EPCS 2006

Model uncertainty and party learning

Ivo Bischoff[†]

Department of Economics, University of Giessen, Licher Strasse 74,

35394 Giessen, Germany[*]

Abstract

This paper develops a computerized model of party competition in a multidimensional policy environment with model uncertainty, in which knowledge is endogenous. The model is used to simulate party learning under different institutional settings and for different forms of learning. The research-guiding question in these simulations is which factors determine the speed by which a society discovers and applies adequate policies. First results are presented.

Key words: party competition, evolutionary economics, genetic algorithms, directional learning, model uncertainty

JEL: D 72, D 83

[†] The author would like to thank Stefan Schäfer and Dorothea Reimer for helpful comments.

[*] E-mail: ivo.bischoff@wirtschaft.uni-giessen.de, Tel. ++49 641 99 22084, Fax. ++49 641 99 22089

1. Introduction

In a complex economic environment with model uncertainty, there exists a large variety of convictions as to how the economy functions and thus which policies are necessary to solve a certain problem, e.g. fight unemployment (e.g., Denzau and North, 1994; Tetlock, 1999). Adequate and inadequate policies can only be distinguished in real experiments (e.g., von Hayek, 1972; Wohlgemuth, 2002). In democracies, it is the political parties that suggest certain experiments (by proposing policy platforms) and the voters who choose between proposed experiments (by voting for one or the other party). This interaction of voters and politicians can be thought of as a dynamic process of knowledge acquisition. The essential question concerning this process is to what extent and at which speed parties learn to apply adequate policies. Using a dynamic model of party learning based on genetic algorithms, Bischoff (2005) shows that adequate economic policies are discovered if the electorate remunerates them.

This paper takes this analysis a few steps further. First, the analysis is concerned with the speed rather than the mere possibility of learning. Second, it will allow for different forms of learning, in particular directional learning. The paper will proceed as follows. Section 2 modifies the spatial theory of voting to capture the behavior of voters and political parties under model uncertainty. Section 3 introduces different forms of learning and develops a dynamic model of party learning in which knowledge is endogenous. A computerized version of this model is used to simulate the process of party learning. Among others, the impact of party motivation, yardstick competition and different forms of learning on the speed of party learning is analysed. The simulation results are presented and discussed in section 4. Section 5 concludes.

2. Elements of an analytical model of party competition in a complex environment

2.1 Voter behavior

Model uncertainty allows for a dual interpretation of heterogeneous preferences among voters (e.g., Swank, 1998). First, following the conventional economic approach, different preferences result from differences in tastes or differences in the way that voters are affected by certain policies individually (e.g., Downs, 1957; Enelow and Hinich, 1984). Second, different preferences can originate in different convictions about which policies are most adequate to solve the current economic problems. This paper assumes that the different preferences result from different convictions about the functioning of the economy, while there is a broad consensus among all voters that unemployment is the most urgent problem public policy has to solve. By this assumption, differences and tastes and the way in which voters are affected by policies individually do not influence their voting decision. Instead, voters are guided by the question which party they consider more successful at reducing unemployment.

Voters can choose between two parties A and B. In their election campaigns, both parties try to convince the voters that their policy platform is more suitable to fight unemployment than the one of the political opponent. The more suitable the individual voter considers the proposed policy platform to fight unemployment, the more likely he is to vote for this party. Other things equal, he will find the argumentation more plausible the more his own convictions have in common with the ones underlying the party's course of argumentation (e.g., Richards, 2001).

Consider a D-dimensional policy-space where each element can be scaled separately and all instruments can be applied in a number of different intensities. In this policy-space, a policy platform can be expressed as a D-dimensional vector. The vector names the intensity with with each of the D instruments is to be applied. As most of the reasoning is qualitative in nature (e.g., Forbus, 1983), these intensities are stated in rough ordinal categories (see table 1). Hereafter, eight categories are differentiated. Let η_t^i denote the policy platform which voter i in t considers adequate. The platforms offered by party A and party B are denoted η_t^A and η_t^B respectively. Other things equal, voter i will vote for party A if:

$$\overline{\eta_t^i \, \eta_t^A} \quad < \quad \overline{\eta_t^i \, \eta_t^B} \tag{1}$$

Thus, the proximity of η_t^i and η_t^A respectively η_t^B is the first determining factor of voter behavior.

