

Using agent-based models to study the
dynamics and (in)stability of financial
markets

Doctoral thesis

IVONNE SCHWARTZ

Supervisor: Prof. Dr. Frank Westerhoff

Submitted in fulfillment of the requirements for the degree of
Doctor rerum politicarum (Dr. rer. pol.) at the Faculty of Social
Sciences, Economics, and Business Administration of the
University of Bamberg.

Bamberg 2024

Diese Arbeit hat der Fakultät Sozial- und Wirtschaftswissenschaften der Otto-Friedrich-Universität Bamberg als kumulative Dissertation vorgelegen.

1. Gutachter: Prof. Dr. Frank Westerhoff
2. Gutachter: Prof. Dr. Christian R. Proaño
3. Gutachter: Prof. Dr. Mishael Milaković

Tag der mündlichen Prüfung: 10.11.2023

Dieses Werk ist als freie Onlineversion über das Forschungsinformationssystem (FIS; <https://fis.uni-bamberg.de>) der Universität Bamberg erreichbar.

Dieses Werk ist durch das deutsche Urheberrecht geschützt. Es steht Ihnen frei, dieses Werk auf jede Art und Weise zu nutzen, die durch die geltende Gesetzgebung zum deutschen Urheberrecht erlaubt ist. Für andere Verwendungszwecke müssen Sie die Erlaubnis der Rechteinhaberinnen und Rechteinhaber einholen.

URN: urn:nbn:de:bvb:473-irb-982101

DOI: <https://doi.org/10.20378/irb-98210>

Von der genannten Lizenzangabe ausgenommen sind folgende Bestandteile dieser Dissertation:

Die Artikel "Estimation of agent-based models: Testing and applying a simulated joint moment approach" (S. 50-94) und "Time is limited on the road to asymptopia: An analysis of the ergodic properties of moment functions in the validation of financial agent-based models" (S. 96-141) stehen unter der CC-Lizenz CC BY.

Lizenzvertrag: Creative Commons Namensnennung 4.0

<https://creativecommons.org/licenses/by/4.0/>



To Elsa, Margarete and all ambitious girls.

Contents

1	Introduction	5
2	Market entry waves and volatility outbursts in stock markets	8
3	Heterogeneous speculators and stock market dynamics: a simple agent-based computational model	28
4	Estimation of agent-based models: Testing and applying a simulated joint moment approach	49
5	Time is limited on the road to asymptopia: An analysis of the ergodic properties of moment functions in the validation of financial agent-based models	95

1 Introduction

The 15th of September, 2008 marks an historic event. Lehman Brothers, the fourth-largest U.S. investment bank at that time, filed for bankruptcy. The subsequent global financial crisis 2007/08 was the biggest post-war era downturn. It has not only triggered a worldwide recession, but also had negative impact on a humanitarian level given roaring unemployment and suicide rates. The crisis further revealed the failure of standard economic models to neither explain nor anticipate these events. Those neoclassical models are still dominating within the economics profession, although they are built on doubtful assumptions of e.g. a perfectly rational and fully informed representative agent maximizing her utility function or the belief that free markets guarantee perfectly efficient price equilibria. In such a perfect world, financial crises simply do not exist. It became, hence, apparent that we need new tools to better understand the complexity of today's economic world. In fact, we observe complex behavior on a macro-level in financial markets. Those emergent phenomena, like bubbles and crashes, volatility clustering and fat tails, are universal among different kinds of financial markets. One approach to study these features is agent-based modeling (ABM). The main idea of ABM is that aggregate behavior is reconstructed from the bottom-up by computing what emerges from the micro-behavior and local interaction of heterogeneous agents. Since human attention capability has been proven to be limited in nature, those agents are supposed to have bounded rationality, a limited amount of processable information and limited computing power. Typically, they make use of simple heuristic rules based on local information. Simulation runs are then used to study the dynamics of the system providing us with artificial laboratories.

