J Econ Interac Coord (2007) 2:125–145
DOI 10.1007/s11403-007-0020-4

REGULAR ARTICLE

An objective function for simulation based inference on exchange rate data

Peter Winker · Manfred Gilli · Vahidin Jeleskovic

Published online: 15 June 2007

Abstract The assessment of models of financial market behaviour requires evaluation tools. When complexity hinders a direct estimation approach, e.g., for agent based microsimulation models, simulation based estimators might provide an alternative. In order to apply such techniques, an objective function is required, which should be based on robust statistics of the time series under consideration. Based on the identification of robust statistics of foreign exchange rate time series in previous research, an objective function is derived. This function takes into account stylized facts about the unconditional distribution of exchange rate returns and properties of the conditional distribution, in particular, autoregressive conditional heteroscedasticity and long memory. A bootstrap procedure is used to obtain an estimate of the variance-covariance matrix of the different moments included in the objective function, which is used as a base for the weighting matrix. Finally, the properties of the objective function are analyzed for two different agent based models of the foreign exchange market, a simple GARCH-model and a stochastic volatility model using the DM/US-$ exchange rate as a benchmark. It is also discussed how the results might be used for inference purposes.

Keywords Indirect estimation · Simulation based estimation · Exchange rate returns

Research has been supported by the DFG grant WI 20024/2-1/2. We are indebted to two anonymous referees of this journal, Leigh Tesfatsion, Patrick Burns and other participants of the CEF'06 conference in Limassol for helpful comments on preliminary versions of this paper.

P. Winker (✉) · V. Jeleskovic
Department of Economics, University of Giessen, Licher Strasse 64, 35394 Giessen, Germany
e-mail: Peter.Winker@wirtschaft.uni-giessen.de

M. Gilli
Department of Econometrics, University of Geneva, and Swiss Finance Institute,
40 Boulevard du Pont-d'Arve, Geneva 4 1211, Switzerland

JEL Classification C14 · C15 · F31

1 Introduction

Exchange rate time series are found to exhibit complex structures which cannot be easily explained by standard economic models or modelled by simple time series models. As an alternative modelling approach, agent based models have been developed starting with the model proposed by Frankel and Froot (1986) with fixed dealer characteristics and Kirman (1991) allowing for endogenous adjustment of behaviour. Since then, a growing number of models has been presented and analyzed with increasing complexity of agents' behaviour and market structures. For recent contributions see the overviews in LeBaron (2006) and Tesfatsion (2006), the references given in Hommes et al. (2005), and Leigh Tesfatsion's website on agent based economics: http://www.econ.iastate.edu/tesfatsi/ace.htm.

The contribution of this paper will not consist in modifying an existing or proposing a new agent based model of foreign exchange rates, but in providing a tool for the evaluation of such models based on simulated indirect inference along the lines proposed by Gilli and Winker (2003).[1] The application of this approach is demonstrated using the models introduced by Kirman (1991, 1993), Lux (1998), and Lux and Marchesi (2000).

Given the variety of agent based models proposed for modelling foreign exchange rates, evaluation becomes increasingly important for assessing their explanatory power, discrimination between models and selection of optimal parameter values (estimation). Evaluation of agent based models can start from different perspectives, in particular from an input or output point of view.[2] This paper concentrates on model output, i.e. the simulated time series. Figure 1 provides an overview on the steps involved in the analysis. The upper part of the figure shows the steps of a typical statistical analysis applied to the foreign exchange market, while the lower part exhibits the proceeding with agent based models. The link from the box for foreign exchange market to the one for the agent based simulation model represents the microfoundation of agents' based simulation models, i.e. the attempt to use reasonable assumptions about agent behaviour and interaction. On the right hand side of the figure, the box marked "comparison" represents the core component of the approach of indirect inference, i.e. the comparison of real time series with those generated from the agent based model for given parameter values. Eventually, this comparison leads to a feedback on the agent based model, e.g. on its parameters. This optimization step will be subject to future research.

It appears to be a standard approach in the literature on agent based models of foreign exchange rates to concentrate the analysis on some statistical properties of simulated time series, e.g. fat tails or conditional heteroscedasticity. However, often

[1] For alternative approaches for models allowing for a closed form solution see Alfarano et al. (2005) and Ahrens and Reitz (2005).

[2] An interesting approach for combining both approaches is presented by Izumi and Ueda (2001). For a more detailed discussion, see again Leigh Tesfatsion's webpage: http://www.econ.iastate.edu/tesfatsi/empvalid.htm.

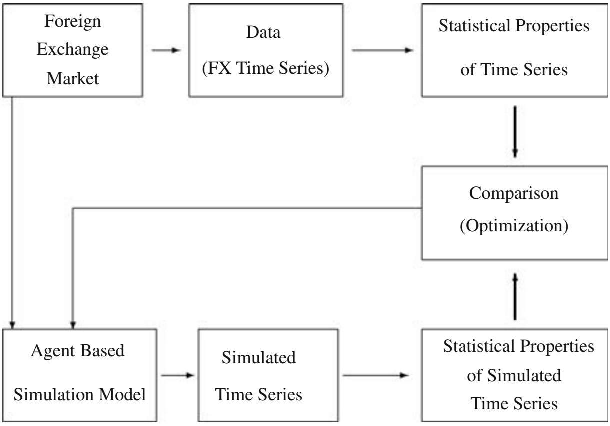

Fig. 1 Indirect inference for agent based models

this analysis is based on a single or a small number of simulated time series and for a given set of model parameters. Furthermore, the statistical properties are analyzed separately. In the following, a systematic approach is proposed taking into account a larger set of time series characteristics and the joint distribution of outcomes from repeated runs of the simulation models.

The selection of time series characteristics used for the evaluation approach is based on previous research on stylized facts of foreign exchange time series (Winker and Jeleskovic 2007, 2006). A single objective function is constructed based on the joint distribution of characteristics obtained from a bootstrap analysis.[3] From repeated simulations of the agent based model, the distribution of this objective function is obtained for different parameter settings. A comparison of this distribution with the distribution obtained from the empirical data using resampling techniques, simulation based inference on the agent based model for given parameter sets becomes feasible.

Finally, the objective function might be used in an optimization approach for indirect estimation of the parameters as proposed in Gilli and Winker (2003). However, this last step of the analysis is not subject of the present contribution. Nevertheless, the descriptive evidence gained on the features of the objective function will be helpful for developing adequate heuristic optimization methods for this purpose.

The paper is organized as follows. Section 2 describes the validation approach and the method of simulated moments in more detail. The construction of the objective function based on a selection of empirical moments and a weighting function resulting from the bootstrap procedure is explained in Sect. 3. Section 4 reports the results of applications of the method to two agent based models of the foreign exchange market, a simple GARCH-time series model, and a stochastic volatility model as a benchmark. A summary of the findings, conclusions and an outlook to future research is provided in Sect. 5.

[3] Alternatively, a multicriteria objective function could be considered (Kalaba and Tesfatsion 1996). Given the resulting increase in computational cost, it is left for future research to follow this approach.

