A simple agent-based financial market model: direct interactions and comparisons of trading profits

Frank H. Westerhoff

University of Bamberg

Department of Economics

Feldkirchenstrasse 21

D-96045 Bamberg

Germany

frank.westerhoff@uni-bamberg.de

Abstract: We develop an agent-based financial market model in which agents follow technical and fundamental trading rules to determine their speculative investment positions. A central feature of our model is that we consider direct interactions between speculators due to which they may decide to change their trading behavior. For instance, if a technical trader meets a fundamental trader and they realize that fundamental trading has been more profitable than technical trading in the recent past, the probability that the technical trader switches to fundamental trading rules is relatively high. Our simple setup is able to replicate some salient features of asset price dynamics.

Keywords: Agent-based financial market models; direct interactions; evolutionary fitness measures; technical and fundamental analysis; stylized facts of financial markets.

JEL classification: G12; G14; G15.

1 Introduction

In the recent past, a number of interesting agent-based financial market models have been proposed. These models successfully explain some important stylized facts of financial markets, such as bubbles and crashes, fat tails for the distribution of returns and volatility clustering. These models, reviewed, for instance, in Hommes (2006), LeBaron (2006), Chen et al. (2009), Lux (2009a) and Westerhoff (2009), are based on the observation that financial market participants use different heuristic trading rules to determine their speculative investment positions. Note that survey studies by Frankel and Froot (1986), Taylor and Allen (1992), Menkhoff (1997), and Menkhoff and Taylor (2007) in fact reveal that market participants use technical and fundamental analysis to assess financial markets. Agent-based financial market models thus have a strong empirical foundation.

As is well known, technical analysis is a trading philosophy built on the assumption that prices tend to move in trends (Murphy 1999). By extrapolating price trends, technical trading rules usually add a positive feedback to the dynamics of financial markets, and thus may be destabilizing. Fundamental analysis is grounded on the belief that asset prices return to their fundamental values in the long run (Graham and Dodd 1951). Buying undervalued and selling overvalued assets, as suggested by these rules, apparently has a stabilizing impact on market dynamics. In most agent-based financial market models, the relative importance of these trading strategies varies over time. It is not difficult to imagine that changes in the composition of applied trading rules – such as a major shift from fundamental to technical trading rules – may have a marked impact on the dynamics of financial markets.

One goal of our paper is to provide a novel view on how financial market participants may select their trading rules. We do this by recombining a number of building blocks from three prominent agent-based financial market models. Let us briefly recapitulate these models:

- Brock and Hommes (1997, 1998) developed a framework in which (a continuum of) financial market participants endogenously chooses between different trading rules. The agents are boundedly rational in the sense that they tend to pick trading rules which have performed well in the recent past, thereby displaying some kind of learning behavior. The performance of the trading rules may be measured as a weighted average of past realized profits, and the relative importance of the trading rules is derived via a discrete choice model. Contributions developed in this manner are often analytically tractable. Moreover, numerical investigations reveal that complex endogenous dynamics may emerge due to an ongoing evolutionary competition between trading rules. Note that in such a setting, agents interact only indirectly with each other: their orders have an impact on the price formation which, in turn, affects the performance of the trading rules and thus the agents' selection of rules. Put differently, an agent is not directly affected by the actions of others.
- In Kirman (1991, 1993), an influential opinion formation model with interactions between a fixed number of agents was introduced. In Kirman's model, agents may hold one of two views. In each time step, two agents may meet at random, and there is a fixed probability that one agent may convince the other agent to follow his opinion. In addition, there is also a small probability that an agent changes his opinion independently. A key finding of this model is that direct interactions between heterogeneous agents may lead to substantial opinion swings. Applied to a financial

market setting, one may therefore observe periods where either destabilizing technical traders or stabilizing fundamental traders drive the market dynamics. Note that agents may change rules due to direct interactions with other agents but the switching probabilities are independent of the performance of the rules.

- The models of Lux (1995, 1998) and Lux and Marchesi (1999, 2000) also focus on the case of a limited number of agents. Within this approach, an agent may either be an optimistic or a pessimistic technical trader or a fundamental trader. The probability that agents switch from having an optimistic technical attitude to a pessimistic one (and vice versa) depends on the majority opinion among the technical traders and the current price trend. For instance, if the majority of technical traders are optimistic and if prices are going up, the probability that pessimistic technical traders turn into optimistic technical traders is relatively high. The probability that technical traders (either being optimistic or pessimistic) switch to fundamental trading (and vice versa) depends on the relative profitability of the rules. However, a comparison of the performance of the trading rules is modeled in an asymmetric manner. While the attractiveness of technical analysis depends on realized profits, the popularity of fundamental analysis is given by expected future profit opportunities. This class of models is quite good at replicating several universal features of asset price dynamics.

Each of these approaches has been extended in various interesting directions. There are also alternative strands of research in which the dynamics of financial markets is driven, for instance, by nonlinear trading rules or wealth effects. For related models see, for instance, Day and Huang (1990), Chiarella (1992), de Grauwe et al. (1993), Li and Rosser (2001), Chiarella et al. (2002), Farmer and Joshi (2002), Li and Rosser (2004), Rosser et al. (2003), de Grauwe and Grimaldi (2006), Westerhoff and Dieci (2006) or

Westerhoff (2008), among many others.

In this paper, we seek to recombine key ingredients of the three aforementioned approaches to come up with a simple model that is able to match the stylized facts of financial markets and that offers a novel perspective on how agents may be influenced in selecting their trading rules. In our model, we consider direct interactions between a fixed number of agents, as in Kirman' approach. However, the switching probabilities are not constant over time but depend on the recent performance of the rules. To avoid asymmetric profit measures, as in the models of Lux and Marchesi, we define a fitness function along the lines of the models of Brock and Hommes, i.e. we approximate the fitness (attractiveness) of a rule by a weighted average of current and past myopic profits. Replication of the dynamics of agent-based models is often a challenging undertaking, which is why these models are sometimes regarded with skepticism. A second goal of our paper is thus to come up with a setting for which replication of our results is rather uncomplicated, even, as we hope, for the (interested) layman.

Our paper is organized as follows. In section 2, we present our approach. In section 3, we show that our model may mimic some stylized facts of financial markets. We also explore how a change in the number of agents and in the frequency of their interactions affects the dynamics. The last section offers some conclusions.

2 A basic model

Let us first preview the structure of our model. We assume that prices adjust with respect to the current excess demand. The excess demand, in turn, depends on the orders submitted by technical and fundamental traders. While technical traders base their orders on a trend-extrapolation of past prices, fundamental traders place their bets on mean reversion. The relative impact of these two trader types evolves over time. We assume that agents regularly meet each other and talk about their past trading performance. As a result, traders may change their opinion and switch to a new trading strategy. In particular, the time-varying switching probabilities depend on the relative success of the rules. Numerical simulations will reveal that the fractions of technical and fundamental trading rules evolve over time, which is exactly what gives rise to interesting asset price dynamics. Now we are ready to turn to the details of the model.

