



**EVOLUTIONARY
AGENT-BASED
POLICY ANALYSIS
IN DYNAMIC
ENVIRONMENTS**

Volker Nannen

promotors: prof. dr. J. C. J. M. van den Bergh
 prof. dr. A. E. Eiben

thesis committee: prof. dr. T. Bartz-Beielstein
 dr. M. C. Schut
 prof. dr. S. Solomon
 prof. dr. J. Treur
 prof. dr. C. A. A. M. Withagen

The research reported in this thesis was supported by the VUBIS project of The Network Institute, Vrije Universiteit, Amsterdam, the interdisciplinary research collaboration between the faculties of Exact Sciences, Social Sciences, and Economic Sciences and Business Administration.

A research visit to the Institute for Scientific Interchange, Torino, was supported by E.C. NEST Contract n. 018474-2 on "Dynamic Analysis of Physiological Networks" (DAPHNet).



Dissertation Series No. 2009-08

The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

© 2009 by Volker Nannen
All rights reserved. Published 2009

Cover design: Marc Ramírez Camps, Barcelona

Printed by Ipskamp Drukkers, Enschede

ISBN: 978-90-9024101-2

VRIJE UNIVERSITEIT

EVOLUTIONARY AGENT-BASED
POLICY ANALYSIS IN
DYNAMIC ENVIRONMENTS

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad Doctor aan
de Vrije Universiteit Amsterdam,
op gezag van de rector magnificus
prof.dr. L.M. Bouter,
in het openbaar te verdedigen
ten overstaan van de promotiecommissie
van de faculteit der Exacte Wetenschappen
op donderdag 16 april 2009 om 10.45 uur
in de aula van de universiteit,
De Boelelaan 1105

door

Volker Nannen

geboren te Norden, Oostfriesland

promotoren: prof.dr. J.C.J.M. van den Bergh
 prof.dr. A.E. Eiben

To Harald & Tjadine

*And crawling on the planet's face
Some insects called the human race
Lost in time, and lost in space
And meaning.*

—Richard O'Brian

CONTENTS

Preface	11
Acknowledgments	13
1 Overview	15
1.1 Background and motivation	15
1.2 Research objectives	17
1.3 Methods	18
1.4 Thesis outline	20
2 Relevance Estimation and Value Calibration of Evolutionary Algorithm Parameters	23
2.1 Background	24
2.2 The algorithm	26
2.2.1 Approaching the maximum entropy distribution	26
2.2.2 Algorithm implementation	27
2.2.3 Interpreting the measurements	29
2.3 Assessing the reliability of REVAC estimates	30
2.3.1 . . . on abstract objective functions	30
2.3.2 . . . on a simple genetic algorithm	35
2.4 Assessing the algorithmic efficiency of REVAC	37
2.4.1 . . . on a simple genetic algorithm	38
2.4.2 . . . on an economic modeling problem	43
2.5 Comparing REVAC to other tuning methods	46
2.6 Conclusions	48
3 A Study of Parameter Relevance in Evolutionary Algorithms	51
3.1 Introduction	51
3.2 Experimental setup	53
3.3 How does the choice of operator contribute to performance?	55
3.4 Which EA component needs the most tuning?	62
3.5 Conclusions	65

4	How to Evolve Strategies in Complex Economy-Environment Systems	67
4.1	Introduction	67
4.2	The economic model	69
4.2.1	General features of the model	69
4.2.2	Strategies, investment, and production	69
4.2.3	The social network	71
4.2.4	The evolutionary mechanism	72
4.3	Experiments	75
4.3.1	Evaluating the initial evolutionary model	75
4.3.2	Evaluating a simplified evolutionary model	76
4.4	Conclusions	79
5	Impact of Environmental Dynamics on Economic Evolution	81
5.1	Introduction	81
5.2	The economic model	84
5.2.1	General features of the model	84
5.2.2	Strategies, investment, and production	84
5.2.3	The evolutionary mechanism of behavioral interactions	86
5.3	The evolutionary dynamics	87
5.3.1	The growth rate of a strategy	87
5.3.2	Efficiency and level sets of investment strategies	89
5.4	Experimental setup	91
5.4.1	The environmental dynamics	91
5.4.2	Implementation details, model calibration, and scaling	94
5.5	Results	96
5.5.1	Economic significance of diversity	96
5.5.2	Policy advice under uncertainty	97
5.5.3	Evolutionary dynamics	99
5.6	Conclusions	100
5.A	Evolution with variable prices	101
6	Policy Instruments for Evolution of Bounded Rationality	109
6.1	Introduction	109
6.2	The economic model	111
6.2.1	General features of the model	111
6.2.2	Strategies, investment, and production	112
6.2.3	Evolution of strategies	115
6.2.4	Policy goals and formulation	117
6.2.5	Model calibration	118
6.3	The evolutionary dynamics	120
6.3.1	Derivation of the growth function	120
6.3.2	Convergence behavior	122
6.4	Policy analysis	124
6.4.1	Experimental setup	124
6.4.2	Evaluating the first best policy, a tax on fossil energy investment	126

6.4.3	Evaluating the <i>prizes</i> policy	126
6.4.4	Evaluating the <i>advertisement</i> policy	130
6.5	Conclusions	130
7	Summary and Conclusions	133
	Symbols	135
	References	137

PREFACE

Little did I know what I was embarking on when I fired off that email to Guszti Eiben, professor of story-writing, suggesting to use evolutionary artificial societies to study economic processes. Not only did he want to do exactly that, but his colleague Jeroen van den Bergh, professor of scientific debate, had already secured funding for a PhD position and, surprise, was willing to take me as a student, despite my ignorance of economic theory. Soon I was running agent-based simulations right and left, leaving far behind what was considered the state of the art in general equilibrium theory. My hard disk was filling with gigabytes of exciting data. Only that they made no sense. The artificial agents were dying. Or their technology went through the roof. Or their welfare exploded. We used different parameters. Different methods. Different ideas. To no avail.

If mainstream economists frown upon evolutionary agent-based simulations, here I am, shaking their hand and hugging their shoulder: they have every reason to do so. These beasts are unruly, unpredictable, incomprehensible, nasty, and mean. Only a madman can believe in them. Or visionaries like Guszti and Jeroen. I myself threw in the towel not only once but twice! Then I had the good fortune to meet Sorin Solomon, professor of good hope, and David Bree, professor of the high seas, who together pulled me back on the firm grounds of wonder and astonishment. You can't really control what you get out of a complex system by what you are putting in. That's why it's called complex. There is some break of causality.

Still, there is no denying it: real economic phenomena emerge from the complex interaction of opinionated and irresolute human beings like me, while the super-hero of rational choice theory really belongs to the comic book. The little headway I could make in taming this raging madness, nothing but charts of shallow waters around an abysmal depth of mystery, is contained within this thesis. I hope that it will be useful to you, reader, if only by warning you of obstacles that lie ahead.

This thesis would not be complete without a word on interdisciplinary research. By this we mean a holistic and systematic research plan that requires scientific advances in unrelated disciplines. I think Sorin has formulated it well: this is where things really get interesting; but unless one enjoys the luxury of a four year PhD contract, it can hardly be done. By and large, research gets funded if it can be published in high ranking journals. Differences in culture, language, style, methods, and expectations make it extremely difficult to write articles that will be accepted by high ranking journals of more than one discipline. From the point of view of academic career building, this cancels the added value of interdisciplinary research. If this is your choice, be prepared for painful misunderstandings, and make sure you can afford to fail.

ACKNOWLEDGMENTS

Thank is due first of all to those that have initiated this research project and brought it to a successful end, Jeroen van den Bergh and Gusztai Eiben. Their guidance and experience is felt throughout the thesis.

Sorin Solomon has repeatedly stepped in and provided perspective, without which I would have failed. David Bree has given much practical advice, and considerably improved the quality of my writing.

I had the fortune to supervise the Master's project of Alex de Landgraaf, and he has contributed Section 2.5. Selmar Smit has contributed to Section 3.3. Martijn Schut, Geoff Nitschke, Joe Wakeling, Jose Ramasco, Simona Cantono, Leanne Ussher, Marco Lamieri, Damien Challet, Gur Yaari, Nadav Shnerb, Tibor Zavadil, Gilles Daniel, Bence Toth, Giancarlo Mosetti, Wendie Shaffer, Tom Erez, and Uzay Kaymak have all given valuable comments.

I am honored to be examined by Cees Withagen, Jan Treur, Martijn Schut, Sorin Solomon, and Thomas Bartz-Beielstein, the learned members of my thesis committee.

Marc Ramírez Camps has captured everything I would have liked to say on economic behavior in his expressive cover art.

The life of a PhD student is a unique mixture of independent scientific investigation, professional creativity, and irresponsible leisure, particularly when you share a flat, or live door to door, with like minded people. I was lucky to share this experience with Stefano Bocconi, Jeroen Groenenboom, Sara Mautino and Davide Moratti, Peter Hofgesang and Maya Joosz, Simona Montagnana and Paolo Zeppini, Marcin Zukowski, Krzysztof Pietrzak, Christian Shaffner and Sonja Peters.

Others have joined my quest for the meaning of life on a more professional level, from an agent-based point of view, always with good humor: Abraham Prieto, Alessandro Cappellini, Anastasia Iartseva, Anna Deters, Christian Tzolov, Diana Mangalagiu, Enrico Scalas, Evert Haasdijk, Federico Divina, Frits Daalmans, Guy Kelman, Jan-Peter Patist, Julia Aronson, Lev Muchnik, Luís Simões, Robert Griffioen, and Samir Hamichi.

Many friends have made my life interesting during these years and joined me for sailing, hiking, or dinner. I take the opportunity to thank each of them: Arjen de Vries, Dorina Jibetea and Miro Kramar, Drahusa Bucsekova, Elena Rangelova, Eli Castejón and Matthias Wozel, Enrico Conti, Falk Bartels and Triu Romero, Farah Berri, Farbod Kayhan, Gabor Maroti, Gea Pesari, Julian Garcia, Levan Tavartkiladze, Lubor Frastia, Maarten Grachten, Meritxell Espín and Sergio Fernández, Nori Jacoby, Roberto Cornacchia, Roby Palermo, Sabine and Torsten Sandmann, Siebo Bogena, Stefan Manegold, Tamás Bíró, Thijs Westerveld, Tslil ben Akiva, Zhao Yang, and the ISI football club in Torino.

I am grateful to Joseph and Yemima Cohen, Andy Pettman, Givi Taktakishvili, Anna Gvasalia and Tamazi Sukhishvili, Manana Tsiramua and Zura Getsadze for their sincere friendship and inspiration.

I am thankful to Maria, Eli, Jaume, Ànnia, Martí, Aina, Oriol, Anna, Abril, Marc, Pau, and Bernat for their warm hospitality and encouragement.

I am indebted to Harald, Tjadine, Edzard, Tido and Bea for their constant support and faith.

I thank Joel for his smiles and laughter.

But most of all I thank Nina for her love.

OVERVIEW

1.1 Background and motivation

Human society is arguably the most complex and fascinating system that can be studied with scientific methods. Composed of intelligent beings with a free will and mind, there seems to be no limit to the ways it can organize itself. From the grassroots democracy of a Swiss canton to the gun culture of a Brazilian shanty town, there are tight networks of personal ties and social interaction through which beliefs are propagated, trust is built, and culture is formed. They deeply embed decision makers at all levels of society, from the director of a primary school to the executive boards and inner circles of large corporations and ruling parties. They certainly reign over the world of academic and religious institutions. No consensus is reached and no rule is established if not through the relevant network.

Social networks are neither regular nor random. They are the result of a development process steered by geographic proximity, shared history, ethnic and religious affiliation, common economic interests, and much else. Yet despite its preeminence and dominance in human society, the scientific study of social networks is relatively young. An early contribution was certainly the “six degrees of separation” by Milgram (1967), but it took another 30 years until the complex nature of social networks came to full attention (Albert and Barabási, 2002; Dorogovtsev and Mendes, 2002), in the sense that the collective behavior of agents in the network is fundamentally different from the average behavior of the individuals that compose it.

A recurrent question in the study of social networks is how they help a society to process information that allows them to react and adapt to a changing environment. Understanding this adaptive learning process is of great economic importance, as it can

help a government to prepare for environmental change, and to correctly anticipate the consequences of public intervention. For example, an appropriate response to climate change and associated problems like flooding and desertification might then be developed, and, if the climate change is caused by human behavior, it might even be curbed.

But as promising as this type of research is, it breaks with important traditional economic assumptions, and requires a new type of economic modeling. Mainstream economic theory, in particular the neoclassical school, is based on two abstractions: the rational agent and the representative agent. The first abstraction is based on the idea that the mathematical concept of a rational strategy approximates real human strategies without significant bias and with an acceptable level of inaccuracy. Provided that this rational strategy exists and that it can be computed, it avoids the uncertainty and ambiguity that is inherent to any cognitive model of human decision making. The second abstraction removes the diversity and social interdependence of a multi-agent systems, either by aggregating all group attributes and projecting them onto a single representative agent, or by observing the emergent behavior of the society and projecting it onto a single representative agent.

The representative agent is in fact required for the rational strategy to be a well defined mathematical object, as the equilibria of even trivially simple multi-agent systems are often not computable, if they exist at all. Together, these abstractions allow for exact solutions in those systems where the dominant strategies can be computed, and their analytic treatment can be used to extrapolate the behavior of a society under given technological or environmental change, as well as the regulative power of a public policy. But as Kirman (1992) has pointed out, “the reaction of a rational representative agent to change need not reflect how the rational individual would respond to change, and the preferences of a representative agent over choices may be diametrically opposed to those of society as a whole.” Furthermore, the idea that rational strategies constitute an unbiased approximation of real human strategies is refuted by a growing and widely recognized body of scientific (experimental and empirical) evidence (e.g., Kahneman et al., 1982; Camerer et al., 2003).

Real collective behavior depends on how information is processed in the social network and how individual beliefs and strategies are adapted over time, yet a policy designed for a rational and representative agent cannot account for the fact that information is not the same for all agents, and is oblivious to the speed and cost of changing a strategy. A policy that is designed with the rational and representative agent in mind cannot be expected to have the intended effects when applied to a real economy (Wegner and Pelikan, 2003). But not only do neoclassical economic models lead to unrealistic conclusions on policies that they can study, they also exclude an entire class of public policies from the analysis, namely those policies that explicitly target the diversity of strategies and social structure. Exemplary reward and punishment of single individuals, as well as the forming and breaking of personal ties between decision makers, such policies have been found to be highly effective policy instruments since the dawn of human statehood. Yet they are inaccessible to neoclassical economic theory. One simply cannot target individual behavior if there is no heterogeneity in the model.

One important aspect of human decision making that cannot be modeled by the representative agent is the evolution of a behavior or strategy as it is passed from agent to agent by way of imitation (Nelson and Winter, 1982; Boyd and Richerson, 1985; Hof-

bauer and Sigmund, 2003). Being able to observe different agents and to imitate the strategy of one of them allows an agent to extend its empirical horizon and to draw on the collective experience of the group. An agent may have no information on the world it lives in, or may not understand it, but if it can imitate the strategy of another agent that fares well it can still be expected to fare likewise. For an agent that has the information and capability to develop a rational strategy, the time and effort needed to do so does not necessarily pay off when a strategy that can readily be imitated achieves the same results. Even the most rational of agents will evolve their strategies by imitation when the utility of an action can only be established empirically through testing, as in such cases an evolutionary approach that varies and recombines a set of candidate solutions often produces the best results.

1.2 Research objectives

Here we investigate the impact of environmental dynamics on social systems with behavioral interactions by means of evolutionary computational experiments. We will study how to model the evolution of investment behavior by imitation in a social network, and we will use the resulting model to generate general insights and methods for the design and evaluation of public policies in an environment that is dynamic. The dynamics can be resource related, ecological, or technological, in which case a policy can guide the process of adaptation. The dynamics can also be policy related, in which case an evolutionary agent-based analysis can help to understand the temporal effects of introducing a policy. The research is fundamental in that it explores the general difficulties and opportunities that arise from applying the methods of evolutionary computation and agent-based modeling to evolutionary and behavioral economics.

Just as the individual agent can be expected to reason in a way that is in its best interest, so a group of agents can be expected to interact in a way that is in their best interest. However, while the individual agent is free to develop a new and elaborate strategy for every situation, the rules by which a group of agents interact can neither change too fast nor can they be too complicated in order to be agreed upon. In fact, the greater the group, the more simple and static the rules need to be in order to be commonly accepted. This implies that while a multi-agent model of social interaction and imitation needs to allow the agents to achieve realistic income growth rates in a wide array of economic situations, it has to be simple and needs to be evaluated against the amount of information that it requires the agents to hold. To quote Einstein, the model we are looking for has to be as simple as possible, but not simpler.

Our first objective is therefore to design a simple and robust agent-based model of imitation in a social network that can be used for an evolutionary policy analysis by numerical simulation. To this end we will try to identify the essential components of an evolutionary algorithm in general, and of evolution by imitation in particular. Our second objective is to use the agent-based model to study how the imitated strategies evolves under different environmental dynamics—perhaps a public policy can use such understanding to optimize the evolutionary mechanism to a particular environmental dynamic. For example, desertification is typically a slow process with long lasting consequences, while a pest can disappear as sudden as it appeared, and each might require its

own mode of adaptation from an agricultural community. Our third and final objective is to study whether a simple and robust evolutionary agent-based model can be used to design and evaluate a new type of public policy that explicitly takes the evolutionary aspects of imitation into account.

1.3 Methods

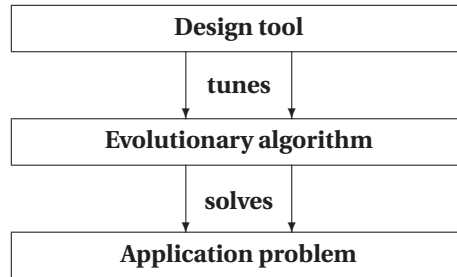
Whether a good strategy is available for imitation depends on how diversity of strategies is maintained in the population and on how information flows through the social network. The survival probability of entities that proliferate by an autocatalytic process, which includes the spread of economic strategies by imitation, depends crucially on the discrete nature of the quantities involved (Shnerb et al., 2000), as well as on the spatial structure of their hosting environment (Lieberman et al., 2005; Louzoun et al., 2007). Evolution of strategies by imitation is therefore best studied in agent-based simulations that pays proper attention to the discrete nature of strategies and agents, as well as the social structure.

According to Lehtinen and Kuorikoski (2007), a major hurdle in introducing agent-based simulations to mainstream economics is the fact that they yield messy data of unclear dependence. While it is possible in principle to assess the importance of any given parameter of a simulation model by running different simulations with one parameter fixed at a time, this is usually impractical because of the amount of computation required, the volume of the resulting data, and interaction between different parameter values. There is a lack of efficient statistical tools that can tell whether parameter values and simulation details are crucial for the results.

This problem is certainly true for a model where economic behavior evolves by imitation in a social network. Evolutionary algorithms form a rich family of stochastic search methods that use the Darwinian principles of variation and selection to incrementally improve a set of candidate solutions. Originally developed to solve computationally hard optimization problems, they can also be used to model real world phenomenon like the evolution of economic strategies. As has first been recognized by Grefenstette (1986), the design of an evolutionary algorithm for a specific application is itself a hard optimization problem that requires its own methodology and tools. This is visualized in the hierarchy of Figure 1.1, where a design tool tunes an evolutionary algorithm, which in turn solves the application problem. In our case the evolutionary algorithm is a model of evolution by imitation, and the application problem is an economic problem like finding an investment strategy that leads to high individual welfare under given technological or environmental dynamics.

Since almost all existing design tools are meant to maximize the performance of an evolutionary algorithm, when it comes to design goals like robustness and simplicity our options are rather limited. Robustness can be achieved by defining a reasonably wide array of application problems, and tuning the evolutionary algorithm such that the agents can adapt, i.e., evolve, reasonably well to all of them. We will frequently use this method. Simplicity is a more challenging design objective. The scientific literature on the design of evolutionary algorithms does not address it, and neither does the literature on the design of experiments. Only in the field of statistical inference has a method-

Figure 1.1: Design hierarchy of an evolutionary algorithm



ological framework developed, the Minimum Description Length principle (Grünwald, 2007). This methodology uses information theory to measure the simplicity of a statistical hypothesis. We will take a similar approach and use information theory to measure the simplicity of an evolutionary mechanism.

As for the application level of Figure 1.1, we will try to strike a balance between the need for a well defined application problem that poses a sufficiently realistic challenge to the agents, and our overall goal to achieve some general understanding on evolutionary agent-based policy analysis in dynamic environments. Global warming is widely considered to be the most acute dynamic economic problem today, and it combines many issues that are difficult if not impossible to address by a neoclassical economic model: it is characterized by a high degree of uncertainty and disagreement with regard to both the cause and the consequences of global warming; the distribution of responsibility (in terms of greenhouse gas emissions) and vulnerability is highly skewed; decision makers are entrenched in established procedures and beliefs; finally, there is no central authority that can punish free-riders. We base the definition of our application problem—that is, of the technological and environmental dynamics that the agents have to adapt to—on the influential work of W. D. Nordhaus who published a series of general-equilibrium economic models of climate policy and global warming, starting with the DICE model (Nordhaus, 1992). For the sake of analytic clarity we remove or simplify those elements of his model that are not essential to our current study.

Unlike Nordhaus, we consider the rationality of all agents to be bounded, and their information to be limited. They can only compare some properties of their fellow agents and use this information to imitate a strategy. Consider a population of several hundred agents, a number that is sufficiently large to allow for a rich social structure and diversity in strategies, income, and wealth. The agents allocate their respective income over a finite number of n investment sectors. Mathematically, such allocations have the precise definition of points in the n -dimensional simplex, and allow for a straightforward formulation of operators for variation and recombination. Standard economic growth and production functions describe how capital accumulates in each sector and contributes to income. These functions are not aggregated: growth and returns are calculated independently for each agent and two agents with different investment strategies can experience very different growth rates and income levels. A non-aggregate model preserves the functional relationship between individual investment strategies and the corresponding economic performance and allows us to model the evolution of strategies by imitation:

an agent can select another agent based on a property that is indicative of its current or future economic performance (the phenotype) and imitate its investment strategy (the genotype), thereby increasing the frequency of the imitated strategy, or at least its proportional input to new strategies if imitation is implemented as recombination of existing strategies.

When working with numerical methods, we have to account for a number of complicating factors that make it difficult to obtain clear and useful results. These include the non-deterministic nature of the evolutionary process, the autocatalytic character of the imitation dynamics, and the large number of free and unspecified parameters. Rather than closely calibrating those parameters that affect our results on a specific set of empirical data, we define broad parameter ranges and collect statistical information over a representative sample of possible economies that fall within these ranges. For example, in order to obtain results that are valid for the general class of scale-free social networks with a high cluster coefficient, we run each computer simulation with a different instance of such a network, and aggregate the statistical data. Likewise, environmental dynamics can be typed among others by how sudden and how frequently the environmental conditions change, and results for specific types of environmental dynamics are based on repeated computer simulations, each with a different realization of the specific type of environmental dynamics. The number of simulations needed to obtain reliable statistical results are determined by standard methods of variance reduction.

The computer programs that simulate our economic models are simple—a few dozen lines of Matlab code that describe basic matrix operations and a simple for-loop. The various growth and production equations that we will use to describe the economic models can each be expressed by a single line of Matlab code. The imitation process, which depends on the local neighborhood structure of the social network, never requires more than a dozen lines of code, even in its most complicated form in Chapter 4. The number of code lines needed to collect and analyze data from the simulations is about ten times more than what is needed to actually run the simulation. The bulk of the coding effort however is not spent on running and analyzing the simulations, but on tuning their free parameters, which in this case has culminated in an independent software solution. All code, together with graphs and annotations, is available at <http://volker.nannen.com/phd-thesis>.

1.4 Thesis outline

The remainder of this thesis is organized as follows. Chapter 2 introduces Relevance Estimation and Value Calibration (REVAC), a numeric method that measures how much the performance of an evolutionary algorithm depends on the correct tuning of its operators and parameters, independent of the actual tuning method. The rationale behind the method is that if parameter values are taken from a probability distribution, the average performance of the resulting evolutionary algorithms can be evaluated against the amount of information—measured in Shannon entropy—that this distribution provides on its random values. To verify the reliability and computational efficiency of REVAC, we test it empirically on abstract objective functions, a simple and well studied genetic algorithm, and an agent-based simulation of our evolutionary economic model.

Chapter 3 uses REVAC to study how the performance of a typical evolutionary algorithm depends on the choice and tuning of its components. This is a novelty in evolutionary computing, where the cost of tuning is normally ignored. We tune a large array of common evolutionary algorithms to optimize four classes of objective function and compare the performance of different evolutionary algorithms before and after tuning, and how this improvement in performance depends on the tuning of a particular component. It turns out that the choice of operator for the selection mechanism typically has the greatest impact on performance, while the tuning of its parameters is of little consequence. Mutation on the other hand depends primarily on tuning, regardless of the operator.

After this preliminary work, Chapter 4 completes our work on the first objective and uses REVAC to develop a simple and robust model of selective imitation in a social network. According to the hypothesis that, by adding extra detail to the imitation mechanism, its adaptive power will either stay the same or increase, two imitation mechanisms are designed. One is rather simple, with free parameters for the selection, recombination, and mutation of strategies, as well as one parameter for the connectivity of the social network. The second model extends the first by using two distinct sets of free parameters for selection, recombination and mutation, one set to define the imitation behavior of rich agents, and one set to define the imitation behavior of poor agents. Both mechanisms are evaluated on an array of different economic environments with non-linear dynamics. REVAC disproves the above hypothesis by showing that for equal amounts of tuning the simpler mechanism consistently outperforms the extended one. As in the previous chapter, the correct tuning of the mutation operator, which maintains diversity in the pool of strategies, emerges as having the biggest impact on the simulation results.

Chapter 5 turns to our second objective. We design a minimal evolutionary mechanism with only one free parameter for the amount of diversity in the pool of strategies and use it to study how the evolutionary system reacts to different environmental dynamics. The analysis of a Cobb-Douglas type economy shows that from an evolutionary perspective only those environmental dynamics matter that affect the production coefficients. We define a number of basic environmental dynamics by varying these coefficients and formulate policy advice for policy makers with different types of risk preference regarding the socially optimal level of diversity.

Our third objective, the design and evaluation of policies that explicitly take the evolution of strategies into account, is addressed in Chapter 6. We use the same evolutionary mechanism as in Chapter 5 to build a simple model of global warming where the goal of the policy maker is to replace a resource with a negative impact on social welfare—fossil energy—by a neutral yet potentially less cost efficient alternative, namely renewable energy. We proceed to formulate two evolutionary policies—*prizes* and *advertisement*—that selectively increase the probability of an agent with a desirable strategies to be imitated, one by increasing the welfare of such an agent, the other by increasing its visibility in the social network. Numerical simulations are used to evaluate their effectiveness over a wide range of values for the additional cost of renewable energy, compared to a standard emission tax.

Chapter 7 concludes. For convenience, a list of all symbolic variables can be found in the appendix.

RELEVANCE ESTIMATION AND VALUE CALIBRATION OF EVOLUTIONARY ALGORITHM PARAMETERS

Abstract

Evolutionary algorithms (EAs) form a rich class of stochastic search methods that use the Darwinian principles of variation and selection to incrementally improve a set of candidate solutions (Eiben and Smith, 2003; Jong, 2006). Both principles can be implemented from a wide variety of components and operators, many with parameters that need to be tuned if the EA is to perform as intended. Tuning however requires effort, both in terms of time and computing facilities.

When resources are limited we are interested to know how much tuning an EA requires to reach an intended performance, and which parameters are most relevant in the sense that tuning them has the biggest impact on EA performance. Likewise, when designing an EA to simulate a real evolutionary process we would like to minimize the dependency of our simulation results on specific parameter values. In this case the amount of tuning required until the simulation behaves as intended indicates how plausible and realistic the simulation really is.

To measure the amount of tuning that is needed in order to reach a given performance, we introduce the REVAC method for Relevance Estimation and Value Calibration. While tuning the EA parameters in an automated and systematic manner, the method provides an information-theoretic measure on how much the tuning of each parameter contributes to overall EA performance. We evaluate its reliability and efficiency empirically on a number of test cases that reflect the typical properties of EA parameter spaces, as well as on evolutionary agent-based simulations. Finally we compare it to another tuning method, meta-GA.

Parts of this chapter have been published in Nannen and Eiben (2007b,a); de Landgraaf et al. (2007).

2.1 Background

One of the big challenges in evolutionary computing is the design and control of evolutionary algorithm (EA) parameters (Eiben et al., 1999). Without exaggeration, one could state that one of the canonical design problems is how to choose the operators for an evolutionary algorithm to ensure good performance. For instance, the question whether crossover is a relevant operator is still open, or rather, the answer depends on the application at hand (Eiben and Smith, 2003). A related issue is the relevance of free EA parameters. Depending on the EA and the problem it is applied to, tournament size can be a highly relevant parameter whose value must be chosen well for good performance, while mutation rate could be less relevant in the sense that its values do not affect EA performance too much. When designing an evolutionary algorithm to model a real evolutionary system, for example in evolutionary economics, one often has to deal with non-standard evolutionary mechanisms. These can include domain specific features of which it is altogether unknown whether the system behavior depends on their correct parameterization.

While the tuning of relevant evolutionary algorithm (EA) parameters is essential to good EA performance, current practice in EA tuning is based on ill-justified conventions and ad hoc methods. In particular, studies on confidence intervals for good parameter values and sensitivity analyzes for parameter robustness are almost non-existent. Part of the problem lies in the fact that most EAs are non-deterministic and path-dependent, in the sense that small changes to the initial conditions can lead to highly divergent results. This makes it difficult to obtain a reliable estimate of EA performance on a given problem. The standard statistical method to reduce variance and improve measurement reliability is measurement replication. With measurement replication, a set of parameter values is chosen, the EA is executed several times with these values on the same problem, and an aggregate performance measure is taken. A classical example of this approach is Analysis of Variance (ANOVA), which provides a clear set of rules how to optimally combine a number of carefully chosen parameter values, how to calculate the number of replications needed to decide whether one combination of values has a significantly better performance than another, and how to infer parameter interaction. An exhaustive overview of how to use ANOVA to tune an EA is given by Czarn et al. (2004).

This approach has a number of disadvantages, particularly when it is applied to an EA with several sensitive parameters. First, the choice of parameter values for the analysis is far from trivial and experiments in this vain often allow for no other conclusion than that a given choice was wrong. Second, the variance of an EA can easily be so high and its distribution so bad-behaved that the number of replications needed to produce significant results is not feasible. Third, there is disagreement in the statistical community on how to treat non-numerical results, for example when an EA does not find an acceptable solution within given computational constraints. Fourth, replications divert computational resources that could otherwise be used to obtain a better cover of the parameter space. This is a serious drawback, since it is virtually impossible to infer from a small number of measurements in a multi-dimensional search space, reliable as they might be, important measures of robustness like sensitivity to small changes and the range of values for which a certain EA performance can be achieved.

Here we propose to use an Estimation of Distribution Algorithm (EDA) to control

the parameters of an evolutionary algorithm: REVAC, which stands for Relevance Estimation and Value Calibration. REVAC is designed to a) tune or calibrate the parameters of an EA in a robust way and b) quantify the minimum amount of information that is needed to tune each parameter. Like a meta-GA (Grefenstette, 1986), it is an evolutionary method, a meta-EDA, that explores the parameter space of an evolutionary algorithm dynamically. Unlike meta-GA, it tunes an evolutionary algorithm on the basis of probability density functions over parameter values, rather than individual parameter values. Starting from a wide distribution over all possible parameter values, REVAC iteratively evaluates and improves the distribution such that it increases the probability of those parameter values that result in good EA performance. To avoid the pitfalls of bad-behaved distributions and non-numerical results, REVAC only uses rank based statistics to decide where to zoom in. Also, instead of investing valuable computational resources in measurement replications, REVAC uses them to get a better cover of the parameter space.

The estimated distributions over each parameter can be used to estimate the relevance of that parameter in an intuitive way. Broadly speaking, a distribution with a narrow peak indicates a highly relevant parameter whose values largely influence EA performance, while a broad plateau belongs to a less relevant parameter whose values do not matter too much. In terms of information theory, the Shannon entropy of a distribution expresses the average amount of information that is needed to specify a value that was drawn from the distribution Shannon (1948). The sharper the peaks of a continuous probability density function, the lower its Shannon entropy, and the less information is needed to specify the values drawn from the distribution. If a distribution over parameter values has maximum entropy for a given level of expected EA performance, then this maximum entropy can be used to calculate the minimum amount of information that is needed to achieve that performance. REVAC forces the distribution it finds to approximate the maximum entropy distribution for a given level of performance by continuously smoothing them between updates, so that their Shannon entropy can be used to estimate the minimum amount of information needed to reach this level of performance. In these terms the objectives of REVAC can be formulated as follows:

- The Shannon entropy of the distribution is as high as possible for a given level of performance,
- The expected performance of the EA in question is as high as possible for a given level of Shannon entropy.

Related work includes meta-GA as an early attempt to automate the tuning of genetic algorithms (Grefenstette, 1986), and Eiben et al. (1999) who established parameter control in EAs as one of the big challenges in evolutionary computing. Czarn et al. (2004) discuss current problems in EA design and use polynomial models of a performance curve to estimate confidence interval for parameter values. François and Lavergne (2001) estimate performance curves for EAs across multiple test cases to measure generalizability. Bartz-Beielstein et al. (2005) uses a Gaussian correlation function to dynamically build a polynomial regression model of the response curve.

The groundwork of statistical experimental design was laid by R. A. Fisher in the 1920s and 1930s. The use of sequential sampling to search for optimal parameter val-

ues were introduced by Box and Wilson (1951). A paradigm shift that emphasizes the robustness of a solution is due to Taguchi and Wu (1980). Estimation of Distribution Algorithms, in particular those based on univariate marginal distributions, to which the present type belongs, were pioneered by Mühlenbein (1997). The relationship between Shannon entropy and EDAs is discussed extensively in Mühlenbein and Höns (2005).

A detailed description of REVAC is given in Section 2.2. In Section 2.3 we use abstract objective functions and a simple genetic algorithm (GA) to verify that REVAC can indeed estimate how much tuning the parameters of an EA need. Section 2.4 uses the same simple GA to evaluate whether REVAC uses the available evaluations of candidate solutions efficiently, without the need for measurement replication. The same section also tests REVAC on an agent-based simulation from evolutionary economy. Section 2.5 compares REVAC to other tuning methods, namely hand tuning and meta-GA. A summary and conclusions can be found in Section 2.6.

2.2 The algorithm

2.2.1 Approaching the maximum entropy distribution

Formally, the tuning of EA parameters to a specific application is itself an optimization problem where the value of an objective function $r = f(\vec{x})$ is maximized. The domain of this objective function are the possible combinations of parameter values \vec{x} for the EA. Its value r , which is also called the response, is the expected performance of the EA on the application problem when executed with these parameter values.

Since the domain of many EA parameters is continuous, the choice of suitable EDAs to tune them is limited. The present algorithm is a steady state variant of the Univariate Marginal Distribution Algorithm (Mühlenbein, 1997). For efficiency, only a single parameter vector is evaluated between every update of the distributions. Given an EA with k parameters, REVAC defines a joint distribution $\mathcal{D}(\vec{x})$ over the space of possible parameter vectors $\vec{x} = \{x_1, \dots, x_k\}$. This joint distribution is composed from a set of independent marginal density functions $\mathcal{D}(\vec{x}) = \langle \mathcal{D}(x_1) \dots \mathcal{D}(x_k) \rangle$ over each parameter. Their Shannon entropy can be used to estimate how much information is needed per parameter to achieve the expected performance of the joint distribution. Let a probability density function \mathcal{D} be defined over a continuous interval $[a, b]$. Its differential Shannon entropy h can be calculated as

$$h(\mathcal{D}_{[a,b]}) = - \int_a^b \mathcal{D}(x) \log_2 \mathcal{D}(x) dx. \quad (2.1)$$

In order to compare the entropy of distributions that are defined over different parameter intervals in a meaningful way, we normalize all parameter intervals to the unit interval $[0, 1]$ before calculating the Shannon entropy. In this way the uniform distribution has a Shannon entropy of zero, and any other distribution has a negative Shannon entropy $h(\mathcal{D}_{[0,1]}) < 0$.

