

**A Model of
Boundedly Rational Consumer Choice**
—
An Agent Based Approach

THOMAS RIECHMANN

Diskussionspapier Nr. 232

ISSN 0949–9962

Mai 2000

Thomas Riechmann, Universität Hannover, Fachbereich Wirtschaftswissenschaften, Institut für
Volkswirtschaftslehre, Abteilung Wachstum und Verteilung, Königsworther Platz 1, 30167 Hannover,
Mail: riechmann@vwl.uni-hannover.de

Abstract

The paper presents an extended version of the consumer choice problem. Different from the standard model, prices are not fixed but arise from Walrasian interactions of total demand and a stylized supply function for each of the goods. Three different types of evolutionary algorithms are set up to answer the question whether agents can learn to solve the problem of extended consumer choice. There are three important answers to this question: a) The quality of the results learned crucially depends on the elasticity of supply, which in turn is shown to be a measure of the degree of state dependency of the economic problem. b) It seems to be relatively easy to adhere to the budget constraint, but relatively difficult to reach an optimum with marginal utility per Dollar being equal for each good. c) Agents equipped with some memory are found to perform notably better than agents without memory.

Zusammenfassung

Gegenstand des Papiers ist die Konsumwahl privater Haushalte. Im Gegensatz zum entsprechenden Standardmodell sind die Güterpreise nicht fest, sondern entstehen jeweils durch Walrasianische Interaktion von aggregierter Nachfrage und einer stilisierten Angebotsfunktion. Es werden drei verschiedene Typen evolutionärer Algorithmen entworfen, um die Frage zu beantworten, in welchem Maße entsprechendes Lernverhalten der Haushalte zur Lösung des Konsumwahlproblems führen kann. Auf diese Frage ergeben sich drei wichtige Antworten: a) Die Qualität der erlernten Verhaltensweisen hängt insbesondere von der Elastizität des jeweiligen Güterangebots ab, welches sich als Maß der Zustandsabhängigkeit der Problemstellung interpretieren läßt. b) Es fällt relativ leicht, die Budgetbeschränkung einzuhalten, wohingegen die Angleichung der Wertgrenznutzen über alle Güter des Konsumbündels sehr schwer zu erlernen ist. c) Agenten, die mit ein wenig Erinnerungsvermögen ausgerüstet sind, erzielen deutlich höheren ökonomischen Erfolg als Agenten ohne Gedächtnis.

Key words: consumer choice, evolutionary algorithms, state dependency

JEL classifications: C63 – D11 – D83

1 Introduction

In introductory courses to microeconomics, when it comes to the problem of consumer choice, a question often heard is ‘Do people *really* behave this way?’ The standard answer to this question is ‘Not *really*, but they can learn to achieve the optimal outcome, anyway’ But, is this really true? Can boundedly rational people learn how to choose their optimal consumption bundle?

This paper tries to answer this question. In order to do this, the standard text book problem of consumer choice is extended: The assumption of fixed prices is dropped, which makes the individual problem even harder to solve. This model is simulated, applying three different learning techniques in form of three different evolutionary algorithms. These kinds of algorithms have often been applied to similar economic problems. Interpretations of evolutionary algorithms as metaphors for various types of learning schemes can be found in e.g. Dawid (1999) or Riechmann (1999b).

It can be shown, that the learnability of an optimal solution of the extended consumer choice problem not only depends on the learning technique chosen, but also on the degree of state dependency of the extended consumer choice problem, which, in turn, can be measured by the slope of the supply functions of the model.

More than this, it can be learned, that it is easier to adhere to the budget constraint than it is to find the optimal consumption bundle.

Apart from answering a core economic problem, this paper shows some new features of evolutionary programming, which have rarely been used in economic modelling before. The most important new feature is the simulation of simultaneous constraint optimization over more than just one variable, including binary coding of the variables and the use of a penalty function.

2 The Economic Model

2.1 The General Structure of the Problem

Let there be a number of n agents (households) facing the standard textbook problem of consumer choice, i.e. selecting a bundle of consumption goods which, under the restriction of a limited budget, maximizes utility.

Agents are assumed to have identical utility functions, which do not change over time t . Utility is derived from consuming a bundle of m different goods indexed k . For conve-

nience, the utility function is assumed to be of Cobb–Douglas type.

$$u_{i,t} = A \prod_{k=1}^m q_{i,k,t}^{\alpha_k}; \quad A > 0; \alpha_k > 0 \forall k = 1, \dots, m. \quad (1)$$

$q_{i,k,t}$ gives the quantity of good k agent i consumes in period t . A and $\alpha_k, k = 1, \dots, m$ are parameters of the model.

The budget M is assumed to be the same for every agent i in every period t , so the budget constraint is given by

$$M \geq \sum_{k=1}^m p_{k,t} q_{i,k,t}. \quad (2)$$

$p_{k,t}$ represents the market price for good k in t .

Thus, every agent i aims to solve the constrained maximization problem

$$\max_{q_{i,k,t}} u_{i,t} \quad \forall k = 1, \dots, m \quad (3)$$

$$\text{s.t.} \quad M \geq \sum_{k=1}^m p_{k,t} q_{i,k,t} \quad (4)$$

$$q_{i,k,t} \geq 0 \quad \forall k = 1, \dots, m \quad (5)$$

Equation (5) gives the usual non–negativity constraints.

Different from the standard textbook model, prices will not generally be held fixed, but will be subject to a stylized Walrasian mechanism. For each good k , in every period t , the price will be determined as the equilibrium price resulting from the interaction of aggregate demand for good k in t and aggregate supply of the good.

Aggregate demand for good k in t , $Q_{k,t}$ is simply the sum of individual demand for k in t :

$$Q_{k,t} = \sum_{i=1}^n q_{i,k,t}. \quad (6)$$

Aggregate supply will be modelled by a time invariant standard supply function for each good k , so that the equilibrium price $p_{k,t}^*$ results as

$$p_{k,t}^* = B_k + m_k \sum_{i=1}^n q_{i,k,t}. \quad (7)$$

2.2 The Basic Model

The basic model is a model of fixed prices. This means that in equation (7) $m_k = 0$, yielding

$$p_{k,t}^* = B_k. \quad (8)$$

Note, that this case does not represent an ‘economic’ problem in the sense of agents’ fitness being state dependent. Each agent’s utility only depends on her own actions, but in no way on the actions of any other agent.

