CHAPTER 1

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; Hofbauer 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 famililiy 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

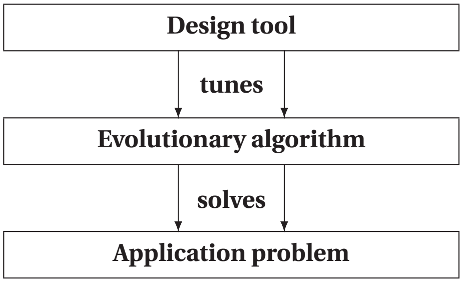

ological framework developed, the Minimum Description Length principle (Grünwald, 2007). This methodology uses information theory to measure the simplicity of a statistical hypothesis. We will take a similar approach and use information theory to measure the simplicity of an evolutionary mechanism.

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 advise 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.

References

Albert, R., Barabási, A.-L., 2002. Statistical mechanics of complex networks. Rev. Modern Phys. 74 (1), 47–97.

Boyd, R., Richerson, P. J., 1985. Culter 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.

Dorogovtsev, S. N., Mendes, J. F. F., 2002. Evolution of networks. Adv. Phys. 51 (4), 1079–1187.

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.

Hofbauer, J., Sigmund, K., 2003. Evolutionary game dynamics. Bull. Amer. Math. Soc. 40 (4), 479.

Kahneman, D., Slovic, P., Tversky, A., 1982. Judgment under Uncertainty: Heuristics and Biases. Cambridge University Press, Cambridge, England.

Kirman, A. P., 1992. Whom or what does the representative individual represent? J. Econ. Perspect. 6 (2), 117–136.

Lehtinen, A., Kuorikoski, J., 2007. Computing the Perfect Model: Why Do Economists Shun Simulation?*. Philos. Sci. 74 (3), 304–329.

Lieberman, E., Hauert, C., Nowak, M. A., 2005. Evolutionary dynamics on graphs. Nature 433 (7023), 312–316.

Louzoun, Y., Shnerb, N. M., Solomon, S., 2007. Microscopic noise, adaptation and survival in hostile environments. Eur. Phys. J. B 56 (2), 141–148.

Milgram, S., 1967. The small world problem. Psychol. Today 2 (1), 60–67.

Nelson, R. R., Winter, S. G., 1982. An Evolutionary Theory of Economic Change. Harvard University Press, Cambridge, MA.

Nordhaus, W. D., 1992. An optimal transition path for controlling greenhouse gases. Science 258 (5086), 1315–1315.

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.

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.