Who Drives the Market? Estimating a Heterogeneous Agent-based Financial Market Model Using a Neural Network Approach

Klein, Achim and Urbig, Diemo

Klein Data Research, Backnang, Germany, Brandenburg University of Technology Cottbus-Senftenberg, Germany

24 June 2008

Who Drives the Market? Estimating a Heterogeneous Agent-based Financial Market Model Using a Neural Network Approach

Achim Klein [1*], Diemo Urbig [2]

April, 2011

Abstract. We propose a method for estimating complex heterogeneous agent-based models, especially their time-varying micro data, based on time-varying real-world macro data . We estimate the model at high frequency without posing simplifying assumptions on the model or the estimation process. We estimate daily time series of market participants' trading strategies, i.e., chartists and fundamentalists, at the S&P 500. For this context, heterogeneous agent-based models which explain macro market behavior by time-varying usage of strategies on the micro level have shown superiority to alternative models. Due to complexity, these agent-based models can hardly be directly estimated. As micro-level data from real stock markets are largely unobservable, model-free estimation methods cannot be applied to map macro to micro variables. Thus, we suggest a combination of both methods in terms of a model-free estimation of the *inverse* of an agent-based model, mapping macro to micro variables, which can then be applied to real-world macro data. Using an artificial neural network we estimate an inverse model of the heterogeneous agent-based financial market model introduced by Lux and Marchesi (1999) and apply it to S&P 500 data. Comparisons with previously estimated yearly time series and with historic events illustrate validity of the estimation results. Our results also contribute to the understanding of theoretical models.

Keywords: *Stock market, heterogeneous agent-based models, indirect model-free estimation, inverse model, trading strategies, chartists, fundamentalists, neural networks.*

JEL classification: C15, C22, C45, C81, G12

1 University of Hohenheim, Faculty of Economic and Social Sciences, Information Systems 2
Schwerzstr. 35, 70593 Stuttgart, Germany
ak@cryptocaptain.com

2 University of Antwerp, Faculty of Applied Economics, Department of Management
Prinsstraat, 13 (Z.107), 2000 Antwerpen, Belgium

* Corresponding author

1 INTRODUCTION

Financial markets with international investors are the backbone of our global economy. Understanding the dynamics of these markets is therefore crucial for understanding this global economy. A first step towards understanding these dynamics is to understand how prices evolve at financial markets. Research on modeling financial markets has made a significant shift during the last two decades. Realizing that Fama's (1970) assumption of homogeneous rational investors and efficient markets is not sufficient to describe real market dynamics (Kirman, 1992), researchers started following new paths of modeling market dynamics (e.g., Kirman, 1993, Brock and Hommes, 1998, and Lux and Marchesi, 1999). These new models are built on the idea that markets are determined by economic agents following heterogeneous strategies and that these strategies do not need to be fully rational (e.g., Shleifer and Summers, 1990, Le Baron, 2006). It was shown that boundedly rational behavior could survive in financial markets in the long run (compare Hommes, 2002). A typology of such strategies used by market participants, which is widely accepted among researchers and which has found theoretical (Aoki, 2002) and empirical support (Shiller, 1984, Menkhoff, 1998, Keim and Madhaven, 1995), is based on the distinction between fundamentalists and chartists respectively noise traders (Shleifer and Summers, 1990). Fundamentalists base their actions on fully rational estimates of true fundamental value of a financial instrument, while noise traders and chartists are less rational, because their actions are based on non-fundamental, imperfect, possibly inaccurate information, i.e., noise[1]. Since noise information can also include extrapolated historic price series (used for trend following) or charts (of e.g., stock price time series) the concept of noise traders also captures chartists (see also Brock and Hommes, 1998) and trend followers. The latter are therefore used as synonyms throughout this paper. By modeling the chartist and fundamentalist strategies as agent behaviors within heterogeneous agent-based market models, this kind of models has been very successful in replicating stylized facts related to real financial markets (compare for example Pagan, 1996, Cont, 2005).

Despite the success of heterogeneous agent-based models these models can hardly be used to directly estimate the micro behavior, e.g., the strategies of participants at real markets, based on observable macro data, e.g., price time series. This is due to complexity

[1] The term noise from an economic perspective was coined by Black (1986) and refers to the opposite of information (actually relevant for valuation of market assets).

reasons. Estimating micro variables for real markets would allow traders and investors to gain a more complete picture of the current market behavior and as a consequence to improve their decision making. Researchers using such heterogeneous agent-based models for estimating real market participants' behaviors – even if not a single participant's behavior but only fractions of participants using a certain type of strategy – need to balance a model's ability to realistically replicate stylized facts with its ability to be estimated (compare Amilon, 2006). Therefore, estimations are typically based on simplifying the model such that the parameters can be directly estimated with standard statistical methods. Estimates are therefore based on clearly defined but nevertheless substantially simplified versions of agent-based models (compare for instance Boswijk et al., 2007).

If identification of a small subset of model parameters is of interest but not the identification of all relevant model parameters, then simplifying the model to be able to completely estimate its set of parameters might go unnecessarily far in balancing complexity with tractability. Relaxing the requirement of a model-based estimation opens the way for using parameter-free estimation models, such as artificial neural networks (Haykin, 1999). These methods are better described as learning methods, because a system, e.g., a neural network, is trained (learns) to replicate a specific mapping behavior. A pre-requisite for such methods is the existence of sets of input-output pairs, which allows learning the inter-relation between known and unknown variables. Applied to our context, this requires data containing actual fractions of chartists and fundamentalists, i.e., micro characteristics, and the resulting market behavior, i.e., the macro behavior. After training the neural network appropriately, one would be able to estimate micro strategies based on market behavior. For our application, however, requiring such training data is a major limitation. We do not know of representative datasets containing such information. Thus, we propose a new method, which combines (1) heterogeneous agent-based models replicating typical behavior of real markets to generate the micro and macro level training data, and (2) a model-free estimation method to estimate fractions of chartists and fundamentalists in real markets based on an inverse model, mapping macro back to micro level data and trained on the training data from step 1. We call this an *indirect agent-based estimation method* and from our point of view this method can be applied to many other scenarios and research questions than the one addressed in this paper. We suggest using complex heterogeneous agent-based models, e.g., the model by Lux and Marchesi (1999), to feed a model-free estimation approach, specifically an artificial neural network, with as realistic pairs of micro and macro data as possible, while the trained neural network is then applied to real macro data in order to estimate the micro data. Note, that adding the agent-based model to supplement the model-free estimation approach makes the method depending on a model and we therefore have a model-based estimation method. However, the complexity of the model is not restricted except by the ability of the model-free estimation approach to be able to capture the dynamics of the agent-based model.

The main contribution of this work is an indirect agent-based estimation approach. It allows estimating micro-parameter time series at high frequency based on an underlying agent-based model of high complexity. Less simplifying assumptions concerning the model or the estimation process have to be applied for this procedure than for other procedure estimating heterogeneous agent-based models. Our work contributes to understanding what kind of micro-level behaviors drive stock markets. Analyzing dependencies between our estimation results and historic market events, we find the fraction of chartists being large at times of crises, crashes, and bubbles. Besides offering a new method for estimating empirical data, this study also contributes to the understanding of theoretical models. By investigating fundamental dependencies in the Lux and Marchesi model by means of sensitivity analysis of the resulting neural network inverse model, price volatility is found to be the key driver. This provides additional support to findings in Lux and Marchesi (2000). Some face validity for final real-world daily estimation results obtained from the S&P 500 is shown by comparing to results of Boswijk et al. (2007). This is the work which comes closest to our approach, albeit their model is simpler and estimation frequency is yearly. We find support for Boswijk et al.'s (2007) key findings of a large fraction of chartists during the end of the 1990s price bubble in technology stocks.

The remainder of this paper is organized as follows. Section 2 reviews related work and discusses the gaps that we are going to close. In Section 3 we propose a general model-free indirect estimation method. Section 4 applies the proposed method to estimate daily time series of fractions of chartist and fundamentalist strategies at the S&P 500. Section 5 discusses the results and addresses strengths and weakness of the proposed method. Section 6 concludes and points out avenues of further research.

2 RELATED WORK

In this work, we propose an indirect model-free estimation of time-varying endogenous daily time series parameters of fractions of chartist and fundamentalist strategies used by market participants. We specifically focus on estimating time series of strategy fractions in stock markets. The estimation is based on heterogeneous agent-based market models according to the behavioral, agent-based approach (Hommes 2006). As heterogeneous agent-based models have a rather complex mapping of to-estimate micro parameters to macro parameters, few attempts have been published. We will discuss a few attempts to estimate chartist-fundamentalist models.

2.1 Estimating agent-based models of financial markets

Vigfusson (1997) parametrically estimates a Markov-switching model version of the statistical chartist-fundamentalist foreign exchange model of Frankel and Froot (1988). Also, Westerhoff and Reitz (2003) directly estimate chartist and fundamentalist coefficients of a statistical exchange rate model of the Smooth Transition Autoregressive family. The model is simple as it does not simulate herding, simulates deterministic traders, uses linear trading rules and the impact of chartists is fixed, only the one of fundamentalists varies. Both, Vigfusson (1997) and Westerhoff and Reitz (2003) do not explicitly use models with distinct individual behavior within a regime. Alfarano et al. (2005) and Alfarano et al. (2006) go beyond this by trying to estimate some parameters of a multi-agent scenario with distinct individual behavior in a modified version of the stochastic chartist-fundamentalist model proposed by (Kirman 1993). However, their version of the model is very simple and also simplifying assumptions are posed during the estimation process.

Our focus on estimating time series is in contrast to Alfarano et al. (2005), who estimate some non time-varying parameters which quantify the overall prevalence of chartists and fundamentalists. Our focus on a stock index is in contrast to Vigfusson (1997) and Westerhoff and Reitz (2003), who estimate foreign exchange models. Boswijk et al.'s (2007) approach comes closest to our aims regarding estimating chartist and fundamentalist *time series* in the S&P 500 *stock* index. Boswijk et al. (2007) estimate a two regime version of the Brock and Hommes (1997, 1998) model (B&H model). Due to a simple noise term, they are able to apply a direct nonlinear least squares regression to the also not too complex single equation time series model. Based on a yearly time series of the S&P 500 and an estimate of fundamental value, derived from quarterly reported dividends, Boswijk et al. (2007) estimate *yearly* time series of coefficients of fractions of chartists and fundamentalists in the *stock* market. Based on these quantitative estimation results, the paper comes close to our approach and we will validate our results against theirs. However, in qualitative terms, like the model of Westerhoff and Reitz (2003), the model of Brock and Hommes (1998) is also rather simply structured as it also does not simulate herding, simulates deterministic traders and uses linear trading rules. Also concerning realism, the version of the model used, exhibits several deficits in resembling real market dynamics. Specifically, it primarily lacks in generating realistic autocorrelation structures of price time series, returns, and squared returns at daily frequency. One of the key findings of Boswijk et al. (2007) is the estimate of a huge fraction of chartists in the late 1990s. The authors claim that this is due to widely adopted trend following behavior, mainly in technology stocks, which led to huge price bubbles. Boswijk et al. (2007) leave it to future research to explore whether their approach would yield similar results at higher frequency, e.g., daily, data.

Due to computational burdens, which are due to complex mappings from to-estimate micro to macro parameters, the *direct* estimation approach, which is employed by all studies presented above, is limited to be applied to rather simple models with rather good tractability. The *indirect* calibration approach (see Fagiolo et al., 2006) is an alternative to overcome this problem. This approach allows for estimating parameters of more complex and presumably more realistic structural models, of which the statistical models are approximations, if they can be derived. This should result in econometric improvements. The term *indirect*, in this case, relates to the fact that parameters of the model are not estimated in the context of its own model, but are picked from unrelated empirical micro-econometric investigations, and/or chosen to guarantee that a simulated model matches some particular, unrelated features of historical macroeconomic data (Fagiolo et al. 2006). Examples of applications to heterogeneous agent-based chartist-fundamentalist models and either foreign exchange or stock market data are given by De Grauwe and Grimaldi (2006), Gilli and Winker (2003), and Amilon (2006). In these approaches, typically, non-time-varying model parameters are calibrated in an optimization process, such that some moments of the simulated model's macro-level parameters match the ones of empirical data. The optimization process is typically driven by a variant of the method of moments, e.g., efficient method of moments (Gallant and Tauchen, 1996). Concerning our aim of estimating time-varying endogenous chartist-fundamentalist time series, this approach is to our knowledge unsuitable, since only non-time-varying parameters can be estimated.

Summarizing and concluding the brief review, both the presented estimation approaches are insufficient for our aim of estimating time-varying endogenous chartist-fundamentalist time series parameters at daily frequency of rather complex and thus presumably pretty realistic models without introducing simplifications which potentially lead to econometric deteriorations. The approaches lack in at least one of the following points, mainly due to the complex micro-macro mapping of heterogeneous agent-based models:

(1) models estimated are simple in their structure (e.g., no herding, deterministic traders, no true multi-agent scenarios with distinct agent individuals within

each trading strategy regime)

(2) substantial simplifications and approximations are applied to the model and the estimation process,

(3) models lack in resembling realistic properties on the macro-level, i.e., stylized facts (such as excess kurtosis, volatility clustering, and price return autocorrelation structures),

(4) estimation frequency of time-varying parameters might be limited due to the availability of fundamental data, and

(5) only non-time-varying parameters are estimated.