For fixed prices the solution to the problem of consumer choice can easily be derived. By standard calculus, the solution can be determined as

$$q_{i,k,t}^* = \frac{\alpha_k}{\sum_{k=1}^n \alpha_k} \frac{M}{B_k}; \quad \forall k; \forall i; \forall t. \quad (9)$$

This result yields two crucial characteristics common to every optimal bundle of consumption goods. The first characteristic is efficiency: The whole budget is being spent. In an optimal situation, (2) becomes binding:

$$M = \sum_{k=1}^m p_{k,t} q_{i,k,t}. \quad (10)$$

The second crucial characteristic of the optimal consumer choice is the fact, that marginal utility per Dollar is the same for every pair of goods k, l ,

$$\frac{\partial u_{i,t}(\cdot) / \partial q_{k,i,t}}{p_{k,t}} = \frac{\partial u_{i,t}(\cdot) / \partial q_{l,i,t}}{p_{l,t}} \quad \forall k, l \in \{1, \dots, m\}. \quad (11)$$

Both of these well known standard results will become of greater importance in the second part of the paper.

2.3 The Enhanced Model

In the enhanced case, prices are no longer fixed, i.e. $m_k \neq 0$ in (7). As a consequence, the problem of consumer choice becomes a problem of state dependent fitness: Aggregate demand now has an influence on the market price and the market price has an influence on every single agent’s economic success.

With the introduction of flexible prices the problem becomes analytically unsolvable. This means that for flexible prices, a general explicit numerical solution analogous to (9) for the fixed price case cannot be computed. Nevertheless, at least for uniformly behaving populations, the characteristics (10) for efficiency and (11) for optimality still apply.

The question if boundedly rational agents are able to achieve relatively good solutions to the enhanced problem can be put in another way: As will be pointed out in greater detail later on (Section 5.1), there are two main influences on an agent's utility: The direct effect of her quantity decision and the indirect effect caused by all agents' demand influencing the price of the goods. In a world of bounded rationality, this indirect effect can be interpreted as a kind of external effect no agent accounts of.¹ Then, it is a question of *how bounded* the rationality of agents really is: If they recognize the external effect, their utility will be greater than if they ignore it. This means that, at least in this model, accounting for the external effect is a sign of rationality.²

3 The Evolutionary Algorithm

3.1 The Basics

An evolutionary algorithm aiming to model the above setting in an agent based manner must above all be capable of coping with two problems: a) The maximization problem in focus involves the optimization of more than just one independent variable, and b) the maximization problem in focus has to be solved subject to a constraint.

Both of the problems have been solved in natural sciences before,³ but up to now there is no economic model making use of these results.

In this paper, a hybrid evolutionary algorithm will be employed, making use of principles from two worlds, from the world of genetic algorithms (GAs, see e.g. Goldberg, 1989) and from the world of evolution strategies (ESs, see e.g. Bäck, Hoffmeister, and Schwefel, 1991). From GAs, the well known operators of crossover and selection/reproduction are used. In economics, they have been broadly interpreted as forms of social learning by communication and interpretation as well as the functioning of the market (Riechmann, 1999b). From evolution strategies, the operator of mutation is used, being interpreted as a form of isolated individual learning by experiment. Moreover, the variables in focus will not be coded as bit strings, but as real valued numbers, which is also a feature of evolution strategies.

¹To be more precise, this is a monetary external effect. Though the characterization of external effects is standard in modern economic literature, it seems interesting and important to still refer to the original work of Buchanan and Stubblebine (1962).

²Note, that this concept is similar to the concept of 'externalities' used in modern endogenous growth theory (Romer, 1986; Lucas, 1988; Barro, 1990). There, so called externalities appear which are recognized by the allmighty social planner but not by the 'normal' individuals populating the economy.

³A summary of results can be found in Michalewicz (1996).

To make things clearer, an agent i will be fully described by her economic plans, i.e. the vector of demanded quantities for each good k in period t , $q_{i,k,t}$. In the simulations, there will only be three different goods available to the agents. Thus, an agent is characterized by a vector $q_{i,t}$,

$$q_{i,t} \in \mathbb{R}^3. \quad (12)$$

An example would be

$$q_{10,5} (11.2; 3.7; 17.0) \quad (13)$$

meaning that agent number 10 plans to consume 11.2 units of the first, 3.7 units of the second and 17 units of the third good in period 5.

3.2 Standard Operators

Due to the change in the representation of the agents, the operators (often called ‘genetic operators’) have to be changed, too. As these operators are used as metaphors for learning techniques, the changed operators still have to support a sensible economic interpretation.

The ‘double operator’ of selection and reproduction, often being interpreted as learning by imitation connected with the role of the market, needs no changes at all. Agents are selected from their population and reproduced into the next one with a probability equal to their relative fitness. Fitness, in this case, equals individual utility. This means, that in this paper, the standard roulette wheel selection operator (Goldberg, 1989) will be used, not one of the selection operators stemming from the tradition of evolution strategies.

Recombination in the form of crossover is usually seen as a metaphor for learning by communication. The change of this operator is quite straightforward. In a real valued rather than a binary representation, agents’ economic strategies can be separated into clearly defined economic substrategies. A substrategy in the current model is the consumption quantity of a single good. This means that in the above example (13) the agent has three substrategies: 11.2 for good one, 3.7 for good two, and 17 for good three. Crossover now works as usual, recombining two agents’ substrategies. The agents involved in crossover are chosen from their population, a crossover point is selected, the vectors of substrategies are cut at the crossover point, the resulting parts are interchanged and put together again, yielding two new strategies. Figure 1 shows an example.

In the world of real valued coding of agents’ strategies, the interpretation of crossover as a form of learning by exchanging substrategies becomes even clearer than in the world of binary coding.