As in Farmer and Joshi (2002), the price adjustment is due to a simple log-linear price impact function. Such a function describes the relation between the quantity of an asset bought or sold in a given time interval and the price change caused by these orders. Accordingly, the log of the price of the asset in period t+1 is quoted as

$$P_{t+1} = P_t + a(W_t^C D_t^C + W_t^F D_t^F) + \alpha_t , \tag{1}$$

where a is a positive price adjustment coefficient, D^C and D^F stand for orders generated by technical and fundamental trading rules, and W^C and W^F denote the fractions of agents using these rules. Excess buying (selling) thus drives prices up (down). Since our model only provides a simple representation of real financial markets, we add a random term to (1). We assume that α is an IID normal random variable with mean zero and constant standard deviation σ^α .

The goal of technical analysis is to exploit price trends (see Murphy 1999 for a practical introduction). Since technical analysis typically suggests buying the asset when prices increase, orders triggered by technical trading rules may be written as

$$D_t^C = b(P_t - P_{t-1}) + \beta_t . \tag{2}$$

The first term of the right-hand side of (2) stands for transactions triggered by an extrapolation of the current price trend. The reaction parameter b is positive and captures how strongly the agents react to this price signal. The second term reflects additional random orders to account for the large variety of technical trading rules. As in (1) we assume that shocks are normally distributed, i.e. β is an IID normal random variable with mean zero and constant standard deviation σ^{β}.

Fundamental analysis (see Graham and Dodd 1951 for a classical contribution) presumes that prices may disconnect from fundamental values in the short run. In the long run, however, prices are expected to converge towards their fundamental values. Since fundamental analysis suggests buying (selling) the asset when the price is below (above) its fundamental value, orders generated by fundamental trading rules may be formalized as

$$D_t^F = c(F_t - P_t) + \gamma_t \,, \tag{3}$$

where c is a positive reaction parameter and F is the log of the fundamental value. Note that we assume that traders are able to compute the true fundamental value of the asset. In order to allow for deviations from the strict application of this rule, we include a random variable γ in (3), where γ is IID normally distributed with mean zero and constant standard deviation σ^{γ}.

For simplicity, the fundamental value is set constant, i.e.

$$F_t = 0 \,. \tag{4}$$

Alternatively, the evolution of the fundamental value may be modeled as a random walk. However, in order to show that the dynamics of a financial market may not depend on fundamental shocks, we abstain from this.

We furthermore assume that there are N traders in total. Let K be the number of technical traders. We are then able to define the weight of technical traders as

$$W_t^C = K_t / N \,. \tag{5}$$

Similarly, the weight of fundamental traders is given as

$$W_t^F = (N - K_t) / N \,. \tag{6}$$

Obviously, (5) and (6) imply that $W_t^F = 1 - W_t^C$.

The number of technical and fundamental trades is determined as follows. As in Kirman (1991, 1993), we assume that two traders meet at random in each time step, and that the first trader will adopt the opinion of the other trader with a certain probability $(1-\delta)$. In addition, there is a small probability ε that a trader will change his attitude independently. Contrary to Kirman's approach, however, the probability that a trader converts another trader is asymmetric and depends on the current and past myopic profitability of the rules (indicated by the fitness variables A^C and A^F, which we define in the sequel). Suppose that technical trading rules have generated higher myopic profits than fundamental trading rules in the recent past. Then it is more likely that a technical trader will convince a fundamental trader than vice versa. Similarly, when fundamental trading rules are regarded as more profitable than technical trading rules, the chances are higher that a fundamental trader will successfully challenge a technical trader. Thus, we express the transition probability of K as

$$K_t = \begin{cases} K_{t-1} + 1 & \textit{with probability} \quad p_{t-1}^{+} = \dfrac{N - K_{t-1}}{N} (\varepsilon + (1-\delta)_{t-1}^{F \to C} \dfrac{K_{t-1}}{N-1}) \\ K_{t-1} - 1 & \textit{with probability} \quad p_{t-1}^{-} = \dfrac{K_{t-1}}{N} (\varepsilon + (1-\delta)_{t-1}^{C \to F} \dfrac{N - K_{t-1}}{N-1}) \\ K_{t-1} & \textit{with probability} \quad 1 - p_{t-1}^{+} - p_{t-1}^{-} \end{cases} \,, \tag{7}$$

where the probability that a fundamental trader is converted into an technical trader is

$$(1-\delta)_{t-1}^{F\to C} = \begin{cases} 0.5+\lambda & for \quad A_t^C > A_t^F \\ 0.5-\lambda & otherwise \end{cases} \tag{8}$$

and the probability that a technical trader is converted into a fundamental trader is

$$(1-\delta)_{t-1}^{C\to F} = \begin{cases} 0.5-\lambda & for \quad A_t^C > A_t^F \\ 0.5+\lambda & otherwise \end{cases}, \tag{9}$$

respectively.

Finally, we measure the fitness (attractiveness) of the trading rules as

$$A_t^C = (\exp[P_t] - \exp[P_{t-1}])D_{t-2}^C + dA_{t-1}^C, \tag{10}$$

and

$$A_t^F = (\exp[P_t] - \exp[P_{t-1}])D_{t-2}^F + dA_{t-1}^F, \tag{11}$$

respectively. Formulations (10) and (11) are as in Westerhoff and Dieci (2006) which, in turn, were inspired by Brock and Hommes (1998). Note that the fitness of a trading rule depends on two components. First, the agents take into account the most recent performance of the rules, indicated by the first terms of the right-hand side. The timing we assume is as follows. Orders submitted in period t-2 are executed at the price stated in period t-1. Whether or not these orders produce myopic profits then depends on the realized price in period t. Second, the agents have a memory. The memory parameter $0 \le d \le 1$ measures how quickly current myopic profits are discounted. For $d = 0$, agents obviously have no memory, while for $d = 1$ they compute the fitness of a rule as the sum of all observed myopic profits.

3 Some simulation results

The dynamics of international financial markets display certain stylized facts (Mantegna and Stanley 2000, Cont 2001, Lux and Ausloos 2002). These features include a random walk-like behavior of prices, the sporadic appearance of bubbles and crashes, excess volatility, fat tails of the distribution of returns, and volatility clustering. To be able to replicate these properties, we have selected the following parameter setting:[1]

$$a = 1,\ b = 0.05,\ c = 0.02,\ d = 0.95,\ \varepsilon = 0.1,\ \lambda = 0.45,$$

$$\sigma^{\alpha} = 0.0025,\ \sigma^{\beta} = 0.025, \text{ and } \sigma^{\gamma} = 0.0025.$$

In the remaining part of the paper, we explore the dynamics of the model for different values of N. In particular, we increase N from 25 to 100 and to 500. In addition, for the case $N = 500$ we consider that there is more than one direct interaction between agents per trading time step.

