 40), and that maximizing the number of schemata being processed is not necessarily useful, or may even be harmful (Fogel, 2000, p. 137).

Only a few of the reviewed ABMs have employed the genotypic approach. Klos (1999), based on Miller (1998), represented agents' PD strategies as finite state machines, which were in turn represented in binary. Here each bit marked the presence or absence of each component of the finite state machine. Other authors have utilized the genotypic approach without relying on binary representations. For example, Premo and Hublin (2008) represented agent's traits as a series of nominal, integer-valued loci. The nature of the trait was defined by the proportions of the different values in the loci. Kachel et al. (2010), also striving for a greater degree of biological fidelity, represented their agents' traits as a result of averaging two real-valued alleles. Each allele was inherited from a different parent. Finally, Hales (2000) used a mixed approach in which part of the phenotype, namely the external marker, was represented with a binary genotype. This multi-dimensional tag space, and the way in which agents traverse it, dramatically changes the dynamics of the model when compared to a one-dimensional tag representation (cf. Riolo, 1997).

However, most studies used the phenotypic approach, as it is perhaps the more natural one for most researchers. The domain of the parameter space has varied depending on the nature of the problem. Macy and Skvoretz (1998), as well as Perez-Losada and Fort (2011) have used binary phenotypes to represent the presence or absence of certain behaviors. More




Figure 2.3: A general schema of survivor and parent selection.

commonly, ordinal scales have been used (e.g. Lake and Crema, 2012), as well as real-valued ones (e.g. Powell et al. 2009, Chiang 2013).

2.5.2 Selection and Survival

Darwinian evolution is driven by the natural selection of genotypes in the form of the survival of the fittest. Evolutionary algorithms acknowledge this by implementing selection and survival mechanisms that determine which individuals are allowed to reproduce and create new phenotypic combinations (as illustrated in figure 2.3). This is done by first assessing the individuals' fitness. Fitness can be objective, as a measure of how well an individual is able to solve a problem within environmental constraints. It can also be implicit, that is, dependent on the behavior of the system as a whole. The ways in which fitness evaluations are conceptualized in agent-based modeling studies varies. In some studies this is done by representing payoffs from social dilemma's through Prisoner's Dilemma games (e.g. Macy and Skvoretz, 1998; Hammond and Axelrod, 2006; Chiang, 2013) or public goods games (e.g. Axelrod, 1984; Bowles and Gintis, 2004). In others agents gain fitness by collecting geographically distributed resources (e.g. Pepper and Smuts, 2000; Kohler et al., 2011;

Lake and Crema, 2012), and some combine multiple elements into a multi-objective fitness evaluation (e.g. Janssen, 2005). Once every individual is associated with a fitness value, the population can be ranked accordingly. There are several ways how to select the subset of the fittest. The simplest one is **truncation** selection, in which the n highest ranked individuals are selected. In **rank-proportional** selection each individual is assigned a probability of selection according to its rank; the probability ratio between subsequent ranks can be tuned with a parameter. In **tournament** selection a subset of individuals of size n is first selected randomly and the most fit individual from it is chosen; this process is then repeated many times. Finally, in **fitness-proportionate** selection every individual's probability of selection is directly proportional to its fitness value.

The consideration of the selection mechanism is important because it greatly affects the selective pressure exerted on the population throughout the simulation. This is in turn related to the phenomenon of genetic drift. Drift is the process through which, under neutral selection, trait variants that are by chance initially over-represented by even the slightest amount compared to others, will eventually take over the entire population. Any form of fitness-biased selection will lessen the effects of genetic drift, however it cannot eradicate it completely. Selective pressure then refers to the usual takeover time in the population. When the takeover time is low, selective pressure is high and vice versa.

There are various reasons why significant drift effects should be avoided. In optimization applications this is known as premature convergence, as it leads to suboptimal solutions. However, in simulations of evolutionary dynamics preventing drift phenomena is important as well, especially if the target system is known to show negligible effects of drift. Other simulation efforts call for the exploration of the degree to which drift affects the system, in which case it can be tuned via parameters of the selection mechanism.

Truncation selection shows the highest rates of selective pressure, leading to extremely rapid convergence, although it can be somewhat controlled by the selection of the truncation point; rank-proportional and tournament selection have a slower convergence overall, but there too the rate can be controlled by choice of their respective parameters (see De Jong,

2005, sec. 6.3.1.2). Fitness-proportionate selection is certainly not without issues either. Once again, in practical applications, it comes with high rates of premature convergence. However, notably in larger populations, selective pressure tapers off as the share of the fittest variants increases (because these variants will have more or less equal fitness). Too low of a selection pressure can also pose a problem, especially in co-evolutionary systems (such as those usually modeled by ABMs), where relative fitness is more important than absolute fitness (Chattoe, 1998). Eiben and Smith (2007, p. 62) point out that tournament selection does in fact select in terms of relative rather than absolute fitness, and is therefore invariant towards any translations of fitness functions, whether objective or implicit.

The second concept complementary to selection is the survival mechanism. The most straight-forward and most commonly used is the generational GA (e.g. Axelrod, 1986; Macy and Skvoretz, 1998; Hales, 2000). Under such a configuration, all of the parent individuals who are selected to reproduce are removed from the simulation and fully replaced by their offspring, while usually keeping population size constant. On the other hand there are so-called steady-state systems in which only a small subset of individuals are removed and replaced by new agents in every step (e.g. de Vos et al., 2001; Janssen, 2005; Kachel et al., 2010). Death in these cases can be either age-based or fitness-based. Arguably, generational methods are prone to higher drift rates, which must be offset by larger population sizes (Sarma & De Jong, 2000). Hancock (2000) also notes that steady-state systems are more adept at dealing with noisy systems.

Not all evolutionary models include both selection and survival methods. In fact, almost a half of the articles surveyed here use no proper selection mechanism, and rely solely on survival mechanisms coupled with fitness-neutral selection to control the parent pool (e.g. Perez-Losada and Fort, 2011; Powell et al., 2009; Saam and Harrer, 1999). Naturally, these models are steady-state systems. Others utilize fitness-biased selection mechanisms. For example Klos (1999) as well as Lake (2001) use tournament selection. However, a surprising number of studies use methods that can be problematic with respect to drift effects. Axelrod (1986), Takahashi (2000) and Lake and Crema (2012) all use some version of truncation

selection. Examples of the use of fitness-proportionate selection include the studies of Hales (2000), Santos et al. (2006) and Hammond and Axelrod (2006).

It is surprising that there is little discussion or explanation regarding choice of selection mechanisms in a number of the surveyed studies, considering the confounding and often dramatic role that drift can play in evolutionary simulations. The lack of fitness-biased selection in steady-state systems can be somehow offset by the non-overlapping nature of the population's generations. However, in non-overlapping configurations which use fitness-proportionate or truncation methods, undesirable selective pressure should ideally be offset either by highly disruptive recombination operators or large population size (Deb, 2000). Unfortunately, this is often not the case (e.g. Hammond and Axelrod, 2006; Perez-Losada and Fort, 2011).

Another interesting aspect of selection is the effect of modeling physical space. Optimization applications recognize this in a number of techniques such as island models, speciation, fitness-sharing, or crowding (see e.g. Eiben and Smith, 2007, sec. 9.3-9.4 for descriptions). Several of the surveyed agent-based modeling studies take geographical effects into consideration as well. For example Pepper and Smuts (2000) test the hypothesis that cooperation can emerge as a result of kin selection (the selection of behaviors which help genetically similar individuals at one's own expense) and the local clustering of trait groups. In other cases, such implicit effects were brought to light *post hoc*, as was the case of Bausch (2015) who showed that the model of Hammond and Axelrod (2006) is sensitive to the practice of placing offspring only in the neighborhoods of their parents. The insight was that the cooperation behavior did not emerge because of the tags as Hammond and Axelrod (2006) argued, but rather simply because agents interacted with their close kin most of the time.

2.5.3 Recombination

Recombination operators control how genes or traits get transmitted from parents to their offspring and how they are modified in the process. The role of recombination is the




Figure 2.4: One-point crossover recombination

exploration of the parameter space by creating new solutions from old ones. It can also help maintain diversity. Not all evolutionary algorithms utilize recombination, in which case offspring are created by (asexually) cloning their parents. However De Jong (2005) notes that cloning is yet another example of a method susceptible to drift. In sexual reproduction models several different recombination techniques are commonly used. The simplest method is one-point crossover in which the parameter vectors of the two parents are split at a randomly chosen position and the complementary parts from opposite parents are then “re-glued” together to create two offspring (as shown in figure 2.4). This can be generalized to n -point crossover in which case there are n cut points. Uniform crossover is another operator which steps through the parameters one by one, while a random trial determines which of the two parent contributes its value in the current position (see figure 2.5 for an illustration). It is important to note that there are trade-offs between the different types of crossover. The n -point version of crossover possesses positional bias, meaning that alleles that are close together have a higher chance of staying together. This has important implications for genotypic representations (Eiben and Smith 2007, p. 49). On the other hand, uniform crossover has distributional bias, meaning that most of the time the ratio of genes inherited from the two parents will be close to even (unless the random trials are explicitly biased). Other possible recombination operators include arithmetic crossover (e.g. averaging of the




Figure 2.5: Uniform crossover recombination

parents' values of real-valued parameters) or multi-parent recombination, which is easily accomplished in the computer and especially useful when simulating cultural dynamics. Genetic programming with its tree-based representations uses a specific set of crossover techniques. The most common one is sub-tree crossover, used by both Edmonds (1997) and Dosi et al. (1999), which acts by switching sub-trees of the two parents at randomly selected nodes (Eiben and Smith 2007, p. 108). Often, GP algorithms would impose some type of limit on the tree size, since this method can lead to undesirable bloating, expending precious computational power (Angeline, 2000, p. 287). Ultimately, what matters most for recombination according to De Jong (2005, p. 185) is that one should choose an internal representation that “in some sense reflects or preserves the meaningful application-dependent sub-components”.

Once again, most of the surveyed agent-based modeling studies choose the path of least resistance and model inheritance via cloning. In some cases this is motivated by the goal of modeling imitation of behaviors rather than true sexual reproduction (e.g. Axelrod, 1986; Powell et al., 2009; Xue, 2011; Lake and Crema, 2012). In many others the phenotype consists of a single value (e.g. a choice between cooperation and defection) and thus cloning is the only possible approach (e.g. de Vos et al., 2001; Bowles and Gintis, 2003; Lake, 2001; Santos and Pacheco, 2006). In some studies no justification for this choice is provided (e.g.

Takahashi, 2000; Hales, 2000; Hammond and Axelrod, 2006; Chiang, 2013).

A few of the studies utilize classic crossover techniques. Kachel et al. (2010) use one point crossover, while Klos (1999) implements two point crossover in his model, and Macy and Skvoretz (1998) use a version of uniform crossover. Yet others use more exotic diversity maintenance techniques. Janssen (2005) does not implement inheritance in his model at all. Instead removed agents are replaced by randomly created individuals. This is a specific choice related to the nature of reputation systems which he sets out to investigate in this model. In Saam and Harrer's (1999) study the focus is on the flow of resources between agents. It is therefore natural that inheritance would act in the form of pooling of the parents' resources. The GP models of Edmonds (1997) and Dosi et al. (1999) both use traditional sub-tree crossover.

2.5.4 Mutation

While recombination serves the purpose of exploration of the solution space, the role of mutation is exploitation, or local search within small regions of the solution space. Whereas recombination operators can have more or less arbitrary arity, mutation operators are usually unary, that is, they act on a single individual at a time, by introducing small stochastic changes to parts of their genotype. The goal is to define mutation in a way that a small change of an internal representation of an object results into a small change in the corresponding external object (De Jong 2005, p.87). This is mostly straight-forward in phenotypic representations, but crucial to understand when dealing with genotypic representations. For real-valued parameters the traditional approach is to introduce small Gaussian perturbations. Integer-valued representations make use of random resetting, which is especially well-suited for nominal variables, or gradual incrementing/decrementing, which is appropriate for ordinal and cardinal variables (Eiben and Smith 2007, p. 43-44). Once again, GP algorithms make use of a set of idiosyncratic mutation strategies such as random leaf growth, random sub-tree shrinkage, or the switching of two random subtrees within a single tree (Angeline 2000, p. 249-250). The chance of a mutation occurring in any given

part of an individual's genotype is usually fairly low, traditionally set at 1% or less. In fact, De Jong (2005, p. 173) notes that in GA-like designs mutation rates higher than 1% result in the system becoming very noisy, and confound other underlying evolutionary dynamics.

Past evolutionary ABMs have employed a variety of mutation techniques. Macy and Skvoretz (1998), Hales (2000) and Axelrod (1986) flip bits in the agent's genotype (without addressing hamming cliff phenomena). Models with nominal phenotypes, such as PD strategies, usually use random resetting (e.g. Bowles and Gintis, 2003; Hammond and Axelrod, 2006). Real-valued problems appropriately utilize perturbation techniques, however, the range of distributions used is broad: from truncated Gaussian (Kachel et al., 2010) and uniform perturbations (Xue, 2011) to Gumbel distributions (Powell et al., 2009). Many models, however, do not implement mutation operators at all (e.g. Saam and Harrer, 1999; Pepper and Smuts, 2000; de Vos et al., 2001; Lake, 2001; Dosi et al., 1999; Santos et al., 2006; Lake and Crema, 2012). A specific application of mutation techniques in certain ABMs is the representation of imperfect imitation. In this case mutation occurs every time imitation is attempted (see Powell et al., 2009; Xue, 2011).

2.5.5 Lamarckian Methods

As noted in section 2.4.4, Lamarckian evolution is arguably an important driver of human cultural dynamics. From a practical point of view nothing prevents a modeler to implement Lamarckian inheritance into an evolutionary computation. For this reason some of the evolutionary ABMs in fact do this to achieve their objectives. Lamarckian effects can be modeled in two flavors: either as horizontal transmission between two specific agents (i.e. imitation or learning) which is propagated into subsequent generations or as an abstract innovation force (essentially a type of in-life mutation which is inherited by offspring). In the case of horizontal transmission, the mechanisms can differ in the way in which individuals select the models that they imitate. This choice can be done in a fitness-biased fashion (by selecting from some subset of fittest individuals) or in a selectively neutral fashion (e.g. based on the possession of certain tags).

Xue (2011) makes use of an innovation operator, while others model horizontal transmission (e.g. Lake 2001, Premo and Hublin 2008, Powell et al. 2009), and in some cases both techniques are combined (see Lake and Crema, 2012). Finally, in some instances the presence of Lamarckian inheritance is implicit, mainly in cases where evolutionary processes are interpreted as imitation and the intended target system has no genetic basis (e.g. Axelrod 1986, Klos 1999).

It is important to note, that the specific choice of a Lamarckian operator also bears consequences with respect to the behavior of the model. Notably, Vaesen (2012) has shown that the conclusion's of Powell's et al. (2009) study do not hold when we relax the assumption that agents will always be able to identify the fittest agent to imitate.

2.5.6 Co-Evolution

Co-evolution is the mutual effect that two or more species exert on each other during the course of their evolution. Strictly speaking, co-evolution only applies to cases in which multiple species are involved (Thompson, 1994), however similar mechanisms, such as intraspecific competition² or evolutionary arms races can also be at play within a single species. There are several types of co-evolution and related phenomena. *Mutualism* is the process whereby two or more species reciprocally affect each other; *host-parasite* and *predator-prey* co-evolution occurs when one species is dependent on the other as a resource, while the other must adapt to evade the predator; finally *competition* (which can be both intraspecific and interspecific) occurs when individuals vie for a shared, limited resource (Thompson, 1994).

In EC co-evolutionary techniques have been used to further increase the efficiency of optimization algorithms (e.g. Hillis 1990; Potter and De Jong, 2000). In simulations of evolutionary dynamics, co-evolution may be inherent, as defined by the scope of the problem. Most of the ABMs reviewed here model some type of competition over limited resources (e.g. Axelrod, 1984; Klos, 1999; Lake, 2001; Chiang 2013). Often it is hard to define whether the

²Competition for a limited shared resource among individuals from a single species.

competition occurs between multiple species or within a single one of them, as many of the models simulate cultural evolution. For example it is difficult to say whether “cooperators” and “defectors” are just two strategies employed by the same species, or whether different utility estimation heuristics evolved independently by competing agents constitute distinct artificial “species” (e.g. Edmonds 1997; Dosi et al., 1999). Only a minority of the models reviewed here do not simulate any aspect of co-evolution, whether it is due to the absence of explicit competition over resources in a single-species model (e.g. Powell et al., 2009; Premo & Hublin, 2009; Kachel et al., 2010) or because the individuals do not interact in any way whatsoever (Xue et al. 2011).

2.6 Discussion

The above survey of agent-based modeling studies that use evolutionary computation techniques shows that there has been great interest in modeling evolutionary dynamics in a rigorous fashion. These efforts have been present for over three decades now and the frequency of such studies has only increased in recent years (of the 23 studies listed in table 2.1, 18 were published in 2000 or later and 7 of them were published in 2010 or later). The methods of evolutionary computation have appeared in ABMs of various target systems. Evolution of biological as well as cultural systems has been modeled. The topic of cultural evolution itself has been studied from the perspective of various practices such as reciprocal cooperation norms, alarm-calling behaviors, grandmothering, or crop diversity. Although these concepts are multi-faceted, they are all studied through the prism of memetic transmission, which stems from sociobiological paradigms emerging in the second half of the twentieth century as a loose analogy to genetic transmission (e.g. Dawkins, 1979; Boyd and Richerson, 1985). Although contentious, it appears that the popularity of such approaches is not fading in the foreseeable future, and thus an influx of new studies utilizing evolutionary algorithms coupled with agent-based simulation is inevitable.

I have stated that the purpose of this study is to confront the existing approaches in

social ABMs employing EC methods with established theory and best practices. The survey has revealed many quality studies that contribute significantly to the scientific knowledge regarding evolutionary processes in social systems, yet apart from a few exceptions, most of them fall short of best practices in one or more aspects. One of these aspects is reproducibility. There are two facets of reproducible research in the computational sciences: a complete and detailed description of the entities and processes present in the model, and the availability of the computer code. Only three of the studies reviewed here have made the computer code publicly available (Premo & Hublin, 2009; Kachel et al., 2010; Chiang, 2013). A few other studies arguably provide enough detail in the description of the model entities and processes, such that a replication of the computer code from scratch would be possible and unproblematic (Pepper & Smuts, 2000; Saam & Harrer, 1999). The remainder of the studies do not make code publicly available and offer only incomplete or ambiguous descriptions of the models. In some studies these ambiguities are minor, such as the absence of commentary on agent activation regimes (e.g. Macy & Skvoretz, 1998; Hales, 2000; Hammond & Axelrod, 2006), or agent state initialization (e.g. Xue et al, 2011). Nevertheless, such minute aspects can still have significant effects on the model (Alizadeh & Cioffi-Revilla, 2015). Other studies fail to fully specify even some of the central mechanisms, such as selection (e.g. Edmonds, 1997; Dosi et al., 2001; Bowles and Gintis, 2003), recombination (e.g. Edmonds, 1997; Kohler et al., 2010), mutation (e.g. Takahashi 2000; Kohler et al., 2010) or specifics of model geography (Powell et al., 2009).

Another aspect of best practices that I discuss here is proper justification of the choice and implementation of various model mechanisms. When EAs are used for optimization, individual measures of performance provide ample justification for the use of different evolutionary mechanisms. In simulation approaches the task is much more difficult, because justification has to include an account of plausibility and an adequate degree of alignment with the target system that is being simulated. For example, when simulating sexual reproduction, two-parent recombination should be used to create new offspring. On the other hand cloning might be a plausible account of cultural transmission in specific cases,

while other instances might call for multi-parent recombination. However, in many of the surveyed studies such accounts are largely absent (e.g. Axelrod, 1986; Takahashi, 2000; Pepper & Smuts, 2000; Hales, 2000; Lake, 2001; Hammond & Axelrod, 2006; Lake & Crema, 2012)

Finally a number of studies fall short in other research practices, such as arbitrary, and unjustified parameter values (Axelrod, 1986; Hammond & Axelrod, 2006) or the selection of inappropriate mechanisms, which have a known confounding effect on the variable being explained (Perez-Losada & Fort, 2011).³

The study of Kachel et al. (2010) on the grandmothering hypothesis is an example of a high quality study that satisfies all of the above mentioned practices. The choice of specific mechanisms is well justified and supported by evolutionary theory as it applies to the problem at hand, the model is described in comprehensive detail using the ODD protocol (Grimm et al., 2006), and the computer code is made publicly available.

To assure the highest possible quality of such future efforts the adoption of a few guiding principles is encouraged. As a first principle, both agent-based modelers as well as EC theorists should actively seek interactions and collaborations among each other whenever the nature of the research lies in the intersection of both topics. It has already been pointed out that the EC community has regrettably little to none interaction with those researchers who focus on the study of evolutionary dynamics (Eiben and Smith, 2007). Other leading figures of evolutionary computation have explicitly called for the development of a unifying EC framework which would, among other things, provide a “means by which outsiders can obtain a high level understanding of EC” (De Jong, 2005, p. 232) and possibly “open the door to interdisciplinary collaboration” (ibid). To the best of my knowledge, there has also been little outreach from the ABM community, apart from rare exceptions. Continued collaboration between EC experts and ABM researchers will only benefit the design of agentized simulations of evolution.

A second, closely related principle is a deeper focus on the theoretical considerations and

³Perez-Losada and Fort attempt to explain cultural drift, but use methods such as cloning, which are known to significantly contribute to drift.

justifications of the choice of specific EC mechanisms. In previous sections I have elaborated on how the implementation of different mechanisms has to be carefully adjudicated to avoid undesirable side-effects. It is therefore surprising to see how so many agent-based modeling studies forego the justification of some of the necessary mechanisms completely. In certain cases it has been proven that the lack of such reflection leads to very idiosyncratic results. For example, Axelrod's (1986) norms game has been replicated by Galan and Izquierdo (2006) using different evolutionary operators with striking results. The outcomes of the simulations were found to be highly sensitive to choice of different selection mechanism, fitness regimes and population sizes. Other studies also defy what is considered good practice in the EC community. The greatest danger seems to be the potentially confounding effect of genetic drift (as noted in section 2.5.2). Although pure genetic drift only occurs in selectively neutral contexts, drift-like phenomena can occur in selective environments as well. This is true, for example, when selection acts on absolute fitness and the variance in the fitness distribution is very low. A number of models use selection mechanisms that rely on absolute fitness, such as roulette-wheel selection, without off-setting its effects with disruptive recombination or mutation methods. Instead, modelers often combine this with other operators which are themselves susceptible to drift phenomena, such as cloning or the complete absence of mutation. In other cases, the exploitative nature of mutation is distorted with the use of bit-flipping in binary genotype representation without addressing hamming cliff issues. There are also many models which simulate insufficiently small populations, numbering in the dozens, or in one extreme case only a handful of agents.

Finally, as a third principle, just as with other important aspects of ABM design, it is necessary to provide all necessary details of the precise implementation of evolutionary mechanisms within the model. For example the ODD protocol (Grimm et al., 2006), has long been recognized as the standard for describing ABM design. In the light of this, it is surprising that some of the surveyed studies have offered only vague descriptions of the evolutionary algorithms powering their simulations. As reviewed above, these details are important when considering the sensitivity of a model, or even the possibility of its

replication.

Even despite some of these shortcomings, the usefulness of such models in offering fresh explanations for complex phenomena and their ability to provide insight into the working of dynamic adaptive systems is invaluable. The design of new agent-based studies following the evolutionary computation approach should be welcomed with enthusiasm, while simultaneously promoting the quest for even sounder designs supported by good practice and theoretical rigor through collaboration with the EC community.

Table 2.1: Comparison of the implementation of evolutionary mechanism and concepts in different ABMs

#	Study	Domain of evolution	Co-evolution	Parameter space domain	Genotype-phenotype mapping	Phenotype	Phenotype change
1	Axelrod	Populations	Intrasp. comp. ⁴	Integer (cardinal)	Binary code	Boldness,vengefulness	Via genotype change
2	Bowles and Gintis	Populations	Intrasp. comp.	Categorical	Direct	Public goods game strategy	Vertical transmission
3	Chiang	Populations	Intrasp. comp.	Real-valued	Direct	Thresholds of cooperation	Vertical transmission
4	de Vos et al.	Populations	Intrasp. comp.	Categorical	Direct	Indirect reciprocity strategy	Vertical transmission
5	Dosi et al.	Mind	Intersp. comp. ⁵	Real-valued, operators	Direct	Estimate of utility function	Vertical transmission
6	Edmonds	Mind	Intersp. comp.	Real-valued, operators	Direct	Estimate of utility function	Vertical transmission
7	Hales	Populations	Intrasp. comp.	Boolean	Binary code ⁶	PD strategy, tag	Via genotype change
8	Hammond & Axelrod	Populations	Intrasp. comp.	Categorical	Direct	Cooperation/defection (PD), tag	Vertical transmission
9	Janssen	Populations	Intrasp. comp.	Categorical, real-valued	Direct	Cooperation/defection (PD), tag, probability of feedback, trust	Learning (neural networks)
10	Kachel et al.	Populations	None	Real-valued	Arithmetic ⁷	Longevity, length of reproductive period	Via genotype change
11	Kohler et al.	Populations	Intrasp. comp.	Categorical	Direct	Public goods game strategy, hierarchical/non-hierarchical	Vertical transmission
12	Klos	Populations	Intrasp. comp.	Boolean	Binary code	Iterated PD strategy	Via genotype change
13	Lake & Crema	Populations	Intrasp. comp.	Integer (ordinal)	Direct	Efficiency of resource extraction	Vertical and horizontal transmissison, innovation
14	Lake	Populations	Intrasp. comp.	Categorical	Direct	Cultural/individual learning	Vertical transmission
15	Macy & Skvoretz	Populations	Intrasp. comp.	Categorical	Direct	PD strategy, tags	Vertical transmission
16	Pepper and Smuts	Populations	Intrasp. comp.	Categorical	Direct	Presence/absence of alarm calling/food restraint	Vertical transmission
17	Perez-Losada & Fort	Populations	None	Boolean	Direct	Types of crops cultivated	Vertical transmission
18	Powell et al.	Populations	None	Real-valued	Direct	Degree of skill in abstract cultural domain	Vertical and horizontal transmission
19	Premo & Hublin	Populations	None	Integer (cardinal)	Arithmetic ⁸	Abstract genetic/cultural traits	Vertical and horizontal transmission
20	Saam & Harrer	Populations	Intrasp. comp.	Real-valued	Direct	Resource acquisition strategy	Vertical transmission
21	Santos et al.	Populations	Intrasp. comp.	Categorical	Direct	Cooperation/defection (PD, stag hunt, snowdrift)	Vertical transmission
22	Takahashi	Populations	Intrasp. comp.	Real-valued	Direct	Indirect reciprocity strategy	Vertical transmission
23	Xue et al.	Populations	None	Real-valued	Direct	Degree of adaptation to environment	Vertical transmission, innovation

Table 2.2: (Continued) Comparison of the implementation of evolutionary mechanism and concepts in different ABMs

#	Fitness	Survival	Selection	Recombination	Mutation
1	Payoffs from public goods game	Generational	Truncation	Cloning	Bit-flipping, 1%
2	Payoffs from public goods game	Steady-state; stochastic death, fitness-based	Not specified ⁹	Cloning	Random resetting, 1%
3	Payoffs from PD	Generational	Tournament ¹⁰	Cloning	Not specified
4	Dead/alive	Steady-state; stochastic death, averted by help from others	Fitness-neutral	Cloning	None
5	Distance from max. utility	Generational	Not specified ¹¹	Sub-tree crossover	None
6	Distance from max. utility ¹²	Generational	Not specified ¹³	Not specified ¹⁴	Not specified ¹⁵
7	Payoffs from PD	Generational	Fitness-proportionate	Cloning	Bit-flipping, 0.1%
8	Payoffs from PD	Steady-state; fixed death rate	Fitness-proportionate	Cloning	Random resetting, 0.5%
9	Payoffs from PD	Steady-state; agents die if reputation/payoffs below threshold	None ¹⁶	None	None
10	Dead/alive	Steady-state; stochastic death, based on longevity trait	Age-dependent	One-point crossover	Truncated Gaussian perturbation, 5%
11	Energy from resources	Age-specific mortality	Age-specific fertility	Not specified	Not specified
12	Payoffs from PD ¹⁷	Generational	Tournament ¹⁸	Two-point crossover	Bit-flipping, 0.5%
13	Energy from resources	Generational	Truncation	Cloning	None
14	Energy from resources	Generational	Tournament	Cloning	None
15	Payoffs from PD	Generational	Fitness-neutral, local	Fitness-biased uniform crossover	Bit-flipping, 1%
16	Energy from resources, death by predation	Generational	Fitness-neutral, local	Cloning	None
17	Age-based	Steady-state; fixed death rate	Fitness-neutral	Cloning	None
18	Degree of skill in abstract cultural domain	Generational	Fitness-neutral	Cloning	Gumbel distributed perturbation, 100 ¹⁹ %
19	Fully determined by genotype	Steady-state; stochastic death, based on current population size	Fitness-proportionate	One-point crossover	Incremental, 0.001%
20	Energy from resources	Generational	Fitness-neutral	Additive	None
21	Payoffs from games	Generational	Fitness-proportionate	Cloning	None
22	Energy from resources	Generational	Truncation	Cloning	Perturbation, distribution not specified, probability varied
23	Distance from optimal fitness	Steady-state; stochastic death, fitness-based	Fitness-neutral ²⁰	Cloning	Truncated uniform perturbation, 100% ²¹

-
- ⁴Intraspecific competition (within species)
- ⁵Interspecific competition (between species)
- ⁶Only the tag part of the genotype is represented in binary.
- ⁷Mean of two real-valued alleles.
- ⁸Each individual allele in a given locus has a fixed fitness contribution to the phenotype.
- ⁹“By the fitness of an agent, we mean the expected number offspring produced by the agent in one period minus the probability the agent dies in that period”. (Bowles and Gintis 2004, p.18)
- ¹⁰Agents imitate fittest neighbor.
- ¹¹“Selection consists of preserving the fittest rules, and discarding the less fit ones”. (Dosi et al. 1999, p. 15)
- ¹²A penalty for model complexity (tree size) is also assessed.
- ¹³“[The agent] then selects the best models in terms of fitness for carrying over in the next period”. (Edmonds 1997, p.6)
- ¹⁴“[The agent] produces some new ones by either combining the previous models with a new operator or by growing a new random one”. (Edmonds 1997, p.6)
- ¹⁵See above.
- ¹⁶There is no reproduction or inheritance in Janssen’s model. Unfit agents are simply removed and replaced by new random ones.
- ¹⁷Relative to neighbors’ payoffs.
- ¹⁸Candidates sampled from Moore neighborhood.
- ¹⁹Represents imperfect imitation.
- ²⁰Random surviving agents is chosen to create a new clone in place of a recently deceased agent.
- ²¹Represents imperfect imitation.