The third standard operator, mutation, has to be changed, too. Mutation, often being interpreted as learning by experiment, has to undergo the most severe changes. Real valued coding certainly does not allow for simple bit flipping. Instead, the mutation operator




Figure 1: Crossover (Example)

from the tradition of evolution strategies can be used. In this tradition, a substrategy $q_{i,k,t}$ is mutated by adding a term v_k to it, where v_k is normally distributed with mean 0 and a finite variance σ^2 , resulting in the new substrategy $\tilde{q}_{i,k,t}$.⁴

$$\tilde{q}_{i,k,t} = q_{i,k,t} + v_k \quad \forall k \in \{1, \dots, m\}; \quad v_k \sim N(0, \sigma^2) \quad (14)$$

By endogenizing the mutation variance σ^2 , this operator could easily be extended in order to represent some kind of meta learning (Riechmann, 1998), but for the clarity of the economic argument this will not be done within this paper.

3.3 Enhanced Operators

In addition to the standard operators, two more ‘enhanced’ operators shall be employed. The first one is the well known election operator (Arifovic, 1994). Election requires two (or more) agents to meet, jointly work out new strategies, evaluate these strategies and then finally decide which strategy to use in reality. Though election, especially in its basic form, is not undisputed in its economic meaning,⁵ it is known to result in stable states of the learning process, which most of the time even represent optimal solutions to the underlying economic problem.⁶ Thus — economically meaningful or not — election represents a good benchmark for the test of the performance of other learning operators.

The second enhanced operator is something more than just an operator, it requires a slight change in the construction and implementation of the agents. For the operator of *preselection*, agents will be equipped with a memory. And although this memory is very limited, it will be shown to help improving agents’ learning performance. Precisely,

⁴In order to make sure the adherence to the non-negativity constraints (7), the pure mutation operator has to be accompanied by some mechanism guaranteeing that $\tilde{q}_{i,k,t} \geq 0$.

⁵For a clear formulation of the points of criticism as well as for some suggestions how to fill the election operator with more economic meaning, see Franke (1997) and Birchenhall et al. (1997).

⁶Riechmann (1999a) provides an explanation to why election tends to result in stable states.

agents will be given the ability to remember one certain strategy together with the level of utility they gained from employing this strategy. The specific strategy an agent remembers is her all time best strategy, i.e. the strategy that has brought her the highest utility during the whole learning process. In preselection an agent chooses from two possible strategies: the strategy she used in the last period (the fitness of which she can remember) and her all time best strategy. She decides by the fitness. The strategy with the higher fitness is chosen as the strategy the following learning processes of the current period (i.e. communication and experiment) are based on.

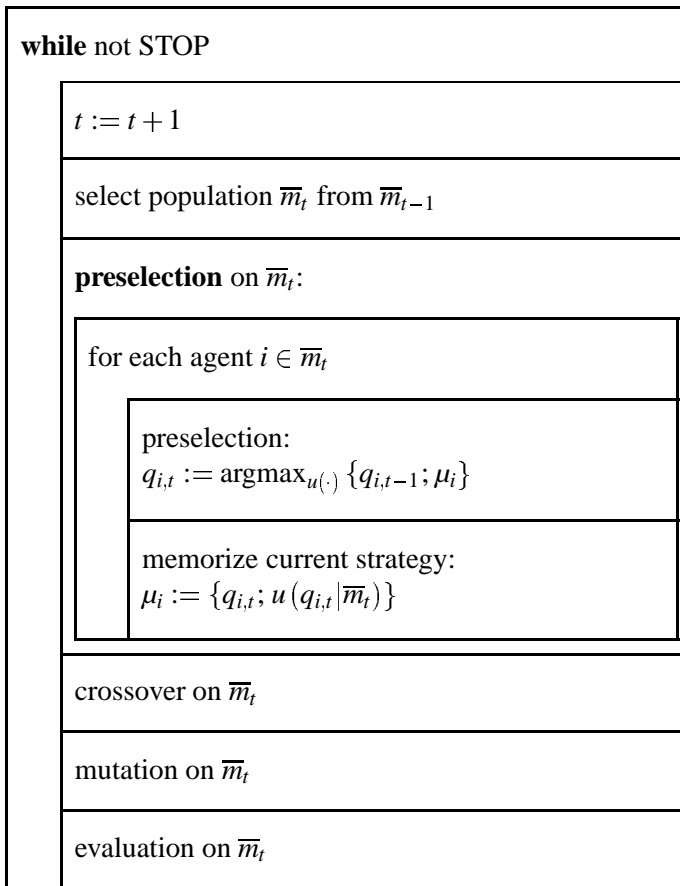


Figure 2: Main Loop of Preselection Algorithm

Note, that for non state dependent problems, like the fixed price problem of consumer choice, the preselection algorithm essentially works like an intra agent operator of elitist selection for the basic strategy in each period. In the simple fixed price model, an agent's fitness or utility is independent of what the others do, i.e.

$$u_i(q_{i,t}|S_t) = u_i(q_{i,t}) \quad \forall S_t \in S'. \quad (15)$$

S_t is the population in Period t , S' is the set of all different populations.

In problems of state dependent fitness, on the contrary, all actions of all the other agents in the population have an influence on an agent's fitness. This means, that for $m_k \neq 0$ in (7), the strategy an agent thinks of as her all time best one, might — due to the current population — not be as good as she thinks. A strategy that once performed brilliantly may perform very poorly in the context of a different population. This is a fact that the agents in this model are assumed to ignore. Agents use the rule of preselection, because they simply do not know that they may be mistaken.

3.4 Coping with the Constraints

The problem in focus is a problem of constrained optimization. Agents do not only have to maximize utility but also have to be careful not to exceed their budget. In the simulations, the budget constraint will not be *directly* accounted for. This means that there is no good the consumed quantity of which serves as some kind of residual. This means that agents *do not* do something like determine the quantity of the last good as all they can afford to buy by the rest of their budget like

$$q_{i,m,t} = \frac{1}{p_{m,t}} \left(M - \sum_{k=1}^{m-1} p_{k,t} q_{i,k,t} \right). \quad (16)$$

Instead, agents freely decide on the quantity of all three goods, at first hand independently of the budget constraint. This also means, that in the fixed price cases, at the time of making their consumer plans, agents do not even know the prices of the goods, which makes their decision problem even harder.