3.1 Setting 1: $N = 25$

In our first experiment, we assume that there are only $N = 25$ agents. Of course, in real markets we usually observe a much larger number of traders. In the first step, it can be assumed that these agents reflect the trading activities of larger trading institutions or of groups of agents who collectively behave in the same manner (think, for instance, of group pressure). However, in the next subsections we increase the number of agents.

The seven panels of figure 1 aim at illustrating what kind of dynamics our model may produce for a limited number of speculators. In the top panel, we see the development of log prices. As can be seen, prices move erratically around their fundamental values. There are periods where prices are close to the fundamental value but occasionally larger bubbles set in. A prominent example is given around time step

[1] Interested readers should note that calibrating agent-based financial market models may be a time-consuming and pain-staking trial and error process. Some initial progress in estimating such models has recently been reported by Westerhoff and Reitz (2003), Alfarano et al. (2005), Boswijk et al. (2007), Manzan and Westerhoff (2007), and Winker et al. (2007).

4000, where the distance between log prices and log fundamental values is about 0.5, implying a substantial overvaluation of about 65 percent.

In the second panel, returns, defined as log price changes, are plotted. Note that extreme price changes are often larger than five percent, although the fundamental value is fixed. A constant fundamental value naturally implies that the entire volatility should be regarded as excess volatility. The third panel depicts the evolution of the weights of technical and fundamental trading strategies. As can be seen, there is a permanent evolutionary competition between the rules. Neither technical nor fundamental trading rules die out over time. We will come back to this soon.

In the two panels below, we characterize the distribution of the returns. Let us start with the left-hand panel. The solid line represents the distribution of the returns of our model, whereas the dotted line visualizes a normal distribution with identical mean and standard deviation. A closer inspection reveals that the distribution of returns of our financial market model has more probability mass in the center, less probability mass in the shoulder parts and more probability mass in the tails than the normal distribution. Estimates of the kurtosis support this view. However, the kurtosis is an unreliable indicator of fat-tailedness.

For this reason, we plot estimates of the tail index in the right-hand panel, varying the number of the largest observations from 0 to 10 percent. For this particular simulation run we obtain a tail index of about 3.7 (using the largest 5 percent of the observations). We found for other simulation runs that the tail index hovers around the range from 3.5 to 4.5, which may be slightly too high on average. Most tail indices estimated from real financial data seem to range between 3 and 4, and are almost always captured by the interval 2 to 5 (e.g. Lux 2009b).

In the last two panels, we plot the autocorrelation functions for raw returns and for absolute returns, respectively. Absence of significant autocorrelation between raw returns suggests that prices advance in a random walk-like manner. Despite the sporadic development of bubbles and crashes, it is thus hard to predict prices within our model. However, the autocorrelation coefficients for absolute returns are clearly significant and decay slowly. The autocorrelation coefficients are even positive for more than 100 lags. This is also in agreement with the second panel, and is a clear sign of volatility clustering, as observed in many real financial markets.

---------- Figure 1 ----------

From figure 1 we can also understand what is driving the dynamics of our model. Comparing the second and the third panel reveals that periods where technical analysis is rather popular are associated with higher volatility. Also, bubbles may be triggered in these periods. The trend-extrapolating (and highly noisy) nature of technical analysis has obviously a destabilizing impact on the dynamics. Note that technical analysis is quite profitable during the course of a bubble. As a result, more traders learn about this due to their interactions with other traders. Since technical analysis consequently gains in popularity, bubbles may possess some kind of momentum. A major shift from technical to fundamental analysis may be witnessed when a bubble collapses. A dominance of fundamental analysis then leads to a period where prices are closer towards fundamental values and where volatility is less dramatic.[2]

---

[2] Why do the weights of technical and fundamental analysis vary so erratically? Since prices fluctuate randomly it is hard for traders to make systematic profits, i.e. the difference in the fitness of the rules is (usually) rather limited, which, in turn, enables “spontaneous” swings in opinion. Put differently, if one of the rules outperformed the other one, it would also dominate the market. In addition, traders may change their opinion independently of market circumstances.

3.2 Setting 2: $N = 100$

Now we turn to the case with $N = 100$ traders. Figure 2 may be directly compared with figure 1, since it is based on the same simulation design. The only difference is that the number of traders is quadrupled. As indicated by the third panel, the popularity of the trading strategies now varies only very slowly over time. Therefore, there are extremely long periods where one or the other trading strategy dominates the market, which has some obvious consequences for the dynamics. For instance, between time steps 1500 and 2700 the majority of traders rely on fundamental analysis, and hence we find a period where prices are more or less in line with fundamental values and where absolute returns are rather low. Afterwards, technical analysis gains in strength and for the next 2000 time steps volatility is elevated. Since the model is calibrated to daily data, 2000 time steps correspond to a time span of about 8 years. Although some stylized facts may still be replicated for $N = 100$ agents, the dynamics of our model appears less convincing than before.

---------- Figure 2 ----------

Apparently, to generate realistic dynamics, the popularity of technical and fundamental trading rules has to vary more quickly, at least from a technical point of view. If there are only 25 traders, it may – in an extreme scenario – only take 25 time steps to accomplish a regime change from pure technical to pure fundamental analysis (or vice versa). An increase in the number of agents naturally increases the duration of such a complete regime switch. As seen in figure 2, regime changes may take a very long time if the number of agents is equal to 100 (of course, internal and external factors delay regime changes). In the next section, we try to show that this is not directly a "problem" of setting the number of agents too high. To achieve a reasonable fit of actual market dynamics with our model, the relation between the number of agents and the number of direct interactions between them per trading time step has to be within certain limits.

3.3 Setting 3: $N = 500$

Let us increase the number of agents up to $N = 500$. In addition, let us assume that there is not only one direct interaction between the agents per trading time step but that there are 20 contacts. Clearly, we now always run the interaction part of the model 20 times before we iterate the trading part of the model. As a result, the whole system may then again complete a full regime turn from pure fundamental to pure technical analysis (or the other way around) within 25 trading time steps.

Figure 3 presents the results. The qualitative similarities between figure 1 and figure 3 are striking. We recover bubbles and crashes, excess volatility, fat tails for the distribution of the returns, absence of autocorrelation for raw returns, and volatility clustering, i.e. our model again mimics key stylized facts of financial markets quite well.

---------- Figure 3 ----------

Two further comments are required. Note first that periods of high volatility may or may not be associated with bubbles and crashes. It may thus happen that prices fluctuate wildly around fundamental values. We consider it interesting that there is not a strict one-to-one relation between high volatility and bubble periods.[3] Finally, although the model once again generates a distribution which deviates from the normal distribution, in the sense that there is more probability mass in its tails, the fat-tailedness

[3] This implies that technical analysis may also outperform fundamental analysis in a non-bubble period; otherwise its weight – which is driven by the agent's learning behavior – would not have increased.

could be stronger. For the underlying simulation run we compute a tail index of 4.3. Other simulation runs generate indices between 3.5 and 4.5, as was the case for $N = 25$ traders.

