International University in Germany

School of Business Administration

A Global Optimization Heuristic for Estimating Agent Based Models

Manfred Gilli and Peter Winker

Working Paper 12/2002

Bruchsal, March 2002

Prof. Dr. Manfred Gilli is professor at the Department of Econometrics at the University of Geneva, Switzerland, where he teaches numerical methods in economics and finance. He is a member of the Advisory Committee of the Society for Computational Economics.

Prof. Dr. Peter Winker is associate professor of Economics and Quantitative Methods at the School of Business Administration at the International University in Germany, Bruchsal.

Contact: Peter.Winker@i-u.de, Phone: +49-7251-700-343, Fax: +49-7251-700-350

A Global Optimization Heuristic for Estimating Agent Based Models

M. Gilli [1]

Département d'Econométrie, Université de Genève, CH–1211 Geneva 4, Switzerland

P. Winker

Department of Economics and Quantitative Methods, International University in Germany, D–76646 Bruchsal

Abstract

We introduce a continuous global optimization heuristic for a stochastic approximation of an objective function, which itself is not globally convex. The objective function arises from the simulation based indirect estimation of the parameters of agent based models of financial markets. The function is continuous in the variables but non–differentiable. Due to Monte Carlo variance, only a stochastic approximation of the objective function is available. The algorithm combines features of the Nelder–Mead simplex algorithm with those of a local search heuristic, called threshold accepting. The Monte Carlo variance of the simulation procedure is also explicitly taken into account. We present details of the algorithm and some results of the estimation of the parameters for a specific agent based model of the DM/US–$ foreign exchange market.

Key words: Global Optimization, Threshold Accepting, Simplex Algorithm, Agent Based Models, Indirect Estimation, Validation.

Email addresses: `Manfred.Gilli@metri.unige.ch` (M. Gilli), `Peter.Winker@i-u.de` (P. Winker).

[1] Supported by the Swiss National Science Foundation (projects 12–5248.97 and 1214–056900.99/1). Thanks to Evis Këllezi for comments and suggestions.

Accepted for publication in Computational Statistics & Data Analysis

1 Introduction

The validation or estimation of agent based models of financial markets by means of simulation methods based on actual data leads to a highly complex optimization problem. We introduce a continuous global optimization heuristic for the resulting globally non convex objective function which can be evaluated, for a given set of parameters, only up to some Monte Carlo variance. Although our first implementation is targeted at the specific problem of indirect estimation, the proposed algorithm might be considered for any continuous non convex optimization problem.

Agent based models gained increasing interest in different areas of economic modelling as reflected in the special issues of *Computational Economics* and the *Journal of Economic Dynamics & Control* (Tesfatsion, 2001*a*, 2001*b*). For the modelling of financial markets, agent based models offer at least two advantages. First, they allow a more realistic modelling of agents' behavior than standard representative agent models (Fama, 1970). At least for specific parameter settings, the results of the simulations of such agent based models show some statistical similarities with actual foreign exchange time series, e.g. fat tails and volatility clustering.

Second, the rapid increase in available computing resources allows for more extensive simulation studies of agent based models. Thus, it became possible to assess the impact of changes in parameter values on the simulated time series. Nevertheless, "validation ... remains a very weak area for the class of models described here" (LeBaron, 2000, p. 698), i.e. a systematic assessment of such simulation models based on actual market data is required.

An indirect estimation approach for validation purposes, discussed in more detail in our companion paper Winker and Gilli (2001) and sketched in Section 2, leads to a highly complex optimization problem. It is based on the comparison of the moments of actual data and simulated moments of the agent based model. This involves the minimization of a stochastic approximation of an objective function $f(\theta)$

$$\min_{\theta \in \Theta} f(\theta)$$

where Θ is a subspace of $\mathbb{R}^d$ describing the possible parameter values, and $f(\theta)$ measures the distance between empirical moments of the financial market data and those of the simulated time series.

Given that the objective function f is not globally convex and can only be approximated up to some Monte Carlo variance, standard optimization algorithms including gradient methods and the Nelder–Mead simplex algorithm (Lagarias *et al.*, 1999) do not produce useful results even if a high number of replications is used for the Monte Carlo simulation step.

Therefore, we suggest an algorithm which is a combination of the Nelder–Mead simplex direct search method (Lagarias *et al.*, 1999) and the threshold accepting optimization heuristic (Dueck and Scheuer, 1990; Winker, 2001). The Nelder–Mead search enables the algorithm to chose efficient steps for a continuous but non differentiable objective function and the threshold accepting strategy avoids the algorithm to be trapped in the many local minima the objective function has due to the simulation variance.

In the Nelder–Mead simplex algorithm each current solution consists of the points of a simplex in the parameter space Θ. Comparing the values of the objective function f at the edges of a simplex, a search direction – pointing towards smaller values of the objective function – is identified. If along this search direction a point with improved value of the objective function can be found, it will replace one of the old simplex edges. Otherwise, reflection and shrink operations concentrate the search inside the existing simplex. While this method appears to be efficient for smooth globally convex functions, in our problem, because of the Monte Carlo variance, the simplex shrinks rapidly and gets stuck in a region close to the starting point.

