Indirect Estimation of the Parameters of Agent based Models of Financial Markets*

Peter Winker† Manfred Gilli‡

February 11, 2002
Second version

Abstract

Agent based models take into account limited rational behaviour of individuals acting on financial markets. The simulation of the behaviour of such heterogenous agents and their interactions generate synthetic price time series. At least for specific parameter settings, these synthetic price time series exhibit marked similarities with actual financial market time series and share some of their statistical characteristics.

In this paper we introduce an indirect estimation approach of the parameters of agent based models. It is based on the comparison of statistical moments of simulated and actual time series. Using a refined global search heuristic, we obtain estimates of some parameters of a specific agent based model for the foreign exchange market. The paper presents details of this estimation approach and first results for the US–$/DM exchange rate.

Keywords: Agent Based Model, Validation, Indirect Estimation, Simulation, Foreign Exchange Market

JEL Codes: C15, G14

1 Introduction

Standard models of financial markets are based on fully rational behaviour of all agents and indicate that markets are efficient (Fama, 1970). Although these mod-

*Paper prepared for WEHIA 2002, (Workshop on Economics and Heterogenous Interacting Agents), Trieste, 30 May – 1 June 2002.

†Department of Economics and Quantitative Methods, International University in Germany, D–76648 Bruchsal, Germany. *Email:* <Peter.Winker@i-u.de>

‡Department of Econometrics, University of Geneva and FAME, CH–1211 Genève 4, Switzerland. *Email:* <Manfred.Gilli@metri.unige.ch>

els are useful for understanding financial markets, most of them lead to predictions in contrast to empirical findings like no trading volume, zero autocorrelation of returns, and price volatility equal or lower than the volatility of the "fundamental value" (Levy, Levy and Solomon, 2000, p. 142). Furthermore, typical statistical characteristics of financial market data like heavy tails and volatility clustering are difficult to explain using these models (Cochrane, 1999). Finally, the assumption of fully rational behaviour of homogenous agents is in stark contrast with empirical and experimental findings (De Bondt, 1998; Cheung and Wong, 2000; Arnswald, 2001).

Agent based models provide an alternative route for modelling financial markets. These models allow for heterogenous agents, and limited rational behaviour like momentum investment (Caginalp, Porter and Smith, 2000), herding or learning.[1] Furthermore, the interactions between agents are taken explicitly into account. The special issues of Computational Economics and the Journal of Economic Dynamics and Control in 2001 devoted to agent based models mirror the growing interest in this type of models, which is not limited to financial market modelling.[2]

Given that the aim of modelling consists in obtaining a simplified, but useful representation of real financial markets, it seems straightforward to compare the results of simulations of agent based models with empirical findings. For this purpose, it is analyzed whether typical features of financial market data like momentum investment (Lettau, 1997), chaotic behaviour (Brock and Hommes, 1997), or heavy tails of return distributions (Lux, 1998, 2000) also appear in the simulated data. In general, however, "validation ... remains a very weak area for the class of models described here." (LeBaron, 2000, p. 698). Neither are parameters of agent based models estimated or calibrated nor is the performance of the models evaluated against actual data. This lack is attributed to the typically large number of parameters describing agent based models, which might allow to fit any feature of real data, and to the complexity of model evaluation for a given set of parameters, which, in general, is not possible analytically.

The contribution of this paper is to indicate, how evaluation of agent based models can be strived for despite these hurdles. To this end, we undertake the estimation of some parameters of a simple agent based model using a simulated indirect estimation method. The estimates are obtained by minimizing a loss function which is based on comparing moments of the simulated data with those of empirical data for the foreign exchange market. Although our application is limited to a single model and a small number of specific features, it seems straightforward to extend the approach to a broader and more comprehensive set

[1]LeBaron (2000) provides an overview on some early approaches and different modelling tracks. See also Levy et al. (2000) for an introduction.

[2]See Tesfatsion (2001a, 2001b).

of data characteristics and, consequently, to a larger set of model parameters. However, as soon as the agent based model does not allow for an analytical solution, available computing resources and optimization techniques may impose a binding constraint on the number of parameters which can be taken into account.

In order to relax this constraint, Gilli and Winker (2002) introduce a global optimization heuristic for the indirect estimation problem. Our first results from applying this algorithm to the simple model of the foreign exchange market presented by Kirman (1991) indicate that – using estimated parameter values – it provides a good approximation of the real data in terms of the two moments underlying the estimation. However, due to the limited and ad hoc choice of moments of the data, which have been taken into account for estimation, the implicit test of the model against the data is not a very restrictive one yet. We provide some further information on the time series properties of the simulated data indicating the route for future research.[3]

The paper is organized as follows. In Section 2, the data are introduced and the specific moments used later for indirect estimation are presented and discussed. A short sketch of the simulation model is provided in Section 3. The method used for indirect estimation of some parameters of this model is described in some detail in Section 4. The results of a first implementation are summarized in Section 5, while Section 6 provides a conclusion and the outlook on further research.