Starting with the uniform distribution, REVAC iteratively refines the distribution by drawing random vectors of parameter values from it, measuring the performance of the EA with these parameter values, and increasing the probability of those regions of the

Table 2.1: Two views on a table of parameter vectors

	$\mathcal{D}(x_1)$	\cdots	$\mathcal{D}(x_i)$	\cdots	$\mathcal{D}(x_k)$
\vec{x}^1	$\{x_1^1$	\cdots	x_i^1	\cdots	$x_k^1\}$
\vdots	\vdots	\ddots	\vdots	\cdots	\vdots
\vec{x}^j	$\{x_1^j$	\cdots	x_i^j	\cdots	$x_k^j\}$
\vdots	\vdots		\vdots	\ddots	\vdots
\vec{x}^n	$\{x_1^n$	\cdots	x_i^n	\cdots	$x_k^n\}$

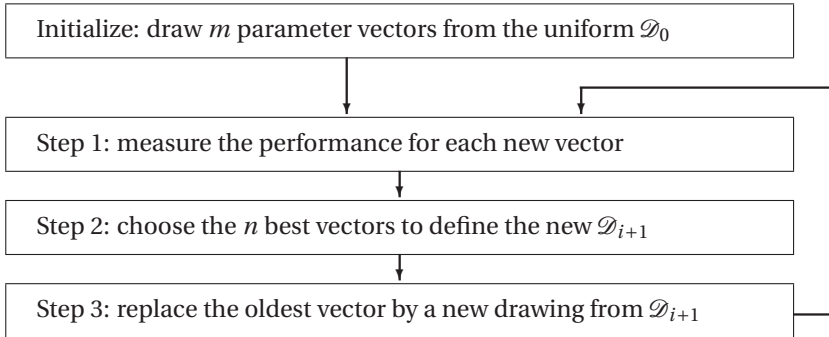
parameter space where a higher EA performance is measured. In this way, new distributions are built on estimates of the response surface that were sampled with previous distributions, each iteration increasing the expected performance of parameter vectors that are drawn from the distribution. To reduce the variance of stochastic measurements and to prevent premature convergence, REVAC continuously smoothes the distribution. It is the unique combination of these two operators, increasing the probability of regions with high performance and smoothing out the resulting probability function, that allows REVAC to approach the maximum entropy distribution for a given level of EA performance.

2.2.2 Algorithm implementation

At each step in the tuning process, REVAC maintains a pool of m vectors of parameter values. From this pool the $n < m$ vectors with the highest measured performance are selected to define the current distribution and to create a single new parameter vector. The new parameter vector always replaces the oldest one in the pool. For a good understanding of how this is done it is helpful to distinguish two views on the n selected parameter vectors as shown in Table 2.1. Taking a *horizontal* view on the table, a row is a vector of parameter values and we can see the table as n of such vectors. Taking a *vertical* view on the table, the i^{th} column shows n values from the domain of parameter i . Each column of Table 2.1 defines a marginal density function and the whole table defines the joint density function.

As can be seen in the diagram of Figure 2.1, REVAC initializes the table of parameter vectors by drawing k vectors from the uniform distribution over the space of possible parameter values. The update process that creates a new table of parameter vectors consists of three basic steps: *evaluating parameter vectors*: Given a vector of parameter values, we can evaluate it by executing the EA with these parameter values and measuring its performance; *updating the probabilities*: Given a set of evaluated parameter vectors, we can calculate the probability that some regions of the parameter space have a higher expected performance than others; *generating parameter vectors*: Given a probability density function over the parameter space, we can draw new parameter vectors proportional to those probabilities.

Figure 2.1: Diagram of the update process



Step one is straightforward. As for step two and three, they can be described from both the horizontal and the vertical perspective of Table 2.1. Looking from the *horizontal* perspective, REVAC can be described as a population based evolutionary algorithm with operators for selection, recombination and mutation. This description of REVAC must not be confused with the EA we are tuning. The population consists of m parameter vectors. It is updated by selecting $n < m$ parent vectors from the old population, which are then recombined and mutated to obtain exactly one child vector every generation. The child vector always replaces the oldest vector in the population.

REVAC uses a deterministic choice for parent selection as well as for survivor selection. The n vectors of the population that have the highest measured performance are selected to become the parents of the new child vector. Recombination is performed by a multi-parent crossover operator, uniform scanning, that creates one child from n parents, cf. Eiben and Smith (2003). The mutation operator—applied to the offspring created by recombination—is rather complicated. It works independently on each parameter i in two steps. First, a mutation interval $[x_a^i, x_b^i]$ is calculated, then a random value is chosen uniformly from this interval. To define the mutation interval for mutating a given x_i^j all other values x_i^1, \dots, x_i^n for this parameter in the selected parents are also taken into account. After sorting them in increasing order, the begin point of the mutation interval or window can be specified as the w -th lower neighbor of x_i^j , while the end point of the interval is the w -th upper neighbor of x_i^j . The new value is drawn from this interval with a uniform distribution. As there are no neighbors beyond the upper and lower limits of the domain, we extend it by mirroring the parent values as well as the mutated values at the limits, similar to what is done in Fourier transformations.

From the *vertical* perspective we consider step two and three as constructing k marginal probability density functions from the columns of Table 2.1 and then drawing a new parameter vector from these distributions. To define a marginal density function $\mathcal{D}(x_i)$, the n values of column i are sorted and arranged such that together with the limits 0 and 1 (the domain of each parameter is scaled to the unit interval) they form $n + 1$ non-overlapping intervals that cover the entire domain. The density over any such

interval $[x_i^a, x_i^b]$ can be defined as

$$\mathcal{D}(x_i) = \frac{1}{(m+1)(x_i^b - x_i^a)}, \quad (2.2)$$

which satisfies $\int_0^1 \mathcal{D}(x_i) dx_i = 1$. This definition of a density function can be extended to allow intervals to overlap, for example by defining intervals between values that are separated by one or two other values. To overcome the problem of missing neighbors at the limits we again mirror all defining parameter values as well as the chosen values at the limits. The further the values that define an interval are separated, the higher the Shannon entropy of the resulting distribution.

In this context, the rationale behind the complicated mutation operator of the horizontal view is that it heavily smoothes the density functions of equation 2.2. Like all evolutionary algorithms, an EDA is susceptible for converging on a local maximum. By continuously smoothing the probability density functions we force them to converge on a maximum of the response surface that lies on a broad hill, yielding robust solutions with broad confidence intervals. But smoothing does more: it allows REVAC to operate under very noisy conditions, it allows it to readjust and relax marginal distributions when parameters are interacting and the response surface has curved ridges, and it maximizes the entropy of the constructed distribution. Smoothing is achieved by taking not the nearest neighbor but the w -th neighbors of x_i^j when defining the mutation interval. Choosing a good value for w is an important aspect when using REVAC. A large w value can slow down convergence to the point of stagnation. A small w value can produce unreliable results. Based on our experience so far, we prefer $w \approx n/10$.

2.2.3 Interpreting the measurements

REVAC, like any EDA, is a random process. The final result is different with every run or tuning session of REVAC. Independently of whether REVAC uses measurement replication, REVAC results can be made more reliable by tuning an EA more than once, and by either choosing the tuned parameter values that resulted in the highest EA performance, or by averaging over the tuned parameter values of several runs of REVAC, as will be explained below. This is indeed a replication of measurements at a higher level. But unlike ordinary measurement replication, which can be too expensive to extract any useful information, REVAC can always provide a first approximation, which can then be refined by repeating the tuning process.

Because REVAC produces a sequence of distributions with slowly decreasing Shannon entropy we use the Shannon entropy of these distributions to estimate the minimum amount of information needed to reach a target performance level. This can be used in several ways. First, it can be used to choose between different sets of EA operators. A set of operators that needs less information to reach a given level of EA performance is easier to tune, more fault tolerant in the implementation, and robuster against changes to the problem definition. Second, it can be used to identify the critical components of an EA. A highly sensitive parameter typically has a sharp peak in the distribution and a low Shannon entropy. When an EA needs to be adjusted to a new problem, sensitive parameters need the most attention, and with this knowledge the practitioner can

concentrate on the critical components straight away. Third, it can be used to suggest values and confidence intervals for the best parameter values. Given a distribution that peaks out in a region of high probability (except for the early stage of the algorithms the marginal distributions have only one peak), we take the 50th percentile (the median) to be the best tuned parameter values, and use the 25th and the 75th percentile of the distribution as confidence interval. That is, every value from this range leads to a high expected performance, under the condition that the other parameters are also chosen from their respective confidence interval. This confidence interval is also useful when we average over several runs of REVAC. As we only want to average over those runs that converged on the same optimum in the parameter space, we take the average of only those REVAC distributions where all medians lie within the 25th and 75th percentiles of the respective distributions of the REVAC run that achieved the best EA performance. REVAC runs that converged on values beyond these intervals are discarded.

Throughout the rest of this doctoral thesis REVAC will use a population of $m = 100$ parameter vectors, from which the best $n = 50$ are selected for being a parent. We smooth by extending the mutation interval over the $w = 5$ upper and lower neighbors. In each run or tuning session REVAC is allowed to evaluate 1,000 parameter vectors.

2.3 Assessing the reliability of REVAC estimates

A real in vivo assessment of REVAC requires that we tune an EA on a set of application problems where the objective function is known, and use this to evaluate the results obtained by REVAC. It is however not feasible to accurately model the response surface of an EA on any non-trivial application problem. According to Czarn et al. (2004), even when working with rather simple application problems and a simple genetic algorithm with only two free parameters, it is difficult to fit anything more sophisticated than a cubic curve to the measured performance. This leaves us with two alternatives: we evaluate REVAC on abstract objective functions with a predefined response surface that is representative for EA tuning problems. This has the added advantage that by abstracting both the application and the algorithm layer the run time of the tuning process is reduced enormously, and that the assessment of REVAC can be based on a large number of measurements. The second alternative is to use an EA and an application problem that have been studied in the literature and to evaluate the REVAC relevance estimates against existing results on the relevance of the EA parameters.

2.3.1 Assessing REVAC reliability on abstract objective functions

In general, we distinguish 3 layers in designing an EA, as shown in Table 1.1 on page 19. For the present assessment, these layers are implemented as follows:

Experimental setup of Section 2.3.1	
design tool	<i>REVAC</i>
evolutionary algorithm	<i>abstract</i>
application problem	<i>abstract</i>

To define abstract response surfaces that resemble the response surface of a typical EA we identify five essential properties: 1) *Low dimensionality*: typically not more than ten parameters need to be tuned. 2) *Non-linearity (epistasis, cross-coupling)*: parameters interact in non-linear ways, which implies that the response curves are non-separable, and that values for one parameter depend on the values of other parameters. 3) *Smoothness*: small changes in parameter values lead to small changes in EA performance. 4) *Low multi-modality (moderate ruggedness)*: the objective function has only one or few significant local optima, i.e., few regions of the parameter space have high EA performance. 5) *Noise*: depending on the application, the performance of an EA can be highly variable and can follow a distribution that is significantly different from the normal distribution.

We present three experiments: two on the accuracy of the relevance estimates (without noise) and one on the resistance of the estimates to noise. In all experiments the abstract objective function simulates the performance of an EA with $k = 10$ parameters, each of which can take values from the range $[0, 1]$.

Experiment 1: hierarchical dependencies. In our first experiment we hardcode pre-defined dependencies between parameters, such that the optimal value of parameter i depends on the current value of $i - 1$ and the utility of tuning parameter i depends on how well parameter $i - 1$ is tuned. The abstract objective function $r = f(\vec{x})$, is defined as the sum

$$r = \sum_{i=1}^{10} r_i \quad (2.3)$$

over ten partial response values r_1, \dots, r_{10} , one for each of the ten parameters. These are calculated as follows. Before a tuning session is started, a single target value t is chosen at random from the range $[0, 1]$ and kept constant throughout the tuning session. When evaluating a parameter vector $\vec{x} = \{x_1, \dots, x_k\}$, the partial response r_1 of the first parameter value x_1 is one minus the distance to the target value,

$$r_1 = 1 - |x_1 - t|. \quad (2.4)$$

The partial response r_i of each consecutive parameter value depends on how close the value x_i is to that of parameter $i - 1$,

$$r_i = r_{i-1}(1 - |x_i - x_{i-1}|). \quad (2.5)$$

We have $x_1 - t \leq 1$ and $x_i - x_{i-1} \leq 1$ for any i . Because r_{i-1} is a factor in the calculation of r_i , the inequality $r_i > r_{i+1}$ always holds, with the effect that the first parameter needs more tuning than the second one, and so forth.

Figures 2.2 and 2.3 shows typical results when using REVAC to tune the 10 parameters to this abstract objective function. Results are from a single run of REVAC. The bar diagram in Figure 2.2 shows the final Shannon entropy per parameter after evaluating 1,000 parameter vectors. The precoded linear relationship in the need for tuning is well approximated. In particular, the order of the first five parameters, which need the most tuning, is correctly identified. The upper left graph of figure 2.3 shows how the measured Shannon entropy of the REVAC distributions over parameter 1, 4, and 7 changes during the tuning session. The other three graphs of the figure show how the 25th, 50th




Figure 2.2: Final Shannon entropy per parameter

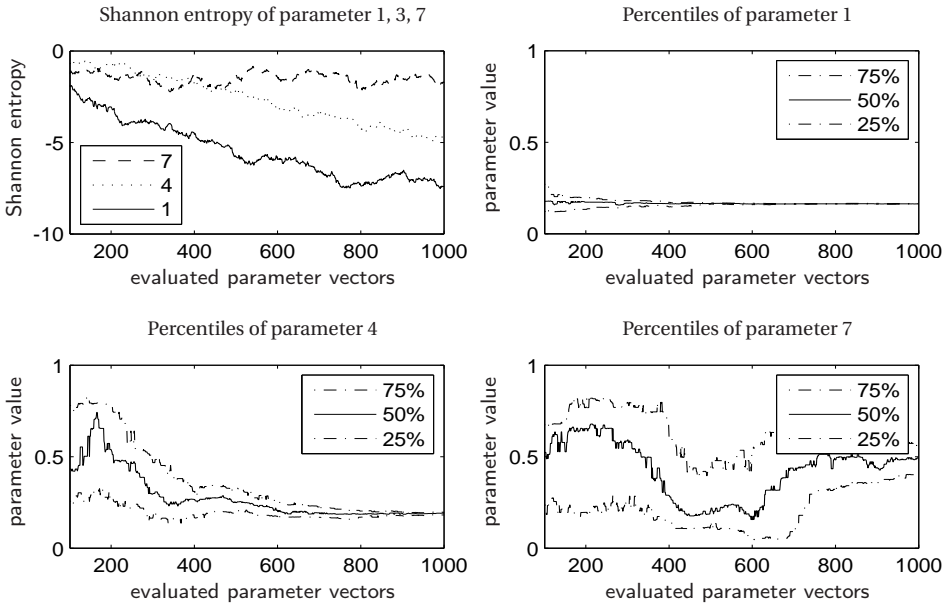


Figure 2.3: Shannon entropy and the percentiles of the REVAC distributions of parameter 1, 3, and 7 during tuning.

and 75th percentile of the same distributions change during tuning. Note that the distance between the 25th and 75th percentile of the distributions behaves similar to the entropy of the same distribution.

Experiment 2: predefined relevance distribution. In this experiment we are interested to see whether REVAC can reveal arbitrary distributions over parameter relevance. To this end we create an abstract objective function with one peak which is placed at random in the 10-dimensional unit space. Again, total response is the sum of the partial response per parameter. This is calculated as one minus the distance of a parameter value x_i to the corresponding peak value t_i , weighted by a predefined vector of target

weights $w = \langle w_1, \dots, w_{10} \rangle$,

$$r = \sum_{i=1}^{10} w_i [1 - (x_i - t_i)]. \quad (2.6)$$

In this way, tuning a parameter with a target weight close to zero has no impact on overall response, while tuning a parameter with a large target weight has a significant impact on the response. We use three sets of target weights to specify three elementary distributions over parameter relevance. The first distribution has two outliers of exceptionally low relevance, while the remaining target weights are equal (normalize w , with $w_1 = 0$, $w_2 = 1$, and $w_3, \dots, w_{10} = 10$). The weights of the second distribution increase linearly over the parameters (normalize w , with $w_i = i$). The weights of the third distribution increase exponentially such that there are outliers with a high relevance (normalize w , with $w_i = i^{10}$). This last distribution represents what is known as *sparsity of effects* in the design of experiments and is the most typical situation when tuning a real EA.

Figure 2.4 shows the target weights (the black columns) together with what REVAC has estimated after 1,000 evaluated parameter vectors (the white columns). Results are from a single run of REVAC. To estimate how relevance is distributed over the parameters we normalize the Shannon entropy of the marginal REVAC distributions, which results in positive values that sum to one. As can be seen, REVAC approximately reproduces the hardcoded order of relevance, in particular with regard to the outliers, but has difficulties when the weights are too similar. When averaging over several REVAC runs (not shown here), the order of the estimated relevance per parameter converges to the order of the predefined target weights.

Experiment 3: measurement noise. In this experiment we study how the reliability of a relevance estimate changes with additive noise of increasing variance. For the abstract objective function we add a noise term η to the objective function 2.6 of experiment 2,

$$r = \sum_{i=1}^{10} w_i [1 - (x_i - t_i)] + \eta. \quad (2.7)$$

To simulate sparsity of effects, the weights w_1, \dots, w_{10} increase exponentially from parameter to parameter, cf. the bottom graph of Figure 2.4. Values for the noise term η are independent and identically distributed. They are drawn from a Pareto distribution

$$P(X > x) = cx^{-\gamma}, \quad x > \log_{\gamma} c \quad (2.8)$$

with exponent $\gamma = 2.3$. The value c controls the variance σ^2 of the noise. Such a distribution is also called a power law distribution and can be found in many physical and social systems. It has frequent outliers, its variance converges rather slowly, and it is generally incompatible with statistical methods that require a normal distribution.

To measure the error of the REVAC estimates we use the mean squared distance between target weights and the corresponding normalized Shannon entropy after evaluating 1,000 parameter vectors. With s_1, \dots, s_{10} the normalized Shannon entropy, the error can be calculated as

$$error = \frac{1}{10} \sum_{i=1}^{10} (s_i - w_i)^2. \quad (2.9)$$

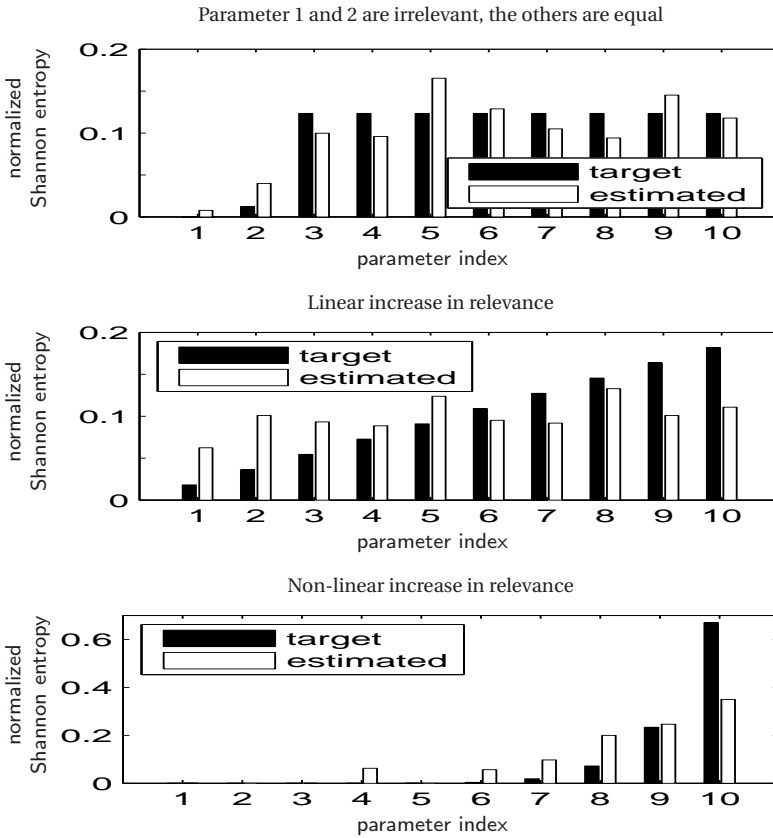


Figure 2.4: Comparing the normalized Shannon entropy per parameter as estimated by REVAC to the target weights of the abstract objective function.

Figure 2.5 plots the measured error for a single run of REVAC against the variance σ^2 of the noise. The mean squared error of the REVAC estimate increases roughly linearly with the variance. Note that the highest value for σ^2 is five, while the objective function itself only takes values from the range $[0, 1]$. The mean squared error hardly exceeds the value 0.1. To compare, the mean squared error between the target weights and a 10-dimensional normalized random vector is 0.29.

The variance of independent and identically distributed noise can be reduced by measurement replication. As seen in Figure 2.6, REVAC estimates can also be improved by taking the average of several REVAC runs that were obtained without measurement replication. Both graphs compare the estimated normalized Shannon entropy after 1,000 evaluated parameter vectors (white columns) to the target weights (black columns). The variance of the noise is $\sigma^2 = 5$. The upper graph is based on a single run of REVAC and the lower graph on 10 runs. The mean squared error of the relevance estimate is 0.022 in the upper graph and 0.011 in the lower graph. This means that under noisy conditions a single run of REVAC without measurement replication can give a quick first approximation that can be refined by further runs if resources permit.

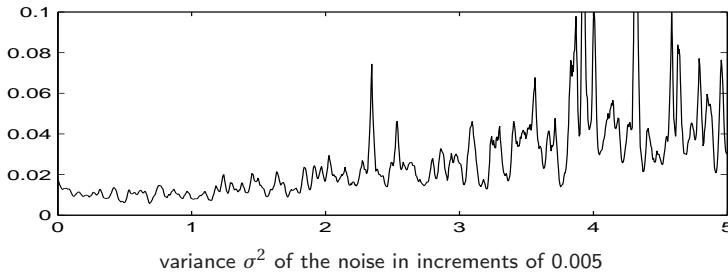


Figure 2.5: Impact of noise on the mean squared error of the normalized Shannon entropy after 1,000 evaluated parameter vectors. The graph is smoothed for readability.




Figure 2.6: Relevance estimates with noise of variance $\sigma^2 = 5$. The upper graph shows a typical estimate from a single run. The lower graph shows the average of 10 typical runs.

2.3.2 Assessing REVAC reliability on a simple genetic algorithm

Experimental setup of Section 2.3.2

design tool	<i>REVAC</i>
evolutionary algorithm	<i>simple genetic algorithm</i>
application problem	<i>standard numerical optimization problems</i>

Here we present the results of tuning an EA on application problems that have been previously studied in the literature, as discussed at the beginning of this Section 2.3. For both the EA and the objective function we rely on Czarn et al. (2004), who use rigorous statistical exploratory analysis to tune a simple genetic algorithm and who compare their results to those of Jong (1975), Schaffer et al. (1989), Grefenstette (1986), and

Table 2.2: REVAC results after 1,000 evaluations

Function & parameters	Optimum value	Confidence interval (25 th and 75 th pctl.)	Shannon entropy	Normalized Shannon entropy	
f_1	p_m	0.012	0.011 – 0.013	-8.6	0.82
	p_c	0.90	0.77 – 0.96	-1.9	0.18
f_2	p_m	0.0146	0.0143 – 0.0148	-9.4	0.82
	p_c	0.82	0.77 – 0.86	-2.1	0.18
f_3	p_m	0.0338	0.0334 – 0.0342	-9.0	0.72
	p_c	0.98	0.82 – 0.99	-3.5	0.28
f_6	p_m	0.0604	0.0635 – 0.0641	-6.9	0.86
	p_c	0.60	0.48 – 0.68	-1.1	0.14

Freisleben and Hartfelder (1993). Specifically, they study the effect of tuning the mutation parameter $p_m \in [0, 1]$ and the crossover parameter $p_c \in [0, 1]$ of a generational genetic algorithm (GA) with 22 bits per variable, Gray coding, probabilistic ranked-based selection, bit flip mutation, single point crossover, and a population of 50 chromosomes. The 4 objective functions for the application layer are standard benchmark problems from Jong (1975) and Schaffer et al. (1989): sphere (f_1), saddle (f_2), step (f_3), Schaffer's f_6 . Their definitions are given in equation 2.10—2.13,

$$f_1(x) = \sum_{i=1}^3 x_i^2, \quad -5.12 \leq x_i \leq 5.12, \quad (2.10)$$

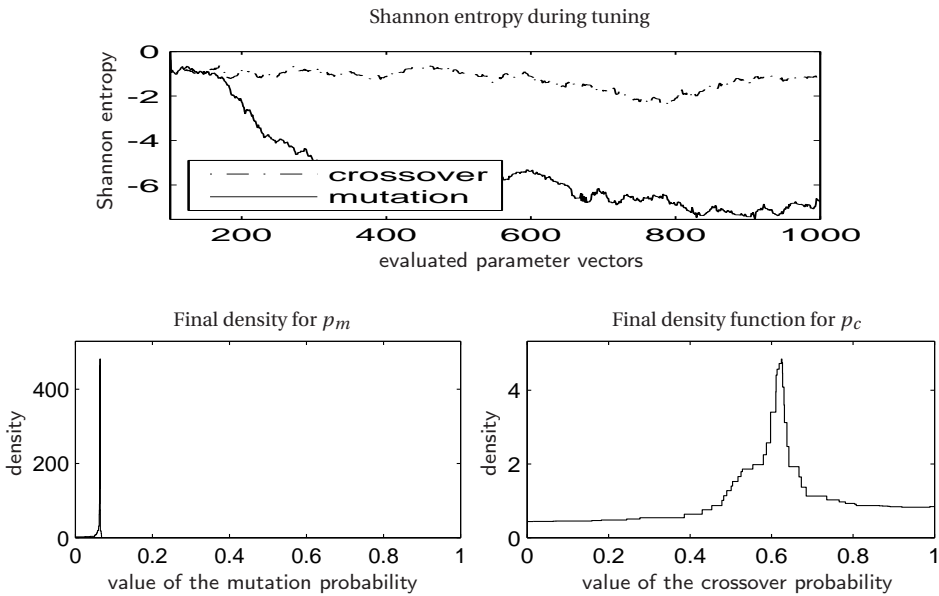
$$f_2(x) = 100(x_2 - x_1^2)^2 + (1 - x_1)^2, \quad -2.048 \leq x_i \leq 2.048, \quad (2.11)$$

$$f_3(x) = \sum_{i=1}^5 \lfloor x_i \rfloor, \quad -5.12 \leq x_i \leq 5.12, \quad (2.12)$$

$$f_6(x) = 0.5 + \frac{(\sin \sqrt{x_1^2 + x_2^2}) - 0.5}{(1 + 0.0001(x_1^2 + x_2^2))^2}, \quad -100 \leq x_i \leq 100. \quad (2.13)$$

Table 2.2 shows the results per objective function after evaluating 1,000 parameter vectors. Figure 2.7 shows the Shannon entropy during tuning and the final distributions after tuning for Shaffer's f_6 . The upper graph shows the Shannon entropy of p_m and p_c during tuning. Only the entropy of p_m decreases significantly, indicating its relevance. The two lower graphs show the final probability density function over the parameter values after evaluating 1,000 parameter vectors. Note the extremely sharp needle for p_m .

These results can be considered from two perspectives, compared with the “usual” GA settings, and with the work of Czarn *et al.* As Table 2.2 shows, the values found by REVAC are consistent with the conventions in evolutionary computing: p_m between 0.01 and 0.1, and p_c between 0.6 and 1.0. On the other hand, a direct comparison with Czarn *et al.* (2004) is difficult because of the different types of outcomes. As for the method, Czarn *et al.* use screening experiments to narrow down the space of feasible parameter settings, partition this space into equally spaced discrete levels, repeatedly measure

Figure 2.7: Tuning Schaffer's f_6

the performance for each level and use ANOVA to calculate the significance of mutation and crossover rates. Then they proceed to approximate the response curve for both parameters by fitting low order polynomials to the performance measures, suggesting to calculate confidence intervals from these approximations. As a main result, Czarn et al. find that the marginal response curves for crossover and mutation are linear and quadratic and that mutation is more significant than crossover. By contrast, REVAC uses no screening and does not partition the parameter space into discrete levels. It studies the complete and continuous parameter space. It can narrow the solution space to any arbitrarily small subspace and can directly read off confidence intervals for any given level of performance. Our global outcomes, however, are in line with those in Czarn et al. (2004): p_m is much more peaked and relevant than p_c .

2.4 Assessing the algorithmic efficiency of REVAC

To reduce the variance in the measured performance of an EA, statistical methods commonly rely on measurement replication. Until some confidence is achieved as to which vectors lead to a higher EA performance, these methods invest valuable computational resources in evaluating the same vectors of parameter values over and over. By contrast, REVAC reduces the variance in the measured performance implicitly through extensive sampling and smoothing. An increase of the pool size m and the number of selected vectors n means that estimated densities are based on a larger and more reliable number of evaluated parameter vectors. An increase of the smoothing parameter w means that the densities are averaged over a larger number of adjacent parameter intervals.

By using sufficiently large numbers for m , n , and w , REVAC aims to both correct for the variance in the measured performance as well as to get a better cover of the parameter space. Here we address the question whether this is indeed achieved, or, conversely, whether measurement replication will improve the quality of REVAC estimates or reduce the computational cost of obtaining them. This question is particularly pressing since REVAC is intended to tune EAs under conditions where established methods like ANOVA are inefficient, and where a maximum of information has to be extracted from every available measurement. We formulate two research questions: First, how does the replication of measurements affect the quality of REVAC estimates? And second, how does the replication of measurements affect the computational efficiency of the REVAC search process?

In order to study the merits of measurement replication for REVAC, a new parameter r for the number of measurement replications is added to step 1 of the REVAC algorithm, cf. Figure 2.1. Upon drawing a new parameter vector \vec{x} from the joint distribution $\mathcal{D}(x)$, it is evaluated r times, and the average result is recorded. We use our standard REVAC implementation with $m = 100$, $n = 50$, $w = 5$. Two sets of experiments are reported: tuning a simple genetic algorithm (GA) on standard numerical optimization problems and the tuning evolutionary mechanism of a complex simulation as part of our research on evolutionary agent-based economics.

2.4.1 Assessing algorithmic efficiency on a simple genetic algorithm

Like in section 2.3.2 we rely on Czarn et al. (2004) for the GA and the objective functions sphere (f_1), saddle (f_2), step (f_3), and Schaffer's f_6 of equations 2.10—2.13. In addition to the two parameters tuned there, mutation $p_m \in [0, 1]$ and crossover $p_c \in [0, 1]$, we also tune the population size of $n \in [10, 200]$ chromosomes, a total of 3 parameters. Figure 2.8 demonstrates how the Shannon entropy and the percentiles of the marginal distributions change during a typical tuning session without measurement replication. Here the step function (f_3) is used. The upper left graph shows the Shannon entropy of all three GA parameters. The other three graphs show the median and the 25th and 75th percentiles per parameter.

Experimental setup of Section 2.4.1

design tool	<i>REVAC</i>
evolutionary algorithm	<i>simple genetic algorithm</i>
application problem	<i>standard numerical optimization problems</i>

The performance measure of the GA that we wish to optimize is the computational cost of maximizing the objective function to which it is applied. This computational cost is the number of fitness evaluations, which in this case is calculated as the population size of the GA times the number of generations that are needed to maximize the objective function. The performance of the GA is maximized when the cost is minimized. When a GA needs 100 generations of 100 individuals or 200 generations of 50 individuals, we will say that it has a cost of 10,000 fitness evaluations. An objective function is considered maximized as soon as one individual of the population encodes a value that is within certain bounds of the best feasible solution. These bounds are chosen such

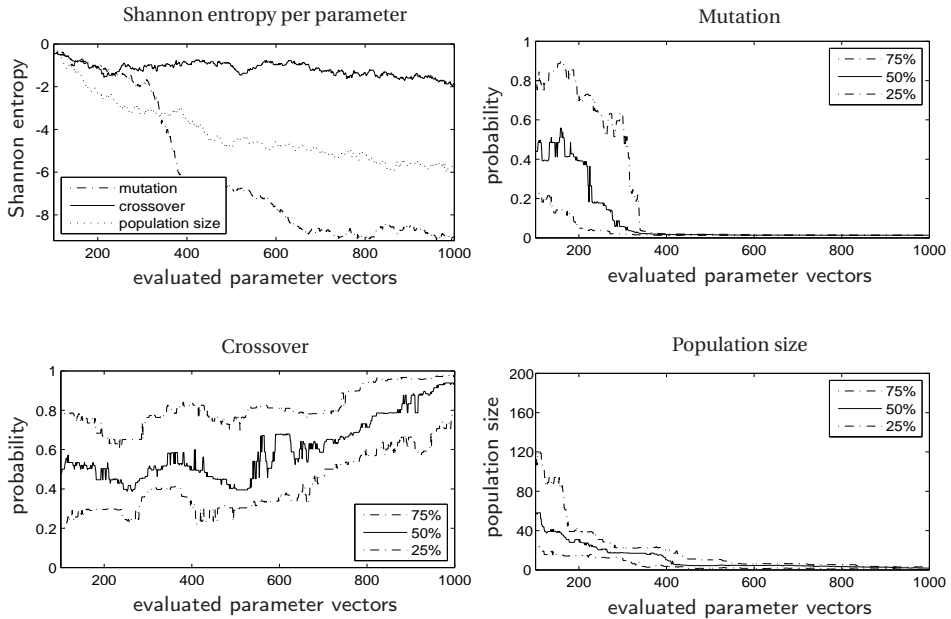


Figure 2.8: Shannon entropy and percentiles of the marginal distributions over 3 GA parameters during a typical tuning session.

that a well tuned algorithm can solve each objective function with a cost of between 5,000 and 10,000 fitness evaluations. If the algorithm does not solve the objective function with a cost of less than 25,000 fitness evaluations (e.g., in 1,000 generations if the population size is 25), execution is aborted and a cost of 25,000 fitness evaluations is recorded. We assess 5 different levels of measurement replication, $r \in \{1, 2, 3, 5, 10\}$. For each level of r and each objective function we run REVAC ten times, each run constituting an independent tuning session of the GA, and we report the average of the ten estimates.

In order to assess the quality of REVAC relevance estimates for each level r we need a reliable target value to compare to. Section 2.3.1 has shown that the average of repeated REVAC estimates without replication converges on the predefined distribution of parameter relevance. Section 2.3.1 has shown that REVAC can tune the simple GA to the four numerical optimization problems and give reasonable relevance estimates. We assume that the average relevance estimate of multiple REVAC runs converges on the correct values and we use these convergent values as the target values. For this reason our target value for each parameter on each objective function is the average Shannon entropy at the end of all REVAC runs with all levels of replications, 50 runs for each objective function. The exact values can be seen in Table 2.3. To assess the quality of the estimates obtained with a given number of measurement replications we use the error or mean squared distance to this target value, cf. equation 2.9. We consider the following four quantities:

- the number of different parameter vectors that REVAC needs to evaluate in order to reach an error < 0.1 with regard to the target values,
- the total number of measurements that are needed to reach an error < 0.1 (i.e., the number of parameter vectors times the number of measurement replications),
- the error after evaluating 1,000 vectors, regardless of the total number of measurements, and
- the error after a total number of 1,000 measurements.

Table 2.4 shows the recorded values for each level of measurement replication. Results are averaged over all objective function. The table clearly shows that a higher number of replications comes with a heavy computational penalty, without leading to a significant improvement in the quality of the relevance estimates. To be precise, there is no observable trend in the error after 1,000 evaluated parameter vectors for $r > 1$, while at this point it is still significantly higher for $r = 1$. As Figure 2.9 indicates, this is can be due to the fact that with fewer overall measurements, after 1,000 evaluated parameter vectors REVAC with $r = 1$ is still converging on the final value. The graph plots the error against the number of evaluated parameter vectors for $r \in \{1, 2, 10\}$, and indeed, only after evaluating about 800 parameter vectors does REVAC with $r = 2$ and $r = 10$ —which makes 1,600 and 10,000 measurements—reach an error that is visibly lower than the final error of REVAC $r = 1$. We conclude that while the evidence regarding $r = 1$ and $r = 2$ is inconclusive, there is no evidence that a number of replication of measurements greater than 2 leads to a significant improvement of the estimate.