Table 1 depicts a possible constellation of η_t^i, η_t^A and η_t^B. The positions are marked by i, A and B respectively. While voters and/or parties agree on the intensity with which some instruments should be applied, they disagree on others. Assuming that voter i considers all instruments equally important, $\overline{\eta_t^i \, \eta_t^A}$ is given by the geometric distance in an eight-dimensional space.

$$\overline{\eta_t^i \, \eta_t^A} \quad = \quad \sqrt{\sum_{d=1}^{D} (\eta_{t,d}^i - \eta_{t,d}^A)^2} \tag{2}$$

$\eta_{t,d}^i$ represents voter i's preferred intensity for instrument d.

Table 1

Due to the complexity of the economic system, however, neither voters nor politicians can anticipate the economic impact of different policy platforms without having seen them in action. Some voters and/or politicians support policies which prove successful while others have misperceived the functioning of the economy and support policies which turn out to be inadequate to fight unemployment. This paper assumes that there exists an ideal policy platform (H^{opt}) which, if implemented, will reduce the unemployment rate to a minimum (U_{min}). Any deviation from H^{opt} will lead to higher unemployment rates. H^{opt} is not known to voters or politicians.

This leads us to the second factor determining the voter's decision: the valence of the parties. Valence measures a party's perceived competence in dealing with a problem on which all voters have the same preferences (e.g., Ansolabere and Snyder, 2000). The degree of valence voters attribute to a party does not depend on the party's current policy platform (e.g., Groseclose, 2001). This paper assumes that the judgement of valence is based on the party's observed success to deal with the problems in the past. The lower the unemployment rate a party managed to achieve in the past, the higher its valence. The parameter K_t^A is introduced to account for the valence of party A:

$$K_t^A = f\left[\left(\frac{U^*}{U_{t-1}^A}\right)^{\kappa}, \left(\frac{U^*}{U_{t-2}^A}\right)^{\kappa}, \cdots\right] \quad (3)$$

K_A is a negative function of the average unemployment rates party A achieved in previous times of government (U_{t-1}^A, U_{t-2}^A, ….). When estimating the valence of a party, voters are assumed to use an expected unemployment rate U* as a point of reference. The parameter κ determines the degree to which differences between U* and (U_{t-1}^A, U_{t-2}^A, ….) influence the perceived valence. The judgement of valence can be thought of as an empirically grounded corrective to the plausibility judgement (e.g., Griffin and Tversky, 2002).

Combining (1) and (3), voter i will prefer party A if

$$\frac{\overline{\eta_t^i\,\eta_t^A}}{K_t^A} < \frac{\overline{\eta_t^i\,\eta_t^B}}{K_t^B} \tag{4}$$

This evaluation function is best understood when taking a voter whose η_t^i is equidistant to η_t^A and η_t^B. In the classical set-up of the spatial model of voting, he will be indifferent between the two parties because he is indifferent between their policy platforms (e.g. Enelow and Hinich, 1984). If, however, $K_t^A > K_t^B$, he accounts for the fact that, in the past, party A was more successful in fighting unemployment than party B was (e.g. Fiorina, 1981). Thus, he votes for party A. The larger κ, the stronger the valence term of both parties depends on their past economic performance and thus the more important the valence for the voters' decisions.

In order to have an attractive counterpart to H^{opt}, the following paper assumes that the distribution of η_t^i among voters is symmetric as defined by Plott (1967). Consequently, there exists a median position η_t^{Med} to which η_t^A and η_t^B would converge if $\kappa = 0$, i. e. the parties' valence did not matter (e.g., Plott, 1967; Kollman, et al., 1992). This paper assumes that the position of η_t^{Med} is not known to the political parties. Due to the very small impact of a single vote on the outcome of elections, voters have no incentives to learn by insight in the sense that they change their convictions about how the economy works. Thus, they stick to their preferred policy platform so that η_t^i and η_t^{Med} (hereafter η^{Med}) will not change systematically across time.