3.4 Robustness of the dynamics

Instead of performing a larger and more sophisticated Monte Carlo study to check the robustness of our results, we restrict ourselves in the following to presenting and discussing additional simulation runs. The reason for doing this is that we strongly believe in the strength of the human eye, which has a remarkable ability to identify both regularities and irregularities in time series. It is also instructive to inspect single simulation runs. Phenomena such as bubbles and crashes or volatility outbursts are infrequent, irregular phenomena, and by measuring them with certain statistics their true nature is at least partially lost. However, we ascertained that a more elaborate statistical analysis would also confirm the robustness of the dynamics.[4]

Figure 4 displays four repetitions of the first three panels of figure 1. The only difference between figure 1 and figure 4 is that we have exchanged the seeds for the random variables. Note that all simulation runs are characterized by an endogenous competition between the trading rules. Volatility clustering is always visible, whereas bubbles and crashes may be absent for longer time periods or may evolve on a smaller scale. However, and this is one of the reasons why we should pay attention to these simulation runs, the panels show us that even after a very long time period without significant mispricing the next bubble may be just about to kick in. This warning may have a philosophical attitude but, given the common sense of policy makers, it seems

[4] Also “modest” changes in the parameter setting do not destroy the model’s ability to mimic actual asset price dynamics reasonably well.

important to us to note that even a stable period of, say, 10 years does not guarantee that the future will also be stable. A major bull or bear market period may just be days away without much forewarning.

---------- Figure 4, 5 and 6 ----------

Figure 5 extends the analysis for $N = 100$ traders. In all simulation runs we see that the degree of volatility clustering is presumably exaggerated. The reason for this is that swings in opinion take too much time. Finally, figure 6 demonstrates that our model may generate realistic dynamics for a scenario with $N = 500$ agents and 20 direct interactions per trading time step.

4 Conclusions

The goal of this paper is to develop a simple agent-based financial market model with direct interactions between the market participants. When the traders meet each other within our model, they compare the past success of their trading rules. Should an agent discover that his opponent uses a more profitable strategy, it is quite likely that he/she will change his/her strategy. Simulations reveal that such a setting may incorporate a permanent evolutionary competition between the trading rules. For instance, there may be periods where fundamental analysis dominates the markets. Prices then fluctuate in the vicinity of their fundamental values. However, at some point in time a major shift towards technical analysis may set in and the market becomes unstable. Besides an increase in volatility, spectacular bubbles and crashes may materialize.

Moreover, we have demonstrated that our model may generate realistic dynamics for a lower or higher number of traders. However, in the latter case we have to increase the number of interactions per trading time step. Otherwise the relative importance of the trading rules is not flexible enough – due to the assumed tandem recruitment process. Of course, one could also consider increasing the number of agents further, say, to 5000 traders. Interesting dynamics may still be recovered as long as the number of contacts between the agents per trading time step is appropriately adjusted.

One interesting extension of the current setup may be to consider that (also) the probability that an agent changes his opinion independently from social interactions is state dependent. One could, for instance, assume that the probability to switch from a technical to a fundamental attitude is relatively high if fundamental analysis outperforms technical analysis. In this sense, the agents would then (also) display some kind of individual economic reasoning behavior.

Finally, we would like to point out that, with a bit of experience, it is quite simple to program our model. It should therefore be possible, even for interested laymen, to reproduce the dynamics of our model. From a scientific point of view, replication of results is important. Everything required for such an exercise is given in our paper.

References

Alfarano, S., Lux, T. and Wagner, F. (2005): Estimation of agent-based models: the case of an asymmetric herding model. Computational Economics, 26, 19-49.

Boswijk, P., Hommes, C. and Manzan, S. (2007): Behavioral heterogeneity in stock prices. Journal of Economic Dynamics and Control, 31, 1938-1970.

Brock, W. and Hommes, C. (1997): A rational route to randomness. Econometrica, 65, 1059-1095.

Brock, W. and Hommes, C. (1998): Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economic Dynamics Control, 22, 1235-1274.

Chen, S.-H., Chang, C.-L. and Du, Y.-R. (2009): Agent-based economic models and econometrics. Knowledge Engineering Review, in press

Chiarella, C. (1992): The dynamics of speculative behavior. Annals of Operations Research, 37, 101-123.

Chiarella, C, Dieci, R, and Gardini, L. (2002): Speculative behaviour and complex asset price dynamics: a global analysis. Journal of Economic Behavior and Organization, 49, 173-197.

Cont, R. (2001): Empirical properties of asset returns: stylized facts and statistical issues. Quantitative Finance, 1, 223-236.

Day, R. and Huang, W. (1990): Bulls, bears and market sheep. Journal of Economic Behavior and Organization, 14, 299-329.

De Grauwe, P., Dewachter, H. and Embrechts, M. (1993): Exchange rate theory – chaotic models of foreign exchange markets. Blackwell: Oxford.

De Grauwe, P. and Grimaldi, M. (2006): Heterogeneity of agents, transactions costs and the exchange rate. Journal of Economic Dynamics and Control, 29, 691-719.

Farmer, D. and Joshi, S. (2002): The price dynamics of common trading strategies. Journal of Economic Behavior and Organization, 49, 149-171.

Frankel, J. and Froot, K. (1986): Understanding the U.S. dollar in the eighties: the expectations of chartists and fundamentalists. Economic Record, 62, 24-38.

Graham, B. and Dodd, D. (1951): Security analysis. New York.

Hommes, C. (2006): Heterogeneous agent models in economics and finance. In: Tesfatsion, L. and Judd, K. (eds.): Handbook of computational economics Vol. 2: Agent-based computational economics. North-Holland, Amsterdam, 1107-1186.

Kirman, A. (1991): Epidemics of opinion and speculative bubbles in financial markets. In: Taylor, M. (ed.): Money and financial markets. Blackwell: Oxford, 354-368.

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

LeBaron, B. (2006): Agent-based computational finance. In: Tesfatsion, L. and Judd, K. (eds.): Handbook of computational economics Vol. 2: Agent-based computational economics. North-Holland, Amsterdam, 1187-1233.

Li, H. and Rosser, B. (2001): Emergent volatility in asset markets with heterogeneous agents. Discrete Dynamics in Nature and Society, 6, 171-180.

Li, H. and Rosser, B. (2004): Market dynamics and stock price volatility. European Physical Journal B, 39, 409-413.

Lux, T. (1995): Herd behavior, bubbles and crashes. Economic Journal, 105, 881-896.

Lux, T. (1998): The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. Journal of Economic Behavior and Organization, 33, 143-165.

Lux, T. (2009a): Financial power laws: empirical evidence, models and mechanisms. In:

Cioffi-Revilla, C. (ed.): Power laws in the social sciences: Discovering complexity and non-equilibrium dynamics in the social universe, in press.