Table 2.3: Average Shannon entropy of the 3 free GA parameters

	Sphere (f_1)	Saddle (f_2)	Step (f_3)	Schaffer's f_6
Mutation	-11.1	-11.3	-10.9	-9.6
Crossover	-1.7	-3.5	-2.2	-0.9
Population size	-5.9	-4.5	-6.2	-1.0

Table 2.4: Quality of the relevance estimate for different numbers of measurement replication. Results are averaged over all objective function.

Number of measurement replications	Number of vectors until error < 0.1	Number of measurements until error < 0.1	Error at 1,000 vectors	Error at 1,000 measurements
1	404	404	0.08	0.09
2	413	826	0.04	0.07
3	741	2,223	0.05	0.23
5	844	4,220	0.04	0.35
10	236	2,360	0.06	0.37

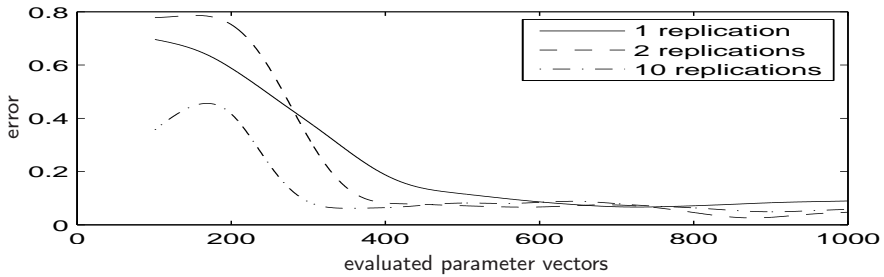


Figure 2.9: Error of relevance estimate for three levels of measurement replication, averaged over all 4 objective function. The lines are smoothed for readability. Note that the x -axis shows the number of evaluated parameter vectors, not the computational cost, which is measured in the total number of measurements.

Table 2.5: Best performance for each objective function, measured as number of fitness evaluations to solution. A lower value is better.

	Sphere	Saddle	Step	Schaffer's f_6
Optimum performance	3,786	2,770	2,107	3,260

To compare the quality of parameter values that REVAC has tuned we need an indication of how well a simple GA can perform on each objective function if properly tuned, i.e., the minimum amount of fitness evaluations that is needed to maximize the objective function. For this purpose we choose from among the 50 REVAC runs per objective function the parameter vector that achieved the best GA performance, and record the average number of fitness evaluations that the GA with this parameter vector needs to maximize the respective objective function. These best performances are shown in Table 2.5. We again consider four quantities:

- the number of parameter vectors that REVAC needs to evaluate in order to bring the computational cost of the GA down to no more than twice the computational cost of the best performance (i.e., 10,000 fitness evaluations if the best GA performance is 5,000 fitness evaluations),
- the number of measurements REVAC needs to perform in order to achieve the same as above,
- the average GA performance after REVAC has evaluating 1,000 parameter vectors, regardless of the number of measurements involved, and
- the average GA performance after REVAC has performed 1,000 measurements.

Table 2.6 and Figure 2.10 show the results. As Figure 2.10 reveals, the performance of the tuned GA on these problems is rather independent from the number of measurement replications employed by REVAC and depends primarily on the number of

Table 2.6: Quality of the tuned parameter values for different levels of measurement replication. Performance is measured in number of fitness evaluations. Results are averaged over all objective functions.

Number of measurement replications	Number of vectors until cost < 2*best	Number of measurements until cost < 2*best	Cost at 1,000 vectors	Cost at 1,000 measurements
1	411	411	9,954	9,789
2	397	795	6,326	7,250
3	241	722	4,783	4,877
5	380	1,901	10,576	10,424
10	277	2,772	9,006	9,072

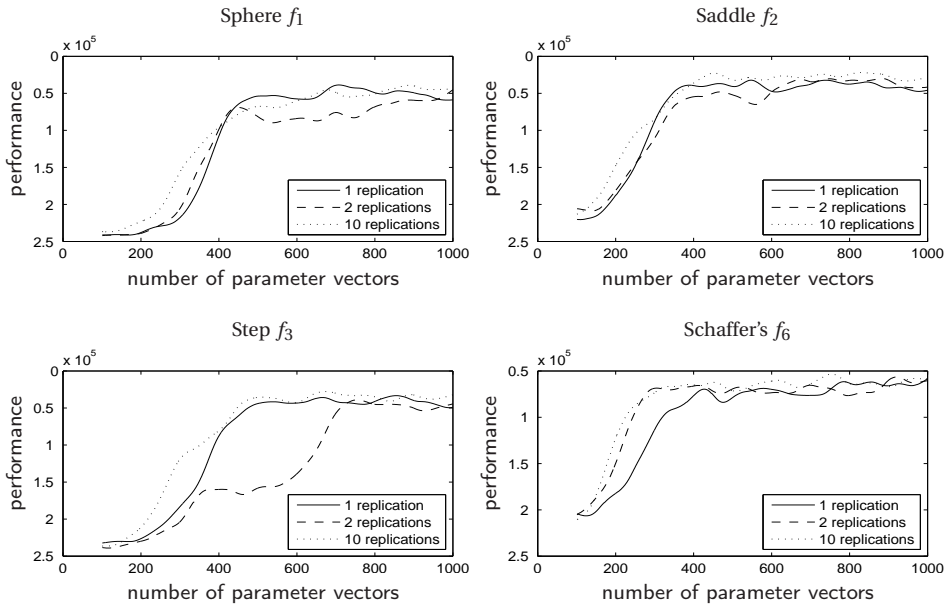


Figure 2.10: EA performance during tuning for three different levels of replications. Performance is calculated as the average number of fitness evaluations of the GA when parameter vectors are drawn from the respective REVAC distributions. The x -axis is counting from top to bottom, so that a higher performance is indeed on top. Graphs are smoothed for readability.

parameter vectors that REVAC has evaluated so far. Note in particular how performance is maximized around parameter vector 400 for all numbers of measurement replication. While measurement replication does not improve the absolute capability of REVAC to tune the parameter values, the performance penalty is huge. The amount of computation needed to reach an arbitrary level of performance increases almost linearly with the level of replication

2.4.2 Assessing algorithmic efficiency on an economic modeling problem

Here we test REVAC as part of our research in evolutionary agent-based economics. The use of evolutionary algorithm to model an economic system comes with a number of unique requirements that warrant a separate verification of whether REVAC can be applied to such modeling problems. These include, but are not limit to, non-standard evolutionary mechanisms, non-linear autocatalytic dynamics, and the need to create stable and realistic system behavior on the population level rather than to find a single optimal solution. Also, economic modeling can force an evolutionary algorithm to include domain specific features of which it is altogether unknown whether the system behavior depends on their correct parameterization.

Experimental setup of Section 2.4.2

design tool	<i>REVAC</i>
evolutionary algorithm	<i>selective imitation in a social network</i>
application problem	<i>dynamic growth and production functions</i>

To describe the experimental setup in a nutshell: 200 agents evolve their investment strategies over a period of 500 time intervals. In each interval each agent invests its current income in a number of economic sectors. The agent's income of the next interval is then calculated according to some production function. The production function changes dynamically, so that the same investment strategy will lead to different growth rates at different points in time. Agents adapt their investment strategies through random mutation and selective imitation in a complex social network. Mutation here is a random change to the way the investment is distributed over the economic sectors. For imitation an agent compares its own economic growth rate to that of its peers in the social network. If a peer has a higher growth rate than that of the comparing agent, the comparing agent can copy the strategy of that peer, wholly or in part, akin to crossover. The evolutionary mechanism will be discussed in detail in Section 4, while Section 5 will elaborate on the growth and production functions that are used here, as well as on the dynamic changes that the agents have to adapt to.

The performance measure of the EA that REVAC has to maximize is the mean log income of all economic agents at the end of a simulation, corresponding to what an economic agent with constant relative risk aversion prefers. Figure 2.11 shows a typical histogram of the performance measure, based on 1,000 runs with identical tuned




Figure 2.11: Histogram of EA performance, based on 1,000 runs.




Figure 2.12: Distribution of the performance after tuning. The y-axis shows the average performance after tuning. The x-axis shows how many REVC runs resulted in an average performance of that level or lower. Note that the median is similar in each graph.

parameter values. The distribution is skewed and has a flat tail, limiting the value of measurement replication. The distribution is not lognormal, but the estimated mean of the logarithmic performance seems to converge faster than the estimated mean of the performance itself and is a more reliable statistic. For this reason we average over the logarithm of the performance measure when we report the performance reached by different sets of tuned parameter values, even though the tuning is done in the original domain.

The algorithm layer has 6 parameters that need to be tuned, corresponding to the simplified evolutionary mechanism at the end of Section 4: *mutation probability*, *mutation variance*, *imitation probability*, *imitation ratio* (how much of the original strategy is preserved), *imitated fraction* (the fraction of well performing peers that are considered for imitation), and the *connectivity* of the social network. For the application layer we consider four different dynamic economic environments: changes occur sudden and with high frequency, sudden and with low frequency, gradual and with high frequency, and gradual and with low frequency.

We use REVC with one, three and ten replications of measurements to tune the algorithm layer to each of the four economic environments. All other REVC parameters are as described before. To improve the reliability of the tuning, we also look into the option of tuning the parameter values several times, choosing those tunings that achieved the highest performance, and averaging over the results. Due to limited computational resources we used different numbers of tunings for each replication scheme: 30 for 1 replication, 10 for 3 replications and 3 for 10 replications.

Figure 2.12 shows the average (log) performance that each tuning achieved during the last 10% of its measurements. Results are sorted per replication scheme to show how the tuned parameter values vary. With only 3 tunings in the case of 10 measurement replications no clear conclusion is possible, but a general trend is visible: the distribution of tuned parameter values is similar for all numbers of replication, with similar mean and variance. The same can be observed for each relevance estimate and each tuned parameter value: all tuning results follow a similar distribution, regardless of the number of measurement replications.

Since not all tuning sessions of REVC achieve the same level of performance, we decide to take only the better 50% and average over the result. To compare REVC with 1 measurement replication and with 3 measurement replications we start by randomly

Table 2.7: REVAC estimates. Average values in bold, followed by the measured variance.

	1 replication		3 replications		10 replications	
	Relevance estimate (absolute entropy)					
Mutation probability	0.6	0.4	0.5	0.2	0.3	0.2
Mutation variance	0.3	0.1	0.2	0.0	1.2	0.3
Imitation probability	0.9	0.3	1.2	0.3	1.1	0.4
Imitation ratio	1.2	0.5	1.8	1.0	1.8	1.0
Imitation fraction	0.8	0.4	0.7	0.1	0.9	0.5
Connectivity	0.1	0.0	0.2	0.1	0.0	0.0
All parameters	3.9	1.0	4.6	1.7	5.3	0.3
	Suggested parameter values					
Mutation probability	0.20	0.03	0.21	0.02	0.45	0.09
Mutation variance	0.28	0.02	0.30	0.03	0.15	0.01
Imitation probability	0.83	0.01	0.86	0.00	0.85	0.01
Imitation ratio	0.90	0.00	0.93	0.00	0.93	0.00
Imitation fraction	0.85	0.01	0.77	0.01	0.83	0.01
Connectivity	0.54	0.04	0.58	0.05	0.51	0.01

selecting 10 out of the 30 runs of REVAC with 1 measurement replication. Of these we take the better 5 and compare their average results to those of the better 5 from the implementation with 3 measurement replications. From the implementation with 10 measurement replications we only use the better 2 tunings. Table 2.7 shows the average relevance estimate (in absolute entropy) for every parameter, and the suggested value for each parameter (the median of the distribution) for one economic environment (sudden, low frequency). Average tuned values for each parameter are shown in bold, followed by the measured variance. Note how the measured variances for the different measurement replications are all of the same order.

To see if each REVAC implementation correctly differentiates between different problems in the application layer we apply each of the four tuned EAs a thousand time to each economic environment and average over the logarithm of the measured performance. This is done separately for each replication scheme. Table 2.8 shows the results. Each row stands for one economic environment and has four entries, showing the results when applying its own set of tuned parameter values and the other three sets of tuned parameter values to that environment. The bold values show the highest value for each row. With correct differentiation we expect to see the highest value for each economic environment when parameters are used that were tuned to that environment. As can be seen, this is almost always the case. The variance of the measured means is below 0.001 and therefore insignificant.

In general one can conclude that there is no significant difference in results obtained with 1, 3, or 10 measurement replication, even though in the case of 1 replication the total number of measurements is significantly smaller. With the exception of one environment, the tuned parameter values perform best on the application to which they

Table 2.8: Performance of the tuned simulation. Each row shows four sets of tuned parameter values applied to the same dynamic environment.

	Gradual, low frequency	Gradual, high frequency	Sudden, low frequency	Sudden, high frequency
1 measurement replication, 10 runs of REVAC				
Gradual, low freq.	2.628	2.619	2.603	2.582
Gradual, high freq.	2.269	2.539	2.524	2.511
Sudden, low freq.	2.686	2.713	2.724	2.727
Sudden, high freq.	2.089	2.233	2.226	2.256
3 measurement replications, 10 runs of REVAC				
Gradual, low freq.	2.610	2.591	2.597	2.584
Gradual, high freq.	2.375	2.531	2.520	2.512
Sudden, low freq.	2.710	2.716	2.733	2.704
Sudden, high freq.	2.102	2.247	2.230	2.258
10 measurement replications, 3 runs of REVAC				
Gradual, low freq.	2.625	2.589	2.595	2.577
Gradual, high freq.	2.202	2.540	2.521	2.502
Sudden, low freq.	2.691	2.712	2.710	2.713
Sudden, high freq.	2.024	2.243	2.202	2.261

were optimized, indicating that REVAC is indeed able to tune parameter values to the problem at hand. One of the design goals of REVAC is to tune parameter values in a robust way so that they work well on problems that are similar to the problem they were tuned on. And indeed, all tuned parameter values achieve good results on all economic environments. To compare, without tuning, the system has a mean logarithmic performance of between 1.7 and 2, depending on the economic environment.

2.5 Comparing REVAC to other tuning methods

In this section we compare the EA parameter values tuned by REVAC to those found by hand-tuning and meta-GA. By hand tuning we mean that the practitioner chooses one or more vectors of parameter values for the EA, evaluates them, and uses the obtained information either to decide on the final parameterization, or to continue and evaluate more vectors of parameter values. This is arguably the most common tuning method even today. Despite the fact that hand tuning can be guided by common wisdom and the extensive experience of a practitioner, it is not effective. Grefenstette (1986) clearly showed that an EA that is tuned by a basic genetic algorithm (GA) outperforms all known hand tuned versions of the EA. Meta-GA tunes an evolutionary algorithm by optimizing a population of parameter vectors through selection, mutation, and recombination. Parent selection is fitness proportional, where the fitness of a vector of parameter values is the mean best fitness returned by the EA of the application layer when executed with

these values. Survival selection is generational, and a population size of 100 parameter vectors is used. To create offspring, one-point crossover is applied with a crossover rate of .5, and thereafter bit-flip mutation with a mutation rate of .001. We use a Gray code to represent the parameter values. The best parameter values are provided by the parameter vector of the last generation of the meta-GA that had the highest fitness.

Experimental setup of Section 2.5	
design tool	<i>REVAC / meta-GA / hand tuning</i>
evolutionary algorithm	<i>Simple GA</i>
application problem	<i>Multi-modal problem generator</i>

To compare REVAC results with those of meta-GA and hand tuning, we follow Eiben et al. (2006) for the application and the algorithm layer. The authors tune a simple genetic algorithm by hand to maximize instances of the multi-modal problem generator (Spears, 2000). They also provide us with a benchmark performance of the hand-tuned algorithm. While the multi-modal problem generator is generally not adequate for assessing the performance of evolutionary algorithms (Lobo and Lima, 2006), Eiben et al. use it specifically to study the effect of different degrees of multi-modality on the performance of the simple GA.

The simple genetic algorithm of the algorithm layer uses a steady-state population model, uniform crossover, bit-flip mutation, tournament parent selection, and delete-worst survival selection. It terminates after 10,000 evaluations. The four free parameters of the algorithm are the crossover rate, the mutation rate, the population size, and the tournament size. The first two parameter values can take values between 0 and 1, encoded with 16 bits. The last two parameter values can take a value between 2 and 1025, encoded in 10 bits.

The multi-modal problem generator works as follows: generate n binary strings of length l to be the local optima. Different local optima are assigned different heights. A point x in the l -dimensional search space is evaluated by first finding the local optimum i with the lowest Hamming distance, i.e., the local optimum that matches x in the largest number of bits. The fitness of x is the fraction of matching bits scaled by the height of i . In case of ambiguity, the highest possible fitness is chosen. Here we use strings of $l = 100$ bits. Eiben et al. define ten different problem classes, each with a different number n of local optima. Those numbers are $\{1, 2, 5, 10, 25, 50, 100, 250, 500, 1,000\}$. The height of the global optimum is 1.

To tune the simple genetic algorithm, both REVAC and meta-GA are allowed 3,000 measurements per tuning session. We run REVAC with 3 measurement replications, so that it has performed 3,000 measurements by the time it has evaluated 1,000 parameter vectors. Table 2.9 shows the best parameter values found by each tuning method. Note that meta-GA and REVAC suggest much larger the values for population size and tournament size than the rather conventional values found by hand-tuning. For both meta-GA and REVAC we find that in two thirds of all solutions the tournament size equals the population size, effectively cancelling tournament selection from the algorithm, something a human designer is not likely to do.

To compare the performance of the best parameter vector of each tuning method, the simple GA is executed 25 times with each best parameter vector on the problem it

Table 2.9: Best parameter values found

	Crossover rate	Mutation rate	Population size	Tournament size
Hand tuned	0.5	0.01	100	2
Meta-GA	0.32	0.017	468	347
REVAC	0.41	0.0043	501	421

was tuned to, for each tuning method. The average performance in terms of mean best fitness is shown in Table 2.10. The highest performance per problem class is printed in bold. The observed differences have no statistical significance. No method performs better than any other on any problem class, which is nicely reflected by the random scatter of bold values over the table.

In our final experiment we measure how robust the tuned parameter vectors are against changes to the problem definition, which in this case amounts to changing the number of local optima. To do so we take a GA with parameter values that are tuned to a problem class with $n = x$ local optima, apply it to a problem class with $n = y$ local optima, and record the mean best fitness. Results are shown in Table 2.11. The average REVAC performance (0.993) seems to be better than that of meta-GA (0.991), but the difference does not have statistical significance.

2.6 Conclusions

In this chapter we introduced and evaluated a customized Estimation of Distribution Algorithm that uses differential Shannon entropy to estimate the relevance of EA parameters. The method searches for high-entropy distributions over the EA parameters that give high probability to parameter values with a high EA performance. Unlike most statistical optimization methods it does not depend on measurement replications to reduce variance. Instead, it reduces variance implicitly by averaging over *adjacent* vectors of parameter values. This allows it to get a good cover of the search space and to extract a high amount of information out of the available measurements. In terms of concrete parameter values, the median of these distributions provides a robust optimum, and the 25th and 75th percentile the confidence interval. The Shannon entropy of these distributions can be used to estimate how much tuning each parameter needs in order to reach a given level of EA performance, independent of the actual tuning method.

The method proves to be able to reproduce the predefined relevance levels of abstract tuning problems to a satisfactory degree, even under high levels of measurement noise from a distribution with far outliers. REVAC results on a simple GA and standard numerical optimization problems are in line with what is reported in the literature. Tests on agent-based simulations of different dynamic economy-environments show that REVAC optimizes a 6-parameter evolutionary algorithm such that each vector of tuned parameter values performs best on the economy-environment it was tuned to, despite a high level of non-Gaussian system noise in the dynamic system. With regard to other tuning methods, we found that the performance of an algorithm tuned by REVAC is roughly comparable to the performance of the same algorithm when tuned by meta-GA.

Table 2.10: Average mean best fitness of the GA tuned with each method

	Number of peaks of the problem class									
	1	2	5	10	25	50	100	250	500	1,000
Hand-tuned	1	1	1	.996	.989	.988	.985	.985	.987	.989
Meta-GA	1	1	.988	.993	.994	.994	.983	.992	.989	.987
REVAC	1	1	1	1	.991	.995	.989	.966	.970	.985

Table 2.11: Mean best fitness when cross-validating the tuned parameter values

	Number of local optima the tuned parameters are applied to										mean
	1	2	5	10	25	50	100	250	500	1,000	
meta-GA											
1	1	1	1	1	.994	1	.999	.986	.986	.988	.995
2	1	1	1	1	1	.980	.977	.981	.991	.998	.993
5	.977	.980	1	.966	.987	.987	.959	.958	.963	.958	.974
10	1	1	1	1	.970	.980	.975	.992	.988	.995	.990
25	1	1	1	1	.994	.983	.989	.986	.994	.987	.993
50	1	1	1	1	1	1	.994	.985	.997	.982	.996
100	1	1	1	1	1	.997	.985	.974	.992	.997	.994
250	1	1	1	1	.987	.997	.995	.970	.986	.970	.990
500	1	1	1	1	.987	.987	.992	.990	.993	.990	.994
1000	1	1	1	.982	1	.980	.985	.995	.992	.982	.992
mean	.998	.998	1	.995	.992	.989	.985	.982	.988	.985	.991
REVAC											
1	1	1	1	1	.990	.990	.998	.990	.994	.991	.995
2	1	1	1	1	.985	.983	.985	.988	.997	.989	.993
5	1	1	1	1	.974	.983	.990	.988	.971	.971	.988
10	1	1	1	1	1	.995	.974	.989	.993	.995	.995
25	1	1	1	1	.990	.998	.994	.994	.988	.996	.996
50	1	1	1	1	1	.975	.990	.991	.993	.995	.994
100	1	1	1	1	.995	.993	.980	.986	.980	.992	.993
250	1	1	1	1	.990	.985	.975	.987	.990	.988	.992
500	1	1	1	.987	.995	.975	.980	.992	.999	.968	.990
1000	1	1	1	1	1	.988	.994	.996	.975	.970	.992
mean	1	1	1	.999	.992	.986	.986	.990	.988	.986	.993

Notes. Rows show the problem class to which the parameters are tuned (labeled by the number of local optima n), columns show the problem class to which the parameters are applied (labeled by the number of local optima n).

A STUDY OF PARAMETER RELEVANCE IN EVOLUTIONARY ALGORITHMS

Abstract

We present an empirical study on the impact of different design choices on the performance of an evolutionary algorithm (EA). Four EA components are considered—parent selection, survivor selection, recombination and mutation—and for each component we study the impact of choosing the right operator and of tuning its free parameter(s). We tune 174 different combinations of EA operators to 4 different classes of fitness landscapes and measure the cost of tuning. We find that components differ greatly in importance. Typically the choice of operator for parent selection has the greatest impact, and mutation needs the most tuning. Regarding individual EAs however, the impact of design choices for one component depends on the choices for other components, as well as on the available amount of resources for tuning.

3.1 Introduction

Evolutionary Algorithms (EA) form a class of search methods that work by incrementally improving the quality of a set of candidate solutions by variation and selection (Eiben and Smith, 2003). The most important *components* of EAs are thus recombination and mutation (umbrella term: variation), parent selection, and survivor selection. To obtain a working EA, each component needs to be instantiated by a specific *operator*, e.g., the one-point crossover operator for the recombination component. Furthermore, an EA has *parameters* that need to be instantiated by a specific *parameter value*, e.g., 0.5

This chapter is an extension of Nannen et al. (2008).

for the crossover rate. In this paper we maintain the distinction between components and parameters and say that the instantiation of EA components by concrete operators specifies a particular EA, e.g., uniform crossover, bit-flip mutation, random parent selection and k -tournament survivor selection. Further details regarding the parameters do not lead to a different EA, only to variants of the one defined by the operators.¹ A complete EA design includes the definition of an EA (operators for its components) and the specification of a particular variant of it (values for its parameters).

Setting EA parameters is commonly divided into two cases, parameter tuning and parameter control (Eiben et al., 1999, 2007). In case of parameter control the parameter values are changing during an EA run. This requires initial parameter values and suitable control strategies, which in turn can be deterministic, adaptive or self-adaptive. The problem of parameter tuning is hard because for any given application there is a large number of options, but only little knowledge about the effect of EA parameters on EA performance. EA users mostly rely on conventions (mutation rate should be low), ad hoc choices (why not use uniform crossover), and experimental comparisons on a limited scale (testing combinations of three different crossover rates and three different mutation rates). Here we address the problem of parameter tuning. Our main research questions are:

1. How does the choice of operator for each component contribute to EA performance? To this end we compare the absolute performance achieved with different combinations of operators.
2. The parameters of which EA component need the most tuning? For this question we measure the amount of information needed to tune the free parameter(s) of each operator (e.g., crossover rate or tournament size).

For a systematic exploration of the space of EA configurations we use exhaustive search for the combination of operators and Relevance Estimation and Value Calibration (REVAC) to tune the free (numeric) parameters. REVAC is an Estimation of Distribution Algorithm (Mühlenbein and Höns, 2005) that tunes an EA by optimizing marginal probability distributions over the free parameters, see Section 2. Starting from a set of uniform distributions and an initial drawing of 100 vectors of random parameter values, REVAC iteratively generates new marginal distributions of increasing expected EA performance by *drawing* a new vector of parameter values from the current distributions, *evaluating* the vector by measuring the performance of the EA with these values, *updating* all marginal distributions based on this evaluation, and *smoothing* the updated distributions. Smoothing is a unique feature of REVAC that forces all marginal distributions to approach the maximum Shannon entropy distribution for a given EA performance. This maximized Shannon entropy is independent from the computational cost of any particular tuning method and can be used as a general estimator of the minimum amount of information required to reach a certain level of EA performance. Hence, it can be regarded as a general indicator of how difficult it is to tune a certain EA parameter, and how relevant it is to overall EA performance.

¹Alternatively, components & operators could also be called symbolic parameters & values, and we could say these values only define different EA variants.

Related work includes the general discussion of EA design (Czarn et al., 2004) and parameter setting (Lobo et al., 2007), in particular within parameter tuning as defined by Eiben et al. (1999), Birattari (2004), and Eiben et al. (2007). Throughout the relevant literature we find that the cost of tuning parameters is largely ignored. Notable exceptions are the theoretical considerations of Oliver et al. (1987) and Goldberg (1989), as well as the systematic parameter sweeps of Jong (1975), Schaffer et al. (1989), and Samples et al. (2005) and the statistical analysis of parameters by François and Lavergne (2001). In the general field of experimental design, a paradigm shift that emphasizes a low cost of tuning over the performance of optimal parameter values was due to Taguchi and Wu (1980). In our field, Freisleben and Hartfelder (1993) propose a meta-GA approach in which both EA components and EA parameters are tuned and shows the importance of the right choice for the GA operators. Samples et al. (2005) show how parameter sweeps can be used for robustness and correlation analysis.

3.2 Experimental setup

For a clear discussion we distinguish three different layers in the analysis of an EA: the problem/application (here: fitness landscapes created by a generator), the problem solver (here: an EA), and the method for tuning the problem solver (here: REVAC). For an unbiased study we use independent software implementations for each layer and combine them through simple interfaces. For the problem layer we use a generator of real-valued fitness landscapes that are formed by the max-set of Gaussian curves in high dimensional Cartesian spaces (Gallagher and Yuan, 2006). Where a Gaussian mixture model takes the average of several Gaussians, a max-set takes their enveloping maximum, giving full control over the location and height of all maxima. For the implementation we followed Rudolph (1992) on rotated high dimensional Gaussians, and used 10 dimensions, 100 Gaussians, and the same distributions over height, location, and rotation of these Gaussians as specified in the exemplary problem sets 1–4 of Gallagher and Yuan (2006). These sets offer an increasing amount of exploitable structure to the EA. Set 1 has the least structure, with peaks of different height scattered at random, while set 4 is the most structured, with peaks that get higher the closer they get to the origin. For each set, different landscapes are created by passing a different random seed to the generator. Initialization of all EA populations is uniform random in the domain of the fitness landscapes. The optimal fitness value is 1 on each problem instance and the condition for successful termination is defined as “fitness > 0.9999 or 10,000 fitness evaluations”.

For the EAs we use the Evolutionary Computation toolkit in Java (ECJ) (Luke), which allows the specification of a fully implemented EA through a simple parameter file, including the choice of operator for each component and the values for the free parameters. The ECJ offers several operators for each EA component, cf. Table 3.1. For any given EA, the population size parameter is always present. Most operators have zero or one free parameter. One operator has 2 free parameters—Gaussian(σ , p) with parameters σ for step size and p for mutation probability, which takes the value 1 in case of Gaussian(σ , 1). Due to technical details of the ECJ, not all combinations of operators are possible. For example, neither fitness proportional nor best selection as parent selection can be combined with tournament selection as survivor selection. In total we have

Table 3.1: EA components, operators, and parameters used in this study

Component	Operator	Parameter(s)
		population size μ
parent selection	tournament	parent tournament size
	best selection	number n of best
	fitness proportional	-
	random	-
survivor selection	generational	-
	tournament	survivor tournament size
	(μ, λ)	λ
	$(\mu + \lambda)$	λ
recombination	random	-
	none	-
	one-point	crossover probability
	two-point	crossover probability
mutation	uniform	crossover probability
	reset (random uniform)	mutation probability
	Gaussian($\sigma, 1$)	step size
	Gaussian(σ, p)	step size, mutation probability

Notes. We follow the naming convention of the ECJ. Arguably, (μ, λ) and $(\mu + \lambda)$ define both parent *and* survivor selection. Here we classify them under survivor selection because that is what the parameter λ influences.

174 combinations of operators, of which 6 with 2, 33 with 3, 65 with 4, 55 with 5, and 15 with 6 free parameters.

The performance of an EA with a given set of parameter values is measured in three different ways: SR (Success Rate, percentage of runs with fitness > 0.9999), MBF (Mean Best Fitness of all runs), and AES (Average number of Evaluations to Solution of successful runs; undefined when SR = 0). Each EA is tuned 5 times on each of the 4 problem sets. During each tuning session on a given set REVAC generates 1,000 different vectors of parameter values. Each vector of values is written to the ECJ configuration file, together with the specification of the operators and the problem generator. The resulting EA is evaluated on 10 different instances of the problem set, generated by different random seeds.

For each REVAC tuning session and each EA, the best performance after n evaluations is the best performance measured after evaluating n vectors of parameter values. The average best performance after n evaluations is averaged over multiple tuning sessions on the same EA. We define *near best performance* as the average best performance after 1,000 evaluations minus 5%². If n is the lowest number of vectors for which the average best performance after n evaluations exceeds this value, then we say that REVAC

²In case of MBF this is calculated after subtracting a performance level of .5, a level that is reached by any reasonably sized population upon random uniform initialization. The maximum possible near best MBF is therefore .975

needs n evaluations to tune the EA to near best performance. Section 3.3 uses this to study the impact of choosing an operator for each component.

In Section 3.4 we analyze the cost and benefits of tuning per EA component. REVAC continuously maximizes the Shannon entropy of the marginal distributions that it optimizes during a tuning session. This maximized Shannon entropy provides a generic information-theoretic measure of the minimum amount of information needed per parameter to reach a given performance level. The differential Shannon entropy h of a probability density function \mathcal{D} over the continuous interval $[a, b]$ is commonly defined as

$$h(\mathcal{D}_{[a,b]}) = - \int_a^b \mathcal{D}(x) \log_2 \mathcal{D}(x) dx. \quad (3.1)$$

The sharper the peaks of a probability density function, the lower its Shannon entropy. In order to compare the entropy of distributions that are defined over different parameter intervals in a meaningful way, we normalize all parameter intervals to the interval $[0, 1]$ before calculating the Shannon entropy. In this way the initial uniform distribution has a Shannon entropy of zero, and any other distribution has a negative Shannon entropy $h(\mathcal{D}_{[0,1]}) < 0$.

3.3 How does the choice of operator per component contribute to performance?

Figure 3.1–3.8 on page 56–61 show the average near best fitness in AES and MBF and the tuning cost in number of evaluations of REVAC needed to reach this fitness. Results in AES are based on those EAs with $SR > 0$ for which the AES could be calculated. Depending on the problem set, 60–70 EAs could be tuned to terminate with success. Results on the MBF are based on all 174 EAs. Each figure contains four scatter plots that show the performance in AES or MBF after tuning, and the cost of tuning, averaged over 5 tuning sessions per EA. The y -axes show the near best performance in AES. The x -axes show the number of REVAC evaluations needed to tune the EA to this performance. Each of the four plots of a figure shows the same EAs but labels them according to the operator choice for a different component. To read the full specification of an EA, one needs to look at the same location in all four plots. Under each AES figure a table (i.e., Table 3.2–3.5) shows the average performance on the same set that was achieved with each operator of each component. The table reports the average success rate, the near best performance, and the cost of finding it, for AES and for MBF respectively. Each value is averaged over all those EAs that use this operator and, in case of AES, terminated with success.