2.2 The behavior of politicians and parties

In this paper, politicians are assumed to strive for political power as a means of reaching higher-level aims. For altruistic and dogmatic politicians, this aim is to pursue the one policy platform which they consider most suitable to fight unemployment. They only derive positive utility from holding office if this policy is put through. Opportunistic politicians reach for political power because this supplies them with income, prestige and power. The utility of being in office does not (primarily) depend on the policy he puts through. Each party can be expected to harbor dogmatic, altruistic and opportunistic politicians.

Each party has a knowledge pool (Π_t^A respectively Π_t^B) of N different policy platforms (e.g. Edmonds, 1999). In addition to the policy platforms itself, the knowledge pool contains information on the expected utility which parties assign to choosing these platform for future election races.

At the beginning of round t, each party chooses the one policy platforms for the next election race from Π_t^A respectively Π_t^B which promises the highest expected utility for the party members. In the terms of evolutionary models of learning especially genetic algorihms, the expected utility represents the fitness of a policy platform (e.g., Riechmann, 2001: 25-26; Frank, 2003). Hereafter, let $\phi_{i,t-1}^{A}$ be the fitness of platform i at the beginning of t. To calculate the fitness, three components are necessary. $\overline{MA(V)_{i,t-1}^{A}}$ stands for the moving average of the share of votes it has attracted in previous election races (provided it was nominated). $\overline{MA(U)_{i,t-1}^{A}}$ represents the moving average of the unemployment rates platform i produced when applied. Finally, γ_A $(1 \geq \gamma_A \geq 0)$ is a weighing factor which represents the degree to which party A is short-sighted and/or opportunistic in its choice of policy platforms.

$$\phi_{i,t-1}^{A} = \gamma_A MA(V)_{i,t-1}^{A} + (1-\gamma_A)\, MA(U)_{i,t-1}^{A} \qquad (5)$$

The higher its fitness, the more likely a policy platform is to be applied. It is important to note that the fitness of a certain policy platform is a measure for the party's expectations.

The platforms chosen then compete with each other for the majority of votes. Hereafter, they will be named η_t^A respectively η_t^B. The winning party puts through the policy platform it offered in the current election race. This platform η_t determines the emerging unemployment rate U_t. The information on η_t^A, η_t^B as well as V_t^A, V_t^B and U_t is stored in a neutral knowledge pool Π_t.

From round t to t+1, both parties learn. Learning is equivalent to changing the knowledge pool Π_t^A respectively Π_t^B. Changes may concern the policy platforms which are available to party or the information about their fitness (e.g. Edmonds, 1999; Riechmann, 2001: 7-11). By this definition, updating the fitness of η_t^A and η_t^B as a consequence of the informational feedback from round t represents a first step of learning. Changes to the policy platforms contained in Π_t^A respectively Π_t^B may result from the insight that there exists a mix of policy instruments which is more successful in fighting unemployment (e.g., Brenner, 1999: 89-91). Opportunistic politicians may also be willing to adopt a new policy platform without believing that it is superior in fighting unemployment. They merely need to expect a higher number of years in office from offering another platform. Consequently, an opportunistic politician will account for the electoral success of a policy platform when looking for platforms to learn from. At the same time, the economic success is important to him as well because, once in office, the probability of being re-elected is higher the lower the unemployment rate during his tenure.[1] Different forms of learning can be applied.

---

[1] In addition, a party can learn on the aggregate level through entries of new members and the exit of former ones. For a detailed description of how parties learn, see e.g. Bischoff (2005).

1) learning by communication

Learning by communication plays an important role in politics. It leads to new policy platforms which represent a synthesis of the policy platforms discussed. Learning by communication does not require new informational feedback from the latest elections but draws on already existing policy platforms. In genetic algorithms, learning by communication is typically modelled in a crossover which represents mating in biology. Therein, two randomly chosen platforms from Π_t^A (parents) produce two potential members for Π_{t+1}^A (children). Each child represents an 8x1 vector where part of the elements stem from parent 1 and part from parent 2 (e.g. Michalewicz, 1996: 89; Riechmann, 2001: 26-27). Learning by communication in party B follows the same procedure. In the internal evaluation by the learning party, the fitness of a new policy platform cannot be lower than the fitness of the superior parent platform. This paper assumes that child platforms are assigned the fitness of the fitter parent platform multiplied by (1+lb). The factor lb ($lb \geq 0$) represents the degree to which the party members expect an increase in utility from applying the child platform instead of the fitter of the two parent platforms.