2 Validation and estimation

One potential goal of the use of agent based models of foreign exchange markets is to "explain" typical characteristics of exchange rate movements. Thus, the following components might be assumed to be given:

1. A realisation of the exchange rate process $\mathbf{X} = \{x_t\}$ for $t = 1, \ldots, T$.
2. An agent based model A, which can be interpreted as a stochastic mapping from the parameter space Θ to the space Υ of exchange rate series of length T: For a given parameter vector $\boldsymbol{\theta} \in \Theta$, the model A generates stochastic realizations $A(\boldsymbol{\theta}) \in \Upsilon$.

The mapping A is stochastic whenever the agent based model involves some stochastic component such as random mutation of agent type or exogenous shocks to the price series. Such components are present in most if not all agent based models of foreign exchange rates including the examples considered in this contribution.

For the foreign exchange market, it is not claimed that agent based models might "explain" the actual level of the exchange rate. They are rather used as a model for describing typical statistical features of foreign exchange rate movements, which are difficult to explain by more traditional time series approaches, e.g. fat tails and long memory of exchange rate returns. Thus, when considering the validity of an agent based model or when searching for good parameters $\theta \in \Theta$, the benchmark is not given by the exchange rate process $\mathbf{X}$, but rather by some specific statistics/moments. The choice of moments for the present application will be discussed in the following Sect. 3. For the moment being, it is assumed that $\mathbf{m} = (m_1, \ldots, m_k)$ are k moments of the exchange rate process to be matched by the agent based model.

Unfortunately, the true vector of moments $\mathbf{m}$ is not known, but has to be estimated based on the available observations $\mathbf{X}$. Let $\mathbf{m}^e = (m_1^e, \ldots, m_k^e)$ denote consistent estimates of $\mathbf{m}$. The variance-covariance matrix of these estimates, $\mathrm{Var}(\mathbf{m}^e)$, might be estimated from $\mathbf{X}$, e.g. by means of bootstrap techniques. Given that some of the moments in the application are conditional moments of exchange rate returns, a block bootstrap method appears adequate. The implementation details are described in Subsect. 3.2. As an example, Fig. 2 shows the distribution of the sum of the ARCH- (α) and GARCH-parameters (β) in a GARCH(1,1)-model fitted to the US-\$/DM-exchange rate for the period 11.11.1991 to 9.11.2000.[4] The plot shows a kernel estimate of the density function based on 10,000 block bootstrap samples.[5]

Given the stochastic nature of A, the variance of some moment m_i might also be used as a moment. Then, a double-bootstrap approach has to be used to obtain the variance estimates for this moment. For the present application, this feature is not considered.

[4] For a low fraction of bootstrap samples, the standard algorithms for the ML-estimation of the GARCH-parameters exhibit numerical problems. For these instances, a local search heuristic has been used to obtain reliable estimates (Winker and Maringer 2007).

[5] All kernel estimates used for graphical presentations in this paper employ the normal kernel with a bandwidth selected to obtain smooth densities.

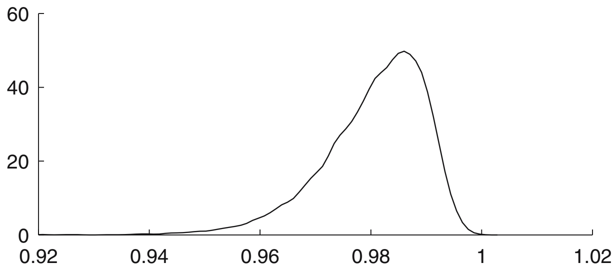

Fig. 2 Block bootstrap distribution of $\alpha + \beta$ in GARCH(1,1)-model

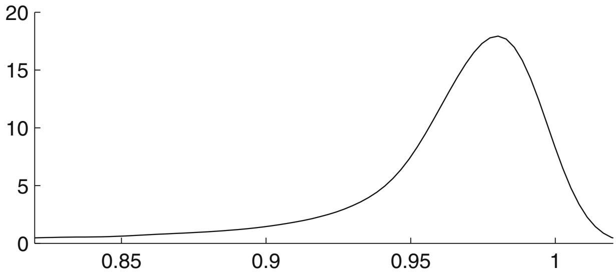

Fig. 3 Distribution of $\alpha + \beta$ in GARCH(1,1)-model for simulated time series from Kirman's model

2.1 Validation approach

From the model point of view, for a given parameter vector $\boldsymbol{\theta}$, each simulation run of A results in a single simulated time series x_t^s, $t = 1, \ldots, T$. For this simulated time series, it is also possible to estimate the vector of moments $\boldsymbol{m}$. Let $\boldsymbol{m}_i^s(\boldsymbol{\theta})$ denote this estimate for simulation i, $i = 1, \ldots, I$. Figure 3 shows the distribution of the sum of the ARCH- (α) and GARCH-parameters (β) in a GARCH(1,1)-model fitted to the simulated time series of an implementation of Kirman's (1991) model for a fixed parameter setting $\boldsymbol{\theta}$.[6] The plot shows a kernel estimate of the density function based on 10,000 simulation runs.

The mean of $\alpha + \beta$ over the bootstrap replications amounts to 0.9804, while the mean for the simulated data in Fig. 3 is slightly lower with 0.9622. However, both distributions exhibit large variance. Thus, the difference in the mean estimators might not be considered as statistically significant at any conventional level. A different result is obtained when considering time series generated from a GARCH(1,1)-model with parameter values set to $\alpha = 0.2$ and $\beta = 0.5$. Figure 4 exhibits the kernel estimate of the density function based on 10,000 simulation runs. In this case, the estimated mean of 0.6923 is obviously significantly smaller than both for the real data and the data obtained from a simulation of Kirman's model.

[6] For this simulation, we used the parameters provided in Table 2 in Subsect. 4.1 below with $\varepsilon = 0.01$ and $\delta = 0.6$.

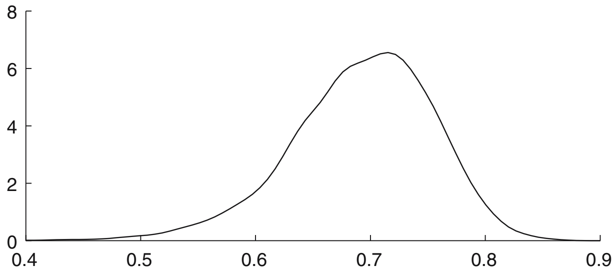

Fig. 4 Distribution of $\alpha + \beta$ in GARCH(1,1)-model for simulated time series for a GARCH(1,1)-process

Generally, for a statistical comparison of the results for the real data based on their bootstrap distribution and the simulated data, two approaches might be taken into account:

1. If the overlap of both density functions is smaller than a marginal level of 1, 5 or 10%, respectively, one might conclude that both distributions result from different data generating processes. This test can be easily evaluated based on the empirical distribution functions for a large enough number of replications/simulations.
2. One might ignore the variance in the simulation model and be interested solely in the means (medians) of simulated moments. Then, it should be checked whether the mean falls in a 90, 95 or 99% confidence interval based on the bootstrap estimates.

Of course, ignoring the (often huge) variance from the simulation model and concentrating on the mean (median) over a large number of replications might not be considered as a preferable approach. However, two arguments can be mentioned in favor of this idea. First, given that the simulated models generate quite different time series for different random number seeds, it might be of only limited interest to learn that almost every outcome might turn up some times, but rather one might be interested in the "typical" behavior as measured by means or medians. Second, when switching from the consideration of a single moment as in the example to $k > 1$ moments, the first approach leads to comparing the overlap of high dimensional density functions. Therefore, for this paper, we follow the second approach, but keep a more detailed analysis of the first method for future research.