Chapter 3: A Model of Co-Evolution of Signs and Cultural Traits

Abstract: The identification of main drivers of stability and change in human cultures remains an open problem. Traditional methods have fallen short of explaining the focal mechanisms responsible for the evolution of culture. Thus, I take an alternative approach and develop an agent-based model capable of simulating cultural dynamics with the use of various underlying mechanisms. The model is inspired by principles of dual inheritance theory. Specifically, I test the mechanisms of indirectly biased transmission and guided variation as hypothesized drivers of the diffusion of cultural traits and the emergence of cultural signs. I show that the resulting distribution of cultural traits and external markers is dependent on the chosen network structure, ratio of vertical and horizontal transmission, and the specific choice of mechanisms.

3.1 Introduction

Cultural practices structure human conduct in contexts ranging from modes of governance and natural resource management to conflict resolution, power relationships and religious practices. Due to their wide-ranging effects and observed persistence, such practices can serve both as insurmountable roadblocks and powerful catalysts for problem solving in different areas of life (e.g. economics, politics or international relations). Consider the example of climate protection: global climate change is considerably driven by human-developed natural resource management practices determined by values, beliefs and social institutions (Tohme, 1992). Another area of concern is opinion radicalization in social groups. Data-driven research shows that the emergence of terrorism is related to differentiation in cultural dimensions and the absence of cultural engagement in individuals (Kluch and Vaux, 2015). If we could properly understand the life-cycle of cultural practices, we would be better equipped to solve complex issues in these areas of interest, as well as others. Here, I outline an effort to develop a powerful explanatory model of the complex long-range dynamics of cultural practices in human societies and present some preliminary results.

I begin by introducing the theoretical background of this work and commenting on previous literature on the topic of modeling cultural evolution in section 3.2. In sections 3.3-3.6, I describe my agent-based model and its implementation in detail. I continue with an overview of the results of model simulations, which I give in section 3.8. Finally, in sections 3.8-3.9, I discuss these findings and offer a conclusion.

3.2 Background

The model presented here is embedded in the dual inheritance theory of cultural evolution which assumes both horizontal (intra-generational) and vertical (inter-generational) transmission of cultural traits in a population of social individuals (Boyd and Richerson, 1985). Dual inheritance theories respond to neo-evolutionary theories (e.g. Wilson, 1975) and hypothesize that culture is transmitted in human populations both vertically—from

generation to generation—by forces that can be modeled with mechanisms similar to biological evolution, as well as horizontally—within a single generation—by social influence and individual learning. Boyd and Richerson (1985, 2005) use genetic evolution partially as a metaphor for the evolution of culture, although they point out major differences. Among some of the important mechanisms in the dual inheritance framework are indirectly biased transmission and guided variation. Transmission of genetic trait variants is said to be indirectly biased when it is driven by preferences for unrelated phenotypic traits. Guided variation is the process of individual in-life adaptation based on the evaluation of self-generated and self-explored trials. I hypothesize that these two mechanisms are important to cultural evolution in human populations and they crucially contribute to the emergence of cultural clusters and cultural signs. To determine what kinds of macro-scale phenomena these mechanisms produce in different qualitative types of human societies I develop and analyze an agent-based model where individuals enter into dyadic interactions with others based on co-evolving preferences for external markers and attempt to solve simple coordination problems. Crucially, I assume that the strategies for solving these problems are selectively neutral, while their interactions are not. In other words each strategy is equally good, but the collaborating individuals will only succeed when their strategies are the same. This is a common occurrence in the realm of culture.

Within computational social science, researchers have also contributed towards explanations of cultural dynamics. One of the most influential ABMs of culture is due to Axelrod (1997). A number of agent-based modelers have since revised and extended Axelrod's original model (e.g. Klemm et al., 2003; Parisi et al., 2003; Centola et al., 2007; Flache and Macy, 2011).

However, these models assume agents possess perfect knowledge of their neighbors cultural “genotype”. Moreover, they only model horizontal transmission of cultural information. Finally, the studies use Prisoner’s Dilemma to model the problematic nature of cooperation, whereas I assume cooperation is unavoidable. Other agent-based models of culture have attempted to circumvent these limitations. One approach to represent agents indirect

inference of others cultural traits is the use of observable tags. Hales (2000) developed a model in which agents played one-shot Prisoner’s Dilemma games only with agents possessing the same tag. Hammond and Axelrod (2006) also made use of tags in their agent-based model of ethnocentrism. However, in these models the tags are assigned meaning *a priori*. Reynolds (1994, 2008) has developed the class of cultural algorithms, introducing vertical transmission. These models are genetic algorithms in which agents evolve based on their performance against a fitness function and a belief space, which consists of the cultural genotypes of top performers. The rest of the population then acquires traits from this space via an influence function. Yet others have explored cooperation using the ultimatum game (Shutters, 2009) or the stag-hunt and snowdrift games (Santos et al., 2006), but without considering some of the other aspects I wish to explore. Furthermore, I take a network approach to modeling the interactions between actors. The network science paradigm holds that network structure is an important determinant of dynamics of certain processes in agent populations. Different network formation models have been proposed over the years. The random graph model was the first such venture (Erdos & Rényi, 1960). In recent years, attempts at more generating more realistic social networks, have resulted in the use of algorithms such as the preferential attachment model (Barabási & Albert, 1999), or the small-world model (Watts & Strogatz, 1998). It is evident from this concise review that computational cultural simulations can be useful in exploring cultural dynamics, however none of the existing models sufficiently address the research objectives that I formulate.

3.3 Overview

In this section I will give a detailed specification of an agent-based model aimed at simulating the co-evolution of cultural signs, or “tags”, as they are often referred to in the literature, and cultural traits possessed by artificial agents. I use the ODD protocol (Grimm et al., 2006) to describe the model. I begin with defining the purpose of the model.

3.3.1 Purpose

The purpose of the model is to *simulate* and *explain* the causes of differentiation of human populations into groups distinctly defined by shared cultural trait variants and observable external markers, and the clustering of such groups in social and physical space. Specifically, the model is designed to test the hypothesis that the phenomenon of cultural clustering and the emergence of meaningful cultural signs is caused in part by the conjunction of two mechanisms: (1) The vertical (parent to child¹) transmission of cultural trait variants indirectly biased by the possession of certain external markers, and (2) the guided variation of partner selection strategies based on past experience.

3.3.2 Entities and Variables

The main entities in the models are the individual agents. Each agent represents a single person. The agents are defined by a collection of state variables (as shown in table 3.1). These can be separated into static variables which are fixed throughout the agent's existence and dynamic variables which may change over time. A description of the static variables is provided first.

- *CulturalTrait*: This variable represents an abstract cultural trait which can take on one of many possible interpretations such as religious affiliation, commodity exchange practices, etc. The value assigned to the agent represents the particular trait variant that the agent possesses, e.g. Islam, as opposed to Christianity or Buddhism (in the case of religion), or barter, as opposed to gifting or currency exchange (in the case of commodity exchange). The only assumption here is that the available variants are selectively neutral when considered independently of each other. This is simply generalizing the statement that a person is no better off by following either Muslim or Christian faith, all else being equal, and presuming no interaction with others. In

¹Here I use the term *parent* in a cultural, rather than biological sense. A cultural parent can be one of many influential persons from the actor's perspective, such as teachers, clergy, and other authority figures, as well as their actual biological parents, friends, or significant others.

Table 3.1: Agent variables

Name	Domain	Scale	Type
<i>CulturalTrait</i>	Integer	Categorical	Static
<i>Tag</i>	Integer	Categorical	Static
<i>Neighborhood</i>	List of agents	Categorical	Static
<i>Fitness</i>	Integer	Cardinal (ratio)	Dynamic
<i>NumGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>NumBad</i>	List of integers	Cardinal (ratio)	Dynamic
<i>FirstGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>FirstBad</i>	List of integers	Cardinal (ratio)	Dynamic
<i>LastGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>LastBad</i>	List of integers	Cardinal (ratio)	Dynamic

reality the individual payoffs resulting from an interaction between two Muslim persons might be different from those between a Muslim and a Christian person, depending on the nature of the interaction. In general, the interactions between the trait variants are therefore selectively biased. This feature of the model will be revisited and explained in full detail in section 3.3.3. It should be noted that the cultural makeup of real individuals does not simply consist of a single trait, but rather of a collection of often inter-dependent traits. Here a single trait is modeled for sake of clarity and simplicity.

- *Tag*: This variable represents an abstract external marker possessed by the agent and directly observable by other agents. Examples of interpretations include morphological features, style of clothing, accents, etc. Once again it is assumed that each variant of the tag is selectively neutral. For example from a functional viewpoint the agent gains no extra benefit from wearing a red shirt instead of a blue shirt (apart from minor exceptions such as use for camouflage).
- *Neighborhood*: A list of other agents with whom the agent is able to interact. This can be viewed in network terms, where the agent represents a single node i from the vertex set V of a graph $G = (V, E)$. For every agent j in i 's neighborhood the unordered pair (i, j) defines an undirected edge in the graph's edge set E . The neighboring relation is

symmetric, i.e if i is in j 's neighborhood then j must be in i 's also.

The following list provides a description of the dynamic agent variables:

- *Fitness*: This variable represents the agent's "cultural" reproductive fitness. A higher fitness value gives the agent a better chance of producing cultural offspring, which inherit its cultural makeup (the "genes" in this model) to a certain degree. Fitness is modified through interaction with other agents. When agents pair with others who possess the same trait variant, they receive a boost to their fitness. When the partner's trait is different the agent's fitness value decreases. It should be noted that because the goal of the built-in EA is not optimization, we do not concern ourselves with issues such as escaping local fitness optima. We are merely interesting in the dynamics, and as far as we know, local optima traps might as well be part of the underlying reality of cultural evolution.
- *NumGood*: A list of length n where n is the number of possible tag variants. The list is used to store the total number of successful interactions with others possessing the associated tag variant throughout the agent's entire lifetime.
- *NumBad*: Analogous to *NumGood*, the difference being that the total numbers of unsuccessful interactions are stored.
- *FirstGood*: A list of length n where n is the number of possible tag variants. The list is used to store the number of steps since the first successful interaction with another agent possessing the associated tag variant.
- *FirstBad*: Analogous to *LastGood*, the difference being that the times since the last unsuccessful interactions are stored.
- *LastGood*: A list of length n where n is the number of possible tag variants. The list is used to store the number of steps since the last successful interaction with another agent possessing the associated tag variant.
- *LastBad*: Analogous to *LastGood*, the difference being that the times since the last unsuccessful interactions are stored.

Apart from the *Fitness* variable, the dynamic variables relate to agents' preferences. The preferences are modified through the mechanism of guided variation, which rests on weighing previous experiences. The above variables code for these experiences.

The second type of entity is the model controller which deploys and manages the processes throughout a simulation. The overview of the processes themselves and their scheduling is given in the following subsection. Table 3.2 provides an overview of the model parameters, while table 3.3 offers brief descriptions.

Table 3.2: Model Parameters

Name	Domain	Scale
<i>Configuration</i>	Integer	Categorical
<i>NumAgents</i>	Integer	Cardinal (ratio)
<i>NumGenerations</i>	Integer	Cardinal (ratio)
<i>NumRounds</i>	Integer	Cardinal (ratio)
<i>NumTags</i>	Integer	Cardinal (ratio)
<i>NumTraits</i>	Integer	Cardinal (ratio)
<i>SuccessPayoff</i>	Integer	Cardinal (ratio)
<i>FailurePayoff</i>	Integer	Cardinal (ratio)
<i>AdjacencyMatrix</i>	Matrix of Booleans	Categorical (Boolean)
<i>MutationRate</i>	Floating-point number	Cardinal (ratio)

Table 3.3: Model parameter descriptions

Name	Description
<i>Configuration</i>	Type of simulation (see Sec. 2.3)
<i>NumAgents</i>	Number of agents in the simulation
<i>NumGenerations</i>	Number of generations executed
<i>NumRounds</i>	Number of rounds in each generation
<i>NumTags</i>	Size of the tag variant set from which each agent is assigned one
<i>NumTraits</i>	Size of the trait variant set from which each agent is assigned one
<i>SuccessPayoff</i>	Fitness payoff for a successful interaction
<i>FailurePayoff</i>	Fitness payoff for an unsuccessful interaction
<i>AdjacencyMatrix</i>	Adjacency matrix of the graph which defines the agents' neighborhoods
<i>MutationRate</i>	Probability of a mutation in a single heritable characteristic of an offspring

3.3.3 Process Overview and Scheduling

The main processes in the model consist of agent activity during which agents choose partners and interact with them. After a certain number of activations the agent population is processed by an evolutionary algorithm, which creates offspring through the recombination and mutation of heritable characteristics of selected parent agents. These processes are repeated for a number of generations.

Algorithm 1 Top-level Simulation Flow

```

1:  $params \leftarrow \{Configuration, NumAgents, NumGenerations, NumRounds, NumTags,$ 
    $NumTraits, SuccessPayoff, FailurePayoff, AdjacencyMatrix, MutationRate\}$ 
2: procedure CULTUREMODEL( $[params]$ )
3:    $agents \leftarrow initializeAgents([params])$ 
4:    $i \leftarrow 1$ 
5:   while  $i \leq NumGenerations$  do
6:      $j \leftarrow 1$ 
7:     while  $j \leq NumRounds * NumAgents$  do
8:        $agent \leftarrow selectRandomFrom(agents)$ 
9:        $potentialPartners \leftarrow getPotentialPartners(agent, agents, Configuration)$ 
10:       $partner \leftarrow getPartner(agent, potentialPartners)$ 
11:       $agent, partner \leftarrow interact(agent, partner)$ 
12:       $j \leftarrow j + 1$ 
13:    end while
14:     $offspring \leftarrow createOffspring(agents, [params])$ 
15:     $agents \leftarrow offspring$ 
16:     $i \leftarrow i + 1$ 
17:  end while
18:  return  $agents$ 
19: end procedure

```

The top-level flow of the model is described in pseudocode in algorithm 1. Each generation consists of $NumRounds * NumAgents$ activations, so that in effect each round consists of $NumAgents$ agent activations. During each activation event one randomly chosen agent is activated. Each agent has an equal probability of being activated. Thus, for each agent the expected value of activations is one per round, although actual results may vary due to the stochasticity involved.

Figure 3.1 shows a rough outline of the logic of agent activity in the model. Once an agent is activated, it will first select a subset of agents from its neighborhood which it designates as potential interaction partners (as shown in algorithm 2). If the model configuration is of

the Unbiased type, then the entire neighborhood is selected. Under any other configuration the selection is done with the use of the *calculateBLA* routine, which calculates a list of base-level activations for each possible tag T . Here the agent takes into account its history of successful interactions. Let us denote the base-level activation for successful interactions with tag T as τ_T^+ . The calculation is as follows:

$$\tau_T^+ = \ln \left[\sum_i^n t_i^{-d} \right] \approx \ln \left[t_n^{-0.5} + \frac{2(n-1)}{\sqrt{t_1} + \sqrt{t_n}} \right] \quad (\text{when } d = 1/2) \quad (3.1)$$

Here t_i is the time elapsed since the i -th interaction with an agent bearing tag T , while $n = \text{NumGood}_T$ is the total number of such experiences, $t_n = \text{LastGood}_T$ is the time since the most recent experience and $t_1 = \text{FirstGood}_T$ is the time since the first experience. Finally, d is the rate of decay. Due to the computational infeasibility of the above relationship when n is large, I implement a widely-used approximation (Anderson & Lebiere, 1998). In line with convention, I use $d = \frac{1}{2}$ (Petrov, 2006).

Once these calculations are carried out, the agent randomly selects a single tag from the set of all tags of type T , such that $\tau_T^+ > 0$ (if there are any)². Finally the agent identifies all of its neighbors that possess this tag (if there are any), and selects them as potential interaction partners.

From these candidates the agent attempts to acquire an actual interaction partner (as shown in algorithm 3). The first step is to simply select a random neighbor from the list of candidates. If the model configuration is of the Unbiased type, then this is all that needs to be done and the selected candidate becomes the agent's partner. Under any other configuration the selected candidate must decide whether it agrees to become the partner. The candidate does this by comparing its base-level activation for the agent's own tag calculated on its successful interactions to the BLA values calculated on its unsuccessful

²the threshold value 0 is used as a default in the ACT-R framework, although it has been pointed out that the retrieval threshold should be estimated for each task separately (Anderson et al., 2004). Here, we do not estimate the threshold, as we rather focus on other parameters. However, we overcome this shortcoming later, in chapter 4.

interactions with the given tag³. If the levels for successful interactions exceed those for the unsuccessful interactions, the candidate agrees to become the agent’s partner. Otherwise, the agent is left without an interaction partner.

Algorithm 3 Selection of Partner

```

1: procedure GETPARTNER(agent, potentialPartners, Configuration)
2:   partner  $\leftarrow$  NULL
3:   if potentialPartners is not empty then
4:     candidate  $\leftarrow$  selectRandomFrom(potentialPartners)
5:     if Configuration = “Unbiased” then
6:       partner  $\leftarrow$  candidate
7:     else
8:       positiveBLA  $\leftarrow$  calculateBLA(candidate, “good”)
9:       negativeBLA  $\leftarrow$  calculateBLA(candidate, “bad”)
10:      if positiveBLA[agent.Tag] > negativeBLA[agent.Tag] then
11:        partner  $\leftarrow$  candidate  $\triangleright$  if candidate has good experience with agent’s tag
        it will agree to become its interaction partner
12:      end if
13:    end if
14:  end if
15:  return partner
16: end procedure

```

The interaction itself proceeds according to the pseudocode laid out in algorithm 4. If the agent was not successful in acquiring a partner it receives the *FailurePayoff*. If the agent was able to acquire a partner the two of them compare their cultural trait variants. If the variants agree both agents receive the *SuccessPayoff*, however, if the variants differ they both receive the *FailurePayoff*. Moreover, under the Lamarckian model configuration the agents update either their *LastGood* and *NumGood* fields, or *LastBad* and *NumBad* fields, depending on whether the interaction was successful or unsuccessful.

After a single generation has elapsed, the evolutionary algorithm is invoked (as shown in algorithm 9). The algorithm creates *NumAgents* offspring, which then replace the entire agent population from the recently executed generation. Although in real-world cases human populations consist of overlapping generations, here, without a clearly defined distribution of longevity and a set of influencing factors, I simply model non-overlapping generations. I introduce overlapping generations in Chapter 6. Reproduction is local, and the network

³this is done in the same way as in equation 3.1, only instead of *NumGood*, *FirstGood* and *LastGood* the values for *NumBad*, *FirstBad* and *LastBad* are used.

Algorithm 4 Agent interaction

```
1: procedure INTERACT(agent, partner, Configuration, SuccessPayoff, FailurePayoff)
2:   if partner is NULL then
3:     agent.Fitness  $\leftarrow$  agent.Fitness + FailurePayoff
4:   else if agent.culturalTrait = partner.CulturalTrait then
5:     agent.Fitness  $\leftarrow$  agent.Fitness + SuccessPayoff
6:     partner.Fitness  $\leftarrow$  partner.Fitness + SuccessPayoff
7:     if Configuration = “Lamarckian” then
8:       agent.LastGood[partner.Tag]  $\leftarrow$  0
9:       partner.LastGood[agent.Tag]  $\leftarrow$  0
10:      agent.numGood[partner.Tag]  $\leftarrow$  agent.numGood[partner.Tag] + 1
11:      partner.numGood[agent.Tag]  $\leftarrow$  partner.numGood[agent.Tag] + 1
12:    end if
13:   else
14:     agent.Fitness  $\leftarrow$  agent.Fitness + FailurePayoff
15:     partner.Fitness  $\leftarrow$  partner.Fitness + FailurePayoff
16:     if Configuration = “Lamarckian” then
17:       agent.LastBad[partner.Tag]  $\leftarrow$  0
18:       partner.LastBad[agent.Tag]  $\leftarrow$  0
19:       agent.numBad[partner.Tag]  $\leftarrow$  agent.numBad[partner.Tag] + 1
20:       partner.numBad[agent.Tag]  $\leftarrow$  partner.numBad[agent.Tag] + 1
21:     end if
22:   end if
23:   return agent, partner
24: end procedure
```

structure of the population is preserved. Thus for each agent occupying a node i one offspring will be produced, which will then take its place on the node i . The parents of this offspring are selected from the set containing the agent currently occupying the node i and its neighbors. The selection mechanism is a tournament of size 3. To select a single parent, three agents are randomly selected from the set of candidates and the one of the three with the highest fitness becomes the parent. Each offspring has two parents ⁴. I use tournament selection for a couple of reasons. First, it represents the imperfect information that agents possess when selecting behavioral models, in the sense that sub-optimal models may be selected. It is also independent of the distribution of absolute fitness, unlike other mechanisms such as fitness-proportionate selection (we only care about relative fitness). Furthermore, I model two-parent reproduction as a baseline, although multi-parent configurations are possible in cultural evolution. I explore this concept in Chapters 5 and 6.

When the parents are selected the offspring is created by performing unbiased uniform crossover on their heritable characteristics (see De Jong, 2005, p.64-65). This includes all of the fields listed in table 3.1, apart from *Neighborhood* which is always defined by the immutable *AdjacencyMatrix* and *Fitness* which is set to 0 for all agents at the beginning of a generation. After the offspring is created it undergoes potential mutation. This consists of changing each of its fields with a 1% probability⁵. The probability is applied element-wise in the case when fields are lists. In the case of cardinal variables the mutations are done in the form of small Gaussian perturbations. If the variable is categorical, mutation is carried out via random resetting of values within the allowed range.

3.4 Design Concepts

- *Sensing*: The agents only sense their direct neighbors. The agents' positions in the

⁴If the candidate set contains two agents only, then no selection takes place and the two automatically become the parents. If the candidate set only contains one agent then the offspring will be its exact copy (barring any subsequent mutations).

⁵As the goal of this EA is not optimization, there is no optimal level of mutation. Similarly, it is beyond the scope of this work to estimate the true rate of mutation of cultural traits in individuals. Therefore, the consensus "appropriate" rate of mutation is applied here.

Algorithm 5 Evolutionary Algorithm

```
1: procedure CREATEOFFSPRING(agents, TournamentSize, MutationRate)
2:    $i \leftarrow 1$ 
3:   offspring  $\leftarrow \emptyset$ 
4:   while  $i \leq \text{length}(\text{agents})$  do
5:     parents  $\leftarrow \emptyset$ 
6:     agent  $\leftarrow \text{agents}[i]$ 
7:     while  $\text{length}(\text{parents}) \neq 2$  do
8:       candidates  $\leftarrow \text{agent}$ 
9:        $j \leftarrow 1$ 
10:      while  $j \leq \text{TournamentSize}$  do
11:        candidates  $\leftarrow \text{candidates} \cup \text{selectRandomFrom}(\text{agent.Neighborhood})$  ▷
12:        including checks for ensuring non-duplicate candidates
13:         $j \leftarrow j + 1$ 
14:      end while
15:      bestCandidate  $\leftarrow \text{tournament}(\text{candidates})$ 
16:      parents  $\leftarrow \text{parents} \cup \text{bestCandidate}$ 
17:    end while
18:    thisOffspring  $\leftarrow \text{uniformCrossover}(\text{parents})$ 
19:    thisOffspring  $\leftarrow \text{mutation}(\text{thisOffspring}, \text{MutationRate})$ 
20:    offspring  $\leftarrow \text{offspring} \cup \text{thisOffspring}$ 
21:     $i \leftarrow i + 1$ 
22:  end while
23:  return offspring
24: end procedure
```

network remain fixed throughout their lifetimes. Before an interaction takes place the agents are only able to directly sense their neighbors' tags.

- *Interaction:* Upon activation an agent is compelled to interact with one of its neighbors. The choice of its interaction partner differs based on the model configuration. In the unbiased configuration the agent simply chooses randomly. The selected partner then always agrees to the interaction. This configuration was implemented as a baseline and a means of comparison against the other configurations.

In the remaining configurations the agent chooses to interact with neighbors based on their preferences for certain tag variants. These preferences are formulated by their prediction and adaptation mechanisms.

The interaction choice must be mutual: When a candidate is selected as a preferred partner by an agent, the candidate must decide whether it is comfortable interacting with the agent as well. Once again, the decision is made by judging the preference for the agent's tag.

The interaction is done by comparing the cultural trait variants of the two agents. The interaction is deemed a success if the two variants are equal, otherwise it is deemed unsuccessful. Thus, one can think of the interaction as an $n \times n$ pure coordination game with n Nash equilibria, where n is the number of trait variants present in the model. The coordination game is a natural choice, because I assume the cultural trait variants represent beliefs, behaviors or knowledge related to a particular domain and that these are functionally equivalent to their alternatives and therefore selectively neutral. If two agents are forced to solve a problem together in the given domain, I assume that they are more likely to succeed if their beliefs and behaviors align. For example, if two agents need to compare measurements it usually does not matter whether they are in metric or imperial units as long as both agents are familiar with the scale being used, and know that it is in fact being used. However, if one of the agents is only familiar with inches, and the other is only familiar with centimeters,

the measurement problem becomes more difficult⁶. Many domains in today’s world are rife with such cultural differences that make effective cross-cultural collaboration difficult, often involving stakes much higher than incorrect measurements⁷.

- *Prediction*: Since the interaction depends on the trait variants of the agents and because these are unknowable by the agents in advance, an agent can only help itself by attempting to predict a neighbor’s trait variant through means other than direct observation. As communication between agents is not allowed in this version of the model, the only information the agents can use to infer the others’ trait variants are the directly observable tags. Thus, the agents use their past experience to form predictions whether the possession of any given tag correlates with the possession of a particular trait variant. This prediction is done via equation 3.1. The relationship and its derived estimate which represent the process of reinforcement learning is taken directly from the ACT-R cognitive architecture (Anderson & Lebiere, 1998). The ACT-R framework is the most influential cognitive architecture, and is well suited towards modeling cognitive phenomena such as memory. The quantity computed in equation 3.1 is very sensitive to recent experiences, while the importance of older experiences progressively decays with time. A new positive (negative) experience with a given tag results in a spike in its base-level activation, meaning that the agent will associate it more strongly with positive (negative) experiences in the future. A lack of positive (negative) experience over an extended period of time will result in the gradual decrease in the tag’s base-level activation and a weakening association with positive (negative) results. Notice that the quantities *NumGood*, *FirstGood*, *LastGood* and the corresponding negative equivalents are inherited by offspring. Taken literally, this would mean that offspring inherit the actual memories and experiences of their ancestors. While this is certainly not possible in the human world, I argue here that this is valid from a modeling perspective. In the model, these values are not used in

⁶Difficulties could also arise if both agents know both scales, but each assume that a different scale is being used by the other.