The choice of operator for the parent selection component has the strongest effect on EA performance. The 16 EAs that are clustered together in the lower left of each plot of Figure 3.1–3.4 display the best performance and the lowest number of evaluations needed to reach this performance. These EAs all use tournament selection for parent selection, either tournament selection or random selection for survivor selection, any recombination operator, and either Gaussian(σ, p) or Gaussian($\sigma, 1$) for mutation. On the other hand, those EAs that have the never terminate with success and generally reach only a minimum MBF (see Figure 3.5–3.8) share one common feature, namely a lack of

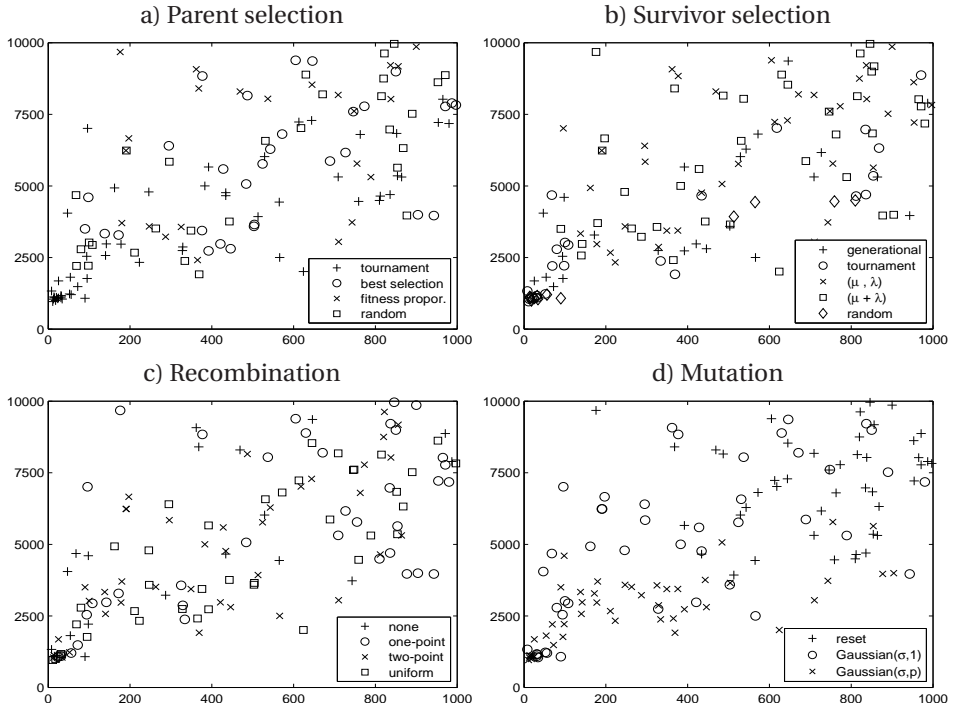


Figure 3.1: Near best performance in AES on set 1 against cost of tuning

Table 3.2: Average near best performance on set 1, per operator

Component	Operator	Success rate	AES	Cost	MBF	Cost
parent selection	tournament	1	3484	344	0.87	281
	best select.	1	5782	556	0.89	279
	fitness pro.	0.64	6582	532	0.87	219
	random	0.54	5611	517	0.77	388
survivor selection	generational	0.5	4224	426	0.8	367
	tournament	1	3320	318	0.87	324
	(μ, λ)	0.97	6259	546	0.89	209
	$(\mu + \lambda)$	0.97	6072	550	0.9	210
	random	0.5	2175	244	0.74	444
recombination	none	0.67	4798	343	0.82	328
	one-point	0.79	5549	548	0.84	289
	two-point	0.81	4727	412	0.85	273
	uniform	0.81	4732	479	0.85	324
mutation	reset	0.78	7168	729	0.83	309
	Gauss. ($\sigma, 1$)	0.78	5001	368	0.85	306
	Gauss. (σ, p)	0.79	2785	285	0.85	287




Figure 3.2: Near best performance in AES on set 2 against cost of tuning

Table 3.3: Average near best performance on set 2, per operator

Component	Operator	Success rate	AES	Cost	MBF	Cost
parent selection	tournament	1	3559	325	0.91	265
	best select.	1	5870	654	0.92	157
	fitness pro.	0.61	6042	590	0.91	162
	random	0.54	5357	479	0.8	324
survivor selection	generational	0.5	4708	427	0.83	248
	tournament	1	3524	181	0.92	255
	(μ, λ)	0.92	5792	595	0.93	201
	($\mu + \lambda$)	1	5643	663	0.93	157
	random	0.5	2529	185	0.77	431
recombination	none	0.67	5046	322	0.85	275
	one-point	0.77	5251	523	0.88	245
	two-point	0.81	4705	515	0.89	232
	uniform	0.81	4559	468	0.89	233
mutation	reset	0.76	7259	689	0.87	239
	Gauss. ($\sigma, 1$)	0.78	4620	430	0.88	232
	Gauss. (σ, p)	0.79	2810	313	0.88	260




Figure 3.3: Near best performance in AES on set 3 against cost of tuning

Table 3.4: Average near best performance on set 3, per operator

Component	Operator	Success rate	AES	Cost	MBF	Cost
parent selection	tournament	1	3526	365	0.92	253
	best select.	1	5636	593	0.93	203
	fitness pro.	0.64	6085	560	0.91	192
	random	0.54	5329	529	0.8	325
survivor selection	generational	0.5	4622	397	0.84	263
	tournament	1	3395	266	0.93	242
	(μ, λ)	0.95	5856	612	0.94	225
	$(\mu + \lambda)$	1	5525	621	0.94	176
	random	0.5	2451	250	0.77	424
recombination	none	0.7	5149	342	0.85	323
	one-point	0.77	5233	522	0.89	245
	two-point	0.81	4590	502	0.89	237
	uniform	0.81	4442	504	0.89	238
mutation	reset	0.76	7283	656	0.87	253
	Gauss. ($\sigma, 1$)	0.79	4532	493	0.89	260
	Gauss. (σ, p)	0.79	2719	309	0.89	250




Figure 3.4: Near best performance in AES on set 4 against cost of tuning

Table 3.5: Average near best performance on set 4, per operator

Component	Operator	Success rate	AES	Cost	MBF	Cost
parent selection	tournament	0.96	5078	408	0.89	319
	best select.	0.8	6782	649	0.89	249
	fitness pro.	0.47	6778	643	0.88	266
	random	0.43	6872	480	0.77	398
survivor selection	generational	0.5	6554	493	0.8	353
	tournament	0.92	5370	298	0.9	313
	(μ, λ)	0.74	6660	631	0.9	265
	$(\mu + \lambda)$	0.74	6321	663	0.9	271
	random	0.5	4019	235	0.73	434
recombination	none	0.6	6446	375	0.81	390
	one-point	0.65	6535	490	0.85	311
	two-point	0.69	5601	545	0.86	283
	uniform	0.71	5785	554	0.86	325
mutation	reset	0.45	9250	630	0.84	323
	Gauss. ($\sigma, 1$)	0.76	5890	514	0.86	301
	Gauss. (σ, p)	0.79	4358	430	0.86	338




Figure 3.5: Near best performance in AES on set 1 against cost of tuning

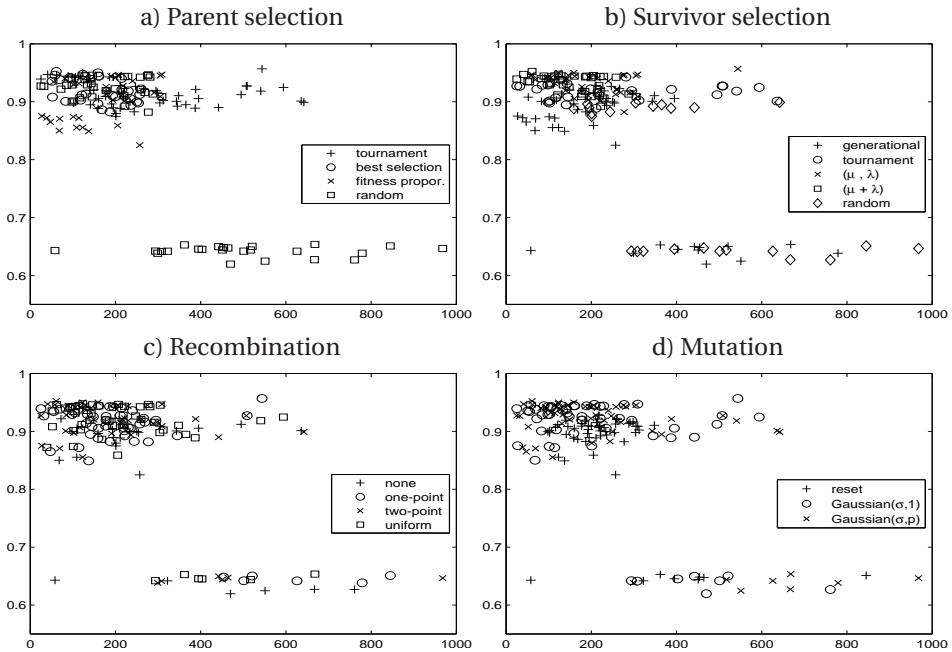


Figure 3.6: Near best performance in AES on set 2 against cost of tuning




Figure 3.7: Near best performance in AES on set 3 against cost of tuning




Figure 3.8: Near best performance in AES on set 4 against cost of tuning




Figure 3.9: Impact of recombination operators on AES and cost of tuning

selection pressure. In particular, EAs with random or fitness proportional selection for parent selection do not terminate with success unless combined with strong survivor selection pressure.

Of the two variation components, the choice of mutation operator has the stronger effect on EA performance, as can be seen from the differences in Table 3.2–3.5. On all our problem sets reset mutation is the worst mutation operator, and non-standard Gaussian(σ, p) mutation is superior to Gaussian($\sigma, 1$) both in terms of performance and in terms of cost of tuning. The latter may come as a surprise, since the additional free parameter for mutation probability increases the parameter search space. We conclude that the tuning cost of different operators is not additive, and that the tuning cost of an operator can only be evaluated in the context of the overall EA composition.

While choosing the recombination operator has the least effect on EA performance, it demonstrates how the choice of operator can depend on the available resources for tuning. Figure 3.9 enlarges the lower left corner of Figure 3.4c, overlaid by four graphs that show the evolution of the average performance of 4 EAs with tournament selection for both parent and survivor selection, Gaussian(σ, p) mutation, and four different recombination operators. 20 tuning sessions were used for each graph. While the tunable recombination operator eventually outperform no recombination, an EA with no recombination consistently outperforms EAs with tunable recombination after about 30–40 parameter vectors have been evaluated, and it has at least average performance for anything under 100 evaluated parameter vectors. We observed this phenomenon over a wide range of operator choices for the other components and over all 4 problem sets. All in all, for recombination, the choice of operator can clearly depend on the amount of effort that can be invested in tuning.

3.4 Which EA component needs the most tuning?

The previous section related the performance of the near best parameter vector to the number of REVAC evaluations needed to find this vector and to achieve this performance. This section takes a rather unconventional approach based on the expected performance when parameter values are drawn from a probability distribution, namely those created by REVAC after 500 evaluations. To calculate the performance gain achieved by tuning, this expected performance is compared to the expected EA performance




Figure 3.10: Correlation between Shannon entropy and standard deviation of tuned parameter values




Figure 3.11: Correlation between performance gain and the sum of the Shannon entropy of all marginal distributions

when parameter values are drawn from the uniform distribution. All results are averaged over 5 REVAC tuning sessions of an EA on each of the 4 problem sets, 20 tuning sessions per EA. In order to extend our analysis to all 174 EAs, we use the Mean Best Fitness that an EA achieves at termination (successful or not), rather than the AES.

Shannon entropy measures the amount of information that a probability distribution provides on its random values. By definition, the lower the Shannon entropy of the maximum entropy distribution that achieves a given expected EA performance, the finer the parameter value has to be tuned in order to achieve that expected performance. This is demonstrated in Figure 3.10. The scatter plot shows the correlation between the Shannon entropy of the marginal distribution over the mutation probability and the standard deviation of the best found parameter values. The x -axis shows the Shannon entropy as estimated by REVAC. The y -axis shows the average (over 4 sets) of the standard deviation of the 5 best values found during the 5 REVAC tuning sessions on each set. The correlation coefficient is 0.9 and the p -value (the probability to observe this or a stronger correlation when the true coefficient is zero) virtually zero. The point here is that if the maximum entropy distribution has a higher Shannon entropy, there is less certainty on the precise parameter value, something that can otherwise be expensive to assess.

Figure 3.11 shows a clear correlation between a gain in expected MBF and the Shannon entropy of the maximum entropy distributions that REVAC has estimated after 500 evaluations. The x -axis shows the average performance gain in percent. The y -axis

Table 3.6: Entropy per EA component, averaged over all EAs

Component & population size		Correlation with MBF gain		Shannon entropy		
		correlation	p-value	max	mean	min
1)	population size	-0.4	0	0	-0.9	-1.7
2)	parent selection	0.2	0.1	-0.1	-0.6	-3.3
3)	survivor selection	-0.1	0.3	0	-0.7	-1.4
4)	recombination	-0.4	0	0	-0.2	-1.2
5)	mutation	-0.6	0	-0.1	-1.2	-4.5
Entire EA		-0.8	0	-0.2	-2.9	-5.1




Figure 3.12: Mean Shannon entropy per component & population size

shows the Shannon entropy of the estimated distributions, summed over all tuneable parameters of the EA. Note that no EA lies above the main diagonal, which shows that there is a minimum information cost for every percent point of gain in expected performance, regardless of the EA specifications. Of those EAs that lie significantly below the diagonal, most use tournament selection for both parent and survivor selection. By 500 REVAC evaluations, their MBF had long been maximized. Further tuning only improved their AES, distorting their performance gain to entropy ratio.

Does the strong correlation between total Shannon entropy and the gain in expected performance carry over to individual EA components? The first two numeric columns of Table 3.6 show the correlation coefficient for each component and its p-value, i.e., the probability to observe this or a stronger correlation coefficient if the true coefficient is zero. Only EAs with a tunable operator were considered for the respective component. The correlation is generally weak, in particular for selection. In other words, the question which component needs tuning in order to improve the performance of a particular EA depends much on the EA in question.

With respect to the average Shannon entropy per component, we see that not all components require the same amount of tuning. The right numeric columns in Table 3.6 show the maximum, mean, and minimum Shannon entropy that we observed for each component (and the population size) when instantiated with an operator that needs tuning. The bar diagram of Figure 3.12 allows a visual comparison of this average mean Shannon entropy. Such a skewed distribution of a need for tuning is commonly known as *sparsity of effects*.

Typically, mutation requires the highest amount of tuning, and recombination the least. This rule has many exceptions, as can be concluded from the low correlation co-

efficient. While the relative order of Shannon entropy per component depends much on the EA in question, consistent patterns can be detected for small groups of EAs. Take for example the two EAs with tournament selection for both parent and survivor selection, Gaussian($\sigma, 1$) mutation and either one-point, two-point or uniform crossover. We find that the Shannon entropy for mutation has the unusually high Shannon entropy of around -2 , while the parent selection operator has a low Shannon entropy below -3 . When combining the same selection operators with other recombination or mutation operators, we find that the Shannon entropy for parent selection is back to normal levels, while it is still comparatively high for mutation. Another example is recombination, which only exhibits a low Shannon entropy for uniform crossover in combination with either $(\mu + \lambda)$, or (μ, λ) . Such irregular patterns are consistent over different problem sets and seem to be inherent to specific combinations of EA components.

3.5 Conclusions

This chapter introduces a novel approach to EA design that emphasizes the cost of tuning. To understand how this cost depends on the choice of operator per EA component, we combined exhaustive search over operators with REVAC for tuning their parameters. Our experiments revealed a number of notable insights.

Our tests confirmed the common wisdom that the choice of operator for one EA component depends on the choice of operator for the other components. Of all components, the choice of operator for parent selection has the biggest impact on EA performance. Furthermore, EAs differ greatly in the amount of tuning needed to reach a given performance, and this tuning cost depends on the overall setup of the EA, rather than the number of free parameters. With regard to recombination, we found that the best EA setup depends on the time and effort one can permit to tune the EA.

To measure the need for tuning per component we use the Shannon entropy of maximum entropy distributions as estimated by REVAC, which expresses the minimum amount of information that is needed to achieve a given expected EA performance. It is a generic information-theoretic measure that is independent of any particular tuning algorithm. Inspired by theoretical considerations, it was validated by a strong correlation with the standard deviation of best solutions found during multiple tuning sessions. Based on this measure we observed that the need for tuning follows a skewed distribution, and that while total Shannon entropy is strongly correlated with performance gain, the correlation per component is weak. The question which component needs the most tuning depends on the precise composition of an EA and can not be answered on a general level. It needs to be addressed by the operational analysis of individual EAs. We recommend that a scientific discussion of individual operators addresses their effect on the overall tunability of an EA and on the need for tuning per component.

Regarding the scope of our results, an empirical study can only use a limited set of test problems, and strictly speaking our findings are only proven for our test problems. However, we consider it unlikely that the complex picture that has emerged here is an artefact of the test problems. What remains to be studied is whether the way in which the need for tuning per component depends on the choice of operator for other components is different on other complex fitness functions.

HOW TO EVOLVE STRATEGIES IN COMPLEX ECONOMY-ENVIRONMENT SYSTEMS

Abstract

An evolutionary model of economic behavior is not plausible if its parameters need excessive tuning. Here we illustrate how Relevance Estimation and Value Calibration (REVAC) can help to find a simple and robust model of an evolutionary system that allows the agents to adapt well to complex environmental dynamics. We apply REVAC to tune two versions of an evolutionary agent-based economic simulation, one where agent behavior is parameterized differently based on relative welfare, and one where there is no such distinction. We find that for equal levels of performance of the evolutionary model, the extra features of the first model increase the overall need for tuning. They should therefore be discarded. We find further that tuning those parameters that control the diversity of strategies is most relevant to the adaptive capabilities of the agents.

4.1 Introduction

One of the canonical challenges in evolutionary computing is to select and tune parameters of an evolutionary algorithm (EA) (Eiben et al., 1999; Eiben and Smith, 2003), i.e., parameters that regulate variation (mutation and recombination), selection, population size, and so on. Often these parameters need to be optimized such that the EA delivers good and robust solutions for a whole family of similar problems. This is true for “traditional” optimization and design applications. For instance, when solving a scheduling

This chapter is an extended version of Nannen and Eiben (2006), which has won a best paper award at the Genetic and Evolutionary Computation Conference 2006 in Seattle.

task with a genetic algorithm, it can be hard to establish good values for the mutation rate, crossover rate, tournament size and population size that give good solutions for all possible problem instances. The problem intensifies in more complex applications like agent-based simulations in the fields of artificial life, artificial societies and evolutionary economics. In such applications evolution is not only the “problem solver” that is expected to lead to “optimality” in some application specific sense. It has to fit in with the general system description and provide a better understanding of the general dynamics of the evolutionary system under investigation.

When used to model real life phenomena, the evolutionary algorithm can include domain specific features that are deemed essential to the simulated evolutionary process. For instance, mating selection can depend on past interactions between individuals, and mutation can be sensitive to environmental factors. When asking whether such features do indeed benefit the evolutionary process in a robust way—e.g., without the need for excessive tuning—common EA wisdom (heuristics and conventions learned over the decades) regarding EA setup is hardly applicable, since this wisdom is mainly based on the traditional task of finding optimal parameter values. In contrast, REVAC provides an information-theoretic measure on how much tuning each parameters of an EA needs so that the EA reaches a given performance, independent of the actual tuning method. This can be used to evaluate the benefits a domain specific feature, as well as to choose between different possible sets of features, from the point of view of robust performance.

We illustrate this robustness test by applying REVAC to the evolution of investment strategies in an economic simulation. We provide a summary of the agent-based application, the non-linear system dynamics that the agents have to adapt to, and the specific evolutionary algorithm which consists of random mutation and selective imitation (recombination) of investment strategies in a social peer network. We describe our initial evolutionary algorithm of 13 parameters, reflecting our best intuitions on the evolutionary dynamics in the given context. Next we describe how REVAC effectively disproves our initial intuitions and leads us to a simplified evolutionary algorithm of 6 parameters that allows the evolutionary agents to reach the same level of aggregate welfare. Because the simpler evolutionary model needs less tuning to achieve the same aggregate welfare, we conclude that its predictive power and general validity have improved.

Some fundamental insights by A. Kolmogorov on the relation between individual data and (probabilistic) sets that contain them were published only recently (Vereshchagin and Vitányi, 2002). Early attempts to relate the generalization power of a statistical model to some practical estimate of algorithmic complexity were based on the number and precision of the parameters involved: first the Akaike Information Criterion (AIC) (Akaike, 1973) and then the Minimum Description Length (MDL) principle (Rissanen, 1978). Later, J. Rissanen, A. Barron and B. Yu developed a version of MDL based on parametric complexity (Barron et al., 1998). All these methods are based on a functional analysis of the statistical model in question. This is not possible here, simply because no analytic tool can tell us how many previously unsolvable problems can be solved by adding feature x to an EA. REVAC is intended to fill this gap by numerical estimation.

4.2 The economic model

4.2.1 General features of the model

The agent-based application treated here is concerned with a finite number of economic agents which may be interpreted as national or regional authorities in charge of domestic energy policy. The agents are challenged to adapt their investment strategies to resource constraints and technological change. The investment strategy of each agent specifies how it allocates its investment over these sectors. Initially all agents use fossil fuel for their energy needs, which has finite supply. In order to sustain economic growth, the agents need to identify a viable source of renewable energy from among a number of nonviable alternatives. Invested capital is non-malleable: once invested it cannot be transferred between sectors. Standard economic growth and production functions describe how capital accumulates in each sector and contributes to income. These functions are not aggregated: growth and returns are calculated independently for each agent. Two agents with different investment strategies can experience very different growth rates and income levels.

The numerical simulations are based on a discrete synchronous time model where the income and strategy of each agent is updated in parallel at fixed time intervals. Each simulation step is divided into two separate update operations: *updating the economy*—each agent invests its income according to its own investment strategy and the individual incomes and growth are calculated by the non-aggregate growth and production and growth functions—and *updating the strategies*, when all agents compare their growth rate with that of their peer group, and when those agents that decide to imitate change their respective strategies simultaneously. Each computer simulation is divided into an initialization phase of 50 time steps during which all strategies are fixed, and a main experimental phase of 500 time steps during which all agents are free to change their strategy. The initialization phase is needed to avoid influencing the simulation results by the choice of initial values. During initialization the simulated economy stabilizes and a “natural” distribution of strategies and growth emerges. All initial strategies are drawn independently at random from the space of possible strategies.

4.2.2 Strategies, investment, and production

The economy has $n = m + 4$ investment sectors: consumption C , general capital K , fossil energy F , one viable renewable energy source R_0 and m nonviable alternative energy sources R_1, \dots, R_m . The number m of nonviable alternatives controls the difficulty of finding the viable source of renewable energy. Formally, an investment strategy $s_a(t)$ of agent a at time t is an n -dimensional vector that specifies what fraction of income the agent invests in the respective sectors,

$$s_a(t) = [0, 1]^n, \quad \sum_i s_{ia}(t) = 1. \quad (4.1)$$

The first fraction $s_{1a}(t) = s_{C,a}(t)$ specifies the fraction of income that is consumed; the second fraction $s_{2a}(t) = s_{K,a}(t)$ specifies the fraction of income that is invested in general capital; the third fraction $s_{3a}(t) = s_{F,a}(t)$ specifies the fraction of income that is in-

vested in fossil energy; and the remaining fractions specify what is invested in the m renewable energy sectors.

Growth in each sector except consumption depends on investment, the learning factor L that reflects the state of technology in that sector and a depreciation which is constant and equal for all sectors and agents. The availability of fossil fuel is physically limited and divided among the agents according to their relative investment. Growth in this sector is therefore at a disadvantage when compared to renewable energy sectors. We use a depreciation δ of .05. The growth functions are

$$\Delta K_a(t) = s_{K_a}(t)Y_a(t)L_K - \delta K_a(t), \quad (4.2)$$

$$\Delta F_a(t) = \frac{s_{F_a}(t)Y_a(t)L_F}{\sum_{b \in P} s_{F,b}(t)Y_b(t)} - \delta F_a(t), \quad (4.3)$$

$$\Delta R_{i_a}(t) = s_{R_{i_a}(t)}Y_a(t)L_{R_i} - \delta R_{i_a}(t) \quad \text{for each } R_i. \quad (4.4)$$

The learning factor L grows endogenously with the log cumulative investment of all agents in that sector, multiplied by sector specific learning rate z_i (for a discussion of the learning function see Nordhaus (2002)). For convenience we use the same depreciation for technology as for capital, $\delta = .05$, implicating that half of all technological achievements become outdated or otherwise irrelevant after 13 to 14 time steps.

$$\Delta L_i(t) = z_i \log(1 + \sum_a s_{i_a}(t)) - \delta L_i(t) \quad \text{for all } i. \quad (4.5)$$

The learning rate z determines how fast a technology develops with investment. To allow a stable economic growth of 2–3% per time step of the simulation we use $z_K = .01$. Given the resource constraint on fossil fuels we need a high $z_F = 1$ so that fossil fuel supplies can initially satisfy rising demand. To make one renewable energy a viable alternative to fossil energy we give it the same learning rate $z_{R_0} = .01$ as general capital. The learning rate of all other renewable energy technologies is so low that any investment in them has no long term effect. $z_{R_1}, \dots, z_{R_m} = .0001$

The domestic income $Y_a(t)$ of agent a at time t is calculated by a Cobb-Douglas type production function with constant returns to scale. In this function fossil energy and renewable energy are perfect substitutes—one can completely replace the other. General capital and the energies are imperfect substitutes—investing everything or nothing in energy will ruin the economy, and the best distribution of investment over general capital and energy depends on the production coefficient α . We set this coefficient to $\alpha = .9$, so that agents have to invest about 10% of their total income in energy in order to achieve healthy growth rates. The production function is

$$Y_a(t) = K_a(t)^\alpha \left(F_a(t) + \sum_{i=0}^m R_{i_a}(t) \right)^{1-\alpha}. \quad (4.6)$$

The welfare of an individual agent a at time step t is measured by its individual investment in consumption $C_a(t)$, which is calculated as

$$C_a(t) = s_{C_a}(t)Y_a(t). \quad (4.7)$$

The aggregate welfare $W(t)$ at the population level is calculated from the discounted mean of the logarithm of individual welfare,

$$W(t) = \frac{d^t}{|P|} \sum_{a \in P} \log C_a(t), \quad (4.8)$$

where $|P|$ is the size of the population and where d is the rate at which future welfare is discounted. We use a discount rate of $d = .97$.

4.2.3 The social network

As has been extensively discussed by Wilhite (2006), agent-based simulation of economic processes needs to give proper attention to the social network. We use a generic class of social networks that reproduce a number of stylized facts commonly found in real social networks, namely small world networks (Erdős and Rényi, 1959) that have a scale-free degree distribution generated by a stochastic growth process with preferential attachment (Barabási and Albert, 1999) and that have a high clustering coefficient \mathcal{C} (Watts and Strogatz, 1998).¹ According to Tomasini (2005), an evolutionary algorithm with spatial structure is of advantage when dealing with dynamics problems. Lieberman et al. (2005) have shown that spatial structures like scale-free networks are a potent selection amplifier for mildly advantageous mutants.

Before the start of each simulation we use a stochastic process to generate a new bi-directional network where the nodes are agents and the edges are communication links. The process assigns to each agent a a set of peers N_a that does not change during the course of the simulation. If agent a is a peer of agent b , then a will consider the income growth rate and the investment strategy of b when choosing an agent for imitation, while b will consider the income growth rate and the investment strategy of a . On the other hand, if a and b are not peers, they will not consider each other for the purpose of imitation. The generating process starts from a circular network where each agent has two neighbors—i.e., average connectivity $k=2$ —and iteratively adds new edges to the network until the desired average connectivity k is reached. The agents for the next new edge are chosen at random with a probability that is proportional to their connectivity (hence the term “preferential attachment”) and their proximity in the network, i.e., the inverse of the minimum number of links to traverse from one agent to the other.

The random way in which the network is created guarantees that the average distance between any two agents is very short, significantly shorter for example than in a regular grid. The preferential attachment leads to a very skewed distribution of peers per agent, with some agents having several times the median connectivity. These well connected agents act as information hubs and dominate the flow of information. A high clustering coefficient implies that if two agents are peers of the same agent, the probability that they are also peers of each other is significantly higher than the probability that two randomly chosen agents are peers. This leads to the emergence of blocks within the social network that exhibit a high level of local interconnectivity.

¹In their seminal paper Watts and Strogatz (1998) define the clustering coefficient \mathcal{C}_i of a node i as the number of all direct links between the immediate neighbors of i divided by the maximum number of links that could possibly exist between them. They define the clustering coefficient \mathcal{C} of the entire network as the average clustering coefficient of the nodes of the network.

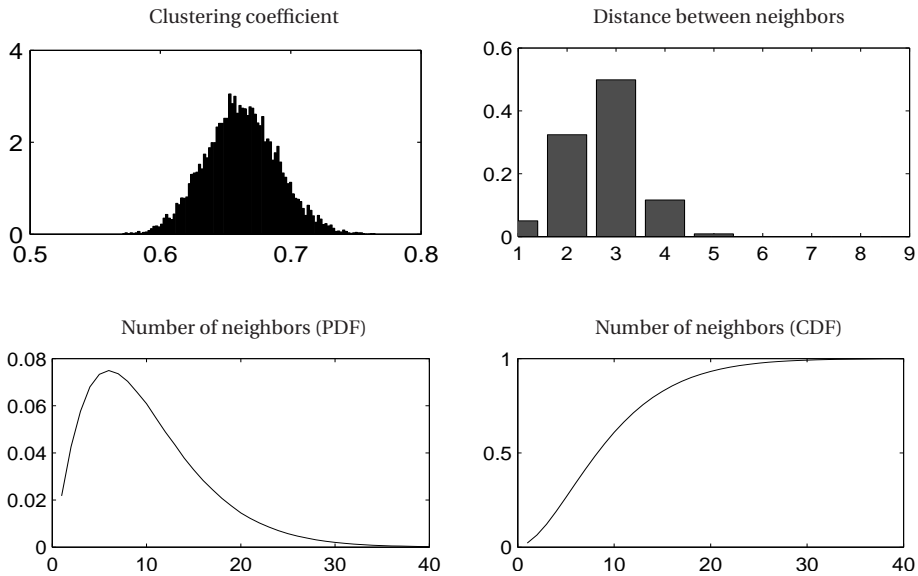


Figure 4.1: Network statistics

Figure 4.1 shows some key statistics collected from 10,000 networks of 200 agents that were created with an average connectivity of $k = 10$. The graphs show a normalized histogram of the clustering coefficient of each network (average .66), a normalized histogram of the distance between any two agents in each network, the probability density function (PDF), and the cumulative density function (CDF) of the number of neighbors per agent in each network. Note the relatively high probability of having 20 or more neighbors when the average connectivity is 10 neighbors. Such significant numbers of highly connected agents do not exist in regular grid networks or random networks of the Erdős-Rényi type, yet their existence in real social networks is well established. They generally act as information or transportation hubs and accelerate the dissemination of goods, viruses and ideas.

4.2.4 The evolutionary mechanism

It is important to note that in this EA agents and strategies are not the same. An agent carries or maintains a strategy, but it can change its strategy and we still consider it as the same agent. This dichotomy is necessary so that we can maintain a social network among the agents, while evolving, i.e., changing, the strategies. Because every agent has exactly one strategy at a time, the active number of active strategies is constant and equals the number of agents.

The first step in determining the selection probabilities is to rank all agents and their peers according to their respective welfare as measured by consumption. Let N_a denote the peers of agent a . The normalized rank $r_a(b) \in (0, 1]$ is the position of agent b among

the group consistent of a and the peers of a , divided by the size of this group,

$$r_a(b) = \frac{|\{c \mid c \in N_a \cup \{a\}, C_c(t) \leq C_b(t)\}|}{|N_a| + 1}. \quad (4.9)$$

If $C_c(t) = C_b(t)$, agents b and c are assigned the same rank. The best agent of a group of peers always has rank 1 while the worst one has rank $(|N_a| + 1)^{-1}$, which is close to zero. The special case $r_a(a)$ is important as it describes how an agent perceives its own economic performance relative to that of its peers. Note that this value does not need to be distributed uniformly over $(0, 1]$ —the fact that the size of N_a is different for different agents leads to a bell shaped distribution, which is skewed when there is correlation between welfare and the size of N_a .

We introduce two probabilistic selection mechanisms, one to decide whether a given strategy will be changed by mutation and one to decide whether it will be changed by selective imitation of a peer in the social network. In terms of traditional evolutionary computing (Eiben and Smith, 2003), imitation corresponds to recombination. However, there is an important difference between imitation as used here and usual recombination in traditional evolutionary computing, where the two recombinants have a symmetrical role: both receive (genetic) information from each other and incorporate it into the offspring. In our imitation mechanism the roles are asymmetrical. One agent imitates the other by receiving its strategy and recombining it with its own. The imitating agent changes its strategy, while the strategy of the imitated agent does not change.

Reflecting our best knowledge and intuition on social systems, we assume that these selection mechanisms depend on relative welfare. They should work differently for agents that have high $r_a(a)$ perceive themselves as rich and for agents that have a low $r_a(a)$ and perceive themselves as poor. We define two different sets of parameters for the selection mechanisms, one for agents with a high self-perception, which we mark with a subscript r for rich, and one for agents with a low self-perception, which we mark with a subscript p for poor. We also introduce two threshold parameters ρ_f and ρ_g to specify whether an agent perceives itself as rich relative to its peers. If $r_a(a) > \rho_f$, an agent perceives itself as rich with regard to mutation. If $r_a(a) > \rho_g$, an agent perceives itself as rich with regard to imitation.

Mutation in our simulation is implemented by Gaussian mutation. That is, an agent mutates its strategy vector by adding a random value drawn from a Gaussian distribution with zero mean. This implies that small mutations are more likely than large ones. The parameters f_p and f_r (for poor and rich agents respectively) specify the probability that an agent will mutate its strategy at each time step of the simulation,

$$P[a \text{ mutates its strategy}] = \begin{cases} f_p & \text{if } r_a(a) \leq \rho_f, \\ f_r & \text{if } r_a(a) > \rho_f. \end{cases} \quad (4.10)$$

The parameters σ_p and σ_r specify the standard deviation of the random value that is added to a mutated strategy. The exact formula for changing the strategy vector $s(t)$ into $s'(t+1)$ is

$$s'_a(t+1) = s_a(t) + \begin{cases} \mathcal{N}(0, \sigma_p) & \text{if } r_a(a) \leq \rho_f, \\ \mathcal{N}(0, \sigma_r) & \text{if } r_a(a) > \rho_f, \end{cases} \quad (4.11)$$

where $\mathcal{N}(0, \sigma)$ denotes a normally distributed random vector with zero mean and standard deviation σ . In order to avoid negative investments we add the additional constraint that $s_a(t) + \mathcal{N}(0, \sigma)$ is non-negative.

Imitation is performed by combining two strategies through linear combination. The resulting vector replaces the strategy of the imitating agent, while the strategy of the imitated agent remains the same. The parameters g_p and g_r specify the probability that an agent will imitate at each time step of the simulation,

$$P[a \text{ imitates}] = \begin{cases} g_p & \text{if } r_a(a) \leq \rho_g, \\ g_r & \text{if } r_a(a) > \rho_g. \end{cases} \quad (4.12)$$

In the event that agent a does imitate it needs to choose one of its richer peers to imitate. The parameters h_r and h_p specify the fraction of rich peers from which the agent chooses a random peer to imitate. That is, a poor agent with $r_a(a) \leq \rho_g$ chooses an agent to imitate according to

$$P[a \text{ imitates } b] = \begin{cases} 0 & \text{if } r_a(b) \leq h_p, \\ \left[(1 - h_p) \times |N_a| \right]^{-1} & \text{if } r_a(b) > h_p, \end{cases} \quad (4.13)$$

and a rich agent with $r_a(a) > \rho_g$ chooses an agent to imitate according to

$$P[a \text{ imitates } b] = \begin{cases} 0 & \text{if } r_a(b) \leq h_r, \\ \left[(1 - h_r) \times |N_a| \right]^{-1} & \text{if } r_a(b) > h_r. \end{cases} \quad (4.14)$$

If a imitates b , then the strategy $s_a(t)$ is linearly combined with $s_b(t)$ into $s'_a(t+1)$ according to

$$s'_a(t+1) = \begin{cases} (1 - w_p) s_a(t) + w_p s_b(t) & \text{if } r_a(a) \leq \rho_g, \\ (1 - w_r) s_a(t) + w_r s_b(t) & \text{if } r_a(a) > \rho_g, \end{cases} \quad (4.15)$$

where w_p and w_r is the weight that is given to the imitated strategy by poor and rich agents respectively. Since the investment fractions are constrained to sum to one, the resulting strategy is normalized,

$$s_a(t+1) = \frac{s'_a(t+1)}{|s'_a(t+1)|}. \quad (4.16)$$

The average connectivity k is the only free parameter of the social network and we will test values of k between 2 and 30. The resulting 13 parameters (1 parameter for average connectivity, 5 for mutation, and 7 for imitation) are shown in the first two columns of Table 4.1. On top of the “traditional” task of finding good values for these parameters we want to know if they are 1) indeed relevant for the evolutionary algorithm and 2) sufficient to tune the system. Here we call a parameter relevant if the aggregate agent welfare depends on the correct tuning of the parameter. Irrelevant parameters should be removed from the model for the sake of analytic clarity and computational stability.

Table 4.1: The 13 parameters of the initial evolutionary model

	Parameter	Shannon entropy	Standard deviation
k	average connectivity	-0.1	0.2
ρ_f	threshold rank for mutation	-0.3	0.4
f_p	P [poor agent mutates its strategy]	-1.0	0.7
f_r	P [rich agent mutates its strategy]	-3.9	1.4
σ_p	mutation variance of poor agent	-1.8	1.4
σ_r	mutation variance of rich agent	-2.2	1.4
ρ_g	threshold rank for imitation	-0.1	0.1
g_p	P [poor agent imitates]	-0.5	0.3
g_r	P [rich agent imitates]	-0.3	0.3
w_p	imitation weight of poor agent	-0.3	0.3
w_r	imitation weight of rich agent	-0.2	0.2
h_p	imitated neighbors of poor agent	-0.2	0.2
h_r	imitated neighbors of rich agent	-0.6	0.7

4.3 Experiments

4.3.1 Evaluating the initial evolutionary model

In the present application we want the economic agents to achieve a high aggregate welfare in a broad set of simulated economic environments. We add three scaling parameters to the agent-based application described in Section 4.2 to define a total of 18 different economic environments. They are

- The number of agents (200 or 2,000).
- The number of investment sectors with nonviable renewable energy technologies (2, 20, or 200). In all cases the simulation has exactly one investment sector with a viable technology and the number of nonviable technologies controls the difficulty of finding this viable technology.
- The vulnerability (“low”, “moderate”, or “high”) of the agent economies to climatic change. Exactly one investment sector leads to climatic change and the agents have to avoid investing in this sector.