To close the outlined gaps and to allow for an improved estimation of endogenous micro-level time series parameters of more complex, more computationally oriented multi-agent models at high frequency, we present *the indirect agent-based estimation approach*.

2.2 Agent-based models of chartists and fundamentalists in financial markets

Within our application context we have decided to use the agent-based model introduced by Lux and Marchesi (1999, 2000), which is introduced below in section 4.1.1. However, alternative models are available in the literature. We focus on pointing out chartist-fundamentalist, behavioral-oriented, dynamic, heterogeneous, agent-based market models. An overview of this category of market models is given by Hommes (2006). The models we discuss have in common that they simulate trading in one risk bearing asset in one market. Furthermore, heterogeneity is limited to a set of different strategies, from which agents should be able to choose in order to maximize profit or utility. Fundamental value is an uniformly (to all traders), exogenously given variable, generated by a preferably non-stationary process. This process is supposed to help in replicating non-stationary price time series of real markets, see for example Hommes (2002). Despite being quite simply structured, the market models reviewed generate quite realistic market dynamics.

Kirman (1993) describes an early simple agent-based regime switching model, involving stochastic mutation and conviction, inspired by ant behavior. The model can be interpreted as a chartist-fundamentalist framework with seminal herding in a true multi-agent scenario, i.e., distinct agents in each regime. Presumably, due to its simplicity, several authors have published on estimating some parameters of variants of the model, e.g., Gilli and Winker (2003) and Alfarano et al. (2005).

The model by Brock and Hommes (1997, 1998) is a simple adaptive belief system, in which agents can choose from a set of different beliefs or predictors of the future price of a risky asset. These beliefs, which reflect the ones held by chartists and fundamentalists, are revised in each period in a boundedly rational way. In contrast to the Kirman model, this model does not model distinct individual agents, but only one representative agent per class of agents. Some variants of this model have been estimated, e.g., Boswijk et al. (2007), De Grauwe and Grimaldi (2006), and Amilon (2006). An extended, more realistic version of the model has been formulated by Hommes (2002). It uses a non-stationary dividend process and a real market maker. Despite this extension, the model lacks in reproducing autocorrelation structures in prices, returns, absolute and squared returns, as well as other statistical properties.

Farmer and Joshi (2000) describe a simple behavioral model with deterministic chartist and fundamentalist strategies. Strategies are not switched according to profit or utility but become active at different points in time. The model has some deficiencies in replicating the statistical market properties of real markets. Also, the authors explore only conformance with some of the stylized facts. Carvalho (2001) extends the model with stochastic (concerning their activity) traders. However, the volatility autocorrelation function is not realistic, as it decays exponentially. To our knowledge, the model has not been estimated yet.

The model that to our knowledge well replicates realistic market behavior and is also in this context better than the other models mentioned above, is the one by Lux and Marchesi (1999, 2000). We base our further work on this model and will describe it in Section 4.

3 AN INDIRECT AGENT-BASED MODEL ESTIMATION APPROACH

In this section we propose a *general*, not application specific approach for estimating (aggregated) micro-level parameter time series of a realistic, highly complex agent-based model at high frequency. In the agent-based model, agents model some micro behavior and on the macro level the model aims on replicating real world behavior of which the micro part is not (completely) observable and thus subject to be estimated. Our general approach will be illustrated in our application context of estimating micro parameters of stock markets. The concrete *application* with a specific agent-based model to this context will be layed out in the next section. With this respect we are to our best knowledge one of the first to utilize the approach in the described context.

3.1 Rationale

First of all, we do not assume a very specific, i.e., simple to specify structural dependency between the (to be estimated) micro level parameters on one hand and the macro level behavior (that can be observed in the real world) on the other hand which could be captured directly by a parametric regression. Actually, we are not interested in the *structure* or functional form of the dependency. Rather, we are only interested in capturing it most *accurately*. This opens the door for utilizing one of the most flexible estimation methods, i.e., a model-free estimation. We thus propose an indirect agent-based estimation perspective using a model-free estimation of an inverse model at its heart.

The inverse model represents the dependency, i.e., it represents a generic mapping from the macro level behavior to the (aggregated) micro level (such as snapshots of the distribution of strategies used by agents on the micro level). By inverse we mean that the inverse model maps from macro to micro parameters instead of mapping from micro to macro parameters (in the underlying agent-based model). For our approach, macro-level parameters are assumed to be observable in the real world. Thus, the inverse model can, once estimated, be used to estimate (aggregated) micro parameters based on empirical real world macro level data as it represents a generic mapping from macro to micro parameters. As we employ an inverse model estimation as an intermediate step, the agent-based model is estimated only indirectly to achieve real world micro parameter estimations.

We propose to estimate the inverse model based on data generated in a preceding step by a simulation process which implements an underlying agent-based model. The agent-based model aims at replicating real world macro behavior by interacting micro level individual agents. This underlying model thus provides corresponding pairs of micro-level and macro-level data, which cannot be gathered empirically in the real world as at least the micro level is usally (not completely) observable. Further, we assume that the macro-level behavior of the model is realistic. This assumption is supported by comparing the model's macro-level behavior with empirical observations.

To realize the model-free estimation of the inverse model, we propose to employ a neural network approach. In terms of statistics, the neural network represents the inverse model of the agent-based model. In mathematical terms, the neural network represents a complex generic non-linear mapping from macro-level parameters to micro-level parameters. The specific mapping is obtained based on very few a-priori assumptions by a supervised iterated learning (i.e., training) process during which the structure of the network (i.e., the model) is adapted. This adaptation is guided by comparison of the mapping which the network represents and the desired mapping, represented by the simulated data which provides pairs of desired input (macro-level parameters) and output (aggregated micro-level) parameters.

3.2 Step by step description of the approach

The proposed indirect estimation approach consists of three major steps that are detailed below in a general, i.e., not application-specific, and pre-formal manner. As with respect to many steps no real theory is available, heuristics have to be applied.

(1) Simulation-based generation of corresponding micro and macro level data using an agent-based model that maps from the micro to the macro level.

  a. Choose an agent-based model A based on two requirements: (1) micro level parameters to be estimated in the real world are being modeled, and (2) macro behavior has been empirically found to be realistic.
  b. Implement and simulate the agent-based model over a pre-defined time span ts with fixed simulation parameters SP to generate corresponding artificial (1) macro level time series $mats_a$, and (2) micro level time series $mits_b$.
  c. Optionally. Aggregate individual agent behavior to aggregated micro level behavior time series $amits_b$(e.g., snapshots of the distribution of agents pursuing different strategies at discrete time intervals).
  d. Optionally. Verify statistical properties of the simulated macro-level time series $mats_a$ against properties of real world time series.
  e. Specify the set DO of desired (aggregated) micro level output parameters do_j for the neural network from the simulation's time series $amits_b$ or $mits_b$.
  f. Determine the set PI of potential macro level input parameters i_k for the neural network by heuristically choosing macro parameters $mats_a$ from the simulation.

(2) Estimation of the inverse model that maps from the macro level (input) to (aggregated) micro level (output) via a neural network approach. Essentially, this step describes neural network training. For more details, see e.g., Haykin (1999) or Masters (1993).

  a. Perform pre-processing of potential input parameter time series $i_k(t)$ according Bishop 1995 and Masters 1993: standardize (use only relative changes, cf. Sarle 2002), smooth (using a method SM, e.g., a moving average; we default to the Hodrick and Prescott filter (Hodrick and Prescott 1997)), and normalize (subtract the mean value from each element of the time series and divide by the standard deviation, cf. Masters 1993).

b. Heuristically determine network building parameters *NBP* as there typically is no theory (e.g., number of hidden neurons, cf. Swingler, 1996). Further parameters are (our default in brackets): number of hidden layers (one), number of neurons per hidden layer (5 for up to 4 input variables), activation function (hyperblic tangent, cf. Kalman and Kwasny, 1992), and training algorithm (bayesian regularization, cf. MacKay, 1992)) for the network training (see section 4 for details).

c. Divide each of the time series $do_j(t)$ and $i_k(t)$ of desired output and potential input parameter time series into three distinct continuous time series of length l_1, l_2, l_3, l_4 and associate the three parts respectively with the following sets (cf. Sarle 2002): (1) *Training set*: used to train the network, i.e., alter its weights while minimizing the error (difference between network output and desired output). (2) *Validation set*: used for early stopping of the training, i.e., if the performance metric on this set has worsened more than a threshold number of times since the last time it has improved, training is stopped. (3) *Test set*: used to adapt the set of input parameters of the network to train by verifying the generalization ability (i.e., robustness) on data not used in the training. (4) *Robustness test set*: For a final test of robustness of the best overall trained network, a separate second independent test set is used.

d. Determine the set $I \subseteq PI$ of macro level input parameters i_k for the inverse model.

    i. Heuristically or randomly choose a subset $I \subseteq PI$ of input parameters from the set *PI* of pre-processed potential input parameter time series.

    ii. Perform training of a neural network (cf. Haykin 1999) using the chosen macro level input parameter time series $i_k(t)$ with $i_k \in I$ and the specified desired (aggregated) micro level output parameter time series $do_j(t)$. Applying the resulting network to the input parameter time series $i_k(t)$ delivers the actual output parameter time series $o_m(t)$.

    iii. Evaluate the pre-determined set of stopping criteria *SC* on the *training set* and *validation set.* In bayesian regularization (used in this work for training), a weighted sum of squared network weights and sum of squared error (cf. Foresee and Hagan 1997) of the (estimated) actual output time series $o_m(t)$ vs. the desired (aggregated) micro level time series $do_j(t)$ is used as objective function. For stopping, typically, the objective function has to be below a pre-determined threshold *goal* on the *training set*, or the objective function's value has increased more than *maxvf* times since the last time it decreased on the *validation set*, or the maximum number of training iterations *maxiter* has to be reached. If none of the stopping criteria (see Table 1 for a full account) has been met, the training algorithm adjusts weights of the neural network based on the objective function's value on the *training set*. Else go to iv).

iv. Evalute the fitness of the trained neural network on the *test set* by determining the pearson correlation coefficient (*pc*) between (estimated) actual output time series $o_m(t)$ vs. the desired (aggregated) micro level time series $do_j(t)$. If *pc* is above a pre-determined threshold (we default to *pc*>0.9), go to i), else go to e).

e. Evaluate robustness of the latest trained neural network inverse model by determining the objective function's value on the *robustness test set*.

(3) Application of the inverse neural network model to real macro data to estimate real (aggregated) micro level time series.

a. Acquire real macro level time series corresponding to the final set I of the neural network inverse model's input parameters $i_k(t)$.

b. Perform pre-processing on the input parameter time series $i_k(t)$: standardize, smooth, and normalize as in step 2.

c. Apply the neural netwok mapping to the set of pre-processed input parameter time series $i_k(t)$. The output of the neural network is the estimated (aggregated) micro level time series $o_m(t)$ for the real word.

All parameters of the step-by-step approach and the allocation of the parameters in our application context of estimating micro parameters of financial markets are listed in Table 1.

- -

Table 1

Subsequently, Figure 1 illustrates our three-step approach within our application context that will be detailed in section 4 of this work.

- -

Figure 1

3.3 Advantages

The major advantage of the proposed method is that the estimation of the inverse model is decoupled from the simulation of the underlying agent-based model. Therefore, the agent-based model can be of high complexity. The "effective" complexity is the complexity that refers to estimating the micro parameters for real data. It is only restricted by the degree to which the model-free estimation approach can capture the complexity of the agent-based model. Utilizing a model-free approach, we, however, do our best in realizing wide instead of narrow limits in terms of complexity. To realize the model-free estimation we propose to employ a neural network approach.

The neural network approach has several advantages compared to other means of estimation. Theoretically, neural networks are able to approximate arbitrary non-linear mappings to arbitrary accuracy (Haykin 1999). As a neural network poses only very little constraints on the structure of the model it represents, it is better in learning (also partly discontinuous) mappings with multiple input parameters, when little is known about the relationship between input and output parameters (Sarle, 2002). Once the network has been trained, the network can be applied to new data very fast, in order to estimate current micro parameters in the real world. In contrast, direct estimation approaches would require costly re-estimating a model on the whole historic data set including the newest part.

Besides estimating characteristics of realistic data, we will illustrate that the neural network can be used to assess the underlying agent-based model itself. For instance, by analyzing the resulting neural network we are able to determine those parameters from the set of input parameters of the neural network that are most relevant within the agent-based model.

Summarizing, we propose a method of *indirect estimation* of an *agent-based model* by means of a *neural network approach*. Our approach allows for estimating time series of agent-based models of high complexity at high frequency without simplifying the model itself or posing constraints on the estimation process.

4 APPLYING THE INDIRECT AGENT-BASED ESTIMATION APPROACH TO THE ESTIMATION OF FRACTIONS OF CHARTISTS

In this section we apply the general approach outlined in the previous section to estimating fractions of chartists and fundamentalists in the S&P 500 index following the three steps. At the heart of our approach, we assume that an agent-based market model closely resembles real market behavior. Therefore, we have to carefully select and verify the agent-based model with respect to its ability to realistically resemble macro features of real market behavior.