⁷To give just one example consider the effect that ideological and religious differences have on the contemporary social and political landscape both on a national as well as international level.

themselves, they are simply a means in expressing the salience of positive/negative associations with tags. Such associations are commonly inherited in cultural evolution through socialization processes (Berger & Luckmann, 1966). The agents thus inherit the means for expressing them. True prediction, in the sense that newly acquired information during an agent’s lifetime (i.e. the agent’s own interactions) is used to update the predicted outcomes, is only present in the “Lamarckian” configuration of the model. As mentioned before, there is no prediction in the unbiased configuration. In a third “genetic” configuration, also devised as a baseline to control for Lamarckian processes in cultural evolution, the prediction calculation does take place, but the input variables are not updated after interaction events. Hence, while offspring still inherit preferences from their parents, they are not plastic and can only change from one generation to another through forces of recombination and mutation.

- *Adaptation* : As described above, one way the agents are able to adapt to a changing environment, particularly in the Lamarckian configuration, is the continual updating of their trait variant predictions for others, by incorporating newly acquired information from recent interactions into their calculations. This type of individual learning falls under the category of what Boyd and Richerson (pp. 6, 82) define as *guided variation*: learning in the form of trial and error in which the trials are self-generated and explored by the individual itself, the results of which can be transmitted to other generations.

Another level of adaptation, present in all model configurations, is that on the inter-generational, evolutionary scale. Because parent selection is fitness-dependent, individuals who do not perform well in the current environment, either because they possess a rare trait variant, an unpopular tag, or because they form predictions that are inconsistent with reality, are less likely to pass on their traits in favor of better-performing individuals.

- *Stochasticity*: There are three stages of the model that are affected by stochasticity. The first stage consists of certain choices the agents make (for example which one neighbor from the possibly many sharing the same preferred tag is selected as a

potential partner in a given round). In this case the stochasticity is simply a proxy for unknown parameters that go beyond the scope of the model. The other two stages include agent initialization and the recombination and mutation phases of the evolutionary algorithm. Here stochasticity serves the purpose of injecting additional diversity into the agent population, as well as simulating evolutionary mechanisms.

- *Emergence*: The model is designed to explore whether and how any regularities in the distribution of cultural trait variants and external markers arise from the decentralized actions of agents possessing limited information and cognitive abilities. The intent is to give a dynamic explanation of how long-term, stable trends emerge from the myopic agent-level adaptations and the non-teleological forces of cultural evolution.
- *Observations*: In accordance with the above stated goals the model output tracks the population-wide distributions of cultural traits variants and tags, as well as a number of statistics derived from them as a function of time and of the simulation parameters.

3.5 Initialization and Inputs

The model is initialized by providing the values of model parameters given in table 3.2. The possible configurations are (a) unbiased, (b) genetic, and (c) Lamarckian. The agents are initialized before the start of the first generation. Their *Fitness* is set to 0 and their neighborhoods are created based on the provided *AdjacencyMatrix*. Tags, denoted by their integer values, are assigned to agents randomly from the uniform distribution over the set of integers from 1 through *NumTags*. In analogous fashion, the cultural trait variants are also assigned randomly from the uniform distribution over the set of integers from 1 through *NumTraits*. It is important to note that tag assignment is not correlated with the assignment of the cultural trait variants in any way during the initialization process. The *NumGood* and *NumBad* values are initialized randomly from the uniform distribution over the integers from the set from 1 through *NumRounds*. Finally, the values of *FirstGood*, *FirstBad*, *LastGood*, and *LastBad* are initialized randomly from the uniform distribution

over the integers 1 through $NumRounds * NumAgents$ (while ensuring a valid sequence of events for each agent). This in essence creates fictional memories for the first generation of agents on a scale naturally tied to their lifetime. The number of generations is given by the *NumGenerations* parameter. Each agent is placed on a single node in a network given by the *AdjacencyMatrix* parameter. This matrix can define any network structure that the user wishes to use. After initialization the model requires no further input from the user.

3.6 Verification and Validation

The model was coded in MATLAB. The model was verified with the use of extensive and detailed code walkthroughs and debugging. Validation of the model was performed to the extent that distribution moments of important network statistics of the simulated agent networks were compared to stylized network characteristics of plausible real-world scenarios, i.e. high modularity, moderate trait diversity, and low tag entropy. Validation in the sense of direct quantitative comparisons with specific empirical observations was not carried out at this point. This is the focus of Chapters 5 and 6. The source code along with documentation can be accessed in the OpenABM model database at the following link:

<https://www.openabm.org/model/5813>.

3.7 Experimental Design

I perform extensive experimentation on the model implemented as described in previous sections by executing a large number of simulations. The simulations were initialized with parameter values as shown in table 3.4. For each individual parameter setting 100 simulation runs were executed.

The three parameters that were varied during experimentation were the *Configuration*, *NumRounds* and the *AdjacencyMatrix*. The configurations were varied to control for the effects of the indirect bias and the guided variation mechanisms. The *AdjacencyMatrix*

⁸The three types of networks used. For each type a concrete adjacency matrix must be provided as input.

Table 3.4: Model parameter values

Name	Value
<i>Configuration</i>	{unbiased, genetic, Lamarckian}
<i>NumAgents</i>	2^{10}
<i>NumGenerations</i>	100
<i>NumRounds</i>	{1, 5, 10, 20, 30}
<i>NumTags</i>	10
<i>NumTraits</i>	10
<i>SuccessPayoff</i>	1
<i>FailurePayoff</i>	-1
<i>AdjacencyMatrix</i>	{complete, random, small-world} ⁸
<i>MutationRate</i>	0.01

parameter was varied to control for different network structures which could hypothetically affect the nature of trait diffusion throughout the population. Finally, the number of rounds per generation was varied to elucidate the effect of the ratio between horizontal and vertical transmission of information in the system. Because time in the model is abstract there is no “natural” length of a generation. It must also be acknowledged that different cultural domains operate at different time scales: some phenomena occur regularly and frequently throughout actors’ lifetimes (for example trading resources with others), while others, such as weddings, or searching for a job, occur infrequently or perhaps only once in a lifetime. The potential effect of the ratio of horizontal to vertical transmission was demonstrated in Santos, et al. (2006) .

Table 3.5: Network parameters

Network type	Parameter	Description	Values
Random	$\hat{k}$	average degree	1, . . . , 5
Small-world	$\hat{k}$	average degree	50, 100, 150, 200
	β	edge rewiring probability	0.01, . . . , 0.1

Complete networks are networks in which every node is linked to every other node. In the case of random and small-world networks further parameters need to be specified. For a random network one must provide the average node degree $\hat{k}$. Then, $\hat{k} = np$, where n is the number of nodes and p is the probability of an edge existing between any given pair of nodes. The Erdős-Rényi algorithm is used to generate the random network (Erdős & Rényi, 1959). The algorithm proceeds by cycling through all possible unordered pairs of nodes and generating an edge between each pair with probability p .

Small-world networks are considered good approximations for many types of social and biological networks (Watts & Strogatz, 1998). Here the Watts-Strogatz algorithm is used to generate them. The algorithm is seeded with the average degree $\hat{k}$ and an edge rewiring probability parameter β . In the first phase the algorithm generates a regular ring lattice of nodes each connected to $\hat{k}$ neighboring nodes, $\hat{k}/2$ on each side. It then cycles through all the edges in the form (n_i, n_j) where $i < j$, and with probability β replaces them with edges of the form (n_i, n_k) where k is taken with uniform probability from all possible values that avoid self-loops and link duplication.

Table 3.5 shows the values of the relevant network parameters that have been tested in conjunction with the other model parameter settings. The values were chosen in an attempt to study those regions that demonstrated the highest rates of sensitivity based on preliminary exploratory simulations.

For each parameter setting, 100 simulation runs were executed. One hundred random seeds were generated and in the case of random and small-world networks 100 different adjacency matrices were generated for each set of network parameters⁹. For comparability purposes the same random seed was used for the k -th run for each parameter setting. In an analogous fashion, the same adjacency matrix A_k was used for the k -th run for each of the parameter settings for a given network type.

In order to perform thorough analysis a number of measures were calculated and collected as output from the simulations. The foremost task is to track the distribution of the cultural

⁹There is only one possible adjacency matrix of a complete graph on n vertices.

trait variants and the tags in the agent population over time. To express the nature of the distribution in relevant terms with a single quantity, the population skewness is calculated. The skewness γ_1 of a random variable X is defined as follows:

$$\gamma_1 = \frac{\mu_3}{\sigma^3}. \quad (3.2)$$

Here μ_3 is the third central moment of the distribution and σ is the standard deviation. Because the tags and trait variants are expressed in terms of categorical values the skewness is measured on the distribution of counts of the respective values¹⁰. A high skewness thus indicates that a small number of variants, perhaps even a single one, dominate the population. A low skewness value, on the other hand, indicates that the variants are more equally spread across the population (i.e. a more culturally diverse society).

The skewness of the agents' fitness distribution over time is also measured. This is an important indicator of the level of selective pressure in the evolutionary algorithm (Eiben & Smith, 2007, p. 59). In the case of fitness, the skewness is calculated directly on the (cardinal) values.

To track whether certain tags become associated with specific trait variants over time a measure of tag “entropy” is calculated. As is the case with Shannon’s Entropy (Shannon, 1948), upon which it is based, it too measures the (un)predictability of specific states. The resulting quantity expresses how well the tag “alphabet” encodes different trait variants. It is defined as follows:

$$\overline{E} = \frac{1}{NumTags} \sum_{j=1}^{NumTags} \frac{E_j}{n_j}, \quad \text{where} \quad E_j = - \sum_{i=1}^{NumTraits} p_i \ln p_i. \quad (3.3)$$

¹⁰Note that because the number of agents and the number of tag types is constant throughout the simulation, the sum of the counts as well as the mean of the counts is also constant. Because of this we can compute an upper bound on the skewness. Assuming we have n agents, the most skewed distribution occurs when one tag is present n times in the population, while the other ones occur zero times. Calculating the moments in equation 3.2 with these values we obtain a skewness of $\gamma_1 = 2.7$ at 2^{10} agents. The same observations apply, of course, to the cultural trait variants.

Here i iterates over the set of trait variants, p_i is the probability of encountering the i -th variant in an agent possessing the tag j , E_j is the entropy of the j -th tag, and finally the size of each tag sub-population, n_j , is used to normalize the values and obtain the average metric entropy $\overline{E}$.

Finally, to explore the degree of clustering of sub-populations defined by specific tags and cultural traits, the modularity of the agent networks is measured in terms of both attributes. Network modularity measures how “neatly” the network decomposes into communities defined on the possession of a shared attribute (Newman, 2006) and is defined as:

$$Q = \frac{1}{2m} \sum_{vw} \left[A_{vw} - \frac{k_v k_w}{2m} \right] \delta(c_v, c_w). \quad (3.4)$$

Here m is the edge count of the network, v, w are nodes, A refers to the adjacency matrix, k_v and k_w are the degrees of the nodes, c_v and c_w are the attributes of the nodes, and δ refers to the Kronecker delta function. Modularity is measured in random and small-world networks only, as it makes little sense to report it in complete networks.

3.8 Results

Figure 3.2 shows simulation results for complete networks. One can observe that there is no discernible trend in the skewness of tag distributions, neither as a function of the transmission ratio, nor as a function of the model configuration. The trait skewness is shown to rise with the transmission ratio, although this increase is not substantial in relative terms. In any case, the distinctive features are a fairly low skewness of the tag distribution throughout, and a very high skewness of the trait distribution which is achieved relatively early on in the simulations. In fact, in most simulations a single trait variant comes to dominate the population after the first few generations, with occasional mutations providing for the small number of other variants. This renders the tag entropy measure much less informative. Unsurprisingly, it remains very low in all cases and it is negatively correlated

with trait skewness (see part (c) of figure 3.2). This owes to the fact that *any* encoding is relatively noiseless when one character becomes predominant in the input stream (i.e. every tag can encode for the same dominant trait variant with relative success). Part (d) shows the total fitness skewness values in different configurations. One may observe that the fitness distributions are for the most part negatively skewed. This may hint at an undesirable lack of selective pressure, as most agents are tightly clustered around average fitness levels with a minority of them trailing behind. Such a fitness distribution may contribute to effects related to the phenomenon known as genetic drift (see section 3.9 for details). This could subsequently explain the rapid convergence to a single trait variant.

Parts (e) and (f) show the most significant differences between the biased and the unbiased configurations. By design, there are no rejected interactions by the agents in the unbiased configuration (since they interact randomly, and thus there is basis on which to reject interactions). On the contrary, in both of the biased configurations the agents reject a significant amount of interaction proposals over the course of the simulations. This is especially pronounced in simulations with shorter generations. This shows that agents are in fact “picky” about the tags, and evolve strong preferences for interacting with agents possessing only certain kinds of tags. Part (f) shows the share of the most common tag type as a function of time for one particular parameter setting. Here it can be seen that the evolved biases towards specific tags coupled with the convergence in the cultural trait result in a single dominant tag emerging, especially in the Lamarckian configuration. However, as the trait population eventually settles on a single variant, selective pressure ceases to act on the tags, and the tag population diversifies once again due to mutation. This is in contrast with the unbiased configuration in which selective pressure never acts on the tags, and thus their distribution remains largely undisturbed.

Next, I turn to the random network configurations. Figures 3.3 and 3.4 show simulation statistics as a function of average degree $\hat{k}$ and the transmission ratio H/V . When $\hat{k}$ is high, the model demonstrates characteristics similar to the complete network configurations. This includes the high trait skewness levels throughout, as well as high tag skewness levels in the

biased configurations coupled with low skewness in the unbiased version. The tag entropy also remains low in all cases, for the same reasons given above. Here I introduce modularity results for the first time: One can note that the values are also very low for high $\hat{k}$. As is the case with tag entropy, this too is a direct consequence of the high skewness levels. When a single trait variant or tag dominates almost the entire population, the modularity equation (3.4) is reduced to a simple difference between the expected and the actual edges in the graph. Because the networks are random these differences are usually close to zero.

As network density decreases along with $\hat{k}$, so do the skewness levels of both the tag and the trait distributions regardless of the type of bias (or lack thereof). On the other hand, tag entropy never drops in this scenario. Finally, the tag and trait modularity values are increased. These results suggest that while on a local level agents tend to share the same characteristics as their neighbors, there is no clear association between specific tags and trait variants across the whole population.

Although all three types of bias result in more or less similar outcomes, differences remain between the configurations in terms of fitness distribution and interaction frequency (as shown in figure 3.3 g-l).

Finally, I shift the focus to the small-world configurations. Figures 3.5 and 3.6 show the simulation results as a function of the rewiring probability β , and average degree $\hat{k}$. The skewness of the trait distribution shows a clear trend that is once again independent of bias configuration: the distributions remain diverse in regions of low average degree and low edge rewiring probability. As network density increases and the networks become more random in their nature (as a result of the rise in edge rewiring), the distribution becomes highly skewed as previously observed. Identical trends hold for trait and tag modularity measures. Thus far, the outcomes are qualitatively aligned with those from the random network configuration. The main difference, however, can be observed in the tag entropy levels. While the unbiased and genetic configurations, in line with previous results, only experience a decrease in entropy in the “straight-forward” cases where trait distribution is highly skewed, the Lamarckian configuration demonstrates fairly low entropy values no

matter how diverse the trait variant population remains. Furthermore, figures 3.7 and 3.8 show the small-world configuration results as a function of rewiring probability β and the transmission ratio H/V . Here one may observe that the sustained decrease in entropy in the Lamarckian configuration is also dependent upon longer generations (i.e. higher transmission ratios). However, apart from the nature of the tag distribution in the case of the Lamarckian configuration, the transmission ratio does not produce any significant variation in model behavior.

3.9 Discussion

The aim of the model was to show the cultural dynamics in networks of agent interaction driven by the mechanisms of indirectly biased transmission and guided variation.

Overall the results show that the social network structure of the agent population plays a significant role in the outcome of cultural evolution. This is especially true in the case of complete and random networks, for which the tracked measures display little sensitivity to the particular bias implementation or to the transmission ratio. There are exceptions to this generalization, for example the trait distribution skewness and tag entropy levels show some, if small, trends with respect to the transmission ratio. One explanation for these changes is the fact that longer generations with more activity will tend to produce more negatively skewed fitness distributions which only exacerbate the effects of genetic drift in trait variants (the tag entropy is then a direct effect of the peaking trait distribution skewness).

Genetic drift is the change in the frequency of a gene variant in a population due to random sampling of genotypes (Masel, 2011). If a random sample contains a disproportionate frequency of a specific variant, this irregularity will only become more pronounced in the long run when sampling iteratively, until eventually the variant takes over the entire population. It should be noted, that when I invoke genetic drift here as a cause for population take-over, what I have in mind is not pure random genetic drift, which is entirely selectively neutral. Rather, the invocation implies that small differences in fitness levels give an initial boost in frequency to a particular variant which is then exacerbated further by random drift in

conjunction with relatively small population sizes. Indeed, pure random drift would not be able to achieve the high levels of skewness alone on the observed time scale. This is demonstrated by the lasting diversity of the tag distribution under the unbiased configuration, which is in fact unaffected by selection biases.

Apart from the minor points sketched out above, network structure, and particularly network density is the main *explanandum*. As the networks grow more dense one may observe an increase in skewness of the trait variant distribution. If small irregularities in trait variant frequency appear locally, they are able to spread faster in a more denser network. Thus, the effects of genetic drift once again become more prominent. Furthermore, in configurations in which the selection of trait variants is biased by possession of tags, the tag distribution becomes very skewed as well. This hints at an association between tags and trait variants existing already during the phase leading up to the convergence to a particular trait variant. As a matter of fact we cannot attribute the skew in the tag distribution to a similar but independent random drift effect, because we see no such convergence in the unbiased configuration on the observed time scale. The variation in fitness distributions and interaction frequency also indicates that significantly different processes may be at play in the unbiased and biased configurations. However, none of this translates into differences in the cultural makeup of the agent populations. The modularity plots show that regardless of bias implementation, the model achieves clustering of tags and trait variants only in sparse networks, which is a relatively “easier” task than in denser networks, as small changes can have a dramatic effect. Similarly, in none of the model configurations, whether biased or unbiased, is the system able to solve the problem of encoding trait variants via tags when the trait variant distribution is diverse.

I note a few similar trends in the small-world configurations. First, the denser the networks are, the more skewed the trait distributions become which in turn promotes low entropy and low modularity. The same can be said for increasing the edge rewiring probability used in generating the small-world networks which essentially makes them more akin to random networks. However, there are two notable departures from previous dynamics

in the small-world case: (1) The agent networks are highly modular both in terms of tags and trait variants even if the average degree is one order of magnitude higher than in the random case and (2) low levels of entropy are maintained even when the trait population is diverse in both biased configurations. Both of these phenomena are made possible by the small-world structure. The individual clusters or “worlds” evolve as in previous networks: the trait population converges on a single variant and a single tag emerges with it. But because the clusters are partially isolated from others, links between them being sparse, as dictated by their small-world nature, they are able to evolve along independent trajectories. This results in a mosaic of communities, each defined by possibly different pairs of prevailing tags and trait variants.

3.10 Conclusion

I have developed an agent-based model of cultural evolution. I carried out a large number of experimental simulations to test the hypothesis that indirectly biased cultural transmission together with guided variation contribute to the emergence of social clusters marked by the possession of specific pairs of cultural traits and external markers. Our simulations have shown some promise for this hypothesis as I was able to generate such clusters under certain parameter regimes. I have found that the chosen network structure and the ratio between horizontal and vertical transmission both play an important role in the resulting distributions of cultural trait variants and external markers. However, more work is needed to account for the effect of the social networks that we have used in our simulations.

The model ultimately contributes to the study of cultural evolution by combining the use of evolutionary algorithms and network analyses to study the interplay of hidden cultural traits, observable external markers and learned cultural preferences, where other models have only focused on a smaller subset of these concepts.




Figure 3.2: Simulation statistics for complete networks as a function of the horizontal/vertical transmission ratio (a-e). Each point in parts a-e represents the average of the sum of the values for the given measure at the end each of the 100 generations. Part f shows the average share of the most common tag as a function of time when $H/V = 20$.




Figure 3.3: Simulation statistics for random networks as a function of the horizontal/vertical transmission ratio and average degree for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show tag skewness. Parts (d)-(f) show trait skewness. Parts (g)-(i) show fitness skewness. Parts (j)-(l) show percentage of rejected interactions.




Figure 3.4: Simulation statistics for random networks as a function of the horizontal/vertical transmission ratio and average degree for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show tag entropy. Parts (d)-(f) show tag modularity. Parts (g)-(i) show trait modularity.




Figure 3.5: Simulation statistics for small-world networks as a function of the edge rewiring probability and average degree for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show trait skewness. Parts (d)-(f) show tag skewness. Parts (g)-(i) show fitness skewness. $H/V = 10$.




Figure 3.6: Simulation statistics for small-world networks as a function of the edge rewiring probability and average degree for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show tag entropy. Parts (d)-(f) show trait modularity. Parts (g)-(i) show tag modularity. $H/V = 10$.




Figure 3.7: Simulation statistics for small-world networks as a function of the edge rewiring probability and transmission ratio for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show trait skewness. Parts (d)-(f) show tag skewness. Parts (g)-(i) show fitness skewness. $\hat{k} = 100$.




Figure 3.8: Simulation statistics for small-world networks as a function of the edge rewiring probability and transmission for the unbiased (left), genetic (center), and Lamarckian configurations (right). Each point represents the average of the sum of the values for the given measure at the end each of the 100 generations. Images show linearly interpolated values across the entire range of tested parameter values. Parts (a)-(c) show tag entropy. Parts (d)-(f) show trait modularity. Parts (g)-(i) show tag modularity. $\hat{k} = 100$.

Chapter 4: Modeling the Co-Evolution of Culture, Signs and Network Structure

Abstract: Culture structures human conduct in countless aspects. Identifying drivers of cultural dynamics in human societies is important if we wish to understand the sources of collaboration and conflict in social interactions. Here I present and test an evolutionary agent-based model of cultural evolution that rests on interplay of unknown cultural traits, directly observable external markers, and the structure of the agents' social networks. I draw from dual-inheritance theory and implement the mechanisms of guided variation and indirectly biased cultural transmission into our model. Crucially, I allow the agents to adjust their social networks based on their cultural preferences. I show that under these assumptions the agent populations evolve into small-world communities of clusters where specific pairs of traits and tags become dominant. I also show that the model is sensitive to the cost of maintaining social ties and the ratio of horizontal and vertical cultural transmission.

4.1 Introduction

Humans are a cultural species, meaning that they can acquire and transmit shared sets of values, knowledge and behaviors during their lifetimes (Boyd & Richerson, 1985). Often the acquired attitudes and behaviors are equivalent to their alternatives from an adaptive point of view. Consider the example of dress codes at receptions. Ultimately it does not matter what the guests wear. The party could be an equal success whether the guests don formal or casual wear. What matters is that the dress code is observed equally by all of the guests. If one guest arrives in a pair of jeans and a T-shirt to a black tie event, the situation might prove equally awkward for both the denim-clad standout as well as the dress code abiding guests. There are two aspects to this problem: one is knowing the dress code, the other is the person's willingness and capacity to abide by it.

Another aspect of culture are sign systems. Languages are the most complex sign systems, but other more rudimentary forms exist as well, such as flags representing nations. Signs can be used to mark the possession of cultural traits that can be difficult to ascertain directly. For example, an accent might be a sign of the speaker's place or class of origin. In fact, the dress code conundrum could be resolved via language, by simply stating it on the invitation. If the invitation is less explicit, one is left to infer the dress code from a variety of available cues, which can act as signs in themselves.

Finally, culture is local in its nature. While their purpose is to generalize, cultural systems vary across physical and social space, forming more or less defined clusters. Thus, while gowns and tuxedos are considered traditional formal wear in the West, this is not necessarily the case in other regions of the world.

In this study I examine the relationship between cultural evolution and the evolution of social network structure. In particular, I investigate the interplay between selectively neutral cultural trait variants, selectively neutral external markers ("tags"), and the maintenance of social ties in environments where collaboration is necessary and coordination of efforts is crucial to success.

I argue that these phenomena are interconnected and have co-evolved over time through

the mechanisms of indirectly biased transmission and guided variation. Transmission of cultural trait variants is defined as indirectly biased when it is driven by preferences for unrelated external markers. Guided variation is the individual in-life adaptation via evaluation of self-generated and self-explored possibilities (Boyd & Richerson, 1985).

Hence, I assume conditions under which it is (a) necessary to coordinate efforts of multiple actors to solve complex problems, (b) costly to maintain meaningful social connections necessary for cooperation, (c) disproportionately difficult for actors to ascertain possession of cultural traits in others directly, (d) possible for actors to direct their behavior based on the possession of observable markers, and I hypothesize that over time such populations will form distinct cultural clusters and meaningful cultural signs will emerge.

To test this hypothesis I build an evolutionary agent-based model under the above assumptions. In the model, every agent possesses a selectively neutral variant of a cultural trait representing equally adept approaches to solving problems in an abstract domain, and a tag which represents the agent’s observable characteristics. The agents are periodically faced with tasks, which can only be carried out by successfully coordinating their efforts with another agent. The agents cannot deduce the specific variants of the trait possessed by others, however they may periodically choose to abandon partners or find new ones.

In the following section I give a brief overview of previous work. Next, in sections 4.3-4.5. I describe the design of the agent-based model in detail. Section 4.6 then comments on the design of the simulation experiments. I present the results of the simulations in section 4.7. Finally, I provide a discussion of the main findings and concluding remarks in section 4.8.

4.2 Background

Many studies have investigated cultural evolution in an agent-based simulation environment (see Bianchi & Squazzoni, 2015) and a number of them have approached this from the perspective of cooperation (e.g. Hales, 2000; Janssen, 2005; Hammond & Axelrod, 2006). The co-evolution of network structure and cooperation has also been explored (Santos et al.,

2006). Most of these studies use the Prisoner’s Dilemma to simulate the interplay between individual and social utility, although others have used the ultimatum game (Shutters, 2009) or the stag hunt and snowdrift games as well (Santos et al., 2006). Moreover, the fact that social networks are malleable on an individual level has been empirically determined (Kossinets & Watts, 2006). I assume some form of cooperation is necessary for the agents’ survival, but I also assume that cooperative behavior is only successful when agents align their strategies, all n of which are equally adept at solving the task at hand. The problem thus becomes an $n \times n$ pure coordination game with n Nash equilibria. The assumption that agents’ strategies are hidden to their interaction partners has been explored in some of the studies, and there have been several models where the connection between hidden traits and tags has been examined (Hales, 2000; Janssen, 2005; Hammond & Axelrod, 2006). However, it is the case in these studies, that the tags are “ready-made” signs, in that agents either recognize them as indicators of group membership (Hales, 2000; Hammond & Axelrod, 2006) or are able to learn a pre-existing relationship between the tag and another trait (Janssen, 2005). I take a different approach and assign the tags randomly and observe whether any signifying quality emerges from the dynamics of the system. The role of indirect bias and guided variation with tags has been studied with mathematical models (Boyd & Richerson, 1985), even on small groups of live subjects (Efferson et al., 2008). Here I add to the study of human cultural dynamics by analyzing their effects via an agent-based simulation on a large social network.

4.3 Overview

In the following sections I give a detailed description of the design of the agent-based model used to test the hypothesis. The description follows the ODD protocol (Grimm et al., 2006).

4.3.1 Purpose

In this chapter I explore the link between evolution of culture and the evolution of social network structure. Particularly, I use an agent-based model to study the interplay between selectively neutral cultural traits, selectively neutral external markers (or ‘tags’ as they are commonly referred to in ABM literature, e.g. Hales, 2000; Janssen, 2005; Hammond & Axelrod, 2006) and the maintenance of costly social ties in an environment which requires cooperation and coordination of efforts. The model is a direct extension of the model described and analyzed in Chapter 3. As in the previous model, agents possess cultural trait variants which are hidden to other agents and which are equivalent to choices in coordination games that the agents play with each other. They also possess tags, which are directly observable by the other agents, and a set of preferences for the tags. However, simulations of the original model demonstrated that the results heavily depend on the chosen network structure of the agent population. Moreover, the original assumption that social networks remain fixed throughout agent’s lifetimes, even for entire generations, is barely realistic in real-world scenarios (for an example one may refer to the evolving co-authorship networks in Barabasi et al., 2002). Therefore, the purpose of this model is to elucidate how the dynamics differ in the case when social networks are allowed to change during simulations, and if it is assumed that the cultural makeup of the agents and their cultural preferences have a direct effect on the deletion and creation of social ties.

4.3.2 Entities and Variables

The model consists of a single type of agents. The agents are defined by their state variables, which are summarized in table 4.1. For the most part, the state variables are identical to those in Chapter 3. Nonetheless, there are three new variables and the *Neighborhood* variable is now dynamic. I comment on these changes in detail:

- *PositiveThreshold* and *NegativeThreshold*: The thresholds refer to the retrieval thresholds in the ACT-R memory model, and thus essentially represent the threshold necessary for recall of particular tags. In the previous model the retrieval threshold for

each tag was simply set to zero for every agent. To eliminate this somewhat arbitrary parameter I now assign each agent its own retrieval threshold. The thresholds are initialized stochastically in a way so that they are on the same scale as the agents' initial base-level activations. Each agent is assigned two different thresholds. The *PositiveThreshold* is used in the creation of new ties, while the *NegativeThreshold* is used in the deletion of ties. The thresholds are inherited by offspring and subject to selection. Thus, they add another dimension to the natural selection mechanism.