The compliance to the budget constraint is secured by using a *penalty function*, which is a standard tool in evolutionary optimization.⁷ Using a penalty function means to decrease an agent's fitness if she breaks the budget constraint. In economic terms, this looks like an extension to the utility function, which may be rather unusual to most theoretical economists.

⁷See e.g. Michalewicz (1996, pp. 321).

In the simulations the following fitness function was used in order to transform utility $u_{i,t}$ into fitness $R_{i,t}$:⁸

$$R_{q_{i,t}} = \begin{cases} U(q_{i,t}) + M - \sum_{k=1}^3 p_{k,t} q_{i,k,t} & \text{for } \sum_{k=1}^3 p_{k,t} q_{i,k,t} > M \\ U(q_{i,t}) & \text{for } \sum_{k=1}^3 p_{k,t} q_{i,k,t} \leq M \end{cases} . \quad (17)$$

Using fitness function (17) means, that an agent *can* break the constraints, but that she should *learn* not to do so.

4 Simulations and Results

For the model described above, simulations were run using three different algorithms, a) the canonical algorithm, using selection/reproduction, crossover and mutation, b) the election algorithm, using selection/reproduction and election, and c) the preselection algorithm, using selection/reproduction, preselection, crossover and mutation. The simulations were run for three different parameter sets which only differ with respect to the elasticity of supply.

All three sets thus share the values for $A = 1$ and for the budget M , $M = 100$, for B_k , $B_1 = B_2 = B_3 = 1$, and for α_k , $\alpha_1 = 0.2$, $\alpha_2 = 0.3$, and $\alpha_3 = 0.5$. The only difference is that for the first set ('FIXPRICE'), $m_1 = m_2 = m_3 = 0$, yielding $p_1 = p_2 = p_3 = B_k = 1$. Thus, FIXPRICE represents the fixed price case or, put in different words, a case of infinitely high elasticity of supply.⁹ The second set ('HIGHELASTICITY') uses $m_1 = m_2 = m_3 = 0.0001$, thus representing a state of high elasticity of supply. The third set ('LOWELASTICITY') uses $m_1 = m_2 = m_3 = 0.1$ which leads to a state of low elasticity of supply.

Simulations were run for populations of $n = 500$ agents and for $t_{\max} = 500$ periods.

⁸Negative fitness values can arise, so that before entering the selection operator, fitness has to be subject to one of the standard positive transfer mechanisms (Goldberg, 1989; Mitchell, 1996).

⁹The definition of 'elasticity of supply' differs between various textbooks of microeconomic theory. This paper makes use of the definition by Henderson and Quandt (1986), who define the price elasticity of supply as the ratio of relative change in the quantity of supply (nominator) to the relative change of the price (denominator).

For this paper this means that the elasticity of supply for good k , ϵ_k , is given by $\epsilon_k = \frac{\partial Q_k / Q_k}{\partial p_k / p_k}$. This definition implies that the change in the quantity is a result of a change in market price. This paper, though, will argue the other way round: A change in price is the reaction on a change in aggregate demand, i.e. a change in quantity.

4.1 Fixed Prices

For the fixed price case, at least for the case of uniform behavior, optimal outcome can be determined analytically. If all agents within a population behave the same, the best strategy according to (9), given parameter set FIXPRICE is

$$q_{1,k,t}^* = 20, \quad q_{2,k,t}^* = 30, \quad q_{3,k,t}^* = 50 \quad \forall \quad k, t. \quad (18)$$

These result represent both, efficiency and optimality. This means that for the results given in (18) the budget is fully spent (i.e. the budget residual is zero) and every agent's utility is as high as possible due to the budget.

The quality of the learning algorithms can be judged by comparing the results learned to the theoretical results given by (18). More than this, it seems appropriate to check if, and if so, how good, the budget constraint is met. In order to do this, the average budget residual ρ_t will be used, which is defined as the average amount of money not spent on consumption throughout the population, i.e.


$$\rho_t = \frac{1}{n} \sum_{i=1}^n \left(M - \sum_{k=1}^3 p_{i,k,t} q_{i,k,t} \right). \quad (19)$$

Representative simulation results for the three respective algorithms are given in the following figures. Each figure consists of two subfigures which give a plot of the population average of the quantities of good one to three (on the left) and a plot of the average budget residual (on the right).


It is easy to see that for all three algorithms, the quantities tend to converge toward the theoretically optimal quantities. The election algorithm seems to perform best whereas the canonical algorithm performs worst. This impression is supported by the results concerning the budget residual. Election converges to a residual equal to zero, preselection tends to oscillate around zero, whereas the canonical algorithm results in relatively large positive residuals. This means that each of the three learning methods is able not to violate the budget constraint, but only learning with election and learning with preselection enables the agents not to waste a part of their income.

4.2 Flexible Prices, High Elasticity

For the case of flexible prices, it becomes impossible to determine an explicit solution analytically even for the case of homogeneous behavior. In order to judge the quality of the simulation results, two measures are introduced. The first one is the average budget residual as defined in equation (19). The second one is the difference between the marginal




(a)Quantities




(b)BudgetResidual

Figure3: CanonicalAlgorithm



(a)Quantities



(b)BudgetResidual

Figure4: Election Algorithm





Figure 5: Preselection Algorithm


utility per dollar for each pair of the three goods. Although it is impossible to solve the problem explicitly, the fact remains that in an optimal solution of the consumer choice problem marginal utility per dollar has to be equal for each good. Thus, a perfect solution should have two characteristics: a) Marginal utility per dollar is the same for each of the three goods, and b) the average budget residual is zero.

The plots of representative simulations for the three algorithms and the HIGH ELASTICITY case (figures 6 to 8) consist of three subfigures each. The first subfigure gives the population average of the quantities of each good. This subfigure only serves as an illustration. This figure does only show the quantities learned. But now, in the case of flexible prices, it does not offer any means for judging the results. The second subfigure gives the average budget residual ρ_t . The third subfigure is a plot of the marginal utility per dollar for each of the three goods.