2) learning by experiment

Second, genetic algorithms incorporate learning by experiment. In the current context, experimentation corresponds to a random change in one or more elements of an existent policy platform. Mutations are the biological equivalent and the mutation rate is the probability with which each of the elements is changed (e.g., Riechmann, 2001: 27).

3) observational learning through imitation

So far, the pool of policy platform from which party A (B) can learn is restricted to Π_t^A (Π_t^B). Observations concerning the behavior and success of other parties is not incorporated, though it can be an important source of learning. The simplest form of observational learning is copying a policy platform which has been successfully applied by another party (e.g., Brenner, 1999: 63-64).

4) directional learning

Observations in round t can initiate directional learning (e.g., Selten, 1998). To illustrate the idea of directional learning, assume that platform X has won the elections in t. It achieved $V_{X,t}$ votes, produced an unemployment rate of $U_{X,t}$ and thus is assigned a new fitness $\phi_{X,t+1}$ Figure 3 illustrates the procedure of directional learning for a policy platform Y whose fitness is lower than that of X ($\phi_{Y,t-1} < \phi_{X,t+1}$). In this case, the policy platform Y might be changed by adjusting the intensities for those policy instruments where Y differs from X. Adjustment means reducing the differences in intensities for a number of instruments by moving the original position of Y closer to X. The new policy platform Y' will be assigned a fitness which is given by the following expression:

$$\phi_{Y,t+1}. = \phi_{Y,t}\ (1+lb) \qquad (6)$$

It does not necessarily mean that $\phi_{Y,t}$ must exceed $\phi_{X,t}$.

Table 2

It is unlikely to assume that all policy platforms in Π_t^A respectively Π_t^B have to be modified in order to become member of Π_{t+1}^A respectively Π_{t+1}^B. In reality, many policy platforms, especially the fitter ones, can be expected to be part of the knowledge pool for many periods without undergoing any changes. In order to account for this fact, this paper assumes that policy platforms from Π_t^A respectively Π_t^B can be passed on to Π_{t+1}^A respectively Π_{t+1}^B

without undergoing major changes (e.g. Frank, 2003; Witt, 2003). The probability of being passed on unchanged is higher the fitter a policy platform is.

2.3 Introducing yardstick competition

In reality, there is a large number of governmental units which simultaneously generate knowledge. Therefore, the informational feedback on which voters and politicians in region 1 can draw on is not restricted to the knowledge generated within this region (e.g., Slembeck, 1997; Schnellenbach, 2004). Instead, they can make use of the feedback generated in all other regions r = 2, … R. Political parties in region 1 can utilize the feedback from other regions to gain information about the electoral and/or economic success of policy platforms which were not tested in region 1. As a consequence, the pool of policy platforms to learn from is not restricted to those held by politicians in region 1 but includes those applied in other regions. For voters, the external feedback provides a yardstick by which they can measure the economic success of their regional government. Technically, they can expect to update their benchmark U* when they see that other regions are more successful in fighting unemployment. This reduces the valence of the current government because, in the light of the external feedback, their performance becomes less impressive. The degree to which voters and politicians can make use of external feedback depends on the degree of interregional homogeneity. The higher the degree of interregional homogeneity, the more valuable the external feedback.

3. A computerized model to simulate party learning

In the current model, party competition is a dynamic process. The dynamics result from two facts. First, both parties learn the effects of different policy platforms on the number of votes and the unemployment rate. By testing and evaluating new policy platforms, they create new knowledge. The path of knowledge-creation depends on the behavior of the parties and thus is endogenous. Second, the environment in which the parties compete changes, because the valence voters assign them depends on the path of learning and is thus endogenous as well. For these reasons, the result of the dynamic process of party competition cannot be foreseen ex ante (e.g., von Hayek, 1972; Wohlgemuth, 2002). In order to predict the outcome of this dynamic process, the process of learning has to be modelled explicitly. This paper develops and applies a computer-based model to simulate the parties' learning process.