The standard threshold accepting algorithm considers only a single element $\theta^c \in \Theta$. In each iteration, a neighbor of this current solution is randomly selected and the value of the objective function is evaluated. If the new point represents an improvement, it is accepted as new current solution. However, even if the objective function value should increase, it will be still accepted as new current solution as long as this worsening of the objective function does not exceed a predefined threshold. Consequently, the threshold accepting algorithm can escape the local minima that result from Monte Carlo variance and non convexity.

Given that the underlying objective function f appears to be rather smooth, we expect improved efficiency of a combined algorithm. For this purpose, we introduce the idea of temporary worsening of the objective function in the simplex algorithm. We will present the details of the algorithm and some results for the specific agent based model introduced by Kirman (1991) and analyzed in an indirect estimation approach by Winker and Gilli (2001) applied to the DM/US–$ foreign exchange market.

The paper is organized as follows: In Section 2 we sketch the underlying problem of validating agent based models and describe the resulting objective function f. The optimization algorithm is described in Section 3. The estimation results of some parameters for a standard agent based model of the DM/US–$ exchange rate are summarized in Section 4, while Section 5 provides a conclusion and the outlook on further research, both with regard to the algorithm and its application for validating agent based models of financial markets.

2 Indirect Estimation of Agent Based Models

We start with a short presentation of the agent based model for which we want to estimate some parameters. Then, we derive the objective function which has to be minimized with respect to the parameters under consideration. More details on the indirect estimation approach can be found in Winker and Gilli (2001).

We use a model originally introduced by Kirman (1991, 1993) that stresses the importance of interaction between heterogenous, not fully rational individuals. The model assumes that the individuals acting on the foreign exchange market differ with regard to their behavior. A first type of agents acts on fundamentals, i.e. they expect an adjustment of prices towards the fundamental value of the asset which is assumed to be known to all agents. The second type of agents follows a momentum strategy or chartist rule. An agent of this type expects price changes to be the same as in the previous period.

Besides this heterogeneity of agents' behavior already introduced in Frankel and Froot (1986), an important feature of the model is that the type of any individual agent can change over time. This may happen for two reasons. Either the agent undergoes a random mutation with probability ε or he is convicted by a direct interaction with a second individual. Conviction is modelled as a discrete process in time. At any time period two randomly selected agents meet. The second one will convince the first one of his point of view, i.e. change the type of the first agent to his type, with a given probability δ. Consequently, due to mutation and conviction, the share of fundamentalists and chartists in the population of agents changes over time. It turns out that this simple model of endogenous agent types coupled with a price process depending on the type of agents in the population is able to generate quite complex dynamics including ARCH–effects and excess kurtosis.

The Algorithm 1 summarizes the price generation process in Kirman's model resulting from n_I interactions of n_A agents of two types, i.e. fundamentalist (F) and chartists (C).

Algorithm 1 Kirman's model.

```
1: for i = 1 : n_I do
2:     Generate q_i, the share of fundamentalists F in the market
3:     Transform share q_i into w_i (perceived weight of fundamentalists in market)
4:     Compute market price p_i by weighting the agents price expectations
5: end for
6: Observe p_i, i = 1 : n_T : n_I     (every n_T interactions between agents)
```

The three steps in the algorithm, i.e. the generation of the share of fundamentalists, the transformation of this share into a perceived weight and the computation of the market price expectation are presented in Algorithm 2. In Statement 3 in Algorithm 2 nothing happens, of course, if agent a_1 and agent a_2 already follow the same strategy.

Algorithm 2 Process defining q_i, the share of fundamentalists.

1: Given the set F after interaction $i-1$ $\quad (A = F \cup C$ and $F \cap C = \varnothing)$
2: Randomly select three agents a_1, a_2, $a_3 \in A$
3: a_2 convinces a_1 with probability δ to follow his strategy (Direct interaction)
4: a_3 changes his strategy with probability ε (Random mutation)
5: $q_i = |F|/|A|$

To compute, in Statement 3 of Algorithm 1, the perceived weight ω_i of fundamentalists at time i in the market, it is assumed that agents have information about q_i in the form of a noisy signal (Kirman, 1991). This signal $\tilde{q}_i$ is supposed to be normally distributed, $\tilde{q}_i \sim N(q_i, \sigma_q^2)$ and the knowledge about the proportion ω_i of fundamentalist at time i in the market is then defined as the following probability

$$\omega_i = P(\tilde{q}_i > 1/2)\,.$$

Finally, in Statement 3 of Algorithm 1, the market price p_i is computed by weighting the agents price expectations. For the fundamentalist the expected change in the price is given by $E^F(\Delta p_i) = \nu(\bar{p} - p_{i-1})$, where $\bar{p}$ is the fundamental value of the underlying asset and ν denotes the error correction term. The chartists are assumed to extrapolate last period prices, i.e. their expectations are described by $E^C(\Delta p_i) = p_{i-1} - p_{i-2}$. Combining the two models yields the market price expectation $E^M(\Delta p_i) = \omega_i\, \nu\, (\bar{p} - p_{i-1}) + (1 - \omega_i)(p_{i-1} - p_{i-2})$. Given a specific utility function and initial wealth of market participants, Frankel and Froot (1986) derive the price formation process as a weighted average of these market expectations and the fundamental value. Thereby, the weight assigned to market expectations depends on the degree of risk aversion implicit in the utility function, which is a further parameter of the model. Kirman (1991) adds an exogenous perturbation $u_i \sim N(0, \sigma_s^2)$ to the price process.