If mean or median of simulated moments exist theoretically, increasing the number of simulation runs allows to control for the approximation error of using the mean of some simulation runs instead of their true values. A quantitative estimate of this dependence will provide useful for indirect estimation approaches as it allows to control the required number of model simulations at each evaluation step.

The validation approach has to be refined when more than a single moment is considered. Although some parameters of the agent based models might be set based on a priori reasoning or empirical evidence (Izumi and Ueda 2001), typically a considerable set of parameters will be subject to calibration/estimation. Consequently, validation and simulation based estimation has also to take into account a large enough number k of moments $\mathbf{m} = (m_1, \ldots, m_k)$. Typically, it cannot be assumed that these moments are distributed independently from each other neither for the empirical nor for the simulated data. Thus, performing the validation analysis moment by moment along the lines described above, is not feasible. Instead, validation has to be based either on the joint distribution of all moments with regard to some multicriteria approach (Kalaba and Tesfatsion 1996) or on a scalar valued objective function depending on all moments. The latter will be the approach discussed in the next subsection in the context of estimation by the method of simulated moments.

2.2 Method of simulated moments

The method of simulated moments puts a slightly different focus on the problem. Now, the parameters of the simulation model $\boldsymbol{\theta}$ are not considered as given, but are subject to estimation. A standard method of moment condition requires that $\mathrm{E}(\boldsymbol{m}^s \mid \boldsymbol{\theta}) = \mathbf{m}$.[7] Consequently, if $\mathbf{m}$ was given, the parameter vector $\boldsymbol{\theta}$ could be estimated replacing the expected value by the mean over the simulated moments. In this case, the simulated moments estimator $\hat{\theta}$ is the solution to

$$\frac{1}{I}\sum_{i=1}^{I}\left[(\boldsymbol{m}_i^s \mid \boldsymbol{\theta}) - \mathbf{m}\right] = \mathbf{0}\,. \tag{1}$$

If the number of moment conditions k exceeds the number of parameters in θ, a generalized method of moments is indicated. Let the $k \times 1$ vector of deviations from the above moment condition be defined as

$$\boldsymbol{G}_I(\boldsymbol{\theta}) = \sum_{i=1}^{I}\left[(\boldsymbol{m}_i^s \mid \boldsymbol{\theta}) - \mathbf{m}\right]\,. \tag{2}$$

Then, minimization of the weighted sum of squares

$$\frac{1}{I}\boldsymbol{G}_I'\mathbf{W}\boldsymbol{G}_I\,, \tag{3}$$

where W is a $k \times k$ positive definite matrix, results in the generalized method of simulated moments estimator $\hat{\boldsymbol{\theta}}$. The intuitive idea for setting the values of the weighting matrix W is to allow larger errors for parameters estimated with a higher degree of uncertainty and vice versa. Choosing W as the inverse of the variance-covariance matrix of $\boldsymbol{G}_I(\boldsymbol{\theta})$ would be an optimal choice (Heij et al. 2004, p. 257). Obviously, $\boldsymbol{G}_I(\boldsymbol{\theta})$ is not available as $\boldsymbol{\theta}$ is not known. Furthermore, estimation of an optimal weighting matrix W is complicated by the fact that $\boldsymbol{G}_I(\boldsymbol{\theta})$ is subject to two types of uncertainty resulting from estimation of empirical moments to be matched $\mathbf{m}$ and simulated moments for a given parameter vector $\mathrm{E}(\boldsymbol{m}^s \mid \boldsymbol{\theta})$. While the uncertainty regarding the simulated moments can be reduced by increasing the number of simulation runs, the uncertainty in the estimation of $\mathbf{m}$ has to be taken into account. Therefore, we propose to use $\mathbf{W} = \mathrm{Var}^{-1}(\mathbf{m}^e)$ in (3).

[7] For an introduction to the method of simulated moments see Stern (2000).

The operationalization of the objective function introduced in (3) requires a selection of moments $\mathbf{m} = (m_1, \ldots, m_k)$. These moments have to be chosen to match relevant characteristics of the time series under study and to contribute to the identification of parameters of agent based models. This choice and the resulting objective function is described in the next section.

3 The objective function

3.1 Selection of moments

The selection of moments has to take into account the eventual goal of the analysis, i.e. evaluation of models, discrimination between models and parameter settings, and indirect estimation of parameters. Consequently, moments have to be chosen in dependence of their ability to contribute to these goals.

We use two criteria for the selection of adequate moments. First, the statistics should be robust. Second, they should exhibit the potential to discriminate between alternative models and/or parameter constellations. The first aim of robustness is assessed by evaluating the statistics for different subsamples of the actual data and by conducting a rolling window analysis.[8]

To what extent the second goal can be achieved depends on three factors. First, the uncertainty of the estimates for the actual data sample is used as a benchmark. This factor is measured by means of the bootstrap variance of the moment under consideration for the real data sample. Obviously, the smaller this variance, the more reliable the estimate. Furthermore, a smaller variance of the estimates increases the probability to reject the null hypothesis that a moment generated by simulation comes from the same distribution as the actual data. Second, the variance of the expectation of the moment estimated from the simulated data should be small in order to be able to discriminate different parameter constellations for the same model. However, this variance can be controlled by repeating the simulation. Thus, this criterion is less relevant for the selection of the moments than the first factor, but it has to be taken into account in the implementation of the simulation model. Thirdly, the power to discriminate amongst modelling alternative depends on the identification of the model given the considered moments. Unfortunately, this aspect can hardly be assessed on a priori grounds due to the highly nonlinear structure of the agent based models considered. Thus, it is advisable to include a large number of moments taking into account all relevant characteristics of the actual data in order to increase the probability of identification of model parameters.

Keeping these criteria in mind, statistics describing the unconditional and the conditional distribution of exchange rate returns are selected. In future applications, this selection might be extended by also taking into account estimates of the sampling variance of these statistics. Furthermore, it is of interest to consider the properties of the different statistics under time aggregation both in the real and the simulated data.

[8] See Winker and Jeleskovic (2006, 2007) for details.

However, this extension is beyond the scope of the present paper and is left for future analysis.

The selection of moments describing the unconditional distribution of exchange rate returns is based on Winker and Jeleskovic (2007). The mean and standard deviation of returns provide standard measures of the overall shape of the unconditional distribution. Higher order moments, in particular the empirical kurtosis appeared to be less robust. This might be a result of rather heavy tails of the return distribution challenging the existence of theoretical higher order moments. Thus, instead of using standard higher order moments,[9] an estimate of the tail index of Pareto-type distributions is used. We consider the Hill estimator of the right tail.[10] The statistic entering our objective function is the mean of this Hill estimator over the 5–10% upper quantiles of the return distribution. Finally, besides considering specific features of empirical and simulated data, the overall shape of the return distributions can be compared. For this purpose, we use the classical Kolmogorov–Smirnov statistic measuring the maximum absolute difference between the cumulated empirical distribution functions of the actual exchange rate returns and the simulated returns, respectively. In particular, this measure provides an upper bound for the deviation at any particular quantile of the distributions.