Lux, T. (2009b): Applications of statistical physics in finance and economics. In: Rosser, J.B., Jr. (ed): Handbook of Research on Complexity. Edward Elgar, Cheltenham, in press.

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

Lux, T. and Marchesi, M. (2000): Volatility clustering in financial markets: a micro-simulation of interacting agents. International Journal of Theoretical and Applied Finance, 3, 675-702.

Lux, T. and Ausloos, M. (2002): Market fluctuations I: Scaling, multiscaling, and their possible origins. In: Bunde, A., Kropp, J. and Schellnhuber, H. (Eds.): Science of disaster: climate disruptions, heart attacks, and market crashes. Springer: Berlin, 373-410.

Mantegna, R. and Stanley, E. (2000): An introduction to econophysics. Cambridge University Press: Cambridge.

Manzan, S. and Westerhoff F. (2007): Heterogeneous expectations, exchange rate dynamics and predictability. Journal of Economic Behavior and Organization, 64, 111-128.

Menkhoff, L. (1997): Examining the use of technical currency analysis. International Journal of Finance and Economics, 2, 307-318.

Menkhoff, L. and Taylor, M. (2007): The obstinate passion of foreign exchange professionals: technical analysis. Journal of Economic Literature, 45, 936-972.

Murphy, J. (1999): Technical analysis of financial markets. New York Institute of

Finance: New York.

Rosser, B., Ahmed, E. and Hartmann, G. (2003): Volatility via social flaring. Journal of Economic Behavior and Organization, 50, 77-87.

Taylor, M. and Allen, H. (1992): The use of technical analysis in the foreign exchange market. Journal of International Money and Finance, 11, 304-314.

Westerhoff, F. (2008): The use of agent-based financial market models to test the effectiveness of regulatory policies. Jahrbücher für Nationalökonomie und Statistik, 228, 195-227.

Westerhoff, F. (2009): Exchange rate dynamics: A nonlinear survey. In: Rosser, J.B., Jr. (ed): Handbook of Research on Complexity. Edward Elgar, Cheltenham, in press.

Westerhoff, F. and Reitz, S. (2003): Nonlinearities and cyclical behavior: the role of chartists and fundamentalists. Studies in Nonlinear Dynamics and Econometrics, 7 (issue 4, article 3).

Westerhoff, F. and Dieci, R. (2006): The effectiveness of Keynes-Tobin transaction taxes when heterogeneous agents can trade in different markets: a behavioral finance approach. Journal of Economic Dynamics and Control, 30, 293-322.

Winker, P., Gilli, M. and Jeleskovic, V. (2007): An objective function for simulation based inference on exchange rate data. Journal of Economic Interaction and Coordination, 2, 125-145.

Legends for figures 1-6

Figure 1: The first three panels show the evolution of log prices, the returns, and the weights of technical trading rules, respectively. The left-hand panel in the fourth line depicts the distribution of the returns (the dotted line gives the corresponding normal distribution), whereas the left-hand panel presents estimates of the tail index. The bottom two panels depict the autocorrelation coefficients of raw and absolute returns, respectively. The simulation is based on 5000 time steps (omitting a longer transient period) and $N = 25$ traders. The remaining parameters are specified in section 3.

Figure 2: The same simulation design as in figure 1, except that we now consider $N = 100$ agents.

Figure 3: The same simulation design as in figure 1, except that we now consider $N = 500$ agents and 20 direct interactions per trading time step.

Figure 4: Four repetitions of figure 1 using different random number streams.

Figure 5: Four repetitions of figure 2 using different random number streams.

Figure 6: Four repetitions of figure 3 using different random number streams.

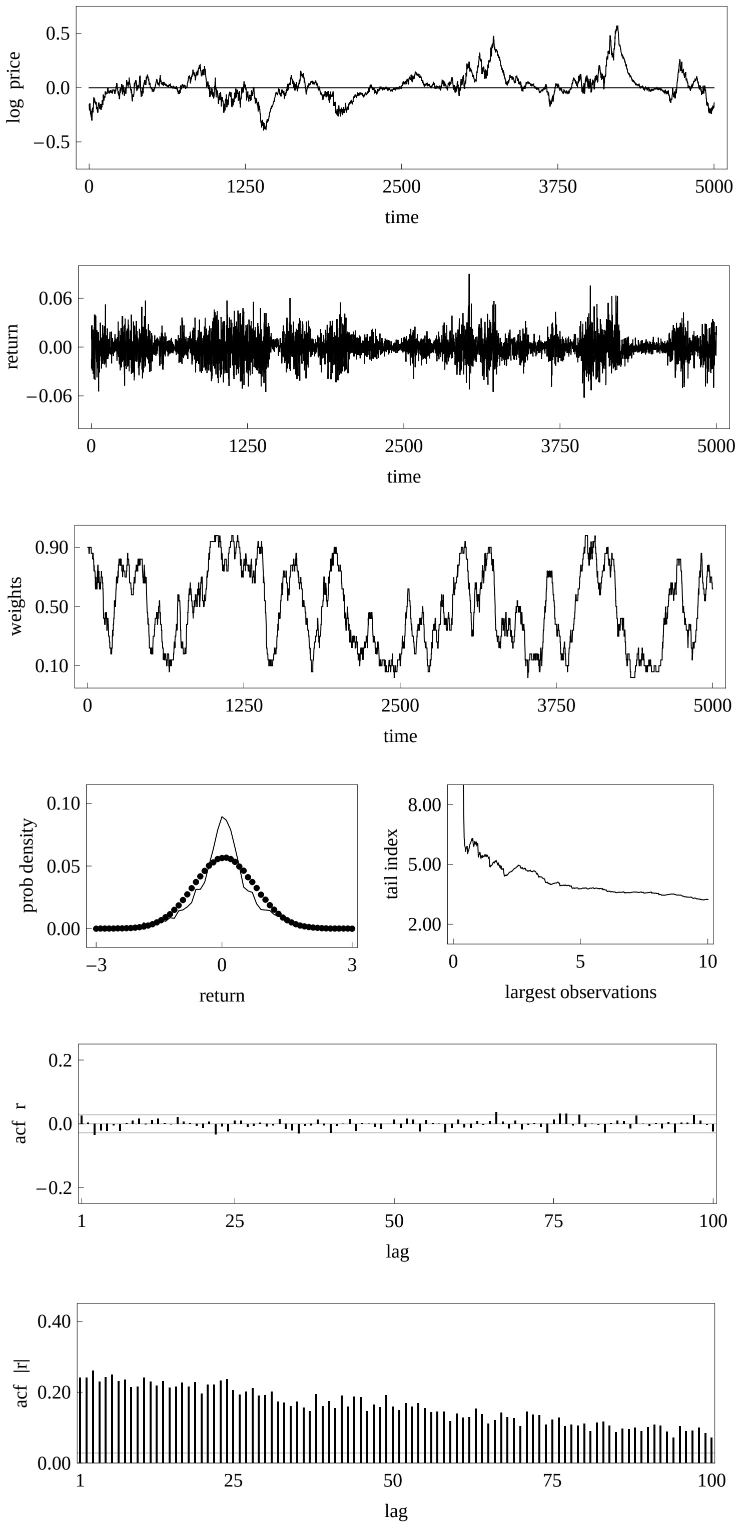
log price
0.5
0.0
−0.5
0
1250
2500
3750
5000
time
return
0.06
0.00
−0.06
weights
0.90
0.50
0.10
prob density
0.10
0.05
0.00
−3
0
3
tail index
8.00
5.00
2.00
0
5
10
largest observations
acf r
0.2
0.0
−0.2
1
25
50
75
100
lag
acf |r|
0.40
0.20
0.00