Starting with the initial parameter set of 13 parameters we search for a distribution over parameter values with a good tradeoff in aggregate welfare and tuning cost for each of the 18 environments. As REVAC updates the marginal distributions over the parameter values, the expected aggregate welfare increases and the Shannon entropy of the marginal distributions decreases almost monotonically. Figure 4.2 illustrates this with graphs from three experiments with 200 agents, low vulnerability and three different numbers of nonviable technologies. The increase in aggregate welfare is greatest at the beginning of a REVAC tuning session, then slows down and comes to a halt after evaluating between 160 and 190 parameter vectors. On the other hand, the Shannon entropy




Figure 4.2: Aggregate welfare and Shannon entropy of the initial model of 13 parameters during a REVAC tuning session.

of the joint distribution that REVAC has estimated decreases linearly until about 200 parameter vectors are evaluated. In most simulated environments it continues to decrease even after that. Visual inspection of the results of all experiments leads us to conclude that the tradeoff between aggregate welfare and Shannon entropy is best after evaluating between 175 and 185 parameter vectors.

Table 4.1 shows the average Shannon entropy per parameter in bits together with the standard deviation. The results are averaged over all 18 simulated environments and over the marginal distributions obtained after evaluating 175–185 parameter vectors. Only 4 parameters show a Shannon entropy of 1 bit or more, which means that the performance of the evolutionary algorithm depends heavily on the correct tuning of these parameters. On the other hand, tuning of the other parameters seems largely irrelevant to aggregate welfare and their number should be reduced. The 4 relevant parameters define the probabilities to mutate a strategy (f_p , f_r) for poor and rich agents and the mutation variance (σ_p , σ_r) for poor and rich agents. REVAC tunes these pairs of parameters to similar values (not shown in the table) and we concluded that they can be combined into one parameter each. These results thoroughly falsify our original hypothesis that agent behavior should depend on relative welfare and that it needs to be tuned by different sets of parameters.

4.3.2 Evaluating a simplified evolutionary model

To verify these conclusions we simplify the evolutionary model by removing all behavioral differences between agents that perceive themselves as rich and agents that perceive themselves as poor. This leaves us with the six parameters shown in Table 4.2: connectivity k , probability to mutate f , mutation variance σ , probability to imitate g ,

Table 4.2: The 6 parameters of the simplified evolutionary model

Parameter	Shannon entropy	Standard deviation	25 th and 75 th percentiles
k average connectivity	-0.1	0.2	5.8–18.2
f $P[a$ mutates its strategy]	-3.2	1.2	0.01–0.07
σ mutation variance	-3.3	1.3	0.02–0.07
g $P[a$ imitates]	-0.5	0.4	0.54–0.88
w imitation weight	-0.3	0.3	0.41–0.88
h threshold rank imitated	-1.0	0.6	0.69–0.93

Notes. Results are averaged over the marginal distributions obtained after evaluating 175–185 parameter vectors. Initial parameter ranges are 0–1, except for connectivity, which has 2–30.

imitation weight w and threshold rank h of rich agents that are considered for imitation.

Figure 4.3 shows the aggregate welfare and Shannon entropy of the simplified model during tuning. Comparing the data visually we find that the best payoff between welfare and Shannon entropy is again achieved after evaluating about 180 parameter vectors. For each parameter of the simplified model Table 4.2 shows the average Shannon entropy of each marginal distribution, its standard deviation, and the 25th and 75th percentile of the marginal distribution at this point in the tuning process—results are averaged over all tested environments and over the marginal distributions obtained after evaluating 175–185 parameter vectors. As with the initial set of parameters, mutation

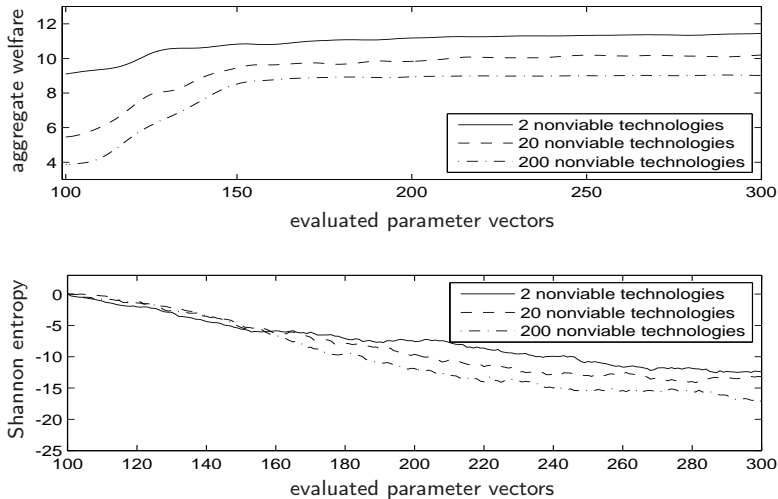


Figure 4.3: Aggregate welfare and Shannon entropy of the simplified model of 6 parameters during a REVAC tuning session.




Figure 4.4: Aggregate welfare as a function of Shannon entropy.

proves to be the most sensitive part of the simplified evolutionary model. Both the probability to mutate a strategy and the mutation variance have to be well tuned in order to allow the agents to adapt to their environment. The tuning of h (the fraction of richer neighbors to be imitated) also shows significant impact on aggregate welfare. The imitation parameters and the average number of neighbors prove to be almost irrelevant.

Figure 4.5 shows how REVAC tunes the parameters of the simplified model, averaged over the 18 simulated environments. Differences in the optimal parameter values (not shown here) are small between environments. For example, an increase in the number of nonviable technologies leads to lower values for the mutation parameters, apparently because mutation becomes riskier. However, such differences show consistently only in a later stage of a REVAC tuning session, typically after evaluating more than 200 parameter vectors. REVAC achieves the best tradeoff in expected aggregate welfare and Shannon entropy after evaluating 180 parameter vectors, and at that point the optimal parameter values are similar for all tested environments. This means that at least for the simulated environments discussed here REVAC tunes the parameters in a stable and consistent way and maximizes the performance of the evolutionary algorithm without compromising general validity.

After tuning the simplified evolutionary model to all 18 simulated economic environments we find that with equal cost of tuning the simplified model consistently achieves a higher aggregate welfare. To illustrate this, Figure 4.4 plots the aggregate welfare against the Shannon entropy that REVAC has estimated for the two evolutionary models. The x -axis shows the Shannon entropy of the joint distribution as it decreases during a REVAC tuning session. The y -axis shows the average aggregate welfare of the simulation when parameter values are drawn from REVAC distributions with the corresponding Shannon entropy. The graphs are based on simulations with 200 agents, 2 nonviable renewable energy technologies and low vulnerability to climatic change. The right cut off of each line marks the Shannon entropy after evaluating 300 parameter vectors. While the aggregate welfare at these cut off points is comparable for the two models, the 13-parameter model needs a significantly larger amount of information to reach it. In general, for each level of Shannon entropy, the aggregate welfare of the simplified 6-parameter model exceeds that of the initial 13-parameter model by about 10%. We conclude that the tradeoff in aggregate welfare to tuning cost is better with the simplified set of evolutionary parameters.




Figure 4.5: Tuning the 6 parameters of the simplified model. 50% of the density is located between the two solid lines, 80% between the dashed lines.

4.4 Conclusions

We illustrated how REVAC can support modeling activities: by showing the experimenter which model details can be considered as irrelevant—at least for the purpose of increasing a particular performance indicator like aggregate welfare. The Shannon entropy of the marginal distributions optimized by REVAC provides us with a useful measure of tuning cost that is independent of the actual tuning method. When comparing the 13-parameter model with the 6-parameter model we find that with equal tuning cost the 6-parameter model consistently outperforms the 13-parameter model by a significant amount of aggregate welfare. We conclude that there is no evidence that agents should condition their evolutionary behavior on relative welfare. Regarding individual parameters we find that the mutation parameters are the most relevant parameters in all versions of the evolutionary model in the sense that tuning them has the biggest effect on aggregate welfare. Tuning the fraction of peers that can be imitated is also important. The details of the social network, in particular the average connectivity, seem to be irrelevant, but warrant further research.

IMPACT OF ENVIRONMENTAL DYNAMICS ON ECONOMIC EVOLUTION: UNCERTAINTY, RISK AVERSION, AND POLICY

Abstract

The general question of how environmental dynamics affect the behavioral interaction in an evolutionary economy is considered. To this end, a basic model of a dynamic multi-sector economy is developed where the evolution of investment strategies depends on the diversity of investment strategies, social connectivity and relative contribution of sector specific investments to production. Four types of environmental dynamics are examined that differ in how gradual and how frequent environmental change occurs. Numerical analysis shows how the socially optimal level of diversity increases with the frequency and rapidity of the changes. When there is uncertainty about which type of environmental dynamics will prevail, the socially optimal level of diversity increases with the degree of risk aversion of the policy maker or the society.

5.1 Introduction

Evolutionary reasoning and agent-based modeling are standard practice in various disciplines, including social sciences (e.g., Binmore, 1994; Galor and Moav, 2002; Tesfatsion, 2006; Mirowski, 2007). A typical evolutionary model uses a population of entities that undergo selection and variation. Although specific domains ask for the development of particular types of model, several common, general questions arise. Here we

A revised version has been published as Nannen et al. (2013).

aim to address one such question, namely how does a dynamic environment influence the behavior of an evolutionary economic system consisting of multiple agents employing different behavioral strategies. The relevance of this question is evident: few economic environments are static.

Generally, one cannot expect evolution in a changing environment to approach a steady state. What matters is not so much how well the agents adapt if given enough time, but how fast they adapt to a new challenge. In a socio-economic context a wide range of environmental variables can be identified: macroeconomic conditions, technological opportunities, policies and institutions, and natural resources. Most studies of social behavior through evolutionary methods have been limited to constant environments, letting selection pressure depend on the population distribution. In part, this allows for analytical treatments, as has been the common approach in evolutionary game theory (both in biology and the social sciences). The addition of a dynamic environment requires a numerical or computational approach. As environmental economics deals with the economic analysis of exploitation of natural resources, abatement of environmental pollution, and human-induced climate change, dynamic environments are prevalent. The evolution of strategies is important when heterogeneous groups of users, polluters, or harvesting strategies are involved (Ostrom, 2000; van den Bergh, 2007). Dynamic environments may cause certain strategies to become evolutionary stable and others to become unstable. We will not only draw upon the social sciences but also make use of certain insights from evolutionary biology. Evidently, many explicit and implicit insights on the influence of environment on evolution are available here.

For our purpose a relevant distinction is between exogenous and endogenous environments. Whereas systems with only exogenous variables are relatively simple, endogenous variables generate complex feedback systems. Unfortunately, most real-world systems studied by biologists and social scientists are of the latter type. Resource dynamics (e.g., Sethi and Somanathan, 1996; Noailly et al., 2003) and dynamic control of a pest population that evolves resistance to pesticides (Munro, 1997) are policy-relevant examples. Another, general example is a coevolutionary system in which two heterogeneous populations cause selection pressure on one another (Epstein and Axtell, 1996). This leads to very complex coevolutionary interactions because the environment of each evolutionary (sub)system is evolving as well. Coevolution thus implies a particular type of dynamic and endogenous environment (Noailly, 2008).

With regard to the evolutionary system, there is a range of theoretical starting points and modeling approaches (Eiben and Smith, 2003; van den Bergh, 2004a). First of all, one can choose to use very theoretical, abstract models of the evolutionary game type. However, adding dynamic environments here will lead to systems that are no longer amenable to analytic solutions. Numeric simulations of multi-agent systems form an alternative to the analytic approach that offer much more flexibility in examining system behavior. They allow a distinction between local and global environments, and between stationary and mobile agents. They further allow to study the influence of population size, and the effects of dynamic environments on group and network formation (Bergstrom, 2002; Henrich, 2004). In addition, different assumptions can be made regarding selection factors and innovation mechanisms (random mutations, deterministic trends, recombination) and bounded rationality of agents (habits, imitation).

In this chapter we investigate the impact of various types of general environmental dynamics on the socially optimal type of behavioral interactions among the agents in the population. We consider the general structure of a dynamic non-aggregate multi-sector economy. Agents have individual investment strategies that specify how they invest their respective income. Their objective is to maximize their individual welfare. They prefer investment strategies which give high welfare. Their rational capabilities are bounded and their information is limited. The only information available to the agents is the investment strategies and the welfare of their fellow agents. The behavioral interactions influence how the agents use this information to evolve their investment strategies through imitation. Our framework postulates that the environmental dynamics are beyond the control of the policy maker, while he or she can regulate (some aspects of) the agent interactions. Various types of government regulation, information and education affect the search for and effectiveness of innovation by economic agents. In particular, by regulating how accurately agents can imitate each other, a policy maker can control the diversity of strategies within the population. Examples of policies that influence diversity are patent and copyright laws, conditions for competition for public R&D funds and subsidies, and the support or enforcement of industry standards.

We will study the effect of diversity on welfare numerically through computer simulations. We will address two research questions. The first is whether it is true that different environmental dynamics require different degrees of diversity for the agents to achieve a high welfare. The second question follows from the fact that environmental dynamics are not only beyond the control of the policy maker, but that they are also uncertain to him. This raises the issue of adequate policies under uncertainty: how do agent interactions that work well for one type of environmental dynamics perform under another environmental dynamics? Depending on the degree of risk aversion of the policy maker or the society, different policies can be recommended.

As for the environmental dynamics, we focus on two general aspects of environmental change: how gradually it occurs, and how often. Gradualness and frequency of change are two aspects of an environmental dynamics that can relatively easily be observed and recognized. Depletion of a mineral resource, for example, typically manifests itself over an extended period of time, while a biotic resource like fish can disappear literally overnight. Or a remote agricultural community is normally exposed to environmental hazards less frequently than one surrounded by a heavily industrialized region. If a policy maker can anticipate these aspects of environmental change, he or she might want to steer behavioral interaction such that economic agents can adapt well.

The remainder of this chapter is organized as follows. Section 5.2 describes production and growth in an economy with a very general structure and presents the evolutionary mechanism of behavioral interaction. In Section 5.3 the relation between an investment strategy and the income growth rate is studied. Section 5.4 describes the experimental setup. Section 5.5 provides simulation results and interpretations. Section 5.6 concludes.

5.2 The economic model

5.2.1 General features of the model

Consider a population of agents with the objective to reach a high level of individual welfare, which can only be achieved by a sustained high income growth rate. Each agent can invest its respective income in a finite number of capital sectors. How it allocates its investment over these sectors is expressed by its individual investment strategy. Invested capital is non-malleable: once invested it cannot be transferred between sectors. Standard economic growth and production functions describe how the invested capital accumulates in each sector and contributes to income. These functions are not aggregated: growth and returns are calculated independently for each agent. Two agents with different investment strategies can experience different income growth rates and income levels.

The agents understand that there is a causal link between an investment strategy and economic performance as expressed by the income growth rate, but they cannot use calculus to find an investment strategy that maximizes the income growth rate. Instead, the agents employ the smartest search method that nature has in store, evolution, and they evolve their investment strategies by imitation with variation. Since they prefer a high income growth rate over a low income growth rate, they imitate the investment strategy of a fellow agent when that fellow agent realizes an income growth rate that is high relative to their own income growth rate and that of their other fellow agents. Imitation is not perfect. Changes that are introduced during imitation guarantee diversity in the pool of strategies and keep the evolutionary search alive. In the terminology of evolutionary theory an agent *selects* another agent based on a property (the *phenotype*) that is indicative of its current economic performance and imitates its investment strategy (the *genotype*) with *variation*.

5.2.2 Strategies, investment, and production

All variables and parameters of the economic model are summarized in Table 5.1. The population approach means that accounting of capital investment, production, and income takes place at the level of individual agents. Let $Y_a(t)$ be the income of agent a at time t and let n be the number of available investment sectors. Formally, the investment strategy $s_a(t)$ of an agent can be defined as an n -dimensional vector

$$s_a(t) = [0, 1]^n, \quad \sum_i s_{ia}(t) = 1. \quad (5.1)$$

The partial strategy $s_{ia}(t)$ —which is the i^{th} element of a strategy—determines the fraction $s_{ia}(t)Y_a(t-1)$ of income that agent a invests in sector i at time t . Each agent must invest its total income in one sector or another, so the partial strategies must be non-negative and sum to one. The set of all possible investment strategies is an $n-1$ dimensional simplex that is embedded in n -dimensional Euclidean space. We call this simplex the *strategy space*.

Capital accumulation in each sector i depends on the sector specific investment of each agent and on the global depreciation rate δ . Depreciation is assumed to be equal for

Table 5.1: Variables and parameters of the model

$ P $	population size
k	average number of neighbors per agent
N	neighbors of an agent
\mathcal{C}	clustering coefficient of the network
n	number of investment sectors
β	scaling factor of production
δ	discount rate
σ	diversity control parameter
K_{ia}	capital that agent a has accumulated in investment sector i
π_i	production coefficient of investment sector i
s_{ia}	fraction of income that agent a allocates to investment sector i
Y_a	net domestic income of agent a
γ_a	income growth rate of agent a

all sectors and all agents. The dynamic equation for non-aggregate growth per sector is

$$K_{ia}(t) = s_{ia}(t)Y_a(t-1) + (1-\delta)K_{ia}(t-1). \quad (5.2)$$

An extended version of this equation that accounts for dynamic prices can be found in the appendix. To calculate the income $Y_a(t)$ from the capital that agent a has accumulated per sector, we use an n -factor Cobb-Douglas production function with a constant elasticity of substitution,

$$Y_a(t) = \beta \prod_i K_{ia}(t)^{\pi_i(t)}, \quad (5.3)$$

where β is a scaling factor that limits the maximum possible income growth rate. The relative contribution of each sector to production is expressed by a dynamic vector of non-negative production coefficients $\pi(t) = \langle \pi_1(t) \dots \pi_n(t) \rangle$. To enforce constant returns to scale, all production coefficients are constraint to add up to one,

$$\pi(t) = [0, 1]^n, \quad \sum_i \pi_i(t) = 1. \quad (5.4)$$

Similar to the strategy space, the set of all possible vectors of production coefficients is an $n-1$ dimensional simplex that is embedded in n -dimensional Euclidean space.

Production coefficients can depend on an array of economic dynamics, like technological development and environmental dynamics. When the technology or the environment changes, the production coefficients can change as well. Progressive desertification of farm land for example increases the dependency of farmers on irrigation. This can be interpreted as an increase of the production coefficient of irrigation, while some or all of the other production coefficients of the agricultural production process would decrease to compensate. Evolutionary economics raises the question of what happens if the production coefficients change. For this reason we model the environmental dynamics as exogenously defined changes in $\pi(t)$. This is a general approach that can also be applied to other economic dynamics such as technological development. Section 5.4

describes how we implement these changes and how we test whether or not they have an impact on behavioral interactions.

To measure how well a population of agents is adapted to a certain economic environment, we use the expected log income $E[\log Y(t)]$ of all agents at time t . Expected log income emphasizes an egalitarian distribution of income. Technically speaking, an economic agent with constant relative risk aversion prefers a society with high expected log growth and high expected log income. Let $|P|$ be the total number of agents in the population P . We calculate expected log growth as

$$E[\log Y(t)] = \frac{1}{|P|} \sum_a \log Y_a(t). \quad (5.5)$$

The individual income growth rate $\gamma_a(t)$ is

$$\gamma_a(t) = \frac{Y_a(t)}{Y_a(t-1)} - 1. \quad (5.6)$$

Expected log income relates to expected log income growth as

$$E[\log Y(t)] = \sum_{i=1}^t E[\log(\gamma(i) + 1)] + E[\log Y(0)], \quad (5.7)$$

where expected log income growth is defined as

$$E[\log(\gamma(t) + 1)] = \frac{1}{|P|} \sum_a \log(\gamma_a(t) + 1). \quad (5.8)$$

5.2.3 The evolutionary mechanism of behavioral interactions

From the point of view of evolutionary modeling, agents and investment strategies are not the same: an agent carries or maintains a strategy, but it can change its strategy and we still consider it to be the same agent (Nowak, 2006). Because every agent has exactly one strategy at a time, the number of active strategies is the same as the number of agents.

To model which agents an agent can imitate we use a generic class of social networks that has been well studied and validated in network theory, namely those that can be generated by a random process with preferential attachment and that have a high clustering coefficient, see Section 4.2.3 on page 71 for details. Before the start of each simulation a stochastic process assigns to each agent a a set of peers N_a that does not change during the course of the simulation. If agent a is a peer of agent b , then a will consider the income growth rate and the investment strategy of b when choosing an agent for imitation, while b will consider the income growth rate and the investment strategy of a . On the other hand, if a and b are not peers, they will not consider each other for the purpose of imitation.

At each time step t an agent may select one of its peers in the social network and imitate its strategy. If that happens, the strategy of the imitating agent changes, while the strategy of the agent that is imitated does not. The choice of which agent to imitate is based on relative welfare as indicated by the current growth rate of income. The imitating agent always selects the peer with the highest current income growth rate. Only

if an agent has no peer with an income growth rate higher than its own, the agent does not revise its strategy.

If imitation was the only mechanism by which agents change their strategies, the strategies of agents that form a connected network must converge on a strategy that was present during the initial setup. However, real imitation is never without errors. Errors are called mutations in evolutionary theory. They are fundamental to an evolutionary process because they create and maintain the diversity on which selection can work. In this model we implement mutation by adding some Gaussian noise to the imitation process. That is, when an agent imitates a strategy, it adds some random noise drawn from a Gaussian distribution with zero mean. This causes small mutations along each partial strategy to be more likely than large ones. The exact formula by which agent a imitates and then mutates the strategy of agent b is

$$s_a(t) = s_b(t-1) + \mathcal{N}(0, \sigma), \quad (5.9)$$

where $\mathcal{N}(0, \sigma)$ denotes a normally distributed n -dimensional random vector with zero mean and standard deviation σ per dimension. Because partial investment strategies have to sum to one, we have to enforce $\mathcal{N}(0, \sigma) = 0$, for example by orthogonal projection of the Gaussian noise term onto the simplex, resulting in the loss of one degree of freedom. The error term is further constraint to leave all partial strategies positive. Needless to say that we do not imply that our boundedly rational agents engage consciously in such mathematical exercise. Subjectively they merely allocate their income such that none is left.

The sum of squares of the n partial errors, i.e., the square of the Euclidean distance covered by the error, follows a chi-square distribution with $n-1$ degrees of freedom and mean $(n-1) * \sigma^2$. In equilibrium, when all agents try to imitate the same perfect strategy, the expected standard deviation of the partial strategies will in fact be σ . Since the parameter σ controls the diversity of the investment strategies, we will call it the diversity control parameter, or simply *diversity*. It is the only free parameter of this evolutionary mechanism and has potential policy implications.

5.3 The evolutionary dynamics

5.3.1 The growth rate of a strategy

If we want to understand the impact of environmental dynamics on how agents evolve their strategies we need to understand if and how these environmental dynamics affect which agents are imitated. Whether the strategy of an agent is imitated depends on whether the agent has a higher income growth rate than those agents it is compared with. We call the mapping from investment strategies to income growth rate the *growth function*. The growth function calculates the equilibrium growth rate that an imitating agent realizes if it holds on to a particular investment strategy. If the growth function maps one strategy to a higher equilibrium growth rate than another strategy, then our evolutionary agents will prefer this strategy over the other strategy and imitate it. In this way the growth function indicates which of any two strategies will survive and propagate. Since it depends only on the order of income growth rates—i.e., which of any

two agents has a higher income growth rate—whether an agent is imitated, the evolutionary dynamics is invariant under any strictly increasing transformation of the growth function. Any two growth functions that are strictly increasing (or decreasing) transformations of each other lead to the same evolutionary dynamics.

Let us start the derivation of the growth function with an analysis of the equilibrium ratio of sector specific capital to income, $K_{ia}(t)/Y_a(t)$, that will be achieved if an agent holds on to a particular strategy. The dynamic equation of this ratio is

$$\begin{aligned} \frac{K_{ia}(t)}{Y_a(t)} &= \frac{s_{ia}(t)Y_a(t-1) + (1-\delta)K_{ia}(t-1)}{(\gamma_a(t)+1)Y_a(t-1)} \\ &= \frac{s_{ia}(t)}{\gamma_a(t)+1} + \frac{1-\delta}{\gamma_a(t)+1} \frac{K_{ia}(t-1)}{Y_a(t-1)}. \end{aligned} \quad (5.10)$$

This equation is of the form

$$x(t) = a + bx(t-1), \quad (5.11)$$

which under the condition $0 \leq b < 1$ converges monotonically to its unique stable equilibrium at

$$\lim_{t \rightarrow \infty} x(t) = a/(1-b).$$

This condition is fulfilled here: investment is always non-negative and sector specific capital cannot decrease faster than δ . With constant returns to scale, income cannot decline faster than capital depreciation, and we have $\gamma_a \geq -\delta$. For the moment, let us exclude the special case $\gamma_a = -\delta$. Then, considering that $0 < \delta \leq 1$, we have the required constraint

$$0 \leq \frac{1-\delta}{\gamma_a(t)+1} < 1 \quad (5.12)$$

and we conclude that the ratio of capital to income converges to

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{K_{ia}(t)}{Y_a(t)} &= \lim_{t \rightarrow \infty} \frac{s_{ia}(t)}{\gamma_a(t)+1} / \left(1 - \frac{1-\delta}{\gamma_a(t)+1}\right) \\ &= \lim_{t \rightarrow \infty} \frac{s_{ia}(t)}{\gamma_a(t)+\delta}. \end{aligned} \quad (5.13)$$

Equation 5.13 describes a unique stable equilibrium to which the economy of an agent converges monotonically. We ignore the limit notation and combine equation 5.13 with equation 5.3 to calculate income at equilibrium as

$$\begin{aligned} Y_a(t) &= \beta \prod_i \left(\frac{s_{ia}(t) Y_a(t)}{\gamma_a(t)+\delta} \right)^{\pi_i(t)} \\ &= \beta \frac{Y_a(t)}{\gamma_a(t)+\delta} \prod_i s_{ia}(t)^{\pi_i(t)}. \end{aligned} \quad (5.14)$$

We can now solve for $\gamma_a(t)$ to derive the growth function

$$\gamma_a(t) = \beta \prod_i s_{ia}(t)^{\pi_i(t)} - \delta. \quad (5.15)$$

Let us return to the special case $\gamma_a = -\delta$. According to equation 5.2, capital per sector decreases at the depreciation rate δ only when it receives zero investment, and it cannot decrease faster. This implies that with constant elasticity of substitution, a growth of $\gamma_a = -\delta$ is only possible if every sector with a positive production coefficient receives zero investment. This implies $s_{ia}(t) = 0$ for at least one partial strategy, and so equation 5.15 holds also for the special case $\gamma_a = -\delta$.

5.3.2 Efficiency and level sets of investment strategies

How does the income growth rate of an imitating agent compare to the income growth rate of a rational agent with perfect information? The term $\prod_i s_{ia}(t)^{\pi_i(t)}$ has a single optimum at $s_a(t) = \pi(t)$, allowing a maximum growth of $\gamma^{opt}(t) = \beta \prod_i \pi_i(t)^{\pi_i(t)} - \delta$. This is the income growth rate that a rational agent with perfect information would expect to achieve. Its exact value depends on the location of the production coefficients in the simplex. In an n -factor economy the term $\prod_i \pi_i(t)^{\pi_i(t)}$ varies between a value of $1/n$ in the center of the simplex where all production coefficients are equal, and a value of one in the corners of the simplex where one sector dominates. In order to remove this variability from the growth function, and to allow an easy comparison with the income growth rate of a rational agent with perfect information, we define the efficiency $\mathcal{E}(s, t)$ of a strategy $s(t)$,

$$\mathcal{E}(s, t) = \prod_i \left(\frac{s_i(t)}{\pi_i(t)} \right)^{\pi_i(t)}. \tag{5.16}$$

The efficiency of a strategy measures the fraction $\gamma_a(t)/\gamma^{opt}(t)$ of optimal growth that an agent achieves with this strategy on given production coefficients, assuming that $\delta = 0$. If one strategy leads to a higher equilibrium growth rate than another, it is also more efficient. Efficiency is therefore a monotonic transformation of the growth function that preserves all information on which agent imitates which other agent, removes the variability due to the location of the optimum on the simplex, and allows us to measure growth in terms of what a rational agent with perfect information would achieve.

Efficiency, like the equilibrium growth rate, is a monotonically decreasing function of the Euclidean distance between the strategy and the production coefficients, $|s_a(t) - \pi(t)|$. This function is not symmetric about the optimum but has different slopes in different directions from the optimum. Figure 5.1 shows how the average efficiency of a




Figure 5.1: Average efficiency as a function of Euclidean distance to the optimum. The x -axis shows the Euclidean distance, the y -axis the corresponding average efficiency. Note the convex shape around the optima.




Figure 5.2: Probability of positive growth when the investment strategy and the production coefficients are chosen independently and at random. The x -axis shows the ratio δ/β . The y -axis shows the corresponding probability that the equilibrium growth rate is positive.

strategy decreases as its Euclidean distance to the optimum increases, when strategies and production coefficients are chosen at random from the simplex. Note the inverse S-shape of the graphs. As the Euclidean distance tends to zero, the gradient approaches zero. This implies that when an evolutionary population of agents converges on the optimum, the differences in growth caused by a small amount of diversity σ around the optimal strategy are negligible.

A set of strategies each with identical equilibrium growth rate, say γ' , is called a level set and forms a contour hypersurface in the strategy simplex. All strategies that are enveloped by this hypersurface have an equilibrium growth rate that is higher than γ' . This inner set is convex (for a related proof see Beer, 1980) and so from equation 5.15 satisfies

$$\prod_i s_{ia}(t)^{\pi_i(t)} \geq \frac{\gamma' + \delta}{\beta}. \quad (5.17)$$

An important level set is $\prod_i s_{ia}(t)^{\pi_i(t)} > \delta/\beta$. This is the set of all strategies that have a positive equilibrium growth rate. Its size is proportional to $P[\gamma > 0 \mid \pi(t)]$, the probability that a random strategy has a positive equilibrium growth rate with given production coefficients. Let $P[\gamma > 0]$ denote the probability that the equilibrium growth rate is positive if both the strategy and the production coefficients are chosen independently at random from the simplex. Figure 5.2 shows how $P[\gamma > 0]$ decreases as δ/β increases, for economies with respectively 2, 4, and 10 investment sectors. The probability tends to zero as δ/β approaches 1. For given δ/β , the probability that the equilibrium growth rate of a random strategy is positive decreases as the number of investment sectors increases.

The parameters δ and β determine the equilibrium growth rate associated with a given hypersurface, as well as the minimum and maximum equilibrium growth rate that can be achieved with given production coefficients. They do not affect the location of the optimum nor the shape of level sets, both of which depend exclusively on the production coefficients. In other words, δ and β define monotonic transformations of the growth function that are irrelevant to the order of equilibrium growth rates and to the understanding of the evolutionary dynamics as a whole. Also, the rate of convergence in equation 5.13 does not depend on the scaling factor β . We will make use of this fact later on in the experimental design where we use a dynamic β for normalization, significantly reducing the variability of the numeric results.

5.4 Experimental setup

5.4.1 The environmental dynamics

The growth function, up to a monotonic transformation, depends on the coefficients of a Cobb-Douglas type production function. When the production coefficients change with the environmental dynamics, strategies that have previously generated a positive income growth rate can now generate a negative income growth rate. Agents that have converged on a strategy that has previously resulted in a high income growth rate can see their income decline and need to adapt their strategies to the new production coefficients. How does the magnitude and duration of this decline depend on the type of environmental dynamics and on the behavioral interactions among the agents? To answer these questions we model the environmental dynamics as exogenously defined changes in $\pi(t)$. That is, the environmental dynamics that change the production coefficients are the independent variable that the policy maker responds to. The parameters of the imitation mechanism are the dependent variables that the policy maker aims to regulate.

We focus on two aspects of environmental dynamics: how gradual the environment changes, and how frequently. In combination they define four types of environmental dynamics: the production coefficients change gradually and with low frequency, gradually and with high frequency, suddenly and with low frequency, and suddenly and with high frequency. We compare these with two control systems: one without imitation and one with imitation and a static environment. Without imitation, with strategies that are randomly distributed over the strategy space and that stay constant throughout the simulation, the income growth rate of most agents is most likely negative, irrespective of the environmental dynamics. Expected log income will decline and welfare at the population level will be at its lowest. On the other hand, in a static environment where strategies evolve they are expected to converge on the optimum strategy and welfare at the population level will be at its highest.

We consider the general case where a change in the production coefficients is defined as the replacement of one vector of production coefficients by another, with each vector drawn independently and at random from the uniform distribution over the simplex $\sum_i \pi_i(t) = 1$. Replacement is instant for a sudden change and by linear transition for a slow change. A sudden change can be modeled by setting the production coefficients of a 2-factor economy to $\pi = \langle .1, .9 \rangle$ up until time t , and to $\pi = \langle .4, .6 \rangle$ from $t + 1$ onwards. Such extreme changes are characteristic of industries that depend on unreliable resources, e.g., a biotic resource susceptible to climate change like forests or fish. A gradual change can be modeled by changing π from $\langle .1, .9 \rangle$ at time t to $\langle .4, .6 \rangle$ at time $t + x$ linearly over x steps, such that

$$\pi(t + j) = \frac{(x - j)\pi(t) + j\pi(t + x)}{x}, \quad 0 \leq j \leq x, \quad (5.18)$$

where the conditions $\sum_i \pi_i = 1$ and $\pi_i \geq 0$ for all i are fulfilled at all times. Such gradual changes are characteristic of industries that depend on reliable resources, e.g., a mineral resource like iron or coal, where known reserves will typically last for decades if not centuries.

Table 5.2: The environmental dynamics

Environmental dynamics	Observable	Example
gradual, low frequency	reliable resource, Kondrat. wave	oil/gas reserves
sudden, low frequency	unreliable resource, Kondrat. wave	climate change
gradual, high frequency	reliable resource, Juglar's cycle	tech. innovations
sudden, high frequency	unreliable resource, Juglar's cycle	biotic resource

We model low frequency changes by starting the transition from one vector of production coefficients to another vector every 50 years, reflecting a Kondratiev type of wave (Kondratiev, 1935), characteristic of industries that are not a driving force of innovation and change only with the general shift in production methods, e.g., forestry. To model high frequency changes the transition starts every 10 years, corresponding to the fast business cycles observed by Clément Juglar (1863), characteristic of industries that invest heavily in research and development. That is, while we acknowledge that technological innovations are driven by research and development, we treat their effect on the production coefficients as exogenous environmental dynamics that the agents of an industry have to adapt to. We do not claim that the cycles observed by Kondratiev and Juglar are caused by this type of exogenous dynamics. We merely use their observations as examples of frequency patterns that can indeed be detected when present.

We consider different sequences of production coefficients as different instances of the same environmental dynamics as long as the individual vectors of production coefficients are replaced with the same gradualness and frequency. Table 5.2 summarizes the environmental dynamics used for the experiments. Figure 5.3 gives graphic examples of production coefficients that are drawn at random according to the specification of each environmental dynamics. Each row of this figure shows five graphs: one each for the time evolution of the four production coefficients of a 4-factor economy, and one area plot that combines the other four graphs into a single graph, stacking the four individual curves one on top of the other (the upper curve has constant value one), with a different shade of gray for the area under each curve.

With regard to the dependent variable under control of the policy maker, the imitation mechanism has one free parameter, diversity σ , and we specify the optimal behavioral interactions as the diversity $\sigma^{opt}(d)$ that maximizes the expected log income of each agent under given environmental dynamics d ,

$$\sigma^{opt}(d) = \underset{\sigma}{\operatorname{argmax}} \operatorname{E}(\log Y(t) | \sigma, d). \quad (5.19)$$

In order to find this optimal value for different environmental dynamics we use repeated numerical simulations with different values of σ and measure the expected log income at the end of each simulation, using standard statistical methods to reduce variance. Having identified the value $\sigma^{opt}(d)$ at which the expected log income is highest under a given environmental dynamics, we proceed to formulate policy advice on the socially optimal level of σ when there is uncertainty over the type of environmental dynamics. To do so we measure the expected log income that an optimal value $\sigma^{opt}(d)$ generates on those environmental dynamics $d' \neq d$ where it is not optimal. We then calculate the value that policy makers with different degrees of risk aversion assign to each $\sigma^{opt}(d)$.