2 Characteristics of financial market data

The statistical description of financial market data is a large and still growing field (Cochrane, 1999; Cont, 1999; Arifovic and Gençay, 2000). Consequently, a large number of specific features of these data has been described already. Given that typical agent based models include a large number of parameters, which have to be set, calibrated or estimated, it seems reasonable to use as much information from empirical data as possible in order to obtain reliable estimates and powerful tests of the models. Therefore, future extensions of the approach presented in this paper will have to take into account a much larger number of moments of financial market data than considered in our first application. We will discuss some possible extensions in Section 5.

Nevertheless, in order to introduce the method of indirect estimation to agent based models of financial market data, we start with a small, ad hoc selected number of moments in order to keep the computational load small and to gain a better understanding of the resulting objective function. Figure 1 provides the daily logarithmic returns of the DM/US–$ exchange rate (implicit EURO/US–$

[3] See also LeBaron (2000, p. 699) for a discussion of more restrictive tests.

rate from January 1999), which serves as empirical example.

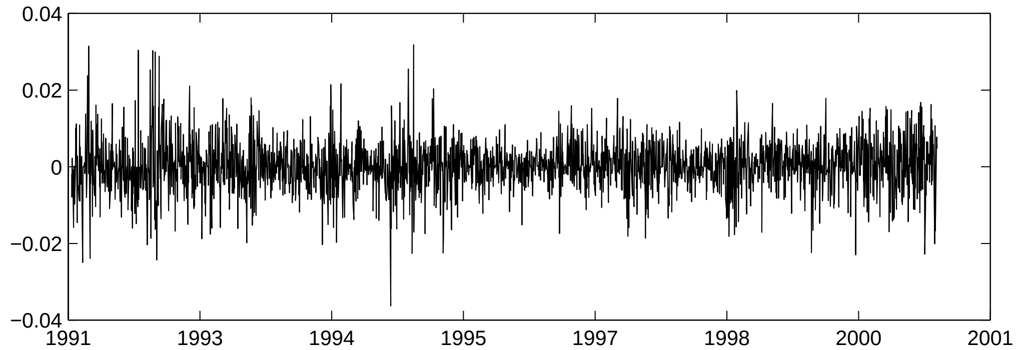

Figure 1: Daily Returns DM/US–$

The high number of large positive and negative returns hints at heavy tails of the return distribution. Consequently, the third moment (kurtosis) of the return distribution is expected to be significantly larger than for a normal distribution.[4] This typical feature of financial market data will serve as one of the moments to be matched in our simulated estimation approach. Besides the leptokurtosis of daily returns, the literature also describes a decreasing leptokurtosis under time aggregation (Lux, 1998). Therefore, Table 1 shows the empirical skewness and kurtosis of the returns for non overlapping time periods of differing length. The result support the view of a decreasing leptokurtosis under time aggregation for the returns of the DM/US–$ exchange rate. We do not take this effect into account for the estimation procedure, but will comment on the issue in our discussion of the results.

Table 1: Impact of time aggregation on the distribution of returns DM/US–$ (11.11.1991 – 8.11.2000)

	Skewness	Kurtosis
daily	0.027	4.983
weekly	0.031	4.106
monthly	0.264	3.248

[4]We are aware of the discussion about the existence of higher order moments (Lux, 2000). Therefore, future analysis will also make use of tail index measures and statistics based on the whole return distribution.

The second characteristic feature of the daily return series, which will be used for the indirect estimation approach, is the time varying volatility. The simplest model for capturing this feature is a standard ARCH–model.[5] Let r_t denote the daily logarithmic return, then an AR(1)–process for r_t with ARCH(1)–effect can be described by

$$r_t = \beta_0 + \beta_1 r_{t-1} + \varepsilon_t, \quad \text{where} \quad Var(\varepsilon_t) = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 \,. \tag{1}$$

In Table 2 some maximum–likelihood estimation results are presented for the coefficients of this model using different sample periods in order to assess the robustness of the estimates.

Table 2: Estimated AR– and ARCH–coefficients

	Sample					
	5.1.71 – 9.11.00	5.1.71 – 31.12.79	1.1.80 – 29.12.89	1.1.90 – 29.12.99	11.11.91 – 9.11.00	11.11.91 – 31.12.98
$\hat{\beta}_1$	0.006	**0.077**	0.024	**-0.050**	**-0.047**	**-0.070**
$\hat{\alpha}_1$	**0.213**	**0.534**	**0.113**	**0.130**	**0.124**	**0.172**

Note: Bold face number indicate significance at the 5%–level.