Figure 1

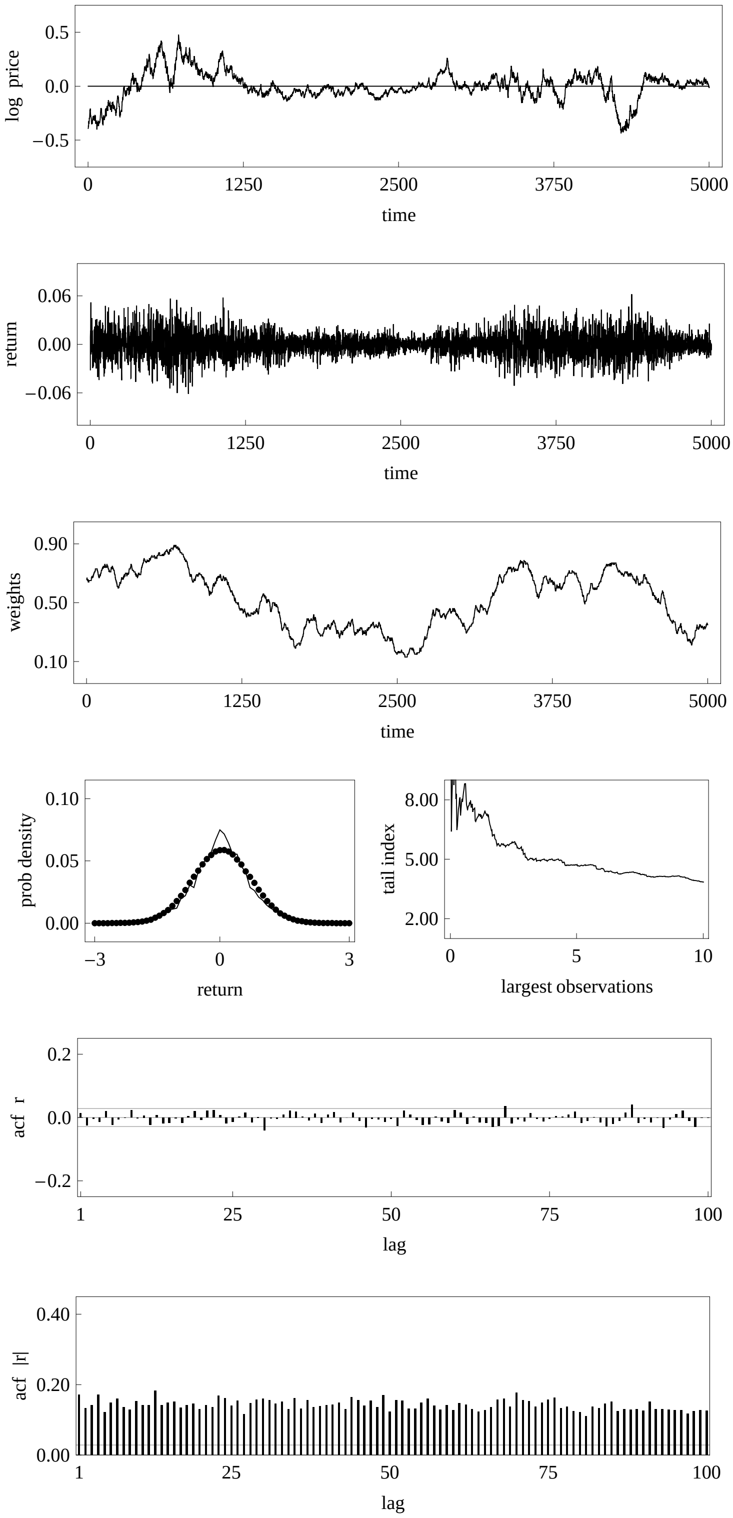
log price
0.5
0.0
−0.5
0
1250
2500
3750
5000
time
return
0.06
0.00
−0.06
0
1250
2500
3750
5000
time
weights
0.90
0.50
0.10
0
1250
2500
3750
5000
time
prob density
0.10
0.05
0.00
−3
0
3
return
tail index
8.00
5.00
2.00
0
5
10
largest observations
acf r
0.2
0.0
−0.2
1
25
50
75
100
lag
acf |r|
0.40
0.20
0.00
1
25
50
75
100
lag