The choice of statistics regarding the conditional distribution relies on the results in Winker and Jeleskovic (2006). As in Gilli and Winker (2003), the parameters of a standard GARCH(1,1) specification appear to be a useful summary statistic for short range dependence in higher moments of the returns. Obviously, the GARCH(1,1) specification is not a perfect model of the data generating process as it neglects long range dependence and dependence at different powers of absolute returns. Nevertheless, it is a robust and easy to calculate measure. The sensitivity analysis in Winker and Jeleskovic (2007) indicates that the sum of the ARCH- and GARCH-parameters, i.e. $\alpha + \beta$ in the usual notation, is a more robust statistic as compared to the individual parameters. Consequently, this sum is used as a base ingredient to our objective function. It might be augmented in future research by more refined measures of conditional heteroscedasticity. As a standard measure of long range dependence, we include the GPH-estimator of the degree of fractional integration.[11] Finally, the (close to) random walk property of exchange rates is taken into account by the ADF-statistic.[12]

Table 1 provides a summary of the statistics used in the objective function including specification details and the values for the benchmark exchange rate time series DM/US-$ for the period 11.11.1991 to 9.11.2000.

[9] Gilli and Winker (2003) included the empirical kurtosis in their first implementation of an indirect estimation procedure. However, this statistic contributed a large amount of Monte Carlo variance to the objective function without being very discriminating between different parameter constellations.

[10] The estimates of the tail index for the right tail appear to be slightly more robust than for the left tail. However, the tail behaviour is quite similar for both tails.

[11] The parameter m of the estimator is fixed at 0.5, i.e. only the first $\sqrt{T}$ frequencies are taken into account for the estimator. For more details, see Winker and Jeleskovic (2007).

[12] The test regression includes a constant, but no lags for the DM/US-$ exchange rate.

Table 1 Moments considered for the objective function

Label	Statistic	Implementation details	Benchmark value ($\mathbf{m}^e$)	Mean over bootstrap replications
m_1	Mean		0.00015	0.00015
m_2	Standard deviation		0.00649	0.00648
m_3	Tail index	Right tail, average $\hat{\alpha}$ (5–10%)	1.68033	1.70010
m_4	Kolmogorov–Smirnov	Empirical distribution as benchmark	NA	0.02461
m_5	GARCH(1,1):	$\alpha + \beta$	0.98877	0.98043
m_6	Fractional integration	GPH estimator for $\mid r_t \mid$ ($m = 0.5$)	0.43823	0.36287
m_7	ADF-test	With drift, no lags	0.01034	−1.28639

Obviously, this list could be easily extended to include further statistics, e.g. the behaviour of moments under temporal aggregation, or estimates of the sampling variance of some statistics.

3.2 Weighting

An estimate $\hat{\Sigma}_{BB}$ of the variance-covariance matrix Var($\mathbf{m}^e$) is obtained from the bootstrap distribution of $\mathbf{m}$ using a window length of 250.[13]

$$\hat{\Sigma}_{BB} = \begin{pmatrix} 1.6\times10^{-8} & -7.4\times10^{-9} & 3.7\times10^{-8} & 2.9\times10^{-6} & 5.4\times10^{-7} & 2.2\times10^{-6} & 4.9\times10^{-5} \\ -7.4\times10^{-9} & 1.5\times10^{-7} & 2.1\times10^{-8} & -1.6\times10^{-5} & -1.2\times10^{-6} & 8.4\times10^{-7} & -4.6\times10^{-5} \\ 3.7\times10^{-8} & 2.1\times10^{-8} & 1.0\times10^{-4} & -2.2\times10^{-7} & -1.1\times10^{-5} & -3.3\times10^{-5} & 6.1\times10^{-4} \\ 2.9\times10^{-6} & -1.6\times10^{-5} & -2.2\times10^{-7} & 0.00649 & 2.8\times10^{-5} & -0.00239 & 0.01045 \\ 5.4\times10^{-7} & -1.2\times10^{-6} & -1.1\times10^{-5} & 2.8\times10^{-5} & 9.6\times10^{-5} & 3.0\times10^{-4} & 0.00149 \\ 2.2\times10^{-6} & 8.4\times10^{-7} & -3.3\times10^{-5} & -0.00239 & 3.0\times10^{-4} & 0.00626 & 0.00615 \\ 4.9\times10^{-5} & -4.6\times10^{-5} & 6.1\times10^{-4} & 0.01045 & 0.00149 & 0.00615 & 0.74257 \end{pmatrix}$$

For the reasons discussed in Subsect. 2.2, we use the estimator

$$\hat{\mathbf{W}}(\mathbf{m}^e) = \hat{\Sigma}_{BB}^{-1}$$

of $\mathbf{W} = \mathrm{Var}^{-1}(\mathbf{m}^e)$ as the weighting matrix in our simulated method of moments approach. Alternative approaches for defining W will be considered in future work. Thus, our objective function for the evaluation of a parameter vector $\boldsymbol{\theta}$ based on simulated data is given by

$$\begin{aligned} f(\boldsymbol{\theta}) &= \frac{1}{I}\hat{\boldsymbol{G}}_I'\hat{\mathbf{W}}(\mathbf{m}^e)\hat{\boldsymbol{G}}_I \\ &= \left[\frac{1}{I}\sum_{i=1}^{I}[(\mathbf{m}_i^s \mid \boldsymbol{\theta}) - \mathbf{m}^e]\right]' \hat{\mathbf{W}}(\mathbf{m}^e)\left[\frac{1}{I}\sum_{i=1}^{I}[(\mathbf{m}_i^s \mid \boldsymbol{\theta}) - \mathbf{m}^e]\right]. \end{aligned} \tag{4}$$

[13] A large window size is required to take into account the long memory properties of the data (Winker and Jeleskovic 2007). The results are quite robust to different window lengths.

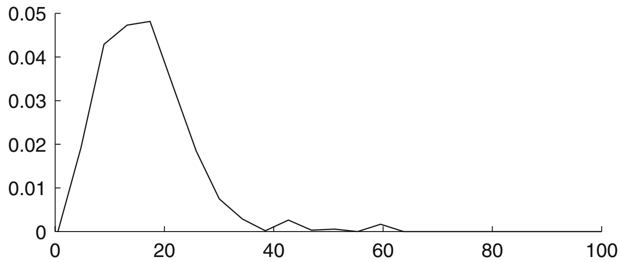

Fig. 5 Distribution of f for block bootstrap samples

For some models, a few simulation runs result in extreme values of $\mathbf{m}_i^s$. These outliers might be due to either the nonstationary characteristics of the data generating process for particular values of $\boldsymbol{\theta}$ or due to numerical problems of calculating the empirical moments. Therefore, we also consider replacing the means over replications in (4) by 95%-trimmed means and medians, respectively.

As a benchmark, Fig. 5 shows a kernal estimate of the distribution of the objective function f obtained for the 10,000 block bootstrap samples which have also been used for calculating $\hat{\Sigma}_{BB}$. Obviously, the quantiles of f obtained from the block bootstrap of the empirical data tend to exceed the quantiles of the χ^2-distribution with seven degrees of freedom. Thus, we will resort to these simulated values for testing specific parameter settings $\boldsymbol{\theta}$ for the simulation models in the following. The empirical 95%-quantile amounts to 35.54 and the 99%-quantile to 49.06.