- *LastOutcome*: This variable stores a pair of values regarding the agent's last interaction—whether it was a success (1) or not (0), and the index of the agent with whom the interaction took place.
- *Neighborhood*: As in the previous model, the *Neighborhood* is a list of other agents with whom the agent is able to interact. In this case, however, the agents may choose to remove agents or add new new ones to their neighborhood, based on their preferences.

Table 4.1: Agent variables

Name	Domain	Scale	Type
<i>CulturalTrait</i>	Integer	Categorical	Static
<i>Tag</i>	Integer	Categorical	Static
<i>PositiveThreshold</i>	Integer	Cardinal (ratio)	Static
<i>NegativeThreshold</i>	Integer	Cardinal (ratio)	Static
<i>Neighborhood</i>	List of agents	Categorical	Dynamic
<i>Fitness</i>	Integer	Cardinal (ratio)	Dynamic
<i>NumGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>NumBad</i>	List of integers	Cardinal (ratio)	Dynamic
<i>FirstGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>FirstBad</i>	List of integers	Cardinal (ratio)	Dynamic
<i>LastGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>LastBad</i>	List of integers	Cardinal (ratio)	Dynamic
<i>LastOutcome</i>	Ordered pair of integers	Boolean/categorical	Dynamic

For a complete description of the other state variables see Chapter 3. The model parameters are summarized in table 4.2. The one new parameter introduced in this version

Table 4.2: Model Parameters

Name	Domain	Scale
<i>Configuration</i>	Integer	Categorical
<i>NumAgents</i>	Integer	Cardinal (ratio)
<i>NumGenerations</i>	Integer	Cardinal (ratio)
<i>NumRounds</i>	Integer	Cardinal (ratio)
<i>NumTags</i>	Integer	Cardinal (ratio)
<i>NumTraits</i>	Integer	Cardinal (ratio)
<i>SuccessPayoff</i>	Integer	Cardinal (ratio)
<i>FailurePayoff</i>	Integer	Cardinal (ratio)
<i>AdjacencyMatrix</i>	Matrix of Booleans	Categorical (Boolean)
<i>MutationRate</i>	Floating-point number	Cardinal (ratio)
<i>MaintenanceCost</i>	Floating-point number	Cardinal (ratio)

of the model is *MaintenanceCost*. This refers to the cost of maintaining links (i.e. social ties) to other agents over time. The parameter value represents the cost of maintaining a single link per one round. The link cost is constant throughout the entire agent network. The total costs for each agent are subtracted from their fitness at the conclusion of every round. The value represents the unavoidable costs of maintaining functional social relationships with other actors (e.g. gifting economies, as described by Mauss, 2000), so that these “connections” (in both the literal and the figurative sense) are subsequently willing to collaborate when the need arises. Once again a detailed description of the remaining parameters may be found in Chapter 3.

4.3.3 Process Overview and Scheduling

The main flow of the model consists of three types of events. First, agents are randomly activated in turn and they decide whether to cut any links to current neighbors and whether to add links to any new ones. Next, agents are randomly activated and they initiate interactions with their neighbors. Finally after a predetermined number of activations the agents are processed by an evolutionary algorithm, which creates offspring through the recombination and mutation of heritable characteristics of selected parent agents. These

processes are repeated for a number of generations.

Algorithm 6 Top-level Simulation Flow

```

1:  $params \leftarrow \{Configuration, NumAgents, NumGenerations, NumRounds, NumTags,$ 
    $NumTraits, SuccessPayoff, FailurePayoff, AdjacencyMatrix, MutationRate,$ 
    $MaintenanceCost\}$ 
2: procedure CULTUREMODEL( $[params]$ )
3:    $agents \leftarrow initializeAgents([params])$ 
4:    $i \leftarrow 1$ 
5:    $d \leftarrow rowMean(AdjacencyMatrix)$ 
6:   while  $i \leq NumGenerations$  do
7:      $j \leftarrow 1$ 
8:     while  $j \leq NumRounds$  do
9:        $k \leftarrow 1$ 
10:      while  $k \leq NumAgents$  do
11:         $agents \leftarrow rewire(agents_k, agents, d)$ 
12:         $k \leftarrow k + 1$ 
13:      end while
14:       $AdjacencyMatrix \leftarrow connectSingletons(AdjacencyMatrix)$ 
15:       $agents \leftarrow applyMaintenanceCosts(agents)$ 
16:       $k \leftarrow 1$ 
17:      while  $k \leq NumAgents$  do
18:         $agent \leftarrow selectRandomFrom(agents)$ 
19:         $partner \leftarrow selectRandomFrom(agent.Neighborhood)$ 
20:         $agent, partner \leftarrow interact(agent, partner)$ 
21:         $k \leftarrow k + 1$ 
22:      end while
23:       $j \leftarrow j + 1$ 
24:    end while
25:     $offspring \leftarrow createOffspring(agents, [params])$ 
26:     $agents \leftarrow offspring$ 
27:     $i \leftarrow i + 1$ 
28:  end while
29:  return  $agents$ 
30: end procedure

```

Figure 4.1 shows a rough outline of the logic of agent activity in the model. Algorithm 6 lays out these processes in more detail. The agents are first initialized by setting the values of their state variables. The exact mechanics of initialization are described in section 4.5. After the initialization phase, rewiring and interaction procedures are carried out $NumRounds$ times. In each round, every agent is first given a chance to modify its neighborhood. Network evolution is a common occurrence on the individual level of social networks (Kossinets & Watts, 2006). The details of the rewiring mechanism are given in algorithm 7. In the unbiased configuration (see section 4.4 for details) the agents use a simple heuristic to update




Figure 4.1: An outline of the logic of agent activity

their network. The agents only “remember” their last interaction. If that interaction was negative, the agent will cut ties with that neighbor. If the last interaction was positive, the agent, emboldened by its recent success, will create an additional connection to a randomly selected agent. In other configurations the decision regarding which ties to cut and whom to connect to is biased by the possession of certain tags. Our preferences for things are arguably shaped by our previous experiences with them, both positive and negative. From a cognitive perspective, both the frequency and recency of those experiences figures significantly in determining the precise levels of preference. To account for this, each agent first calculates its so-called base-level activation for each of the possible tag variants. The base-level activation τ_T^- for unsuccessful interactions with tag T is calculated as follows:

$$\tau_T^- = \ln \left[\sum_i^n t_i^{-d} \right] \approx \ln \left[t_n^{-0.5} + \frac{2(n-1)}{\sqrt{t_1} + \sqrt{t_n}} \right] \quad (\text{when } d = 1/2) \quad (4.1)$$

Here t_i is the time since the i -th unsuccessful interaction with an agent possessing tag T , while $n = NumBad_T$ is the total number of these experiences, $t_n = LastBad_T$ is the time since the most recent experience and $t_1 = FirstBad_T$ is the time since the first experience. Finally, d is the rate of decay. Due to the expensive nature of the above calculation for large n , I implement a well-known approximation (Anderson & Lebiere, 1998). In line with convention, I use $d = \frac{1}{2}$ (Petrov, 2006). The agent then retrieves those tags whose base-level activation exceed its *NegativeThreshold* value and subsequently cuts ties to any neighbors possessing such tags. Ties are deleted unilaterally—the neighbor has no say in the decision. To create new connections, the agent’s base-level activation of positive experiences is calculated¹. The agent retrieves all tags whose base-level activation exceed its *PositiveThreshold* value. After this a random subset of agents possessing any one of these tags is selected. The size of this subset is given so that the number of candidates together with the size of the agent’s current neighborhood does not exceed the average degree of the

¹this is done in the same way as in equation 4.1, only instead of *NumGood*, *FirstGood* and *LastGood* the values for *NumBad*, *FirstBad* and *LastBad* are used

AdjacencyMatrix used during initialization². Unlike link deletion, creation is a mutual act. Thus, every candidate will perform the same base-level activation calculation as the agent who proposed the connection. If the activation of the agent's tag is above the candidate's *PositiveThreshold* value, the candidate will accept the proposal and the connection will be created.

Algorithm 7 Network Rewiring

```

1: procedure REWIRE(agent, agents, d)
2:   if Configuration = "Unbiased" then
3:     if LastOutcome = 0 then
4:       neighbor  $\leftarrow$  agent.LastOutcome[2]
5:       agent.Neighborhood  $\leftarrow$  agent.Neighborhood  $\setminus$  neighbor
6:       neighbor.Neighborhood  $\leftarrow$  neighbor.Neighborhood  $\setminus$  agent
7:     else if LastOutcome = 1  $\wedge$  agent.Neighborhood < d then
8:       newNeighbor  $\leftarrow$  selectRandomFrom(agents)
9:       agent.Neighborhood  $\leftarrow$  agent.Neighborhood  $\cup$  newNeighbor  $\triangleright$  while making
       sure this is not a self-loop or duplicate link
10:      newNeighbor.Neighborhood  $\leftarrow$  newNeighbor.Neighborhood  $\cup$  agent
11:    end if
12:  else
13:    negativeBLA  $\leftarrow$  calculateBLA(agent, "bad")
14:    retrievedTags  $\leftarrow$  {tag | negativeBLAtag < agent.NegativeThreshold}
15:    agent.Neighborhood  $\leftarrow$  agent.Neighborhood  $\setminus$  {i | i  $\in$  agent.Neighborhood  $\wedge$  i.Tag  $\in$ 
       retrievedTags}
16:    positiveBLA  $\leftarrow$  calculateBLA(agent, "good")
17:    retrievedTags  $\leftarrow$  {tag | positiveBLAtag > agent.PositiveThreshold}
18:    candidates  $\leftarrow$  {i | i  $\in$  agents  $\wedge$  i.Tag  $\in$  retrievedTags}
19:    candidates = randomSample(candidates, size = d - |agent.Neighborhood|)
20:    for candidate  $\in$  potentialNeighbors do
21:      candidatePositiveBLA  $\leftarrow$  calculateBLA(candidate, "good")
22:      if candidatePositiveBLAagent.Tag > candidate.PositiveThreshold then
23:        agent.Neighborhood  $\leftarrow$  agent.Neighborhood  $\cup$  candidate
24:        candidate.Neighborhood  $\leftarrow$  candidate.Neighborhood  $\cup$  agent
25:      end if
26:    end for
27:  end if
28:  return agents
29: end procedure

```

Once the rewiring stage is finished the population enters the interaction stage. This consists of *NumAgents* activation events, during which a single agent is activated. All of the agents have an equal probability of being activated. Thus, for each agent the expected value

²This is done to prevent bloating of the network in terms of density

of activations is one per round, although this will vary due to the stochasticity involved. A single interaction between two agent proceeds as outlined in algorithm 8. The two agents compare their cultural trait variants. If the variants agree both agents receive the *SuccessPayoff*; if the variants differ they both receive the *FailurePayoff*. Also, under the Lamarckian model configuration the agents update either their *LastGood* and *NumGood* fields, or *LastBad* and *NumBad* fields, based on whether the interaction was a success or a failure.

Algorithm 8 Interaction

```

1: procedure INTERACT(agent, partner, Configuration, SuccessPayoff, FailurePayoff)
2:   if agent.culturalTrait = partner.CulturalTrait then
3:     agent.Fitness  $\leftarrow$  agent.Fitness + SuccessPayoff
4:     partner.Fitness  $\leftarrow$  partner.Fitness + SuccessPayoff
5:     if Configuration = “Lamarckian” then
6:       agent.LastGood[partner.Tag]  $\leftarrow$  0
7:       partner.LastGood[agent.Tag]  $\leftarrow$  0
8:       agent.numGood[partner.Tag]  $\leftarrow$  agent.numGood[partner.Tag] + 1
9:       partner.numGood[agent.Tag]  $\leftarrow$  partner.numGood[agent.Tag] + 1
10:    end if
11:  else
12:    agent.Fitness  $\leftarrow$  agent.Fitness + FailurePayoff
13:    partner.Fitness  $\leftarrow$  partner.Fitness + FailurePayoff
14:    if Configuration = “Lamarckian” then
15:      agent.LastBad[partner.Tag]  $\leftarrow$  0
16:      partner.LastBad[agent.Tag]  $\leftarrow$  0
17:      agent.numBad[partner.Tag]  $\leftarrow$  agent.numBad[partner.Tag] + 1
18:      partner.numBad[agent.Tag]  $\leftarrow$  partner.numBad[agent.Tag] + 1
19:    end if
20:  end if
21:  return agent, partner
22: end procedure

```

At the end of every generation the evolutionary algorithm is used to re-seed the model with a new population of offspring agents (as shown in algorithm 9). In the EA each offspring is created by the recombination of two selected parent agents. Reproduction is local and the existing network structure is preserved during this phase. Specifically, for each node in the agent network one offspring will be created. Upon completion of the EA phase this offspring will replace the agent currently occupying said node. For each node, the parents are selected from a set of candidates which consists of the current occupant of the node and all of its

immediate neighbors. Each parent is selected via a tournament of size 3: Three agents are chosen from the candidate set with uniform probability and the fittest of the three becomes the parent³.

Once the selection phase is carried out the offspring is produced via unbiased uniform crossover (see De Jong, 2005, p. 64-65) acting on the heritable characteristics of the parents. The heritable characteristics include all of the fields listed in table 4.1, apart from *Fitness* which is reset to zero for all offspring, and the *LastOutcome* field which is left empty at the beginning of the generation. Once the offspring are generated they undergo mutation. Mutation is applied to each field individually with a 1% probability. Mutation is carried out element-wise if the fields are lists. The mutation itself consists of a “small” change in the value of the field. Cardinal variables are modified by introducing small Gaussian perturbations to the current value. Categorical variables are modified by resetting the current value to another value in the allowed range with uniform probability.

4.4 Design Concepts

- *Sensing*: The agents can only sense their direct neighborhood. Their neighborhoods may, however, change over time. Until an interaction between an agent and its neighbor takes place the agent is only able to sense the neighbor’s tag.
- *Prediction*: The agents have an interest to surround themselves with neighbors who possess the same trait variant as them, otherwise interactions will result in decreasing fitness. However, because the agents cannot sense the others’ trait variants directly, they must decide on their preferred neighbors based on other criteria. These criteria vary based on the model configuration.

In the Lamarckian configuration of the model, the agents use the directly observable tags and their past experiences to formulate predictions regarding correlations between the possession of a specific tags and specific trait variants. This prediction is governed

³If there are only two candidates, then no tournament takes place and the two automatically become parents. If there is only one candidate then the offspring will be an exact copy of that candidate (barring any subsequent mutations).

Algorithm 9 Evolutionary Algorithm

```
1: procedure CREATEOFFSPRING(agents, TournamentSize, MutationRate)
2:    $i \leftarrow 1$ 
3:   offspring  $\leftarrow \emptyset$ 
4:   while  $i \leq \text{length}(\text{agents})$  do
5:     parents  $\leftarrow \emptyset$ 
6:     agent  $\leftarrow \text{agents}[i]$ 
7:     while  $\text{length}(\text{parents}) \neq 2$  do
8:       candidates  $\leftarrow \text{agent}$ 
9:        $j \leftarrow 1$ 
10:      while  $j \leq \text{TournamentSize}$  do
11:        candidates  $\leftarrow \text{candidates} \cup \text{selectRandomFrom}(\text{agent.Neighborhood})$  ▷
12:        including checks for ensuring non-duplicate candidates
13:         $j \leftarrow j + 1$ 
14:      end while
15:      bestCandidate  $\leftarrow \text{tournament}(\text{candidates})$ 
16:      parents  $\leftarrow \text{parents} \cup \text{bestCandidate}$ 
17:    end while
18:    thisOffspring  $\leftarrow \text{uniformCrossover}(\text{parents})$ 
19:    thisOffspring  $\leftarrow \text{mutation}(\text{thisOffspring}, \text{MutationRate})$ 
20:    offspring  $\leftarrow \text{offspring} \cup \text{thisOffspring}$ 
21:     $i \leftarrow i + 1$ 
22:  end while
23:  return offspring
24: end procedure
```

by equation 4.1. The relationship and the derived estimate represent the process of reinforcement learning and is adopted from the ACT-R cognitive framework (Anderson & Lebiere, 1998). The base-level activation represents the ability to retrieve the associated “chunk” from memory when given a cue (such as the question “which tags do I associate with my own trait variant?”, or its converse). A chunk, in this case a particular tag, is retrieved when its base-level activation surpasses a certain threshold, which results in the agent perceiving it as associated with the given trait variant. The quantity is very sensitive to recent experiences, while the importance of older experiences progressively decays with time. A new positive experience with a given tag results in a spike in its base-level activation. A lack of experience over an extended period of time will result in the gradual decrease in the tag’s base-level activation (see figure 4.2 for an illustration). It should be noted that the offspring inherit their parents’ values in all of the fields that keep track of past experiences. A literal interpretation of this would suggest that the offspring inherit their ancestors’ memories. However, the correct interpretation in this case is that the parents’ attitudes and associations are being inherited. These can in fact become internalized by offspring through processes of socialization (Berger & Luckmann, 1966). The inputs being inherited and stored do not matter themselves, what matters is that they produce the same result in equation 4.1.

To control for the effects of guided variation and indirect bias I introduce two other model configurations. In the “genetic” configuration, I eliminate the plasticity of the tag associations. Thus, agents’ preferences for the individual tags are set at birth and they do not change throughout their lives as a result of interaction. The preferences can only be modified from one generation to another through forces of recombination and mutation.

Finally I devise an unbiased configuration, one in which the agents do not take tags into consideration at all. However, the agents still display some ability to maximize their interaction utility, as described in section 4.3.3. This version of agent behavior




Figure 4.2: Example of a base-level activation for some input over time. The spikes in the chart coincide with instances of processing the input, followed by gradual decay.

was chosen because it did not involve any tag-related bias, nor any complex predictive calculations, while still allowing for some network rewiring dynamics.

- *Adaptation:* The agent population as a whole adapts on an inter-generational scale through evolutionary processes. Parent selection is biased by fitness and therefore under-performing agents, e.g. those who possess a locally rare trait variant, an unpopular tag, or whose learned predictions do not align well with their current environment, do not get a chance to pass on their traits in favor of fitter individuals.

Moreover, on the individual and intra-generational level the agents adapt by choosing whom to “neighbor” with. A creation of a link between two agents represents the initiation of a relationship which allows either of them to demand collaboration from the other, knowing that the other will do the same in the future. For the link to endure, the agents must periodically signal their commitment to the relationship. This is what the maintenance cost represents: temporal, emotional, even financial investments such as small talk, phatic communication, gift exchange etc. Link deletion,

on the other hand, signifies a divestment from a particular social relationship.

In the Lamarckian configuration, the agents adapt individually by updating their predictions after every interaction, incorporating the most recent experiences into their calculations. This sort of in-life adaptation based on self-generated and self-explored possibilities is essentially what Boyd and Richerson (1985) define as *guided variation*.

- *Interaction*: An interaction between two agents consists of the two comparing their cultural trait variants. If their variants are equal, the interaction is successful. If the variants differ, it is unsuccessful. Thus the interaction is essentially an $n \times n$ pure coordination game with n Nash equilibria (where n is the number of possible trait variants). The trait variants represent cultural preferences in a given domain, and we assume that when two actors collaborate in an attempt to solve a problem in this domain, the probability of a successful outcome is increased when their preferences are aligned. Moreover, I assume that it does not matter on which approach they coordinate as they are all equally equipped to solve the problem. For example when two individuals compare measurements it does not matter whether they are in inches or centimeters as long as both use the same scale. Miscommunication could, however, prove disastrous, as it did, for example, in the case of the Mars Climate Orbiter probe crash (NASA, 1999). Similarly, if two people go out for dinner together the evening will be more enjoyable in the end, if their expectations of who pays for whom are the same.

Finally, it should be noted, that unlike in chapter 3, the agents do not directly choose their interaction partners. Instead, the agents interact with a randomly selected neighbor. Thus, the strategic component shifts from choosing partners from a fixed neighborhood, to choosing the neighbors themselves.

- *Stochasticity*: The model is affected by stochasticity in two ways. It affects agent decision-making: the choice of interaction partner from among neighbors etc. Here stochasticity stands in for unknown parameters beyond the scope of the model. Second, stochasticity is present during the initialization and EA phases. In this case stochasticity

is a tool for boosting and maintaining diversity in the agent population.

- *Emergence*: The purpose of the model is to track how the distributions of external markers and cultural trait variants, and the dynamic social network structure are affected by choices made by boundedly-rational agents on the individual level.
- *Observations*: The model tracks the distributions of the tags and the cultural trait variants and their change in time, as well as a number of measures derived from these distributions.

4.5 Initialization and Inputs

The model is fully determined by giving the values of model parameters shown in table 4.2 as input. The agent population is initialized at the onset of the simulation. Every agent has its *Fitness* set to zero, and its neighborhood initially set according to the provided *AdjacencyMatrix*. Furthermore, every agent is assigned a single tag, denoted by an integer value, chosen with uniform probability from the set of integers from 1 through *NumTags*. Equivalently, every agent is assigned one cultural trait variant, chosen with uniform probability from the set of integers from 1 to *NumTraits*. It should be noted, that for every agent the choice of tag is in no way correlated to the choice of cultural trait variant during the initial assignment. The values of *NumBad* and *NumGood* are assigned with uniform probability from the set of integers from 1 through *NumRounds*. The variables *FirstGood*, *FirstBad*, *LastGood* and *LastBad* are initialized with values chosen with uniform probability from the set of integers from 1 through *NumRounds*NumAgents*. This has the effect of creating random fictional memories for the newly created agents. The variable ranges are chosen so that the memories are on the scale of an agent's lifetime.

4.6 Verification and Validation

The model was developed and implemented in the MATLAB computing environment. The model was verified by way of detailed and comprehensive code walkthroughs and debugging.

Validation of the model was undertaken in the sense that distribution of significant network characteristics of the evolved social networks were compared to generally accepted attributes of empirical social networks, such as high clustering coefficients, short path lengths, high levels of modularity, etc. Validation of the code in the sense of direct quantitative comparisons to relevant sources of empirical data was not performed at this point, as this is the focus of Chapters 5 and 6. The source code together with documentation can be accessed in the OpenABM model database at the following link: <https://www.openabm.org/model/5815>.

4.7 Experimental Design

I carry out a large number of simulation experiments on the model which is implemented as described in previous section. The simulations were initialized with parameter values as shown in table 6.3. For each parameter setting 100 simulations were run.

Table 4.3: Model parameter values

Name	Value
<i>Configuration</i>	{unbiased, genetic, Lamarckian}
<i>NumAgents</i>	2^{10}
<i>NumGenerations</i>	100
<i>NumRounds</i>	{1, 5, 10, 20}
<i>NumTags</i>	10
<i>NumTraits</i>	10
<i>SuccessPayoff</i>	1
<i>FailurePayoff</i>	-1
<i>AdjacencyMatrix</i>	$\{\hat{k} \hat{k} = 2^i, i = 2, \dots, 6\}^4$
<i>MaintenanceCost</i>	$\{c c = \frac{2^i}{100}, i = 0, \dots, 9\}$
<i>MutationRate</i>	0.01

I varied the *Configuration*, *NumRounds*, *AdjacencyMatrix* and *MaintenanceCost* parameters. The configurations were varied to control for effects of the indirect bias and guided variation mechanisms, which I set out to study. The average degree of the

⁴The average degree of the network. For each value of $\hat{k}$ a specific adjacency matrix must be provided.

AdjacencyMatrix was varied to control for a diverse range of possible initial conditions—I wanted to see whether populations that are densely connected evolve into states different from those that are connected only sparsely. For each value of $\hat{k}$, I created 100 different adjacency matrices using the Erdős-Rényi algorithm. The algorithm assumes that $\hat{k} = np$, where n is the number of nodes and p is the probability of an edge existing between any given pair of nodes. The network is then generated by iterating over all unordered pairs of nodes and creating an edge between each pair with probability p . I varied the number of rounds in one generation to test the model’s sensitivity to the ratio of horizontal and vertical transmission. The reasoning behind this is the abstract definition of time in the model coupled with the fact that cultural domains may operate at significantly different time scales. For example people tend go out for dinner with their friends many times during their lives, but most will only get married once. The frequency of the cultural practice certainly adds to what is at stake. I make a similar argument for varying the *MaintenanceCost*. Different types of cooperation require a different type of relationship. For example a person might ask even some of their casual friends to help them move, and they might accept. On the other hand when asking for a personal loan, people usually turn to family or their closest friends—relationships in which they tend to be heavily personally invested. I executed 100 simulations with different random seeds for each parameter setting.

A number of measures were calculated from the simulation outputs for purposes of analysis. The population skewness of the trait variant distribution and the tag distribution were calculated, giving us a single quantity that summarizes the nature of the distributions. Skewness γ_1 of a random variable is defined as:

$$\gamma_1 = \frac{\mu_3}{\sigma^3}. \quad (4.2)$$

Here μ_3 is the third central moment of the distribution and σ is the standard deviation. Because the tags and trait variants are categorical variables, skewness is measured on the distribution of counts of the values. Low skewness suggests that each variant is represented

more or less equally in the population, while high skewness suggest that one or perhaps a few variants dominate the population.

I also measure tag entropy (based on information entropy) to give a sense of how well the tag types align with the trait variant population. I define it in the following way:

$$\overline{E} = \frac{1}{NumTags} \sum_{j=1}^{NumTags} \frac{E_j}{n_j}, \quad \text{where} \quad E_j = - \sum_{i=1}^{NumTraits} p_i \ln p_i. \quad (4.3)$$

Here i iterates over the set of trait variants, p_i is the probability of encountering the i -th variant in an agent possessing the tag j , E_j is the entropy of the j -th tag. The size of each tag sub-population, n_j , is used to normalize the values, which gives us the average metric entropy $\overline{E}$.

Furthermore, to express the degree of clustering of tag types and trait variants on the networks I measure network modularity. The quantity represents how neatly the graph decomposes into communities defined by a common attribute (Newman, 2006). It is defined as follows:

$$Q = \frac{1}{2m} \sum_{vw} \left[A_{v,w} - \frac{k_v k_w}{2m} \right] \delta(c_v, c_w). \quad (4.4)$$

Here m is the edge count of the network, v, w are nodes, A refers to the adjacency matrix, k_v and k_w are the degrees of the nodes, c_v and c_w are the attributes of the nodes, and δ refers to the Kronecker delta function.

Finally, to track the evolution of the network structure itself I measure two important graph metrics. One of them is the average local clustering coefficient. It represents how tightly-knit and inter-connected the network is. It is defined in the following way:

$$\overline{C} = \sum_{i \in V} C_i \quad \text{where} \quad C_i = \frac{\text{number of closed triplets containing } i}{\text{number of triplets containing } i}. \quad (4.5)$$

Here i is a single node in the network. The second network measure is the average shortest path length $\bar{L}$. It simply measures the average number of edges needed to travel between every pair of nodes i, j when traveling along the shortest route. Pairs of nodes that are mutually unreachable are excluded from the calculation. Ultimately, I measure the clustering coefficients and the path lengths as their ratio to the clustering coefficient and average path length achieved on a regular ring lattice of the same size and same average degree. This essentially measures the “small-worldness” of the network. Small-world networks are defined by clustering coefficients relatively similar to those of the regular lattices, while having significantly shorter path lengths (Watts & Strogatz, 1998).

4.8 Results

Figure 4.3 shows the differences in evolutionary dynamics using the three different model configurations with a single choice of parameters. Here, I initialize the agent population on a random network with average degree $\hat{k} = 8$, link maintenance cost $c = 0.16$ and 10 rounds of interaction per generation. Figures 4.3(a),(b) illustrate the distribution of tags and trait variants over time. I observe that the unbiased configuration results in the most heavily skewed trait distribution, actually drifting away towards a single variant. The tag population is fairly similar in all cases, remaining relatively equally distributed. This is to be expected in the unbiased configuration, because there is no selective pressure exerted on the tags. However, the lack of skew in tag distribution in the other configurations requires further analysis.

Figure 4.3(c) shows tag entropy values for the different configurations. While the Lamarckian configuration is able to evolve into a state with low entropy, meaning that each tag more or less faithfully encodes for a specific cultural trait variant, the unbiased configuration shows even lower values of entropy. This is surprising at first glance, as the tags are an afterthought in the unbiased configuration, possessing no real significance. However, it is important to note, that this is a result of the convergence in the trait variant distribution.

Once there is only one trait variant present it can be encoded equally well by any of the tags.

I thus further analyze the interplay of tags and cultural trait variants by measuring the modularity of the agent networks in terms of both attributes. Figures 4.3(d),(e) show that in the biased configurations I observe the emergence of communities that are well defined by both their tags and the trait variants that they share. Such communities do not appear in the unbiased configuration.