Based on these simulations, different questions can be answered: First, it may be of interest to know under which circumstances parties can be expected to learn adequate policies (i.e. policies which are close to H^{opt}) at all. Bischoff (2005) has shown that the essential precondition for successful learning is that voters remunerate good economic policies. In technical terms, this means that the parameter κ is sufficiently high. Only in this case can we expect a majority of voter to vote for an economically successful party even though, on plausibility grounds, its policy platform is less appealing to them than other (less adequate) policy platforms. At the same time, the degree of opportunism, represented by the parameter γ, is only of minor importance. The path of learning proved sensitive to the forms of learning the parties apply.

The section question which can be addressed using a computerized model of party competition is which factors determine the progress of learning in the early rounds. This question is of particular interest in a world in which H^{opt} is changing over time. For this purpose, the following section will model the process of learning in a way which allows for a direct correspondence between one round in the simulated learning process and an election term in real time. Given this set-up, it is possible to make predictions concerning the speed of learning and identify factors that determine how quickly a system learns. In addition, the current paper allows for directional learning as an alternative form of learning. Simulations will compare the speed of learning of GA-learning and directional learning.

3.1 Set-up of the political arena

Let us assume that there are D = 8 possible policy instruments (see table 1). A policy platform can be thought of as an 8x1 vector in which the first element states the intensity for instrument 1, the second element the intensity for instrument 2 and so on. For illustrative reasons, the following analysis will reinterpret this vector. Accordingly, there are only two policy instruments which can be fine-tuned in their intensity. Each policy instrument can be applied with an intensity of 0 to $8^{L/2}$-1 = 4095. The transformation of the 8x1 vector is described in figure 1. It ensures that any possible 8x1 policy vector has a unique representative in the two-dimensional policy space.

Figure 1

In total, 4096^2 = 16.8 million different policy platforms can be held. Thus, the number of possible policy platforms remains large and the decision making problem the parties face is complex. At the same time, this reinterpretation makes it possible to project the policy platforms into a two-dimensional Cartesian system and visualize the path of learning (see figure 2). The policy space is a square with the following corners {(0,0); (0,4095); (4095,4095); (4095,0)}. Voters prefer a large number of different policy platforms within the boundaries of this squarewith the median voter preferring the policy platform η^{Med} at (2740,2740). The optimal platform H^{opt} is located at (1370,1370). In round 1, Π_t^A and Π_t^B are centered around the points (3000,1000) respectively (1000,3000). The heterogeneity of platforms is moderate within both parties while the average platform of party A differs distinctly from that of party

B. In this initial set-up, both politicians and voters entertain biased beliefs about how the economy works and which policies are suitable to minimize unemployment.

Figure 2

The unemployment rate U_t depends on the distance between η_t and H^{opt} in the two-dimensional policy space:

$$U_t \quad = \quad \alpha \left(\overline{\eta_t \; \mathrm{H}^{opt}} \right)^2 + U_{\min} \qquad \alpha > 0 \tag{6}$$

Given this equation, the problem of finding H^{opt} does not seem complex in the two-dimensional policy space. Especially for parties which learn directionally, H^{opt} seems easy to discover. This course of argumentation does, however, ignore the fact that directional learning takes place in the eight-dimensional policy space. In this policy space, directional learning on the level of the single policy instrument does not correspond to a smooth movement towards H^{opt} in the two-dimensional policy space. The effects of changing the intensity of one policy instrument depend on the intensity with which the other instruments are applied. If they harmonize, the new platform will be more successful, if not, it will be less successful. Thus, a party would need to know the transformation in figure 1 and know about the smooth path on the two-dimensional level in order to be able to discover the path to H^{opt}.

3.2 The learning process

The computerized model of learning for over 100 rounds. Each round contains the following steps:

Step 1: Electoral competition

The policy platforms η_t^A and η_t^B are chosen. A total of 1001 voters make their choice according to the procedure described above.

In round t = 1, $K_t^A = K_t^B = 1$ and each policy platform within Π_t^A respectively Π_t^B is assigned a fitness which assumed that it reaches 50 % of all votes and produced an unemployment rate of U*.

