CHAPTER 4

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 where 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 $R_1, \ldots, R_m$. 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,

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

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 deprecation 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 deprecation δ of .05. The growth functions are

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

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

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

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 deprecation 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.

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

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. $z_{R_1}, \ldots, z_{R_m} = .0001$

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

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

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

$$C_a(t) = s_{Ca}(t)\,Y_a(t). \tag{4.7}$$

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

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

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 $\mathscr{C}$ (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 bidirectional 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.

[1] In their seminal paper Watts and Strogatz (1998) define the clustering coefficient $\mathscr{C}_i$ 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 $\mathscr{C}$ of the entire network as the average clustering coefficient of the nodes of the network.

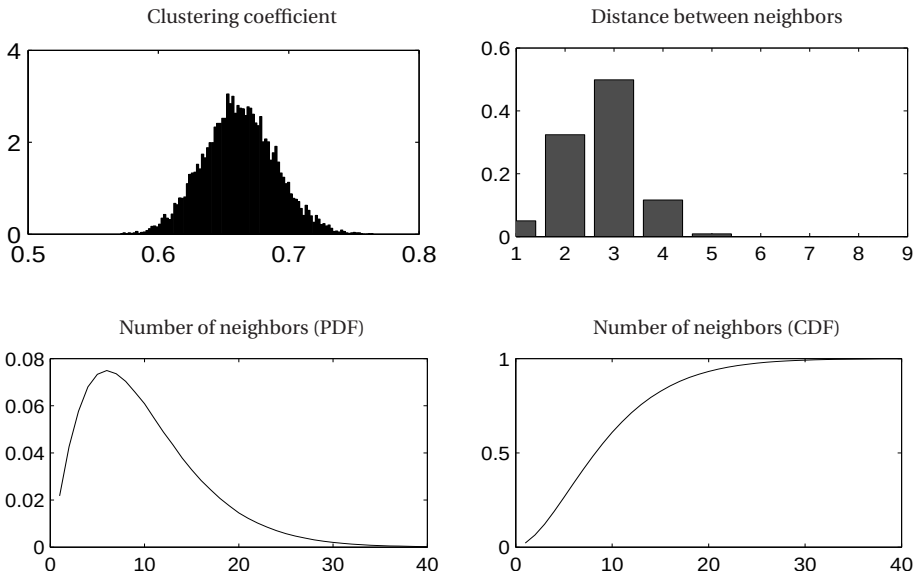

Figure 4.1: Network statistics

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,

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

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,

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

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

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

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,

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

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 $r_a(a) \le \rho_g$ chooses an agent to imitate according to

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

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

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

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

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

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,

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

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

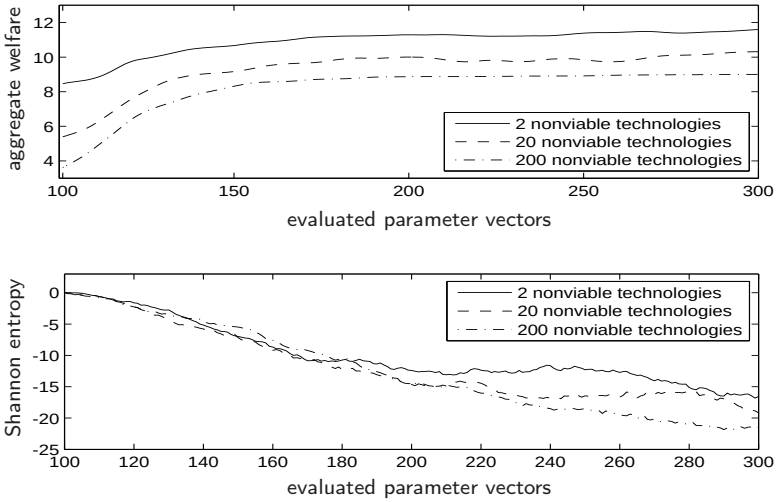

Figure 4.2: Aggregate welfare and Shannon entropy of the initial model of 13 parameters during a REVAC tuning session.

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 25^{th} and 75^{th} 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

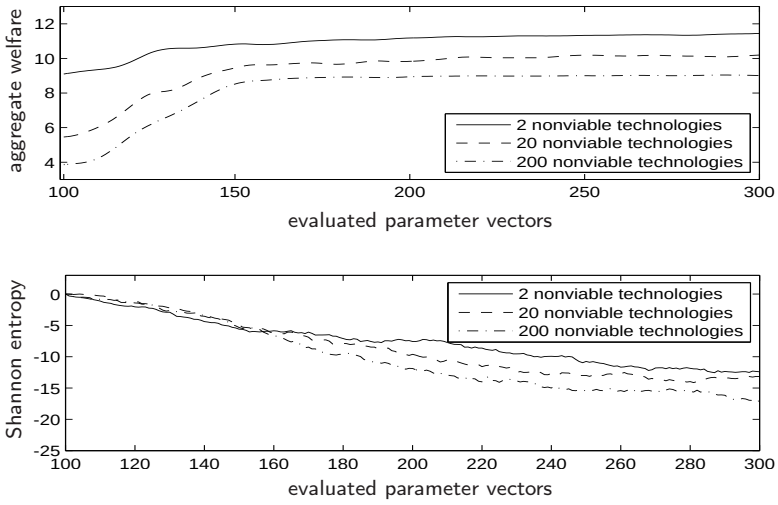

Figure 4.3: Aggregate welfare and Shannon entropy of the simplified model of 6 parameters during a REVAC tuning session.