One drawback of ABM is that the modeler is left with enormous degrees of freedom in choosing the types of agents and their behavioral rules. Consequently, existing models often differ in their assumptions and are perceived as black boxes making it hard to clearly detect which macro pattern is a result of which micro property. Therefore, we need models that have convincing building blocks, generate realistic dynamics and produce reasonable stories. Furthermore, to increase their acceptance among economists and policy-makers, we need powerful yet simple models that are econometrically tested and verified. Estimating the parameters of an ABM is, yet, a challenging and computationally demanding task. Often, the dynamic properties

cannot be studied analytically. Their high dimensionality and inherent non-linearity further complicate any estimation approach.

The goal of this thesis is two-folded. The first two papers are dedicated to foster the research on agent-based financial market models by studying two newly developed models that aim at explaining a broad number of stylized facts of financial markets. Chapters 2 and 3 seek to answer how behavioral patterns of speculators may help to explain complex phenomena of financial markets.

The first paper in *Chapter 2* studies how market entry decisions of single market participants are subject to herding behavior and market risk. The paper shows, both analytically and numerically, that speculators' market entry and exit behavior can explain volatility clustering. In fact, herding-induced market entry waves lead to an further increased excess demand triggering long-lasting periods of high volatility. This then turns into higher stock market risk. Speculator's evaluation of heightened stock market risk drives them out of the market and volatility eventually decreases. Simulations show that the model is able to produce prominent statistical properties of financial markets including bubbles and crashes, excess volatility, fat-tailed return distributions and serially uncorrelated price changes.

The second paper in *Chapter 3* introduces a simple agent-based financial model in which the trading behavior of heterogeneous interacting speculators causes bubbles and crashes, excess volatility, serially uncorrelated returns, fat-tailed return distributions, and volatility clustering. Speculators are heterogeneous in their trading behavior since they individually derive trading signals from fundamental and technical analysis. Consequently, stock prices fluctuate constantly around their fundamental values and are excessively volatile. If speculators, however, collectively react to similar trading signals, heterogeneity spontaneously vanishes and extreme returns emerge. Instead of exogenous sunspots, the model endogenously generates short-lived periods in which speculators' behavior is coordinated causing market turmoil. Periods of high volatility are long-lasting since speculators persistently receive strong trading signals due to past price movements. Numerical simulations also reveal that circuit breakers may effectively combat market turbulences.

The second aim of this thesis is to deepen the understanding of estimation and validation approaches when applied to ABM. With the rise of computational power, simulation-based methods have become a popular choice. Their basic idea is to define an objective function over a parameter vector that minimizes the distance between a set of empirical and simulated moments. This intuitive and transparent approach makes simulated moment estimators quite appealing. Although they have well-understood mathematical properties, their regularity conditions are not fully met

by most ABM leading to biases in the parameter estimates. Chapters 4 and 5 of this thesis are going to shed light on the potential difficulties when applying simulated moment estimators to validate agent-based financial market models.

The paper of *Chapter 4* builds on the work of the previous chapter as it takes its theoretical model and estimates it. In fact, the goal here is to study the performance and properties of a simulated joint moment estimator as guidance for future research. We test its ability to recover parameters consistently and efficiently using Monte Carlo simulations. We focus on the replicability of the study by offering a transparent, not overly technical and easily computable estimation framework. The computational burden of the repetitive Monte Carlo runs is reduced by the use of machine-learning surrogates. As it turns out, this allows to explore the parameter space at low costs.

The goal of the final paper in *Chapter 5* is to further add to the blooming research on the estimation of agent-based financial market models. This paper deals with the impact of broken ergodicity on the convergence of a simulated moment estimator and its single moment functions. Ergodicity is, besides stationarity, a precondition for any estimation-related approach. Its absence, hence, violates regularity conditions and results in uncertainty in terms of biases in the estimates. Therefore, we systematically study how to reduce these biases when the computational resources are limited. We take two prototype financial market ABM for which we run various Monte Carlo experiments. We find that for most moment functions the convergence times are infeasibly long, thus leaving us in pre-asymptopia. Choosing an efficient mix of ensemble size and simulated time length can help guiding validation efforts through this jungle of uncertainty.