Step 2: Policy application

The winning platform η_t is applied and leads to the unemployment rate U_t.

Step 3: Party learning

Out of the N elements in Π_t^A respectively Π_t^B, the bxN fittest platforms are passed on to Π_{t+1}^A respectively Π_{t+1}^B without any changes. The remaining (1 – b)xN elements, are passed to the intermediate population Π_t^{Ai} respectively Π_t^{Bi}.

<u>In the case of learning by communication (crossover):</u>

In a random process, N elements from Π_t^A respectively Π_t^B are selected as parent platforms. The probability of selection is a positive function of the individual platform's fitness. A single platform can be selected more than once. Parent platforms engage in a randomly matched crossover to produce N children. These are passed to the intermediate population Π_t^{Ai} respectively Π_t^{Bi}. The children are assigned a fitness according to the procedure described above.

In the case of directional learning:

In a random process, N elements from Π_t^A respectively Π_t^B are selected as candidates for directional learning. All platforms which have a lower fitness than η_t can learn from η_t directionally. For a learning platform η_i^A, each of the eight intensities is changed with a probability of $p^{dir} = 0.5$. The new intensity of η_i^A for policy instrument d is calculated as follows:

$$\eta_{i,d}^{A} = round\left[(\eta_{i,d}^{A} + \lambda(\eta_{t,d} - \eta_{i,d}^{A})\right] \tag{7}$$

where λ is a random number which is equally distributed in the interval [0,1].

In the case of observational learning through imitation:

Up to N previously applied policy platforms from Π_t are passed to Π_t^{Ai} and Π_t^{Bi}. If $t > N$, the policy platforms are chosen in a random process where the probability that a certain η_t is chosen is a positive function of its performance in t. Their fitness is calculated using their electoral and economic success as documented in Π_t.

All N policy platforms from Π_t^A respectively Π_t^B are subject to mutations. For each policy platform and each of the eight policy instruments, the probability of mutation is $p^M = 0.05$. In case of a mutation, the intensity of the corresponding policy instrument is multiplied by the factor f [$f \sim No(1,\sigma^f)$]. The resulting new intensity is rounded to the full number; the boundaries [0, 7] are respected. After passing the mutation operator, the mutated policy platforms are passed to Π_t^{Ai} respectively Π_t^{Bi}. Their fitness remains unchanged.

At the end of this process the intermediate population Π_t^{Ai} respectively Π_t^{Bi} contains up to (3+b)xN policy platforms. In a random procedure, (1 – b)xN strategies are selected to become members of Π_{t+1}^{A} respectively Π_{t+1}^{B}. The probability of a single platform from Π_t^{Ai} respectively Π_t^{Bi} to be selected is a positive function of its fitness. It is made sure that the new knowledge pools do not contain duplicates.

Step 4: Voters' learning

In the current model, voters learn only rudimentary by updating their perceived valence for the parties. If party A is in government, K_{t+1}^{A} is updated as follows:

$$K_{t+1}^{A} = \frac{1}{2} \cdot \left[K_t^A + \left(\frac{U*}{U_t^A} \right)^{\kappa} \right] \tag{8}$$

If party B wins the election, the update is accordingly.

Step 5: Updating Π_t

Finally, the common knowledge pool is updated.

3.3 Scenarios to be simulated

Based on the set-up described above, a computerized model of party learning is used to simulate a large variety of possible scenarios. The scenarios differ with respect to the values the parameters within the model are assigned. Table 3 gives an overview over the most important parameters and their standard values. By changing these values, a great variety of quantitatively different scenarios can be produced. At the same time, especially the value of κ is sufficiently high to create an environment in which learning pays in the long run. This means that the majority of voters will prefer H^{opt} over η^{Med} once they have repeatedly observed the corresponding unemployment rates and adjusted their expected valence accordingly.

Table 3

Next to quantitatively different scenarios, simulations can be made for scenarios which differ qualitatively. For instance, qualitatively different scenarios apply different forms of learning or allow for yardstick competition while others do not.