To make the evolution of the network clearer I inspected the changes in the social networks of the agents visually over the course of the simulations. Figure 4.4 shows illustrative examples of the network evolution. I observe that in the early stages, the networks become very sparse in both cases which causes social interaction to decrease dramatically. The behavior in the two configurations diverges as the simulation progresses, and the differences become clear. The unbiased configuration results in a single giant component; meanwhile the Lamarckian configuration evolves a network organized into clear communities marked by possession of distinct tags and trait variants. Furthermore, each tag locally strongly correlates with a single trait variant (even if the associations are not perfect). The presence of clearly defined clusters suggests the “small-worldness” of these networks. To confirm this, I compute the average clustering coefficients and the average path lengths of the networks. Figures 4.3(f),(g) show these measures, and I observe that while all three configurations show significantly shorter path lengths than regular lattices, the Lamarckian configuration demonstrates significantly higher clustering coefficients than the other two. This indicates a small-world quality of the networks when using the Lamarckian configuration. Meanwhile, the other configurations result in networks that are more random in nature.

Having sufficiently illustrated the temporal dynamics on a single example, I now continue with assessing the sensitivity of these outcomes to the choices of parameters as described in section 4.7. Figure 4.5 shows the dependence of some of the measured quantities on the maintenance cost and the transmission ratio. I observe that the Lamarckian configuration demonstrates higher clustering coefficient values than the other configurations across a




Figure 4.3: Model statistics for $\langle k \rangle = 8, c = 0.16$ and $H/V = 10$.

major part of the parameter space, especially in cases with very low maintenance costs. In similar fashion, the Lamarckian configuration evolves networks with consistently shorter path lengths, while in other configurations the path lengths increase either with generation length (i.e. the genetic configuration) or with maintenance cost (i.e. the unbiased model). Moreover, both of the biased versions display consistently higher rates of tag modularity than the unbiased configuration. Once again, this is especially pronounced in low maintenance cost regimes. A similar narrative applies to trait modularity, although it is worth noting that in high maintenance cost regions the Lamarckian model now performs worse than the unbiased model.


Finally, figure 4.6 shows the sensitivity of the model behavior as a function of the maintenance cost and the average degree of the initial network. I note that in the case of the Lamarckian configuration the clustering coefficient increases with average degree. Similarly, the path length increases with the average degree, this time in all three configurations. On the other hand, the change in average degree does not seem to matter for the modularity values of any of the evolving networks. Neither tag nor trait modularity is significantly




(a)



(b)



(c)



(d)

Figure 4.4: Visualizations of the agent network in the Lamarckian bias (left) and unbiased (right) configurations after 5, 20 and 100 generations. The border color of a node determines its external marker, while the inner color determines its cultural trait variant. $\langle k \rangle = 8, c = 0.16, H/V = 10$.




Figure 4.5: Model statistics as a function of maintenance cost and interaction/transmission ratio (H/V) for the unbiased (left), genetic bias (centre), and Lamarckian bias (right). Parts (a)-(c) show the average local clustering coefficients at the end of runs compared to regular lattices. Parts (d)-(f) show average path lengths at the end of runs compared to regular lattices. Parts (g)-(i) show the sum of tag modularity values at the end of each generation. Parts (j)-(l) show the sum of trait modularity values at the end of each generation. $\hat{k} = 32$.

affected.

4.9 Discussion

The goal of this work was to derive the effect of indirectly biased cultural transmission and guided variation dynamics on the co-evolution of hidden cultural traits, observable external markers and network structure in multi-agent populations.

I have shown that indirectly biased transmission coupled with guided variation can have a significant effect on social network structures in the long run. This is especially true when the two mechanisms work in tandem, that is if the bias acquired through individual learning is inherited by offspring, and when the cost of social tie maintenance is low. In the early stages, we do not observe significant differences in system dynamics. The random assignment of tags and trait variants together with the random placement of agents onto nodes of the network results in a low probability of successful interaction at the onset of the simulation. Thus, in the biased configurations, forces of evolution will tend to select agents that minimize social interaction by cutting their ties as a result of their strict activation thresholds. So far this does not differ from the unbiased configurations where agents simply cut ties with neighbors after every failed interaction. Both of these processes explain the abrupt decrease in network density in the first several generations. The network dynamics diverge after this point. The low network density produces a very localized environment in which adaptations cannot spread far very quickly. Thus, in the biased versions, small connected groups of agents converge on a single trait variant and a single tag, undisturbed by what is happening elsewhere. I argue that the convergence in trait variants is due to drift effects, which become exacerbated in small populations, and most importantly due to kin selection (Hamilton, 1963). As offspring are always born in their parents' vicinity, agents become surrounded by neighbors that are increasingly homogeneous with respect to agent attributes. If, by chance, the dominant phenotype possesses a strong preference for its own tag, the dynamics are only reinforced. When the preference for a particular tag is established early within the group, it grows by attaching to agents possessing the same




Figure 4.6: Model statistics as a function of maintenance cost and average degree of the seed network for the unbiased (left), genetic bias (centre), and Lamarckian bias (right). Parts (a)-(c) show the average local clustering coefficients at the end of runs compared to regular lattices. Parts (d)-(f) show average path lengths at the end of runs compared to regular lattices. Parts (g)-(i) show the sum of tag modularity values at the end of each generation. Parts (j)-(l) show the sum of trait modularity values at the end of each generation. $H/V = 10$.

tag, while drift effects correct for any disturbances in its trait distribution. These processes emerge in parallel within the many disjoint network components, and thus multiple clusters defined by a dominant tag-variant pair emerge. Due to these processes, the tags become imbued with meaning, transforming into cultural signs representing possession of specific variants. I further note that due to the local nature of network coalescence the meaning is also local: the same variant can be represented by different tags in different clusters, and different variants can be represented with identical tags in separate communities.

In the unbiased configuration the network evolves very differently. The lack of distinct tag groups is self-explanatory, but the absence of clustering of trait variants warrants more attention. I argue that the main driver of the network structure in this case are drift effects. In the absence of bias, link creation between pairs of agents is random. If one particular trait variant owns a slight edge in frequency, agents possessing that variant will have a better chance of meeting their counterparts. As this trait group grows so does its advantage, attracting and infecting smaller components more easily.

Additionally, low tie maintenance costs were a significant factor in exacerbating the clustering phenomena. When costs were high the “small-worldness” of the biased networks was diminished and trait variant distributions displayed higher takeover rates. I argue that the increase in costs amplifies the effects of drift at the expense of kin selection. As the maintenance cost becomes high relative to interaction payoffs, mistakes—nearly ubiquitous at the onset—become cruelly punished. This results in a more or less uniform distribution of fitness. The higher costs also push the agent population to become more guarded, keeping their neighborhoods smaller as a result. Both of these facts contribute to higher rates of drift and lower rates of modularity and clustering.

Today the effect of cultural differences is as significant as ever. We see conflicts emerging at the boundaries of otherwise fairly segregated cultural clusters as a result of disjoint sets of belief or values, whether it is the current tension between the Judeo-Christian West and the Muslim world, or the political and cultural divide between so-called red states and blue states in the United States. Meanwhile, in the cores of these clusters individuals become more and

more isolated from outside influences in what amounts to “echo-chambers” (Quattrocioni et al., 2016). As new modes of networking appear, they quickly become pervaded by our own cultural matrices. Does culture factor in transactions between strangers on apps such as *AirBnB*, *Uber* or *Craigslist*? How does large-scale collaboration on open-source software projects emerge? I believe that my model could be applied to such real-world scenarios and validated with the help of empirical data.

The simulations analyzed in this study have provided considerable insight into the effect of indirectly biased transmission and guided variation on the evolution of network and community structure. However, questions worthy of further investigation remain. It is natural to ask how and to what extent does kin selection affect the observed dynamics. It is therefore worthwhile considering studies in which the effects of kin selection would be clearly isolated from other factors. Furthermore, I did not attempt to adjudicate whether the mechanism of bias and the proposed form of individual learning are evolutionarily viable and robust. In the future, it will be important to identify conditions from which bias-driven strategies appearing as random mutations are able to “infect” populations while themselves resisting invasion.

Finally, unlike in other tag-based models, where external markers were essentially pre-fabricated signs of group membership, my model contributes to the line of cultural evolution ABMs by showing that populations can cluster into communities defined by possession of distinct cultural traits and tags, even in the case when tags are assigned randomly, and homophily is not assumed.

Chapter 5: Cultural Aspects of Political Polarization: Simulating U.S. Congress Co-Sponsorship Networks

Abstract: Political polarization is becoming a prominent issue in the United States and across the world. Here, I attempt to explain the root causes of political polarization from an evolutionary, cultural perspective. I adapt and modify my ABM of cultural evolution in an attempt to show the effect of indirectly biased transmission, guided variation and social network structure acting on agents' political ideologies. I use longitudinal congressional data on bill co-sponsorship to validate the model. I show that the hypothesized mechanisms of cultural evolution result in ideological dynamics that are closely aligned with the empirical data.

5.1 Introduction

One does not need to search too hard for evidence of growing concern over increasing political polarization in the United States. Today, the issue is front and center on the pages of newspapers, both national and local, on cable television news programs, even within the increasingly relevant Twittersphere (e.g. Conover et al., 2011; Prior, 2013). The rising polarization of U.S. politics has been a topic of academic research for over three decades, dating back to Poole and Rosenthal's (1984) study, which analyzed the variance of liberal-conservative positions of U.S. politicians between 1959 and 1980. Interest in the issue of polarization and heightened partisanship has only been exacerbated by some of the most recent events in U.S. politics.

The concept of polarization has many operational definitions. Bramson et al. (2016) give examples of several of these: A distribution of a variable can be identified as polarized based on metrics including the range of the distribution, its variance, as well as group-based measures such as overlap (in the distributions over the distinct groups in the population), or a combination thereof. However, all of these definitions rely solely on the distribution(s) of the actors' positions on some pre-defined scale. What these measures do not take into account is the attitudes of the actors towards others holding different positions. Here, I would like to measure both of these aspects. I thus define polarization as both the distance between the positions of a population of actors and as the extent of tolerance towards other positions.

The model rests on a number of assumptions which I now enumerate explicitly:

1. I assume that legislators socialize and interact with their colleagues on a regular basis, and that they are able to initiate social interactions with any other member of the House. These interactions might be political or entirely non-political in their nature. I further assume that maintaining meaningful social connections is costly and that each connection is associated with a considerable expense of time, effort and/or resources.
2. I posit that legislators interact strategically, in that they form preferences for interaction

partners based on past experience.

3. I assume that for a variety of reasons and under a wide range of circumstances it might be difficult or too expensive for a legislator to become aware of a fellow representative's ideological position prior to their interaction taking place.
4. I conjecture that in the absence of knowledge regarding ideological positions of others, legislators will choose interactions based on other known information, notably party affiliation.
5. I assume that legislators will recruit potential co-sponsors solely from the ranks of their active social connections.
6. Finally, I presume that legislators possess a certain degree of tolerance towards ideological positions other than their own, and that they will only co-sponsor legislation that is within their bounds of tolerance. Similarly sponsors will only accept co-sponsors from within their tolerated range.

Based on these assumptions I form hypotheses about the model, and thus by extension about the nature of the social and ideological dynamics in the U.S. House of Representatives:

- H_1 : The model based on the above listed assumptions and implemented as described in the Methods section reproduces the co-sponsorship networks empirically observed in the U.S. House of Representatives over the past four decades. In other words, bill co-sponsorship is driven by the complex inter-dependencies of social network structure, party affiliation, ideology and their evolution in time.
- H_2 : The reliance on party affiliation as a sign of closeness at the expense of underlying ideological considerations leads to increasing polarization, as evidenced by lower levels of tolerance to other opinions and deepening chasms in the ideological positions among legislators.

I also devise and later test three alternative hypotheses:

- A_1 : Legislators do not form social connections (which serve as a necessary foundation for potential political cooperation) with their colleagues based on ideological position or party affiliation, but rather solely on the relationships that they share with others.
- A_2 : Legislators do not form social connections with their colleagues based on party affiliation, rather, they keep track of individual interactions, and disassociate themselves from those with whom they have had negative experiences on the past.
- A_3 : Legislators do not form social connections with their colleagues based on party affiliation, instead they are aware of everyone’s exact ideological positions, and associate themselves only with those with whom they are close from an ideological perspective.

I begin with a brief overview of related work in section 5.2. In section 5.3, I specify the model in full detail. I then extensively describe the experiments that I have executed on the model in section 5.4. Next, in section 5.5, I present the results and in section 5.6 I conclude with a discussion of findings as they relate to the assumptions and hypotheses stated above.

5.2 Background

Although some argue that the idea of polarization is more an issue of perception and access to information rather than a real problem (see Fiorina, 2006), many studies have yielded quantifiable findings, suggesting that the increase in polarization constitutes an objective reality. The evidence ranges from surveys administered to voters (e.g. Abramowitz and Saunders, 2007) to quantitative analyses of co-sponsorship networks (e.g. Fowler, 2006a; Fowler, 2006b; Zhang et al., 2008) and roll-call voting patterns in both chambers of the U.S. Congress (Andris et al., 2015).

The causes of political polarization have also been extensively debated. Dalton (2008) has highlighted the polarizing effects of party system quantity and quality. Moreover, Baldassari and Gelman (2008), as well as Lachat (2008) have shown that increasing partisanship among the political elites has an effect on the polarization of the entire voter population, while others

point out the significance of media bias on the ideological structure of society (Bernhardt et al., 2008), or the clash of different moral communities (Alizadeh et al., 2016).

Here, I focus solely on U.S. political elites, namely the congressmen and congresswomen of the U.S. House of Representatives, and I attempt to quantify and explain the causes of rising polarization with the use of an-agent based model. The basis of the model rests on the complex web of social interactions among the legislators and the way in which they drive aggregate ideological trends over a prolonged period of time. I specifically single out a special type of interaction among legislators: the co-sponsoring of bills. Bill sponsorship and co-sponsorship is considered an important aspect of political activity among legislators. Campbell (1982) notes that most legislators along with their staff expend a considerable amount of time and effort to draft and distribute “Dear Colleague” letters, which are traditionally used to recruit potential co-sponsors for the legislator’s own legislature. Co-sponsorship has been often viewed as a low-cost opportunity for legislators to signal their policy stance to their constituency or the median voter in their district (e.g. Mayhew, 1974; Campbell, 1982; Kessler & Krehbiel, 1996). The effect of co-sponsorship on legislators’ success (ultimately measured in terms of re-election) has also been studied. Campbell (1982) notes, that co-sponsorship can be used by legislators to political advantage in seeking re-election. Furthermore, Ansolabehere and Jones (2010) find that most constituents hold beliefs about the ideological positions of their representatives and that they hold them accountable based on these beliefs. In fact, research has shown that co-sponsorship is more prevalent among legislators that are relatively more electorally vulnerable (Koger, 2003). Thus, I take co-sponsorship activity as a measure of electoral success.

The model represents the latest addition to a growing tradition of agent-based modeling of political and legislative processes. Kollman et al. (1992) started this trend with their model of two-party elections in a spatial environment. They represent parties as boundedly-rational agents who adapt their strategies in an effort to capture votes in a changing electoral landscape. Subsequently, Laver (2005) extended the agent-based approach to multi-party cases. Laver and Sergenti (2011) have further extended this line of research and developed

a comprehensive agent-based modeling framework of party competition. The concept of party strategy is also explored in Fowler and Laver’s (2008) study, which pits competing party decision rules in a tournament, akin to Axelrod’s PD tournament (1984). Others have focused on voters as agents and specifically on voter behavior and decision-making (Bendor et al., 2003; Fowler, 2006c).

5.3 Methods

I design a model of congress bill co-sponsorship and of the resulting political and ideological landscapes, which rests on the co-evolution of social networks, ideological positions and party affiliation as cultural signs. For the purpose of testing this model, I devise an agent-based simulation whose foundation is based on the framework described in Chapter 4. There are several changes from the original design, most of them relatively minor. The one significant change lies in the way in which cultural trait variants are represented. While the nature of these was thought to be categorical in the original model (as described in section 3.3.2), here I assume that the variants represent different measurable levels of the absence or presence of a particular trait. Specifically, because I am modelling the behavior of legislators, I take them to represent their ideological position on a one dimensional liberal–conservative scale. I also fix the concept of “tag” to represent the agent’s party affiliation, as this is the single most defining observable marker of a political actor.

5.3.1 Entities and Variables

The model consists of 438 agents, each of them representing a member of the U.S. House of Representatives. The agents are very similar to those in Chapter 4 (see table 5.4). The only change in their design is the addition of a new *Tolerance* variable, which I here denote as t . This is because I now assume that trait variants are cardinal, so we can draw comparisons between them. For example, if we enumerate the variants with the integers 1 through n , we may say that two agents that possess variants 3 and 7 are closer to each other in this aspect than a pair of agents with variants 2 and 11. Here, the trait

variants refer to the agent’s position on the liberal–conservative spectrum and I assume that agents have a natural tendency to cooperate with those who are closer to them on this spectrum rather than those that are further away. An agent’s tolerance value t thus defines a range $[v - t, v + t] \cup [min_v, max_v]$ where v is the value of the agent’s own trait variant and min_v, max_v are the minimum and maximum trait variant values respectively. The agent is then willing to cooperate with others whose trait variants fall within this range. Furthermore, unlike in Chapters 3 and 4, I will also assume that the distribution of trait variants is biased to a certain extent by the possession of specific tags. Figure 5.2 shows an illustration of this.

Table 5.1: Agent variables

Name	Domain	Scale	Type
<i>CulturalTrait</i>	Integer	Cardinal (interval)	Static
<i>Tag</i>	Integer	Categorical	Static
<i>Tolerance</i>	Integer	Cardinal (ratio)	Static
<i>PositiveThreshold</i>	Integer	Cardinal (ratio)	Static
<i>NegativeThreshold</i>	Integer	Cardinal (ratio)	Static
<i>Neighborhood</i>	List of agents	Categorical	Dynamic
<i>Fitness</i>	Integer	Cardinal (ratio)	Dynamic
<i>NumGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>NumBad</i>	List of integers	Cardinal (ratio)	Dynamic
<i>FirstGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>FirstBad</i>	List of integers	Cardinal (ratio)	Dynamic
<i>LastGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>LastBad</i>	List of integers	Cardinal (ratio)	Dynamic

For a complete description of the other agent state variables see Chapter 3. Table 5.2 summarizes the global model parameters. Apart from *RewiringLambda* the other parameters have been introduced in the original model and their description can be found in Chapter 3. Furthermore, for the specific purposes of this version, I have fixed several quantities which previously figured in the model as parameters. Here, the number of agents remains fixed at 438 (the number of representatives in the House) and the number of tags remains fixed at 2 (the number of parties represented in the House; by convention, I group all Independents with

Table 5.2: Model Parameters

Name	Domain	Scale
<i>NumTraits</i>	Integer	Cardinal (ratio)
<i>SuccessPayoff</i>	Integer	Cardinal (ratio)
<i>FailurePayoff</i>	Integer	Cardinal (ratio)
<i>AdjacencyMatrix</i>	Matrix of Booleans	Categorical (Boolean)
<i>MutationRate</i>	Floating-point number	Cardinal (ratio)
<i>MaintenanceCost</i>	Floating-point number	Cardinal (ratio)
<i>RewiringLambda</i>	Integer	Cardinal (ratio)

the Republicans). In later sections I comment on how I derive the number of generations, rounds and trait variants. The new *RewiringLambda* parameter defines the typical rate at which agents delete and create links in their social networks. In the original model agents simply cut the links with agents possessing disliked tags all at once (similarly they filled all their available slots for new links all at once with a random sample of agents possessing preferred tags). Since in this case we are limited to only two tags, the mechanism appears to become both unrealistic and highly disruptive. Instead, I define the number of links that an agent will delete or create in any given round to be a random variate obtained from a Poisson distribution with parameter λ which is given by the value of the agent’s *RewiringLambda* parameter¹.

5.3.2 Process Overview and Scheduling

The model is divided into a certain number of “generations”. Each generation represents one two-year electoral period. A generation is then divided into a certain number of rounds. During every round, each legislator is given a chance to act. The order in which legislators are activated is randomized at the beginning of every round. When legislators are activated, they perform two types of actions. They will *interact* by reaching out to another legislator

¹I choose the Poisson distribution to model this behavior because it best defines the probability of a given number of independent discrete events occurring in a fixed time interval.




Figure 5.1: Illustration of the function of the tolerance attribute (T). In the above example the top case results in a failed interaction, while the bottom case results in a success.




Figure 5.2: An example of how the ideological position of legislators could be biased by party affiliation.

within their social network and asking them to co-sponsor their bill². They will also *network* with other legislators. The networking aspect of legislator activity can be further broken down into two equally important parts. First, legislators may choose to *cut* previously existing ties to others. Second, legislators may choose to *create* new connections with others in the House. At the end of each generation the House is re-populated by members new and old, whose characteristics are generated or modified by means of an evolutionary algorithm.

A legislator will interact with one randomly chosen network connection each round. The nature of the interaction process between two legislators in the model is simple. The legislators check each others ideological positions p_i, p_j against their tolerance levels t_i, t_j . If $|p_i - p_j| < \min(t_i, t_j)$ then the interaction is successful, meaning that the legislator agrees to co-sponsor the other's bill (as illustrated in figure 5.1). Both legislators then receive the *SuccessPayoff*. If the interaction is unsuccessful, then both legislators receive the *FailurePayoff*. This represents the idea that co-sponsorship increases the chances of a legislator's re-election, as hypothesized by Campbell (1982).

The flow of the networking process is illustrated in Algorithm 10. As in the original model (as described in section 3.3), the decision regarding which ties to cut and whom to connect to is biased by the possession of particular tags. In this context it simply means that certain legislators will prefer to maintain connections with Democrats, while others will prefer to connect with Republicans, and yet others will be equally prone to connect with members of either party³. To determine these preferences, I use the ACT-R memory model (Anderson & Lebiere, 1998) and calculate the base-level activations for both tags for each legislator individually. The base-level activation τ_P^- for a legislator's unsuccessful interactions with party P is calculated as follows:

$$\tau_P^- = \ln \left[\sum_i^n t_i^{-d} \right] \approx \ln \left[t_n^{-0.5} + \frac{2(n-1)}{\sqrt{t_1} + \sqrt{t_n}} \right] \quad (\text{when } d = 1/2) \quad (5.1)$$

²The actual nature and identity of the bill is beyond the scope of the model and I leave this concept in the abstract simply as "a bill".

³To exhaust the list of possibilities, this also allows for legislators who dislike both parties and would simply prefer to go at it alone and forego connecting with anyone at all.

Algorithm 10 Network Rewiring

```
1: procedure REWIRE(agent, agents, rewiringLambda)
2:   negativeBLA  $\leftarrow$  calculateBLA(agent, “bad”)
3:   retrievedTags  $\leftarrow$  {tag | negativeBLAtag < agent.NegativeThreshold}
4:   n  $\leftarrow$  randomPoisson(rewiringLambda)
5:   dislikedNeighbors  $\leftarrow$  {i | i  $\in$  agent.Neighborhood  $\wedge$  i.Tag  $\in$  retrievedTags}
6:   neighborsToCut  $\leftarrow$  randomSample(dislikedNeighbors, size = n)
7:   agent.Neighborhood  $\leftarrow$  agent.Neighborhood  $\setminus$  neighborsToCut
8:   positiveBLA  $\leftarrow$  calculateBLA(agent, “good”)
9:   retrievedTags  $\leftarrow$  {tag | positiveBLAtag > agent.PositiveThreshold}
10:  candidates  $\leftarrow$  {i | i  $\in$  agents  $\wedge$  i.Tag  $\in$  retrievedTags}
11:  n  $\leftarrow$  randomPoisson(rewiringLambda)
12:  candidates = randomSample(candidates, size = n)
13:  for candidate  $\in$  potentialNeighbors do
14:    candidatePositiveBLA  $\leftarrow$  calculateBLA(candidate, “good”)
15:    if candidatePositiveBLAagent.Tag > candidate.PositiveThreshold then
16:      agent.Neighborhood  $\leftarrow$  agent.Neighborhood  $\cup$  candidate
17:      candidate.Neighborhood  $\leftarrow$  candidate.Neighborhood  $\cup$  agent
18:    end if
19:  end for
20:
21:  return agents
22: end procedure
```

Here t_i is the time elapsed since the i -th unsuccessful interaction with any legislator from party P , while $n = \text{NumBad}_T$ is the total number of such experiences, $t_n = \text{LastBad}_T$ is the time since the most recent experience, and $t_1 = \text{FirstBad}_T$ is the time since the first experience. Finally, d is the rate of decay. Due to the computational complexity of the quantity when n is large, I implement a sufficiently accurate approximation (Petrov, 2006). By convention, I use $d = \frac{1}{2}$. If the base-level activation of a party drops below the legislator’s *NegativeThreshold* it will designate it as disliked.

Once the base-level activations for negative interactions are calculated, a random variate n is generated from the Poisson distribution with $\lambda = \text{rewiringLambda}$. The legislator then randomly selects n of its existing connections from the disliked party (or parties) and cuts ties to them.

In a similar manner I compute base-level activations for successful interactions and compare them to the legislator’s *PositiveThreshold* to designate preferred parties. The legislator then selects a random sample of colleagues from the preferred parties, the size of

which is once again determined by a random Poisson variate. Unlike cutting ties, creating them is a mutual act. Thus, in the same fashion, a candidate will only agree to create the new connection if the proposing legislator is a member of one of the candidate’s preferred parties.

At the end of each round everyone’s fitness levels are adjusted by the maintenance costs for their network connections. Each legislator incurs a fixed cost per connection per round, the value of which is given by the *MaintenanceCost* parameter. At the end of each generation a new House is populated. This is done by considering each “seat”, that is, each node in the network separately. For each seat we take into account the current legislator along with its network connections. From this set we then randomly select a subset, whose size is given by the value of the *TournamentSize* parameter. This subset is then ranked by the legislator’s fitness values. The seat is then filled with a “new” legislator whose attributes at first exactly match those of the top-ranked member of the selected subset. Finally, I introduce mutation. This has the effect that each attribute of the legislator is modified with 1% probability. In the case of party affiliation this modification amounts to switching parties. In the case of cardinal variables the mutation is done by introducing small Gaussian perturbations to the current value⁴. Mutation excludes fitness which is always set to zero at the beginning of a generation and the legislator’s neighborhood, as I fix the global network structure when transitioning between generations.

I measure the success of legislators by their fitness, because it represents their ability to gain co-sponsors for their own bills as well as their willingness to co-sponsor others’ legislation. As noted in the Introduction, bill co-sponsorship has an effect on the continued success of a legislator.

The mechanism for creating “new” legislators can then be interpreted in two ways. The clone that is created can either represent the old legislator being re-elected for the new two-year period, or it could represent a new legislator being elected for the first time while adopting the attributes of a previously successful legislator⁵. However, I remain agnostic

⁴For discrete variables the perturbations are taken from discretized Gaussian distributions.

⁵The fact that representatives can be (and very often are) elected for multiple congresses , i.e. survive for

regarding interpretation in the case of individual seats.

The mutation then represents both the natural shift in positions that sometimes occurs over the course of a politician’s career as well as the imperfections in attempts at mimicking others.

5.3.3 Initialization and Inputs

The model is fully specified by providing the values of model parameters shown in table 5.2 as input. The initial agent population is generated at the beginning of the simulation. The fitness of every legislator is set to zero and their social networks are instantiated based on the provided *AdjacencyMatrix*. Party affiliations are distributed uniformly among the legislators at first. The ideological trait variants are then instantiated based on the party affiliations. As in a real-world scenario, I assume that Republicans will be more conservative on average, and that they will never take extremely liberal positions. Similarly, I expect Democrats to take more liberal positions on average, while never occupying the conservative extremes of the ideological scale. However, I assume that there will be some overlap in the middle of the spectrum. Specifically, for a scale with n levels, where n is the most liberal, I randomly assign each Republican an ideological variant from the uniform distribution on the interval $[1, \frac{2n}{3}]$. Similarly, each Democrat is randomly assigned a variant from the uniform distribution on the interval $[\frac{n}{3}, n]$. The values of *NumBad* and *NumGood* are assigned with uniform probability from the set of integers from 1 to *NumRounds*. The variables *FirstGood*, *FirstBad*, *LastGood* and *LastBad* are initialized with values taken from the uniform distribution on the set of integers from 1 to *NumRounds*NumAgents*. In essence, this creates random fictional memories for the newly created legislators. The ranges of the relevant variables are chosen so that the memories are on the scale of electoral periods. Finally, each legislator is assigned a *Tolerance* value from the uniform distribution on the interval $[0, n - 1]$.

multiple generations, is the reason why there is no recombination in this version of the EA.

5.4 Verification and Validation

The model code has been thoroughly verified by performing code walkthroughs and standard debugging practices. The model was developed in the MATLAB computing environment. To elicit external validity, the model results are compared with the empirical evidence taken from the House co-sponsorship dataset. The process of model validation has been carried out by comparing the real-world co-sponsorship network statistics with the simulated networks, and is thoroughly described in the Results and Discussion sections.