The research on the empirical validation and estimation of agent-based financial market models is a very active field where many different approaches have emerged during the last decade. They range from information-theoretic similarity-based measures, over likelihood-based methods to a whole bunch of Bayesian applications. So far, the research community has not reached a state of consensus. This is partly caused by the immense computational burden that most sophisticated techniques suffer from. The thesis is, hence, dedicated to focus on the efficiency and feasibility of estimation methods allowing them to run on standard computing resources without the need for high-computing clusters. This thesis' ultimate goal is to advocate for inclusive and fully accessible research which is one key to unlock the full potential of agent-based modeling.

2 Market entry waves and volatility outbursts in stock markets

This chapter contains joint work with Noemi Schmitt and Frank Westerhoff which has already been published in the *Journal of Economic Behavior & Organization*, 2018, Vol. 153: 19-37: <https://doi.org/10.1016/j.jebo.2018.03.022>. Frank Westerhoff conceived the original idea of the paper and supervised the project. Noemi Schmitt performed the analytical analysis of the model. Ivonne Schwartz worked out the numerical simulations for the stochastic dynamics of the model. All three authors contributed to equal parts to the discussion of the results and the writing of the final manuscript. In the following, the published version will be included.



Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jeboMarket entry waves and volatility outbursts in stock markets[☆]

Ivonne Blaurock, Noemi Schmitt, Frank Westerhoff*

University of Bamberg, Department of Economics, Germany



ARTICLE INFO

Article history:
Received 8 August 2017
Revised 12 March 2018
Accepted 25 March 2018

JEL classification:
C63
D84
G15

Keywords:
Stock markets
Heterogeneous speculators
Exponential replicator dynamics
Herding behavior
Stylized facts

ABSTRACT

We develop a simple agent-based financial market model in which speculators' market entry decisions are subject to herding behavior and market risk. In addition, speculators' orders depend on price trends, market misalignments and fundamental news. Using a mix of analytical and numerical tools, we show that a herding-induced market entry wave may amplify excess demand, triggering lasting volatility outbursts. Eventually, however, higher stock market risk reduces stock market participation and volatility decreases again. Simulations furthermore reveal that our approach is also able to produce bubbles and crashes, excess volatility, fat-tailed return distributions and serially uncorrelated price changes. Moreover, trading volume is persistent and correlated with volatility.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

The goal of our paper is to develop a simple agent-based financial market model to explain a number of important stylized facts of stock markets. In particular, we analytically and numerically demonstrate that speculators' market entry and exit behavior may give rise to volatility clustering. Our model's key features and its main implications may be summarized as follows. We assume that there is a market maker who adjusts stock prices with respect to speculators' orders, which, in turn, use technical and fundamental trading rules to determine their trading behavior. Speculators' market entry decisions depend on two socio-economic principles. First, speculators are subject to herding behavior and increasingly enter the stock market as the number of active speculators increases. Second, speculators react to stock market risk. The higher the past volatility of the stock market, the lower the probability that a speculator will enter the stock market. As it turns out, the stock market is relatively stable if the number of active speculators is low. Since stock market risk is then perceived as negligible, more and more speculators become active. Consequently, excess demand increases, the market maker adjusts stock prices more strongly and volatility picks up. Due to the increase in stock market risk, stock market participation eventually decreases again. Confronted with a lower excess demand, the market maker needs to adjust stock prices less strongly.

[☆] Presented at the 23rd International Conference on Computing in Economics and Finance, June 28–30, 2017, New York City, United States of America, and at the Summer School in Economics and Finance, July 17–21, Alba di Canazei, Italy. We thank the participants, in particular Roberto Dieci and Jan Tuinstra, for their helpful comments. Our paper also benefited from valuable feedback of an anonymous referee.

* Corresponding author.

E-mail address: frank.westerhoff@uni-bamberg.de (F. Westerhoff).