4 Applications

In order to assess the properties of the objective function $f(\boldsymbol{\theta})$ introduced in the previous section we consider the application to two agent based models of the foreign exchange market. We are interested in three aspects of the objective function. First, for a given model and a given parameter set $\boldsymbol{\theta}$, we are interested in the dependence of the objective function values on the number of replications I of the model, i.e. at which computational cost the Monte Carlo variance can be controlled. Second, for a given model, it is analyzed whether the objective function can discriminate amongst different parameter settings, i.e. whether it can be assumed that the model parameters can be identified based on the objective function. This condition is central for a successful application of optimization techniques for the estimation of the model parameters as proposed by Gilli and Winker (2003). Third, the objective function might be used to compare different models. Thus, our interest is to see to what extent the two models analyzed differ with regard to our objective function.

4.1 Kirman

The first test case uses a model originally introduced by Kirman (1991, 1993) and considered by Gilli and Winker (2003) in an indirect estimation approach. The central ingredients of the model are the interaction between heterogenous, not fully rational individuals, and the endogenous evolution of the shares of agent types in the population over time. Individuals acting on the foreign exchange market are characterized as fundamentalists and chartists, respectively. Fundamentalists expect an adjustment of prices towards the fundamental value of the asset which is assumed to be known to all market participants. Chartists follow a momentum strategy, i.e., they expect the price change to be the same as in the previous period. The type of any individual can change for two reasons: random mutation with probability ε and conviction by direct interaction with a second individual of different type with probability δ.[14]

An implementation of Kirman's model can be fully described by a parameter vector $\boldsymbol{\theta}$ summarizing all model parameters (including ε and δ), a time path for the fundamental value and initial values for the price and the shares of agent types. Then, for a given random number vector $\mathbf{r}$, the model will provide a realization of a price time series, which can be used to calculate simulated moments $\mathbf{m}^s(\boldsymbol{\theta}, \mathbf{r})$.

Since we are not interested in the value of the objective function for a specific realization, but solely as a function of the models' parameters $\boldsymbol{\theta}$, the analysis has to be repeated for a large number of Monte–Carlo replications with differing random number vectors $\mathbf{r}_i$, $i = 1, \ldots, I$, each resulting in at least T simulated prices.[15] From the resulting Monte–Carlo distribution of simulated moments, we obtain an estimate of $f(\boldsymbol{\theta})$ based on the mean, trimmed mean or median, respectively. Of course, this estimate will be subject to some remaining Monte–Carlo sampling variance depending on the number of replications performed.

Algorithm 1 summarizes the approximation of the objective function over a grid of relevant values of the two central parameters δ and ε. Kirman (1991) shows that for his model the equilibrium distribution of the process for opinion formation and transmission of agents depends only on the conviction rate δ and the mutation rate ε.

Algorithm 1 Calculating the objective function.

for $j = 0 : n_j$ **do**
  Set $\delta = \delta_l + (\delta_u - \delta_l) \times j/n_j$
  for $l = 0 : n_l$ **do**
    Set $\varepsilon = \varepsilon_l + (\varepsilon_u - \varepsilon_l) \times l/n_l$
    Set $\boldsymbol{\theta} = (\delta, \varepsilon, \boldsymbol{\theta}_1)$
    for $i = 1 : I$ **do**
      Generate random vector $\mathbf{r_i}$
      Simulate the model
      Calculate $\mathbf{m}_i^s(\boldsymbol{\theta}, \mathbf{r_i})$)
    end for
    Calculate estimate of $f(\boldsymbol{\theta})$
  end for
end for
Plot $f(\boldsymbol{\theta})$ against δ and ε

[14] For further details on the model specification and model parameters see Kirman (1991) and Gilli and Winker (2003).

[15] If the model simulates several interactions per trading day, this fact has to be taken into account when determining the length of r_i.

Table 2 Parameter settings for Kirman's model

Label	Interpretation	Value
n_A	Number of agents	100
n_I	Number of interactions	2350×50
n_T	Number of interactions per trading day	50
ν	Adjustment speed in fundamentalists expectations	0.08
σ_s	Standard deviation of price shocks	$m_2^e/50$
σ_q	Standard deviation of noise in majority assessment	0.070
c	Price adjustment	0.5
δ	Probability for direct interaction	$[0.25, \ldots, 0.40]$
ε	Probability for random mutation	$[0.005, \ldots, 0.060]$

In particular, the relation between both parameters defines whether the process will exhibit bubble-like dynamics or not. Finally, $\boldsymbol{\theta}_1$ denotes all other relevant parameters of the model, i.e., $\boldsymbol{\theta} = (\delta, \varepsilon, \boldsymbol{\theta}_1)$.

Table 2 provides an overview on the other relevant parameters of the model and their numerical values for the present implementation. The fundamental price is assumed to be constant and equal to the mean of the empirical data.[16]

The results of the simulation with regard to the two parameters ε and δ are shown in Fig. 6. For these results, the number of replications I has been set to 100. From the left plot it becomes obvious that both parameters exhibit a pronounced influence on the objective function f. In particular, too large values of the mutation rate and too small values of the conviction rate result in simulated time series which differ considerably from the real exchange rate series with regard to the moments considered in this application.

Based on the findings in the left plot, a second simulation has been run on a smaller subset of the parameter space comprising those values which appeared most favourable in the first run. From the results in the right subplot it becomes obvious that although small values of ε result in better properties of the simulated time series, the optimal value is strictly positive, while for δ optimal values might be close to 0.35. Finally, one should note that given the fixed values of all other parameters of the model, the lowest values of the objective function f are of the order of magnitude of 200. Thus, compared to the bootstrap distribution of f for the real data, the model with the parameter vectors considered and the assumption about a constant fundamental price has to be rejected. However, other parameter vectors and/or a different assumption about the fundamental price might lead to different conclusions.

[16] We also experimented with alternative assumptions about the fundamental price. Assuming a random walk model does not improve the overall fit of the model, but results in "best" values of ϵ and δ close to zero. However, assuming a moving average over the last 200 observations reduces the value of the objective function although not below the critical value at the 99%-level (49.06).