Although the estimated coefficients exhibit some variation depending on the sample chosen, the estimated ARCH(1)–effect α_1 is always significant, typically in a range between 0.1 and 0.25. Consequently, the time series generated by the agent based models should replicate this feature. For the further analysis, the ten year sample from 11.11.1991 to 9.11.2000 is fixed.

It should be kept in mind that neither the estimates of kurtosis nor the ARCH–model provide a satisfying description of the exchange rate time series or the return time series.[6] Nevertheless, these features appear to be significant and robust. Hence, they may contribute to the estimation of parameters of agent based models (Gouriéroux and Monfort, 1996). As long as our estimates are based on only these two moments of the data, they can hardly be highly discriminatory. Consequently, in order to discriminate between several agent based models using simulation based indirect estimation, further moments of the exchange rate time series have to be taken into account in future applications.

[5]An introduction to ARCH–models for financial market data is provided by Gouriéroux (1997). Again the simple ARCH–model can be replace by more complex GARCH–models or nonlinear models like MS–AR (Dewachter, 2001) or SETAR. Furthermore, statistics based on the conditional distribution of returns will be used in future implementations.

[6]In fact, the coefficient of determination for the ARCH–model is smaller than 0.002.

3 A simple agent based model

From the large number of agent based models proposed to mirror characteristic properties of financial markets, this paper deals only with a single model. This model introduced by Kirman (1991, 1993) stresses the importance of interaction between heterogenous, not fully rational individuals. Of course, future research will take into account further models of this type, but also models concentrating on individual learning, as one of the ultimate goal of our analysis is to discriminate between different modelling approaches based on observed data.

The model assumes that two groups of individuals act on the foreign exchange market. The members of the first group act on fundamentals, while the second group follows a simple chartist rule based on last period return. Besides this heterogeneity in agent behaviour, an even more important feature of the model is that the type of an individual can change over time. This may happen due to a random mutation with probability ε or as a result of direct interaction with a second individual. Therefore, it is assumed that if two agents meet randomly on the market place, the second one will convince the first one of his point of view with a given probability δ. In this case, the total share of fundamentalists and chartists in the population will change. As our results presented in Section 5 will show, these simple model of heterogenous agents with interaction is able to generate quite complex dynamics of the time series of daily logarithmic returns including ARCH–effects, excess kurtosis and decreasing leptokurtosis under time aggregation.

For the application on the foreign exchange market, it is assumed that fundamentalists expect the exchange rate S_t to return gradually to its fundamental value $\bar{S}$ which is assumed to be known with certainty. Consequently, the expected change of the exchange rate for fundamentalists is given by $E^f[\Delta S_{t+1}] = \nu(\bar{S} - S_t)$. The chartists are assumed to extrapolate last period returns, i.e. their expectations are given by $E^c[\Delta S_{t+1}] = S_t - S_{t-1}$.

In order to derive a time series for the exchange rate, some additional assumptions are required. In contrast to Frankel and Froot (1986), who first used a similar model of the foreign exchange market, and De Grauwe and Dewachter (1993) the weights of fundamentalists and chartists are not fixed. They evolve according to the shares of fundamentalists and chartists in the population represented by a set of n agents. However, instead of using these shares directly, Kirman (1991) propose to use the assumption that all agents have some knowledge about the share of fundamentalists. Consequently, their behaviour adjusts to the expected majority. As the signal on this majority is subject to some error, this assumptions leads to a higher probability of large shares of either fundamentalist views or chartist views dominating the market. As in Frankel and Froot

(1986) the market price is calculated as a weighted mean of market expectation

$$E^m[\Delta S_{t+1}] = w_t E^f[S_{t+1}] + (1 - w_t) E^c[\Delta S_{t+1}] \,,$$

where w_t is the share of agents expecting fundamentalists to be dominating, and fundamental value. This market outcome can be motivated based on the assumption of risk averse agents.

Of course, this model can be extended and modified in several directions.[7] For example, the probability δ of convincing another trader could be made dependent on past success. Furthermore, the market price could be obtained from explicitly modelled demand and supply schedules of individual agents. The number of different agent types can be increased (Lux, 1998) etc. However, since our goal in this paper is to present simulation based indirect estimates of model parameters, we stick to the basic model and leave a comparison with more sophisticated variants for future research.

4 Indirect estimation

Besides the two parameters ε and δ already introduced, a view further parameters have to be set in order to describe the agent based model completely. Table 3 provides an overview on these parameters.