<https://doi.org/10.1016/j.jebo.2018.03.022>

0167-2681/© 2018 Elsevier B.V. All rights reserved.

We also show that the repeated inflow and outflow of speculators along with their heterogeneous trading behavior may produce bubbles and crashes, excess volatility, serially uncorrelated returns and a fat-tailed return distribution. Moreover, trading volume displays significant memory effects and is strongly correlated with volatility. Keeping track of the individual speculators' wealth dynamics reveals that heterogeneity among speculators may be a persistent phenomenon of financial markets, i.e. neither do a few speculators accumulate all the wealth and dominate the market nor does a substantial fraction of speculators go bankrupt and vanish from the market.

Our paper adds to the burgeoning stream of literature on agent-based financial market models (see Chiarella et al., 2009a; Hommes and Wagener, 2009; Lux, 2009 for surveys). Within these models, speculators apply technical and fundamental trading rules to determine their orders. Technical trading rules (Murphy, 1999) are usually based on trend extrapolation and tend to destabilize the dynamics of financial markets. In contrast, fundamental trading rules (Graham and Dodd, 1951) bet on mean reversion, exercising a stabilizing impact on the dynamics of financial markets. Models by Day and Huang (1990), De Grauwe et al. (1993), Brock and Hommes (1998), LeBaron et al. (1999), Farmer and Joshi (2002), Chiarella et al. (2007), Franke and Westerhoff (2012) and Jacob Leal and Napoletano (2017), for example, show that (non-linear) interactions between speculators relying on technical and fundamental trading rules can produce dynamics which resembles the dynamics of actual financial markets quite closely. Without question, this line of research helps us to improve our understanding of the functioning of financial markets. For instance, agent-based financial market models reveal that a bubble may emerge if speculators forcefully rely on technical analysis while a crash can be set in motion if speculators put more weight on fundamental analysis. Such a time-varying impact of technical and fundamental trading rules can also produce volatility clustering. Financial markets tend to be relatively stable when speculators prefer fundamental analysis but turn wilder when speculators opt for technical analysis.

Herding behavior plays a prominent role in a number of agent-based financial market models. In Alfarano and Lux (2007); Kirman (1993); Lux and Marchesi (1999), speculators' herding behavior influences whether they choose technical or fundamental trading rules to determine their orders. Cont and Bouchaud (2000) and Stauffer et al. (1999) assume that speculators' herding behavior influences whether they are optimistic or pessimistic. Bischi et al. (2006) show that complex asset price dynamics may emerge if speculators mimic the buying and selling behavior of other speculators. LeBaron and Yamamoto (2008) study imitation behavior which results from social learning and show that it may be responsible for long memory effects in trading volume and volatility. Tedeschi et al. (2012) develop a model in which speculators imitate the behavior of more successful speculators. In Schmitt and Westerhoff (2017), speculators' herding behavior may lead to changes in the heterogeneity of trading rules applied. Compared to these models, we assume in our paper that speculators' herding behavior affects their stock market participation.

In fact, empirical evidence suggests that stock market participation changes over time and is influenced by social interactions. Most importantly for our approach, Hong et al. (2004, 2005), Brown et al. (2008) and Shiller (2015) report that households and professional investors regard a stock market as increasingly attractive the more of their peers participate in it. Surprisingly, there are only a few agent-based models which explicitly study speculators' market entry and exit behavior. Alfi et al. (2009b, 2009c, 2009a) show that agent-based models with a fixed number of speculators may lose their ability to produce realistic dynamics if the number of speculators is set either too high or too low. Against this background, they endogenize the number of speculators and explore under which conditions the model dynamics may self-organize such that the number of active speculators approaches a level which generates realistic dynamics. Iori (1999, 2000) develops a more involved agent-based simulation framework with heterogeneous interacting agents. Due to trade frictions, such as trading costs or information processing constraints, speculators may become inactive. However, communication and imitation among speculators may lead to a spontaneous spark in stock market participation and elevate price fluctuations. To study the effects of transaction taxes, Westerhoff and Dieci (2006) develop a model in which speculators have the choice between technical trading, fundamental trading and being inactive. Speculators' choices depend on the past profitability of these alternatives. Schmitt and Westerhoff (2016) show that although speculators' inflow and outflow may create bubbles and crashes, their market entry and exit behavior is not subject to herding effects.