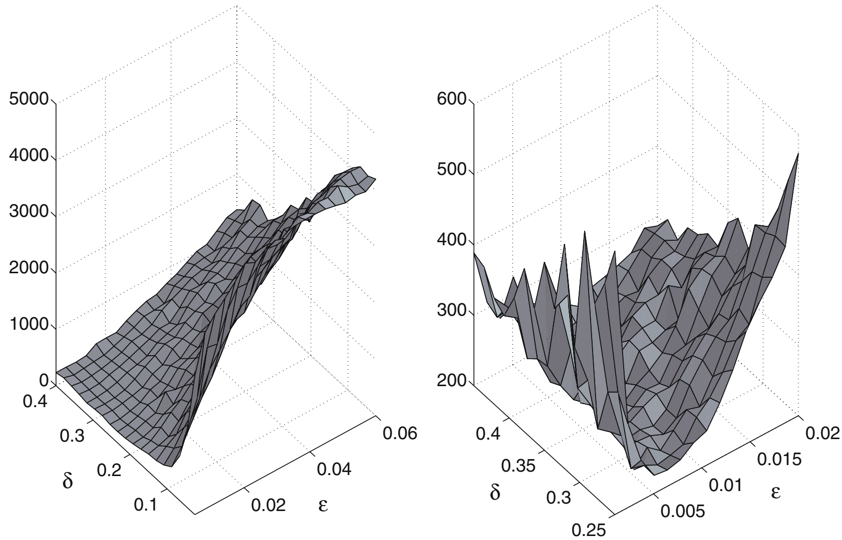

Fig. 6 Simulated response surface $f(\boldsymbol{\theta})$ for Kirman's model as a function of ε and δ

4.2 Lux

The models introduced by Lux (1998), Lux and Marchesi (1999) and Lux and Marchesi (2000) also build on the work by Frankel and Froot (1986). A fixed number of traders N is endogenously divided in three groups, fundamentalists (N_f) trading on mean reversion with regard to the fundamental price (p_f), optimistic chartist traders (N_+) and pessimistic chartist traders (N_-). Switching between these groups depends on passed success with the respective strategies. Lux and Marchesi (2000) argue that the main stylized facts of financial markets can be mimicked by this micro-simulation model.

The endogenous switching within the group of chartists is modelled as depending on the market sentiment and the current price trend. Market sentiment is measured through a so called *opinion index* x which is defined as

$$x = \frac{N_+ - N_-}{N_c}\,, \quad x \in [-1, 1]\,. \tag{5}$$

The direction of the price change (in continuous time) $\dot{p} = dp/dt$ is the second component which determines the transition probabilities to switch from the pessimistic group to the optimistic one (π_{-+}) and vice versa (π_{+-}):

$$\pi_{-+} = v_1 \left(\frac{N_c}{N} e^{U_1}\right); \quad \pi_{+-} = v_1 \left(\frac{N_c}{N} e^{-U_1}\right) \quad \text{where } U_1 = \alpha_1 x + \alpha_2 \frac{\dot{p}}{v_1}\,. \tag{6}$$

The model parameters α_1 and α_2 describe the sensitivity with regard to both influence factors, while v_1 represents a measure for the frequency of revaluation of opinion.

It is assumed that chartists buy or sell always a fixed number of units depending on the expected price change. The strategy of fundamentalists is to buy (sell) when the price at time t is below (above) the fundamental price p_f. The probability of switching between the group of fundamentalists and chartists depends on realized excess profits. The excess profit per unit for an optimistic chartist is given by $(r + dp/dt)/p - R$, where r is nominal dividends of the asset and R is average real returns from other investment. The motive of a pessimistic chartists, on the other hand, is to avoid losses by searching for other investment opportunities. Hence, his excess profit is given by $R - (r + dp/dt)/p$. Excess profit of fundamentalists is driven from the deviation between p_f and p at time t: $s \mid (p - p_f)/p \mid$, where s is a discount factor for expected excess profits which are realized only when the price has returned to its fundamental value ($0 < s < 1$).[17] The transition probabilities between fundamentalists and chartists are defined as follows:

$$\pi_{+f} = v_2 \left(\frac{N_+}{N} e^{U_{2,1}} \right); \quad \pi_{f+} = v_2 \left(\frac{N_f}{N} e^{-U_{2,1}} \right) \tag{7}$$

$$\pi_{-f} = v_2 \left(\frac{N_-}{N} e^{U_{2,2}} \right); \quad \pi_{f-} = v_2 \left(\frac{N_f}{N} e^{-U_{2,2}} \right) \tag{8}$$

where

$$U_{2,1} = \alpha_3 \left[\left(r + \frac{\dot{p}}{v_2} \right) / p - R - s \left| \frac{p_f - p}{p} \right| \right] \tag{9}$$

$$U_{2,2} = \alpha_3 \left[R - \left(r + \frac{\dot{p}}{v_2} \right) / p - s \left| \frac{p_f - p}{p} \right| \right]. \tag{10}$$

Lux (1998) assumes that prices react sluggishly to the presence of non-zero net demand. The actual prices are set by an auctioneer. Lux and Marchesi (2000) take over this hypothesis and adjusted the model for the case of discrete time. Excess demand of chartists (ED_c) and fundamentalists (ED_f), respectively, is given by:

$$ED_c = (N_+ - N_-)t_c\,; \quad ED_f = N_f \gamma (p_f - p). \tag{11}$$

Thereby, t_c denotes the fixed number of units that chartists either buy or sell. $\gamma > 0$ is the reaction speed of fundamentalists to price deviation from the fundamental value. Total excess demand is given by $ED = ED_c + ED_f$.

It is assumed that the reaction of the market maker is a stochastic process itself which is driven by the reaction speed β and some imprecision [modelled by a small noise term $\mu \sim N(0, \sigma_\mu^2)$].[18] Furthermore, Lux and Marchesi (2000) assume that the market price increases or decreases by a fixed tick size of one cent ($\Delta p = \pm 0.01$)

---

[17] In contrast, chartists' excess profits are immediately realized.

[18] The imprecision might stem from additional liquidity traders in the market with stochastic excess demand or an imprecise perception of the excess demand by the market maker.

according to the following probabilities of price changes:

$$\pi_{\uparrow p} = \max[0, \beta(ED + \mu)] \tag{12}$$
$$\pi_{\downarrow p} = -\min[0, \beta(ED + \mu)]\,. \tag{13}$$

This tick by tick movements of the price happen in micro-intervals during a trading period. The recorded price at the end of the trading period is given by the accumulated price changes. Lux and Marchesi (2000) found that a micro-time interval of length $\Delta t = 0.01$ works well for phases of "normal" market dynamics, while the phenomenon of volatility bursts requires a finer grid. Consequently, they propose to increase the number of micro-intervals per trading day by a factor of 5 ($\Delta t = 0.002$) whenever the frequency of price changes becomes higher than on average. For the present implementation, we used a version of the model without this endogenous changing of the frequency of price changes.[19]

In the simulation, Lux and Marchesi (2000) exclude absorbing states by enforcing that the number of agents of each type does not fall below four. The present implementation also enforces this lower bounds as the model exhibits quite unwelcome features when allowing to move into the absorbing states.

Table 3 summarizes the parameter settings used for the simulations presented in the following. The settings follow closely parameter set I in Lux and Marchesi (2000, p. 689). However, the implementation does not allow for a switch in the time increment, i.e. the number of microintervals per trading day $n_{\Delta t}$.[20] As for Kirman's model, the fundamental price is assumed to be constant and equal to the mean of the empirical data.[21]

The simulation of the objective function over a grid of values of two parameters keeping all other parameters fixed follows the same procedure as described in Algorithm 1 for Kirman's model. Given the higher complexity of the Lux model, a selection of only two key parameters for the simulations is less obvious. We are mainly interested in two aspects, namely the transition of agents' type and the pricing mechanism. With regard to the transition probabilities, the parameters ν_1 and ν_2 determining the frequency of revalution of opinion appear to have the strongest influence.[22] Given that the expected gains of the fundamentalist trading rules are realized only after achiev-

[19] The described method for determining price movements as a slow adjustment process results in the possibly unwelcome feature that the market will never be exactly in an equilibrium. In principle, the model could be modified to allow for market clearing in each period by solving $ED(p) = 0$ for the actual price p. See Alfarano et al. (2005) for a similar approach. However, this modification of the model is left for future research.