Next, we describe the estimation of an inverse model of the agent-based financial market model based on a neural network. Since there are many different ways an artificial neural network can be structured, e.g., the number of neurons and the way these neurons are linked, and trained, e.g., the different training algorithms, we heuristically optimize the corresponding parameters. Furthermore, there are many possible macro variables that could be used as input for mapping to the micro-level variables, e.g., the price volatility, historic price intervals of different lengths, or price trends. We present just one out of many possible models. While we selected this model based on heuristically optimizing the mapping quality (as is explained below) and are therefore sure to have selected a reasonable one, further research might focus on the heuristic optimization and model selection step more thoroughly. In this study we focus on presenting the basic idea and providing of proof of concept of our estimation apprach.

Finally, we apply the neural network to real macro data in order to estimate corresponding micro level characteristics. These characteristics are then related to historic market events, which provide some face validity for the estimated strategy data.

4.1 Simulation-based generation of corresponding micro and macro level data using an agent-based model

4.1.1 The Lux and Marchesi agent-based financial market model

From the set of agent-based models describing financial markets, the model introduced by Lux and Marchesi (1999, 2000) is among those that get closest in resembling the statistical properties of time series of real markets. These specific statistical properties are the so-called stylized facts of real markets – concerning details we refer to Pagan (1996), Cont (2005), and Hommes (2002). Lux and Marchesi (L&M) focused their work on demonstrating agreement of their chartist-fundamentalist model with realistic price series (see Lux and Marchesi, 1999, 2000, and Chen et al., 2000). To our best knowledge, there are no published deficiencies in this context. However, Lux and Marchesi have not published on all possible aspects, e.g., autocorrelation functions. But considering the information available, this model is more realistic in contrast to other models considered. Comparing to Brock and Hommes (1997, 1998), L&M additionally implement stochastic traders (concerning switching of strategies) and model herd behavior of chartists. The herding aspect can be seen in line with Kirman's model (Kirman 1993). Nevertheless, L&M has a bit more detailed structures, as it for instance differentiates optimistic and pessimistic chartists. It also focuses more on replicating the stylized facts of real markets. Comparing to Farmer and Joshi (2000), the L&M model seems to be more realistic as it allows for varying fractions of used strategies due to dynamic alterations of beliefs and also models herd behavior. To our knowledge, L&M has not been estimated yet, which might be because of it being a bit more computationally oriented and more complex than the other models pointed out and thus being less tractable.

The L&M model will be subsequently described based on Lux and Marchesi (1999, 2000). The model basically has two groups of traders: the fixed number of traders in the market (N) is split into groups of fundamentalists (f) and chartists (c) with $n_c(t)$ and $n_f(t)$ denoting the numbers of agents in each group and $N=n_c+n_f$. The fundamentalist group expects the price to follow the fundamental value of the asset and thus buys (sells) when the actual market price is believed to be below (above) the fundamental value. Chartists follow price trends as a proxy for chartist practices, and also consider the behavior of other traders as a source of information, which results in a tendency towards herding behavior. Furthermore, the model distinguishes between optimistic (+) and pessimistic (-) chartists (who respectively believe in rising or declining markets) and the time-varying number of agents in their groups are denoted by $n_+(t)$ and $n_-(t)$ with $n_c(t)=n_+(t)+n_-(t)$. Optimists will buy additional units of the asset, whereas the pessimists will sell part of their actual holdings of the asset (Lux and Marchesi 1999, 2000).

The main building blocks of the L&M model according to Lux and Marchesi (1999, 2000) are movements of individuals from one group to another together with the (exogenous) changes of the fundamental value and the (endogenous) price changes resulting from the agents' market operations. Switching between groups, i.e, in all six directions between optimistic chartists, pessimistic chartists, and fundamentalists, occur with a certain endogenous and time-varying probability.

Switching between optimistic and pessimistic chartists is governed by the prevailing price trend $(dp/dt)/p$ and the majority opinion (determined by relating the sizes of each of the groups by $x=(n_+-n_-)/n_c$). The switching probabilities from the optimistic to the pessimistic group and vice versa within a small time increment Δt are given by $\pi_{+-}\Delta t$ and $\pi_{-+}\Delta t$ with

$$\pi_{+-} = v_1 \frac{n_c}{N} e^{U_1}, \qquad \pi_{-+} = v_1 \frac{n_c}{N} e^{-U_1} \quad \text{and with} \quad U_1 = \alpha_1 x + \frac{\alpha_2}{v_1} \frac{dp/dt}{p}$$

Parameters v_1, α_1, and α_2 are measures of frequency of revaluation of opinion and the importance of majority opinion and trend respectively. Switching between chartists and fundamentalists are driven by the difference between momentary profits g of individuals of each of the groups: $g_+ = \frac{a+\left(\frac{1}{v_2}\right)\left(\frac{dp}{dt}\right)}{p} - R$, $g_- = R - \frac{a+\left(\frac{1}{v_2}\right)\left(\frac{dp}{dt}\right)}{p}$, $g_f = s\left|\frac{p_f-p}{p}\right|$ with a being nominal dividends, R being average real risk-adjusted return available from other investments, p_f being fundamental value, and s is a discount factor for fundamentalists' profits with $s<1$. The switching probabilities are determined with:

$$\pi_{+f} = v_2 \frac{n_+}{N} e^{U_{2,1}}, \qquad \pi_{f+} = v_2 \frac{n_+}{N} e^{-U_{2,1}}, \qquad \pi_{-f} = v_2 \frac{n_-}{N} e^{U_{2,2}}, \quad \pi_{f-} = v_2 \frac{n_-}{N} e^{U_{2,2}}$$

with $U_{2,1} = \alpha_3\left(g_+ - g_f\right)$ and $U_{2,2} = \alpha_3\left(g_- - g_f\right)$ with parameters α_3 and v_2 being measures of sensitivity regarding profit differences between chartists and fundamentalists and frequency of revaluation of opinion (Lux and Marchesi 1999, 2000).

Changes of the fundamental value are given exogenously. The log of the fundamental value is determined as $ln\left(p_{f,t}\right) = ln\left(p_{f,t-\Delta t}\right) + \varepsilon_t \Delta t$ with ε_t being identically and independently distributed according to a Normal distribution with mean zero and (time-invariant) variance. Finally, price changes are determined by a virtual Walrasian market maker, which clears excess demand (or supply) $ED=ED_f+ED_c$ with excess demand for chartists $ED_c=(n_+-n_-)t_n$ and fundamentalists $ED_f=n_f\gamma(p_f-p)/p$ with t_n being the constant average trading volume per transaction and γ being a parameter for the strength of reaction on differences between p and p_f. Probabilities for the price to increase (decrease) by a small percentage $\Delta p = \pm\, 0.001p$ during a time increment Δt are given by $\pi_{\uparrow p} = max\,[0,\beta(ED+\mu)]$ and $\pi_{\uparrow p} = -\,min\,[\beta\,(ED+\mu)\,,0]$ with β being a parameter for the price adjustment speed and μ is a small random component (Lux and Marchesi 1999, 2000).

4.1.2 Implementing and verifying the properties of the model

We re-implemented the L&M model in the Java-based agent-based simulation environment RePast[2]. We verify the model and its implementation by comparing the statistical properties of our implementation of the model with published properties and with properties of real markets as for instance reported by Pagan (1996) and Cont (2001). Table 2 lists all selected properties for the L&M model's original implementation, as published in Lux and Marchesi (1999, 2000) and in Chen et al. (2000), as well as for our re-implementation, and for real markets. Clearly, we re-produce the results of L&M at large. Concerning realism, we additionally investigated two properties (see Table 2), which L&M have to our knowledge not published for their implementation. Again, our results are in good agreement with empirically investigated properties of real markets.

Table 2

Figure 2 plots the time series of the price and of the total chartist population in our market simulation. Prices are almost unpredictable, as the time series of price returns exhibits almost no self similarity (see Table 2). Price fluctuations alternate from tranquil to turbulent (i.e., highly volatile) in an intermittently fashion. As price fluctuates much more than fundamental value (which is normally distributed) as indicated by excessive kurtosis (see Table 2), price volatility is considered excessive. Also, the property of price volatility clustering can be observed, as big price changes occur in timely clusters (see Figure 2). This finding is supported by the time series' of volatility (i.e., absolute price returns) exhibiting significant self similarity (see Table 2). Also, the distribution of price returns exhibits heavy tails (see Table 2). All these properties can also be observed in real markets (Cont, 2005).

Figure 2

We further investigated autocorrelation functions of (raw, absolute, squared) price returns (see Figure 3), which L&M to our knowledge have not published. For absolute and squared price returns the function begins at significantly positive levels (up to 0.4) for small lags and then decreases slowly for larger lags. This is in agreement with the properties of real markets (Cont 2001). We observe one shortcoming though, concerning realism of the properties of the generated time series, which also has not been published by L&M: the autocorrelation function of raw price returns starts slightly negative for lags of 1-4 days (see Figure 3) and then approaches zero. Realistically, the autocorrelation function of raw returns

[2] http://repast.sourceforge.net/

would start slightly positive (at about the same quantity) for small lags and then would approach zero (for an autocorrelation function of raw S&P 500 index returns compare with Hommes, 2002).

Figure 3

4.1.3 Generating corresponding artificial micro and macro level data for training the neural network

Using Lux and Marchesi's (1999) simulation parameters, we generated time series sampled on a daily basis from the market simulation (with *ts*=9950 days in total length). These time series were used in all further work.

Since we aim in identifying current aggregate fractions of strategies used by market participants, these parameters of our market simulation comprise the set of potential output parameters for the network to train. Therefore, either the daily percentage of fundamentalists *fund(t)* or of all chartists *chart(t)* active in the market were in question. Both parameters are being computed based on all samples of the numbers of agents in the chartist (n_c) and fundamentalist (n_f) group within each day *t* during the simulation. Because of *fund(t)=1-chart(t)*, we choose *chart(t)* to be our sole output parameter. Further differentiating chartists into optimists and pessimists was ruled out, because we would not gain much information as the time series of these two parameters are strongly positively correlated with a correlation coefficient of 0.91.

As the potential input to the neural network we considered a set of different variables: price *p(t)*, historical price intervals of various lengths τ: $p(t)-p(t-\tau)$, price returns $r(t)=ln(p(t)) - ln(p(t-1))$, price volatility as the square of the price returns: $r(t)^2$, relative price trends $rpt(\tau) = \frac{p(t)-p(t-\tau)}{\tau \cdot p(t)}$ of length τ and differences in relative price trends of different lengths, e.g., $rpt(\tau_1)-rpt(\tau_2)$.

According the general approach outlined in section 3.2, we divided all time series (length *l*of continuous days in brackets) into (1) training set (l_1=5000), (2) validation set (l_2=500), (3) test set (l_3=946), and (4) robustness second test set (l_4=946).

4.2 Estimation of the inverse model that maps from the macro to aggregated micro level via a neural network approach

We now proceed to estimating the inverse model of the agent-based model via a neural network approach. A neural network is represented by a set of neurons which can be activated at different levels and that are linked with one another. Input neurons become activated at a specific level that depends on the input to the neural network; their activation is transferred to other neurons, which themselves become activated. These neurons then link to other neurons or output neurons. The activation of output neurons represents the output of the neural network. A neural network training algorithm adjusts the weights describing how much of a neuron's activation is transferred to another neuron in order to make the network linking desired outputs to pre-specified inputs. The complex weighted network is an implicit representation of the relation between input and output. For more detailed information on neural networks we refer to Haykin (1999).

To create a neural network which well represents the mapping from input to output parameters, structural parameters of the network, training algorithm parameters, and input parameters have to be selected and resulting networks' quality has to be compared based on the selections. This comparison is based on two criteria: (1) a normalized mean squared error (NMSE)[3] and (2) the Pearson correlation (*pc*) between network output and desired values (we aim at values of $pc > 0.9$ on test data that is not part of data used to train the network). As there is no theory, we heuristically optimized the structure of the network including the number of neurons, the training algorithm, but also the selection of macro variables used to map to the fraction of chartists based on the mentioned measures.

4.2.1 Network and training algorithm

As the basis for our neural network estimation procedure we use a three layer perceptron with a feed-forward network topology. Choosing the right number of neurons on the hidden layer is critical in obtaining good training results (Masters 1993). Choosing too less, the network might not be able to capture the inter-relationships properly. Choosing too much, training times increase and the net might be prone to overfitting, i.e., greatly reduced generalization ability. As there is no theory but only rules of thumb for choosing the number of hidden neurons (cf. Swingler, 1996), we used 5 neurons in the default case and in the case

---

[3] the normalized mean squared error (NMSE) between actual outputs o and desired outputs d, summed up for every neuron j (with N being the number of output neurons) and every pattern i in the data set, which comprises K tuples in total:

$$NMSE = \frac{\sum_{j=1}^{N}\sum_{i=1}^{K}(o_{ij}-d_{ij})^2}{\sum_{j=1}^{N}\frac{K\cdot\sum_{i=1}^{K}d_{ij}^2-\left(\sum_{i=1}^{K}d_{ij}\right)^2}{K}}$$

of more than 4 input variables (equivalent to four input neurons) we used the number of input neurons plus one as the number of neurons on the hidden layer. We chose the hyperbolic tangent as the activation function for neurons. It is from a class of frequently used functions (Masters, 1993) and Kalman and Kwasny (1992) argue that it is even the ideal function. With this setup, we should be able to meet the requirements of a multi-layer-perceptron with continuous non-linear activation functions in the neurons of one hidden layer – with a sufficiently high number of neurons – which enables us to approximate arbitrary non-linear continuous inter-relationships to arbitrary accuracy (Haykin, 1999). According to Masters (1993) even discontinuities can be tolerated.