Table 1 provides an overview on the most important parameters of the model. In order to keep the problem computationally feasible, we decided to start with the estimation of only the three parameters δ, ε and σ_q describing the dynamics of agent types. The other parameters have been kept fixed for the present first implementation according to their values given in Table 1.

Table 1
Parameters of the agent based model.

Label	Interpretation	Value
n_A	Number of agents	100
n_I	Number of interactions	50 000
n_T	Number of interactions per trading day	50
ν	Adjustment speed in fundamentalists expectations	0.045
σ_s	Standard deviation of price shocks	0.25
δ	Probability for direct interaction	
ε	Probability for random mutation	
σ_q	Standard deviation of noise in majority assessment	

It should be noted that the goal of the present paper is not a discussion of the agent based model itself, but of a method for estimating the parameters based on actual data. Algorithm 3 resumes the main steps of the indirect estimation procedure. For this first implementation, the moments of the actual data to

Algorithm 3 Indirect estimation procedure

```
1: Give x^(0) ∈ R^n initial vector of parameters to be estimated
2: while not converged do
3:    Determine successive vectors x (defined by the minimization algorithm)
4:    for each x do
5:       Evaluate objective function f(x) =| k_d^ag − k_d^emp | + λ | α_1^ag − α_1^emp |
6:    end for
7: end while
```

be matched are the estimated ARCH(1)–effect α_1^{emp} and the empirical kurtosis k_d^{emp} of the daily logarithmic returns of the DM/US-$ exchange rate over the sample period 11.11.1991 – 8.11.2000. The objective function to be minimized is

$$f = \mid k_d^{\text{ag}} - k_d^{\text{emp}} \mid + \lambda \mid \alpha_1^{\text{ag}} - \alpha_1^{\text{emp}} \mid , \tag{1}$$

where k_d^{ag} and α_1^{ag} denote the moments of the time series generated from the agent based model and the weight $\lambda = 15$ given to the second component was chosen ad hoc based on the relative magnitude of the two moments. Of course, the algorithm presented in the sequel is flexible enough to be applied to objective functions based on different moments of the data, e.g. estimates of SETAR–models (Killian and Taylor, 2001) or quantiles of the whole (conditional) distribution of returns. This will be a subject of our future research.

3 Algorithm

3.1 Simulation and Monte Carlo Variance

Let us first explain in more detail how the stochastic approximation of the objective function is obtained in Statement 5 of the indirect estimation procedure Algorithm 3. This is illustrated with the Algorithm 4 given hereafter.

Algorithm 4 Stochastic approximation of the objective function $f(x)$.

```
Given x ∈ R^n a particular value of the vector of parameters to be estimated
for j = 1 : R do
    Generate random sequences for price simulation
    Simulate price path p^(j), compute returns r^(j) and corresponding α̂^ag_{1_j} and k̂^ag_{d_j}
end for
Compute α̃^ag_1, k̃^ag_d, the truncated means of α̂^ag_{1_j} and k̂^ag_{d_j}, j = 1, ..., R
Compute f̃(x) =| k̃^ag_d − k^emp_d | + λ | α̃^ag_1 − α^emp_1 |
```

Figure 1 shows the resulting stochastic approximation $\tilde{f}$ of the objective function f against the parameters ε and δ for $\sigma_q = 0.20$ and holding all the other parameters of the model fixed at their values shown in Table 1.

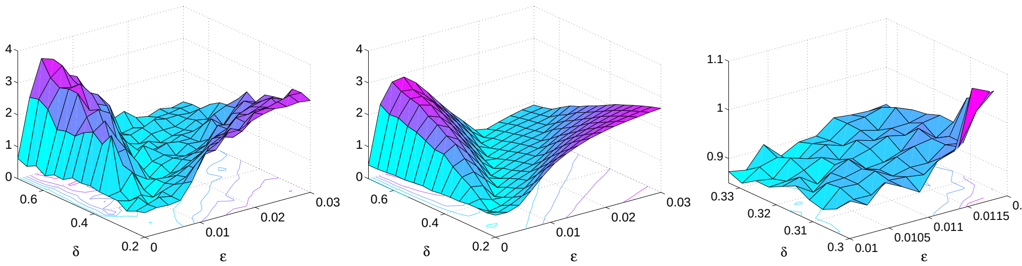

Fig. 1. Simulated values of f against ε–δ–grid. Left panel $R = 200$, central panel $R = 10\,000$, right panel $R = 10\,000$ with finer grid.

The left plot provides the results for $R = 200$ Monte Carlo replications. The considerable Monte Carlo variance of the estimates of $f(\varepsilon, \delta)$ is evident, while using $R = 10\,000$ replications leads to a much smoother plot (central panel). In fact, regression analysis indicates that the Monte Carlo variance shrinks at the usual rate of $\sqrt{R}$. However, the estimates of f are still not smooth enough to allow the application of standard optimization tools. This is made visible with the plot in the right panel of Figure 1 where we use a ten times finer grid to represent $\tilde{f}$ computed with $R = 10\,000$ replications.