Figure 5.3: Environmental dynamics: changes in the production coefficients. Each row illustrates a different type of environmental dynamics. The sequences of production coefficients are chosen at random. The x -axis shows the 500 time steps (initialization and main experimental phase). The y -axis of the four graphs positioned at the left of each row shows the value of one particular production coefficient in a 4-factor economy. The single graphs at the right are area plots that stack the values of the same four production coefficients one on top of the other, with a different shade of grey under each curve.

5.4.2 Implementation details, model calibration, and scaling

The numerical simulations are based on a discrete synchronous time model where the income and strategy of each agent is updated in parallel at fixed time intervals. We consider each time step t to simulate one financial quarter. As no significant financial market requires a publicly traded company to publish financial results more than 4 times a year, we consider it the limit of feasibility to account for growth and to review an economic strategy as often as 4 times a year. Most economic agents will alter their strategy less often. Each simulation step is divided into two separate update operations: *updating the economy*—each agent invests its income according to its own investment strategy and the individual incomes and growth are calculated by the non-aggregate growth and production and growth functions—and *updating the strategies*, when all agents compare their income growth rate with that of their peer group, and when those agents that decide to imitate change their respective strategies simultaneously.

Each computer simulation spans 500 time steps, simulating 500 financial quarters or 125 years. These are divided into an initialization phase of 100 time steps or 25 years, and a main experimental phase of 400 time steps or 100 years. An initialization phase is needed to avoid influencing the simulation results by an arbitrary choice of initial values. During initialization the simulated economy stabilizes and a “natural” distribution of strategies and growth emerges. Initial conditions are always defined in the same way: all strategies and the initial production coefficients are drawn independently at random from the simplex. The production coefficients are kept static throughout the initialization phase but the agents can imitate in the same way as they do during the main experimental phase, with the same σ . During the 400 time steps (100 years) of the main experimental phase the agents have to adapt to the dynamic changes in the production coefficients. To avoid any initialization effect, the increase in log income is measured from the beginning of the main experimental phase.

Numerical methods are inherently constraint by the availability of computational resources. The computational complexity of multi-agent systems typically scales at least polynomially with system size. The accepted method is to extensively study a system that is large enough to incorporate all the essential ingredients of the model, and to only increase the system size to test whether the obtained results are scalable. Here the main experiments are based on an economy of 200 agents and a 4-factor economy. Sensitivity and scalability are tested with 1,000 agents and with a 10-factor economy. To understand whether the results depend on the specific implementation of the evolutionary mechanism we also test more sophisticated implementations: one version where each agent imitates with probability .1 at every step—as opposed to probability one in the main experiment—and one version where imitation is partial, such that a new strategy is a linear combination of the imitated strategy (with weight .1) and the strategy of the imitating agent (with weight .9). As before, σ controls the standard deviation of the normally distributed errors per partial strategy.

Recall that the rate of capital depreciation δ (equation 5.2) and the scaling factor β (equation 5.3) of the production function have no effect on the evolutionary dynamics and the adaptive behavior of the agents. For the present model we set $\delta = .01$ per time step—about 4% per year—for all sectors. We use a dynamic β to reduce variability of the numeric results. As seen in Section 5.3, different vectors of production coefficients

Table 5.3: Values of economic parameters

population size $ P $	200
investment sectors n	4
duration of the initialization phase	100 time steps
duration of the main experimental phase	400 time steps
capital depreciation δ	.01
dynamically normalized scaling factor $\beta(t)$	$.015 \prod_i \pi_i(t)^{-\pi_i(t)}$
average network connectivity	10

have different optimal equilibrium growth rates. To study how efficient the agents adapt to different vectors of production coefficients, we correct for this variability in optimal growth by dynamically normalizing the scaling factor β . To keep the optimal income growth rate at a value of .005 (i.e., an income growth rate of about 2% per year), we let the scaling factor $\beta(t)$ depend on the vector of production coefficients,

$$\beta(t) = .015 \prod_i \pi_i(t)^{-\pi_i(t)}. \quad (5.20)$$

With this normalization the equilibrium growth rate of all strategies is constraint to the range $[-.01, .005]$, where the minimum of $-.01$ is realized when $s_{ia}(t) = 0$ for some positive π_i and where the maximum of .005 is realized when $s_a(t) = \pi(t)$. Numerical tests show that with these parameter values the probability that a random strategy has a negative equilibrium growth rate on random production coefficients is about .65.

To model which agents an agent can imitate we use a generic class of social networks that has been well studied and validated in network theory, namely those that can be generated by a random process with preferential attachment and that have a high clustering coefficient, see Section 4.2.3 on page 71 for details. Here we use an average connectivity of $k = 10$. In a population of 200 agents this value results in a highly connected network—the average distance between any two agents in the network is 2.7—while maintaining the overall qualities of a complex network.

To improve the general validity of our results we use large number of numerical simulations where—rather than closely calibrating those factors that affect the evolutionary dynamics on a specific economy—we define broad parameter ranges and collect statistical information over a representative sample of different possible economies that fall within these ranges. For example, in order to obtain results that are valid for the general class of scale-free social networks with a high cluster coefficient, each simulation is based on an independent random instance of the social network. Likewise, in order to obtain general results for specific environmental dynamics, each simulation uses an independent random sequences of production coefficients, which are replaced according to the gradualness and frequency of the respective environmental dynamics. The number of simulations needed to obtain reliable statistical results is determined by standard methods of variance reduction. The values of all economic parameters are listed in Table 5.3.

5.5 Results

5.5.1 Economic significance of diversity

Figure 5.4 shows the expected log income of an agent for different values of the diversity control parameter σ , for the two control systems and the four types of environmental dynamics. 50,000 simulations are used for each graph. 500 different values of σ from the range $[0, .5]$ are evaluated, and results are averaged over 100 simulations per value. The plots are smoothed with a moving average with a window size of ten values.

In the first graph of Figure 5.4—the control system without imitation—expected log income is uniformly negative for all levels of σ . This graph is based on a static environment, but the same is observed for any environmental dynamics. All other graphs of Figure 5.4 show systems with imitation and there is a clear functional relation between the value of σ and log income. For each system there is a single optimum $\sigma^{opt}(d)$ that maximizes log income under the given environmental dynamics d , exact values are given in Table 5.4. The value of $\sigma^{opt}(d)$ is higher for more frequent changes than for less frequent changes, and higher for sudden changes than for gradual changes. Its value is lowest in the static environment. Further to this, the graphs show a clear pattern in the relationship between σ and expected log income: the slope to the left of the optima, i.e., for small values of σ , is much steeper than to the right, where σ is large. We will revisit this fact in our discussion of policy advice under uncertainty.

Our first research question can now be answered: almost any level of diversity σ will allow the evolutionary agents to reach a positive log income under any environmental dynamics, yet a unique optimum where log income is highest can be identified for each environmental dynamics. So while it is not mandatory to define policies that effect σ , in the sense that imitating agents can almost always return to positive growth, it is optimal in the sense that there can be a significant gain in expected log income.




Figure 5.4: Expected log income as a function of the diversity parameter σ . The x -axes show the diversity σ , the y -axes the resulting log income.

Table 5.4: Optimal level of diversity σ for each environmental dynamics

Environmental dynamics	Optimal σ
static	.005
gradual, low frequency	.008
sudden, low frequency	.046
gradual, high frequency	.057
sudden, high frequency	.123

5.5.2 Policy advice under uncertainty

When the type of environmental dynamics that affects an economy is unknown, the optimal diversity σ depends on the risk preference of the policy maker. We consider four types of risk preference: extreme risk seeking (maxmax), modest risk seeking (max average), modest risk aversion (minimax regret), and extreme risk aversion (maxmin). We calculate the optimal value for each preference from a generalization table where each σ that is optimal under one environmental dynamics is applied to the other tested environmental dynamics, including the static environment. The result is shown in Table 5.5. Each row shows the expected log income for the same environmental dynamics but different σ , each column shows results for the same σ but different environmental dynamics. Each entry is averaged over 10,000 simulations (with different instances of the social network and different random sequences of production coefficients). The values in the diagonal are highest for each row, confirming that the optimal value is indeed the best choice for a given environmental dynamics.

For each risk preference, each tested σ can now be associated with an expected value, and the σ with the best such value is considered optimal for the type of risk preference. This is shown in Table 5.6. Each row shows the value associated with each σ under a given risk adversity, with the optimal value in *Italic type*. A risk seeker looks at the highest expected log income that each tested σ has achieved under the different environmental dynamics, and chooses the highest of these. Risk neutrality means choosing the σ that maximizes the average expected log income over all environmental dynamics. Under minimal regret the σ is chosen that minimizes the greatest possible difference between actual log income and the best log income that could have been achieved. Minimal regret first calculates the maximum possible regret for each σ and all environments, and then chooses the σ that minimizes this maximum. Risk aversion means choosing the σ that promises the highest minimum log income under any environmental dynamics.

The numerical results clearly show that under uncertainty the optimal value of σ rises with the degree of risk aversion. This is in line with our earlier observation about the functional relationship between σ and expected log income: the gradient is steeper for lower values of σ than for higher values, which makes higher values of σ the safer bet. These observations are confirmed by the control experiments that test for sensitivity and scalability and that use alternative implementations of the imitation mechanism.

Table 5.5: Expected log income when a value of σ that is optimal under one environmental dynamics is applied to other environmental dynamics

Environmental dynamics that σ is applied to	Optimal σ	Environmental dynamics that σ was optimized for				
		static	gradual, low freq.	sudden, low freq.	gradual, high freq.	sudden, high freq.
static	.005	1.996	1.995	1.952	1.934	1.794
gradual, low frequency	.008	1.988	1.991	1.969	1.959	1.880
sudden, low frequency	.046	1.117	1.250	1.412	1.406	1.331
gradual, high frequency	.057	0.550	0.721	1.377	1.385	1.308
sudden, high frequency	.123	-0.042	0.092	0.529	0.548	0.596

Notes. Each column shows the expected log income when the diversity σ that is optimal for one type of environmental dynamics is applied to another type of environmental dynamics. Each row shows the expected log income when agents adapt to a specific environmental dynamics with a value σ that is optimal for another dynamics.

Table 5.6: Optimal policy advice under uncertainty and different degrees of risk aversion

Type of policy maker or society	Optimal σ	Environmental dynamics that σ was optimized for				
		static	gradual, low freq.	sudden, low freq.	gradual, high freq.	sudden, high freq.
risk seeking	.005	<i>1.996</i>	1.995	1.969	1.959	1.880
risk neutral	.008	1.122	1.210	<i>1.448</i>	1.446	1.382
minimal regret	.046	0.835	0.664	0.067	<i>0.062</i>	0.202
risk averse	.057	-0.042	0.092	0.529	0.548	<i>0.596</i>

Notes. Each entry is calculated from a column of Table 5.5 and shows the value that a risk preference assigns to a particular diversity σ . Each row shows the value that is optimal under that risk preference in *Italic type* (minimum for minimal regret, maximum for the others).

We also tested more sophisticated imitation mechanisms where either only a random selection of 10% of all agents would imitate per step, or where imitation was partial, such that a new strategy is a linear combination of the imitated strategy and the strategy of the imitating agent (again with normally distributed errors per partial strategy). We also made the selection process—the choice of which agent to imitate—dependent on income instead of growth. The arrangement of optimal values $\sigma^{opt}(d)$ is similar for each environmental dynamics. The gradient is always steeper for small values of σ than for large ones. Under uncertainty the optimal value of σ always increases with the degree of risk aversion.

5.5.3 Evolutionary dynamics

Figure 5.5–5.10 show the evolution of some key statistics during the 500 time steps of the simulation. Figure 5.5 discusses the control systems without imitation (a static environment is used). Figure 5.6 discusses the control system with imitation and a static environment. The remaining four figures show the four types of environmental dynamics where agents imitate and where changes occur gradually and with low frequency (Figure 5.7), gradually and with high frequency (Figure 5.8), suddenly and with low frequency (Figure 5.9), and suddenly and with high frequency (Figure 5.10).

Each figure contains six statistics that describe the economic performance of the agent population, the heterogeneity of their strategies, and the relevance of connectivity in the social network at each of the 500 time steps of the simulation. These statistics are averaged over 10,000 simulations. Each simulation uses a different instance of the social network and a different sequence of vectors of production coefficients. Two area plots on top of each group of six illustrate how the production coefficients change under the respective environmental dynamics. Each area plot shows a single different random sequence of production coefficients. To further ease the analysis, in Figure 5.7–5.10 dotted vertical lines are inserted into each area plot and each of the six statistics to show the points in time where the transition to a new set of production coefficients starts.

All statistics react visibly to any change in the production coefficients. This is particularly interesting for those environmental dynamics where change occurs gradually, because when a new vector of random production coefficients is introduced only the momentum changes, not the rate of change. And yet there is a clear and strong economic response to this change in momentum. Note that many statistics have reached some sort of equilibrium after the 100 steps of the initialization phase.

Of the six statistics, the first four visualize the economic performance of the agents. The first statistic shows the average efficiency (see equation 5.16) over all strategies, allowing a direct comparison to what rational agents with perfect information would achieve. The second statistic shows the behavior of the Gini coefficient, a measure of how egalitarian the accumulated capital is distributed. The third statistic shows average log income, which generally behaves as expected: after each change the income level drops temporarily, only to grow continuously thereafter. The fourth statistic shows average log growth, which falls dramatically immediately after a change, as most strategies become obsolete, but peaks within no less than ten time steps after the change, indicating that the recovery process of our evolutionary economy starts almost immediately after a change. The fifth statistic measures the variance of partial strategies within the population and shows how the heterogeneity of strategies is affected by a change in production coefficients. As discussed in Section 5.2.3, at equilibrium the square root of this variance approaches the value of the diversity control parameter σ . The sixth and final statistic measures the covariance between log income and connectivity, to emphasize the effect of a skewed distribution of connectivity on the evolutionary process. It shows how the correlation between log income and network connectivity rises each time that a new change in the production coefficients is initiated. Evidently the highly connected agents are among the first to learn and profit from the improved strategies.

When read in combination with statistic one through four on economic performance, the fifth statistic on the variance of partial strategies allows us to identify the different

phases of the adaptive evolutionary search process. After each onset of a new environmental change the standard deviation of partial strategies peaks for a brief period, then drops rapidly, and finally returns slowly to its former state. This shows that in the immediate aftermath of a change those agents that were previously most successful lose their attractive power—average efficiency is at its lowest—and that the diversity of the pool of strategies increases significantly. This is the early phase of unstructured exploration. As the agents evaluate new strategies, some agents are more successful than others—average efficiency rises again—and get heavily imitated, leading to a rapid decline in diversity. This is the second phase of structured, directed search. What is remarkable is that during this second phase the average efficiency declines for a second time and reaches a low point between ten to twenty time steps after the onset of the environmental change. Finally, during the last phase of exploration, the agents seem to finally settle into the new order. Average efficiency increases again, and as more and more agents approach the (moving) optimum, they diversify around it.

5.6 Conclusions

We have studied the general question of how different types of environmental dynamics affect behavioral interaction in an evolutionary economy. For this purpose a simple model of evolutionary formation of investment strategies through variation and selection was presented. Variation occurs when an agent replaces its own strategy by that of another agent (imitation) in an imperfect way. Selection occurs when an agent bases its choice to imitate another agent on some property of the other agent, here individual income growth. The evolutionary mechanism has one free parameter that controls diversity by defining how closely agents imitate each other. This parameter has a clear policy dimension as there are various laws and regulations that regulate how closely agents imitate each other.

If agents in an economy with a Cobb-Douglas type production function use relative income growth rate to determine which agent to imitate, the evolutionary dynamics are governed by the equilibrium growth rate of a strategy. This equilibrium growth rate is uniquely determined by the production coefficients of the Cobb-Douglas function. Modeling environmental dynamics as dynamic changes in these production coefficients enables us to study the impact of such environmental dynamics on the optimal behavioral interactions. This is a general approach that can be applied to model technological or macroeconomic dynamics as well as environmental hazards like (climate change induced) desertification and diseases.

We specified four different types of environmental dynamics that differ in the gradualness and frequency of change. We further specified one control system without imitation and one control system with imitation and a static environment. To achieve general results that are valid for a broad class of economies, all numerical results were based on large number of computer simulations, each with different instances of those factors that affect the evolutionary dynamic.

Our first research question was whether or not different values for the diversity control parameter are optimal under different environmental dynamics. We established that for almost all tested values of this parameter and all tested environments the agents

quickly adapt and find strategies with a positive equilibrium growth rate. We found further that each environmental dynamics has a unique optimal diversity that maximizes log income. This optimum increases with the frequency and rapidity of the changes.

Our second research question was whether or not different policies can be defined for different degrees of risk aversion when the type of environmental dynamics is unknown. Here we found that if there is uncertainty about the environmental dynamics, the optimal value for σ increases with the degree of risk aversion. The generality of our findings were confirmed by control experiments that tested for scalability and sensitivity of the economic parameters and that used alternative implementations of the imitation mechanism.

Various types of public policies directly and indirectly affect the imitation behavior of economic agents, the diversity of their investment strategies, and their ability to adapt to a changing environment. Numeric simulations of stochastic multi-agent systems can be used to evaluate such policies even when there is uncertainty on the specific nature of the environmental dynamics. Despite, or rather because of, their stochastic nature they can identify the preferred policy under a particular degree of risk aversion.

Appendix 5.A Evolution with variable prices

The interested reader will be curious to know how variable prices affect the evolutionary process. Regardless of the market structure and price formation, equation 5.2 for non-aggregate growth per investment sector i can be extended to include a dynamic price $p_i(t)$,

$$K_{ia}(t) = \frac{s_{ia}(t)Y_a(t-1)}{p_i(t)} + (1-\delta)K_{ia}(t-1). \quad (5.21)$$

The ratio of capital to income (equation 5.13) now converges to

$$\lim_{t \rightarrow \infty} \frac{K_{ia}(t)}{Y_a(t)} = \lim_{t \rightarrow \infty} \frac{s_{ia}(t)/p_i(t)}{\gamma_a(t) + \delta}. \quad (5.22)$$

The existence of this limit and the speed of convergence depend on the behavior of $p_i(t)$. If the price converges, the ratio of capital to income will converge as well. In that case the growth rate at equilibrium is

$$\gamma_a(t) = \beta \prod_i p_i(t)^{-\pi_i(t)} \prod_i s_{ia}(t)^{\pi_i(t)} - \delta. \quad (5.23)$$

That is, as long as the market structure does not prevent the capital-income ratio to converge in reasonable time, variable prices have a similar effect on the evolutionary process as the scaling factor. Both are monotonic transformations of the growth function that do not affect the evolutionary dynamics.




Figure 5.5: Average time evolution of an economy *without imitation* (the environment is static).

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The *x*-axis shows the 500 time steps while the *y*-axis shows the respective statistics.




Figure 5.6: Average time evolution of an economy with imitation and a *static environment*.

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x -axis shows the 500 time steps while the y -axis shows the respective statistics. Note how most statistics have stabilized during the first 100 steps of the initialization phase.

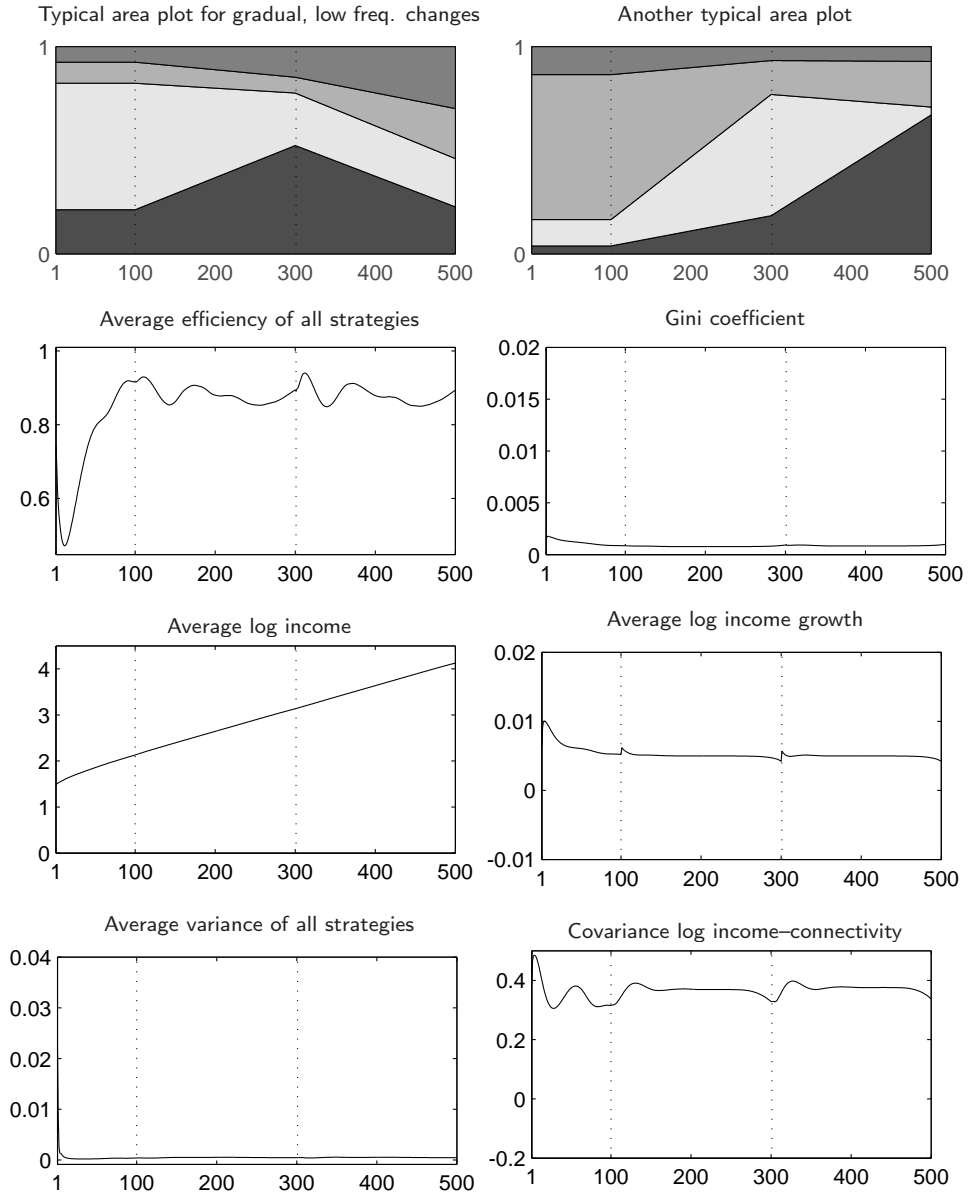


Figure 5.7: Average time evolution of an economy with imitation and a dynamic environment characterized by *gradual, low frequency* changes.

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The *x*-axis shows the 500 time steps while the *y*-axis shows the respective statistics.

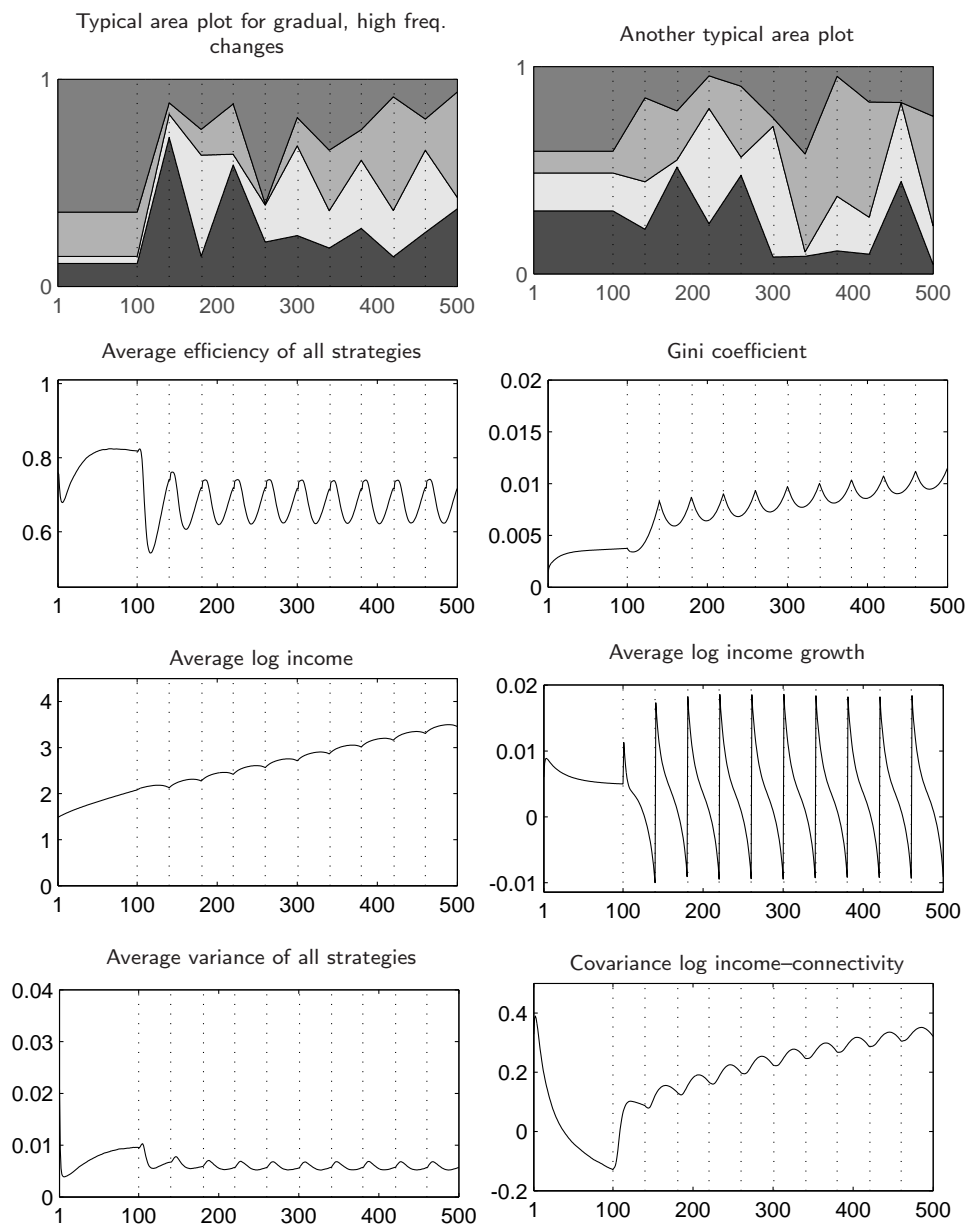


Figure 5.8: Average time evolution of an economy with imitation and a dynamic environment characterized by *gradual, high frequency* changes.

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x -axis shows the 500 time steps while the y -axis shows the respective statistics.

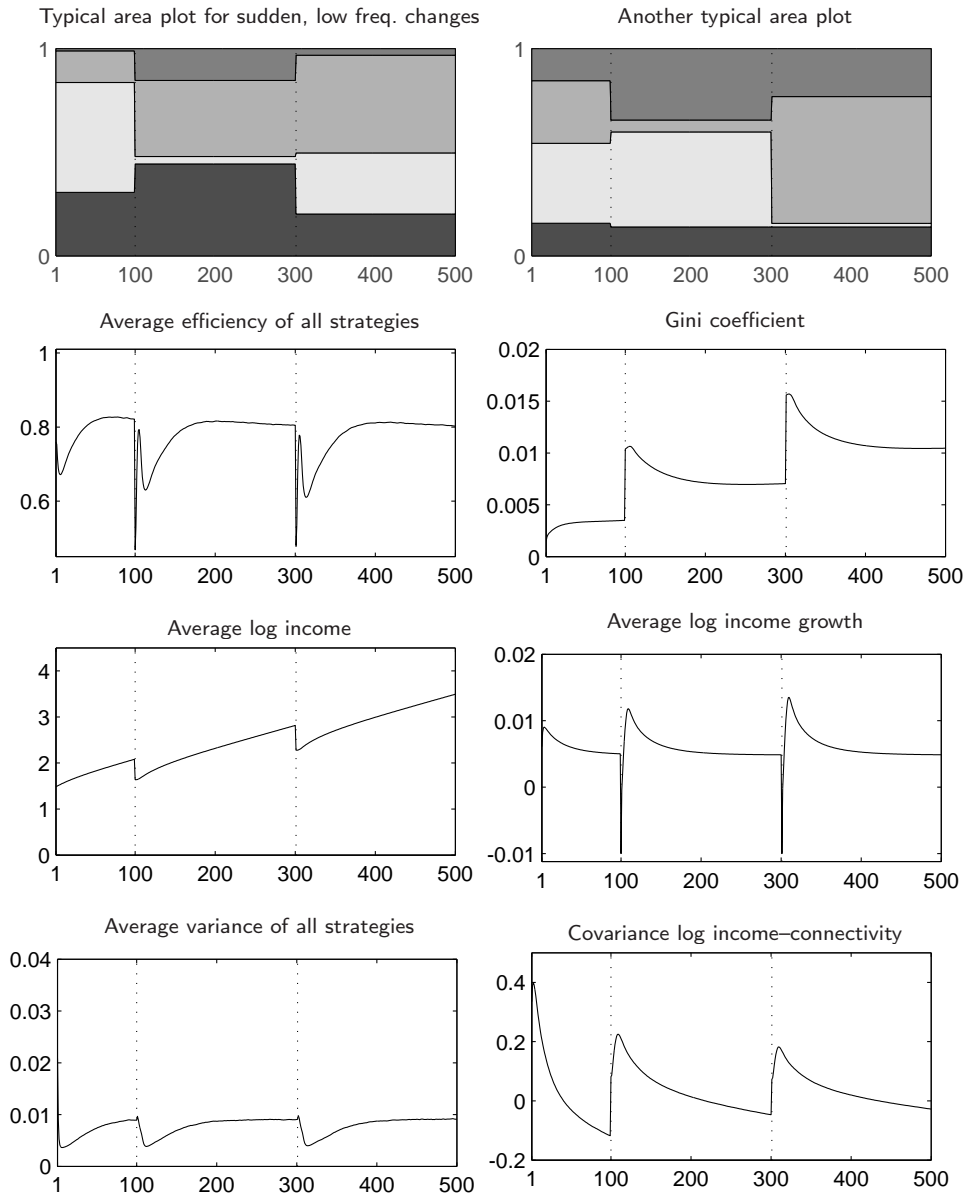


Figure 5.9: Average time evolution of an economy with imitation and a dynamic environment characterized by *sudden, low frequency* changes.

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x -axis shows the 500 time steps while the y -axis shows the respective statistics.




Figure 5.10: Average time evolution of an economy with imitation and a dynamic environment characterized by *sudden, high frequency* changes.

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x-axis shows the 500 time steps while the y-axis shows the respective statistics.

POLICY INSTRUMENTS FOR EVOLUTION OF BOUNDED RATIONALITY: APPLICATION TO CLIMATE-ENERGY PROBLEMS

Abstract

We demonstrate how an evolutionary agent-based model can be used to evaluate climate policies that take the heterogeneity of strategies of individual agents into account. An essential feature of the model is that the fitness of an economic strategy is determined by the relative welfare of the associated agent as compared to its immediate neighbors in a social network. This enables the study of policies that affect relative positions of individuals. We formulate two innovative climate policies, namely *prizes*, altering directly relative welfare, and *advertisement*, which influences the social network of interactions. The policies are illustrated using a simple model of global warming where a resource with a negative environmental impact—fossil energy—can be replaced by an environmentally neutral yet less cost-effective alternative, namely renewable energy. It is shown that the general approach enlarges the scope of economic policy analysis.

6.1 Introduction

The analysis of the economic impact of climate change and climate policy is dominated by neoclassical general equilibrium and growth models. Some models in this vein, which have played a prominent role in the IPCC and international policy debates, are: DICE (Nordhaus, 1991, 1994), RICE (Nordhaus and Yang, 1996), ENTICE (Popp, 2004), CETA (Peck and Teisberg, 1993), MERGE (Manne, 1992) and FUND (Tol, 1995). Kelly and Kolstad (1999) and van den Bergh (2004b) present brief accounts and evaluations. Al-

This chapter has been published as Nannen and van den Bergh (2010).

though these models have generated many insights, they do not represent the full range of possible approaches nor of the questions that can be addressed. They omit certain elements in their description of reality: out of equilibrium processes, choice between multiple equilibria (path-dependence), structural changes in the economy due to innovations, and the influence of income or welfare distribution on strategies. In addition, the available models assume representative agents, rational behavior, perfect information, and an aggregate production function. This approach allows for exact solutions, but it also limits the type of policies that can be studied. For example, they cannot study the effects of information provision, or of exemplary reward and punishment.

Here we present a model that starts from a set of alternative feasible assumptions offered by evolutionary economics (Nelson and Winter, 1982; Dopfer, 2005; Witt, 2008) and by agent-based computational economics (Tesfatsion, 2006; Levy et al., 2000; Epstein and Axtell, 1996). Evolutionary modeling has gained some popularity in economics, but most studies in this vein lack a policy dimension. This holds especially for applications of pure evolutionary game theory (Friedman, 1998). However, agent-based models applied to economics have rarely addressed public policy issues, and if they have done so, only in a way that does not fully exploit the policy potential offered by an evolutionary model (Janssen and Ostrom, 2006). The present study adds a policy angle to an agent-based evolutionary approach, focusing on the opportunities that an evolutionary system offers for policy design and analysis. This approach recognizes evolution in the economy rather than emphasizing the use of evolutionary algorithms to optimize non-evolutionary complex systems (e.g. Janssen et al., 2004).

The evolutionary agent-based model developed here addresses policy in a setting of global warming. The latter is endogenous to the model and depends on the source of energy used by agents in the model. These agents may be interpreted as national or regional authorities in charge of the energy policy of an independent economy. Global warming is assumed to have a negative effect on social welfare. The overall goal at the global level is to replace a resource with a negative impact on social welfare—fossil energy—by a neutral alternative, namely renewable energy. On an individual level the alternative comes with no economic advantage, and possibly even with a disadvantage, so that there is no incentive to adopt it. A complicating factor is that there is no central authority that can formulate and enforce a policy. Climate policies are usually based on international agreements, and compliance by countries is voluntary.

Each agent is modeled individually and agents are assigned only limited information and boundedly rational capabilities. Their objective is assumed to be to reach a high level of individual welfare. The only information available to the agents are the investment strategies and the income growth rates of their fellow agents. The agents believe that there is a strong causal link between these two variables. Since they prefer a high over a low income growth rate, they imitate the investment strategy of a fellow agent when that fellow agent realizes an income growth rate that is high relative to their own income growth rate and that of their other fellow agents. That is, they imitate an investment strategy when they believe that it causes a relatively high income growth rate. This approach is inspired by findings on relative welfare and income comparison effect of happiness or “subjective well-being” studies (e.g. Ferrer-i Carbonell, 2005; Frank, 1987).

Imitation is never perfect. Small errors are introduced during the imitation process that lead to slightly different variants of the same investment strategy. While these errors

are necessary to maintain diversity within the pool of strategies and to allow a population of imitating agents to find and converge on the individually optimal strategy, they also form a hitherto unexploited opportunity for the policy maker. If the desirable variant is given a selective advantage over the undesirable variant, the first will diffuse faster and will ultimately be used more often. For example, a policy that aims for agents to adopt a greener investment strategy—in the sense that they invest less in fossil energy—can do so by convincing at least some agents that the greenest variants in the current pool of strategies lead to a relatively high income growth rate. As these strategies are imitated, the errors guarantee that some of the new variants will be greener still. An evolutionary policy can thus “breed” a new strategy by progressively giving a selective advantage to the most desirable variants.

As has been extensively discussed by Wilhite (2006), agent-based simulation of economic processes needs to give proper attention to the social network. Communication links between economic agents, individuals and institutions, are neither regular nor random. They are the result of a development process that is steered by geographic proximity, shared history, ethnic and religious affiliation, common economic interests, and much else. The social networks used for this study reproduce a number of stylized facts that are commonly found in real social networks: the small world effect (Erdős and Rényi, 1959), a high clustering coefficient (Watts and Strogatz, 1998), and a scale-free degree distribution (Barabási and Albert, 1999). Modeling the evolution of strategies in such complex social networks allows us to formulate economic policies that exploit the effect of social visibility on the diffusion of a desirable strategy.