The training of a network alters the weights such that the objective function is optimized. The initial weights are randomly determined and subsequently adapted by a training algorithm. Back-propagation can be considered as the standard algorithm (Reed and Marks, 1999). Because this algorithm converges very slowly, we selected the Gauss-Newton approximation to Bayesian regularization (GNBR) – an algorithm for neural network training proposed by Foresee and Hagan (1997), which combines Bayesian regularization (MacKay, 1992) with the Levenberg-Marquardt training algorithm (Hagan and Menhaj, 1994). By utilizing numerical optimization techniques, the Levenberg-Marquardt algorithm provides for much faster convergence on networks with no more than a few hundred weights, which applies to our case[4]. Additionally, Bayesian regularization improves generalization of the trained networks by constraining the size of the network weights. For this, according to Foresee and Hagan (1997), the objective function becomes $F = \alpha E_D + \beta E_W$, where E_W is the sum of squares of the network weights and E_D is the sum of squared errors. The parameters α and β are determined automatically by the GNBR algorithm, such that generalization of the trained network is optimized. This forces the network response to be smoother and less susceptible to problems related to overfitting.

4.2.2 Input and output data

The potential set of input parameters has been defined in section 4.1.3 for the given

[4] The GNBR algorithm can be adjusted by choosing a set of training parameters. We set the maximum number of training iterations to *maxiter*=500, the performance goal for the objective function on the training set as *goal*=0.001, the Marquardt adjustment parameter μ to 0.005, the decrease factor for μ as μ_{dec}=0.1, the increase factor for μ as μ_{inc}=10, and the maximum value for μ as μ_{max}=10^{10}, the maximum number of validation failures to *maxvf*=5, and the minimum performance gradient equal to *mingrad*=10^{-10}.

output parameter of the fraction of chartists. To obtain good results in training the neural network according to the metrics, input and output data first needs to be pre-processed appropriately (Bishop 1995, Masters 1993). Following (Sarle 2002), we standardize all input parameters by using only relative changes.

Taskaya-Temizel and Ahmad (2005) point out that neural networks have difficulties modeling non-stationary processes. This is the case, for instance, if the time series of an input parameter contains trends. Thus, trends usually are removed from these time series. As Taskaya-Temizel and Ahmad (2005) point out, differencing is a method to remove trends from time series. A further, more advanced method to remove trends is the Hodrick and Prescott (HP) filter (Hodrick and Prescott 1997). The HP filter is a standard tool used for smoothing macro-economic time series which are published quarterly. By adjusting the filter-specific smoothing constant, we are able to apply it to our data of daily frequency. The smoothing due to this filter, i.e., having training data not containing large extreme values or discontinuities, further improves the training quality of the network (Sarle 2002; Reed and Marks 1999). Best results were achieved using the HP filter with a smoothing constant of λ=100000, having tested several values in the range between 5000 – 500 million.

By further normalizing all time series of input and output parameters by subtracting the mean value from each element of the time series and dividing by the standard deviation, we can eliminate negative effects which result from differing offsets and scaling (Masters 1993).

On heuristically optimizing the set of input parameters for the given output parameter (fraction of chartists) of the neural network to train with respect to the defined metrics, only small sets of input parameters yielded good results on smoothed input data. In fact, price volatility ($r(t)^2$) as sole input parameter produced the best result, see Figure 4 (NMSE=0.08, pc=0.96 on the out of sample test data set), with additional input parameters such as relative price trends even worsening results. This result is most likely due to the fact that the underlying market model reproduces excessive and clustered volatility as features of real markets as a central feature of the model (Chen et al. 2000).

Figure 4

We verified robustness of the results of the best trained network by applying it to another separate data set that has not been used for training nor for evaluating training results. While absolute values of the performance metrics (NMSE=0.2; pc=0.91) differ from the first data set, results are still good, see Figure 5.

Figure 5

4.3 Application of the inverse neural network model to real macro data to estimate real aggregated micro level time series

After training the neural network and thus generating a model free estimation of the inverse model of the agent-based financial market model introduced by Lux and Marchesi (1999), we now apply the neural network to real stock market data in order to estimate the daily fraction of active market participants that are chartists. Like Boswijk et al. (2007) we chose the S&P 500 as a reasonable proxy for a real market. As the neural network does not require the fundamental value as input parameter, applying it to real market data is straightforward and does not require additional estimations. We compute the daily price volatility on a daily S&P 500 price time series (averaging daily open, high, low, and close prices) and apply the same normalizing and smoothing methods, i.e., the HP filter, as we did on the training data. Figure 6 shows the resulting estimated daily time series of the fraction of chartists in the S&P 500 stock market index.

Figure 6

According to our estimation results, the fraction of chartists in a market is high at times of crises, crashes, and price bubbles. The case of a bubble has been exemplarily illustrated above with the tech bubble, of which in Figure 7 the culmination is marked by point (3) and the final burst by (5). Examples of crises are points (1), which marks the East Asian financial crisis (Radelet, 1998), and (2) which marks the Russian financial crisis (Kharas et al., 2001). An example for a small crash is September 11 in 2001 (4) and a big one in October 1987 (not depicted) with more than 80% of chartists.

Figure 7

Firstly, a market strategy founded explanation for the not-obvious relationship between fundamental market events like a currency crisis, we have pointed out, and a large fraction of chartists is that first fundamentalists react on news, thus inducing price changes. Chartists by definition then identify these price changes, perceive them as trends, and follow them by opening positions in the market, thus reinforcing the trend. Due to further trend-following and herd-behavior, the fraction of chartists increases as chartists create a self-fulfilling prophecy and a self-reinforcing process. As chartists do not consider the fundamental value in their trading decisions, this kind of process also is the driving force of deviations of the market price from the fundamental value, for which we used the term bubble, above.

Secondly, the relationship between a large fraction of chartists and the market events we point out: crises, crashs, and price bubbles can be explained model intrinsically by the predominance of high price volatility which accompanies this kind of events (for example for the case of the 1987 crash, compare Schwert 1990). Price volatility is also a central feature of the market model of Lux and Marchesi (1999) and is also the central parameter we use for estimating fraction of strategies used in a market. As Lux and Marchesi (1999) point out that a large fraction of chartists is active at periods of high volatility in their model, it therefore is reasonable that we identify a large fraction of chartists during the market events mentioned.

However, a small fraction of chartists arises after the burst of a bubble, i.e., deviations of the price from the value are reduced by a growing fundamentalist regime. Also, the fraction of chartists shrinks when there are no clearly identifiable trends and low volatility.

5 DISCUSSION

In the last section, we applied and empirically validated our method of indirect estimation of agent-based models at daily frequency. We generated micro- and macro-level market data using a simulation based on our re-implementation of the L&M model which has been verified to closely resemble properties of real markets. Using the HP filter for pre-processing this data and the GNBR training algorithm, we made the neural network accurately and robustly learn an inverse model which maps macro-level parameters back to micro-level aggregated strategy parameters. We applied the inverse model to S&P 500 data to estimate daily time series of fractions of strategies used by market participants. Finally, we validated our results based on empirical background information on the events that have affected the usage of strategies in the stock market. In this section we discuss our estimation results and will provide additional analyses for the estimated agent-based model.

5.1 Comparison with previous estimations of yearly time series

Boswijk et al. (2007) is one of the more recent publications which by estimating yearly fractions of chartists and fundamentalists in the S&P 500 comes close to our estimation objectives for stock market strategies. Boswijk et al. use a statistical direct estimation approach and therefore had to stick to a relatively simple model. In contrast, we use an inverse model indirect estimation approach which makes very few assumptions on the model to be estimated and allows for using a more complex and more realistic model at higher frequency. Figure 8 compares estimation results of both approaches.

Figure 8

We reproduce the high percentage of chartists at times of the technology bubble (Western, 2004) at the end of the 1990s, which Boswijk et al. point out as a central finding in their paper. Despite not reproducing absolute values, we roughly reproduce relative changes, which is more important. In this context, Boswijk et al. (2007) leave it to future work whether they would find similar results using their approach at higher data frequencies. As we estimate at daily frequency in contrast to Boswijk et al. who estimate at yearly frequency, our curve of the estimated fraction of chartists is more detailed and naturally shows more diffentiated patterns. At large our curve supports the findings of Boswijk et al. of large fractions of chartists starting from 1996/1997. In contrast to Boswijk et al., our chartist curve is more volatile with some periods of low estimated chartists' activity in between. Also, our curve starts to decrease 1.5 years later, marking the end of the technology bubble. This is due to different market models we apply. The B&H model which Boswijk et al. use, defines chartists to believe in continuation of the increase in the deviation of the market price vs. the fundamental value. As the fraction of fundamentalists in the B&H model increases, they drive prices more to the fundamental value, thus bursting the bubble. In contrast, the model of L&M, which we use, seems more accurate by differentiating optimistic and pessimistic chartists, which are defined to consider upward and downward trends respectively and also model herd behavior. In this case, pessimistic chartists collectively strengthen downward trends, by selling themselves, in addition to fundamentalists, which initially started to sell. Thus, the total fraction of chartists stays large, until the selling-off has come to an end in 2003.

Considering the pre-technology bubble time period during 1992-1995, we estimate the fraction of chartists to be very low (<10%) while Boswijk et al. estimate considerable fractions (40-80%) to be active. When considering the price chart of the S&P 500 it becomes clear that during this time period the price development was rather sideways in comparison to more erruptive, up- and downtrending, and breakout-like price developments after 1996 which are accompanied by higher volatility. Considering the L&M model's definition of chartists which is based on trend following and herding, it seems fair that during the 1992-1995 time period the fraction of chartists was estimated comparatively low.

Concluding, we believe that the model of L&M seems more appropriate, also by being more realistic, as we point out in section 4. Additionally, our method of estimation provides for significantly more information value due to much higher data frequency. Also, once trained, the neural network can be applied to arbitrary markets with little effort.

Boswijk et al. (2007) claim that more work is needed to investigate robustness of their empirical finding that behavioral heterogeneity of market participants explains financial market data. Behavioral heterogeneity manifests in time-varying fractions of different strategies used by market participants. In this sense, our work also contributes to this research by investigating higher frequency data than Boswijk et al do. As we estimate strongly time-varying fractions of chartist and fundamentalist strategies at daily frequency, our work gives additional support to their claim.

5.2 Limitations

While we believe that this study provides an innovative method for estimating agent-based models, we are aware that the method as well as the specific way in which we apply the method in our example, has weaknesses arising from various decisions made during the study. We hope that this discussion can help other researchers in further developing and applying the method.

5.2.1 Global versus local smoothing

To interpret the estimated chartist time series correctly, one has to bear in mind the global HP filter used for data pre-processing. This filter was chosen because of yielding much better training results than a local filter (e.g., a moving average) which uses only historic data. However, the global filter considers always the whole time series on determining a single processed value of a time series, thus enriching every element of the time series *before* the end with future information. Considering the current point in time, i.e., the end of the time series, networks using input data which has been smoothed using a global and a local filter respectively, yield almost identical results. But as new elements of the input time series become available, older elements of the estimated chartist time series (using globally filtered input data) are subject to adaptations. To illustrate this effect and thus be better able to interpret estimated chartist time series, we subsequently show the difference between chartist time series based on globally smoothed inputs vs. locally smoothed inputs in four different sections: flat, ascending, descending, up-and-down.

Figure 9 shows the overall picture with all four sections over a total time period of 231 days. The chartist time series have been obtained with the best neural network, trained on globally smoothed data, given (1) the same globally smoothed input time series using the HP filter and (2) locally smoothed input time series. For local smoothing of a time series $x_1..x_t..x_n$, we utilized the weighted moving average $MA_t = \frac{\sum_{i=1}^{20} w_i \cdot x_{t-i+1}}{\sum_{i=1}^{20} w_i}$ calculated over a window of 20 time steps. The weights are calculated as parabolic weights $w_i = a(1-i^2)$ with a=0.7 and $i \in \{-1, -1+1/19, -1+2/19, ..., 0\}$.

Figure 9

The following four figures show in detail for each of the four sections how the chartist time series based on globally smoothed inputs changes in already available historical sections of the time series when new information becomes available, i.e., time progresses and new data becomes available for the time series. This is in contrast to the time series with locally smoothed inputs which does not change.

For the flat section, displayed by the following figure, there is almost no difference between the two time series, compare Figure 10.

Figure 10

The ascending section in Figure 11 shows that the ends of both time series are located almost at the same point. As the time series expand (sub-figures on the right hand side), the time series with globally smoothed input was determined with more future input, if we take point 71 as reference point. This new ascending information on the time series subject to global smoothing (in right hand side sub-figures) leads to an earlier ascend in comparison to the left hand sub-figures which do not yet contain the ascent. That is, the ascent on the time series with globally smoothed inputs is “moved” into the past, i.e., it starts earlier, as at the end of the time series more ascending points are added.

Figure 11

The descending section in the Figure 12 shows that on the time series with global inputs, the descent is "moved" a bit into the future as the time series is getting longer (from left to right hand side sub-figures) considering a certain reference point in comparison to the locally smoothed time series.

Figure 12

Figure 13 shows an up-and-down-movement of the time series based on locally smoothed inputs. The difference to the time series with globally smoothed input is that the latter does not take the first small hike as quickly. Also, the third sub-figure shows that small intermediate hikes in the middle of the time series are smoothed out by the curve with globally smoothed inputs, as globally there has become enough information available that the hike has been intermediate. This is only known when considering the whole time series which the global filter does. When only incomplete information is available (sub-figure 2), the global filter would like the local filter anticipate an ascent.