Table 3: Parameters of the agent based model

Label	Interpretation	Value
n	number of agents	100
$iter$	number of interactions	50 000
$inter$	number of interactions per trading day	50
ν	adjustment speed in fudamentalists' expectations	0.045
σ_s^2	variance of price shocks	0.25
c	weighting factor in price function	0.5

Of course, all of them could be incorporated in the estimation approach described in this section. However, as the resulting optimization problem is non standard and requires the use of a new global optimization routine, we restricted ourselves to the addition of a third parameter, namely the variance of errors in the majority assessment σ_q^2. For the resulting three dimensional parameter space, it is still possible to obtain a good approximation of the objective function on a

[7] E.g. Chiarella, Dieci and Gardini (2001).

grid of the parameters values using a reasonable amount of computer resources.[8] This grid allows to assess the results generated from the optimization algorithm, in particular, whether the algorithm converges to a solution in the vicinity of the global optimum or rather to some local minimum as it might result from Monte Carlo variance. Given the good performance of the algorithm on this problem, we plan to increase the dimensionality of the parameter space after some further tuning of the algorithm and the selection of an objective function taking into account more moments of the data.

As pointed out in Section 2, this first implementation points at matching the observed ARCH(1)-effect $\alpha_{1_{\text{emp}}}$ and the observed kurtosis $k_{d_{\text{emp}}}$, measured at a daily time scale, of the daily logarithmic returns of the DM/US-$ exchange rate. The objective function measures the discrepancy between these empirical moments and the mean of the estimated moments $\overline{k}_d$ and $\overline{\hat{\alpha}}_1$ for the data simulated for a specific parameters triple $(\varepsilon, \delta, \sigma_q^2)$. Consequently, the parameter estimates $\hat{\varepsilon}$, $\hat{\delta}$ and $\hat{\sigma}_q^2$ are defined as the solution to the following minimization problem:

$$\min_{\varepsilon,\delta,\sigma_q^2} f = \mid \overline{k}_d - k_{d_{\text{emp}}} \mid + \lambda \mid \overline{\hat{\alpha}}_1 - \alpha_{1_{\text{emp}}} \mid ,$$

where $\overline{k}_d$ and $\overline{\hat{\alpha}}_1$ are the mean values obtained from *nrep* simulations of the model. In order to make these statistics more robust we delete the first and the last 10% of results in the tails. The scaling parameter λ was chosen equal to 15 to correct for the different scalings of $k_{d_{\text{emp}}}$ and $\alpha_{1_{\text{emp}}}$. Algorithm 1 describes the estimation procedure in more detail.

The criterion for selecting the successive vectors in Step 3 of Algorithm 1 is its crucial ingredient besides an efficient simulation setup. It has to taken into account two problems exhibited in Figure 2. The figure shows the stochastic approximation of the objective function f against the parameters ε and δ holding all the other parameters of the model fixed (for $\sigma_q^2 = 0.04$).

The left plot provides the results for $nrep = 200$ Monte Carlo replications. The considerable Monte Carlo variance of the estimates of $f(\varepsilon, \delta, \sigma_q^2)$ is evident and causes any standard optimization procedure to stop in a local minimum close to the starting vector. The same holds true for the rather sophisticated Nelder–Mead simplex direct search method (Lagarias, Reeds, Wright and E.Wright, 1999). Although the approximations become smoother as the number of replications is increased to $nrep = 10\,000$ as indicated in the central panel,[9] this is not sufficient to make ordinary search techniques work well, as there is still some Monte Carlo variance left as the plot in the right panel of Figure 2 shows, where we use a ten

[8]The approximations of the objective function have been computed with $nrep = 500$ 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.

[9]In fact, the Monte Carlo variance shrinks at the usual rate of $1/\sqrt{nrep}$.

Algorithm 1 Indirect estimation procedure

Give $x^{(0)} \in \mathbf{R}^n$ vector of starting values of parameters to be estimated
while not converged **do**
  Determine successive vectors x (defined by optimization algorithm)
  for each x **do**
    Initialize random variable generators
    for $i = 1 : nrep$ **do**
      Generate random sequences for price simulation
      Simulate price path $p^{(i)}$ and returns $r^{(i)}$
      Compute $\hat{\alpha}_{1,i}$ and $k_{d,i}$
    end for
    Truncate tails (10%) of the distribution of $\hat{\alpha}_{1,i}$ and $k_{d,i}$, $i = 1, \ldots, nrep$
    Compute means $\overline{\hat{\alpha}}_1$ and $\overline{k}_d$ of truncated distributions
    Evaluate objective function $f = \mid \overline{k}_d - k_{d_{\text{emp}}} \mid + \lambda \mid \overline{\hat{\alpha}}_1 - \alpha_{1_{\text{emp}}} \mid$
  end for
end while