[20] This more general formulation of the model will be subject to further analysis.

[21] As for Kirman's model, we also tested the model using a random walk model for the fundamental price, which did not improve the overall fit of the model. In contrast to Kirman's model, the "best" values of the parameters (ν_1,ν_2) and (ν_2,γ) were not very much affected. Assuming a moving average over the last 200 observations as fundamental price, again reduces the value of the objective function although not below the critical value at the 99%-level (49.06).

[22] However, a final decision has to be based on the simultaneous optimization with regard to all parameters which has to be left for future research.

Table 3 Parameter settings for Lux's model

Label	Interpretation	Value
N	Number of agents	500
T	Number of trading days	2 350
$n_{\Delta t}$	Number of microintervals per trading day	100
ν_1	Frequency of revaluation [eq. (6)]	3
ν_2	Frequency of revaluation [Eqs. (7)–(10)]	2
α_1	Weight factor in transition equations [eq. (6)]	0.6
α_2	Weight factor in transition equations [eq. (6)]	0.2
α_3	Weight factor in transition equations [eqs. (9)–(10)]	0.5
β	Auctioneer's reaction speed	6
R	Average real return per trading day	0.0004
s	Discount factor for expected gains of fundamentalists	0.75
γ	Reaction strength of fundamentalist's excess demand	0.01
t_c	Trading volume of chartists	0.02
σ_μ	Standard deviation of excess demand noise	0.05
$\Delta p_{\min}$	Tick size	0.003

ing the fundamental price, the parameter of the reaction strength γ with regard to fundamentalist's excess demand should also exert a strong influence.[23]

The results of two simulations are presented in Fig. 7. While the left hand plot shows the results over a grid of values for ν_1 and ν_2, the right hand plot exhibits values of the objective function over a grid of values of γ and ν_2.

The left hand plot exhibits only a small part of the results obtained with regard to ν_1 and based on 50 simulations of the model at each grid point. Nevertheless, it might help to realize that the parameter ν_1 appears to be not well identified at least for the fixed values of all other parameters, while for ν_2 a clear minimum can be spotted. The question of identification of ν_1 will have to be considered again for optimized parameter values.

The plot on the right hand side is based on only 20 model simulations at each grid point. Nevertheless, a minimum of the objective function might be spotted for some inner values of the parameter set. However, at least based on only 20 model simulations per grid point, it appears that the objective function exhibits several local optima in this area. This feature of the objective function together with the Monte Carlo variance introduced by the indirect estimation approach will have to be considered in the estimation step following the proposal made by Gilli and Winker (2003). However, this part of the analysis is left for future research.

For both parameter settings considered, the minimum value of f is above 150 and 400, respectively. Consequently, as for Kirman's model, the model together with the

[23] An alternative would consist in varying the auctioneer's reaction speed β.

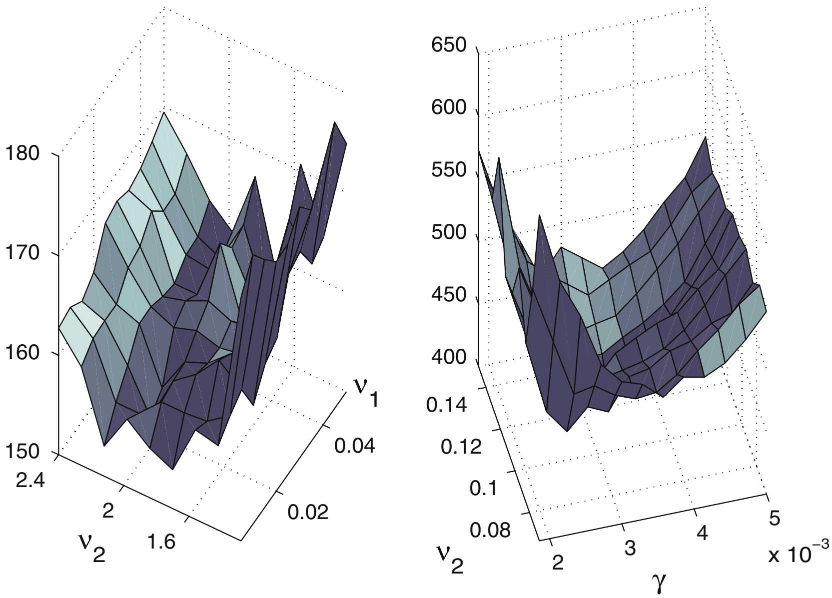

Fig. 7 Simulated response surface $f(\boldsymbol{\theta})$ for Lux's model as a function of (ν_1, ν_2) and (γ, ν_2)

parameter settings analyzed and the assumptions about the fundamental price has to be rejected.

4.3 Time series models

The simulation based approach introduced in Sect. 2 is not restricted to applications of agent based models. In fact, in recent years, financial market data have been subject to the analysis with an increasing number of time series models. Often, it has been claimed that these models can mimic some of the pertinent properties of financial market time series. When using time series models for the analysis of financial market data, two aspects deserve attention. First, being pure statistical instruments, this type of model might help to characterize or even forecast financial market time series, but it is not well suited to provide an economic rational for specific properties as is the case for agent based models. Obviously, this potential drawback of pure time series models cannot be resolved by the present approach. Second, it has to be checked to what extent the time series models exhibit all characteristics observed for real financial time series. The approach presented in this paper is well suited to answer this question.

GARCH(1,1) As a first simple example, Fig. 8 presents in the left panel the outcome of simulating a simple GARCH(1,1)-model for different values of the parameters α and β, while the initial variance has been set to the unconditional empirical variance of returns.

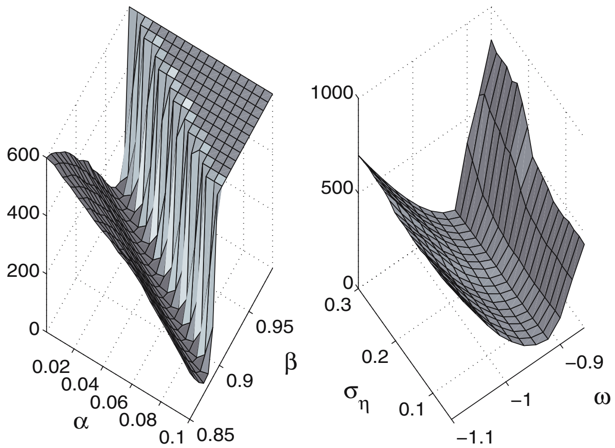

Fig. 8 Simulated response surface $f(\boldsymbol{\theta})$ for GARCH(1,1) and SV models

For values of α and β corresponding to nonstationarity, the objective function has been cut off at a value of 600 for the graph. The plot demonstrates that a GARCH(1,1)-model is best suited to reproduce the characteristics of empirical exchange rate data for values of α and β which in sum are close to unity. However, the parameters appear to be not well identified although a GARCH(1,1)-model is included as a component in the objective function f. This result deserves further attention. Finally, the best value of f obtained in this simulation exercise amounts to 33.24 which is below the 95% level of the bootstrap distribution. Thus, we are not able to reject the null hypothesis that the empirical data are generated by a GARCH(1,1)-model with the parameters corresponding to this minimum value of f. This is in contrast to our findings for the agent based models. However, two aspects deserve observation in this context. First, the agent based models contain many more parameters than the ones used in our first simulation analysis. Thus, it cannot be excluded that there exist parameter settings for the agent based models which also result in a non rejection of the null hypothesis. Second, as pointed out before, given that a time series model does not provide a deep economic intuition about the process we study, agent based models might be still interesting even if—in their current version—they are less well suited to mimic empirical data.