Figure 13

Concluding, when interpreting the chartist time series based on globally smoothed inputs to the neural network, one has to bear in mind that compared to the one based on locally smoothed inputs,

- it is almost identical in flat sections.
- it "moves" ascends into the earlier parts of the time series, when the ascend increases as new information becomes available at the end of the time series.
- it "moves" descends into the latter parts of the time series, when the descend increases as new information becomes available at the end of the time series.
- it smoothes out small intermediate hikes.
- it has almost the same ending points, which implies, that it "reacts" (almost) as quickly to newly developing ascends at the end of the time series in the same way.

5.2.2 Systematic exploration of the inverse model, i.e., the neural network

A weakness of neural networks as a method to estimate the dependency between variables is the fact that the estimation result is implicit in the network's weights. For instance, in regression analysis one can conclude, for instance, from an insignificant coefficient of an independent variable that the dependency between this and the dependent variable is not likely to be linear. In contrast, for neural networks one can (in most cases) not derive such direct conclusions about how the dependency looks like. However, we suggest that the network can be used to systematically explore the estimated dependency.

In the case of the best neural network trained on smoothed data, price volatility is the only input parameter, which makes it easy to explore the estimated model. Figure 14 plots how the fraction of chartists depends on the price volatility. To generate this figure we systematically vary the input to the network and plot the corresponding output.[5] Basically the figure shows that the fraction of chartists increases with an increase in price volatility – the dependency is, however, non-linear.

Figure 14

Applying the mapping that the neural network represents to real market data, i.e., price volatility time series provides hints to the non-linear dependency between price volatility in real markets and the estimated fraction of chartists in this market (see Figure 15).

Figure 15

5.2.3 Limitations of agent-based model plus model-free estimation method

As discussed in the introduction, direct estimations of agent-based models frequently need to employ a set of substantially simplifying assumptions. Our combination of a model-free estimation with an agent-based model does not require such strict assumptions. The quality of our estimation approach is however also limited by various factors.

If the agent-based model does not well represent reality, then the estimation procedure will not be able to estimate realistic output values for the real-world data. As such, our estimation method is model-based and more specifically it is based on the agent-based model. Concerning the L&M model which we utilize, we observed a shortcoming concerning autocorrelations of price returns. For small lags, the autocorrelation curve of raw price returns starts at a slightly negative level, as opposed to a slightly positive one, which would be realistic (Hommes 2002). We leave it to future research to employ other, more realistic, and

[5] If one had more input variables, one could keep all variables except one constant and plot the resulting out put (we did this for some networks, but do not report the results here). Similarly one could generate plots of potential interaction effects.

also more complex models.

Despite neural networks are considered to be able to represent complex dependencies, we have already mentioned limitations such as difficulties to represent non-stationary processes. The estimation method is therefore also restricted by the limitations of the neural network approach. Note however, that these limitations are far less restricting than the assumptions usually made for directly estimating agent-based models.

Furthermore, the application of the neural network approach and especially the network training process is a heuristic optimization process – estimation results usually satisfy a specific minimum threshold with respect to the quality they reflect the dependency, but they are likely to be sub-optimal. Despite having achieved pretty good results concerning accuracy and robustness, better results are still possible. In this context, we leave it to future research to explore other training algorithms, other, more elaborate optimization metrics, and other data pre-processing methods. One could even go further and apply other model-free estimation procedures in combination with an agent-based model.

5.2.4 Simple neural network

Note that the best neural network that we have found is a very simple one, which utilizes only one input variable, which is price volatility. There are two points to be discussed with respect to this observation.

First, the fact that we have not restricted our model to one variable but that this model emerged out of a larger set of models including many different variables indicates that the dependency is indeed as simple as we found. Identifying price volatility to be the most important parameter of the model with regard to the fractions of strategies used is in agreement with Lux & Marchesi (1999) who also point out volatility, i.e., clustering and persistence of volatility, to be a central feature of the model. In this sense, we give additional support for this claim. Following up on this insight, future research might even be able to analytically derive a closed form solution between the price volatility and the fraction of chartists in the L&M model.

Second, the selection of the best neural network is based on practices that despite being based on previous research remain heuristics that may or may not converge to the best solution. Thus, further research is needed to independently run and replicate these analyses to gain a better understanding of the robustness of these results.

5.2.5 More than two strategies and other variables of interest

The final neural network that we employed in this study is simplistic in terms of (1) the number of input (macro) variables, and (2) the number of output, i.e., target variables to which the input parameters are being mapped to. We use a single input variable, price volatility, to map to a single output variable, the fraction of chartists. Neural networks can also be used to map to more than a single output variable. Such analyses would exploit the capabilities of neural networks to a larger extent. We believe that the simple example discussed in this paper is able to trigger interest in expanding the complexity.

In the context of financial markets, for instance, potential avenues of future research are to estimate a more diverse set of strategies or to further increase the frequency of the estimation data, e.g., using intra-daily data. Finally, other types of strategies in other types of markets, e.g., the fraction of carry traders in foreign exchange markets could be estimated.

6 CONCLUSION

We have proposed a general method for estimating micro-parameter time series of complex agent-based models using an inverse model that is being estimated based on a model-free estimation method. We utilize a neural network approach to estimate the inverse model based on data generated by a simulation of an agent-based model. The inverse model maps macro behavior back to aggregated individual micro-level behavior, which in our study are fractions of strategies used by market participants. The main advantage of the approach is the model-free non-parametric regression of the neural network which poses only little assumptions on the mapping. By applying the neural network to real world data, we are able to indirectly estimate micro-parameter time series parameters of the underlying agent-based model of high complexity at high frequency without having to apply any simplifying assumptions concerning the model or the estimation process.

We have applied the proposed estimation method to the L&M model of a financial market to estimate fractions of strategies used by market participants in the S&P 500 at daily frequency. We empirically validated our estimated daily time series of the fraction of chartists by comparing it with previous results by Boswijk et al. (2007) and by interpreting the results based on past distinct events such as booms and crashes at financial markets. We find that the fraction of chartists is large at times of crises, crashes, and price bubbles. The fraction is low at times of sideways markets with no distinct price trends and low volatility.

Future research might exploit the estimated information about fractions of strategies employed in the markets in *predicting prices* and giving *indications for market inefficiencies*.

For instance, as trend following chartists might induce and sustain trends in prices, at these times price time series should be persistent, i.e., exhibit self similarity, which can be estimated for instance by the Hurst parameter (Rose, 1996; Clegg, 2006). Thus, in case of an increasing fraction of chartists, (1) a continuation of price trends can be expected, and (2) predictability of prices should increase. This is due to the self-enforcing, positive feedback process induced by trend followers. During these times, trend following becomes a self-fulfilling prophecy and it is rational to adopt a trend following strategy in trying to earn excess returns (DeLong et al. 1990), (Shleifer & Summers 1990). Investor George Soros (1987) for example has apparently successfully applied this strategy in real markets by betting on future crowd behavior. Thus, as the estimated fraction of chartists starts to increase rapidly from a low level, one could utilize a trend following strategy in anticipation of future trend following. Thus, a proof of concept trading system, by earning above buy-and-hold or random strategies returns, could provide indications for market inefficiencies on empirical market data.

More generally, we believe that our method of indirect estimation of an agent-based model could be used in other domains as well to estimate micro-level parameters, provided that an appropriate, realistic agent-based model is available or can be constructed. The combination of a model-free estimation method with simulations of complex agent-based models provides a method that comes closer to what one would call an "estimation of an agent-based model" than any previous suggestions. If researchers follow up on this path, the robustness of the approach could be validated on a broader scale.

7 REFERENCES

Alfarano, S., T. Lux and F. Wagner (2006): "Estimation of a Simple Agent-Based Model of Financial Markets: An Application to Australian Stock and Foreign Exchange Data," *Physica A Statistical Mechanics and its Applications*, 370 (Oct.), 38-42.

Alfarano, S., F. Wagner and T. Lux (2005): "Estimation of Agent-Based Models: The Case of an Asymmetric Herding Model," *Computational Economics*, 26, 19-49.

Amilon, H. (2006): "Estimation of an Adaptive Stock Market Model with Heterogeneous Agents," *Riksbank Research Paper* (16), 19 October, Sveriges Riksbank (Central Bank of Sweden).

Aoki, M. (2002): "Open Models of Share Markets with Two Dominant Types of Participants," *Journal of Economic Behavior and Organization*, 49, 199-216.

Bishop, C. M. (1995): *Neural Networks for Pattern Recognition*, Oxford: Oxford University Press.

Black, F. (1986): "Noise," *Journal of Finance*, 41(3), 529-543.

Boswijk, H. P., C. H. Hommes and S. Manzan (2007): "Behavioral Heterogeneity in Stock Prices," *Journal of Economic Dynamics and Control*, 31(6), June, 1938-1970.

Brock, W. A. and C. H. Hommes (1998): "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model," *Journal of Economic Dynamics and Control*, 22(8), 1235-1274.

Carvalho, R. (2001): "The Dynamics of the Linear Random Farmer Model," Lissabon: Instituto Superior Técnico.

Chen, S.-H., T. Lux and M. Marchesi (2000): "Testing for Non-Linear Structure in an Artificial Stock Market," *Journal of Economic Behavior and Organization*, 46 (2001), 327-342.

Clegg, R. (2006): "A Practical Guide to Measuring the Hurst Parameter," *International Journal of Simulation: Systems, Science & Technology*, 7(1), 3-14.

Cont, R. (2001): "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues," *Quantitative Finance*, 1, 223-236.

Cont, R. (2005): "Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models," in *Long memory in economics*, ed. by A. &. T. Kirman, G., Springer.

De Grauwe, P. and M. Grimaldi (2006): "Exchange Rate Puzzles: A Tale of Switching Attractors," *European Economic Review*, 50(1), January, 1-33.

DeLong, J. B., A. Shleifer, L. H. Summers and R. J. Waldmann (1990): "Positive Feedback

Investment Strategies and Destabilizing Rational Speculation," *The Journal of Finance*, 45(2), 379-395.

Fagiolo, G., P. Windrum and A. Moneta (2006): "Empirical Validation of Agent-based Models: A Critical Survey," in *LEM paper series*, 2006/14, Pisa, Italy: Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, May.

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

Farmer, J. D. and S. Joshi (2000): "The Price Dynamics of Common Trading Strategies," *Journal of Economic Behavior and Organization* (30/10).

Foresee, F. D. and M. T. Hagan (1997): "Gauss-Newton Approximation to Bayesian Regularization.", International Joint Conference on Neural Networks, vol. 3, Piscataway, NJ: IEEE Press, 1930-1935.

Frankel, J. A. and K. A. Froot (1988): "Chartists, Fundamentalists and the Demand for Dollars," *Greek Economic Review*, 10(1), June, 49-102.

Gallant, R. and G. Tauchen (1996): "Which Moments to Match?" *Economic Theory*, 12, 657-681.

Gilli, M. and P. Winker (2003): "A Global Optimization Heuristic for Estimating Agent-based Models," *Computational Statistics & Data Analysis*, 42(3), 299-312.

Gopikrishnan, P., M. Meyer, L. Amaral and H. Stanley (1998): "Inverse Cubic Law for the Distribution of Stock Price Variations," *European Physical Journal B*, 3(2), 139-140.

Hagan, M. and M. Menhaj (1994): "Training Feedforward Networks with the Marquardt Algorithm," *IEEE Transactions on Neural Networks*, 5(6), Nov, 989-993.

Haykin, S. (1999): *Neural Networks: A Comprehensive Foundation, 2nd ed.*, Upper Saddle River, NJ: Prentice Hall.

Hodrick and Prescott (1997): "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit and Banking*.

Hommes, C. H. (2002): "Modeling the Stylized Facts in Finance Through Simple Nonlinear Adaptive Systems," *Proceedings of the National Academy of Sciences of the United States of America*, 99 (05/14), 7221-7228, Washington.

Hommes, C. H. (2006): "Heterogeneous Agent Models in Economics and Finance," in *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, ed. by L. Tesfatsion and K. Judd, Elsevier Science, 1109-1186.

Kalman, B. L. and S. C. Kwasny (1992): "Why Tanh? Choosing a Sigmoidal Function.", International Joint Conference on Neural Networks, Baltimore, MD.

Kaufman, P. J. (1998): *Trading Systems and Methods, 3rd Edition*, John Wiley and Sons.

Keim, D. B. and A. Madhavan (1995): "Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Traders," *Journal of Financial Economics*, 37(3), 371-398.

Kharas, H., B. Pinto and S. Ulatov (2001): "An Analysis of Russia's 1998 Meltdown: Fundamentals and Market Signals," *Brookings Papers on Economic Activity* (1 (2001)), 1-67.

Kirman, A. (1992): "Whom or What Does the Representative Individual Represent?" *Journal of Economic Perspectives*, 6(2), Spring, 117-136.

Kirman, A. (1993): "Ants, Rationality, and Recruitment," *Quarterly Journal of Economics*, 108, 137-156.

LeBaron, B. (2006): "Agent-Based Computational Finance," in *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, ed. by L. Tesfatsion and K. l. Judd, Elsevier Science, 1187-1233.

Lux, T. and M. Marchesi (1999): "Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market," *Nature*, 397 (11/02), 498-500.