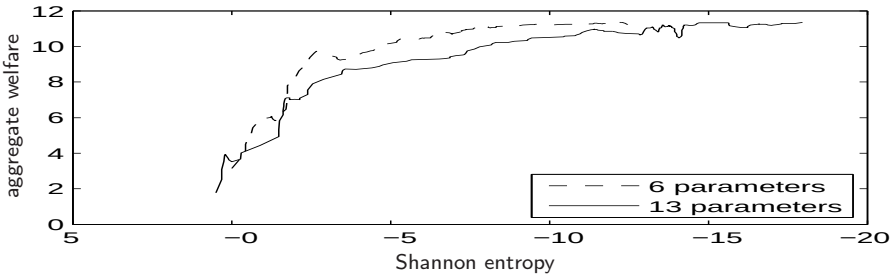

Figure 4.4: Aggregate welfare as a function of Shannon entropy.

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.

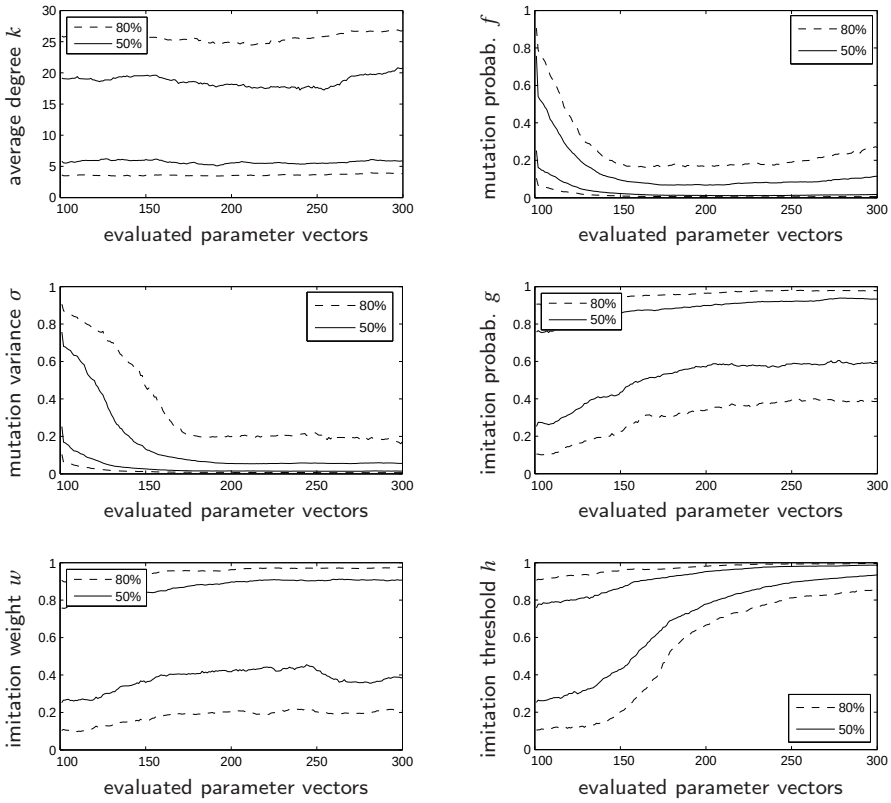

Figure 4.5: Tuning the 6 parameters of the simplified model. 50% of the density is located between the two solid lines, 80% between the dashed lines.

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.

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.

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.

Eiben, A. E., Hinterding, R., Michalewicz, Z., 1999. Parameter control in evolutionary algorithms. IEEE Trans. Evol. Comput. 3 (2), 124–141.

Eiben, A. E., Smith, J. E., 2003. Introduction to Evolutionary Computing. Springer, Berlin / Heidelberg.

Erdős, P., Rényi, A., 1959. On random graphs I. Publ. Math. Debrecen 6, 290–297.

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

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.

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.

Rissanen, J. J., 1978. Modeling by shortest data description. Automatica 14 (5), 465–471.

Tomasini, M., 2005. Spatially Structured Evolutionary Algorithms. Springer, Berlin / Heidelberg.

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.

Wilhite, A., 2006. Economic Activity on Fixed Networks. In: Tesfatsion, L., Judd, K. L. (Eds.), Handbook of Computational Economics. Elsevier, pp. 1013–1045.