Furthermore, the estimation problem is not restricted to the two dimensional problem represented in Figure 1, but, in general, will include some more parameters of the agent based model. For the current application it includes the parameter σ_q besides ε and δ.

The shape of the objective function suggests the need for a global optimization technique for our minimization problem. Unfortunately, the evaluation (approximation) of the objective function is computationally expensive: For n_I being the number of simulated interactions and R the number of Monte Carlo replications, the complexity for approximating f is $\mathcal{O}(n_I R)$. In our application we have $n_I R \propto 10^7$. Therefore, we have to think of an optimization approach which keeps the number of evaluations as low as possible. For this purpose, we consider the simplex search method which was first introduced by Spendley *et al.* (1962) and has then been modified by Nelder and Mead (1965), while Lagarias *et al.* (1999) provide results about convergence of the method. The simplex search approach enables the algorithm to choose efficiently search steps. However due to the rough surface of the objective function the simplex search would soon get trapped in one of the many local minima. Therefore, the simplex search will be combined with a threshold accepting strategy.

3.2 The Simplex Search

Figure 2 illustrates the fundamental idea of the simplex search for an objective function with $n = 2$ variables. The objective function is evaluated for $n+1$ points $x^{(1)}$, $x^{(2)}$ and $x^{(3)}$, defined by the vertices of a simplex (a triangle for $n = 2$). We then construct a new simplex, adjacent to the old one, which is closer to the minimum of the objective function. The adjacent simplex is obtained by projecting the vertex $x^{(3)}$ corresponding to the worst function value through the mean $\bar{x}$ of the other $x^{(i)}$, $i = 1, \ldots, n$ vertices. The new simplex is then given by the vertices $x^{(1)}$, $x^{(2)}$ and x^{R}.

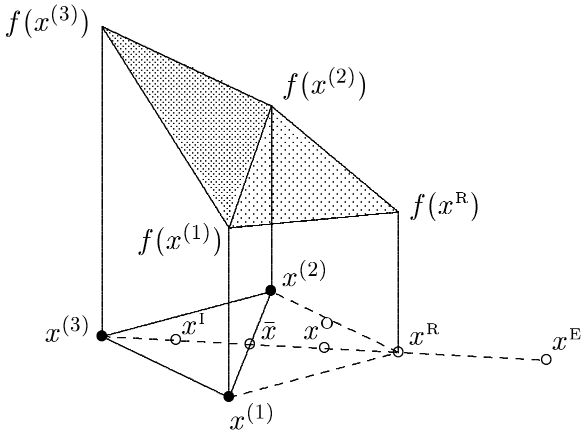

Fig. 2. Step in the simplex search algorithm.

Let us now give a general formalization of the algorithm. We denote $x^{(i)}$ the n-dimensional vector of the coordinates of vertex i of the simplex, where the vertices are numbered such as to verify $f(x^{(1)}) \leq f(x^{(2)}) \leq \ldots \leq f(x^{(n+1)})$. In order to compute the appropriate displacement of the simplex we need additional points defined in Table 2.

Table 2
Points computed for a given simplex in Algorithm 5.

$\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x^{(i)}$	Mean
$x^{\mathrm{R}} = (1+\rho)\,\bar{x} - \rho\, x^{(n+1)}$	Reflection
$x^{\mathrm{E}} = (1+\rho\xi)\,\bar{x} - \rho\xi\, x^{(n+1)}$	Expansion
$x^{\mathrm{O}} = (1+\psi\rho)\,\bar{x} - \psi\rho\, x^{(n+1)}$	Out-contraction
$x^{\mathrm{I}} = (1-\psi\rho)\,\bar{x} + \psi\rho\, x^{(n+1)}$	In-contraction
$x^{(i)} = x^{(1)} - \sigma(x^{(i)} - x^{(1)}) \quad i = 2,\ldots,n+1$	Shrink simplex toward vertex $x^{(1)}$

Using these additional points, the overall simplex search procedure is given by Algorithm 5. If the reflection point x^{R} represents an improvement over the so far best value of the objective function $f(x^{(1)})$, the algorithm tries to find an even better value by expanding the search in the direction defined by x^{R}. If expansion is successful, x^{E} will replace the worst vertex of the current simplex, otherwise the original reflection point x^{R} will be used. It is also used to replace the worst vertex, if it represents an improvement compared to $f(x^{(n)})$, i.e. the vertex with the second highest value of the objective function.

Algorithm 5 Simplex search.