A closer look at the figures confirms the results of the FIXPRICE case. The election algorithm performs best. Marginal utilities per Dollar seem to converge completely while the budget residual vanishes. The preselection algorithm generates slightly worse results. While the budget residual fluctuates around zero, marginal utilities per Dollar become similar but not equal. The canonical algorithm performs worst. Marginal utilities per Dollar become similar, but the budget is never fully spent.



(a)Quantities




(b)BudgetResidual




(c)MarginalUtilityperDollar


Figure6: CanonicalAlgorithm–HighElasticityofSupply



(a)Quantities




(b)BudgetResidual




(c)MarginalUtilityperDollar


Figure7: ElectionAlgorithm–HighElasticityofSupply



(a)Quantities



(b)BudgetResidual



(c)MarginalUtilityperDollar

Figure8:PreselectionAlgorithm–HighElasticityofSupply

4.3 Flexible Prices, Low Elasticity

The third case in focus, flexible prices with a low elasticity of supply, causes quite different results. It can be seen that the canonical algorithm does not reach a sensible outcome. The budget constrained is never met while marginal utilities per Dollar never get close to each other. In the LOWELASTICITY case, a sensible consumer choice cannot be learned by use of the simple canonical algorithm learning rules.

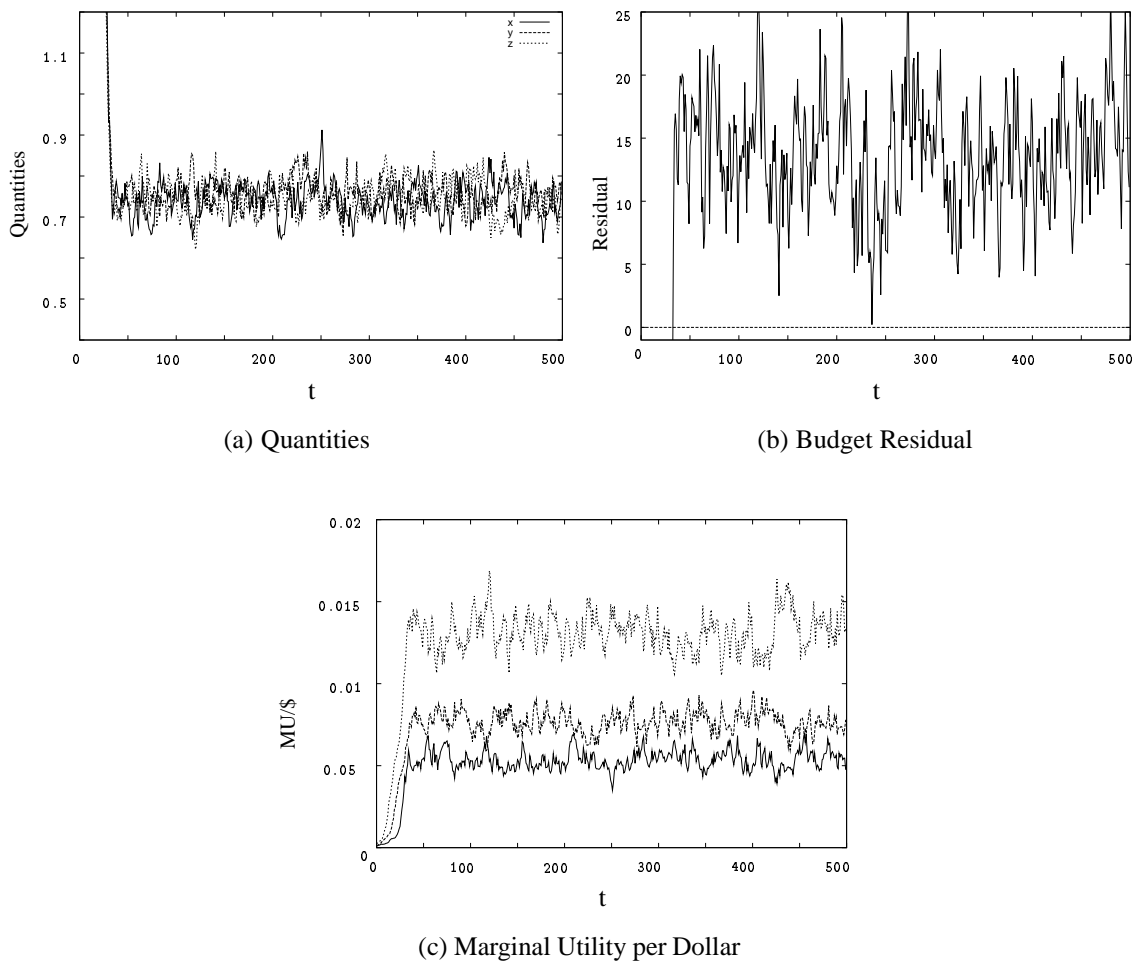


Figure 9: Canonical Algorithm – Low Elasticity of Supply

Even worse is the result for the election algorithm. In each of the simulations for this

case, with only one explicitly shown in this paper, the model collapsed. Election leads to totally unrealistic cases of violation of the budget constraint while marginal utilities per Dollar do not converge at all. For the case of flexible prices and a low elasticity of supply, election is far from leading to any kind of sensible consumer choice.