Figure 2

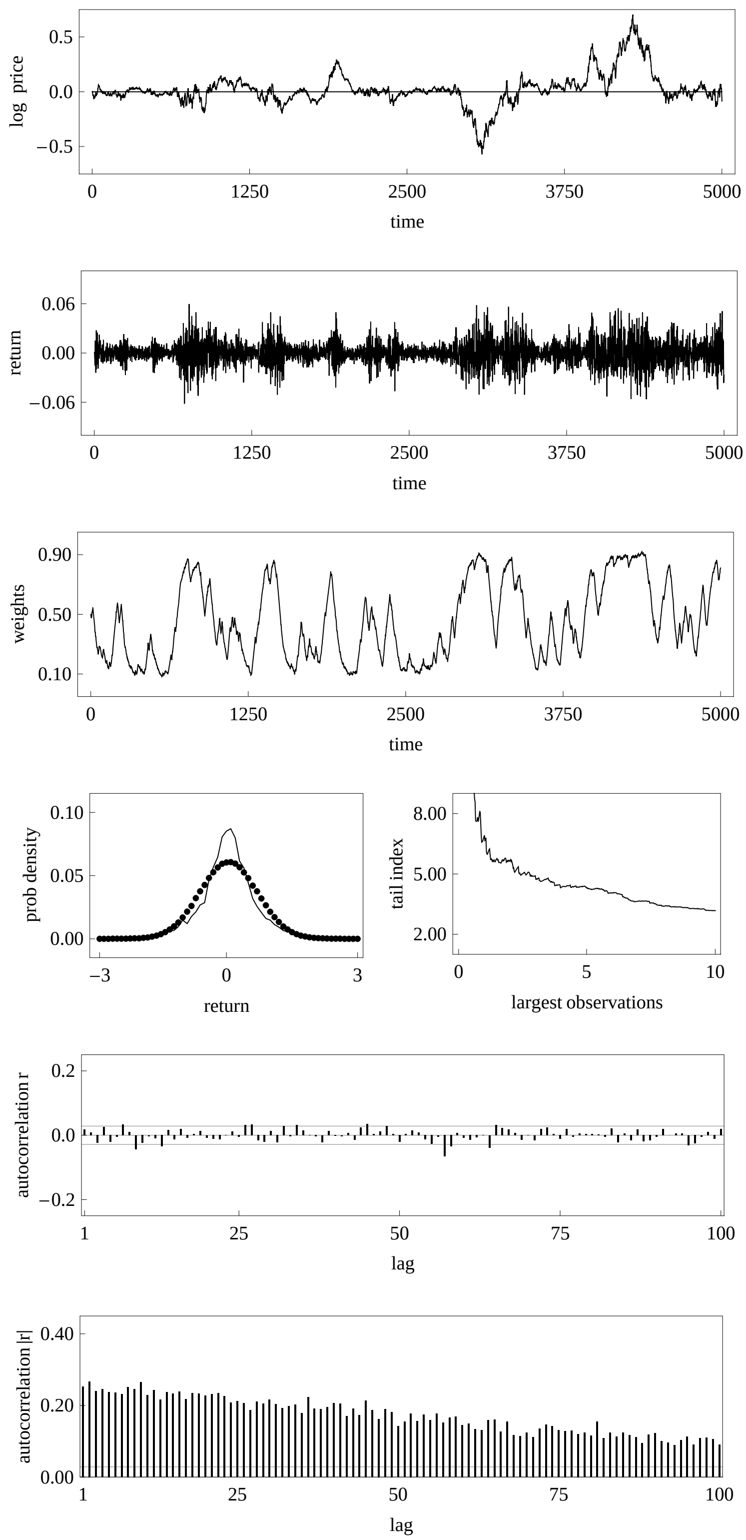
log price
0.5
0.0
−0.5
0
1250
2500
3750
5000
time
return
0.06
0.00
−0.06
time
weights
0.90
0.50
0.10
time
prob density
0.10
0.05
0.00
−3
0
3
return
tail index
8.00
5.00
2.00
0
5
10
largest observations
autocorrelation r
0.2
0.0
−0.2
1
25
50
75
100
lag
autocorrelation |r|
0.40
0.20
0.00
lag

Figure 3

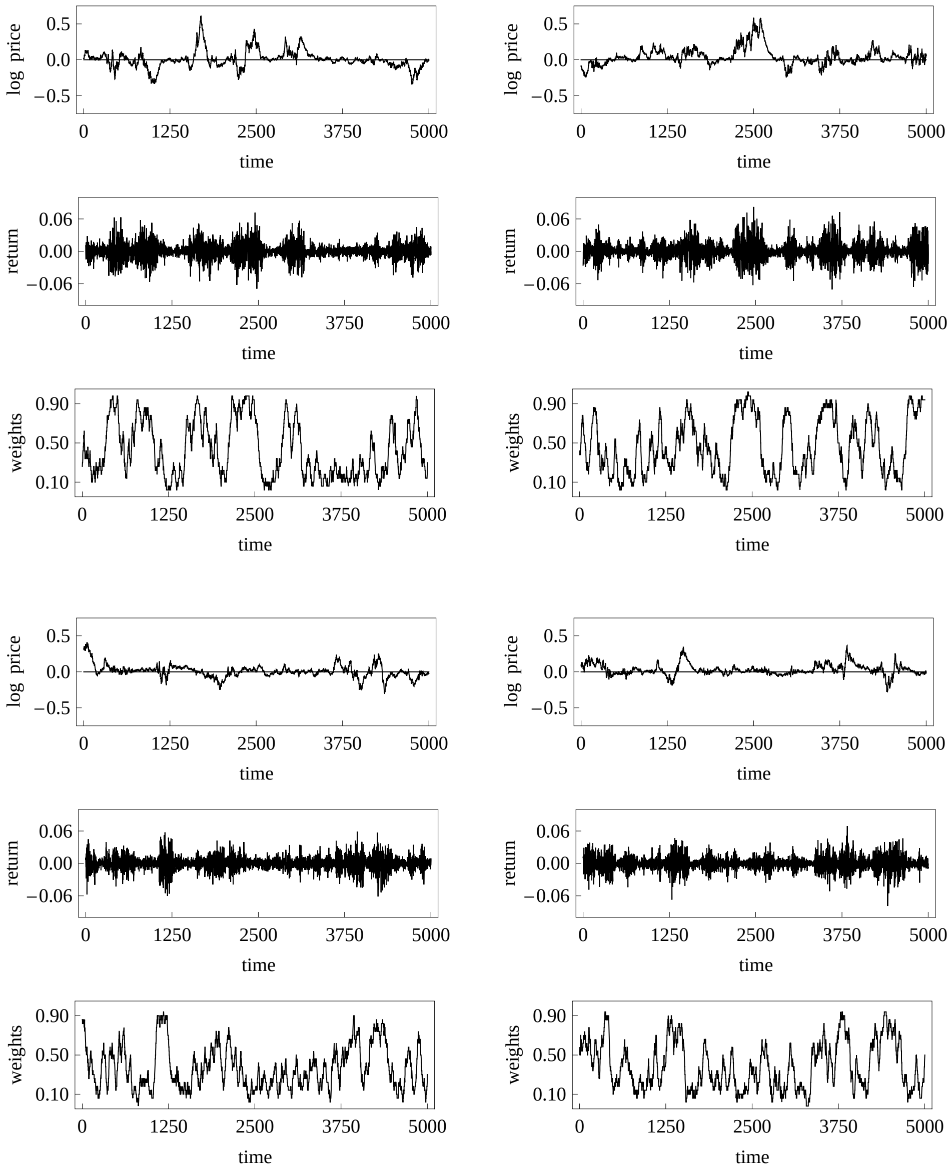
log price
return
weights
time
0
1250
2500
3750
5000
0.5
0.0
−0.5
0.06
0.00
−0.06
0.90
0.50
0.10