Our approach differs to these contributions in several dimensions. One advantage of our model is that its deterministic skeleton allows us to derive a number of analytical insights which make the model's functioning and the origin of volatility clustering rather transparent. For instance, our model possesses a steady state in which prices reflect their fundamental values and in which all speculators are active. We analytically show that this steady state becomes unstable (via a Neimark–Sacker bifurcation) if speculators strongly extrapolate past price trends. Simulations reveal that the dynamics we then observe are characterized by alternating periods of high volatility, pushing destabilizing speculators out of the stock market, and periods of low volatility, attracting destabilizing speculators to the stock market. The same forces are at work in a stochastic version of our model which is able to mimic a number of important time series properties of stock markets. It is important to note that our results are not driven by speculators who constantly lose money or by speculators who become very rich.

The rest of our paper is organized as follows. In Section 2, we present our simple agent-based financial market model. In Section 3, we study the properties of the model's deterministic skeleton. In Section 4, we illustrate that the model's stochastic version is able to replicate a number of important stylized facts of stock markets. In Section 5, we conclude our paper and point out some avenues for future research.

2. A simple agent-based financial market model

The key elements of our agent-based financial market model may be summarized as follows. We consider a stock market which is populated by a single market maker and a time-varying number of heterogeneous speculators. A market maker adjusts the stock price with respect to speculators' orders which, in turn, depend on the stock market's price trend, its misalignment and current fundamental news. The probabilistic market entry decision of a speculator is repeated at the beginning of each trading period. We assume that the probability that a given speculator will enter the market increases with current stock market participation and decreases with current stock market risk. Since the total number of speculators is fixed, the number of active speculators follows a binomial distribution. As we will see, a gradual inflow and outflow of speculators may lead to alternating periods of high and low volatility.

Let us turn to the details of our model. We assume that the stock market's log fundamental value follows a random walk. To be precise, the stock market's log fundamental value in period $t + 1$ is given by

$$F_{t+1} = F_t + n_{t+1}. \quad (1)$$

Fundamental shocks n_t are normally distributed with mean zero and constant standard deviation σ^n . Note that fundamental shocks represent the only extrinsic force in our model.

Following Day and Huang (1990), a market maker adjusts the stock price using a log-linear price-adjustment rule, i.e.

$$P_{t+1} = P_t + a \sum_{i=1}^{N_t} D_{t,i}, \quad (2)$$

where P_t stands for the log of the stock price at time t , a is a positive price adjustment parameter, N_t represents the number of active speculators, and $D_{t,i}$ denotes the order placed by an active speculator.¹ Accordingly, the market maker increases the stock price if buying exceeds selling, and vice versa.

As in Chiarella and Iori (2002); Chiarella et al. (2009b); Pellizzari and Westerhoff (2009), the order placed by an active speculator i depends on a linear blend of technical and fundamental trading signals. In addition, speculator i 's trading behavior is influenced by the arrival of new information. The order placed by speculator i is formalized as

$$D_{t,i} = b_{t,i}(P_t - P_{t-1}) + c_{t,i}(F_t - P_t) + d_{t,i}(F_t - F_{t-1}). \quad (3)$$