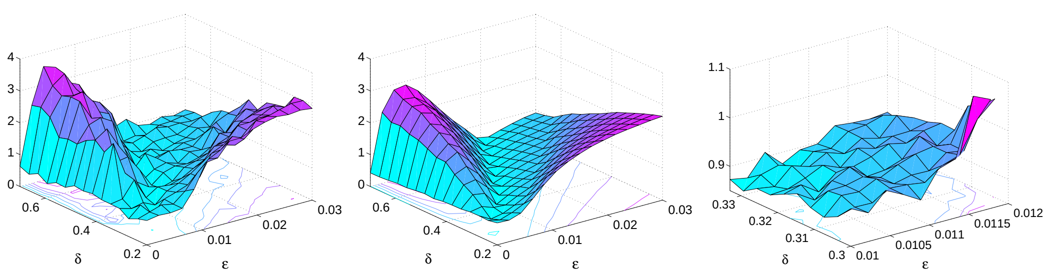

Figure 2: Simulated values of f against ε–δ–grid. (Left panel $R = 200$, central panel $R = 10\,000$), right panel $R = 10\,000$ with finer grid.

times finer grid to represent the objective function computed with $nrep = 10\,000$ replications. On the other hand, computational requirements are already quite high when using $nrep = 10\,000$ replications.

Besides this problem caused by the fact that only a stochastic approximation to the objective function is available, the objective function itself does not appear to be globally convex in the parameter space. Since estimation in our implementation includes three parameters, but might include even more in the future, an adequate optimization routine should also be able to tackle this situation.

For these two reasons, Gilli and Winker (2002) introduce a global optimization heuristic for the indirect estimation problem based on 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 and because it is not globally convex. We will not present details of the optimization algorithm in this paper,[10] but present first results in the next section.

5 Results and further research

Figure 3 shows the grid plot of the objective function in the ε–δ and ε–σ_q^2 subspace, while Figure 4 shows the δ–σ_q^2 subspace. For all three plots, the value of the third parameter is fixed at the grid point closest to the optimum value obtained through the application of the global optimization heuristic. This optimal parameter values [11] are $\hat{\varepsilon} = 0.0001$, $\hat{\delta} = 0.3037$ and $\hat{\sigma_q^2} = 0.05$ with $f = 0.229$.

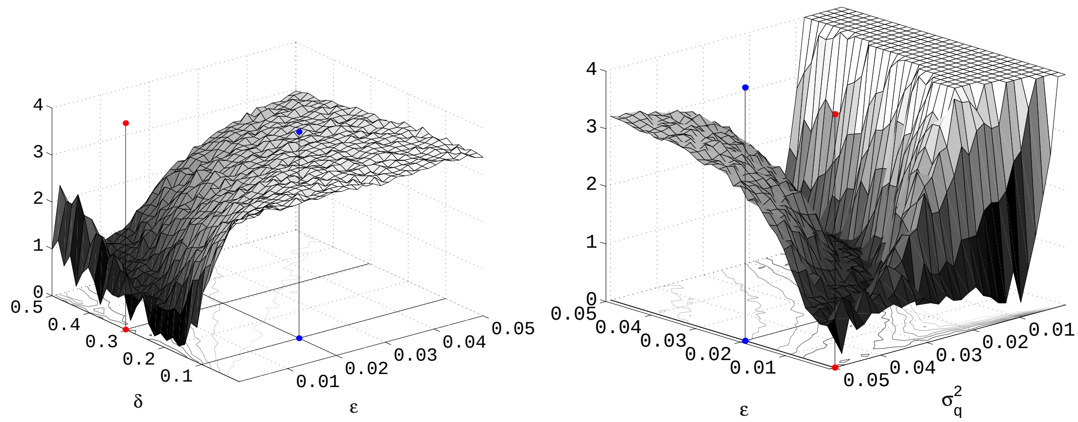

Figure 3: Approximations of the objective function in the ε–δ, ($\sigma_q^2 = 0.05$) and ϵ–σ_q^2 ($\delta = 0.3065$) grid.

Let us now turn to some analysis of this result. First, Kirman (1993) uses Markov chain theory to demonstrate that the process will exhibit large shares of fundamentalists and chartists, respectively, with high probability, if $\varepsilon < (1-\delta)/(n-1)$. For our estimates we find $\varepsilon = 0.0001 < 0.007 = (1-0.3037)/(100-1) = (1-\delta)/(n-1)$. 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.

[10] The paper Gilli and Winker (2002) can be obtained in electronic format from the authors on request.