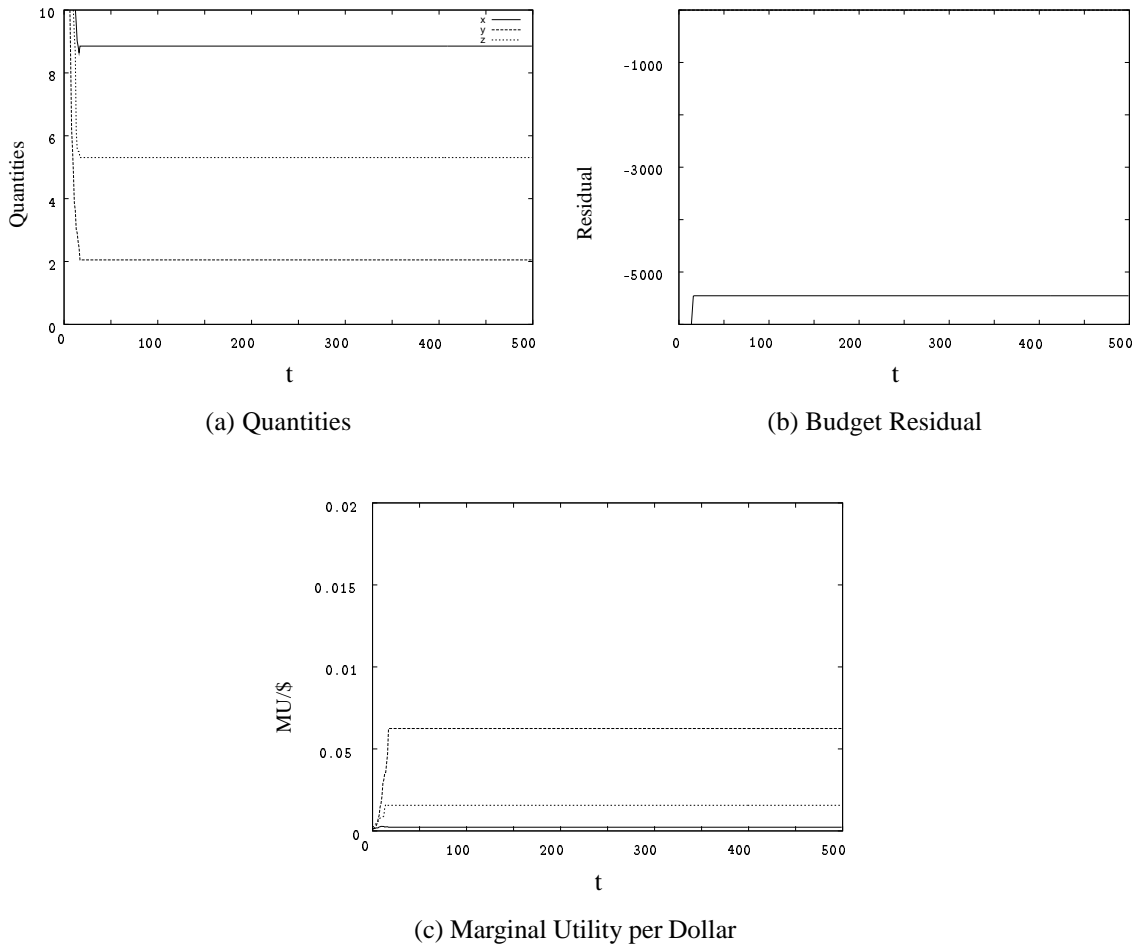




Figure 10: Election Algorithm – Low Elasticity of Supply


The only sensible results are achieved by the preselection algorithm, which causes the marginal utilities per Dollar to become at least slightly similar to each other while the budget constraint is not violated too severely.



(a)Quantities



(b)BudgetResidual



(c)MarginalUtilityperDollar

Figure 11: Preselection Algorithm—Low Elasticity of Supply

4.4 Summary of Results

At this early stage of the paper, the results are essentially based on two sources: The plots of the simulations shown above and some statistics, which are up to now only based on one run of each simulation only. The statistics can be found in Appendix A.

Concerning the budget constraint it can be found, that it seems to be relatively easy to learn not to violate it. For the FIXPRICE case as well as for HIGHELASTICITY, the canonical algorithm results in states which are significantly positive, while election is significantly very close to zero and preselection is even significantly equal to zero.

For the choice of the correct quantity or the best marginal utility per Dollar, for FIXPRICE and HIGHELASTICITY, the results from the plots do not look too bad, but from the statistics it can be found, that significances are quite poor. From the three different algorithms, the canonical one shows the highest degree of fluctuations (i.e. the highest variance) and consequently the highest (though still very poor) significance of reaching the optimal result, while election produces nearly no fluctuations at all, resulting in the least significance of an optimal outcome.

For LOWELASTICITY, only the preselection algorithm leads to sensible results at all.

All in all, it seems as if the preselection algorithm performs best over all types of economic situations. This means, that a little memory to the past notably improves learning abilities.

5 Conclusions

5.1 The Influence of State Dependency

As a conclusion to be drawn from the simulation results it can be found that it seems to be much harder to learn a sensible consumer choice in a situation with low elasticity of supply than in situation with high (or even infinitely high) elasticity.

The reason for this is the following. In each turn of the algorithm (in each market period) the resulting market prices and by that, the utility gained, reveal some information to each agent. This information is information about the quality of her last period consumption plan. *Ceteris paribus*, the plan was a good plan if utility is high and so the plan should not be changed too much for the next period. If, on the contrary, utility was not very high, the plan obviously was a bad one and thus should be changed. Unfortunately, this argument is only a *ceteris-paribus* argument. The information revealed actually consists of two parts which cannot be distinguished. Surely the information is information about the quality of the agent's plan. But it is only information about the quality of the

agent's plan in the context of the plans of all other agents in the population. This means, that from period to period the quality of an agent's plan can be changed either by a change in the plan itself or by a change in the plans of the rest of the population. While there certainly is a *direct* impact of the agent on her economic success, there is also an indirect impact caused by the rest of the population. This indirect impact has often been called *state dependency* of agents' fitness.

State dependency can be found to cause noise in the part of the information most valuable for the agent, i.e. the information about the quality of her plan. The stronger the state dependency, the stronger is the noise, the less valuable is the information to the agent, and thus the more complicated is the consumption decision. Thus it should be shown that less elasticity of supply means more state dependency in the problem of consumer choice in order to give a reason why less elasticity seems to complicate the problem.