The first component of (3) reflects speculator i 's technical trading (Murphy, 1999). Speculator i receives a buying (selling) signal if prices increase (decrease). Parameter $b_{t,i} > 0$ defines how strongly speculator i reacts to the price signal. The second component of (3) formalizes speculator i 's fundamental trading (Graham and Dodd, 1951). Since $c_{t,i}$ is a positive reaction parameter, speculator i obtains a buying signal when the market is undervalued and a selling signal when it is overvalued. The third component of (3) indicates that speculator i also reacts to the arrival of new information (Pearce and Roley, 1985). Positive news stimulates buying orders while negative news triggers selling orders. Of course, reaction parameter $d_{t,i}$ is also positive. In the following, we assume that reaction parameters $b_{t,i}$, $c_{t,i}$ and $d_{t,i}$ are uniformly distributed, i.e. $b_{t,i} \sim U(b - \beta, b + \beta)$, $c_{t,i} \sim U(c - \gamma, c + \gamma)$ and $d_{t,i} \sim U(d - \delta, d + \delta)$, with $b > \beta \geq 0$, $c > \gamma \geq 0$ and $d > \delta \geq 0$. Hence, all speculators follow their own time-varying trading strategy.

Before we continue with the description of our approach, let us derive a convenient model property. First, inserting (3) in (2) reveals that

$$\begin{aligned} P_{t+1} &= P_t + a \sum_{i=1}^{N_t} (b_{t,i}(P_t - P_{t-1}) + c_{t,i}(F_t - P_t) + d_{t,i}(F_t - F_{t-1})) \\ &= P_t + a(P_t - P_{t-1}) \sum_{i=1}^{N_t} b_{t,i} + a(F_t - P_t) \sum_{i=1}^{N_t} c_{t,i} + a(F_t - F_{t-1}) \sum_{i=1}^{N_t} d_{t,i}. \end{aligned} \quad (4)$$

Recall next that the sum of independently uniformly distributed random variables follows a uniform sum distribution.² Defining $\sum_{i=1}^{N_t} b_{t,i} = B_t$, $\sum_{i=1}^{N_t} c_{t,i} = C_t$ and $\sum_{i=1}^{N_t} d_{t,i} = D_t$ yields $B_t \sim \text{USD}(N_t, \{b - \beta, b + \beta\})$, $C_t \sim \text{USD}(N_t, \{c - \gamma, c + \gamma\})$ and $D_t \sim \text{USD}(N_t, \{d - \delta, d + \delta\})$, respectively, and we can rewrite (4) as

$$P_{t+1} = P_t + a(B_t(P_t - P_{t-1}) + C_t(F_t - P_t) + D_t(F_t - F_{t-1})). \quad (5)$$

Apparently, our setup has the convenient property that it is not necessary to evaluate the trading rules of all N_t active speculators, each consisting of three different components, to simulate its dynamics. We simply need to generate three uniform sum distributed random variables. Moreover, the means and variances of the three uniform sum distributed random variables are given by $\mu_B = bN_t$, $\mu_C = cN_t$, $\mu_D = dN_t$, $\sigma_B^2 = \frac{N_t}{3}\beta$, $\sigma_C^2 = \frac{N_t}{3}\gamma$ and $\sigma_D^2 = \frac{N_t}{3}\delta$, respectively. In particular, note that the means of the uniform sum distributed random variables increase with the number of active speculators, i.e. if there is an inflow of speculators, there is, on average, stronger trend extrapolation trading, a stronger mean reversion behavior and

¹ For notational convenience, we use the index $i = 1, 2, \dots, N_t$ to refer to active speculators in trading period t . Clearly, index i does not stand for a specific speculator.

² A uniform sum distribution, also called an Irwin-Hall distribution, approaches a normal distribution as the number of added random variables increases.

a stronger reaction to new information. It is easily imaginable that this will have a destabilizing impact on the model dynamics, at least for some parameter constellations. We also remark that the variances of the three uniform sum distributed random variables vanish if β , γ and δ approach zero.