Construct vertices $x^{(1)}, \ldots, x^{(n+1)}$ of starting simplex
repeat
  Rename vertices such that $f(x^{(1)}) \leq \ldots \leq f(x^{(n+1)})$
  if $f(x^{\mathrm{R}}) < f(x^{(1)})$ **then**
    if $f(x^{\mathrm{E}}) < f(x^{\mathrm{R}})$ **then** $x^* = x^{\mathrm{E}}$ **else** $x^* = x^{\mathrm{R}}$
  else
    if $f(x^{\mathrm{R}}) < f(x^{(n)})$ **then**
      $x^* = x^{\mathrm{R}}$
    else
      if $f(x^{\mathrm{R}}) < f(x^{(n+1)})$ **then**
        if $f(x^{\mathrm{O}}) < f(x^{(n+1)})$ **then** $x^* = x^{\mathrm{O}}$ **else** shrink
      else
        if $f(x^{\mathrm{I}}) < f(x^{(n+1)})$ **then** $x^* = x^{\mathrm{I}}$ **else** shrink
      end if
    end if
  end if
  if not shrink **then** $x^{(n+1)} = x^*$ (Replace worst vertex by x^*)
until stopping criteria verified

If the reflection point x^{R} provides only an improvement compared to the worst vertex, i.e. $f(x^{\mathrm{R}}) < f(x^{(n+1)})$, but $f(x^{\mathrm{R}}) \geq f(x^{(n)})$, an outside contraction is tested, i.e. moving only a bit in direction of x^{R}. If x^{R} is even worse than $x^{(n+1)}$ an inside contraction is tested. If outside or inside contraction fail to result in an improvement over $f^{(n+1)}$, the simplex shrinks. The conditions for shrinking of the simplex are met for a smooth function if the simplex moves closer to a local minimum. However, in case of the stochastic approximation used in the current application, shrinking may also result due to Monte Carlo variance. Consequently, the simplex tends to shrink very fast and the algorithm gets stuck without approaching any local or global minimum.

3.3 Threshold Accepting

Threshold accepting (TA) was introduced by Dueck and Scheuer (1990) and is one among the many existing heuristic optimization techniques used with increasing success in various disciplines (e.g. Pardalos *et al.* (2000); Pardalos and Resende (2002)). The TA algorithm has the advantage of an easy parameterization, it is robust to changes in problem characteristics and works well for many problem instances. An extensive introduction to TA is given in Winker (2001). TA can be considered as a deterministic analog to simulated annealing (Kirkpatrick *et al.*, 1983), where the stochastic acceptance criterion of simulated annealing is replaced by a deterministic rule. TA is a refined local search procedure which escapes local minima by always accepting solutions which are not worse than a current solution by more than a given threshold τ. A typical implementation uses an a priori number of rounds n_R and explores the local structure of the objective function with a fixed number of steps n_S during each round. The threshold τ is decreased successively for each round and reaches the value of zero in the last round.

If we formalize the optimization problem as $f : \mathcal{X} \to \mathbb{R}$ where $\mathcal{X}$ is either a continuous or a discrete set and where we may have more than one optimal solution defined by the set

$$\mathcal{X}_{\mathrm{min}} = \{x \in \mathcal{X} \mid f(x) = f_{\mathrm{opt}}\} \qquad \text{with} \qquad f_{\mathrm{opt}} = \min_{x \in \mathcal{X}} f(x)\,,$$

the TA heuristic described in Algorithm 6 will hopefully, after completion, provide us with a solution $x^* \in \mathcal{X}_{\mathrm{min}}$ or a solution close to an element in $\mathcal{X}_{\mathrm{min}}$. In fact, Althöfer and Koschnick (1991) show that the algorithm converges to a solution close to a global optimum with probability approaching one with the number of iterations growing to infinity. The complexity of the algorithm is $\mathcal{O}(\sum_{r=1}^{n_R} n_{S_r})$ with the factor of proportionality depending on the complexity of evaluating f.

Given the high computational cost of obtaining a high quality approximation to f by means of a high number R of Monte Carlo replications of the

Algorithm 6 Threshold accepting algorithm.

```
1: Initialize n_R, n_{S_r} and the sequence of thresholds τ_r, r = 1, 2, ..., n_R
2: Generate starting point x_old ∈ 𝒳
3: for r = 1 to n_R do
4:   for i = 1 to n_{S_r} do
5:     Generate x_new ∈ 𝒩_{x_old}    (neighbor of x_old)
6:     if f(x_new) < f(x_old) + τ_r then
7:       x_old = x_new
8:     end if
9:   end for
10: end for
```

agent based model, the optimization heuristic must resort to a limited number of function evaluations and try to hold the requirements for R as small as possible. Given the apparent smoothness of the objective function f, it seems promising to combine ideas of simplex search with the TA approach. Consequently, we use the following strategies.

First, in the beginning of the optimization procedure a coarse approximation $\tilde{f}$ to the objective function f is employed requiring only a small number of replications R in the Monte Carlo simulation procedure (Algorithm 4). The resulting high variance of $\tilde{f}$ is compensated by larger values of the threshold parameter τ_r, which is chosen proportionally to the estimated Monte Carlo variance of $\tilde{f}$. As the algorithm proceeds, τ_r is reduced and R is increased allowing for a more accurate estimation of f. The unavoidable Monte Carlo variance is still taken into account for by the threshold height τ_r, which also allows to escape local minima of the non globally convex objective function f.

Second, instead of working on a single current solution, the TA algorithm is applied to a simplex. In order to obtain a local improvement of the solution, two methods are employed which are chosen with a given probability ξ. The first strategy consists of a random shift of the simplex in the parameter space, where the length and the direction of the shift are chosen randomly. This neighbor solution is accepted, if the best vertex of the new simplex is not worse than the best vertex of the previous one by more than the given threshold τ_r. The second method, which is used with probability $(1-\xi)$ is a simplex search step as described in Algorithm 5. However, in the Statements 4, 5, 7, 10, 11 and 13 of the simplex search algorithm (5), where we compare two function values, we augment the function on the right–hand side by the threshold τ in order to allow for a temporary worsening of the simplex. Both methods are implemented by replacing Statements 5–8 in the original TA algorithm 6 by the code given in Algorithm 7.