The frequency with which the strategy of an agent is imitated depends on two factors: the relative welfare of an agent, as observed from the income growth rate, and the position of an agent in the social network. A policy that is aimed at increasing the selective advantage of a particular strategy can do so in two ways, namely by changing the relative welfare of an exemplary agent that uses such a strategy, or by changing the position of such an agent in the social network. We formulate policies for both alternatives, and compare their effect to that of a standard tax on fossil energy. One policy, *prizes*, increases the relative welfare of exemplary agents by awarding them a monetary prize. The other policy, *advertisement*, increases the social visibility of exemplary agents by increasing their connectivity in the social network. Policy tools that increase social visibility include, for instance, sponsorship of industrial fairs and scientific venues, and the production and distribution of educative material.

The remainder of this chapter is organized as follows. Section 6.2 presents the climate-economy model, including the evolutionary mechanisms of strategy formation and diffusion, and formulates the climate policies. Section 6.3 studies the convergence behavior of the resulting evolutionary process. Section 6.4 evaluates the climate policies using numerical simulations with the climate-economy model. Section 6.5 concludes.

6.2 The economic model

6.2.1 General features of the model

The present economic model is formulated in order to study the effectiveness of regulatory public policies when economic behavior evolves through social interactions. The

approach focuses on climate policies and energy investment strategies, but the model can easily be adapted to other problems. Each agent controls an independent economy with its own supply and production. The agent formulates a strategy to invest current domestic income in different sectors. The returns for each independent economy are then calculated from standard economic growth and production functions. Some allocations give higher returns than others, and the goal of the agents is to find a strategy that can realize a high level of individual welfare.

The present model is loosely based on the influential work of W. D. Nordhaus, who published a series of general-equilibrium economic models of climate policy and global warming, starting with the DICE model (Nordhaus, 1992). From this model all economic factors that were not essential to the current study were removed, in particular elements relating to labor, technical details of global warming, and resource constraints. The reason is that our model aims to be illustrative rather than to accurately replicate reality. Moreover, simplification here allows for additional complexity in the module describing the evolution of strategies.

A fundamental difference between the present evolutionary agent-based approach and the general-equilibrium approach of Nordhaus is that here agents do not make perfectly rational decisions that are based on perfect knowledge. Instead, agents evolve their strategies through random mutation and selective imitation in a social network. Moreover, while here agents are homogeneous in terms of production functions, initial strategies and initial income, they are heterogeneous in their placement in the social network and the information they receive, and their strategies and income quickly diverge.

The numerical simulations are based on a discrete synchronous time model where the economy and strategy of each agent are updated in parallel at fixed time intervals. A time step is divided into two separate update operations: 1) updating the economy: each agent invests its income according to its own investment strategy and the individual returns are calculated by standard growth and production functions; 2) updating the strategies: all agents compare their strategies and those agent that decide to imitate change their strategies simultaneously. Each policy is evaluated over a period of 400 time steps, simulating 400 quarters or 100 years, a period that is sufficiently long to study the long-term effects of a policy on climate and welfare. As no significant financial market requires a publicly traded company to publish financial results more than 4 times a year, we consider a quarter to be the limit of feasibility to account for growth and to review an economic strategy. Given habitual behavior and organizational routines (Nelson and Winter, 1982), most economic agents will in fact review their strategy less often.

6.2.2 Strategies, investment, and production

All parameters of the economic model are summarized in Table 6.1. Our basic model of energy investment consists of three investment sectors: general capital K , fossil energy F , and renewable energy R . Here, the capital accumulation in an energy sector includes technology, infrastructure, and licenses for production, distribution, and consumption of a particular form of energy. Let $Y_a(t)$ be the income of agent a at time t . Formally, the

Table 6.1: Economic and climate variables

a, b	individual agents	k	average number of neighbors per agent
N	neighbors of an agent	\mathcal{C}	clustering coefficient of the network
s	investment strategy	σ	mutation variance of strategies
\mathcal{F}	fitness of a strategy	Q	income without global warming
Y	net domestic income	γ	income growth rate
K	general capital sector	α	Cobb-Douglas exponent
F	fossil energy sector	τ	tax on fossil energy investments
δ	capital discount rate	\mathcal{T}	revenue of the fossil energy tax
R	renewable energy sector	c	additional cost of renewable energy
G	greenhouse gas level	ϕ	breakdown fraction of greenhouse gas
ν	scale of climatic damage	ε	environmental tax on income
β	scaling factor	\mathcal{E}	fund financed by the environmental tax
t	time step		

investment strategy $s_a(t)$ of an agent can be defined as a three dimensional vector

$$s_a(t) = [0, 1]^3, \quad \sum_{i \in \{K, F, R\}} s_{ia}(t) = 1. \quad (6.1)$$

The non-negative partial strategy $s_{ia}(t)$ determines the fraction of the previous period's income $Y_a(t-1)$ that agent a invests in sector i at time t . All partial strategies are constrained to add up to one. The set of all possible investment strategies is a two dimensional simplex (i.e., a triangular surface) embedded in a three dimensional Euclidean space.

Invested capital is non-malleable: once invested it can not be transferred between sectors. The accumulation of capital in each sector depends on individual investment and the global depreciation rate δ , which is constant and equal for all sectors and all agents. In the case of fossil energy the investment can be reduced by a regulatory tax τ on investments in the fossil energy sector. This tax is defined as a fraction of fossil energy investments before taxes, so that a tax of $\tau = 100\%$ doubles the cost of all expenditures on production, distribution, and consumption of fossil energy. In this way, if an agent's total spending on fossil energy is $x = Y_a(t-1)s_{F,a}(t)$, then an amount of $\frac{x}{1+\tau}$ is indeed invested, while the remaining $\frac{x\tau}{1+\tau}$ is paid as a tax. The revenue

$$\mathcal{T}(t) = \frac{\tau}{1+\tau} \sum_a Y_a(t-1)s_{F,a}(t) \quad (6.2)$$

of this tax is recycled and distributed evenly among all agents.

To model a competitive disadvantage for renewable energy—for example through a higher cost of technology, production, or storage—we introduce an additional cost c for renewable energy, representing the difference between the costs of renewable and fossil energy. In analogy to the fossil energy tax τ , we express this additional cost in percent of the unit cost of fossil energy before taxes, i.e., renewable energy is twice as expensive as fossil energy before taxes when $c = 100\%$. The difference equations for non-aggregate

growth per sector are then

$$\Delta K_a(t) = Y_a(t-1)s_{K,a}(t) - \delta K_a(t-1) \quad (6.3)$$

$$\Delta F_a(t) = \frac{Y_a(t-1)s_{F,a}(t)}{1+\tau} - \delta F_a(t-1) \quad (6.4)$$

$$\Delta R_a(t) = \frac{Y_a(t-1)s_{R,a}(t)}{1+c} - \delta R_a(t-1). \quad (6.5)$$

We proceed by first calculating the income of agent a as if there had been no global warming, and then by accounting for global warming. We calculate $Q_a(t)$ from the returns of the individual capital sectors by a Cobb-Douglas type production function with constant returns to scale and constant elasticity of substitution,¹

$$Q_a(t) = \beta(K_a(t))^\alpha (F_a(t) + R_a(t))^{1-\alpha}, \quad (6.6)$$

where β is a scaling factor. In this production function fossil energy and renewable energy are assumed to be perfect substitutes: one can completely replace the other. General capital and combined energy are assumed to be imperfect substitutes. Production requires both types of input, and only a specific combination will lead to a high production level.

Global warming is commonly defined as the increase of global mean temperature above the pre-industrial mean, due to an increased level of atmospheric greenhouse gases $G(t)$. The dynamics of the greenhouse gas effect includes many local and global subsystems, resulting in complex and chaotic dynamics that allow for a range of possible climate scenarios (e.g., Stainforth et al., 2005). Here we just include a simple feedback loop that captures one of the main characteristics of greenhouse gas induced economic damage: a long delay between action and reaction that spans several decades. We do so by modeling the level of atmospheric greenhouse gases as a result of only two factors: cumulative fossil energy consumption by economic agents, which we assume to be equal to the total amount of capital accumulated in the fossil energy sector, and a natural breakdown fraction ϕ ,

$$\Delta G(t) = \sum_a F_a(t) - \phi G(t-1). \quad (6.7)$$

We assume that renewable energy does not contribute to global warming. We further pose the relationship between economic damage, global warming, and economic damage to be linear, scaled by a factor ν . The net income $Y_a(t)$ of an agent a can then be calculated as

$$Y_a(t) = Q_a(t) [1 - \nu G(t)] + \frac{\mathcal{T}(t-1)}{|P|}. \quad (6.8)$$

¹Instead of including the fossil energy tax τ and the additional cost c of renewable energy in the growth functions, they might be incorporated in the production function,

$$Q_a(t) = \beta(K_a(t))^\alpha \left(\frac{F_a(t)}{1+\tau} + \frac{R_a(t)}{1+c} \right)^{1-\alpha}.$$

where $\mathcal{F}(t-1)$ are the revenues from the regulatory tax τ , distributed with one time step delay among the $|P|$ agents of the population. The growth rate $\gamma_a(t)$ of the income of agent a is

$$\gamma_a(t) = \frac{Y_a(t)}{Y_a(t-1)} - 1. \quad (6.9)$$

6.2.3 Evolution of strategies

From the point of view of evolutionary modeling, agents and investment strategies are not the same: an agent carries or maintains a strategy, but it can change its strategy and we still consider it to be the same agent (Nowak, 2006). Because every agent has exactly one strategy at a time, the number of active strategies is the same as the number of agents.

To model which agents an agent can imitate we use a generic class of social networks that has been well studied and validated in network theory, namely those that can be generated by a random process with preferential attachment and that have a high clustering coefficient, see Section 4.2.3 on page 71 for details. Before the start of each simulation a stochastic process assigns to each agent a a set of peers N_a that does not change during the course of the simulation. If agent a is a peer of agent b , then a will consider the income growth rate and the investment strategy of b when choosing an agent for imitation, while b will consider the income growth rate and the investment strategy of a . On the other hand, if a and b are not peers, they will not consider each other for the purpose of imitation.

At each time step an agent may select one of its peers in the social network and imitate its strategy. If that happens, the strategy of the imitating agent changes, while the strategy of the imitated agent does not. The choice of which agent to imitate is based on relative welfare as indicated by the current growth rate of income. Note that the relation between income $Y_a(t)$ and growth $\gamma_a(t)$ is

$$Y_a(t) = Y_a(0) \prod_{i=1}^t [\gamma_a(i) + 1]. \quad (6.10)$$

The imitating agent always selects the peer with the highest current income growth rate. Only if an agent has no peer with an income growth rate higher than its own, the agent does not revise its strategy.

If imitation were the only mechanism by which agents change their strategies, the strategies of agents that form a connected network must converge on a strategy that was present during the initial setup. However, real imitation is never without errors. Errors are called mutations in evolutionary theory. They are fundamental to an evolutionary process because they create and maintain the diversity on which selection can work. In this model we implement mutation by adding some Gaussian noise to the imitation process. That is, when an agent imitates a strategy, it adds some random noise drawn from a Gaussian distribution with zero mean to each partial strategy. This causes small mutations along each partial strategy to be more likely than large ones. The exact formula by which agent a imitates and then mutates the strategy of agent b is

$$s_a(t) = s_b(t-1) + \mathcal{N}(0, \sigma), \quad (6.11)$$

where $\mathcal{N}(0, \sigma)$ denotes a normally distributed three dimensional random vector with zero mean and standard deviation σ per dimension. Because partial investment strategies have to sum to one, we have to enforce $\mathcal{N}(0, \sigma) = 0$, for example by orthogonal projection on the simplex, resulting in the loss of one degree of freedom. The error term is further constraint to leave all partial strategies positive. Needless to say that we do not imply that our boundedly rational agents engage consciously in such mathematical exercises. Subjectively they merely allocate their income such that none is left.

In order to measure the impact of an individual agent on the evolution of strategies at the population level, we need to introduce the concept of fitness. In analogy with biology, where fitness usually measures an individual's capability to reproduce, we define the fitness of an economic agent as the frequency with which it is imitated. In the model that has been presented so far, the frequency with which agent a is imitated is fully determined by the income growth rate of a and its first and second degree neighbors. Further degrees do not matter. First degree neighbors are relevant because only direct neighbors consider a for imitation. Second degree neighbors are relevant as they are the agents that a competes with. Agent a will only be imitated by agent b if a has a higher income growth rate than b and all other neighbors of b (who are second degree neighbors of a).

This functional relationship can be expressed by a fitness function. Let $\{N_b \cup b\}$ be the set consisting of agent b and its peers, i.e., those agents with which agent b compares its income growth rate. Let $\gamma_{N_b \cup b}^{max}(t)$ be the income growth rate of the fastest growing agent in this set at time t ,

$$\gamma_{N_b \cup b}^{max}(t) = \operatorname{argmax}_{c \in \{N_b \cup b\}} \gamma_c(t). \quad (6.12)$$

Then the fitness $\mathcal{F}_a(t)$ of agent a at time t is

$$\mathcal{F}_a(t) = \sum_{b \in N_a} \begin{cases} 1 & \text{if } \gamma_a(t) = \gamma_{N_b \cup b}^{max}(t), \\ 0 & \text{otherwise.} \end{cases} \quad (6.13)$$

Or, in set notation:

$$\mathcal{F}_a(t) = |\{b | b \in N_a \wedge \gamma_a(t) = \gamma_{N_b \cup b}^{max}(t)\}|. \quad (6.14)$$

In this function the fitness of an agent is bounded by the number of its neighbors. An agent a_1 who has just one neighbor and has the highest income growth rate among the neighbors of that neighbor has a fitness of one, whereas an agent a_{10} who has ten other neighbors and whose income growth rate is highest among the neighbors of just two of them has a fitness of two, even if in absolute terms a_1 has a much higher income growth rate than a_{10} . We see that the principal factors that determine the fitness of an agent are relative welfare as indicated by the current growth rate of income as well as the number of agents it communicates with. This gives us two different means by which a policy can regulate the evolution of economic strategies: either by changing the income growth rate of some agents, depending on the desirability of their current strategies, or by changing their connectivity in the social network, again depending on the desirability of their current strategies.

6.2.4 Policy goals and formulation

The goal of the policies that are being studied here is to let the economic agents reach a high social welfare. Assuming that fossil fuel consumption has a negative economic impact because of the associated global warming, a successful policy has to reduce consumption in fossil fuels but without considerably reducing social welfare, such that the social costs of implementing the policy do not outweigh the social benefits from a reduction in global warming.

We will study three policies, starting with a tax τ on fossil energy investments. This is the first best policy under traditionally assumed conditions (rational agents, perfect markets), and we study it here in the context of imperfect information and bounded rationality. It is a regulatory and not a revenue raising tax and is defined as a fixed percentage on all investments in fossil energy, cf. equation 6.2, 6.4, and 6.8. We compare this standard policy with two novel policies that take advantage of the evolutionary process by increasing the fitness of those agents that invest a larger fraction of their income in renewable energy. These policies increase the fitness of an agent either by increasing its income growth rate, or by increasing its visibility in the social network. The rationale is that, if we increase the fitness of agents that use certain strategies, these strategies will be employed more frequently.

Under the first policy, agents pay a tax that is proportional to their investment in fossil fuel. This tax makes investment in fossil energy economically less attractive. However, since the incentive not to comply with this policy is also proportional to their investment in fossil fuel, the effect of this policy depends much on the existence of a central authority that can enforce it. The second policy studied here, *prizes*, increases the fitness of agents that invest a larger fraction of their income in renewable energy by awarding them a monetary prize, financed by a global tax that is payed by all agents. That is, it is not important who pays the tax, as long as someone pays it, for example those agents that are most affected by global warming. This does not entirely solve the problem of compliance, but makes it less acute. The third policy, *advertisement*, increases the fitness of agents that invest a higher fraction of their income in renewable energy by increasing their social visibility, i.e. their connectivity in the social network. No compliance is required.

The *prizes* policy gives a monetary prize to those agents who invest the largest fraction of income in renewable energy, increasing their relative welfare, and with that their fitness. This prize is financed by an environmental tax ε on production $Q_a(t)$. Since this is a revenue raising tax to finance the policy and not a regulatory tax that depends on individual behavior, it has the same level for each agent. Let $\mathcal{E}(t)$ be the size of the environmental fund at time t :

$$\mathcal{E}(t) = \sum_a Q_a(t) [1 - \nu G(t)] \varepsilon. \quad (6.15)$$

At each time step, the q agents that invest the highest fraction of their income in renewable energy are each awarded an equal share $\mathcal{E}(t-1)/q$, such that under the *prizes* policy the net income becomes

$$Y_a(t) = Q_a(t) [1 - \nu G(t)] [1 - \varepsilon] + \begin{cases} \mathcal{E}(t-1)/q & \text{if } a \text{ is awarded a prize,} \\ 0 & \text{otherwise.} \end{cases} \quad (6.16)$$

Table 6.2: Free policy parameters

<i>Fossil energy tax</i>	τ	tax on fossil energy investments
<i>Prizes</i>	q	number of agents that receive a prize
	ε	tax on income to finance the prize
<i>Advertisement</i>	q	number of agents that are advertised
	p	probability that an agent is reached by advertisement (the simulations use a fixed value of $p = .25$)

To give an example: if the income tax ε is 1%, and 10 out of 200 agents are selected to receive a prize, then under the assumption that their income does not deviate significantly from the average income, it is raised by about 20%. If the majority of agents receive a prize, the tax to finance the prizes is in effect a selective punishment of those agents that invest relatively much in fossil fuels.

The *advertisement* policy increases the social visibility of those agents that invest the largest fraction of their income in the renewable energy sector, increasing the number of agents that consider the advertised agents when deciding whom to imitate. At each time step the q agents that invest the largest fraction of their income in renewables are selected to be advertised. The advertised agents are temporarily added to the group of neighbors of some other agents, so that these other agents consider the advertised agents when deciding whom to imitate. Advertisement does not oblige an agent to consider an advertised agent. Instead, its success rate depends on the resources invested in the campaign. For simplicity we assume that whether agent a considers the advertised agent b for imitation is an independent random event for each a , b and t and has probability p . We ignore the cost of advertisement and assume a success rate of just $p = .25$.

To give an example, let the average number of neighbors per agent before advertisement be $k = 10$ and let $q = 8$ agents be selected for advertisement. On average, each agent can now choose between $k + q * p = 10 + 8 * .25 = 12$ neighbors when deciding whom to imitate. If an agent imitates, chances are one in six that it imitates the strategy of an advertised agent, provided that the income of the advertised agents does not deviate significantly from that of the other agents. The free parameters of each policy are listed in Table 6.2.

6.2.5 Model calibration

The free parameters of the economic model are calibrated such that global warming has a significant negative welfare effect, emphasizing the need for policies. The calibrated values of all free economic parameters are summarized in Table 6.3. A fixed number 200 of agents is used in all simulations; this is approximately the number of independent states and a rough approximation of the number of agents with an independent energy policy. The quarterly capital discount rate is $\delta = .01$. The exponent of general capital in the production function is $\alpha = .9$, and the exponent of the combined energy sector is $1 - \alpha = .1$. In this way income is highest when 90% of an agent's capital is in the general sector and 10% in the two energy sectors. The scaling factor β of the production function

Table 6.3: Calibrated parameter values of the economic model

k	network connectivity	10	\mathcal{C}	clustering coefficient	.66
σ	mutation variance	.02	α	Cobb-Douglas exponent	.9
δ	capital discount rate	.01	ϕ	breakdown of greenhouse gases	.01
β	scaling factor	.021	ν	scale of climatic damage	.00007

is calibrated such that the calibrated economic model without climate damage has an economic growth rate of about 2% per annum. The breakdown fraction ϕ of greenhouse gas and the sensitivity ν to global warming are calibrated such that without any climate policy the greenhouse gas emissions reduce the per annum growth rate by an order of magnitude over the 100 years of the simulation, consistent with the studies reported in Section 6.1. For the social network we use an average connectivity of $k = 10$. In a population of 200 agents this value results in a highly connected network—the average distance between any two agents in the network is 2.7—while maintaining the overall qualities of a complex network.

The mutation variance σ is the only free parameter that regulates the evolutionary mechanism. Small values of σ slow down the discovery of a good strategy. Large values prevent convergence. A good value of σ lies somewhere in between. Figure 6.1 shows how the average income of the agents depends on σ . The x -axis shows different values for σ . The y -axis shows the average income that a population of 200 agents realize after 400 time steps. Each measurement point in the graph is averaged over 100 simulations. The initial strategy of each agent is chosen at random. There are no taxes, global warming has no effect, and the additional cost of renewable energy is $c = 100\%$. Under these conditions the optimal strategy that maximizes the income growth rate of an individual agent is $\langle s_K, s_F, s_R \rangle = \langle .9, .1, 0 \rangle$. The graph shows that average income is maximized for a value of approximately $\sigma \approx .02$, and for this reason we use a value of $\sigma = .02$ in the remainder of this study.

In order to avoid any dependency of the simulation results on initial conditions, the numeric simulations are divided into an initialization phase and a main experimental phase. During the initialization phase certain parameters of the evolutionary economy will converge regardless of the initial conditions, contributing to the general validity of the numerical results. The initialization effect is visualized in Figure 6.2, which shows how the average investment strategy converges on an equilibrium. 800 time steps are




Figure 6.1: Effect of the mutation variance on economic performance




Figure 6.2: Convergence of the average investment strategy

shown. Initial strategies are chosen at random from the two dimensional simplex, and so at $t = 1$ the average strategy is $\langle s_K, s_F, s_R \rangle = \langle 1/3, 1/3, 1/3 \rangle$. The average strategy at $t = 800$ is $\langle .9, .09, .01 \rangle$. Results are averaged over 10,000 simulations.

Note that while it takes the agents only about a dozen time steps to learn to invest some 90% of their investments in general capital, they need about 200 time steps to become sufficiently sensitive to the difference in cost between the two energy sectors and to differentiate their energy investments. From $t = 200$ to $t = 400$ the convergence on the final strategy can be seen to follow a damped oscillation pattern. The full effect of a policy can only be established if it is introduced after the system without policy has reached equilibrium. It takes 400 time steps for the system without a policy to converge, and so we base the numeric evaluation of climate policies on simulations that consist of an initialization phase of 400 time steps during which no policies are applied, followed by a main experimental phase of 400 time steps during which policies are applied and evaluated. In particular, the tax τ on fossil energy investments and the environmental tax ε are always zero up until $t = 400$. Only from $t = 401$ they take the value assigned to them by the respective policy.

6.3 The evolutionary dynamics

6.3.1 Derivation of the growth function

An important prerequisite for regulating an evolutionary system is to understand its dynamics. Here we are primarily interested in what strategy the agents will converge on. With regard to global warming we are further interested in whether the evolutionary agents can converge on a globally optimal strategy, rather than on individually optimal

strategies.

While the fitness function describes how the imitation of a strategy (the genotype) depends on welfare as indicated by the income growth rate (the phenotype), the *growth function* describes how a strategy determines the income growth rate of the agent that carries it. The growth function calculates the economic utility of a strategy as the equilibrium growth rate to which the income growth rate of an agent converges if it holds on to that particular strategy. Derivation of the growth function is essential for an understanding of the evolutionary dynamics. We will base it on an analysis of the ratio of sector specific capital to income.

To start with, equation 6.4 and 6.5 can be combined to express the difference equation of a combined energy sector $E = F + R$,

$$\Delta E_a(t) = Y_a(t-1) \left(\frac{s_{F,a}(t)}{1+\tau} + \frac{s_{R,a}(t)}{1+c} \right) - \delta E_a(t-1), \quad (6.17)$$

where the combined energy investment strategy of an agent is $s_{F,a}(t) + s_{R,a}(t) = 1 - s_{K,a}(t)$. Let $r_a(t)$ be the fraction of $1 - s_{K,a}(t)$ that is invested in fossil energy, and $1 - r_a(t)$ the fraction that is invested in renewable energy,

$$r_a(t) = \frac{s_{F,a}(t)}{1 - s_{K,a}(t)}. \quad (6.18)$$

This enables us to rewrite equation 6.17 as

$$\Delta E_a(t) = Y_a(t-1) [1 - s_{K,a}(t)] f(r_a, t) - \delta E_a(t-1), \quad (6.19)$$

where $f(r_a, t)$ stands for

$$f(r_a, t) = \frac{r_a(t)}{1+\tau} + \frac{1-r_a(t)}{1+c}. \quad (6.20)$$

We collapse the scaling factor and the economic effect of global warming into a single factor $\zeta(t)$,

$$\zeta(t) = \beta [1 - vG(t)]. \quad (6.21)$$

Next we combine the calculation of income (equation 6.8) with the production function (equation 6.6) and simplify it by ignoring the additive term $\frac{\mathcal{F}(t-1)}{|P|}$, which is identical for all agents,

$$Y_a(t) = \zeta(t) K_a(t)^\alpha E_a(t)^{1-\alpha}. \quad (6.22)$$

We now use equation 6.3 and 6.9 to calculate the difference equation of the ratio of general capital to income as

$$\begin{aligned} \frac{K_a(t)}{Y_a(t)} &= \frac{Y_a(t-1) s_{K,a}(t) + (1-\delta)K_a(t-1)}{(\gamma_a(t)+1) Y_a(t-1)} \\ &= \frac{s_{K,a}(t)}{\gamma_a(t)+1} + \frac{1-\delta}{\gamma_a(t)+1} \frac{K_a(t-1)}{Y_a(t-1)}. \end{aligned} \quad (6.23)$$

This dynamic equation is of the form

$$x(t) = a + bx(t-1), \quad (6.24)$$

which under the condition $0 \leq b < 1$ converges monotonically to its unique stable equilibrium at

$$\lim_{t \rightarrow \infty} x(t) = a/(1-b). \quad (6.25)$$

In a model without global warming this condition is normally fulfilled: investment is always non-negative and sector specific capital cannot decrease faster than δ . Excessive economic damage caused by global warming, $vG(t)$, does theoretically allow for $\gamma_a \leq -\delta$. However, the social and political ramifications of such a catastrophic decline go beyond the scope of this model. Hence, with the restriction that this model only covers the case $\gamma_a > -\delta$, and considering that $0 < \delta \leq 1$, we have the required constraint for convergence

$$0 \leq \frac{1-\delta}{\gamma_a(t)+1} < 1. \quad (6.26)$$

We conclude that the ratio of general capital to income converges to

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{K_a(t)}{Y_a(t)} &= \lim_{t \rightarrow \infty} \frac{s_{K,a}(t)}{\gamma_a(t)+1} / \left(1 - \frac{1-\delta}{\gamma_a(t)+1}\right) \\ &= \lim_{t \rightarrow \infty} \frac{s_{K,a}(t)}{\gamma_a(t)+\delta}. \end{aligned} \quad (6.27)$$

Equation 6.27 describes a unique stable equilibrium to which the ratio of general capital to income converges monotonically. A similar result can be obtained for the energy sector:

$$\lim_{t \rightarrow \infty} \frac{E_a(t)}{Y_a(t)} = \lim_{t \rightarrow \infty} \frac{[1-s_{K,a}(t)]f(r_a, t)}{\gamma_a(t)+\delta}. \quad (6.28)$$

Ignoring the limit notation we combine equation 6.27 and 6.28 with equation 6.22 to calculate income at equilibrium as

$$\begin{aligned} Y_a(t) &= \zeta(t) \left(\frac{Y_a(t-1) s_{K,a}(t)}{\gamma_a(t)+\delta} \right)^\alpha \left(\frac{Y_a(t-1) [1-s_{K,a}(t)] f(r_a, t)}{\gamma_a(t)+\delta} \right)^{1-\alpha} \\ &= \zeta(t) \frac{Y_a(t-1)}{\gamma_a(t)+\delta} s_{K,a}(t)^\alpha [1-s_{K,a}(t)]^{1-\alpha} f(r_a, t)^{1-\alpha}. \end{aligned} \quad (6.29)$$

Solving for $\gamma_a(t)$ yields the growth function

$$\gamma_a(t) = \zeta(t) s_{K,a}(t)^\alpha [1-s_{K,a}(t)]^{1-\alpha} f(r_a, t)^{1-\alpha} - \delta. \quad (6.30)$$

6.3.2 Convergence behavior

We can now address the question whether evolutionary agents can be expected to converge on the globally rather than on the individually optimal strategy. In equation 6.27 and 6.28 neither the rate of convergence nor the equilibrium itself depend on the value of $\zeta(t)$. In equation 6.30 we find that $\zeta(t)$ is a multiplicative factor that does not change the relative order of the equilibrium growth rate of individual strategies. Since the fitness




Figure 6.3: Growth effect of investment in general capital. The production coefficient is $\alpha = .9$.

of an agent depends on the order of income growth rates, the fitness function is invariant under such a monotonic transformation. In other words, $\zeta(t)$ does not change the likelihood of a particular strategy to be imitated. This means that global warming has no effect on the evolutionary process: agents must not be expected to show any type of behavioral response to the economic effects of global warming and are not likely to choose the globally over the individually best strategy.

To answer the question of which strategy the agents will converge on, the growth function of equation 6.30 can be decomposed into a term that describes the effect of income allocation to general capital on growth, and a term that describes the growth effect of the allocation of the remaining income over the two energy sectors. The dependency of the equilibrium growth rate on the general capital allocation as seen in equation 6.30 is given by the term

$$s_{K,a}(t)^\alpha [1 - s_{K,a}(t)]^{1-\alpha}, \quad (6.31)$$

which depends exclusively on the constant production coefficient α . This term is maximized for $s_{K,a}(t) = \alpha$, which implies that the optimal allocation to the combined energy sector is $s_{F,a}(t) + s_{R,a}(t) = 1 - \alpha$. As can be seen in Figure 6.3, the growth effect is a concave function of $s_{K,a}(t)$ with an extended region around the maximum that has a gradient close to zero.

The effect of $r_a(t)$ on the income growth rate is via $f(r_a, t)^{1-\alpha}$ which, from equation 6.20, is

$$f(r_a, t)^{1-\alpha} = \left(\frac{r_a(t)}{1+\tau} + \frac{1-r_a(t)}{1+c} \right)^{1-\alpha}. \quad (6.32)$$

Figure 6.4 shows that this function is flat when $\tau = c$ and otherwise concave. For $\tau > c$ the term is maximized when $r_a(t) = 0$, and for $\tau < c$ it is maximized when $r_a(t) = 1$. For a given value of $r_a(t)$ the curvature increases with $|\tau - c|$. For a given value of $|\tau - c|$ the curvature increases with the distance to the maximum. Unlike term 6.31, the curvature at the maximum is not zero. Maximizing this type of growth function poses no challenge to a (collective) learning mechanism. It has a single global optimum, no local optima, and a distinct slope that increases with distance to the optimum. Even the simplest of hill climbing algorithms can find and approach this optimum. Learning mechanisms will differ mostly in the speed of convergence.




Figure 6.4: Growth effect of different investment allocations. The production coefficient is $\alpha = .9$ and the total energy allocation is $1 - s_{K,a}(t) = .1$. In the upper graph the tax on fossil energy investments is $\tau = 0$. In the lower graph the additional cost of renewable energy is $c = 100\%$.

Regarding the speed with which our evolutionary agents converge on the optimum strategy, we must bear in mind that evolutionary agents will only converge on the individually optimum strategy if there is sufficient selection pressure. Figure 6.2 shows that the speed of convergence gradually decreases as the optimum strategy is approached. The previous discussion has shown that the slope of the growth function monotonically decreases as the maximum is approached. This apparent correlation between the speed of convergence and the slope of the growth function can be explained from the fact that the actual income growth rate of an agent only approximates the equilibrium growth rate of its strategy. Due to this inaccuracy, the selective advantage of an investment strategy over a variant with lower equilibrium growth rate diminishes as the difference in equilibrium growth rates decreases. So as the slope of the growth function decreases around the optimum, the selection pressure among variants decreases as well, with the important consequence that the evolutionary economy potentially never converges and never reaches equilibrium.

6.4 Policy analysis

6.4.1 Experimental setup

We use numerical simulations to determine how sensitive the three policies of Section 6.2.4 are to a particular choice of values for their free parameters (cf. Table 6.2), and how sensitive they are to a particular choice of value for the cost of renewable energy. We measure this sensitivity with regard to how effective each policy is in regulating the




Figure 6.5: Evolution of the calibrated economy

economic behavior in an evolutionary economy, which in this model means to reduce global warming and increase social welfare.

To reduce the variance of the simulation results, we replicate each simulation 1,000 times for each tested level of free parameters and cost of renewable energy. In order to obtain results that are valid for the general class of scale-free social networks with a high clustering coefficient, each replication uses a different random instance of such a network, so that the results are valid for our general class of social networks but do not depend on a specific choice of network. Also, at the start of each replication the agents are initialized with random strategies that converge during the initialization phase of 400 time steps. During the following 400 time steps, the main experimental phase, the policy is applied. We report the average value at time $t = 800$ of three key statistics: global temperature, average income, and average income growth rate. We also report the average energy allocation at $t = 800$.

Figure 6.5 shows the evolution of the three key statistics when no policy is applied, for an additional cost of renewable energy $c = 0\%$ and $c = 100\%$. These are the two systems against which the policies are evaluated. For each policy and for each parameter scan, the graph will include the same statistic for a system without policy.

6.4.2 Evaluating the first best policy, a tax on fossil energy investment

When fossil energy and renewable energy are perfect substitutes, investment in the more cost-effective energy sector generates a higher income growth rate for an investing agent. A rational agent is expected to use the strategy with the highest equilibrium growth rate and to invest exclusively in the more cost efficient energy sector, even if the difference is very small: if $\tau < c$, a rational agent invests only in fossil energy. If $\tau > c$, it invests only in renewable energy. If $\tau = c$, it is indifferent between the two energy sectors. This does not hold for evolutionary agents, which converge on a strategy only if there is sufficient selection pressure. When the cost difference between fossil and renewable energy is small, the resulting difference in the equilibrium income growth rate is also small. Since the actual income growth rate only approximates the equilibrium growth rate to a certain degree, small difference in equilibrium growth rate are harder to detect than large ones.

Figure 6.6 shows the economic effect at $t = 800$ of different levels of a tax on fossil energy investment. Here the additional cost of renewable energy is $c = 100\%$. The average energy allocation is a smooth function of the cost difference of the two energy sectors, and hence of the slope of equation 6.32. The curves can best be described as two symmetric sigmoids that cross each other at about $\tau = 125\%$. In other words, the evolutionary agents are indifferent between the two energy sectors at a tax level of $\tau = 125\%$. For a rational agent as described above, we would observe two step functions that cross each other at the point where both energy sectors carry the same cost, i.e., $\tau = 100\%$. Figure 6.7 allows for a similar observation for different levels of the cost of renewable energy when the tax is $\tau = 100\%$. Here the two curves (approximate sigmoids) of the energy allocation cross each other at a tax level of $c = 80\%$. Again, for a rational agent as described above, we would observe two step functions that cross each other when the additional cost of renewable energy is equal to the tax, i.e., at $c = 100\%$.

That the observed crossover points deviate significantly from the point $\tau = c$ where rational agents are indifferent is due to a particular type of lock-in or memory effect of an evolutionary system. During the initialization phase no policy is applied and due to its selective advantage the agents converge on a strategy that allocates energy investments to fossil energy. During the main experimental phase a tax on fossil energy investment changes the selective advantage in favor of renewable energy. As the agents move towards the new optimum, the slope of the growth function decreases to the point that the selection pressure becomes insignificant. For all practical purposes, the convergence comes to a halt somewhere between the old and the new optimum.

In both Figure 6.6 and 6.7 the increase in global temperature generally reflects the investment in fossil energy, and this increase is in turn reflected in the average income growth rate and the average income. All statistics change monotonically as a function of $\tau - c$. The higher the tax on fossil energy investment, the lower the global temperature and the higher the average income growth rate and the average income.