It can be shown that the gradient m_k of the respective supply function (7) is a measure of state dependency. A sketch of the argument runs as follows: The utility function¹⁰ (1) can be rewritten as integrating the budget constraint (2)¹¹, using $q_{i,m}$ as a numeraire, yielding

$$V_i = V(q_{i,1}, q_{i,2}, \dots, q_{i,m-1}, M, p_1, p_2, \dots, p_m). \quad (20)$$

A change in utility can be written as

$$dV_i = \underbrace{\sum_{k=1}^{m-1} \frac{\partial V_i}{\partial q_{i,k}} dq_{i,k}}_{\text{direct effect}} + \frac{\partial V_i}{\partial M} dM + \underbrace{\sum_{k=1}^m \frac{\partial V_i}{\partial p_k} dp_k}_{\text{indirect effect}}. \quad (21)$$

Let us assume the budget M to be constant (this assumption fits with the simulations carried out), so that $dM = 0$. Consequently, (21) becomes a bit simpler, resulting in

$$dV_i = \underbrace{\sum_{k=1}^{m-1} \frac{\partial V_i}{\partial q_{i,k}} dq_{i,k}}_{\text{direct effect}} + \underbrace{\sum_{k=1}^m \frac{\partial V_i}{\partial p_k} dp_k}_{\text{indirect effect}}. \quad (22)$$

From this, it can be recognized that a change in an agent's utility can be caused by two effects, by the direct effect of the agent changing her quantities, and by the indirect effect of a change in prices.

¹⁰For notational convenience, the period index t is omitted.

¹¹An important prerequisite for this is the constraint being an equation. As most of the simulations in fact do result in a situation where the budget is fully spent, for (2), equality is being assumed.

The interesting aspect is the change of the prices. From (7) it can be deduced that a change in the price of good k is

$$d p_k = \sum_{j=1}^n \frac{\partial p_k}{\partial q_{j,k}} d q_{j,k} = m_k \sum_{j=1}^n d q_{j,k}. \quad (23)$$

It can thus be seen that a change in the price of a good is caused by a change in the demand for this good of one or more of the agents.

A change in utility becomes

$$d V_i = \underbrace{\sum_{k=1}^{m-1} \frac{\partial V_i}{\partial q_{i,k}} d q_{i,k}}_{\text{direct effect}} + \underbrace{\sum_{k=1}^m \left(\frac{\partial V_i}{\partial p_k} m_k \sum_{j=1}^n d q_{j,k} \right)}_{\text{indirect effect}}. \quad (24)$$

Abbreviating the change of aggregate demand for good k , $\sum_{j=1}^n d q_{j,k}$ as $d Q_k$, (24) can be simplified:

$$d V_i = \underbrace{\sum_{k=1}^{m-1} \frac{\partial V_i}{\partial q_{i,k}} d q_{i,k}}_{\text{direct effect}} + \underbrace{\sum_{k=1}^m \left(m_k \frac{\partial V_i}{\partial p_k} d Q_k \right)}_{\text{indirect effect}}. \quad (25)$$

Focusing on the indirect effect, it is now easy to see that for each market k , it is the parameter m_k that decides the impact of the indirect effect on an agent's utility.

If, for example, $m_k = 0$, there is no indirect effect at all. Only the agent herself has an influence on her utility. In other words: For $m_k = 0$, there is no state dependency. Notice, that this case is the case of fixed prices, which leads to relatively good results for all three types of algorithms.

If, on the contrary, m_k is very high, there is also a large influence of all other agents on each agent's utility. This is the case of high state dependency. This case is equal to the case of flexible prices and low elasticity of supply. In this case, due to the high degree of state dependency, the noise in the information caused is strong and consequently the learning results are relatively bad.

Summarizing, the higher m (i.e. the lower elasticity of supply), the higher is the impact of state dependency on the change of each agent's utility. This makes it harder for an agent to recognize the impact of her own consumer plan on her economic success which in turn makes it harder to learn a sensible solution to the consumer choice problem.

5.2 The Influence of Different Learning Schemes

5.2.1 Election and the Concept of Potential Fitness

The algorithm using the election operator seems to perform very good in the FIXPRICE and in the HIGHELASTICITY cases, whereas in the LOWELASTICITY case performance is extremely poor. In order to find out the reasons for this behavior, it is appropriate to recall the central working principle of election. During election, two agents meet and jointly try to find a new strategy. They do this by, among others, calculating a so called potential fitness for the newly created strategies, which should help to find out which strategy is the best one. The calculation of potential fitness requires the knowledge about all the influences on future economic success of the new strategies. As it is impossible to know about all these influences, all these influences are assumed to be unchanged since the last period of time.

This means, that the concept of potential fitness is basically a concept of ‘ceteris paribus’ fitness: An agent calculates the potential fitness of her strategy assuming all the other agents will not change their behavior. This means that agent i , while finding potential fitness, assumes that $d q_{j,k} = 0 \forall i = \{1, \dots, n\}, i \neq j$. In other words, most of the indirect influence on actual fitness is neglected. This, of course, is absolutely correct for situations without state dependency, like the FIXPRICE case, and this is still quite good for low degrees of state dependency like in the HIGHELASTICITY case. But, the more important the indirect effect, the more severe becomes the difference between actual fitness and potential fitness. In situations with high state dependency, e.g. in the LOWELASTICITY case, this may lead to systematically wrong strategy choices, as can be seen from the simulation results in figure 10.

5.2.2 Preselection

In contrast to election, preselection is not the end of the learning process in a period, but the beginning. Whereas in the process of election, the strategy resulting from the election process is the one to be applied in the market, the strategy resulting from preselection is subject to learning by communication (crossover) and experiment (mutation), before it is used at the market. This means that in situations with only little state dependency, preselection – like election – has the advantage of choosing between two strategies (the all time best and last strategy used last period), but the ‘pure’ result of the preselection process can be slightly changed during the following two learning steps, which may be a disadvantage. But, these following two learning steps become the great advantage of preselection in situations with high degrees of state dependency. This means, that different from election learning, preselection learning does not get stuck in strategies that are suc-

cessful only due to potential but not to actual fitness, but can still change the preselected strategy in each period.

6 Summary

The paper employs three different learning algorithms in order to find out, if boundedly rational agents can learn to choose the optimal consumption bundle. In the model, the problem is complicated by allowing for flexible prices and allowing for violation of the budget constraint.

It can be found that it is relatively easy to learn not to break the budget constraint, but that it seems to be quite complicated to find the optimal consumption bundle. Simulation results show that the problem becomes even worse if the elasticity of supply decreases. It is shown that a decrease in the elasticity of supply means an increase in the degree of state dependency. For some types of algorithm, state dependency works like an external effect: Agents do not include it into their calculations for their future strategy, thus being badly mistaken in each future period of time.