Regardless of the scenario developed, the essential aim is to observe the speed with which the society discovers and applies adequate policy platforms. The adequacy of a policy platform is measured by the unemployment rate it produces. The lower the unemployment rate is, the more adequate the policy platform. Consequently, the course and speed of learning can be analysed by tracing the unemployment rates across rounds of the computer simulation.

Due to the importance of chance for the path of learning, a single run of the simulation model yields a learning process which is unique but not necessarily representative. In order to arrive at a reliable picture of the typical learning process for a given scenario, a large number of simulations need to be conducted. This paper simulates 500 learning processes for every single scenario.

4. Preliminary results

Before applying the computerized model to trace the process of learning in different scenarios, the model needs to be calibrated. In the current context, calibration means testing whether the current model can reproduce the results of previous analyses (e.g., Kydland and Prescott, 1996). In particular, if we abstain from the valence-term and set $K_t^A = K_t^B = 1$, does η_t converge to η^{Med} as predicted by theory (e.g.,Plott, 1967) and as observed in the analysis of party learning by Kollman et al. (1992)? The answer is affirmative for all models of learning. Thus, the model can be used for the purpose described above.

The current paper presents work in progress. For this reason, I cannot delivers final results at this stage of the project. First results indicate that there are some interesting differences between the different forms of learning, especially between GA-learning and directional learning. Furthermore, the results indicate that for low values of N, the marginal benefit from increasing N is large while for larger values, a further increase in the size of the knowledge pool does not speed up the process of learning. The degree of opportunism (i.e. γ) seems to have a negative effect on the speed of learning. The learning benefit lb seems to have a positive impact on the speed of learning because it leads to more new policy platforms being tested. If lb = 0, parties are less likely to choose new elements from Π_t^A respectively Π_t^B.

In the next two months, I will run the missing simulations and evaluate them. I am confident that I can present a mature version of the paper which thoroughly reports on the results of my simulations by the end of February, 2006.

5. Conclusion

At this stage, final conclusions cannot be drawn.

References

Ansolabehere, S. and Snyder, J. M.(2000). Valence politics and equilibrium in spatial election models. Public Choice 103 (1-2), p.327-336.

Bischoff, I. (2005). Party competition in a complex environment – Do politicians learn good economic policies? Paper presented at the EPCS 2005 in Durham, UK and at the ISNIE 2005 in Barcelona, Spain.

Brenner, T. (1999). Modelling Learning in Economics. Cheltenham, Edward Elgar.

Denzau, A. T. and North, D. C. (1994). Shared Mental Models: Ideologies and Institutions. Kyklos Vol. 47, p.3-31.

Downs, A. (1957). An economic theory of democracy, New York.

Edmonds, B. (1999). Modelling Bounded Rationality in Agent-based Simulations Using the Evolution of Mental Models. in: Brenner, T. (Ed.). Computational Techniques for Modelling Learning in Economics. Boston usw., Kluwer Academic Publishers: 305-332.

Enelow, J. M. and Hinich, M. J. (1984). The spatial theory of voting. Cambridge, Cambridge University Press.

Fiorina, M. P. (1981). Retrospective Voting in American National Elections. New Haven/London, Yale University Press.

Forbus, K. D. (1983). Qualitative Reasoning About Space and Motion. In: Gentner and Stevens (Eds): Mental models. Hillsdale etc. Erlbaum, p.53-73.

Frank, J. (2003). Natural selection, rational economic behaviour, and alternative outcomes of the evolutionary process. Journal of Socio-Economics 32, p.601-622.

Griffin, D. and Tversky, A. (2002). The weighing of evidence and the determinants of confidence. in: Gilovich, T. et al. (Eds.): Heuristics and biases. The psychology of intuitive judgement. Cambridge, Cambridge University Press. p.230-249.

Groseclose, T. (2001). A model of candidate location when one candidate has a valence advantage. American Journal of Political Science 45, p.862-886.

Kollman, K., Miller, J. E. and Page, S. E. (1992). Adaptive Parties in spatial elections. American Political Science Review 86 (4), p.929-937.

Kydland, F. E. and Prescott, E. C.: The Computational Experiment. Journal of Economic Perspectives, 10 (1996), p.69-85.