6.4.3 Evaluating the prizes policy

In a model of rational expectations, a prize that is awarded only to selected agents introduces complex social dynamics that can be highly sensitive to initial conditions or, worse, intractable (Challet et al., 2005). In this evolutionary model, agents do not make

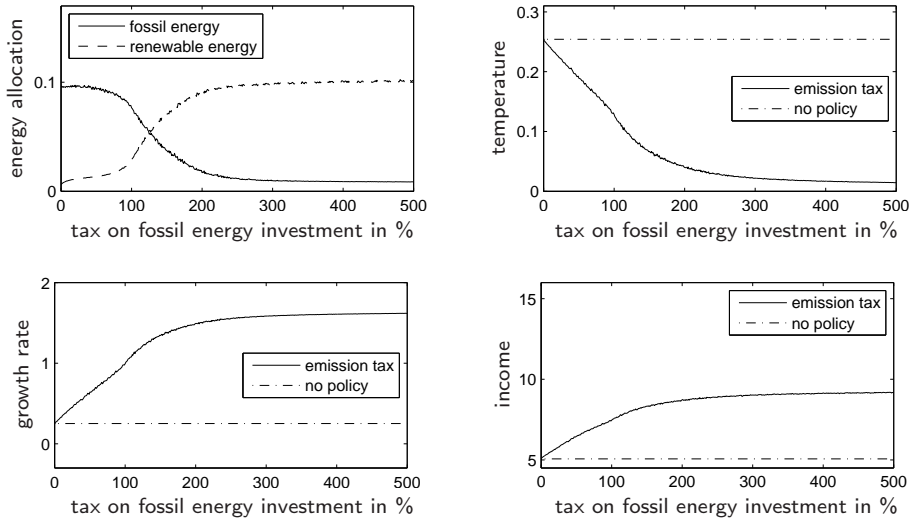


Figure 6.6: Effect of a tax on fossil energy investment for different tax levels, at $t = 800$. The additional cost of renewable energy is $c = 100\%$.




Figure 6.7: Effect of a tax on fossil energy investment for different levels of cost of renewable energy, at $t = 800$. The tax on fossil energy investment is $\tau = 100\%$. Results are averaged over 1,000 simulations for every cost increment of 10 percent points.

choices in anticipation of a prize. Instead, they choose to imitate an agent after a prize is given, based on relative welfare. The evolutionary impact of a prize is a simple function of its effect on the relative welfare.

Figure 6.8 shows how the economic impact of a prize varies with the number q of rewarded individuals. Here the income tax that finances the prize is $\varepsilon = 1\%$ and the additional cost of renewable energy is $c = 100\%$. When $q = 3$ agents are rewarded, both the average investment in fossil fuel and the global temperature are minimized, and average growth and income are maximized. For values of $10 \leq q \leq 50$ the agents weakly prefer renewable energy, and for all values of $q \leq 160$ a significant improvement in income growth rate and income can be observed, compared to the system without policies. Very high values of q have no positive economic effect, and we conclude that the evolutionary system is not showing the same positive response to selective punishment as it shows to selective reward.

Figure 6.9 shows the policy effect for different levels of income tax, for an additional cost of renewable energy $c = 100\%$ and $q = 3$ rewarded agents. Investments in renewable energy increase monotonically as the tax increases, and the global temperature decreases. The positive welfare effect however peaks at a tax of 3% and declines for higher tax levels. Figure 6.10 shows how the policy effect varies with the cost of renewable energy, for $q = 3$ rewarded agents and an environmental tax on income $\varepsilon = 1\%$. With the chosen parameters the policy proves to be effective for an additional cost of renewable energy of up to 100%. While the global temperature and the average income growth rate are positively affected even by higher cost levels, average income approaches that of the system without policy.

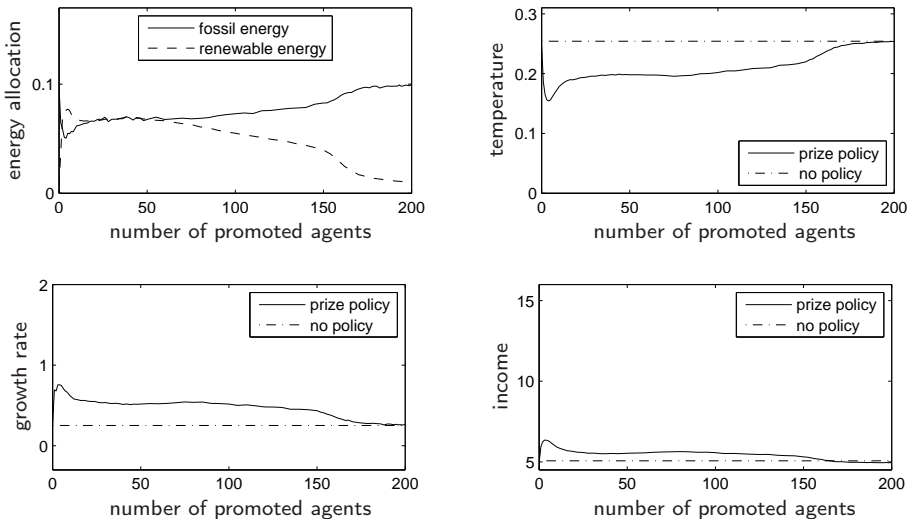


Figure 6.8: Effect of the *prizes* policy for different numbers of promoted agents, at $t = 800$. The additional cost of renewable energy is $c = 100\%$. The environmental tax on income is $\varepsilon = 1\%$

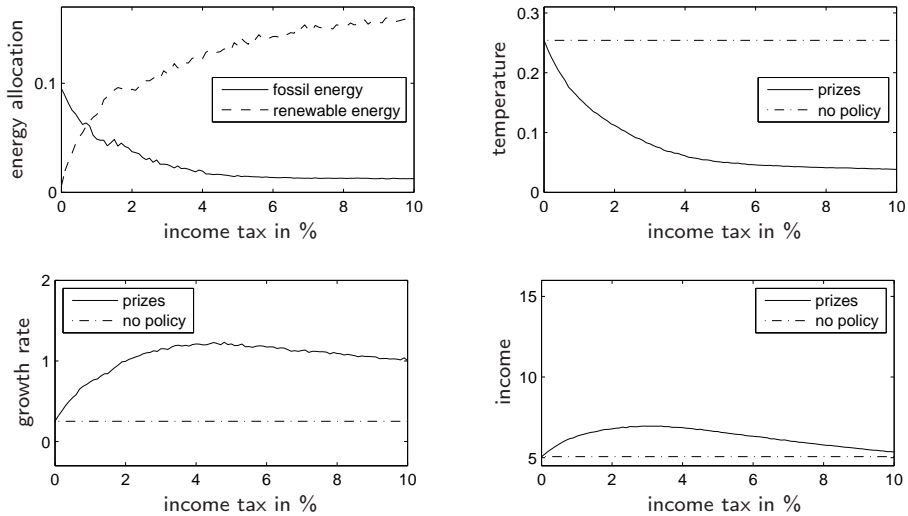


Figure 6.9: Effect of the *prizes* policy for different levels of the environmental tax on income, at $t = 800$. The additional cost of renewable energy is $c = 100\%$, the number of promoted agents is $q = 3$.




Figure 6.10: Effect of the *advertisement* policy for different levels of cost of renewable energy, at $t = 800$. The number of promoted agents is $q = 3$, the environmental tax on income is $\varepsilon = 1\%$.

6.4.4 Evaluating the *advertisement* policy

The social effect of *advertisement* can not be understood correctly without the concept of evolutionary fitness. No money is being transferred and there is no increase in information. All that is changed is the number of other agents that consider an advertised agent for imitation.

Figure 6.11 shows how the economic impact of *advertisement* varies with the number q of advertised agents. The additional cost of renewable energy is $c = 0\%$. A broad range of values for the number q of advertised agents proves to be effective, peaking in the region of ten to forty agents, and decreasing slowly as the maximum of $q = 200$ is approached, at which point the network is fully connected. Figure 6.12 shows how the effect of the advertisement policy varies with the cost of renewable energy. The number of advertised agents is $q = 10$. The policy proves to be effective only up to an additional cost of renewable energy of 1%. Beyond this point global temperature and average income approach the levels without policy, and the income growth rate becomes even lower than without policy. In other words, *advertisement* is only effective when the slope of the growth function is small (cf. equation 6.32) and the selection pressure to invest in the more cost efficient fossil energy sector is negligible.

6.5 Conclusions

An agent-based simulation of an economic process facilitates the study of climate policies under conditions that are difficult if not impossible to study in equilibrium type of models with representative and rational agents. The agent-based approach describes agents that are heterogeneous in their strategies and assets and reflect bounded rationality. This allows for the implementation and study of selective policies that treat agents differently depending on their behavior. The particular form of the agent-based model employed here is an evolutionary model of strategy formation in a social network.

The approach was applied to model investment choices by individual agents in general, fossil energy, and renewable energy capital, as part of a simple economic model with global warming feedback. Use of fossil energy is the individually optimal strategy, but causes global warming and a decline in social welfare that calls for a climate policy. As there is no central authority to enforce a climate policy, compliance is a problem. A social-evolutionary module describes selective imitation and random mutation of investment strategies. The probability that an agent is imitated depends on relative welfare and social connectivity, with relative welfare measured by the relative growth rate of individual income in the individual's (peer) network.

How an investment strategy translates into relative welfare as indicated by the income growth rate is described by a growth function. In this economic model the growth function is a concave function with a single optimum, in principle an easy optimization problem for any type of learning algorithm. However, as the strategies converge on the optimum, selection pressure decreases, the convergence process slows down, and the evolutionary economy potentially never reaches equilibrium. The growth function shows further that global warming has no effect on relative welfare, and the evolutionary agents therefore can not be expected to choose the globally over the individually optimal strategy.

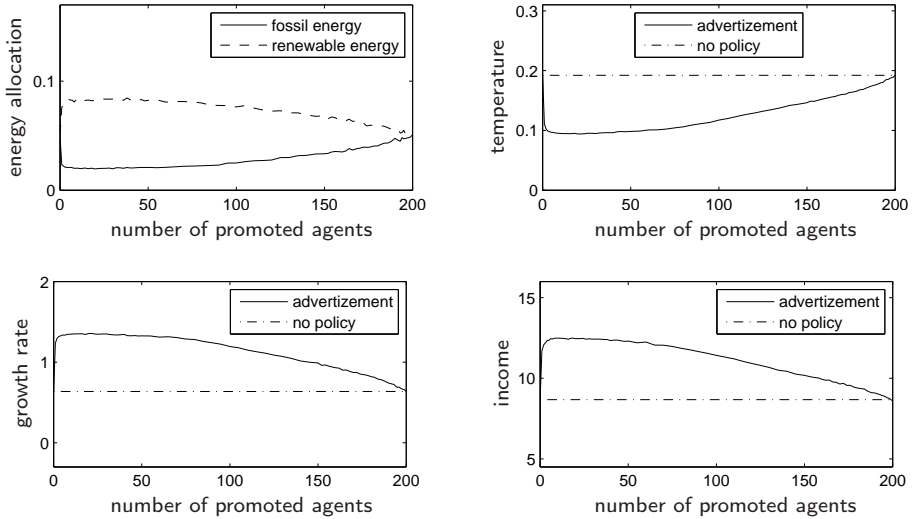


Figure 6.11: Effect of the *advertisement* policy for different numbers of promoted agents, at $t = 800$. The additional cost of renewable energy is $c = 0\%$.




Figure 6.12: Effect of policy 3 for different levels of cost of renewable energy, at $t = 800$. The number of promoted agents is $q = 10$.

Two selective policies were formulated that take heterogeneity of the strategies and of the social connectivity of individual agents into account. They influence the evolutionary formation of strategies by increasing the probability of desirable strategies to be imitated. Numerical simulations compared both policies with that of a standard regulatory tax on fossil energy investment, measuring how effective each policy is in reducing global warming and increasing social welfare. One selective policy, *prizes*, regulates relative welfare positions and causes agents with a desirable strategy to be ranked higher by their neighbors. The other selective policy, *advertisement*, regulates social visibility so that agents with a desirable strategy are seen by more agents. With regard to effectiveness, the regulatory tax on fossil energy investment depends on the compliance of the big polluters. *Prizes* depends on the compliance of at least some agents that pay into the environmental fund, for example those agents that suffer most from global warming. *Advertisement* does not depend on enforcement.

Both *prizes* and the tax on fossil energy investment were found to be effective over a wide range of values for the additional cost of renewable energy, with a gradual decrease in effectiveness as this cost increases. Numerical evaluation of the tax on fossil energy investment has shown that due to lock-in, the tax level at which evolutionary agents are indifferent between the two energy investment sectors is significantly higher than the tax level at which their costs are equal. This can be seen to reflect a tax on a lock-in externality. *Prizes* have shown that an evolutionary system is far less responsive to selective punishment than to a prize. *Advertisement* only works well when the cost difference between the two energy sectors is very small and the selection pressure to invest in fossil energy is very low.

The evolution of economic strategies and the dynamics of global warming are far more complex than expressed here, but one may expect that selective policies have the same qualitative effect. The effect of *prizes* is similar to that of a regulatory tax on fossil energy investment, but depends less on the compliance of the big polluters. When the costs of fossil and renewable energy are nearly equal, economic fairs and conferences, scholarships, awards for outstanding contributions, and publication of informative material represent relatively inexpensive policy tools that do not depend on an enforcing authority and that can have a significant effect on combating global warming.

SUMMARY AND CONCLUSIONS

This doctoral thesis combines methods from several scientific disciplines in order to allow for an evolutionary agent-based policy analysis in dynamic environments. We started with a fundamental study on how to build a simple and robust model of economic evolution. We evaluated an array of conventional evolutionary algorithms for their simplicity and robustness and proceeded to build a simple evolutionary mechanism where economic behavior evolves by imitation in a social network. We then used such a model in numeric agent-based simulations to investigate the evolutionary process under different environmental dynamics. Finally, we have formulated and evaluated environmental policies that explicitly take the evolution of economic behavior into account.

In order to build an objectively simple evolutionary mechanism, we introduced and evaluated a numerical method, Relevance Estimation and Value Calibration. It measures the minimum amount of information that is needed to tune an evolutionary algorithm such that it reaches a desired level of performance, and how this performance depends on the correct tuning of individual algorithm parameters. In order to do so it uses probability distributions over parameter values during the tuning process. Probability distributions are well suited to measure information. The method can in principle be applied to fields other than evolutionary computing.

We applied the method to an array of conventional evolutionary algorithms, and found that the need for tuning is distributed in a highly skewed way over the different algorithm components. Typically the tuning of the mutation operators—which maintain diversity—has the highest impact on algorithm performance. This general result was confirmed for models of economic evolution, in simulations where the agents have to adapt under the complex dynamics of climatic and technological change. The same simulations also revealed that when extra detail is added to an evolutionary model, tuning of the model becomes less effective, and the capability of the agents to adapt is re-

duced. This is a further argument to keep the model of imitation as simple as possible.

We proceeded to build a simple evolutionary mechanism with only one free parameter for the diversity of strategies, and studied the welfare effect of changing this parameter under different environmental dynamics. The dynamics were defined with regard to two widely recognized and easily measurable aspects of environmental change, namely how gradual and how frequent it occurs. Extensive numeric simulations identified the level of diversity that leads to a welfare distribution which is socially optimal in terms of constant relative risk aversion. This optimum level is different for different environmental dynamics, and based on our simulations we formulated policy advice on the socially optimal level of diversity when the environmental dynamics are unknown. It emerged that in general, a higher degree of risk aversion calls for a higher degree of diversity. When the precise nature of the environmental dynamics is known, optimizing the diversity to the particular dynamics allows for a significant increase in social welfare. This possibility to improve social welfare by optimizing social imitation to a particular environmental dynamics or to a particular degree of risk aversion constitutes a new opportunity for public intervention that has not previously been recognized.

Finally we applied our methods to a simple model of global warming where the policy maker wants to encourage the agents to replace fossil energy, which has a negative environmental impact, by renewable energy, which is environmentally neutral yet less cost-efficient. Numerical evaluation of a regulatory tax on investments in fossil energy revealed that due to lock in, the tax level for which the evolutionary agents are indifferent between the two technologies differs significantly from what can be concluded from a model with rational and representative agents. As a consequence, the level of such a regulatory tax has to be significantly higher if agents are to be convinced to abandon fossil fuels for good. We designed and evaluated two novel public policies—*prizes* and *advertisement*—that selectively increase the probability of environmentally friendly strategies to be imitated. Both policies are easier to enforce than a regulatory tax, which depends on whether the worst offenders can be persuaded to comply—a difficult task if they have to pay the highest tax. With *prizes*, policy enforcement has shifted to finding enough donors that finance a prize for the best behaving agent. Numerical evaluation showed that the effect of *prizes* on welfare and global warming is similar to that of a regulatory tax. *Advertisement* proved to work well only when the cost difference between fossil and renewable energy is small, but has the unique advantage that it does not depend on a central authority to enforce it.

We conclude that a model of economic evolution can be designed that is simple and robust in an objective way, and that the simple evolutionary mechanism of such a model is sufficient to allow a community of agents to adapt well to different environmental dynamics, even when their rational capabilities are bounded and their information is limited. Despite the inherent randomness of a simulated evolutionary process, robust results can be obtained that are valid over a large number of different environmental and social conditions, pointing to their general validity outside the tested conditions.

Such models do not only lead to different predictions with regard to established policy tools like a regulatory tax, but they open the door for new policy instruments that regulate the selective advantage of economic behavior.

SYMBOLS

Symbol	Meaning	Chapter
a, b	individual economic agents	4–6
α	production coefficient of the Cobb-Douglas production function	4, 6
β	scaling factor of the production function	5, 6
C	consumption	4
c	additional cost of renewable energy	6
\mathcal{D}	distribution over parameter values	2, 3
d	dynamics that change the production function	5
δ	rate of depreciation for capital and technology	4–6
E	energy investment sector	6
\mathcal{E}	net effect of a strategy on growth;	5
ε	revenue tax on income to fund climate policy	6
	fund financed by the environmental tax	6
F	capital accumulation in the fossil energy sector	4, 6
\mathcal{F}	fitness of a strategy/agent	6
f	P [agent mutates its strategy]	4
ϕ	breakdown fraction of greenhouse gases	6
G	level of atmospheric greenhouse gases	6
g	P [agent imitates]	4
γ	income growth rate of an agent	5, 6
h	Shannon entropy of a distribution;	2, 3
	threshold rank of imitated peers	4
i	index of a capital sector	4–6
K	capital accumulated in an investment sector	4–6
k	number of parameters;	2
	number of peers of an agent in the social network	4–6
L	technology	4
m	number of parameter vectors that form the REVAC table;	2
	number of renewable energy sectors	4
N	neighbors of an agent in the social network	4–6
n	number of parameter values that define a distribution;	2
	number of capital sectors	4, 5
P	population of economic agents	4–6
p	price of investment;	5
	probability that an agent is reached by advertisement	6
π	production coefficient of the Cobb-Douglas production function	5

Q	income without global warming	6
q	number of agents that receive a prize or are advertised	6
R	capital accumulation in the renewable energy sector	4, 6
r	rank of a strategy in a group of peers	4
ρ_f	threshold rank for mutation	4
ρ_g	threshold rank for imitation	4
s	investment strategy of an agent	4–6
σ	variance of the mutation operator or diversity control parameter	4–6
t	unit of time (usually financial quarter) in the discrete time model	4–6
τ	regulatory tax on investments in fossil energy	6
ν	sensitivity of economic growth to G	6
W	welfare of a population	4
w	amount of smoothing applied to a distribution; imitation weight	2 4
\vec{x}	vector of parameter values	2
Y	income of an agent	4–6
z	learning rate	4
ζ	factor that combines several monotonic transformations of the production function	6

REFERENCES

- Akaike, H., 1973. Information theory and an extension of the maximum likelihood principle. In: 2nd Int. Symp. Inf. Theory, ISIT. Tsahkadsor, Armenian SSR, pp. 267–281.
- Albert, R., Barabási, A.-L., 2002. Statistical mechanics of complex networks. *Rev. Modern Phys.* 74 (1), 47–97.
- Barabási, A.-L., Albert, R., 1999. Emergence of scaling in random networks. *Science* 286 (5439), 509–512.
- Barron, A. R., Rissanen, J. J., Yu, B., 1998. The minimum description length principle in coding and modeling. *IEEE Trans. Inform. Theory* 44 (6), 2743–2760.
- Bartz-Beielstein, T., Lasarczyk, C. W. G., Preuss, M., 2005. Sequential parameter optimization. In: *IEEE Congr. Evol. Comput., CEC'05*. IEEE, Edinburgh, UK, pp. 773–780.
- Beer, G., 1980. The Cobb-Douglas Production Function. *Math. Mag.* 53 (1), 44–48.
- Bergstrom, T. C., 2002. Evolution of Social Behavior: Individual and Group Selection. *J. Econ. Perspect.* 16 (2), 67–88.
- Binmore, K., 1994. *Game Theory and the Social Contract. Vol.1 - Playing Fair, and Vol.2 - Just Playing*. MIT Press, Cambridge, MA.
- Birattari, M., 2004. The problem of tuning metaheuristics as seen from a machine learning perspective. Ph.D. thesis, Université Libre de Bruxelles.
- Box, G. E. P., Wilson, K. B., 1951. On the experimental attainment of optimum conditions. *J. R. Stat. Soc. B* 13 (1), 1–45.
- Boyd, R., Richerson, P. J., 1985. *Culture and the Evolutionary Process*. University of Chicago Press, Chicago.
- Camerer, C. F., Loewenstein, G., Rabin, M. (Eds.), 2003. *Advances in Behavioral Economics*. Princeton University Press, Princeton, NJ.
- Challet, D., Marsili, M., Zhang, Y.-C., 2005. *Minority Games: Interacting Agents in Financial Markets*. Oxford University Press, Oxford.
- Czarn, A., MacNish, C., Vijayan, K., Turlach, B., Gupta, R., 2004. Statistical Exploratory Analysis of Genetic Algorithms. *IEEE Trans. Evol. Comput.* 8 (4), 405–421.

- de Landgraaf, W. A., Eiben, A. E., Nannen, V., 2007. Parameter calibration using meta-algorithms. In: IEEE Congr. Evol. Comput., CEC'07. IEEE, Singapore, pp. 71–78.
- Dopfer, K. (Ed.), 2005. *The Evolutionary Foundation of Economics*. Cambridge University Press, Cambridge, England.
- Dorogovtsev, S. N., Mendes, J. F. F., 2002. Evolution of networks. *Adv. Phys.* 51 (4), 1079–1187.
- Eiben, A. E., Hinterding, R., Michalewicz, Z., 1999. Parameter control in evolutionary algorithms. *IEEE Trans. Evol. Comput.* 3 (2), 124–141.
- Eiben, A. E., Michalewicz, Z., Schoenauer, M., Smith, J. E., 2007. Parameter control in evolutionary algorithms. In: Lobo, F. G., Lima, C. F., Michalewicz, Z. (Eds.), *Parameter Setting in Evolutionary Algorithms*. Studies in Computational Intelligence. Springer, Berlin / Heidelberg, pp. 19–46.
- Eiben, A. E., Schut, M. C., de Wilde, A. R., 2006. Is self-adaptation of selection pressure and population size possible?—A case study. In: *Parallel Probl. Solving Nat.*, PPSN IX. Springer, Reykjavik, Iceland, pp. 900–909.
- Eiben, A. E., Smith, J. E., 2003. *Introduction to Evolutionary Computing*. Springer, Berlin / Heidelberg.
- Epstein, J. M., Axtell, R., 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press, Cambridge, MA.
- Erdős, P., Rényi, A., 1959. On random graphs I. *Publ. Math. Debrecen* 6, 290–297.
- Ferrer-i Carbonell, A., 2005. Income and well-being: an empirical analysis of the comparison income effect. *J. Public Econ.* 89 (5-6), 997–1019.
- François, O., Lavergne, C., 2001. Design of evolutionary algorithms—a statistical perspective. *IEEE Trans. Evol. Comput.* 5 (2), 129–148.
- Frank, R. H., 1987. *Choosing the Right Pond: Human Behavior and the Quest for Status*. Oxford University Press, Oxford.
- Freisleben, B., Hartfelder, M., 1993. Optimization of genetic algorithms by genetic algorithms. In: Albrecht, R. F., Reeves, C. R., Steele, N. C. (Eds.), *Artificial Neural Networks and Genetic Algorithms*. Springer, Berlin / Heidelberg, pp. 392–399.
- Friedman, D., 1998. On economic applications of evolutionary game theory. *J. Evol. Econ.* 8 (1), 15–43.
- Gallagher, M., Yuan, B., 2006. A general-purpose tunable landscape generator. *IEEE Trans. Evol. Comput.* 10 (5), 590–603.
- Galor, O., Moav, O., 2002. Natural selection and the origin of economic growth. *Q. J. Econ.* 117 (4), 1133–1191.

- Goldberg, D. E., 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- Grefenstette, J. J., 1986. Optimization of control parameters for genetic algorithms. *IEEE Trans. Syst. Man and Cybernet* 16 (1), 122–128.
- Grünwald, P. D., 2007. *The Minimum Description Length Principle*. MIT press, Cambridge, MA.
- Henrich, J., 2004. Cultural group selection, coevolutionary processes and large-scale cooperation. *J. Econ. Behav. Organ.* 53 (1), 3–35.
- Hofbauer, J., Sigmund, K., 2003. Evolutionary game dynamics. *Bull. Amer. Math. Soc.* 40 (4), 479.
- Janssen, M. A., Anderies, J. M., Walker, B. H., 2004. Robust strategies for managing rangelands with multiple stable attractors. *J. Environ. Econ. Manag.* 47 (1), 140–162.
- Janssen, M. A., Ostrom, E., 2006. Governing social-ecological systems. In: Tesfatsion, L., Judd, K. L. (Eds.), *Handbook of Computational Economics*. Elsevier, pp. 1465–1509.
- Jong, K. A. D., 1975. An analysis of the behaviour of a class of genetic adaptive systems. Ph.D. thesis, University of Michigan.
- Jong, K. A. D., 2006. *Evolutionary Computation: A Unified Approach*. MIT Press, Cambridge, MA.
- Juglar, C., 1863. *Crises Commerciales*. Imprimerie de Veuve Berger-Levrault, Strasbourg.
- Kahneman, D., Slovic, P., Tversky, A., 1982. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, England.
- Kelly, D. L., Kolstad, C. D., 1999. Integrated assessment models for climate change control. In: Folmer, H., Tietenberg, T. (Eds.), *International yearbook of environmental and resource economics 1999/2000: A survey of current issues*. Edward Elgar Cheltenham, UK, pp. 171–197.
- Kirman, A. P., 1992. Whom or what does the representative individual represent? *J. Econ. Perspect.* 6 (2), 117–136.
- Kondratiev, N. D., 1935. The Long Waves in Economic Life. *Rev. Econ. Stat.* 17 (6), 105–115.
- Lehtinen, A., Kuorikoski, J., 2007. Computing the Perfect Model: Why Do Economists Shun Simulation?*. *Philos. Sci.* 74 (3), 304–329.
- Levy, H., Levy, M., Solomon, S., 2000. *Microscopic Simulation of Financial Markets: From Investor Behavior to Phenomena*. Academic Press, New York.
- Lieberman, E., Hauert, C., Nowak, M. A., 2005. Evolutionary dynamics on graphs. *Nature* 433 (7023), 312–316.

- Lobo, F. G., Lima, C. F., 2006. On the Utility of the Multimodal Problem Generator for Assessing the Performance of Evolutionary Algorithms. In: Proc. Genet. Evol. Comput. Conf., GECCO'06. ACM, Seattle, Washington, pp. 1233–1240.
- Lobo, F. G., Lima, C. F., Michalewicz, Z. (Eds.), 2007. Parameter Setting in Evolutionary Algorithms. Studies in Computational Intelligence. Springer, Berlin / Heidelberg.
- Louzoun, Y., Shnerb, N. M., Solomon, S., 2007. Microscopic noise, adaptation and survival in hostile environments. *Eur. Phys. J. B* 56 (2), 141–148.
- Luke, S., 2007. The ECJ Evolutionary Computation System. Tech. rep., Department of Computer Science George Mason University.
- Manne, A. S., 1992. Buying Greenhouse Insurance: The Economic Costs of CO2 Emission Limits. MIT Press, Cambridge, MA.
- Milgram, S., 1967. The small world problem. *Psychol. Today* 2 (1), 60–67.
- Mirowski, P., 2007. Markets come to bits: Evolution, computation and markomata in economic science. *J. Econ. Behav. Organ.* 63 (2), 209–242.
- Mühlenbein, H., 1997. The equation for response to selection and its use for prediction. *Evol. Comput.* 5 (3), 303–346.
- Mühlenbein, H., Höns, R., 2005. The estimation of distributions and the minimum relative entropy principle. *Evol. Comput.* 13 (1), 1–27.
- Munro, A., 1997. Economics and biological evolution. *Environ. Resour. Econ.* 9 (4), 429–449.
- Nannen, V., Eiben, A. E., 2006. A method for parameter calibration and relevance estimation in evolutionary algorithms. In: Proc. Genet. Evol. Comput. Conf., GECCO'06. ACM, Seattle, Washington, pp. 183–190.
- Nannen, V., Eiben, A. E., 2007a. Efficient relevance estimation and value calibration of evolutionary algorithm parameters. In: IEEE Congr. Evol. Comput., CEC'07. IEEE, Singapore, pp. 103–110.
- Nannen, V., Eiben, A. E., 2007b. Relevance Estimation and Value Calibration of Evolutionary Algorithm Parameters. In: Proc. Int. Jt. Conf. Artif. Int., IJCA'07. AAAI Press, Hyderabad, India, pp. 975–980.
- Nannen, V., Smit, S. K., Eiben, A. E., 2008. Costs and Benefits of Tuning Parameters of Evolutionary Algorithms. In: Parallel Probl. Solving Nat., PPSN X. Springer, Dortmund, Germany, pp. 528–538.
- Nannen, V., van den Bergh, J. C. J. M., 2010. Policy instruments for evolution of bounded rationality: Application to climate–energy problems. *Technol. Forecast. Soc. Change* 77 (1), 76–93.

- Nannen, V., van den Bergh, J. C. J. M., Eiben, A. E., 2013. Impact of environmental dynamics on economic evolution: A stylized agent-based policy analysis. *Technol. Forecast. Soc. Change* 80 (2), 329–350.
- Nelson, R. R., Winter, S. G., 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.
- Noailly, J., 2008. Coevolution of economic and ecological systems. *J. Evol. Econ.* 18 (1), 1–29.
- Noailly, J., van den Bergh, J. C. J. M., Withagen, C. A., 2003. Evolution of harvesting strategies: replicator and resource dynamics. *J. Evol. Econ.* 13 (2), 183–200.
- Nordhaus, W. D., 1991. To Slow or Not to Slow—The Economics of the Greenhouse Effect. *Econ. J.* 101 (407), 920–937.
- Nordhaus, W. D., 1992. An optimal transition path for controlling greenhouse gases. *Science* 258 (5086), 1315–1315.
- Nordhaus, W. D., 1994. *Managing the global commons: The economics of the greenhouse effect*. MIT Press, Cambridge, MA.
- Nordhaus, W. D., 2002. Modeling Induced Innovation in Climate Change Policy. In: Grubler, A., Nakićenović, N., Nordhaus, W. D. (Eds.), *Modeling Induced Innovation in Climate Change Policy*. Resources for the Future Press, pp. 259–290.
- Nordhaus, W. D., Yang, Z., 1996. A Regional Dynamic General-Equilibrium Model of Alternative Climate-Change Strategies. *Amer. Econ. Rev.* 86 (4), 741–765.
- Nowak, M. A., 2006. *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, Cambridge, MA.
- Oliver, I. M., Smith, D. J., Holland, J. R. C., 1987. A study of permutation crossover operators on the traveling salesman problem. In: *Proc. 2nd Int. Conf. Genet. Algorithms Genet. algorithms appl.* Cambridge, Ma, pp. 224–230.
- Ostrom, E., 2000. Collective Action and the Evolution of Social Norms. *J. Econ. Perspect.* 14 (3), 137–158.
- Peck, S. C., Teisberg, T. J., 1993. *Optimal CO2 Emissions Control with Partial and Full World-wide Cooperation: An Analysis Using CETA*. Electric Power Research Institute, Palo Alto, Calif.
- Popp, D., 2004. ENTICE: endogenous technological change in the DICE model of global warming. *J. Environ. Econ. Manag.* 48 (1), 742–768.
- Rissanen, J. J., 1978. Modeling by shortest data description. *Automatica* 14 (5), 465–471.
- Rudolph, G., 1992. On correlated mutations in evolution strategies. In: *Parallel Probl. Solving Nat., PPSN II*. Springer, Brussels, Belgium, pp. 107–116.

- Samples, M. E., Daida, J. M., Byom, M., Pizzimenti, M., 2005. Parameter sweeps for exploring GP parameters. In: Proc. Genet. Evol. Comput. Conf., GECCO'05. ACM, Washington, DC, pp. 1791–1792.
- Schaffer, J., Caruana, R., Eshelman, L., Das, R., 1989. A study of control parameters affecting online performance of genetic algorithms for function optimization. In: Proceedings of the third international conference on genetic algorithms. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 51–60.
- Sethi, R., Somanathan, E., 1996. The Evolution of Social Norms in Common Property Resource Use. *Amer. Econ. Rev.* 86 (4), 766–788.
- Shannon, C. E., 1948. A mathematical theory of communication. *Bell Syst. Tech. J.* 27, 379–423– 623–656.
- Shnerb, N. M., Louzoun, Y., Bettelheim, E., Solomon, S., 2000. The importance of being discrete: life always wins on the surface. *Proceedings of the National Academy of Sciences* 97 (19), 10322–10324.
- Spears, W. M., 2000. *Evolutionary Algorithms: The Role of Mutation and Recombination*. Springer, Berlin / Heidelberg.
- Stainforth, D. A., Aina, T., Christensen, C., Collins, M., Faull, N., Frame, D. J., Kettleborough, J. A., Knight, S., Martin, A., Murphy, J. M., 2005. Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature* 433 (7024), 403–406.
- Taguchi, G., Wu, Y., 1980. *Introduction to Off-Line Quality Control*. Central Japan Quality Control Association, Nagoya, Japan.
- Tesfatsion, L., 2006. Agent-based computational economics: A constructive approach to economic theory. In: Judd, K. L. (Ed.), *Handbook of Computational Economics*. Elsevier, pp. 831–880.
- Tol, R. S. J., 1995. The climate fund: Sensitivity, uncertainty, and robustness analysis. Tech. Rep. W95/02, Institute for Environmental Studies, Vrije Universiteit, Amsterdam.
- Tomasini, M., 2005. *Spatially Structured Evolutionary Algorithms*. Springer, Berlin / Heidelberg.
- van den Bergh, J. C. J. M., 2004a. Evolutionary modelling. In: Proops, J., Safonov, P. (Eds.), *Modelling in Ecological Economics*. Edward Elgar Publishing, Cheltenham, pp. 9–35.
- van den Bergh, J. C. J. M., 2004b. Optimal climate policy is a utopia: from quantitative to qualitative cost-benefit analysis. *Ecol. Econ.* 48 (4), 385–393.
- van den Bergh, J. C. J. M., 2007. Evolutionary thinking in environmental economics. *J. Evol. Econ.* 17 (5), 521–549.

- Vereshchagin, N. K., Vitányi, P. M. B., 2002. Kolmogorov's Structure Functions with an Application to the Foundations of Model Selection. In: 43rd Annu. IEEE Symp. Found. Comput. Sci., FOCS'02. IEEE, Vancouver, BC, pp. 751–760.
- Watts, D. J., Strogatz, S. H., 1998. Collective dynamics of 'small-world' networks. *Nature* 393 (6684), 440–442.
- Wegner, G., Pelikan, P., 2003. Introduction: evolutionary thinking on economic policy. In: Pelikan, P., Wegner, G. (Eds.), *The Evolutionary Analysis of Economic Policy (New Horizons in Institutional and Evolutionary Economics)*. Edward Elgar Publishing, Cheltenham, pp. 1–14.
- Wilhite, A., 2006. Economic Activity on Fixed Networks. In: Tesfatsion, L., Judd, K. L. (Eds.), *Handbook of Computational Economics*. Elsevier, pp. 1013–1045.
- Witt, U., 2008. What is specific about evolutionary economics? *J. Evol. Econ.* 18 (5), 547–575.