Lux, T. and M. Marchesi (2000): "Volatility Clustering in Financial Markets: A Micro-Simulation of Interacting Agents," *International Journal of Theoretical and Applied Finance*, 3, 675-702.

MacKay, D. J. C. (1992): "Bayesian Interpolation," *Neural Computation*, 4(3), May, 415-447.

Masters, T. (1993): *Practical Neural Network Recipes in C++*, Academic Press.

Menkhoff, L. (1998): "The Noise Trading Approach - Questionnaire Evidence from Foreign Exchange," *Journal of International Money and Finance*, 17(3), 547-564.

Pagan, A. (1996): "The Econometrics of Financial Markets," *Journal of Empirical Finance*.

Radelet, S. and J. Sachs (1998): "The East Asian Financial Crisis: Diagnosis, Remedies, Prospects," *Brookings Paper on Economic Activity*, 28(1 (1998)), 1-74.

Reed, R. and R. Marks, II (1999): *Neural Smithing: Supervised Learning in Feedforward Artificial Neural Networks*, Cambridge, MA: The MIT Press.

Roche, A., G. Malandain, X. Pennec and N. Ayache (1998): "Multimodal Image Registration by Maximization of the Correlation Ratio," 3378, Sophia Antipolis: Unité de recherche INRIA Sophia Antipolis.

Rose, O. (1996): "Estimation of the Hurst Parameter of Long-Range Dependent Time Series," in *Research Report Series*, University of Würzburg, Institute of Computer

Science, 137, Würzburg.

Sarle, W. S., saswss@unx.sas.com (2002): "Neural Network FAQ," ftp://ftp.sas.com/pub/neural/FAQ.html (2009-11-20).

Schwert, G. W. (1990): "Stock Volatility and the Crash of '87," *Review of Financial Studies*, 3(1), 77-102.

Shiller, R. J. (1984): "Stock Prices and Social Dynamics." *Brookings Papers on Economic Activity*, 2, 457-510.

Shleifer, A. and L. H. Summers (1990): "The Noise Trader Approach to Finance," *Journal of Economic Perspectives*, 4(2), 19-33.

Soros, G. (1987): *The Alchemy of Finance*, New York: Simon and Schuster.

Spanos, A. (2006): "The Curve-Fitting Problem, Akaike-Type Model Selection, and the Error Statistical Approach," in *Working Paper*, Virginia Tech, 17 October.

Swingler, K. (1996): *Applying Neural Networks: A Practical Guide*, London: Academic Press.

Taskaya-Temizel, T. and K. Ahmad (2005): "Are ARIMA Neural Network Hybrids Better Than Single Models?", International Joint Conference on Neural Networks, Montreal, Canada.

Triola, M. F. (2003): *Elementary Statistics, 9th Edition*, Addison Wesley.

Vigfusson, R. (1997): "Switching Between Chartists and Fundamentalists: A Markov Regime-Switching Approach," *International Journal of Finance & Economics*, 2(4), October, 291-305.

Westerhoff, F. H. (2005): "Technical Analysis Based on Price-Volume Signals and the Power of Trading Breaks."

Westerhoff, F. and S. Reitz (2003): "Nonlinearities and Cyclical Behavior: The Role of Chartists and Fundamentalists," *Studies in Nonlinear Dynamics & Econometrics*, 7(4).

Western, D. (2004): *Booms, Bubbles and Busts in the US Stock Market*, New York: Routledge.

Table 1: Parameters (and their symbols) of our indirect neural network-based estimation approach and the allocation of the parameters in our application of the approach to estimating micro parameters of the Lux and Marchesi (1999, 2000) financial market model.

Parameter	Symbol	Allocation
Agent based model	*A*	Lux-marchesi financial market model (Lux and Marchesi, 1999, 2000)
Simulation time span	*ts*	*ts* = 9950 days
Set of simulation parameters	*SP*	According to Lux and Marchesi (1999, 2000), see also section 4.1.
Micro time series	$mits_a$	Not applicable, as Lux and Marchesi, (1999, 2000) only model groups of agents, not individuals.
Macro time series	$mats_b$	Price *p*, and fundamental value p_f (see Lux and Marchesi, 1999, 2000)
Aggregated micro time series	$amits_b$	Daily fraction of chartists *chart(t),* daily fractions of pessimistic chartists *pess(t)* and optimistic chartists *opt(t),* and daily fraction of fundamentalsists *fund(t)*
Desired (aggregated) micro level output parameters do_j	$do_j \in DO$	Daily fraction of chartists *chart(t)*
Potential macro level input parameters	$i_k \in PI$	price *p(t)*, historical price intervals of various lengths τ: $p(t)-p(t-\tau)$, price returns $r(t)=ln(p(t)) - ln(p(t-1))$, price volatility as the square of the price returns: $r(t)^2$, relative price trends $rpt(\tau) = \frac{p(t)-p(t-\tau)}{\tau \cdot p(t)}$ of length τ and differences in relative price trends of different lengths, e.g., $rpt(\tau_1)-rpt(\tau_2)$.
Smoothing method for pre-processing of input parameter time series $i_k(t)$	*SM*	Hodrick and Prescott filter (Hodrick and Prescott 1997)
Set of neural network building parameters	*NBP*	Number of hidden layers = 1; Number of hidden neurons = 5; activation function =

		hyperbolic tangent; training algorithm = Gauss-newton approximation to bayesian learning (Foresee and Hagan, 1997); Objective function $F = \alpha E_D + \beta E_W$, where E_W is the sum of squares of the network weights and E_D is the sum of squared errors (Foresee and Hagan, 1997); Marquardt adjustment parameter μ = 0.005; decrease factor for μ as μ_{dec} = 0.1; the increase factor for μ as μ_{inc} = 10
Length of the time series in the training set (l_1), validation set (l_2), test set (l_3), and robustness test set (l_4)	l_1, l_2, l_3, l_4	l_1=5000; l_2=500; l_3=946; and l_4=946
Set *I* of (actual) macro level input parameters *i*	$i_k \in I \subseteq PI$	Price volatility $r(t)^2$
Actual micro level output parameter time series	$o_m(t)$	Estimated daily fraction of chartists $\text{chart}_{\text{est}}(t)$
Set of stopping criteria for the training algorithm	*SC*	Maximum number of training iterations to *maxiter* = 500; performance goal for the objective function on the training set *goal* = 0.001; maximum value for μ is $\mu_{max} = 10^{10}$; maximum number of validation failures *maxvf* = 5; minimum performance gradient *mingrad* = 10^{-10}
Fitness criterion on the *test set*: pearson correlation coefficient with respect to actual $o_m(t)$ vs. desired output $do_j(t)$ time series.	*pc*	*pc*>0.9

Table 2: Statistical properties of real markets, published properties of the L&M market model, and properties of our re-implementation of the L&M model.

Statistical property	Real markets	L&M model	Our results
Self-similarity Hurst parameter of raw price returns	0.5 [1)]	0.48	0.42
Hurst parameter of absolute price returns	>0.6 [1)]	0.85	0.89
Hurst parameter of raw returns of fundamental value	0.5 [2)]	0.51	0.52
Hurst parameter of absolute returns of fundamental value	0.5 [2)]	0.49	0.5
Exponent α for the tail of the unconditional pareto distribution of absolute price returns (covering 30% of the largest observations of the time series)	2-5 [3)] ~3 [4)]	2.64	2.42
Exponent α for the tail of the unconditional pareto distribution of absolute price returns (covering 10% of the largest observations of the time series)	2-5 [3)] ~3 [4)]	- [5)]	2.83
Kurtosis of the distribution of price returns	>3 [6)]	- [5)]	9.65

Notes.
1) Westerhoff (2005:10)
2) Fundamental value is assumed to be randomly distributed, thus exhibiting no self-similarity, thus H=0.5.
3) Cont (2001:224)
4) Gopikrishnan et al. (1998:139)
5) Not published by L&M to our knowledge.
6) According to Cont (2001) kurtosis is positive and excessive, i.e., >3, as the normalized kurtosis of the normal distribution is 3. E.g., they determine kurtosis for 5 minute price increments for S&P 500 futures to be 18.95.

Figure 1: The proposed three-step indirect model-free estimation approach for micro parameters of complex heterogeneous agent-based models, illustrated by the case of an agent-based financial market model.

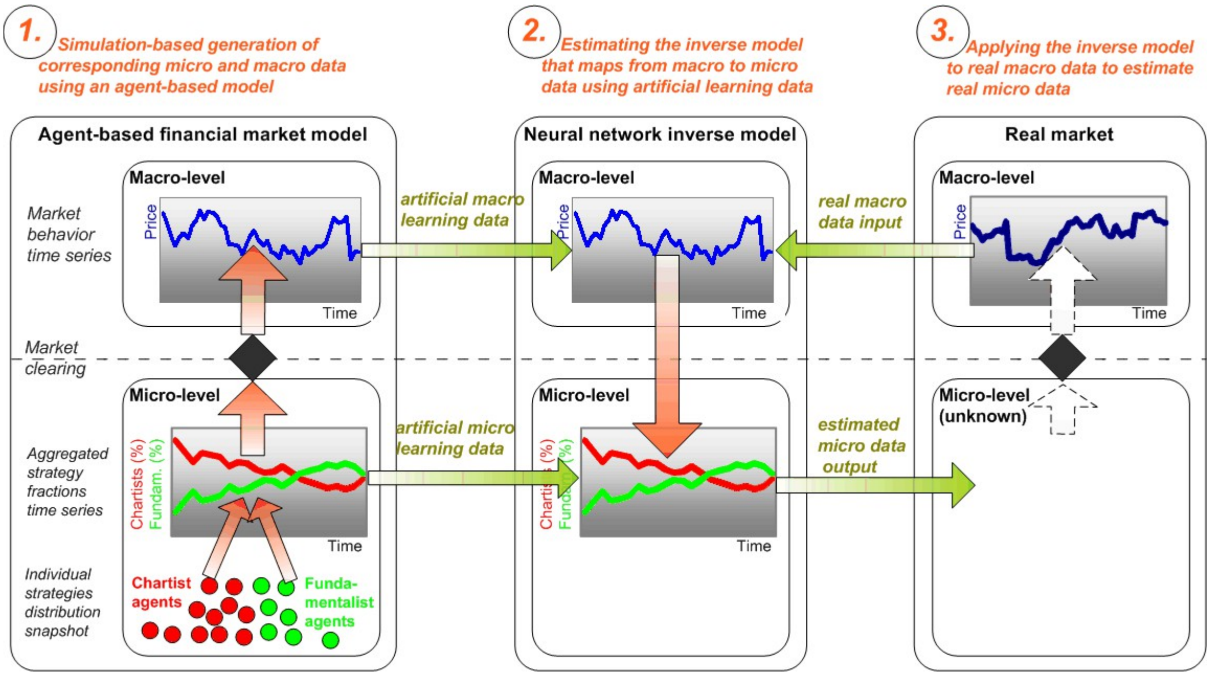

Figure 2: Time series (9950 days) of price vs. total fraction of chartists (optimists + pessimists) in percent.

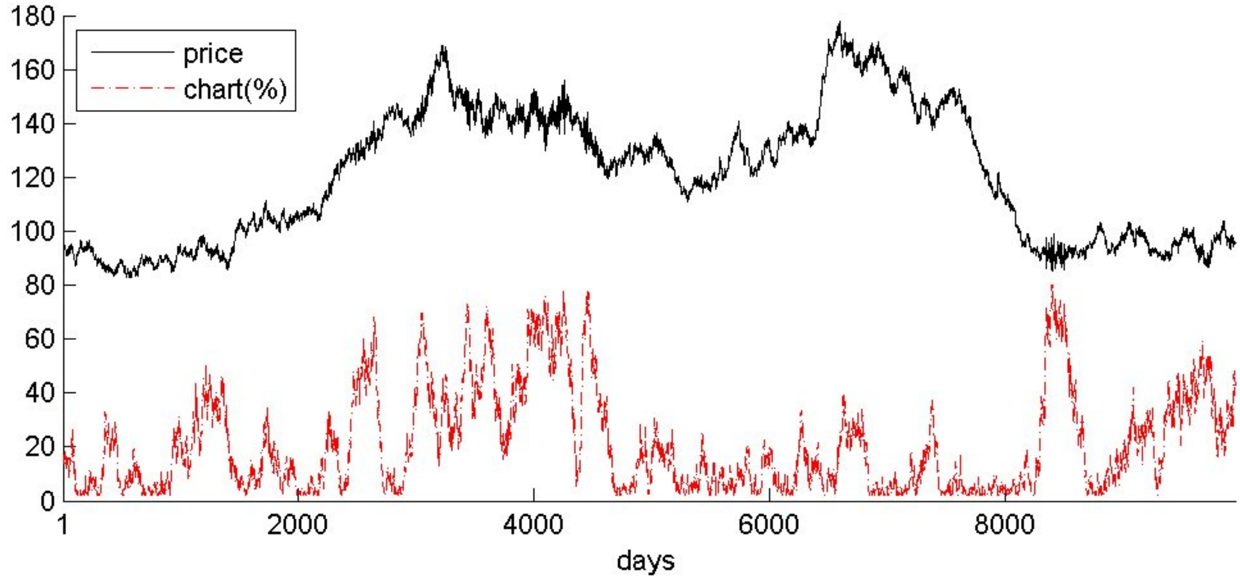

Figure 3: Autocorrelation function of (raw, absolute, and squared) price log returns.