Michalewicz, Z. (1996). Genetic Algorithms + Data Structures = Evolution Programs, 3rd Edition, Berlin etc, Springer.

Plott, C. R. (1967). A notion of equilibrium and its possibility under majority rule. American Economic Review 57 (September), p.787-806.

Richards, D. (2001). Coordination and Shared Mental Models. American Journal of Political Science 45(2), p.259-276.

Riechmann, T. (2001). Learning in Economics. Analysis and Application of Genetic Algorithms, Heidelberg and New York, Physika Verlag.

Schnellenbach, J. (2004). Dezentrale Finanzpolitik und Modellunsicherheit. Tübingen, Mohr Siebeck.

Selten, R. (1998). Features of experimentally observed bounded rationality. European Economic Review 42: 413-436.

Simon, H. A. (1956). Theories of decision-making in economics and behavioral science. American Economic Review 48 (19), p.253-283.

Slembeck, T. (1997). The Formation of Economic Policy: A Cognitive-Evolutionary Approach to Policy-Making. Constitutional Political Economy 8, p.225-254.

Swank, O. (1998). Towards an economic theory of party ideology. European Journal of Political Economy 14, p.223-240.

Tetlock, P. E. (1999). Theory-Driven Reasoning about Plausible Pasts and Probable Futures in World Politics: Are We Prisoners of Our Preconceptions? American Journal of Political Science 43(2), p.335-366.

Von Hayek, F. A. (1972). Die Theorie komplexer Phänomene. Tübingen, Mohr Siebeck.

Witt, U. (2003). Economic policy making in evolutionary perspective. Journal of Evolutionary Economics 13, p.77-94.

Wohlgemuth, M. (2002). Evolutionary Approaches to Politics. Kyklos 55, p.223-246.

Table 1 Preferred policy platforms in a multidimensional policy space

Instrument (d = 1, … D)	Intensity of application							
	0	1	2	3	4	5	6	7
1. (e.g. qualification schemes)		i		A		B		
2. (e.g. improved infrastructure)					B	i	A	
3. (e.g. labor market deregulation)		A,B	i					
4. (e.g. lower income and profit taxes)				B	i,A			
5. (e.g. deficit spending)	A				B	i		
… … … … … …		B		A			i	
… … … … … …	A	i	B					
D. (e.g. protectionist measures)				A	i,B			

Table 2: Directional learning

Instrument	Intensity of application							
	0	1	2	3	4	5	6	7
1	X	Y’	←	Y				
2				X		Y		
3			X, Y					
4			Y			X		
5				X			Y	
6					Y	X		
7		Y	→	Y’		X		
8		X		Y’	←	←	Y	

Table 3: Parameters of the computerized learning model

Parameter	Standard value	Description
κ	0.4	Exponent in the valence function; determines the importance voters assign to the economic performance of a party
γ	0.5	Relative weight of the electoral success of a policy platforms when calculating its fitness
N	20	Size of the knowledge pool of a party
p^M	0.05	Probability that a single intensity of a policy platform is subject to mutations
σ^f	0.05	Factor determining the amplitude of mutations
p^{dir}	0.5	Probability of a single intensity of a policy platform to be changed in the process of directional learning
B	0.5	Share of policy platforms in Π_t^A (Π_t^B) which are passed to Π_t^A (Π_t^B) unchanged
Lb	0	expected learning benefit which parties add to the fitness of new policy platforms which emerge from crossover or directional learning
U_{min}	0.03	Minimum unemployment rate if $\eta_t = H^{opt}$
U*	0.15	Benchmark unemployment (= U_t for central position of η_1^A or η_1^B
α	0.7×10^{-4}	Scalar in the unemployment function (6)

Figure 1: Transformation of 8x1 policy vector

4 0 7 5 1 6 2 3 original 8x1 vector

Zeile 1	Zeile 2	Position
4 x 8^3	1 x 8^3	1
0 x 8^2	6 x 8^2	2
7 x 8^1	2 x 8^1	3
5 x 8^0	3 x 8^0	4
= 2109	= 915	

2x1 policy vector

(2109, 915)

Figure 2: Starting point of party learning in the two-dimensional policy space