5.5 Experimental Design

The parameter values that I have tested are shown in table 5.3. The number of generations is equal to the number of congresses between the tracked period from 1973 to 2008 (93rd to 110th Congress). I set the number of rounds by taking the average number of bills sponsored and co-sponsored by a single legislator over the entire period and round to the nearest hundred.

Table 5.3: Model parameter values

Name	Value
<i>NumGenerations</i>	18
<i>NumRounds</i>	300
<i>NumTraits</i>	9
<i>SuccessPayoff</i>	{fixed, uniform}
<i>FailurePayoff</i>	{fixed, uniform}
<i>AdjacencyMatrix</i>	random, $\hat{k} = 70$
<i>MutationRate</i>	0.01
<i>MaintenanceCost</i>	0.1, 0.2, $\dots$, 1
<i>RewiringLambda</i>	1, 5, 10, 15, $\dots$, 50

I chose to divide the ideological scale, somewhat arbitrarily, into 9 levels. I believe this is more than enough granularity and it allows for the levels to be split evenly among Republican-only, Democrat-only and mixed types.

The choice of the payoff parameters is important, because it does have the potential to significantly drive the dynamics of the model. One natural choice is to fix the payoffs to a value that remains constant for all interactions regardless of the actors involved. We may therefore choose +1 for the *SuccessPayoff* and -1 for the *FailurePayoff*. However, it should be noted that under this payoff schedule agents positioned closer to the middle of the ideological spectrum as well as more tolerant agents should expect to collect more positive payoff assuming a uniformly random distribution of agents on a random network (see figure 5.3). It may well be the case that such types of agents are predisposed to fare better in the scenarios that I simulate, but we cannot be certain of it. One way to circumvent this artifact is to normalize an agent's payoffs from a given interaction based on its tolerance and ideological variant; simply dividing the unit payoff by its expected payoff under the fixed schedule. This choice is not entirely free of assumptions either. The normalization scheme posits that any single instance of co-sponsorship activity is relatively less valuable to more tolerant and more ideologically centric legislators. Because it is difficult to speculate on the relative payoffs of co-sponsorship to different types of legislators, I decide to test both payoff schedules.

For the *AdjacencyMatrix* of the initial social network I choose one that represents an unweighted, undirected, random Erdős-Rényi network with average degree $\hat{k} = 70$. It should be noted that the social network of House members in the beginning of the tracked period was probably not random. However, because it is close to impossible to determine its structure precisely, I choose to proceed with the random network as a baseline. The value of $\hat{k}$ was chosen based on the average number of unique cosponsors per legislator in the 93rd congress. It is possible, and in fact probable, that an average legislator has other connections in the House apart from those that actively cosponsor their bills. Once again, the accurate numbers are unknown, but since any cosponsor should be considered a social tie, this value practically serves as a lower bound.

I vary the values of the *MaintenanceCost* parameter because I believe that it has a significant affect on the dynamics of system. I thus test values ranging from close to




Figure 5.3: Expected relative payoffs on a random network as a function of tolerance and ideological positions

negligible up to 1, at which point the potential payoff from any given social connection cannot exceed the cost of maintaining it under either of the payoff schedules. Similarly, I vary the *RewiringLambda* parameter within what I believe is a range of sensible values⁶.

The mutation rate is set to 1% as is standard practice in evolutionary modeling of social systems (as shown in section 2.5.4).

Finally, to be able to compare the performance of the party-based model I devise and introduce three simple alternative models of co-sponsorship activity which will serve as benchmarks. These are based on the the alternative hypotheses stated in section 5.1. The models are identical to the party-based model, apart from the networking mechanisms.

In the model based on A_1 the agents' choice of connections depends on ties shared with others. The agents do not form preferences based on party affiliation, rather they cut ties with others once the ratio of common connections drops below a certain threshold value. Vice versa, they consider new connections only if the ratio of common connections exceeds

⁶At $\lambda = 50$ an agent would, on average, cut ties with and create new connections to about 11% of the entire House during any given rewiring action.

a certain threshold. To be exact, an agent cuts ties with someone if $r > F(-w, s)$ and they attempt to connect to someone if $r < F(w, s)$. Here, r is a random variate from the uniform distribution on the interval $[0, 1]$, w is a weight, and s is the ratio of common connections between the two agents. The function $F(x, y)$ then gives the value of the cumulative distribution function of the truncated normal distribution on the interval $[0, 1]$ at the point y , with $\mu = 0, \sigma = 1 + x$ when $x < 0$ and $\mu = 1, \sigma = 1 - x$ when $x > 0$. Figure 5.4 shows the shapes of the cumulative distribution functions for several values of w . The weight w controls the direction and the strength of the attracting force of shared connections. As w gets closer to 0, the probability of cutting and creating ties becomes more proportional to the ratio of common connections. As w increases towards 1 an agent will be less likely to create new connections unless the ratio of shared connections gets progressively closer to 1. It will also be less likely to cut existing ties unless the ratio drops progressively closer to 0. On the other hand as w decreases to -1 the agent will require a progressively smaller ratio of shared connections to become likely to attempt creating a new tie, and a progressively larger ratio of shared connections to be likely to keep an existing tie. I test values ranging from $w = 0$, representing an undirected, proportional social force, to $w = -0.9$, representing a strong positive social force. I will refer to this model as the *social* model.

A third model, based on A_2 has the agents employ a strictly *individual* approach to rewiring their social networks. In this version the agents simply keep a record of their interactions with every single legislator separately (which is arguably still somewhat feasible with at most 437 connections). If they have a negative experience with a certain legislator, they cut ties with them and never re-connect again during the current generation. They create new ties randomly at rate $\lambda = \text{RewiringLambda}$.

Finally, the model based on A_3 assumes that the agents have perfect information regarding the position of all other legislators on the ideological spectrum. I thus allow the agents to re-connect at once to their preferred partners based on their tolerance and ideological position. I then calculate all of the payoffs and maintenance costs based on a hypothetical 300 round generation and finally re-seed the agent population using the evolutionary algorithm




Figure 5.4: Examples of the the social force function F for different values of x .

to complete a single generation⁷. I refer to this as the *ideological* model.

To test the second hypothesis I devise two measures of polarization, each capturing a different aspect of the concept. The first measure is the average pairwise ideological distance between legislators. I simply consider every pair of legislators (not just those connected to each other), mark their distance on the ideological spectrum and calculate the mean. This measure is sensitive to both the variance and the range of the distribution of ideological variants, both of which increase with the average pairwise distance. The second measure I use is average tolerance. Here, I take the tolerance of each legislator as defined in section 5.3.1 and once again take the population mean.

Our analysis is performed on the co-sponsorship network of the legislators. It should be noted that this is different from the social network of the legislators. The co-sponsorship network is a directed, weighted network and it is generated through events occurring during the simulation: If there is a successful interaction between two actors (i.e. one has agreed to co-sponsor the other's bill), a link with a weight of 1 initiating at the co-sponsor and

⁷Because the networking phase is executed all at once, there is no use for the *RewiringLambda* parameters, and thus in this case I only vary *MaintenanceCost* when performing the sensitivity analysis.

terminating at the sponsor is created. If a link between the two already exists in that direction, then the weight is increased by 1.

I then calculate several network statistics on the co-sponsorship network. The first of these is the average local *weighted clustering coefficient* (Barrat et al., 2004). This is an extension of the standard clustering coefficient measure devised by Watts and Strogatz (1998), applied to weighted networks. In our case, when measuring the local clustering coefficient of any given actor I only consider the incoming links. Thus, I effectively measure the probability that a pair of legislators who co-sponsor bills by the same legislator also have an active co-sponsoring relationship between them. In the remainder of this paper I will refer to this measure simply as the clustering coefficient.

Furthermore, I compute the modularity of the weighted co-sponsorship network. As in the above case this extends the standard modularity calculation by considering the weights of the links per Newman's (2004) approach. I also track the number of unique co-sponsors for each legislator. In essence, this is the average in-degree of the unweighted version of the co-sponsorship network (i.e. all weights are made equal to 1 prior to the calculation). For the sake of simplicity, I refer to this statistic as the average degree. Finally, I measure the average shortest path lengths between pairs of actors in the unweighted, undirected version of the co-sponsorship network.




Figure 5.5: Cosponsorship networks in the 95th Congress. Part (a) shows the empirical networks. Remaining parts show networks simulated by different models under the fixed payoff scheme: (b) party-based ($RewiringLambda = 20$, $cost = 0.5$), (c) social ($w = -0.5$, $cost = 0.5$), (d) individual ($cost = 0.5$), (e) ideological ($RewiringLambda = 20$, $cost = 0.5$). In the simulated networks the different shades of red and blue reflect the agents' positions on the ideological spectrum.




Figure 5.6: Cosponsorship networks in the 108th Congress. Part (a) shows the empirical networks. Remaining parts show networks simulated by different models under the fixed payoff scheme: (b) party-based ($RewiringLambda = 20$, $cost = 0.5$), (c) social ($w = -0.5$, $cost = 0.5$), (d) individual ($cost = 0.5$), (e) ideological ($RewiringLambda = 20$, $cost = 0.5$). In the simulated networks the different shades of red and blue reflect the agents' positions on the ideological spectrum.

5.6 Results

Table 5.4: Co-sponsorship network statistics in the U.S. House of Representatives

Congress	Avg. In-Degree	Avg. Path Length	Clustering Coeff.	Modularity
93rd	70	1.95	0.41	0.15
94th	79	1.89	0.44	0.13
95th	93	1.83	0.47	0.14
96th	111	1.76	0.50	0.15
97th	132	1.72	0.56	0.13
98th	157	1.65	0.63	0.12
99th	171	1.61	0.67	0.12
100th	174	1.60	0.67	0.13
101st	84	1.58	0.68	0.14
102nd	172	1.61	0.66	0.14
103rd	144	1.67	0.59	0.19
104th	105	1.77	0.47	0.23
105th	127	1.73	0.54	0.21
106th	151	1.67	0.60	0.18
107th	143	1.68	0.59	0.17
108th	147	1.67	0.60	0.19

I first visualize and inspect the networks evolved under the four different models and compare them to the true co-sponsorship networks in different U.S. congresses. Figures 5.5 and 5.6 show illustrative snapshots of the network evolution. It is immediately clear that each model produces different types of network structures. However, beyond this it is difficult to ascertain the degree to which these networks vary and in which particular aspects. It is equally impossible to tell which of the models produces networks that most closely resemble the data, or even if any of the models come close to re-creating the true co-sponsorship network. I thus turn towards a more rigorous and quantitative analysis of the networks.

Figure 5.7 shows how the networks created by the four different models differ in four important network characteristics. The violin plots show the distributions of resulting values obtained from simulations across the entire search space. Despite this aggregation, one will




Figure 5.7: Violin plots of the distribution of measured network statistics for the four models across the entire search space.

notice that there exist clear-cut quantitative differences among the four models. Furthermore, we note that the party-based model stands out the most: It vastly outperforms the other three models in terms of error in the case of average node degree and clustering coefficient. It also fares better than the other models with respect to network modularity, as only the ideological model achieves similarly low error levels. The only case in which the party-based model is outperformed by any of the remaining models is average shortest path length, where only the social model edges it out. However, I believe it is safe to conclude that overall the party-based model performs substantially better than any of the baseline models in terms of goodness of fit with the historical data. Moreover, we can consider this result highly robust, because it holds regardless of the position in the search space.

To gauge how the goodness of fit varies through time I present results for a single point in the search space for each of the four models. Figure 5.8 shows the temporal variation of the tracked measures. Once again we note that all three baseline models perform poorly on one or more of the four measures. Meanwhile, not only does the party-based model perform the best in terms of error, it also seems to pick up on some of the trends present in the data. For example, the clustering coefficient of the simulated networks exhibits a distinct, increasing trend after the first handful of generations, similar to the empirical data (although the trend appears noticeably earlier in the data). In the same vein, the simulations appear to foretell the increase in empirical network modularity in the middle stages of the observed period, before eventually leveling off.

From here on I therefore only consider the party-based model. Figures 5.9 and 5.10 show the sensitivity of the model to its location in the search space. In the case of the uniform payoff schedule, we observe that the model is especially sensitive to tuning the parameter λ . As we increase λ , the error in average path length also increases, while the error in clustering coefficient decreases. Meanwhile the error in average degree becomes more prominent as we move toward both extremes of the tested λ values. The fit in terms of network modularity does not appear to be as sensitive to either of the varied parameters, apart from very small regions in the extremes of the search space.




Figure 5.8: Time series of the measured network statistics for the four models at a single point in the search space ($RewiringLambda = 20$, $MaintenanceCost = 0.5$, $SocialWeight = -0.5$, uniform payoff schedule).

When employing the fixed payoff schedule one may notice that lower λ values contribute to poorer fit of the model in terms of average degree as well as clustering coefficient. Furthermore, we see that error in average path lengths increases with both λ and the maintenance cost. In the case of modularity it is difficult to discern any clear trend as either of the parameters are varied. Finally, we observe that the fixed payoff model variation shows better goodness of fit with respect to average degree and average path lengths. Conversely, the uniform payoffs perform better with respect to modularity, while the difference in clustering coefficient fit is negligible.




Figure 5.9: Gradient fields of the measured network statistics for the party-based model under the uniform payoff scheme as a function of λ and t .

Finally, I comment on how the position in the search space affects our previously defined measures of partisanship and polarization. Figures 5.11 and 5.12 display the dependence of the resulting average tolerance and average pairwise ideological distances on the link maintenance cost and the λ parameter. We first note that, under the uniform payoff regime, the average tolerance is below the expected value of the initial uniform distribution, regardless of position in the search space. Furthermore, as the cost decreases and as λ increases the actors become less tolerant on average. On the other hand, the populations become invariably closer on average in the ideological space compared to the initial distribution of ideologies⁸.

⁸Assuming a population that is uniformly distributed on an ideological spectrum with levels from 1 to 9, the average pairwise distance will be approximately 2.97.




Figure 5.10: Gradient fields of the measured network statistics for the party-based model under the fixed payoff scheme as a function of λ and c .

However, the actors tend to become further apart on average, as both λ and link maintenance costs increase.

The situation is different when considering the fixed payoff regime. Here the population ends up being substantially more tolerant than in the previous case, although it is still below the expected value for a uniform distribution. The average pairwise distance remains low, albeit somewhat higher than under the uniform payoff regime.




Figure 5.11: Gradient fields of average tolerance and average pairwise distance for the party-based model under the uniform payoff scheme as a function of λ and c .

5.7 Discussion

The findings from our experiments show that a model in which legislators rely on party affiliation as a proxy for complex ideological positions in hopes of collaborating on legislation results in co-sponsorship networks that closely resemble the empirical evidence. Moreover, models in which legislators rely on strategies that take into account attributes other than party affiliation show poorer fit with the actual co-sponsorship networks found in the U.S. House of Representatives. This tells us that party affiliation is indeed an important driver of the collaboration dynamics in the House, even in those cases where certain legislators across the aisle might be ideologically closer than many fellow party members.

There are a few potential reasons why party affiliation is such a strong driving force. As I eluded earlier, it could be because close evaluation of others' positions is too expensive and legislators thus resort to sticking with party lines as a cheap and dirty heuristic. It may also be the case that legislators are forced to obey party lines by mechanisms out of their own control, such as party leadership (Rohde, 2010), electoral change (Jacobson, 2000), etc. Adjudicating these underlying causes is, however, beyond the capability of the model and thus beyond the scope of this study.




Figure 5.12: Gradient fields of average tolerance and average pairwise distance for the party-based model under the fixed payoff scheme as a function of λ and c .

Furthermore, the party-based model shows the system evolving towards low levels of ideological tolerance among the body of legislators under the uniform payoff schedule. On the other hand, the average ideological distance decreases towards minimal levels. This suggests that increased partisanship leads towards a certain type of polarization, but actually suppresses others. When fixed payoffs are assumed ideological proximity tends to be high and tolerance remains moderate. This indicates that ideological tolerance will suffer on the aggregate level in contexts where the act of co-sponsorship is individually more valuable to legislators who assume more extreme positions. This notion is supported by the findings of Kessler and Krehbiel (1996) who report that policy extremists on both sides of the spectrum are more likely to support legislation in its initial phases.

The higher aggregate tolerance levels under the fixed payoff schedule confirm our intuition supported by the calculus of expected payoffs in a randomly networked body of legislators: more tolerant and middle-of-the-road actors will fare better. In the case of uniform payoffs, less tolerant legislators end up faring better, although from a mathematical point of view they should hold no advantage over others in a randomly connected body of legislators. This reflects on the importance of the network structure itself. As the networks become significantly




Figure 5.13: Violin plots of ideological trait assortativity distributions across the entire search space over time in the party-based model.

clustered and modular in terms of ideology (see figure 5.13), the payoff expectations change: steadfast adherence to party lines rewards intolerant attitudes over more open-minded ones.

In general, the combination of low tolerance and close ideological proximity is intriguing, yet plausible. Downs (1957) shows from an economic perspective, that convergence towards the ideological center is an equilibrium state in a two-party system. Furthermore, comments on the ideological similarity between the two parties are well-documented, dating as far back as the 19th century, when Tocqueville (2000 [1835]) proclaimed that while America has had great parties, it no longer has any, as they do not adhere to particular principles or generalities.

The findings hold interesting implications for further research. Several factors contributing to the model remain unexplored. Here, I have represented ideological position as a simple one-dimensional scale. In reality ideology is almost certainly multi-dimensional. It is therefore natural to wonder how the dimensionality of ideology representation would affect the behavior of the model. Furthermore, the model framework allows researchers to consider a wide range of hypothetical scenarios. For example, this creates the opportunity to the

test how the quantitative and qualitative characteristics of party systems contribute to polarization dynamics in legislative bodies.

Finally, it is worth noting that although this model does not take voter behavior or constituency factors into account, these are probably also important factors which may contribute to or feed back into the dynamics of partisanship and polarization.

5.8 Conclusion

I have used an agent-based model of cultural evolution to explain the dynamics of bill cosponsorship and, by extension, of political polarization in the U.S. House of Representatives.

I have found that the proposed model based on the imperfect knowledge of actors regarding the ideological dispositions of others and their reliance on party affiliations as an alternative heuristic to garner support for proposed legislation performs well against empirical evidence. Moreover, I have determined that it fares better compared to several alternative models. As the model pertains to our definitions of political polarization, I gathered mixed results, finding that the proposed mechanisms lead to decreased tolerance, as well as decreased ideological distance. Overall, I believe that the model provides increased confidence in the hypothesis that bill co-sponsorship, and the underlying social and cultural processes which lead to it, contribute to certain aspects of polarization in the U.S. House of Representatives.

Chapter 6: Cultures of Collaboration: Simulating Scientific Co-Authorship Networks

Abstract: Modern science has become a team effort. Collaboration and co-authorship is the new standard in most disciplines. Pathways of collaboration between scientists have been studied previously. Here I take a novel approach where I combine the agent-based modeling paradigm and a cultural, dual inheritance perspective to develop a model of the dynamics of scientific collaboration. I study the performance of the model by comparing it to an extensive empirical dataset on academic publishing. I show that there are significant differences in model performance across different disciplines, suggesting that the cultures of collaboration may differ across fields.

6.1 Introduction and Background

Scientific collaboration is a complex phenomenon that is crucial to the effective diffusion, communication and exchange of scientific knowledge. Scientific collaboration has arguably only increased in complexity and effect since the advent of modern information technology. It is perhaps for this reason, that research into networks of scientific collaboration has surged in frequency in recent decades. One of the most direct manifestations of the pathways of scientific collaboration is the co-authorship of scientific literature. The study of scientific co-authorship networks has originated in the social sciences (e.g. Endersby et al. 1996; Moody, 2004) but has since been joined by efforts in the natural sciences as well (e.g. Barabasi et al., 2002; Newman, 2004). The focus in most studies has always been the analysis of the structure of the co-authorship networks, whether quantitative or qualitative. Barabasi et al. (2004) have compared networks generated from different datasets and noted the scale-free structure of such co-authorship networks. They then applied the preferential attachment model to successfully generate similar structures. Newman (2004) has further corroborated the scale-free nature of co-authorship networks in his study. Many other efforts to probe the underlying mechanisms of scientific collaboration through network analysis have appeared since: from analyzing the choice of data (De Stefano et al., 2013) and network definitions (De Stefano et al., 2011), through community detection approaches (Perianes-Rodrigues et al., 2010), qualitative approaches (Velden et al., 2010), to investigating the effects of network position on scientific performance (Abbasi et al., 2011, Abbasi et al., 2012). Explanatory models of scientific collaboration have also been proposed. Barabasi et al.'s (2002) application of the preferential attachment mechanism is perhaps the most influential. Morris and Goldstein (2007) have developed a mathematical model based on concept of weak ties (Granovetter, 1973), while De Lange and Glanzel (1997) use a series expansion approach.

These research efforts have two important elements in common: First, the focus is almost exclusively on network concepts or social forces as causal explananda, and second,

the studied networks are either historical snapshots or cumulative graphs. I take a new approach to the study of collaboration. I study the issue of scientific collaboration as a multi-faceted one, and perceive co-authorship networks as a result of a combination of social mechanisms, cultural mechanisms, network structure effects and temporal (evolutionary) effects. Furthermore, I am interested in the analysis of co-authorship networks as constantly evolving structures, such, where pathways of collaboration can be forged, but also abandoned over time. I also approach the issue from an agent-based modeling perspective. Agent-based models are well-suited to explaining trends emanating from individual decisions as well as elucidating gradual change over time (Gilbert & Troitzsch, 2005).

I thus propose and test a model of scientific co-authorship. In this model I assume that actors (scientists) make choices regarding collaboration with others based on both their *cultural* and *social* preferences, which are in turn affected by their current position in the network and past experiences of their own, as well as those of their previous collaborators. To test the model I develop an agent-based simulation.

The *cultural* aspect of the model rests on the logic of dual inheritance theory, specifically on the mechanisms of indirectly biased transmission and guided variation; both being suggested drivers of cultural evolution (Boyd & Richerson, 1985). Here, indirectly biased transmission refers to the evaluation, adoption and subsequent diffusion of specific cultural behaviors and attitudes on the basis of possession of initially unrelated external markers. Meanwhile, guided variation refers to the evaluation and adoption of cultural behaviors based on self-generated and self-explored alternatives. I introduce these concepts into the model, because I posit (a) that cultural considerations, such as the actor's approach to research organization and management, communication, or writing, not to mention the choice of scientific paradigm (from potentially many within any given discipline) are all important determinants of successful collaborations, as measured by the actors' satisfaction with the process and the end result, and (b) that these cultural factors are often unknown to others beforehand. Therefore, I assume, that the actors are often forced to rely on externally observable markers, or cultural signals, such as institutional affiliation, rank, past publication

record, etc., as proxies for actual cultural behavior.

The *social* aspect of the model then rests on the observation that actors in social networks tend to cluster together (Watts & Strogatz, 1998), i.e. the probability of two actors collaborating increases with the degree to which their ego-based networks overlap. In fact, today we see many scientific collaborations where the individual actors do not necessarily know each other, and participate in the effort together solely on the basis of a shared acquaintance that also happens to be part of the project.

Finally, I model temporal and structural effects via an *evolutionary algorithm*. Evolutionary algorithms are used both as optimization tools, but also as a means to simulate the evolutionary dynamics of social systems (De Jong, 2005, p. 28). In this case, I assume that current actors may leave the network entirely (e.g. retirement), and that new actors may enter it at any point. Furthermore, I assume that new actors joining the network adopt cultural behaviors and preferences of current actors to a certain extent (e.g. graduate students from their advisors, post-docs from PIs, or junior faculty from senior faculty), and that successful, high-performing actors gain relatively more “disciples” on average than others. Finally, I assume that the maintenance of active scientific collaborations is costly in terms of time, effort and resources.

To test the model I compare its output to empirical data on scientific publishing. Specifically, I use the Microsoft Academic Graph¹ (Sinha et al., 2015) which is a large database that includes information on over 126 million scientific publications, written by over 114 million different authors, from over 50,000 fields, over the course of more than 150 years. Despite certain shortcomings of the dataset, such as the extent of missing, incorrect, or duplicate data, it correlates well with other major publishing databases (e.g. CORE, Scimago, or Mendeley), and is currently considered the most comprehensive publicly available dataset of its kind (Herrmannova & Knoth, 2016).

In the following sections I first describe the agent-based implementation of our model in full detail (section 6.2). I then proceed to outline the experiment design and the data

¹I will further refer to the Microsoft Academic Graph by the abbreviation MAG.

manipulation process (6.4). Finally, I report the observed results and discuss the most important findings (6.6-6.7).

6.2 Methods

I devise a model of the evolution of academic co-authorship networks which rests on *cultural* (or *institutional*) forces, as well as *social* forces. To test the model I carry out multi-agent simulations that are based on the model introduced in Chapter 4 and extended further to control for the specifics of the academic publishing context. As in the original cultural evolution model, the agents possess one of many possible variants of a trait that is unknowable to other agents *a priori*; similarly, the agents possess one of several possible *tags*, i.e. directly observable external markers. The agents seek to collaborate with others, but are only successful if their trait variants match those of their partners. The agents are therefore forced to rely on the tags to select adequate partners, and as a result they form distinct preferences for different tags over time.

There are two major departures from the original model in this version. The first change is related to population size. Unlike in the original model, here the population sizes are not necessarily constant. In fact, the populations keep growing as time progresses. This is true not only of the agent population, but also of the populations of possible trait variants and tags. As time progresses I assume that new variants and tags will be “discovered” or “invented”. In the case of scientific collaboration this might represent the founding of new institutes and journals, or the development of new paradigms and methodologies. The second major change comes with the addition of a *social attraction* mechanism, which seconds the cultural mechanism of indirect bias as a driving force for the agent network evolution. I will discuss these additions in greater detail later in this section.

6.2.1 Entities and Variables

The model consists of a number of agents who are interested in collaborating to produce value. The agents are defined by a list of state variables. These are presented in table

6.1. Most of these are identical to the agent state variables described in Chapter 4. I have already discussed the *CulturalTrait* and *Tag*. As before, the *PositiveThreshold* is defined as the minimum base-level activation that a tag has to clear, for it to be considered by the agent when creating new ties. Similarly the *NegativeThreshold* is the minimum base-level activation a tag has to clear, for it to be considered when deleting ties. The *Neighborhood* is simply the set of all of the other agents that connect to the ego via an immediate link and *Fitness* is the sum of all successful interactions during an agent’s lifetime minus the number of unsuccessful interactions. Here, fitness represents an author’s ability to collaborate with others and thus publish articles. The implied assumption is that the more published the author, the more successful they will be in attracting disciples. The *LastOutcome* variable keeps track of whether the most recent interaction was successful or not. The sole addition to the list of state variables is the variable *YearsLeft*. Unlike before, this version of the model is non-generational. Instead, new agents are introduced into the simulation and old agents leave on a step-by-step basis. This is certainly the case in scientific communities, where individuals have finite careers. The agent’s lifetimes are initiated at their “birth” into the *YearsLeft* variable, and its value is decremented every step (or “year”) to keep track of the agent’s remaining lifetime, before it is retired from the population. Finally, as in the original model, the remaining variables figure in calculating the preference levels for different tags (their meaning is described in subsection 6.2.2).

Table 6.2 lists the global model parameters. I experiment with two configurations of the model. One in which both the *social* and the *cultural* mechanism are in play, which I will refer to as the *biased* configuration, and a baseline configuration, in which these mechanisms are omitted. I refer to the baseline as the *unbiased* configuration. I explain the differences between the two configurations in subsection 6.2.2. The *InitPopsize*, *InitTraits*, and *InitTags* give the initial number of agents, available trait variants and tags at the beginning of the simulation. The *PopGrowth*, *TraitGrowth*, and *TagGrowth* variables control the rates at which new agents, trait variants or tags are introduced into the simulation at each step. As before, the *AdjacencyMatrix* represents the initial configuration of the agent social network

Table 6.1: Agent variables

Name	Domain	Scale	Type
<i>CulturalTrait</i>	Integer	Categorical	Static
<i>Tag</i>	Integer	Categorical	Static
<i>PositiveThreshold</i>	Integer	Cardinal (ratio)	Static
<i>NegativeThreshold</i>	Integer	Cardinal (ratio)	Static
<i>Neighborhood</i>	List of agents	Categorical	Dynamic
<i>Fitness</i>	Integer	Cardinal (ratio)	Dynamic
<i>LastOutcome</i>	Ordered pair of integers	Boolean/categorical	Dynamic
<i>YearsLeft</i>	Integers	Cardinal ratio	Dynamic
<i>NumGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>NumBad</i>	List of integers	Cardinal (ratio)	Dynamic
<i>FirstGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>FirstBad</i>	List of integers	Cardinal (ratio)	Dynamic
<i>LastGood</i>	List of integers	Cardinal (ratio)	Dynamic
<i>LastBad</i>	List of integers	Cardinal (ratio)	Dynamic

and the *SuccessPayoff* and *FailurePayoff* give the fitness increments (decrements) for each successful (unsuccessful) interaction between two agents. Similarly, the *MaintenanceCost* refers to the cost the agent bears every step for maintaining a single link to another agent and the *MutationRate* defines the probability with which agent states are stochastically modified after each step. The value of *SocialWeight* indicates the relative strength of the social attraction mechanism in the model. Finally, the *TimeMatrix* holds information on the next scheduled activation of links between agents: the value stored in position (i, j) refers to the next activation of the link between agents i and j . The *TagInnovation* and *TraitInnovation* variables give the probabilities of adopting newly discovered trait variants and tags by surviving agents. The *Lifespans* distribution is used for sampling lifetimes when agents are created. Similarly, the *ActivationIntervals* distribution is used to sample times remaining until the next activations of links in the agent networks.