Can boundedly rational agents learn the optimal consumer choice? — They can, if the problem is not too complicated. And they do even better if they have some memory of the past.

This is all of the message this paper can give, but at least this seems to be a better story to tell the first year students if they ask again ...

A Data

The following tables show data from one randomly chosen simulation each. All simulations were run for 1 000 periods with a population size of $n=500$. The data contains information about rounds 501 to 1 000 to eliminate possible startup effects of the algorithms.

FIXPRICE

The t -values are test statistics of an approximate Gauss tests¹² for $q_1 = 20$, $q_2 = 30$, $q_3 = 50$ and $\rho = 0$, respectively.

¹²This is a t -test for large populations.

		Canonical	Election	Preselection
q_1	Mean	20.2484	20.0492	20.2077
	Var.	0.515554	0.000617538	0.0153379
	t	7.73661	44.2923	37.4965
q_2	Mean	30.8866	30.0306	29.9948
	Var.	1.36406	0.00149178	0.00520296
	t	16.9738	17.7022	-1.62135
q_3	Mean	47.4426	49.9195	49.7968
	Var.	1.33563	0.00358468	0.00447201
	t	-49.481	-30.0562	-67.9504
ρ	Mean	1.4224	0.000674724	0.000770113
	Var.	0.453434	3.88333 E(-7)	0.000613616
	t	47.2334	24.2108	0.69517

HIGHELASTICITY

The t-value is a test statistic of an approximate Gauss test for $\rho = 0$, the values t_{12} , t_{13} , and t_{23} are test statistics of approximate Gauss test for $MU_1/\$ = MU_2/\$$, $MU_1/\$ = MU_3/\$$ and $MU_2/\$ = MU_3/\$$, respectively.

		Canonical	Election	Preselection
$MU_1/\$$	Mean	0.178931	0.185862	0.189568
	Var.	0.0000532438	8.59688 E(-7)	7.43989 E(-6)
$MU_2/\$$	Mean	0.178914	0.183026	0.181765
	Var.	0.0000914054	2.21833 E(-7)	3.54032 E(-6)
$MU_3/\$$	Mean	0.196372	0.184637	0.183881
	Var.	0.0000684217	2.27176 E(-7)	6.94431 E(-7)
t_{12}		0.0316065	60.9781	52.6552
t_{13}		-35.3568	26.2744	44.5869
t_{23}		-30.8784	-53.7592	-22.9925
ρ	Mean	4.1415	0.0108742	-0.0000875281
	Var.	0.63692	0.0000392457	0.00494539
	t	116.038	38.8137	-0.0278312

LOWELASTICITY

		Canonical	Election	Preselection
$MU_1/\$$	Mean	0.00531371	0.000225087	0.00901293
	Var.	3.19915 E(-7)	0	2.0274 E(-7)
$MU_2/\$$	Mean	0.00768015	0.00624466	0.00837344
	Var.	6.2859 E(-7)	0	9.85724 E(-8)
$MU_3/\$$	Mean	0.0131581	0.00156332	0.0080648
	Var.	1.04098 E(-6)	0	4.02153 E(-8)
t_{12}		-54.3326		26.0501
t_{13}		-150.36		43.0121
t_{23}		-94.7983		18.5254
ρ	Mean	13.3108	-5453.53	-0.538288
	Var.	26.2627	0	4.07375
	t	58.0792		-5.96352

References

- Arifovic, Jasmina (1994). Genetic Algorithm Learning and the Cobweb-Model. *Journal of Economic Dynamics and Control*, **18**, 3–28.
- Bäck, Thomas, Hoffmeister, Frank, and Schwefel, Hans-Paul (1991). A Survey of Evolution Strategies. In: *Proceedings of the 4th International Conference on Genetic Algorithms*, edited by Belew, Richard K. and Booker, Lashon B., pp. 2–9, San Mateo, CA. Morgan Kaufmann.
- Barro, R. J. (1990). Government Spending in a Simple Model of Endogenous Growth. *Journal of Political Economy*, **98**, S103–S125.
- Birchenhall, Chris, Kastrinos, Nikos, and Metcalfe, Stan (1997). Genetic Algorithms in Evolutionary Modelling. *Journal of Evolutionary Economics*, **7**, 375–393.
- Buchanan, James M. and Stubblebine, Wm. Craig (1962). Externality. *Economica*, **29**, 371–384.
- Dawid, Herbert (1999). *Adaptive Learning by Genetic Algorithms*. Springer, Berlin, Heidelberg, New York, 2nd edn.

- Franke, Reiner (1997). Behavioural Heterogeneity and Genetic Algorithm Learning in the Cobweb Model. Discussion Paper 9, IKSF — Fachbereich 7 — Wirtschaftswissenschaft. Universität Bremen.
- Goldberg, David E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison–Wesley, Reading, Massachusetts.
- Henderson, James M. and Quandt, Richard E. (1986). *Microeconomic Theory. A Mathematical Approach*. McGraw–Hill, New York, 3rd edn.
- Lucas, Robert E. (1988). On the Mechanics of Economic Development. *Journal of Monetary Economics*, **22**, 3–42.
- Michalewicz, Zbigniew (1996). *Genetic Algorithms + Data Structures = Evolution Programs*. Springer, Berlin, Heidelberg, New York, 3rd edn.
- Mitchell, Melanie (1996). *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, London.
- Riechmann, Thomas (1998). Learning How To Learn. Towards an Improved Mutation Operator within GA Learning Models. In: *Computation in Economics, Finance and Engineering: Economic Systems*, Cambridge, England. Society for Computational Economics.
- Riechmann, Thomas (1999a). Genetic Algorithm Learning and Evolutionary Games. *Journal of Economic Dynamics and Control*, forthcoming.
- Riechmann, Thomas (1999b). Learning and Behavioral Stability — An Economic Interpretation of Genetic Algorithms. *Journal of Evolutionary Economics*, **9**, 225–242.
- Romer, Paul M. (1986). Increasing Returns and Long Run Growth. *Journal of Political Economy*, **94**, 1002–1037.