This first implementation of the hybrid algorithm requires the choice of some parameters like the number of rounds and the number of Monte Carlo replications for the different rounds. However, we assume that it will be possible

Algorithm 7 Step i in round r for TANM.

Draw uniform random variable u
if $u < \xi$ **then**
  Generate $x_{\text{new}} \in \mathcal{N}_{x_{\text{old}}}$ with random shift
  if $f(x_{\text{new}}) < f(x_{\text{old}}) + \tau_r$ **then** x_{new} is accepted
else
  Generate $x_{\text{new}} \in \mathcal{N}_{x_{\text{old}}}$ with simplex search and check whether it is accepted
end if
if x_{new} is accepted **then**
  $x_{\text{old}} = x_{\text{new}}$
end if

to transfer ideas from TA implementations for optimization problems on discrete sets to the current implementation allowing for an automatic data driven parameter setting.

4 Estimation Results for an Agent Based Model

We will now present the results for the indirect estimation of the three parameters ε, δ and σ_q. The estimated parameters minimize the objective function defined in (1). In order to benchmark the performance of our optimization heuristic we computed the approximations of the objective function for a grid with 32 steps in each of the directions of the three parameters. The approximations of the objective function on this grid have been computed with $R = 500$ repetitions and the overall computing time for the $\sim 10^{12}$ prices (500×32^3 price path with 50 000 interactions) was 7 days on a Pentium III 600 MHz PC. The plot of the objective function in the ε–δ and ε–σ_q grid is given in Figure 3 and the plot in the δ–σ_q grid is given in the left panel of Figure 4.

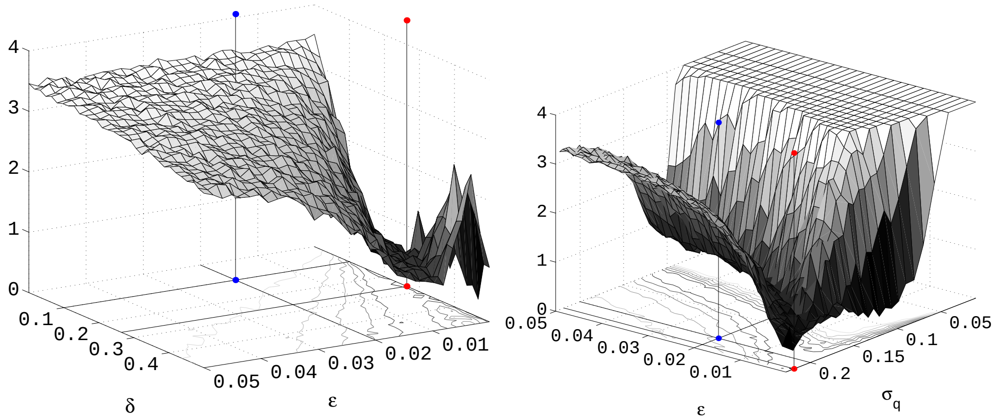

Fig. 3. Approximations of the objective function in the ε–δ, ($\sigma_q = 0.219$) and ϵ–σ_q ($\delta = 0.264$) grid.

For the TA implementation, we choose $n_R = 3$ for the number of rounds, and

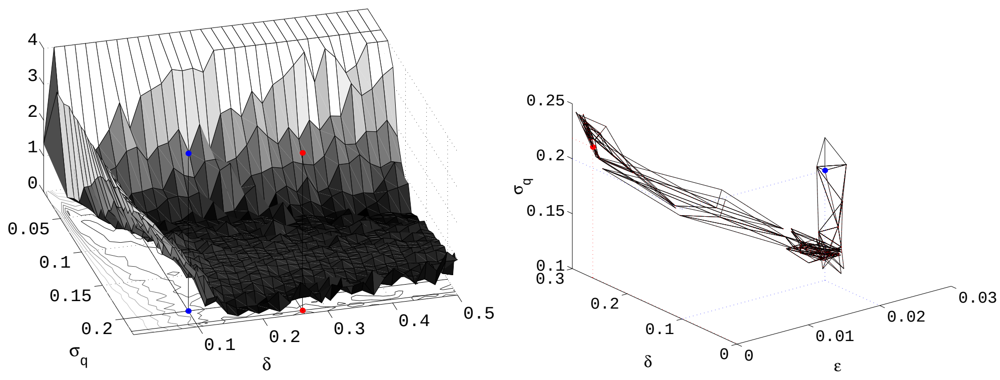

Fig. 4. Left panel: Approximations of the objective function in the δ–σ_q ($\varepsilon = 0.0001$) grid. Right panel: Evolution of the simplex in the parameter space.

a variable number of steps $n_{S_1} = 50$, $n_{S_2} = 20$ and $n_{S_3} = 10$. The number of repetitions R for the approximation of the objective function evolves as follows in the different rounds: $R_1 = 100$, $R_2 = 200$, and $R_3 = 500$. The value for ξ, i.e. the probability for a random shift in Algorithm 7, was set to 0.15. Consequently, with a probability of 85% the construction of an adjacent simplex was performed.