Let us now return to our model. At the beginning of each trading period, speculators decide whether to enter the stock market. We assume that speculators' probabilistic market entry decisions are influenced by two socio-economic principles. In line with empirical evidence reported by [Hong et al. \(2004, 2005\)](#), [Brown et al. \(2008\)](#) and [Shiller \(2015\)](#), speculators regard a stock market as increasingly attractive when more speculators are already active. A similar herding perspective is taken in [Iori \(1999, 2000\)](#). Moreover, speculators' market entry decisions also depend on market circumstances: the higher the stock market risk, the less attractive the stock market appears to be. As in [Alfi et al. \(2009b, 2009c, 2009a\)](#), stock market risk is represented by the stock market's volatility

$$V_t = mV_{t-1} + (1 - m)(P_t - P_{t-1})^2, \quad (6)$$

where $0 \leq m < 1$ is a memory parameter controlling how strongly current and past price changes affect volatility. We summarize both socio-economic principles by the following relative fitness function

$$A_t = hN_{t-1} - vV_t, \quad (7)$$

where h and v are positive parameters. Accordingly, market participation is regarded as increasingly attractive the more speculators are active in the market and less attractive the higher the stock market's past volatility.

We use exponential replicator dynamics ([Hofbauer and Sigmund, 1988](#); [Hofbauer and Weibull, 1996](#)) to model speculators' probabilities of entering the market. The probability that a speculator will enter the stock market can thus be written as

$$W_t = \frac{W_{t-1}}{W_{t-1} + (1 - W_{t-1}) \exp[-\lambda A_t]}, \quad (8)$$

where parameter $\lambda > 0$ denotes speculators' intensity of choice. Note that the exponential replicator dynamics term has three important properties. First, speculators' probabilities of entering the market depend positively on the stock market's relative fitness. The higher the stock market's relative fitness, the more probable it is that speculators will enter the market. Second, an increase in λ implies that speculators react more sensitively to the stock market's relative fitness. If speculators' intensity of choice approaches zero, they have a 50% probability of entering the market. If speculators' intensity of choice goes to plus infinity, the probability that they will enter the market is either 100% if the herding component dominates the risk component or zero percent otherwise. Third, market entry probabilities display a mild form of inertia. If W_{t-1} is either close to zero or close to one, market entry probabilities depend less strongly on the stock market's relative fitness.³

Obviously, the number of active speculators is binomially distributed, i.e.

$$N_t \sim B(N, W_t), \quad (9)$$

where $N > 0$ denotes the total number of speculators. As is well known, the mean and variance of the number of active speculators are given by NW_t and $NW_t(1 - W_t)$, respectively.

3. Analysis of the model's deterministic skeleton

In this section, we explore the model's deterministic skeleton. In [Section 3.1](#), we derive the model's dynamical system and analyze under which conditions the model's steady states are locally asymptotically stable. In [Section 3.2](#), we introduce a base parameter setting to explain the functioning of our deterministic model. In [Section 3.3](#), we study how the model's parameters affect its global dynamics. In [Section 3.4](#), we show that the model may also give rise to coexisting attractors and produce bubbles and crashes. In [Section 3.5](#), we briefly discuss the dynamics of our model for an alternative parameter setting.

3.1. Dynamical system, steady states and local stability

By setting $\beta = \gamma = \delta = \sigma_n = 0$ and introducing the auxiliary variable $\tilde{P}_t = P_{t-1}$, we can summarize our model by the four-dimensional nonlinear map

$$X : \begin{cases} P_{t+1} = P_t + aN_t \{b(P_t - \tilde{P}_t) + c(F - P_t)\} \\ \tilde{P}_{t+1} = P_t \\ V_{t+1} = mV_t + (1 - m)(P_{t+1} - P_t)^2 \\ N_{t+1} = N \frac{N_t}{N_t + (N - N_t) \exp[-\lambda(hN_t - vV_{t+1})]} \end{cases}. \quad (10)$$

³ [Dindo and Tuinstra \(2011\)](#) provide a deeper discussion of exponential replicator dynamics. Further economic examples in this direction include [Bischi et al. \(2015\)](#), [Kopel et al. \(2014\)](#) and [Schmitt et al. \(2017\)](#).