Table 6.2: Model Parameters

Name	Domain	Scale
<i>Configuration</i>	Integer	Categorical
<i>InitPopsize</i>	Integer	Cardinal (ratio)
<i>InitTraits</i>	Integer	Cardinal (ratio)
<i>InitTags</i>	Integer	Cardinal (ratio)
<i>PopGrowth</i>	Floating-point number	Cardinal (ratio)
<i>TraitGrowth</i>	Floating-point number	Cardinal (ratio)
<i>TagGrowth</i>	Floating-point number	Cardinal (ratio)
<i>SuccessPayoff</i>	Integer	Cardinal (ratio)
<i>FailurePayoff</i>	Integer	Cardinal (ratio)
<i>AdjacencyMatrix</i>	Matrix of Booleans	Categorical (Boolean)
<i>TimeMatrix</i>	Matrix of integers	Cardinalratio
<i>MaintenanceCost</i>	Floating-point number	Cardinal (ratio)
<i>SocialWeight</i>	Floating-point number	Cardinal ratio
<i>TraitInnovation</i>	Floating-point number	Cardinal ratio
<i>TagInnovation</i>	Floating-point number	Cardinal ratio
<i>Lifespans</i>	Histogram of integers	Cardinal (ratio)
<i>ActivationIntervals</i>	Histogram of integers	Cardinal (ratio)
<i>MutationRate</i>	Floating-point number	Cardinal (ratio)

6.2.2 Process Overview and Scheduling

The model processes can be broken into three distinct parts. This includes *neighborhood maintenance* along with *interaction* done by the agents, and the *evolutionary algorithm* which controls the nature of the agent, trait and tag populations over time. In each step every agent is activated and gets a chance to maintain its neighborhood, and potentially to interact with one other agent. The order in which agents are activated is randomized at the beginning of each step.

The *neighborhood maintenance* phase is carried out similarly as in Chapter 4. First, each agent determines its sets of preferred and disliked tags. Additionally, the agent then computes a *social* score for each of its current neighbors. The social score largely depends on the degree to which the neighborhoods of the two agents overlap. If a current neighbor fails to reach a certain threshold degree of overlap *and* it possesses a disliked tag, the agent will

cut ties to it. In a similar fashion, the agent will propose to create a new link to anyone with whom it is not yet connected, if their neighborhoods overlap sufficiently *and* the other agent possesses a preferred tag. Once more, link creation is a mutual act, and thus, the link is created only if the other agents accepts by following the same protocol. At the conclusion of this phase we adjust the agents' fitness values by the maintenance costs, which is determined by the number of their connections.

The preference for a given tag is determined by calculating its base-level activation value. This quantity is taken from the ACT-R memory model (Anderson & Lebiere, 1998). The base-level activation τ_T^+ for successful interactions with tag T is calculated as follows:

$$\tau_T^+ = \ln \left[\sum_i^n t_i^{-d} \right] \approx \ln \left[t_n^{-0.5} + \frac{2(n-1)}{\sqrt{t_1} + \sqrt{t_n}} \right] \quad (\text{when } d = 1/2) \quad (6.1)$$

Here t_i is the time elapsed since the i -th successful interaction with an agent bearing tag T , while $n = NumGood_T$ is the total number of such experiences, $t_n = LastGood_T$ is the time since the most recent experience, and $t_1 = FirstGood_T$ is the time since the first experience. Finally, d is the rate of decay. Due to the computational complexity of the above relationship for large n , I use the above approximation (Petrov, 2006). In line with convention, I use $d = \frac{1}{2}$. Dislike for a given tag is calculated in the same way, using unsuccessful interactions as input. Figure 6.1 shows an example of the change in base-level activation for some input over time.

The *social* score of an ordered pair of agents (i, j) is computed in multiple steps. First, the overlap $\mathcal{O}$ in their neighborhoods is defined as:

$$\mathcal{O} = \frac{|N_i \cap N_j|}{|N_i|}.$$

Here, $|N_i|$ is the size of i 's neighborhood. Next, we take the truncated normal distribution on the interval $[0, 1]$ with $\mu = 0, \sigma = 1 + w$, if $w < 0$, and $\mu = 1, \sigma = 1 - w$, if $w > 0$. Finally,




Figure 6.1: Example of a base-level activation for some input over time. The spikes in the chart coincide with instances of processing the input, followed by gradual decay.

we take the value s of the CDF of the chosen distribution at $\mathcal{O}$ (Figure 6.2 shows examples of the CDFs for several values of w). This gives us the social score.

When deleting links, we take $w = -\text{SocialWeight}$ and the social condition for deletion is met when $s < r$, where r is a random draw from the uniform distribution on the interval $[0, 1)$. When creating new links, we set $w = \text{SocialWeight}$, and the condition is met when $s \geq r$.

The *interaction* phase of the simulation is executed exactly as in Chapter 4: The two agents simply compare their trait variants, and if they match they both increment their fitness by the *SuccessPayoff*. If the trait variants do not match, the fitness of both agents is decremented by the *FailurePayoff*.

Once every agent has taken its turn in the current step, the evolutionary algorithm is invoked. First, the number of new trait variants, tags and agents is calculated. Next, any agents scheduled for removal are removed from the simulation. Finally, new agents are created and added to the network.




Figure 6.2: Examples of the the social force function F for different values of x .

Reproduction is performed locally. For each new agent that is set to enter the population a random node location in the network is chosen. A set of candidates for that location is assembled by taking the agent occupying the chosen node itself along with all of its immediate neighbors. Once the candidate set is defined, we calculate the mean and standard deviation of the fitness distribution within this set. Next, we select those candidates whose fitness is at least one standard deviation above the set mean to become the *parents* of the new agent. The agent is created by performing uniform attribute-wise crossover on the parent set, i.e. for each attribute the new agent copies its value from one of the parents, chosen with uniform probability. I model multi-parent recombination, because in cultural evolution it is possible for offspring to have more (or conversely less) than two parents. The truncation point is chosen so that agents always select above average models, in a way that makes the selection mechanism less sensitive to the distribution of absolute fitness (we are only interested in relative fitness). Finally, mutation is introduced by modifying each agent attribute by a small amount with probability equal to *MutationRate*. For categorical variables mutation amounts to uniformly random switching. In the case of cardinal variables

mutation is carried out by adding small perturbations sampled from Gaussian distributions. The *Fitness* of a new agent is set to zero. Once the agent is created it is automatically linked to its parents.

The links between surviving agents remain intact during the evolutionary procedure. However, the trait variants and tags of surviving agents may be modified with a small probability equal to *TraitInnovation* and *TagInnovation* respectively, to reflect occasional in-life adoption of new tags or variants.

6.2.3 Initialization and Inputs

A complete specification of the model is given by providing the values of the model parameters in table 6.2. The agents are placed on nodes of the network specified by the initial *AdjacencyMatrix*. The trait variants and tags are assigned to agents uniformly from the initial distributions, and independently of each other. Each agent is assigned a lifetime by sampling the *Lifespans* distribution. The *TimeMatrix* tracks the times to the next scheduled activation of specific links in the network (i.e. interactions between specific pairs of agents). Once a link is activated and the interaction between agents i and j takes place, the value of $TimeMatrix_{i,j}$ is re-seeded with a random variate from the *ActivationIntervals* distribution.

6.3 Verification and Validation

The model was implemented in code in the MATLAB computing environment. The model code has been verified by performing thorough code walkthroughs and standard debugging practices. To ascertain the external validity of the model, the resulting distributions of important characteristics of the simulated networks are compared to empirical observations extracted from the House co-sponsorship dataset. The process of model validation has been carried out by comparing the empirical co-sponsorship networks to the simulation results, and is described in greater detail in the Results and Discussion sections.

6.4 Experiment Design and Data Processing

The tested model parameter values are listed in table 6.3. I tested the sensitivity of the model to four variables: the model configuration (biased, unbiased), the cost of maintaining social relationships, the weight of the social mechanism and, crucially, the scientific *field*.

To control for differences between scientific disciplines I randomly chose the fields of Economics and Artificial Intelligence. The only requirements for the choices of fields were that they are not too obscure (i.e. that the number of authors and publications in the field is substantial) and that they have a significant history. The MAG database holds records of over 50,000 scholarly publications by thousands of different authors in both Economics and Artificial Intelligence. However, certain differences remain: Artificial Intelligence (first appearing in the database in 1946) is younger than Economics (first appearing in 1931), and AI is currently just reaching its peak volume and popularity, while Economics has been long considered an established discipline.

Working with the MAG, I first interrogated the *Fields of Study* table and filtered out all of the fields whose name included the strings “economics” or “artificial intelligence” respectively (ignoring case). I then filtered the *Papers* table to include only those publications whose *Field* value matched one of those extracted in the previous step. Finally, I then joined the *Papers* and the *Authors* tables on the *Title*, *Author* and *Year* fields, to create a table of unique Paper–Author–Year triplets for each of the two field groups.

The population growth rates were established by analyzing the totality of authors and their appearances in the MAG for the two respective fields. The counts of new author appearances in given years were then fitted with exponential models for each field separately, using non-linear least squares. These models were then used to generate the number of new agents introduced into the simulation at each time step. I took a similar approach when modelling the lifespans and activation intervals for the two fields. I measured the career span of each author by noting the time elapsed between their first and last publications in the given field. I also measured the interval lengths between successive co-authored publications for all pairs of co-authors in the respective fields. I then created histograms

from the observed values (binned by year). I ran the simulations for 70 steps, which is equal to the number of years for which Artificial Intelligence (the “younger” of the two fields) appears in the MAG prior to 2015 (the last complete year in the database at the time of experimentation).

The initial population size was set, somewhat arbitrarily to 8 agents. In both fields, there were only a few active authors in the beginning years. However, I chose a slightly higher number, because for lower initial population sizes the number of runs that went “extinct” prematurely in the first few steps (due to stochasticity in the agent lifespans coupled with the chosen growth models) was inconveniently high. However, preliminary observations concluded that final results were not significantly affected by slight changes in the initial population size. The values of the *InitTraits*, *InitTags*, *TraitGrowth* and *TagGrowth* variables were also chosen with a degree of arbitrariness. I initially intended to deduce these values from analyzing the MAG as well (by observing variables such as the numbers of unique journals, institutional affiliations, etc.), however I soon learned that the records for the relevant fields were insufficiently reliable, with an excessive number of duplicities, missing values and errors. I was thus forced to provide my own values, which I believe represent reasonable estimates (I posit that growth in cultural trait variants is faster than in relevant external markers, such as institutional affiliation, rank, possible publication venues, etc.). I carried out 100 simulations for each tested parameter combination.

6.5 Results

I measure three important attributes of the agent networks as they evolve over time: the average local clustering coefficients, average node degree, and the average shortest path lengths. I then compare the simulated values with the values observed in the data extracted from the MAG. Figure 6.3 shows the longitudinal trends in different networks statistics for the Economics and the AI networks. Figures 6.4 and 6.5 show the root-mean-squared errors between the simulated and empirical time series for each of the three attributes as a function

Table 6.3: Model parameter values

Name	Value
<i>Configuration</i>	biased, unbiased
<i>NumSteps</i>	70
<i>InitPopsize</i>	8
<i>InitTraits</i>	3
<i>InitTags</i>	3
<i>PopGrowth</i>	Economics, AI
<i>TraitGrowth</i>	0.04
<i>TagGrowth</i>	0.02
<i>TraitInnovation</i>	0
<i>TagInnovation</i>	0
<i>Lifespans</i>	Economics, AI
<i>ActivationIntervals</i>	Economics, AI
<i>SuccessPayoff</i>	1
<i>FailurePayoff</i>	-1
<i>AdjacencyMatrix</i>	random
<i>MutationRate</i>	0.01
<i>MaintenanceCost</i>	(0, 1]
<i>SocialWeight</i>	(−1, 0]

of *MaintenanceCost*, *SocialWeight* and the model configuration. The root-mean-squared error is measured only after the first 20 steps (years), once the empirical populations are large enough and the trends settle sufficiently (see figures 6.6 and 6.7 for comparison of the actual time series). We first notice that there are significant differences in both the performance and the sensitivity of the models between the two fields. One may observe, for example, that while the Economics simulations seem to be sensitive only to the value of *MaintenanceCost*, the AI simulations show sensitivity with respect to both the *MaintenanceCost* and the *SocialWeight*. Furthermore, the scale of the errors differs as well: where the average degree RMSE ranges from negligible levels to over 20 in the case of Economics, it remains between 0.7 and 2.1 for the AI runs. Similar discrepancies can be seen in the clustering coefficients and the average path lengths of the networks.

As we focus on the AI simulations, we note that in the biased model the error in terms




Figure 6.3: Observed statistics for the empirical co-authorship networks in given years.




Figure 6.4: Gradient fields of the RMSE in terms of (a) average degree, (b) clustering coefficient and (c) average path length relative to the Economics data for the biased (right) and the unbiased (left bar) model configurations as a function of social weight and maintenance cost.

of the clustering coefficient as well as the average path length decreases with both the *SocialWeight* and the *MaintenanceCost*. A similar trend appears at first in the case of average degree, however it becomes noticeably reversed at the furthest extremes of the tested ranges. Nonetheless, we observe that the unbiased model is able to follow the empirical trends in network evolution just as well, if not better, in terms of average path length and average degree. It also performs well in terms of clustering coefficient, relative to the biased model.




Figure 6.5: Gradient fields of the RMSE in terms of (a) average degree, (b) clustering coefficient and (c) average path length relative to the AI data for the biased (right) and the unbiased (left bar) model configurations as a function of social weight and maintenance cost.

The situation is different when considering the Economics simulations. First, we notice that in certain regions of the search space the biased model produces a better bit fit than the unbiased configuration in terms of clustering coefficient and average path length. It also performs similarly well in specific regions in terms of average degree.

Moreover, I point out that the biased model seems to achieve the best goodness of fit in the critical region of $MaintenanceCost \approx 0.6$. In fact, one may observe distinct trifurcations in the resulting network characteristics (see figure 6.6), with one of the branches following the empirical trends significantly better than the others. As we move away from this region into lower values of $MaintenanceCost$, the error increases both in terms of clustering coefficient as well as average path length. In the opposite direction, as $MaintenanceCost$ increases, the clustering coefficient error increases once again. Meanwhile, the unbiased model does not seem to be significantly sensitive to changes in the $MaintenanceCost$ parameter; this observation holds in the AI simulations as well.

6.6 Discussion

The model results are illustrative in multiple aspects. First, they suggest that different scientific fields are marked by different sets of mechanisms driving collaboration. Our




Figure 6.6: Time series of the (a) average degree, (b) clustering coefficient and (c) average path length for the Economics biased model runs (blue) and the Economics data (red).




Figure 6.7: Time series of the (a) average degree, (b) clustering coefficient and (c) average path length for the AI biased model runs (blue) and the AI data (red).

model performed very differently when applied to the Economics networks rather than the Artificial Intelligence networks. In the case of the Artificial Intelligence populations, neither the cultural nor the social mechanisms showed any significant effect on the resulting co-authorship networks. The parameters that seemed to matter were the cost of maintaining relationships and the weight of social forces. Conversely, the Economics populations showed good fit with the biased model, but demonstrated sensitivity only to costs of collaboration and not with respect to the social weight. This result further carries the implication that in certain scientific populations the hypothesized social and cultural mechanisms do potentially

guide the dynamics of the system.

The fact that the AI and Economics cultures is different, are reflected in the empirical measures of their network structures. This is not as surprising, considering that the field of Artificial Intelligence is rooted mainly in computer science, while Economics is mostly categorized among the social sciences. The cultural differences between the social and the natural, or exact sciences have been noted before. For example, the rate of co-authorship (as opposed to sole authorship) is significantly higher in the natural sciences (Moody, 2004). This is perhaps reflected in the higher observed clustering coefficient and the higher average observed degree in the AI networks. Moreover, computer scientists are specific in that they give disproportionately more weight to conference articles rather than journal articles, relative to other disciplines. Other cultural differences might be at play, such as thresholds for co-authorship, or the number of appropriate publication venues.

I also suspect that another reason for the seemingly different modes of collaboration is the time of maturation of both fields. As Economics matured earlier, any culture that has been established in the field in the past is potentially deeply rooted in the ways collaborations are forged. Meanwhile, the younger field of AI could be potentially more influenced by the recent proliferation of modern channels of communication and knowledge exchange.

Although the model that I have developed and analyzed has shown some promise, I remain wary to draw any hard conclusions, as I was forced to leave a large region of the parameter space unexplored. This was mostly due to the lacking availability of data. Although the MAG dataset is currently the most extensive and complete database of its kind, it still has its shortcomings. This is reflected in our limited ability to operationalize concepts such as the growth in cultural and external markers, or the rates of innovation. I believe that, as the Information Age matures, availability of data of this nature will only increase, in turn extending the possibilities for testing and validating models such as the one presented here.

6.7 Conclusion

The aim of this work was to explain the dynamics of scientific collaboration from a perspective of cultural evolution. I have developed and tested an agent-based model of scientific collaboration founded on social, cultural, as well as evolutionary principles and the complex interactions among them. I have validated the performance of our model against extensive empirical co-authorship data in two scientific disciplines. The results show that there are large differences in the underlying mechanisms of collaboration between scientific fields. They also demonstrate that under certain assumptions the proposed model fits well with the empirical data. However, further refinement of the model with respect to its parameters is needed, resulting in an antecedent need for higher quality and higher resolution data. Nevertheless, the model shows promise, as it has shown a good fit with the empirical evidence in the case of the Economics networks.

Chapter 7: Conclusion

7.1 Main Contributions

In this work I contribute to the study of culture by developing a formal framework for cultural evolution in the form of an agent-based model, its implementation in computer code, and its thorough and rigorous description, aided by the use of tools such as the ODD protocol. In this way, the model stands out from a series of agent-based models of culture, in that it is more formal, more likely reproducible, and more thoroughly justified in the implementation of particular assumptions and mechanisms (cf. Chapter 2).

Furthermore, I offer a stronger argument for the validity of dual inheritance theory and some of its postulated drivers of cultural evolution, namely indirectly biased transmission and guided variation. Dual inheritance theory is both established and formalized in terms of mathematical models and system dynamics. In this work, I have extended the validation effort to account for the complex webs of interaction in populations of social actors and demonstrated the emergence of empirically observed macro-phenomena from individual-level behavioral dynamics.

Finally, I show that my model of cultural evolution, which uses only two simple mechanisms, is robust in that it can explain empirical evidence coming from qualitatively different domains of human culture. Where other cultural models are specialized and narrow-focused, my model proves to be more general.

7.2 Discussion

The diffusion of cultural traits and the emergence of cultural signs is a complex phenomenon whose root causes are potentially numerous and multi-faceted. For this reason it is crucial to

adopt a complex systems, bottom-up approach to the modeling of its dynamics. In this work, I have taken advantage of several distinct computational social science approaches to better understand the problem. I have used agent-based modeling, evolutionary computation, network analysis, as well as statistical and Big Data analysis methods, coupled with a theoretical perspective adopted from social and biological anthropology.

Specifically, I focus on the effect of the use of external markers as signals for the underlying cultural attitudes and behaviors in agents' decisions regarding cooperation, coupled with the dynamics of cultural inheritance and individual learning. This work is a first step towards a more formal, computational study of the dynamics of the diffusion and reproduction of cultural values, beliefs and behaviors.

I set out to answer how the mechanisms of indirectly biased transmission and guided variation affect the distribution of cultural traits and external markers in a multi-agent population in the long run. The research in Chapter 3 shows that they indeed affect the dynamics significantly, and that both mechanisms contribute to the emergence of clustered and modular communities defined in terms of shared tags and trait variants. This effect was visible on small-world networks, however, not so much on other types of networks.

This prompted me to investigate how culture and network structure co-evolve. When I adjusted the model, to allow agents to dynamically adjust their neighborhoods, I found once again, that the two mechanisms were sufficient to produce social network that plausibly resemble real-world social networks. Moreover this resemblance vanished when the mechanisms were absent from the model.

I then sought to show whether a further modified version of the ABM could explain the dynamics of political polarization in the U.S. House of Representatives from a cultural perspective. The model demonstrated good fit with empirical evidence (bill co-sponsorship data from 1973-2008), and although it showed increasing levels of polarization in terms of a decline in average population tolerance it also showed, somewhat counter-intuitively, a decline in average ideological distance within the agent populations.

Finally, I set out to explore whether a modified version of the cultural evolution model

could be applied to explain the dynamics of scientific collaboration. I have compared the model’s performance to empirical evidence extensive records of academic publishing in the disciplines of Economics and Artificial Intelligence. The results showed that the cultures of collaboration are potentially very different across disciplines.

7.3 Summary of Results

The work presented here has revealed some interesting findings regarding the evolutionary drivers of cultural dynamics in groups of humans. Chapter 2 provided a broad survey of Evolutionary Computation methods and their use in agent-based models. To my knowledge it is the only such survey focusing on this topic, which is increasingly attracting more interest within computational social science. Chapters 3 and 4 then presented a detailed agent-based implementation of the dual-inheritance model of cultural evolution focused on the mechanisms of indirectly biased transmission and guided variance. In Chapter 4 these mechanisms were coupled with the agents’ ability of adjusting their social ties. Chapter 3 showed that the resulting distributions of cultural traits and tags heavily depend on the chosen network structures. In Chapter 4, I further demonstrated that in a dynamic social environment, the above mechanisms resulted into plausible configurations of social linkages, closely resembling small-world networks. Under a wide range of parameter settings, the model showed that the postulated evolutionary mechanisms lead to the association of certain tags with specific tags, thus giving rise to cultural signs.

Chapters 5 and 6 then extended the original model and validated it against empirical datasets. Chapter 5 extends the model to account for the ideological position of actors and uses a historical dataset (1973-2008) of bill co-sponsorship in the U.S. House of Representatives to test its performance. Experimentation has shown that the cultural evolution model performs better in reproducing the empirical networks than several alternatively formulated models. It also provides an explanatory framework for the increasing polarization and partisanship in U.S. politics. Finally, Chapter 6 extends the model to account for growing populations of actors and uses a scientific co-authorship dataset to validate the findings.

Interestingly, I find that the model performance varies significantly across different disciplines, suggesting that the drivers of scientific collaboration may vary from one field to another.

7.4 Limitations and Future Work

There are a few limitations to the current work. First, in the original model I never explore the effects of genetic drift and of kin selection, which, I believe, potentially contribute significantly to the resulting dynamics. Next, the robustness and viability of the different interaction strategies (unbiased, genetic, Lamarckian) has not been determined. This could potentially be done by simulating mixed populations with multiple strategies. In the polarization study, I use a one dimensional ideological scale, although recently multi-dimensional scales have been preferred in political science. It would be beneficial to test the sensitivity of the model to the dimensionality of the scale. As a final drawback, I was unable to properly calibrate several of the untested parameters in the collaboration study. This was limited by the granularity of the available empirical data.

Moreover, apart from the two datasets used in this dissertation, it will be interesting to see what other domains of culture are susceptible to the same vein of explanations that the cultural evolution model affords. Other possible sources of data include sharing economies, financial markets, communication networks, open-source collaborations, etc.

Other considerations are physical and environmental conditions limiting or supporting the spread of culture. For this reason it is worth examining the dynamics of the current model in a spatially explicit environment on different scales.

Furthermore, in the current model I have only considered the effect of a single cultural trait and a single external marker. In reality, every actor possesses a large number of these. Moreover, there are significant interactions among the individual traits. The effect of the interactions on the resulting population-level dynamics is potentially considerable. Thus, assessing the interplay of actor networks and cultural trait networks is a natural next step in this line of research.

7.5 Conclusion

I have carried out an extensive review of literature concerning Evolutionary Computation in ABMs of social systems. Subsequently, I developed an agent-based modeling framework of cultural evolution based on the co-evolution of hidden traits, independent external markers and network structure that shows qualitative alignment with empirical social networks. This agent-based framework is publicly available for replication and further use by other researchers. Finally, I applied the framework to the cases of political polarization in Congress as well as scientific collaboration, yielding promising results and different degrees of validation. The work contained herein is original in the sense that it presents a novel agent-based model founded on a previously unexplored combination of theoretical concepts relating to cultural evolution, and it benefits a broader audience by providing a unique view on some of the current cultural issues our society faces.

Appendix A: Model Version Overview

Table A.1: Overview of differences in implementation of selected model concepts in individual model versions

	Chapter 3	Chapter 4	Chapter 5	Chapter 6
Pop. size	fixed	fixed	fixed	variable (growth)
Network	fixed	dynamic	dynamic	dynamic
Cultural trait	categorical	categorical	cardinal	categorical
No. of variants/tags	fixed	fixed	fixed	variable (growth)
Survival	generational	generational	generational	age-based
Selection	tournament	tournament	tournament	local truncation
Recombination	2-parent	2-parent	cloning	multi-parent

Appendix B: Mathematical Representation of the Basic Agent-Based Model

In this appendix I outline a mathematical representation of the original agent-based model introduced in Chapter 3, with some commentary. In its most abstract form the model is an evolutionary algorithm EA , composed of its basic operators, acting on a sequence of coordination games between pairs of agents i and j :

$$EA = Mutation \circ Recombination \circ Selection \circ Survival \circ Evaluation(\{Coord(i, j)\}_{t=1}^{t_{max}}) \quad (B.1)$$

Here t denotes the time step and t_{max} is the number of rounds in each generation. In each time step the pair i, j is chosen as follows:

$$i = X(\omega) \quad \text{where} \quad X \sim U(A) \quad (B.2)$$

$$j = P_i^t(\{T_k\}_{k \in N_i}) \quad (B.3)$$

Here ω is an outcome of sampling from the random variable X distributed uniformly over the set of agents A . P_i^t is then the preference function of agent i in time step t acting on the set of tags T_k , possessed by some agent k in the agent's neighborhood N_i .

Furthermore, if we let Tr denote the trait variants of agents i and j , the payoffs from the coordination game can be expressed as:

$$Coord(i, j) = \begin{cases} [1, 1] & \text{if } Tr_i = Tr_j \\ [-1, -1] & \text{otherwise} \end{cases} \quad (\text{B.4})$$

Moreover, the functional P_i^t maps from the set of i 's base-level activations for the tags T_k at time t in its neighborhood to a single tag T :

$$P_i^t : \{BLA_{T_k}^t\}_{k \in N_i} \mapsto T \quad (\text{B.5})$$

$$T = X(\omega) \quad \text{where} \quad X \sim U(\{T_k | BLA_{T_k} > \tau\}) \quad (\text{B.6})$$

Here τ is some given threshold. As described in Chapter 3, the base-level activation for tag T is defined as follows:

$$BLA_T = \ln \left[\sum_{l=1}^n t_l^{-d} \right] \quad (\text{B.7})$$

Here t_l is the time elapsed since the agent's l -th positive experience with tag T , the total number of such interactions is n , and d is a decay parameter. Thus, the state of the model M at time t can be perhaps most tellingly written as:

$$M(t) = EA \left(\left\{ Coord(i, P_i^t(\{T_k\}_{k \in N_i})) \right\}_{t=1}^{t_{max}} \right) \quad (\text{B.8})$$

This notation reveals the model's reliance on the evolutionary algorithm, the coordination between the agents, and most importantly the guided variation mechanism (here represented by the time dependent preference mapping as a function of previous experiences) and the indirect bias mechanism (here represented by the set of tags as the only *direct* inputs to the preference mapping).