The choice of the threshold sequence τ_r, $r = 1, \ldots, n_R$ has to be made in relation with the Monte Carlo variance of the objective function, which is a function of the number of repetitions R_r used in the simulation which approximates the objective function. This functional dependency has been estimated in a preliminary step using empirical estimates of the Monte Carlo variance $\hat{\sigma}_{f_{R_k}}$ for different numbers of replications R_k. We fitted the model

$$\hat{\sigma}_{f_{R_k}} = \alpha \, R_k^{\beta} + \varepsilon_k$$

to our data. As expected, our estimates for β were very close to and not significantly different from $-1/2$ and for α we obtained estimates very close to 1.3. For a given number of Monte Carlo replications R_r, the threshold value is then determined automatically with the following relation

$$\tau_r = \gamma \, \alpha \, \sqrt{R_r} \qquad r = 1, \ldots, n_R \, ,$$

where only γ has to be set for the implementation. The parameter γ determines how likely the TA algorithm will leave a local minima caused by Monte Carlo variance in a single iteration. The value of γ was set to 2 for the present implementation.

Further implementation details, which have not been mentioned in Section 3, are that we have to check whether the simplex moves out from the feasible region for the parameters and the fact that we must prevent the simplex from expanding too large and shrinking too small. In the application we allow the simplex to expand by a factor of 2 and shrink by a factor of 0.2 from its initial volume. A future refinement of the algorithm will adjust these factors in dependence of the shrinking Monte Carlo variance as the algorithm proceeds.

The best results obtained with these settings are $\hat{\varepsilon} = 0.0001$, $\hat{\delta} = 0.264$ and $\hat{\sigma_q} = 0.219$ with $\tilde{f} = .279$. The computing time for this solution is less than 1 hour on a Pentium III 600 MHz PC. The starting values for the minimization have been $\varepsilon^{(0)} = 0.02$, $\delta^{(0)} = 0.10$ and ${\sigma_q}^{(0)} = 0.20$. Starting values and solutions have been marked in the three plots of the objective function in Figures 3–4.

Asymptotic results on the convergence of Markov chains provided in Kirman (1993) allow for the conclusion, that the process generated from the agent based model exhibits large shares of fundamentalists and chartists, respectively, with high probability, if $\varepsilon < (1-\delta)/(n-1)$. For our estimates we find $\varepsilon = 0.0001 < 0.0074 = (1-0.264)/(100-1) = (1-\delta)/(n-1)$. Thus, our estimates indicate that, in fact, the foreign exchange market can be better characterized by switching moods of the investors than by assuming that the mix of fundamentalists and chartists remains rather stable over time.

We also experimented the algorithm with different starting values and observed that it always ended in the deep portion of the objective function, not necessarily for the same parameter values, as this region is quite flat. In fact, a standard approach when working with optimization heuristics like TA is to use several restarts with different initializations of the random number generator and different starting values (Winker, 2001, pp. 129ff).

Figure 4 and 5 provide some insight about how the algorithm works. The right panel of Figure 4 visualizes the sequence of adjacent simplexes the algorithm constructs in the search to the minimum. In Figure 5 we plot the corresponding values taken by the objective function. Notice that the TA strategy allows to

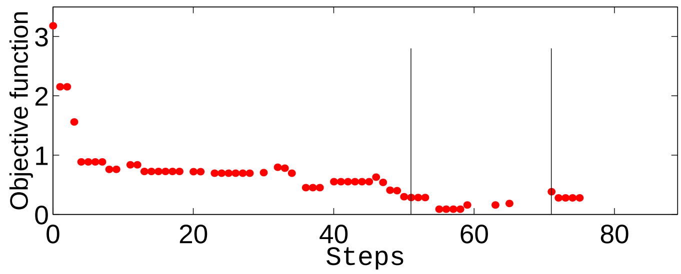

Fig. 5. Value of objective function in the succeeding steps.

accept intermediate solutions for which the objective function increases. Moreover, in each new round, starting at steps marked in Figure 5 with a vertical line, we increase the number R of replications. This reduces the variance of the Monte Carlo approximation of the objective function, making it smoother and thus reducing the depth of local minima due to this variance. This explains why the value of the objective function can increase when we move from one round to another.

5 Conclusion and Further Research

The global optimization heuristic introduced in this paper allows for the indirect simulation based estimation of agent based models of financial markets. It takes into account the fact that moments of the series generated from the simulation models can only be approximated using Monte Carlo methods. The specific features of the threshold accepting algorithm, implemented to overcome local minima due to Monte Carlo variance, also help to obtain close approximations of global minima for the non globally convex objective function.

The present application is restricted to three parameters for a specific model. Further research will concentrate on higher dimensional problems. At the same time, parameter settings of the algorithm will be made data dependent to allow for easy application. We will also implement a restart option based on ideas of uniform design to improve the robustness of the results.