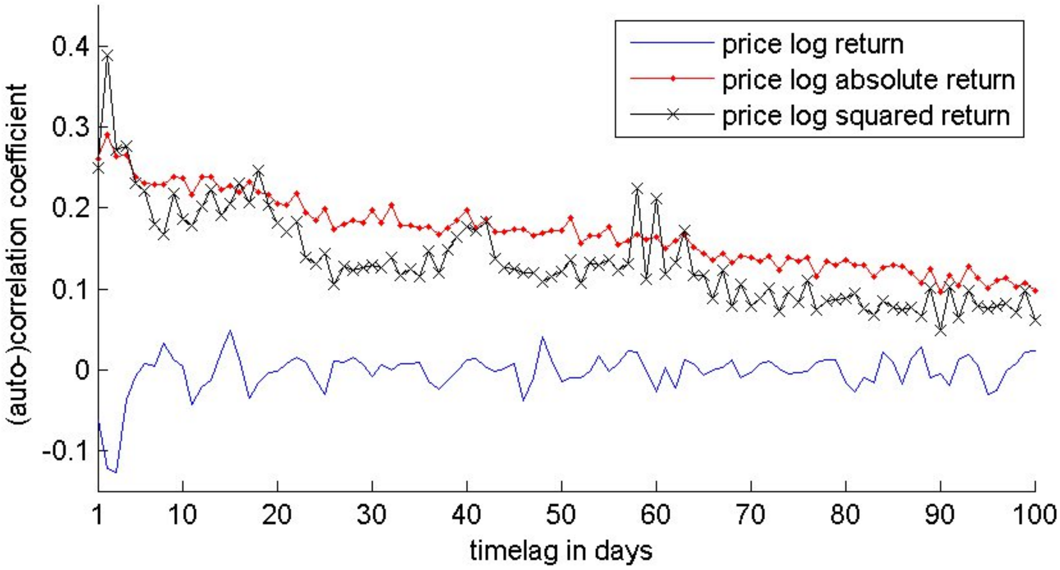

Figure 4: Time series of the actual (network output) vs. desired (target) fraction of chartists on the test data set, determined by the best neural network, trained on HP filter smoothed input data.

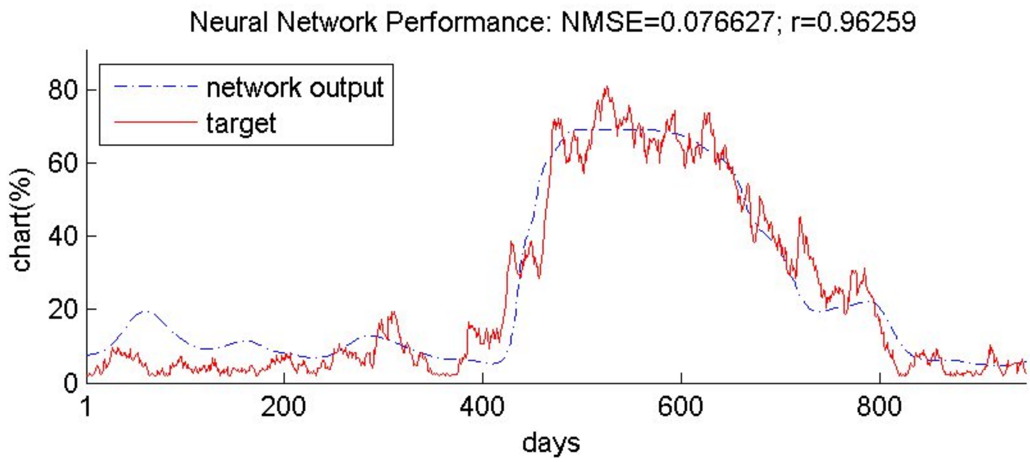

Figure 5: Verification of robustness: Time series of actual (network output) vs. desired (target) fraction of chartists, obtained on a second separate test data set.

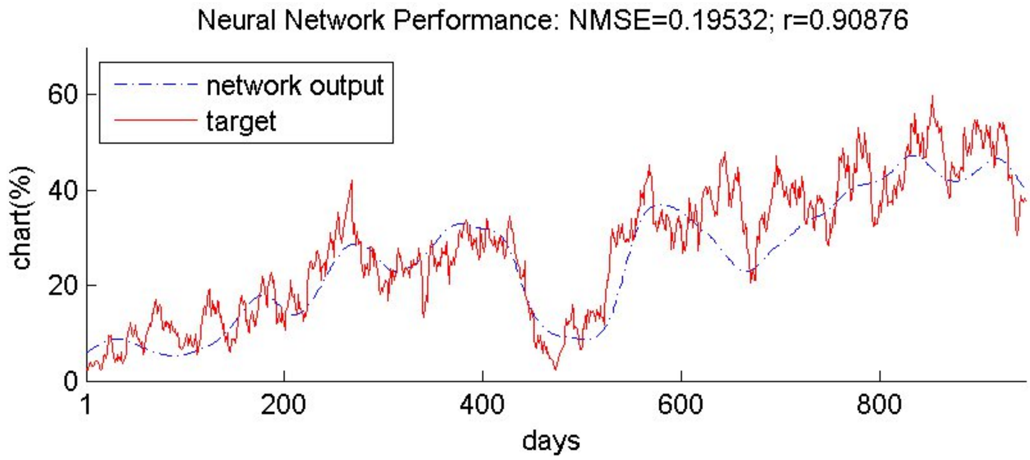

Figure 6: S&P 500 vs. estimated daily fractions of chartists.

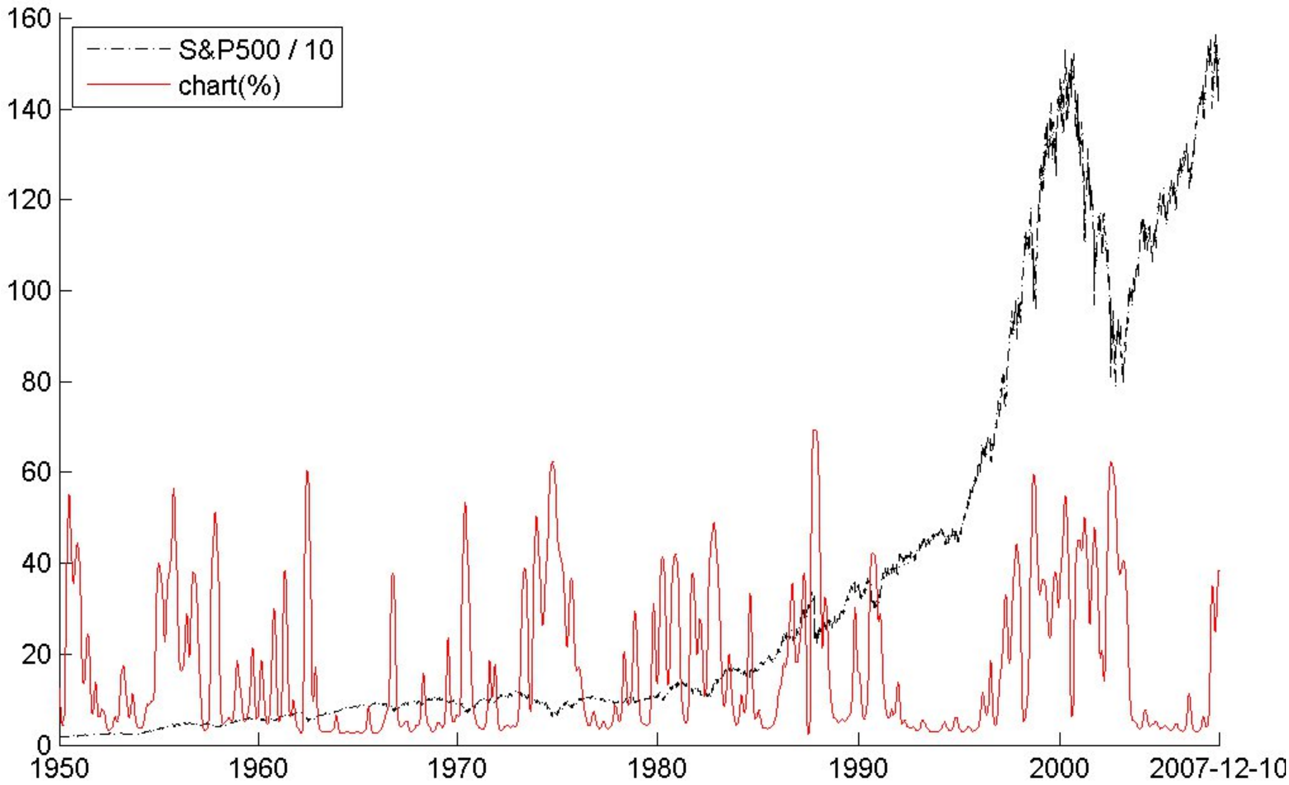

Figure 7: S&P 500 vs. estimated fraction of chartists.

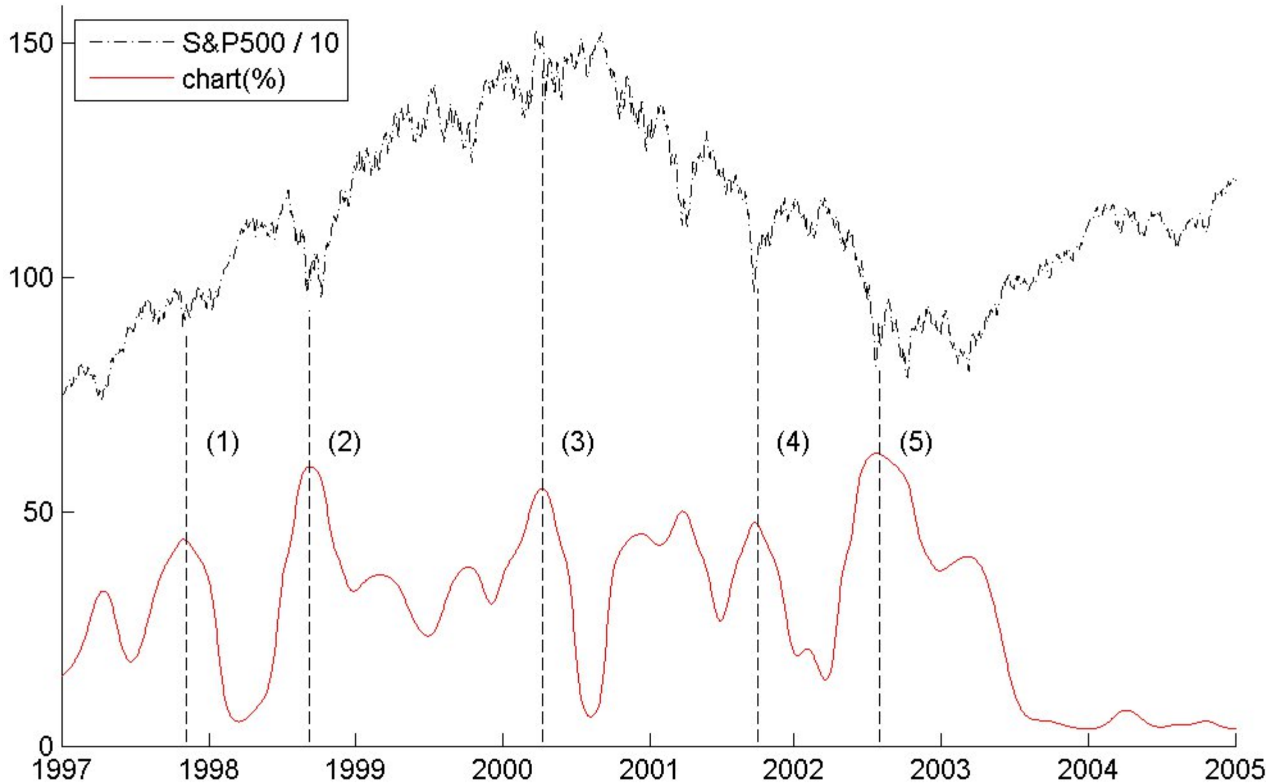

Figure 8: S&P 500 (daily time series) vs. fraction of chartists: our curve vs. Boswijk et al's estimation.