Appendix C: List of Abbreviations

- ABM** – agent-based model
- ACT-R** – Adaptive Control of Thought–Rational
- AI** – Artificial Intelligence
- CSS** – Computational Social Science
- CDF** – cumulative distribution function
- EA** – evolutionary algorithm
- EC** – Evolutionary Computation
- ES** – evolutionary strategy
- GA** – genetic algorithm
- GP** – genetic programming
- IEEE** – Institute of Electrical and Electronics Engineers
- JASSS** – Journal of Artificial Societies and Social Simulation
- MAG** – Microsoft Academic Graph
- MASON** – Multi-Agent Simulator of Neighborhoods
- ODD** – Overview, Design Concepts, Details (protocol)
- PD** – prisoner’s dilemma
- RMSE** – root-mean-square error

References

- Abbasi, A., Altmann, J., & Hossain, L. (2011). Identifying the Effects of Co-authorship Networks on the Performance of Scholars: A Correlation and Regression Analysis of Performance Measures and Social Network Analysis Measures. *Journal of Infometrics*, 5(4), 594-607.
- Abbasi, A., Chung, K. S. K., & Hossain, L. (2012). Egocentric Analysis of Co-authorship Network Structure, Position and Performance. *Information Processing and Management*, 48(4), 671-679.
- Abramowitz, A., & Saunders, K. (2008). Is Polarization a Myth? *The Journal of Politics*, 70(2), 542-555.
- Alizadeh, M., & Cioffi-Revilla, C. (2015). Activation Regimes in Opinion Dynamics: Comparing Asynchronous Updating Schemes. *Browser Download This Paper*.
- Alizadeh, M., Weber, I., Cioffi-Revilla, C., Fortunato, S., & Macy, M. (2016). Psychological and Personality Profiles of Political Extremists.
(Available at: <https://arxiv.org/abs/1704.00119>)
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An Integrated Theory of the Mind. *Psychological Review*, 111(4), 1036.
- Anderson, J. R., & Lebiere, C. (1998). *The Atomic Components of Thought*. Mahwah, NJ: LEA Publishers.
- Andris, C., Lee, D., Hamilton, M., Martino, M., Gunning, C. E., & Selden, J. A. (2015). The Rise of Partisanship and Super-Cooperators in the U.S. House of Representatives. *PLoS ONE*, 10(4), e0123507.

- Angeline, P. J. (2000). Parse Trees. In T. Bäck, D. Fogel, & T. Michalewicz (Eds.), *Evolutionary Computation 1: Basic Algorithms and Operators*. Bristol, United Kingdom: Institute of Physics Publishing.
- Ansolabehere, S., & Jones, P. E. (2010). Constituents Responses to Congressional Roll-Call Voting. *American Journal of Political Science*, 54(3), 583-597.
- Arthur, B. W. (1994). Inductive Reasoning and Bounded Rationality. *American Economic Review*, 84, 406-411.
- Atran, S. (2001). The Trouble with Memes: Inference Versus Imitation in Cultural Creation. *Human Nature*, 12, 351-381.
- Axelrod, R. (1984). *The Evolution of Cooperation*. New York, New York: Basic Books.
- Axelrod, R. (1986). An Evolutionary Approach to Norms. *The American Political Science Review*, 80(4), 1095-1111.
- Axelrod, R. (1997). The Dissemination of Culture: A Model with Local Convergence and Global Polarization. *Journal of Conflict Resolution*, 41(2), 203-226.
- Axelrod, R., & Hamilton, W. D. (1981). The Evolution of Cooperation. *Science*, 211, 1390-96.
- Bäck, T. (1996). *Evolutionary Algorithms in Theory and Practice*. New York, New York: Oxford University Press.
- Bäck, T., Fogel, D., & Michalewicz, T. (2000). *Evolutionary Computation 1: Basic Algorithms and Operators*. Bristol, United Kingdom: Institute of Physics Publishing.
- Baldassari, D., & Gelman, A. (2008). Partisans without Constraint: Political Polarization and Trends in American Public Opinion. *American Journal of Sociology*, 114(2), 408-446.
- Barabási, A.-L., & Albert, R. (1999). Emergence of Scaling in Random Networks. *Science*, 286(5439), 509-512.
- Barabási, A. L., Jeong, H., Neda, Z., Ravasz, E., Schebert, A., & Vicsek, T. (2002). Evolution of the Social Network of Scientific Collaborations. *Physica A: Statistical Mechanics and its Applications*, 311(3-4), 590-614.

- Barrat, A., Barthélemy, M., Pastor-Satorras, R., & Vespignani, A. (2004). The Architecture of Complex Weighted Networks. *Proceedings of the National Academy of Sciences of the United States of America*, 101(11), 3747-3752.
- Bausch, W. (2015). The Geography of Ethnocentrism. *Journal of Conflict Resolution*, 59(3), 510-527.
- Bendor, J., Diermeier, D., & Ting, M. (2003). A Behavioral Model of Turnout. *American Political Science Review*, 97(2), 261-280.
- Berger, P. L., & Luckmann, T. (1966). *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. Garden City, NY: Anchor Books.
- Bernhardt, D., Krasa, S., & Polborn, M. (2008). Political Polarization and the Electoral Effects of Media Bias. *Journal of Public Economics*, 92(5-6), 1092-1104.
- Bianchi, F., & Squazzoni, F. (2015). Agent-based Models in Sociology. *Wiley Interdisciplinary Reviews: Computational Statistics*, 7, 284306.
- Bourdieu, P. (1984). *Distinction: A Social Critique of the Judgement of Taste*. London, UK: Routledge.
- Bowles, S., & Gintis, H. (2004). The Evolution of Strong Reciprocity: Cooperation in Heterogeneous Populations. *Theoretical Population Biology*, 65, 17-28.
- Boyd, R., Bowles, S., & Gintis, H. (2010). Coordinated Punishment of Defectors Sustains Cooperation and Can Proliferate when Rare. *Science*, 328, 617-620.
- Boyd, R., & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago, Illinois: University of Chicago Press.
- Boyd, R., & Richerson, P. J. (2005). *Not By Genes Alone: How Culture Transformed Human Evolution*. Chicago, Illinois: University of Chicago Press.
- Bramson, A., Grim, P., Singer, D. J., Fisher, S., Berger, W., Sack, G., & Flocken, C. (2016). Disambiguation of Social Polarization Concepts and Measures. *Journal of Mathematical Sociology*, 40(2), 80-111.
- Campbell, J. E. (1982). Cosponsoring Legislation in the U.S. Congress. *Legislative Studies Quarterly*, 7(3), 415-422.

- Castelfranchi, C., Conte, R., & Paolucci, M. (1998). Normative Reputation and the Costs of Compliance. *Journal of Artificial Societies and Social Simulation*, 1(3).
- Cegielski, W. H., & Rogers, J. D. (2016). Rethinking the Role of Agent-Based Modeling in Archaeology. *Journal of Anthropological Archaeology*, 41, 283-298.
- Centola, D., Gonzalez-Avella, J. C., Eguiluz, V. M., & San Miguel, M. (2007). Cultural Drift, and the Co-Evolution of Cultural Groups. *Journal of Conflict Resolution*, 51(6), 905-929.
- Chattoe, E. (1998). Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes? *Journal of Artificial Societies and Social Simulation*, 1(3).
- Chiang, Y.-S. (2013). Cooperation Could Evolve in Complex Networks when Activated Conditionally on Network Characteristics. *Journal of Artificial Societies and Social Simulation*, 16(2).
- Cioffi-Revilla, C. (2017). *Introduction to Computational Social Science*. New York, New York: Springer.
- Conover, M., Ratkiewicz, J., Francisco, M. R., Gonçalves, B., Menczer, F., & Flammini, A. (2011). Political Polarization on Twitter. *ICWSM*, 133, 89-96.
- Dalton, R. (2008). The Quantity and Quality of Party Systems: Polarization, Its Measurement, and Its Consequences. *Comparative Political Studies*, 41(7), 899-920.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford, UK: Oxford University Press.
- Deb, K. (2000). Introduction to selection. In T. Bäck, D. Fogel, & T. Michalewicz (Eds.), *Evolutionary Computation 1: Basic Algorithms and Operators*. Bristol, United Kingdom: Institute of Physics Publishing.
- Deffuant, G., Huet, S., & F., A. (2005). An Individual-based Model of Innovation Diffusion Mixing Social Value and Individual Benefit. *American Journal of Sociology*, 110, 1041-1069.
- De Jong, K. A. (2005). *Evolutionary Computation: A Unified Approach*. Cambridge, Massachusetts: MIT Press.
- De Lange, C., & Glänzel, W. (1997). Modelling and Measuring Multilateral Co-authorship

- in International Scientific Collaboration. Part I. Development of a New Model Using a Series Expansion Approach. *Scientometrics*, 40(3), 593–604.
- De Stefano, D., Fuccella, V., Vitale, M. P., & Zaccarin, S. (2013). The Use of Different Data Sources in the Analysis of Co-authorship Networks and Scientific Performance. *Social Networks*, 35(3), 370–381.
- De Stefano, D., Giordano, G., & Vitale, M. P. (2011). Issues in the Analysis of Co-authorship Networks. *Quality and Quantity*, 45, 1091–1107.
- de Vos, H., Smaniotto, R., & Elsas, D. A. (2001). Reciprocal Altruism under Conditions of Partner Selection. *Rationality and Society*, 13(2), 139–183.
- Dosi, G., Marengo, L., Bassanini, A., & Valente, M. (1999). Norms as Emergent Properties of Adaptive Learning: The Case of Economic Routines. *Journal of Evolutionary Economics*, 9, 5–26.
- Downs, A. (1957). An Economic Theory of Political Action in a Democracy. *Journal of Political Economy*, 65(2), 135–150.
- Edmonds, B. (1997). *Modelling Bounded Rationality using Evolutionary Techniques* (Discussion Paper). Centre for Policy Modelling, Metropolitan University of Manchester. (<http://www.cfpm.org/papers/mbruet.pdf>)
- Efferson, C., Lalive, R., & Fehr, E. (2008). The Coevolution of Cultural Groups and Ingroup Favoritism. *Science*, 321, 1844–1849.
- Eiben, A. E., & Smith, J. E. (2007). *Introduction to Evolutionary Computing*. New York, New York: Springer.
- Endersby, J. W. (1996). Collaborative Research in the Social Sciences: Multiple Authorship and Paper Credit. *Social Sciences Quarterly*, 77, 375–392.
- Epstein, J., & Axtell, R. (1996). *Growing Artificial Societies*. Cambridge, Massachusetts: MIT Press.
- Erdős, P., & Rényi, A. (1959). On Random Graphs I. *Publicationes Mathematicae*, 6, 290–297.
- Erdos, P., & Rényi, A. (1960). On the Evolution of Random Graphs. *Publ. Math. Inst.*

- Hung. Acad. Sci*, 5(1), 17–60.
- Fiorina, M. P., Abrams, S. J., & Pope, J. C. (2006). *The Polarization of American Politics, 2nd Edition*. New York, NY: Pearson Longmann.
- Flache, A., & Macy, M. W. (2011). Local Convergence and Global Diversity: From Interpersonal to Social Influences. *Journal of Conflict Resolution*, 55(6), 970-995.
- Fogel, D. (1995). *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*. Piscataway, New Jersey: IEEE Press.
- Fogel, D. (2000). Real-valued vectors. In T. Bäck, D. Fogel, & T. Michalewicz (Eds.), *Evolutionary Computation 1: Basic Algorithms and Operators*. Bristol, United Kingdom: Institute of Physics Publishing.
- Fogel, L., Owens, A., & Walsh, M. (1966). *Artificial Intelligence through Simulated Evolution*. New York, New York: John Wiley and Sons.
- Fowler, J. H. (2006a). Connecting the Congress: A Study of Cosponsorship Networks. *Political Analysis*, 14, 456-487.
- Fowler, J. H. (2006b). Habitual Voting and Behavioral Turnout. *Journal of Politics*, 68(2), 335–344.
- Fowler, J. H. (2006c). Legislative Cosponsorship Networks in the US House and Senate. *Social Networks*, 28, 456-465.
- Fowler, J. H., & Laver, M. (2008). A Tournament of Party Decision Rules. *Journal of Conflict Resolution*, 52(1), 68–92.
- Fracchia, J., & Lewontin, R. C. (1999). Does Culture Evolve? *History and Theory*, 38, 52-78.
- Gabora, L. (2011). Five Clarifications about Cultural Evolution. *Journal of Cognition and Culture*, 11, 61-83.
- Galan, J. M., & Izquierdo, L. R. (2005). Appearances Can Be Deceiving: Lessons Learned Re-Implementing Axelrod’s ‘Evolutionary Approach to Norms’. *Journal of Artificial Societies and Social Simulation*, 8(3).
- Geertz, C. (1973). *The Interpretation of Cultures* (Vol. 5019). New York, NY: Basic books.

- Gilbert, N., & Troitzsch, K. G. (2005). *Simulation for the Social Scientist. Second Edition*. Maidenhead, England: Open University Press.
- Goldberg, D. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. New York, New York: Addison-Wesley.
- Gould, S. J. (1981). *The Mismeasure of Man*. New York, NY: W.W. Norton and Company.
- Gould, S. J., & Lewontin, C. R. (1971). The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 205(1161), 581-598.
- Gramsci, A. (2000). *The Antonio Gramsci Reader: Selected Writings 1916-1935* (D. Forgacs, Ed.). New York, NY: New York University Press.
- Granovetter, M. S. (1973). The Strength of Weak Ties. *American Journal of Sociology*, 78(6), 1360-1380.
- Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., ... DeAngelis, D. (2006). A Standard Protocol for Describing Individual-based and Agent-based Models. *Ecological Modeling*, 198(1), 115-126.
- Hales, D. (2000). Cooperation without Space or Memory: Tags, Groups and the Prisoner's Dilemma. In S. Moss & P. Davidsson (Eds.), *Multi-Agent-Based Simulation*. Berlin: Springer-Verlag.
- Hall, B. K., & Hallgrímsson, B. (2008). *Strickberger's Evolution*. New York, New York: Jones and Bartlett Publishers.
- Hamilton, W. D. (1963). The Evolution of Altruistic Behavior. *American Naturalist*, 97(896), 354-356.
- Hammond, R. A., & Axelrod, R. (2006). The Evolution of Ethnocentrism. *Journal of Conflict Resolution*, 50, 926-936.
- Hancock, P. J. B. (2000). A Comparison of Selection Mechanisms. In T. Bäck, D. Fogel, & T. Michalewicz (Eds.), *Evolutionary Computation 1: Basic Algorithms and Operators*. Bristol, United Kingdom: Institute of Physics Publishing.

- Henrich, J., & Boyd, R. (2001). Why People Punish Defectors: Weak Conformist Transmission Can Stabilize Costly Enforcement of Norms in Cooperative Dilemmas. *Journal of theoretical biology*, 208(1), 79–89.
- Henrich, J., Boyd, R., & Richerson, P. J. (2008). Five Misunderstandings About Cultural Evolution. *Human Nature*, 19, 119–137.
- Henrich, J., Boyd, R., & Richerson, P. J. (2012). The Puzzle of Monogamous Marriage. *Phil. Trans. R. Soc. B*, 367(1589), 657–669.
- Heppenstall, A. J., Crooks, A. T., See, L. M., & Batty, M. (2011). *Agent-based Models of Geographical Systems*. Springer Science & Business Media.
- Herrmannova, D., & Knoth, P. (2016). An Analysis of the Microsoft Academic Graph. *The Magazine of Digital Library Research*, 22(9-10).
- Hillis, W. D. (1990). Co-evolving Parasites Improve Simulated Evolution as an Optimization Procedure. *Physica D: Nonlinear Phenomena*, 42(1-3), 228–234.
- Hobbes, T. (2013 [1651]). *Leviathan*. Folkestone, United Kingdom: Renaissance Books.
- Holland, J. (1962). Outline for a Logical Theory of Adaptive Systems. *JACM*, 9, 297–314.
- Holland, J. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor, Michigan: University of Michigan Press.
- Holland, J. (1992). *Adaptation in Natural and Artificial Systems, 2nd Edition*. Cambridge, Massachusetts: MIT Press.
- Jacobson, G. C. (2000). Party Polarization in National Politics: The Electoral Connection. In *Polarized Politics: Congress and the President in a Partisan Era* (Vol. 5, pp. 17–18).
- Janssen, M. (2005). Evolution of Cooperation when Feedback to Reputation Scores is Voluntary. *Journal of Artificial Societies and Social Simulation*, 9(1).
- Kachel, A. F., Premo, L. S., & Hublin, J. J. (2010). Grandmothering and Natural Selection. *Proceedings of the Royal Society B Biological Sciences*, 278(1704).
- Kertzer, D. I. (1989). *Ritual, Politics, and Power*. New Haven, Connecticut: Yale University Press.

- Kessler, D., & Krehbiel, K. (1996). Dynamics of Cosponsorship. *The American Political Science Review*, 90(3), 555-566.
- Klemm, K., Eguiluz, V. M., Toral, R., & Miguel, M. S. (2003). Role of Dimensionality in Axelrods Model for the Dissemination of Culture. *Physica A: Statistical Mechanics and Its Applications*, 327(1-2), 1-5.
- Klos, T. B. (1999). Decentralized Interaction and Co-Adaptation in the Repeated Prisoners Dilemma. *Computational & Mathematical Organization Theory*, 5(2), 147-165.
- Kluch, S. P., & Vaux, A. (2017). Culture and Terrorism: The Role of Cultural Factors in Worldwide Terrorism (1970-2013). *Terrorism and Political Violence*, 29(2), 323-341.
- Koger, G. (2003). Position Taking and Cosponsorship in the U.S. House. *Legislative Studies Quarterly*, 28(2), 225-246.
- Kohler, T. A., Cockburn, D., Hooper, P. L., Bocinsky, R. K., & Kobti, Z. (2012). The Coevolution of Group Size and Leadership: An Agent-based Public Goods Model for Prehispanic Pueblo Societies. *Advances in Complex Systems*, 15(1-2), 1-29.
- Kollman, K., Miller, J. H., & Page, S. E. (1992). Adaptive Parties in Spatial Elections. *American Political Science Review*, 86(4), 929-937.
- Kossinets, G., & Watts, D. J. (2006). Empirical Analysis of an Evolving Social Network. *science*, 311(5757), 88-90.
- Koza, J. (1992). *Genetic Programming*. Cambridge, Massachusetts: MIT Press.
- Lachat, R. (2008). The Impact of Party Polarization on Ideological Voting. *Electoral Studies*, 27(4), 687-698.
- Lake, M. W. (2001). The Use of Pedestrian Modelling in Archaeology, with an Example From the Study of Cultural Learning. *Environment and Planning B: Planning and Design*, 28, 385-403.
- Lake, M. W. (2014). Trends in Archaeological Simulation. *Journal of Archaeological Method and Theory*, 21(2), 225-287.
- Lake, M. W., & Crema, E. R. (2012). The Cultural Evolution of Adaptive-trait Diversity

- when Resources are Uncertain and Finite. *Advances in Complex Systems*, 15(1-2), 1-19.
- Laland, K. N., & Brown, G. R. (2002). *Sense & Nonsense: Evolutionary Perspectives on Human Behavior*. Oxford, UK: Oxford University Press.
- Laver, M. (2005). Policy and the Dynamics of Political Competition. *American Political Science Review*, 99(2), 263–281.
- Laver, M., & Sergenti, E. (2011). *Party Competition: An Agent-based Model*. Princeton University Press.
- Lewontin, R., Kamin, L., & Rose, S. (1984). *Not in Our Genes: Biology, Ideology, and Human Nature*. New York, NY: Pantheon Books.
- Macy, M., & Skvoretz, J. (1998). The Evolution of Trust and Cooperation Between Strangers: A Computational Model. *American Sociological Review*, 63, 638-660.
- Masel, J. (2011). Genetic Drift. *Current Biology*, 21(20), R837R838.
- Mattias, K. E., & Whitley, L. D. (1994). Transforming the Search Space with gray Coding. In *Proceedings of the First IEEE Conference on Evolutionary Computation. IEEE World Congress on Computational Intelligence* (Vol. 1, p. 513-518).
- Mauss, M. (2000). *The Gift: The Form and Reason for Exchange in Archaic Societies*. WW Norton & Company.
- Mayhew, D. R. (1974). *Congress: The Electoral Connection, Yale Studies in Political Science*, vol. 12. New Haven, CT: Yale University Press.
- Miller, J. H. (1996). The Co-evolution of Automata in the Repeated Prisoner's Dilemma. *Journal of Economic Behavior and Organization*, 29(1), 87-112.
- Mitchell, M. (1998). *An Introduction to Genetic Algorithms*. Cambridge, Massachusetts: MIT Press.
- Moody, J. (2004). The Structure of a Social Science: Disciplinary Cohesion from 1963 to 1999. *American Sociological Review*, 69, 213-238.
- Moravec, H. (1987). Human Culture: A Genetic Takeover Underway. In C. G. Langton (Ed.), *Artificial Life: The Proceedings of an Interdisciplinary Workshop on the Synthesis and*

- Simulation of Living Systems, Los Alamos, New Mexico, September 1987* (Vol. VI, p. 167-199). Redwood City, CA: Addison-Wesley.
- Morris, S. A., & Goldstein, M. L. (2007). Manifestation of Research Teams in Journal Literature: A Growth Model of Papers, Authors, Collaboration, Coauthorship, Weak ties, and Lotka's Law. *Journal of the Association for Information Science and Technology*, 58(12), 1764-1782.
- NASA. (1999). *Mars Climate Orbiter Mishap Investigation Board Phase I Report* (Technical Report). (Retrieved 10/23/17 from: http://sunnyday.mit.edu/accidents/MCO_report.pdf)
- Newman, M. E. J. (2004a). Analysis of Weighted Networks. *Physical Review E*, 70, 056131.
- Newman, M. E. J. (2004b). Coauthorship Networks and Patterns of Scientific Collaboration. *Proceedings of the National Academy of Sciences of the United States of America*, 101(1), 5200-5205.
- Newman, M. E. J. (2006). Modularity and Community Structure in Networks. *Proceeding of the National Academy of Sciences USA*, 103(23), 8577-8696.
- Parisi, D., Cecconi, F., & Natale, F. (2003). Cultural Change in Spatial Environments: The Role of Cultural Assimilation and Internal Changes in Cultures. *Journal of Conflict Resolution*, 47(2), 163-179.
- Pepper, J. W., & Smuts, B. (2000). The Evolution of Cooperation in an Ecological Context: An Agent-based Model. In G. G. J. Kohler T. A. (Ed.), *Dynamics in Human and Primate Societies: Agent-Based Modelling of Social and Spatial Processes* (p. 45-76). New York, New York: Oxford University Press.
- Perez-Losada, J., & Fort, J. (2011). Spatial Dimensions Increase the Effect of Cultural Drift. *Journal of Archaeological Science*, 38(6), 1294-1299.
- Perianes-Rodriguez, A., Olmeda-Gomez, C., & Moya-Anegon, F. (2010). Detecting, Identifying and Visualizing Research Groups in Co-authorship Networks. *Scientometrics*, 82(2), 307-319.

- Petrov, A. A. (2006). Computationally Efficient Approximation of the Base-Level Learning Equation in ACT-R. In D. Fun, F. Del Missier, & A. Stocco (Eds.), *Proceedings of the Seventh International Conference on Cognitive Modeling*. Trieste, Italy: Edizioni Goliardiche.
- Poole, K. T., & Rosenthal, H. (1984). Culture War? The Myth of a Polarized America. *The Journal of Politics*, 46(4), 1061-1079.
- Potter, M. A., & De Jong, K. A. (2000). Cooperative Coevolution: An Architecture for Evolving Coadapted Subcomponents. *Evolutionary computation*, 8(1), 1-29.
- Poundstone, W. (1992). *Prisoner's Dilemma: John von Neumann, Game Theory, and the Puzzle of the Bomb*. Oxford, United Kingdom: Oxford University Press.
- Powell, A., Shennan, S. J., & Thomas, M. G. (2009). Late Pleistocene Demography and the Appearance of Modern Human Behavior. *Science*, 324, 1298-1301.
- Premo, L. S., & Hublin, J.-J. (2009). Culture, Population Structure, and Low Genetic Diversity in Pleistocene Hominins. *Proceedings of the National Academy of Science*, 106(1), 33-37.
- Prior, M. (2013). Media and Political Polarization. *Annual Review of Political Science*, 16, 101-127.
- Quattrociocchi, W., Scala, A., & Sunstein, C. R. (2016). *Echo Chambers on Facebook* (Discussion Paper). (Available at SSRN: <https://ssrn.com/abstract=2795110>)
- Rappaport, R. A. (1967). Ritual Regulation of Environmental Relations among a New Guinea People. *Ethnology*, 6(1), 17-30.
- Rechenberg, I. (1965). Cybernetic Solution Path of an Experimental Problem. In *Library Translation 1122*. Farnborough: Royal Aircraft Establishment.
- Reynolds, R., Peng, B., & Ali, M. Z. (2008). The Role of Culture in the Emergence of Decision-making Roles: An Example Using Cultural Algorithms. *Complexity*, 13(3), 27-42.
- Reynolds, R. G. (1994). An Introduction to Cultural Algorithms. In *Proceedings of the 3rd annual conference on evolutionary programming*. Singapore: World Scientific

Publishinge.

- Riolo, R. (1997). *The Effects of Tag-Mediated Selection of Partners in Evolving Populations Playing the Iterated Prisoner's Dilemma* (Working Paper No. 97-02-016). Santa Fe, NM: Santa Fe Institute. (<http://www.cfpm.org/papers/mbruet.pdf>)
- Rohde, D. W. (2010). *Parties and Leaders in the Postreform House*. University of Chicago Press.
- Rousseau, J. J. (1968 [1762]). *The Social Contract*. New York, New York: Penguin Classics.
- Saam, N. J., & Harrer, A. (1999). Simulating Norms, Social Inequality, and Functional Change in Artificial Societies. *Journal of Artificial Societies and Social Simulation*, 2(1).
- Sahlins, M. (1960). *Evolution and Culture*. Ann Arbor, MI: University of Michigan Press.
- Sahlins, M. (1976). *The Use and Abuse of Biology: An Anthropological Critique of Sociobiology*. Ann Arbor, MI: University of Michigan Press.
- Santos, F. C., Pacheco, J., & Lenaerts, T. (2006). Cooperation Prevails when Individuals Adjust Their Social Ties. *PLoS Computational Biology*, 2, e140.
- Sarma, J., & De Jong, K. (2000). Generation Gap Methods. In T. Bäck, D. Fogel, & T. Michalewicz (Eds.), *Evolutionary Computation 1: Basic Algorithms and Operators*. Bristol, United Kingdom: Institute of Physics Publishing.
- Saussure, F. (1986 [1916]). *Course in General Linguistics* (C. Bally, A. Sechehaye, & A. Riedlinger, Eds.). Chicago, Illinois: Open Court.
- Schelling, T. C. (1971). Dynamic Models of Segregation. *Journal of Mathematical Sociology*, 1, 143-186.
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *AT&T Technical Journal*, 27(3), 379-423.
- Shutters, S. (2009). Strong Reciprocity, Social Structure, and the Evolution of Fair Allocations in a Simulated Ultimatum Game. *Computational & Mathematical Organization Theory*, 15(2), 64-77.
- Simon, H. A. (1996). *Sciences of the Artificial, 3rd Edition*. Cambridge, Massachusetts:

- MIT Press.
- Sinha, A., Shen, Z., Song, Y., Ma, H., Eide, D., Hu, B. J., & Wang, K. (2015). An Overview of Microsoft Academic Service (MAS) and Applications. In *Proceedings of the 24th International Conference on World Wide Web (WWW 15 Companion)*. New York, NY: ACM.
- Takahashi, N. (2000). The Emergence of Generalized Exchange. *The American Journal of Sociology*, 105(4), 1105-1134.
- Thompson, J. N. (1994). *The Coevolutionary Process*. University of Chicago Press.
- Tocqueville, A. (2000). *Democracy in America*. Chicago: University of Chicago Press.
- Tohme, P. (1992). *Cultural Development and Environment*. Paris: UNESCO.
- Velden, T., Haque, A., & Lagoze, C. (2010). A New Approach to Analyzing Patterns of Collaboration in Co-authorship Networks: Mesoscopic Analysis and Interpretation. *Scientometrics*, 85, 219-242.
- Watts, D. J., & Strogatz, S. H. (1998). Collective Dynamics of “Small-world” Networks. *Nature*, 393(6684), 440-442.
- Weber, M. (1978). *Economy and Society: An Outline of Interpretive Sociology* (Vol. 1). University of California Press.
- Williams, G. C. (1957). Pleiotropy, Natural Selection, and the Evolution of Senescence. *Evolution*, 11(4), 398-411.
- Williams, R. (1978). *Marxism and Literature*. Oxford, UK: Oxford Paperbacks.
- Wilson, E. O. (1975). *Sociobiology: The New Synthesis*. Cambridge, MA: Harvard University Press.
- Xue, J. Z., Costopoulos, A., & Guichard, F. (2011). Choosing Fitness-Enhancing Innovations Can Be Detrimental under Fluctuating Environments. *PLoS ONE*, 6(11), e26770.
- Zhang, Y., Friend, A. J., Traud, A. L., Porter, M. A., Fowler, J. H., & Mucha, P. J. (2008). Community Structure in Congressional Cosponsorship Networks. *Physica A: Statistical Mechanics and its Applications*, 387(7), 1705-1712.

Curriculum Vitae

Peter Revay has earned bachelor's degrees in Mathematics and Sociology from the Masaryk University in Brno, Czech Republic in 2012. He then attended University of Vermont where he received his Master of Science in Mathematics along with a Certificate in Complex Systems in 2014. In 2015 he has also received a Master of Science in Sociology from the Masaryk University. Peter started his PhD in Computational Social Science at the George Mason University in 2014. He has served as Graduate Research Assistant on multiple projects for Dr. Claudio Cioffi-Revilla.