[11] Computing time is less than 1 hour on a Pentium III 600 MHz PC.

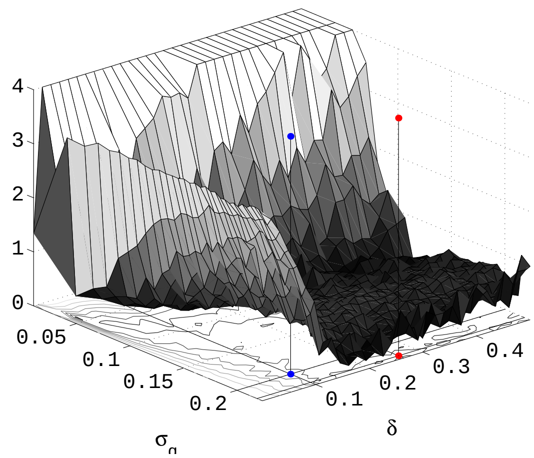

Figure 4: Approximations of the objective function in the δ–σ_q^2 ($\varepsilon = 0.0001$) grid.

Second, we are also interested to what extent the price and return series generated from the agent based model using these optimal parameter values resemble actual market data. For this purpose, we rerun the simulation using a series derived from the theory of purchasing power parity[12] instead of a constant fundamental value, and $\sigma_s^2 = 0.003$ in order to obtain similar orders of magnitude. Figure 5 shows the daily logarithmic returns of the US–\$/DM exchange rate in the upper panel and a typical simulated returns series in the lower panel.

A first look indicates that both series exhibit volatility clustering to a relevant extent. For the real data, the estimated ARCH(1)–effect amounts to 0.130, while it is 0.111 for the simulated series. The frequency of large changes is also high for both series leading to a kurtosis of 4.98 for the actual data and 4.39 for the simulated series, i.e. both return series have heavy tails. While a close resemblance of ARCH(1)–effect and kurtosis should have been expected given that these moments have been used for estimating some crucial parameters of the agent based models, it is interesting to see that the heavy tails tend to disappear under time aggregation not only for the actual data, but for the simulated data where kurtosis for weekly returns is 3.86 (4.11 for the actual data), and for monthly returns it amounts to 3.19 (3.25).

While this property of the simulated returns looks promising, a closer look at the whole distribution of returns draws a slightly different picture. Figure 6 provides a QQ–plot of the actual returns against the simulated returns. If the distribution of the simulated returns mimics the distribution of actual returns one

[12] This series was generated from the terms of trade series published by the Deutsche Bundesbank.

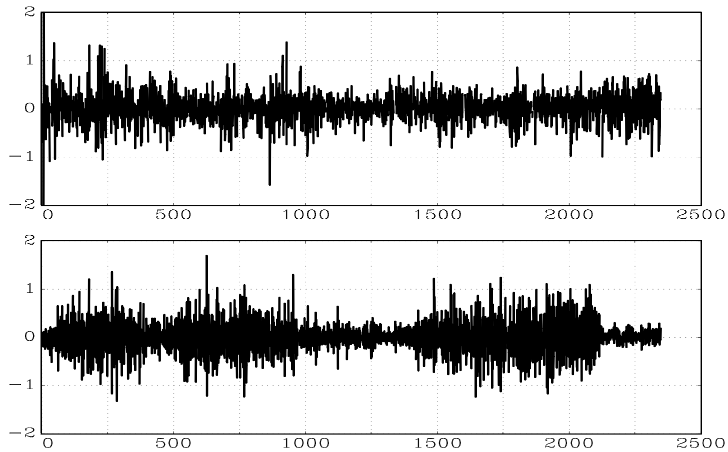

Figure 5: Daily Returns DM/US–$ (upper panel) and Simulated Returns (lower panel)

to one, a straight line should result. From the plot it becomes obvious that this requirement is met basically for “normal” returns in the range of -1% to 1%, but that the tail behaviour differs. However, it should be noted that the frequency of returns exceeding 1% in absolute terms is around 0.6% for the actual data sample. Nevertheless, some improvement in matching the unconditional distribution of returns might result from replacing the kurtosis in the objective function by a measure based on the QQ–plot, i.e. on the whole distribution function.

A further departure of the simulated time series from actual data can be detected by looking at the price levels or, alternatively, at the conditional distribution of returns given past returns. Here, the estimated ARCH–model provides a first hint with an estimated AR–term of -0.43 for the simulated data as compared to the -0.05 for actual data. This strong negative autocorrelation of the simulated returns becomes apparent in the upper right panel of Figure 7, which shows a scatter plot of the simulated daily returns against the simulated daily returns lagged one day. The negative autocorrelation is clearly visible, while a similar plot for the actual return data in the upper left panel does not exhibit such a relationship. The lower panel provides information on fifth order autocorrelation, which does not appear to be substantial both for the actual and the simulated data. Finally, these plots also allow to detect the tendency for volatility clustering.