Stochastic volatility model As a second example of a time series model, we consider the stochastic volatility model for the time series of returns r_t given by van der Sluis (1997, p. 79):

$$\begin{aligned} r_t &= \sigma_t \varepsilon_t \\ \ln \sigma_t^2 &= \omega + \gamma \ln \sigma_{t-1}^2 + \sigma_\eta \eta_t \\ \varepsilon_t, \eta_t &\overset{iid}{\sim} N(0,1), \quad t = 1, \ldots, T\,. \end{aligned} \tag{14}$$

The estimation of (14) for our empirical data provided estimates $\hat{\omega} = -0.968$, $\hat{\gamma} = 0.903$ and $\hat{\sigma}_\eta = 0.159$. For the simulations, we set γ to its estimated value and fixed $\ln \sigma_0^2$ to its expected value $\omega/(1-\gamma)$ for given values of ω and γ. The right hand plot in Fig. 8 exhibits the response surface of f with regard to ω and σ_η. While ω appears to be well identified by the indirect estimation approach, the variation of f with regard to σ_η is less well pronounced. Nevertheless, we can spot a minimum of f with function value of 62.13. Consequently, for this parameter setting, the null hypothesis that the empirical data have been generated by model (14) has to be rejected.[24] However, it is of interest to note that the minimum of f is obtained for parameter settings close to the estimated ones. The indirect estimation approach appears to work well for this model.

5 Conclusion

Statistics used for the description of the properties of exchange rate return data are selected based on their robustness. Weighting the selected statistics according to their estimated sampling variance results in an objective function of a generalized method of moments type. This objective function is evaluated for two different agent based models and two time series models of the foreign exchange market to demonstrate its ability to differentiate between different parameter settings and models.

The simulation results indicate that the proposed objective function is well suited to exclude large parts of the parameter space for all models considered. For the time series models, the optimal values are obtained for parameter values close to those resulting from a direct estimation approach even for the stochastic volatility model which is not considered explicitly in the objective function. Nevertheless, the considerable sampling variability of some components of the objective function preclude the identification of a very narrow region of "optimal" parameter values. In fact, for some parameters of the agent based models, individual identifiability is not obvious from our results. Nevertheless, it is possible to reject the null hypothesis that the observed data have been generated with any of the agent based models for some of the considered parameter settings.

Obviously, these results have to be considered as a first step of the analysis. First, the selection of adequate statistics/moments has to be continued. In particular, the aspect of behaviour under temporal aggregation will be taken into account in future research.[25] Second, further agent based models and alternative parameter settings will be analyzed. Third, given the high dimensionality of the parameter space of agent based models, the type of grid search presented in the present contribution is not well suited to search for globally optimal parameter values. Thus, the eventual goal consists in using the methodology introduced in this paper for an indirect estimation procedure of the models' parameters along the lines suggested by Gilli and Winker (2003), i.e. using heuristic optimization methods. Besides the modified Nelder-Mead

[24] The stochastic volatility (SV) model is not considered in the objective function. However, both GARCH and SV model allow for an ARMA-presentation resulting in quite similar patterns of the autocorrelation function (Carnero et al. 2004). Therefore, one might conclude that the SV model is not unduly put at disadvantage in the present framework.

[25] For some results see Winker and Jeleskovic (2006, 2007).

simplex method proposed in their contribution, alternative optimization routines might also be considered.

References

Ahrens R, Reitz S (2005) Heterogeneous expectations in the foreign exchange market—evidence from daily DM/US Dollar exchange rates. J Evol Econ 15(1):65–82

Alfarano S, Wagner F, Lux T (2005) Estimation of agent-based models: the case of an asymmetric herding model. Comput Econ 26(1):19–49

Carnero MA, Peña D, Ruiz E (2004) Persistence and kurtosis on GARCH and stochastic volatility models. J Financ Econom 2:319–342

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

Gilli M, Winker P (2003) A global optimization heuristic for estimating agent based models. Comput Stat Data Anal 42(3):299–312

Heij C, de Boer P, Franses PH, Kloek T, van Dijk HK (2004) Econometric methods with applications in business and economics. Oxford University Press, Oxford

Hommes C, Huang H, Wang D (2005) A robust rational route to randomness in a simple asset pricing model. J Econ Dyn Control 29(6):1043–1072

Izumi K, Ueda K (2001) Phase transition in a foreign exchange market—analysis based on an artificial market approach. IEEE Trans Evol Comput 5(5):456–470

Kalaba R, Tesfatsion L (1996) A multicriteria approach to model specification and estimation. Comput Stat Data Anal 21:193–214

Kirman A (1991) Epidemics of opinion and speculative bubbles in financial markets. In: Taylor M (ed) Money and financial markets. Macmillan, New York, pp 354–368

Kirman A (1993) Ants, rationality, and recruitment. Q J Econ 108:137–156

LeBaron B (2006) Agent-based computational finance. In: Tesfatsion L, Judd KL (eds) Handbook of computational economics: agent-based computational economics. Handbooks in economics, vol 2. North-Holland, Amsterdam, pp 1188–1233

Lux T (1998) The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. J Econ Behav Organ 33:143–165

Lux T, Marchesi M (1999) Scaling and criticality in a stochastic multi-agent model of a financial market. Nature 397:498–500

Lux T, Marchesi M (2000) Volatility clustering in financial markets: a microsimulation of interacting agents. Int J Theoret Appl Finance 3:675–702

Stern S (2000) Simulation-based inference in econometrics: motivation and methods. In: Mariano R, Schuermann T, Weeks MJ (eds) Simulation-based inference in econometrics. Cambridge University Press, Cambridge, pp 9–37

Tesfatsion L (2006) Agent-based computational economics: a constructive approach to economic theory. In: Tesfatsion L, Judd KL (eds) Handbook of computational economics: agent-based computational economics. Handbooks in economics, vol 2. North-Holland, Amsterdam, pp 831–882

van der Sluis PJ (1997) Emmpack 1.01: C/C++ code for use with Ox for estimation of univariate stochastic volatility models with the efficient method of moments. Stud Nonlinear Dyn Econom 2(3):77–94

Winker P, Jeleskovic V (2006) The unconditional distribution of exchange rate returns: statistics, robustness, time aggregation. Technical report WP008-06, University of Essex, Colchester

Winker P, Jeleskovic V (2007) Dependence of—and long memory in—exchange rate returns: statistics, robustness, time aggregation. Technical report WP011-07, University of Essex, Colchester

Winker P, Maringer D (2007) The convergence of optimization based estimators: theory and application to a GARCH model (forthcoming)