A further field of future research will consist of using the algorithm for validation of agent based models against actual data. For this purpose, we have to consider further moments of the time series. Then, it might become possible to discriminate among different agent based models and improve their structure. Still, citing LeBaron (2002): "...this field is only in its infancy, and much remains to be done."

References

Althöfer, I. and K.-U. Koschnick (1991). On the Convergence of 'Threshold Accepting'. *Applied Mathematics and Optimization* **24**, 183–195.

Dueck, G. and T. Scheuer (1990). Threshold Accepting: A general purpose algorithm appearing superior to Simulated Annealing. *Journal of Computational Physics* **90**, 161–175.

Fama, E. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance* **25**, 383–417.

Frankel, J. A. and K. Froot (1986). The dollar as a speculative bubble: A tale of fundamentalists and chartists. Technical Report 1845. NBER. Cambridge, MA.

Killian, L. and M. P. Taylor (2001). Why is it so difficult to beat the random walk forecast of exchange rates?. Technical Report 88. European Central Bank. Frankfurt.

Kirkpatrick, S., C. D. Jr. Gelatt and M. P. Vecchi (1983). Optimization by Simulated Annealing. *Sience* **220**, 671–80.

Kirman, A. (1991). Epidemics of opinion and speculative bubbles in financial markets. In: *Money and Financial Markets* (M. Taylor, Eds.). Macmillan. pp. 354–368.

Kirman, A. (1993). Ants, rationality, and recruitment. *Quarterly Journal of Economics* **108**, 137–156.

Lagarias, J. C., J. A. Reeds, M. H. Wright and P. E. Wright (1999). Convergence behavior of the Nelder–Mead simplex algorithm in low dimensions. *SIAM J. Optimization* **9**, 112–147.

LeBaron, B. (2000). Agent–based computational finance: Suggested readings and early research. *Journal of Economic Dynamics & Control* **24**, 679–702.

LeBaron, B. (2002). Building financial markets with artifical agents: Desired goals, and present techniques. In: *Computational Markets* (G. Karakoulas, Eds.). MIT Press. p. to appear.

Nelder, J. A. and R. Mead (1965). A Simplex Method for Function Minimization. *Computer Journal* **7**, 308–313.

Pardalos, P. M. and M. Resende (2002). *Handbook of Applied Optimization*. Oxford University Press.

Pardalos, P. M., H. Tuy and E. Romeijn (2000). Recent developments and trends in global optimization. *Journal of Computational and Applied Mathematics* **124**(1–2), 209–228.

Spendley, W., G. R. Hext and F. R. Himsworth (1962). The Sequential Application ot Simplex Designs in Optimization and Evolutionary Operation. *Technometrics* **4**, 441–452.

Tesfatsion, L. (2001*a*). Introduction to the computational economics sepcial issue on agent–based computational economics. *Computational Economics* **18**, 1–8.

Tesfatsion, L. (2001*b*). Introduction to the sepcial issue on agent–based computational economics. *Journal of Economic Dynamics & Control* **25**(3–4), 281–293.

Winker, P. and M. Gilli (2001). Indirect estimation of the parameters of agent based models of financial markets. Technical Report 03/2001. School of Business Administration. International University in Germany, Bruchsal.

Winker, Peter (2001). *Optimization Heuristics in Econometrics*. Wiley. Chichester.

The following papers have been published since 2001:

01/2001	Tomás Bayón, Jens Gutsche, Hans H. Bauer	*Equity Marketing: Gegenstand, Konzepte, Methoden*
02/2001	Horst Entorf, Peter Winker	*The Economics of Crime: Investigating the Drugs-Crime Channel*
03/2001	Peter Winker, Manfred Gilli	*Indirect Estimation of the Parameters of Agent Based Models of Financial Markets*
04/2001	Florian v. Wangenheim, Tomás Bayón, Lars Weber	*Der Einfluß von persönlicher Kommunikation auf Kundenzufriedenheit, -bindung und Involvement: Design und Ergebnisse einer empirischen Studie im privaten Strommarkt**
05/2001	Peter Winker	*Applications of the Optimization Heuristic Threshold Accepting in Statistics*
06/2001	Florian v. Wangenheim, Tomás Bayón	*Satisfaction, Loyalty and Word-of-Mouth Within a Firm's Customer Base: Differences Between Stayers, Switchers and Referral Switchers*
07/2001	Martin Beck, Peter Winker	*Modelling Spillovers and Feedback of International Trade in a Disequilibrium Framework*
08/2001	Oliver Hülsewig, Peter Winker, Andreas Worms	*Bank Lending in the Transmission of Monetary Policy: A VECM Analysis for Germany*
09/2001	Florian v. Wangenheim, Tomás Bayón	*Direct and Moderated Effects of Customer Satisfaction on two dimensions of loyalty in a Business-to-Business Context*
10/2002	Tomás Bayón, Florian v. Wangenheim	*Valuation of Customers' Word-of-Mouth Referrals: Approach and First Result*
11/2002	Jan Doppegieter, Erwin Rode	*Capitalisation Rates and Property Yields: An Analysis of the South African Commerical Property Market*
12/2002	Manfred Gilli, Peter Winker	*A Global Optimization Heuristic for Estimating Agent Based Models*

All papers are available for download at http://www.i-u.de/departments/ba/research.htm

* Available in German only