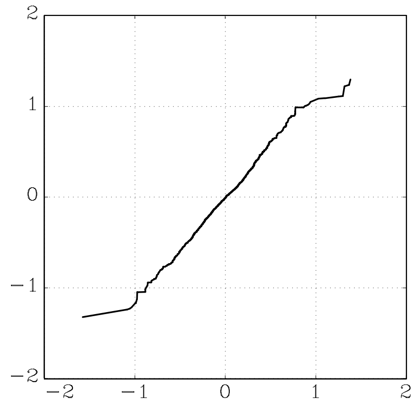

Figure 6: QQ–Plot of Daily Returns DM/US–$ and Simulated Returns

Summarizing the findings for a single realization of the price series resulting from the agent based models using estimated values for three parameters, it can be stated that the moments, which have been integrated in the objective function, are matched quite well as expected. Some further characteristics of the actual data are also reflected in the simulated ones, e.g. the disappearance of excess kurtosis under time aggregation. However, the simulated data exhibit some features in marked contrast to the actual data, e.g. the strong negative autocorrelation.

These findings also define the next step of our analysis. First, we will replace the ad hoc chosen objective function by a new one taking into account the whole (conditional) distribution(s) of returns. Second, we will increase the efficiency of the global optimization heuristic used for the indirect estimation approach to allow for the estimation of all central parameters of the agent based model. Third, based on the new results obtained we can suggest changes of the model and compare its performance with other types of agent based models of the foreign exchange market.

6 Conclusion

Previous research using different kinds of agent based models indicated that empirical features of financial market data, which appear difficult to motivate using

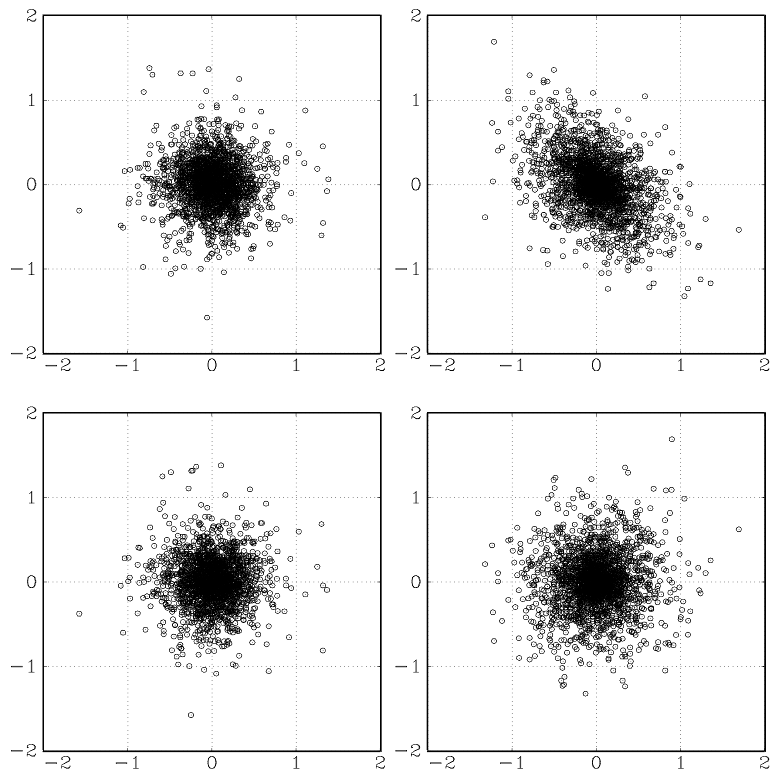

Figure 7: Conditional Distributions of Daily Returns DM/US–$ and Simulated Returns

standard efficient market models, can be replicated within this setting. However, to our knowledge, a thorough validation of these models has not yet been undertaken (LeBaron, 2002).

Using a classic agent based model of the foreign exchange market, this paper demonstrates how the estimation of parameters, in principle, can be performed using indirect simulation methods. However, the non convexity of the objective function used for this purpose precludes the efficient use of standard optimization tools. The global optimization heuristic introduced in our companion paper Gilli and Winker (2002) allows to tackle this problem. We find that the model with some estimated parameters is able to replicate the moments of the data which have been integrated in the objective function, but behaves quite differently, in particular, with regard to a strong negative autocorrelation of returns.