Figure 4

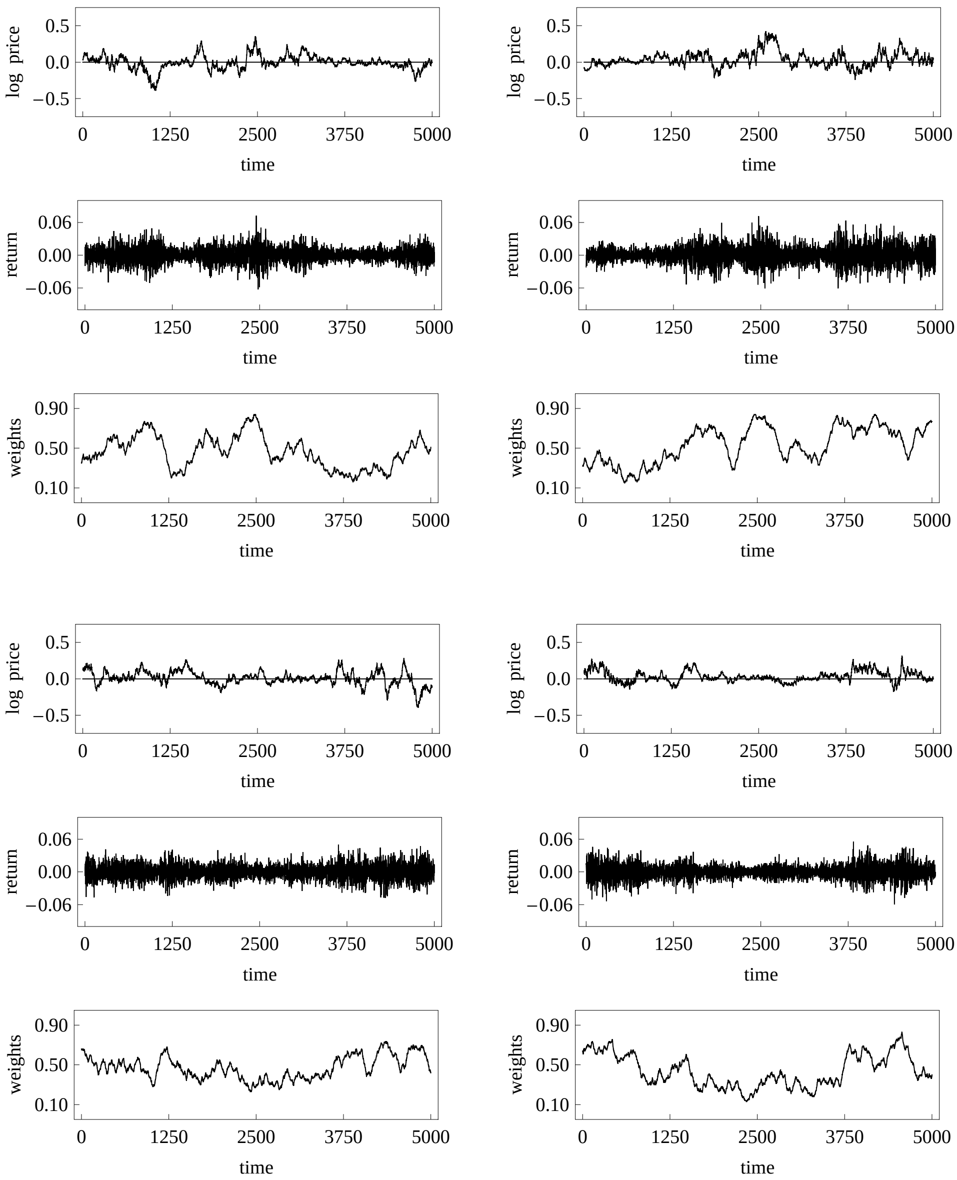
log price
0.5
0.0
−0.5
0
1250
2500
3750
5000
time
return
0.06
0.00
−0.06
weights
0.90
0.50
0.10

Figure 5

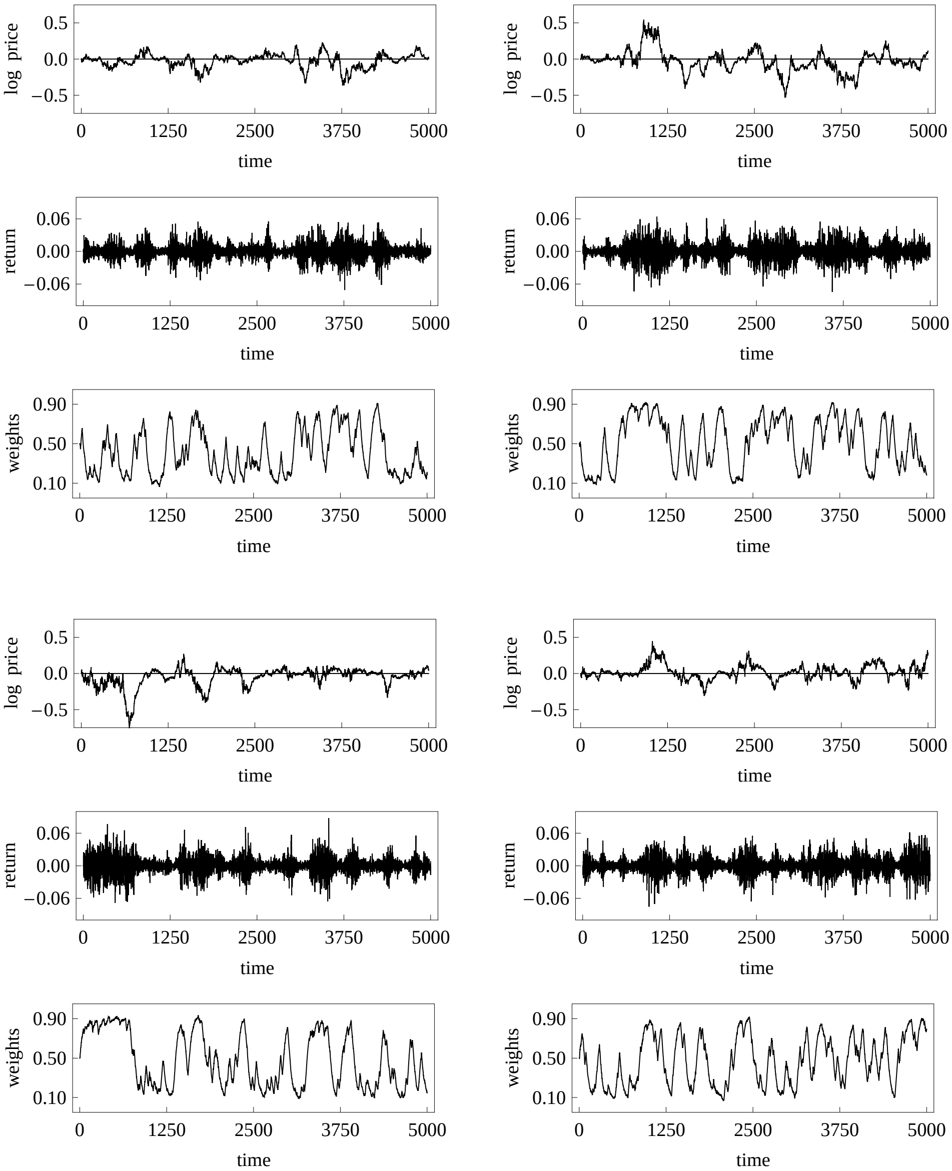
log price
0.5
0.0
−0.5
0
1250
2500
3750
5000
time
return
0.06
0.00
−0.06
weights
0.90
0.50
0.10

Figure 6