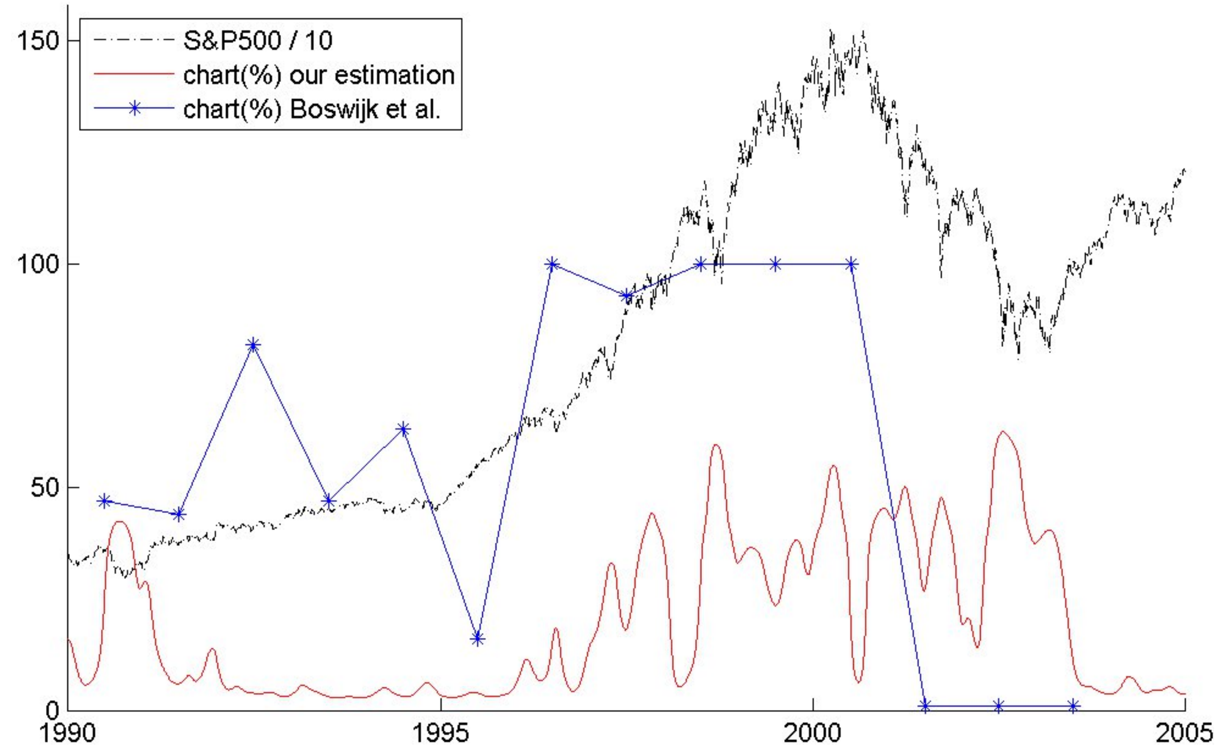

Figure 9: Overview of all four sections comparing neural network chartist fraction output based on (1) locally smoothed input time series parameters vs. (2) globally smoothed input tine series parameters (dashed line).

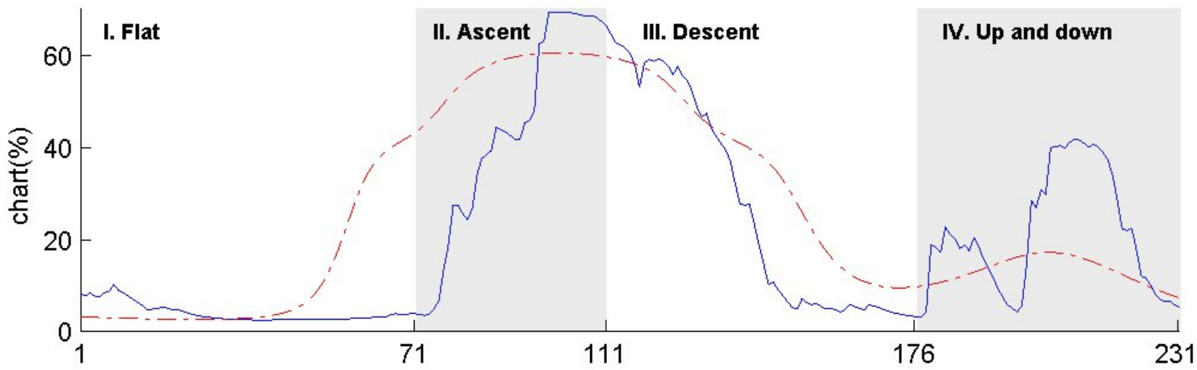

Figure 10: Section I (Flat): network output based on (1) locally smoothed input vs. (2) globally smoothed input (dashed line)

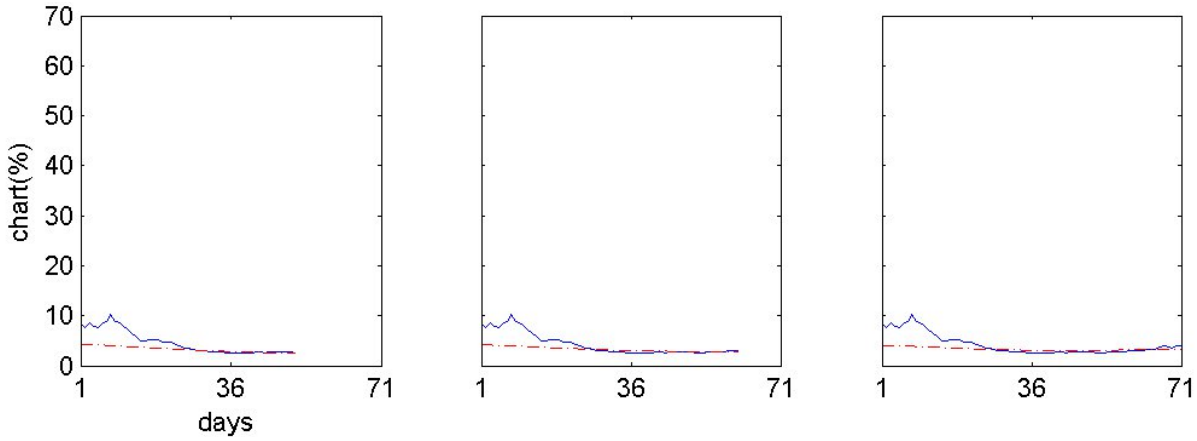

Figure 11: Section II (Ascent): network output based on (1) locally smoothed input vs. (2) globally smoothed input (dashed line)

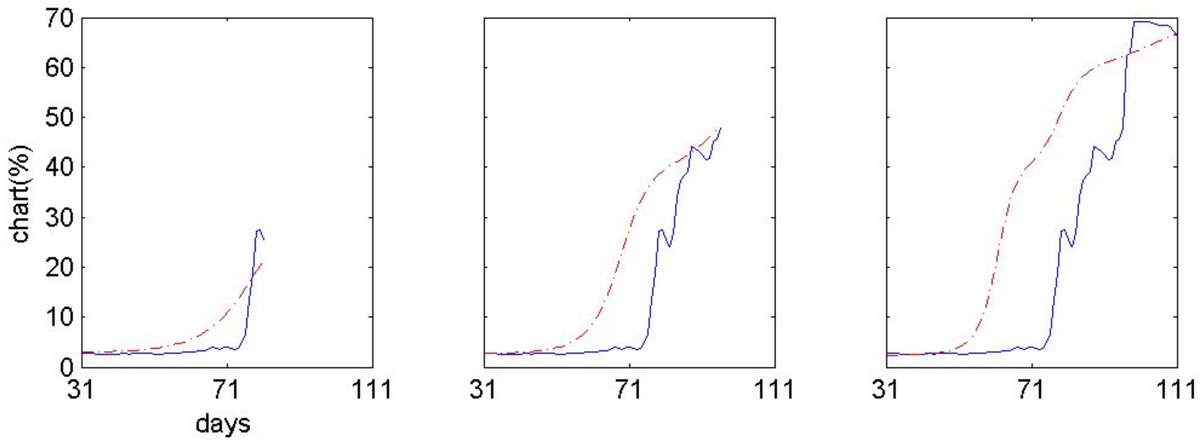

Figure 12: Section III (Descent): network output based on (1) locally smoothed input vs. (2) globally smoothed input (dashed line)

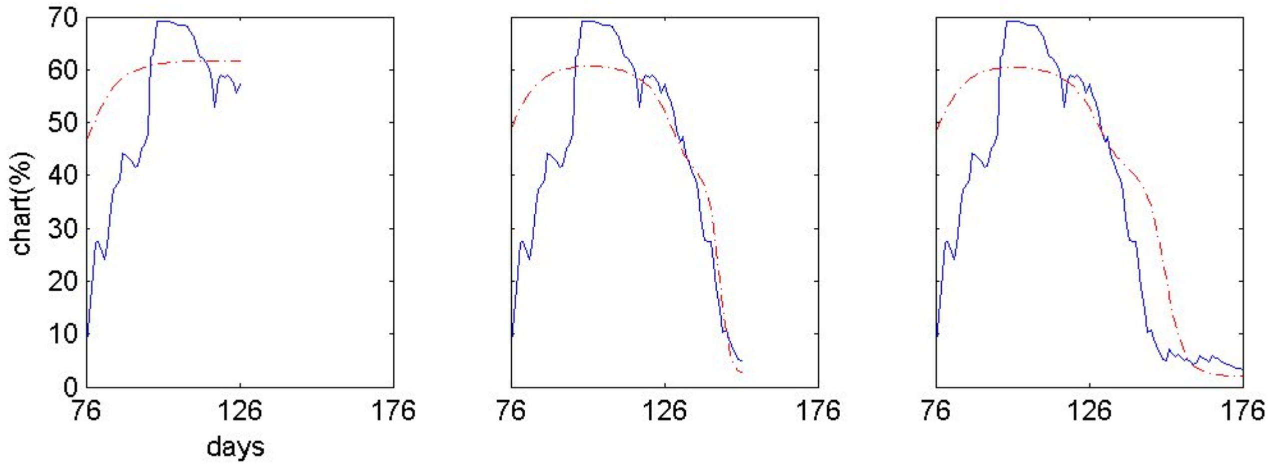

Figure 13: Section IV (Up and down): network output based on (1) locally smoothed input vs. (2) globally smoothed input (dashed line)

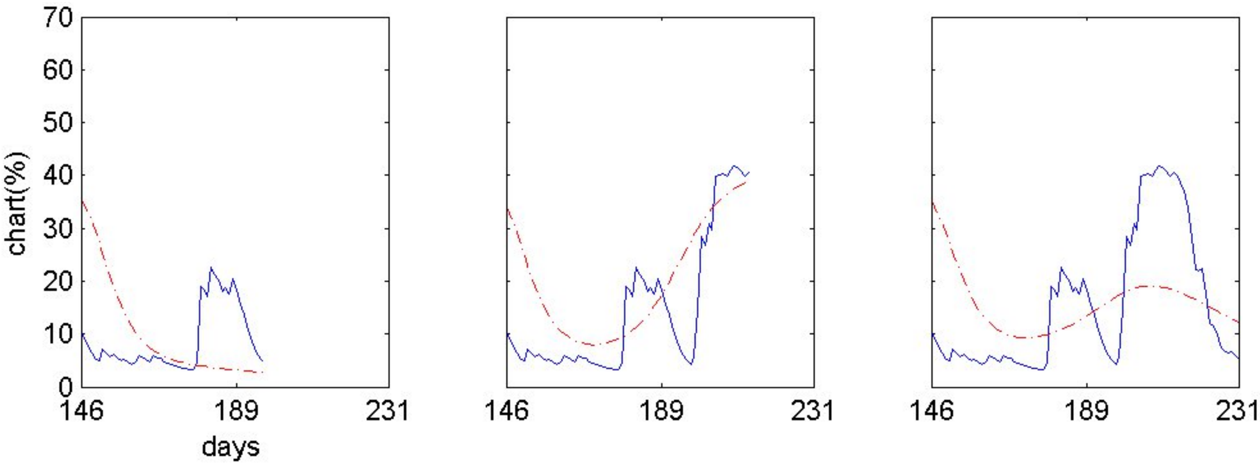

Figure 14: Dependency of the chartist fraction of the best neural network trained on smoothed input data from the sole input parameter, price volatility r(t)²

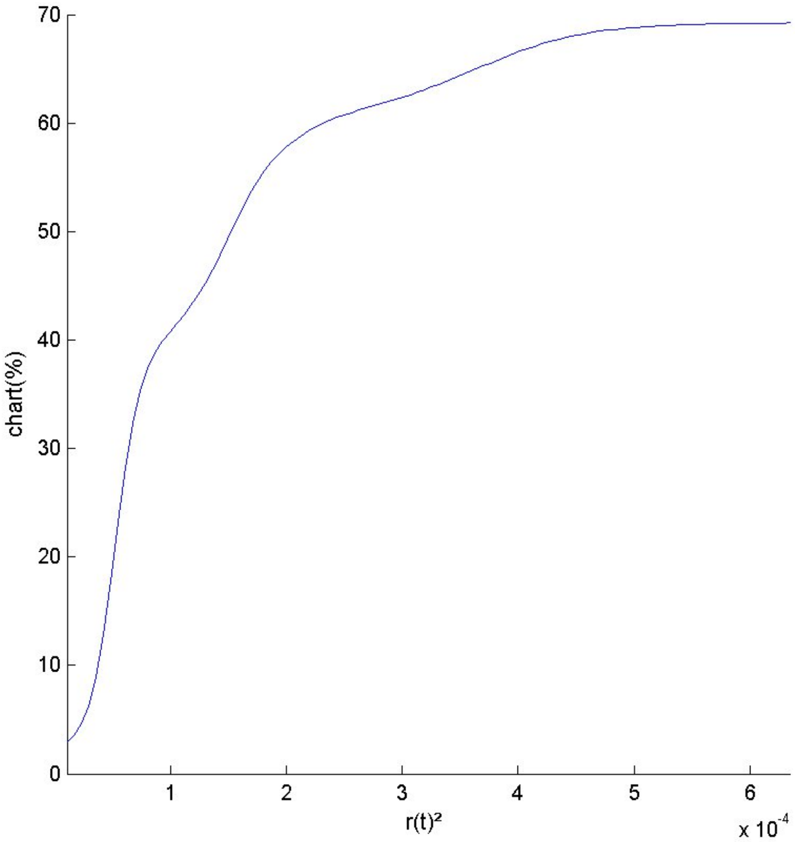

Figure 15: Dependency of the chartist fraction on daily price volatility for the S&P 500