Therefore, our next steps will consist in choosing a more comprehensive objective, and applying the technique presented in this paper to a larger set of model parameters. Furthermore, we will extend our analysis to different kinds of agent based models. We expect, that using a large enough number of empirical moments, a discrimination between agent based models should become possible. However, citing again LeBaron (2002): "this field is only in its infancy, and much remains to be done."

References

Arifovic, J. and R. Gençay, "Statistical properties of genetic learning in a model of exchange rate," *Journal of Economic Dynamics and Control*, 2000, *24* (5–7), 981–1006.

Arnswald, T., "Investment Behaviour of German Equity Fund Managers," Technical Report 08/01, Deutsche Bundesbank, Frankfurt 2001.

Brock, W. A. and C. H. Hommes, "A Rational Route to Randomness," *Econometrica*, 1997, *65* (5), 1059–1095.

Caginalp, G., D. Porter, and V. Smith, "Momentum and Overreaction in Experimental Asset Markets," *International Journal of Industrial Organization*, 2000, *18*, 187–204.

Cheung, Y.-W. and C. Y.-P. Wong, "A Survey of Market Practioners' Views on Exchange Rate Dynamics," *Journal of International Economics*, 2000, *51*, 401–419.

Chiarella, C., R. Dieci, and L. Gardini, "Asset Price Dynamics in a Financial Market with Fundamentalists and Chartists," *Discrete Dynamics in Nature and Society*, 2001, p. forthcoming.

Cochrane, J. H., "New Facts in Finance," Technical Report 7169, NBER, Cambridge, MA 1999.

Cont, R., "Statistical properties of financial time series," Technical Report urlwww.cmap.polytechnique.fr/ rama/papers/, Centre de Mathématiques Appliquées, Ecole Polytechnique, Palaiseau, France 1999.

De Bondt, W. F. M., "A portrait of the Individual Investor," *European Economic Review*, 1998, *42*, 831–844.

De Grauwe, P. and H. Dewachter, "A Chaotic Model of the Exchange Rate: The Role of Fundamentalists and Chartists," *Open Economies Review*, 1993, *4*, 351–379.

Dewachter, H., "Can Markov Switching Models Replicate Chartist Profits in the Foreign Exchange Market?," *Journal of International Money and Finance*, 2001, *20*, 25–41.

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

Fama, E., "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance*, 1970, *25*, 383–417.

Frankel, J. A. and K. Froot, "The Dollar as a Speculative Bubble: A Tale of Fundamentalists and Chartists," Technical Report 1845, NBER, Cambridge, MA 1986.

Gilli, M. and P. Winker, "A Global Optimization Heuristic for Estimating Agent Based Models," Technical Report, Department d'Économetrie, Université de Geneève 2002.

Gouriéroux, C., *ARCH Models and Financial Applications*, New York: Springer, 1997.

____ **and A. Monfort**, *Simulation–Based Econometric Methods*, Oxford: Oxford Unviersity Press, 1996.

Kirman, A., "Epidemics of Opinion and Speculative Bubbles in Financial Markets," in M. Taylor, ed., *Money and Financial Markets*, Macmillan 1991, pp. 354–368.

____ , "Ants, Rationality, and Recruitment," *Quarterly Journal of Economics*, 1993, *108*, 137–156.

Lagarias, J. C., J. A. Reeds, M. H. Wright, and P. E.Wright, "Convergence Behavior of the Nelder-Mead Simplex Algorithm in Low Dimensions," *SIAM J. Optimization*, 1999, *9*, 112–147.

LeBaron, B., "Agent–based Computational Finance: Suggested Readings and Early Research," *Journal of Economic Dynamics & Control*, 2000, *24*, 679–702.

——, "Building Financial Markets With Artifical Agents: Desired Goals, and Present Techniques," in G. Karakoulas, ed., *Computational Markets*, MIT Press 2002, p. to appear.

Lettau, M., "Explaining the Facts with Adaptive Agents: The Case of Mutual Fund Flows," *Journal of Economic Dynamics and Control*, 1997, *21*, 1117–1147.

Levy, M., H. Levy, and S. Solomon, *Microscopic Simulation of Financial Markets*, San Diego: Academic Press, 2000.

Lux, T., "The Socio–Economic Dynamics of Speculative Markets: Interacting Agents, Chaos, and the Fat Tails of Return Distributions," *Journal of Economic Behavior & Organization*, 1998, *33*, 143–165.

——, "On Moment Condition Failure in German Stock Returns: An Application of Recent Advances in Extreme Value Statistics," *Empirical Economics*, 2000, *25*, 641–652.

Tesfatsion, L., "Introduction to the Sepcial Issue on Agent–Based Computational Economics," *Journal of Economic Dynamics and Control*, 2001a, *25* (3–4), 281–293.

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