Since we set the scaling parameters a and λ to 1, the dynamics depends solely on seven parameters: b, c, N, h, v, F and m .⁴ Straightforward computations reveal that our dynamical system may give rise to two steady states, namely

$$X_1^* = (P^*, \tilde{P}^*, V^*, N^*) = (F, F, 0, N) \quad (11)$$

and

$$X_2^* = (P^*, \tilde{P}^*, V^*, N^*) = (P^*, \tilde{P}^*, 0, 0). \quad (12)$$

As can be seen, the steady-state price is given by the fundamental value, the stock market's volatility is zero and all speculators are active at X_1^* , while X_2^* has zero active speculators, an indeterminate price and also a volatility of zero. Since the second steady state is economically uninteresting, we will focus now on X_1^* , which we also call the fundamental steady state of our model.

To determine the stability of X_1^* , we derive the characteristic polynomial from the Jacobian matrix of (10), i.e.

$$J(X_1^*) = \begin{pmatrix} 1 + bN - cN & -bN & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & m & 0 \\ 0 & 0 & 0 & e^{-hN} \end{pmatrix}, \quad (13)$$

and obtain

$$(e^{-hN} - z)(m - z)(z^2 + z(cN - bN - 1) + bN) = 0. \quad (14)$$

The steady state is locally asymptotically stable if the eigenvalues of the polynomial are less than one in modulus (see, e.g. Gandolfo, 2009 or Medio and Lines, 2001). It is easy to see from (14) that the first two eigenvalues are given by $z_1 = e^{-hN}$ and $z_2 = m$. Since we assume that $h > 0, N > 0$ and $0 \leq m < 1$, we have $|z_1, z_2| < 1$. However, the eigenvalues of $z^2 + z(cN - bN - 1) + bN$ are less than one in modulus if and only if

$$c > 0, \quad (15)$$

$$c < c_c = \frac{2}{N} + 2b \quad (16)$$

and

$$b < b_c = \frac{1}{N} \quad (17)$$

simultaneously apply. Recall that c is a positive reaction parameter, which is why condition (15) is always fulfilled. According to (16), the fundamental steady state becomes unstable if c crosses c_c , a situation which leads to a flip bifurcation and the onset of a period-two cycle. If b exceeds its critical value b_c , condition (17) is violated, which is associated with a Neimark–Sacker bifurcation, i.e. the emergence of a cyclical motion. In economic terms, these two conditions imply that the steady state becomes unstable if speculators react to market misalignments or to price trends too strongly. Note that (16) and (17) also depend on the total number of speculators. Hence, the stock market also becomes unstable if N increases.

To visualize our analytical results, we depict in Fig. 1 combinations of b and c for which the model's fundamental steady state is locally asymptotically stable. Since the two black lines represent stability conditions (16) and (17), the model's fundamental steady state is always locally asymptotically stable for parameter combinations within these two lines. As indicated by the arrows, an increase in parameter b may cause a loss of stability via a Neimark–Sacker bifurcation while an increase in parameter c may cause a loss of stability via a flip bifurcation. Moreover, the gray shaded area indicates the parameter space for which the steady state becomes unstable if the number of speculators increases from N to N' .

For the sake of completeness, note that the Jacobian matrix of (10) at the second steady state is given by

$$J(X_2^*) = \begin{pmatrix} 1 & 0 & 0 & c(F - P^*) \\ 1 & 0 & 0 & 0 \\ 0 & 0 & m & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (18)$$

from which the characteristic polynomial

$$-z(m - z)(1 - z)^2 = 0 \quad (19)$$

can be derived. Therefore, we obtain $z_1 = 0, z_2 = m$ and $z_{3,4} = 1$, which implies that the second steady state is always unstable.

⁴ Note that N_t is given by NW_t , i.e. we focus in this section on the mean dynamics of the active number of speculators. Such a procedure is common in this line of research, see, for instance, Sandholm (2015). In fact, numerical experiments confirm that our analytical results predict the properties of the non-mean dynamics quite well. The same is true for the simulation results presented in Section 3. For simplicity, we call N_t the active number of speculators.




Fig. 1. Combinations of b and c for which the fundamental steady state is locally asymptotically stable. The two black lines represent stability conditions (16) and (17), i.e. $c = \frac{2}{N} + 2b$ and $b = \frac{1}{N}$, respectively.

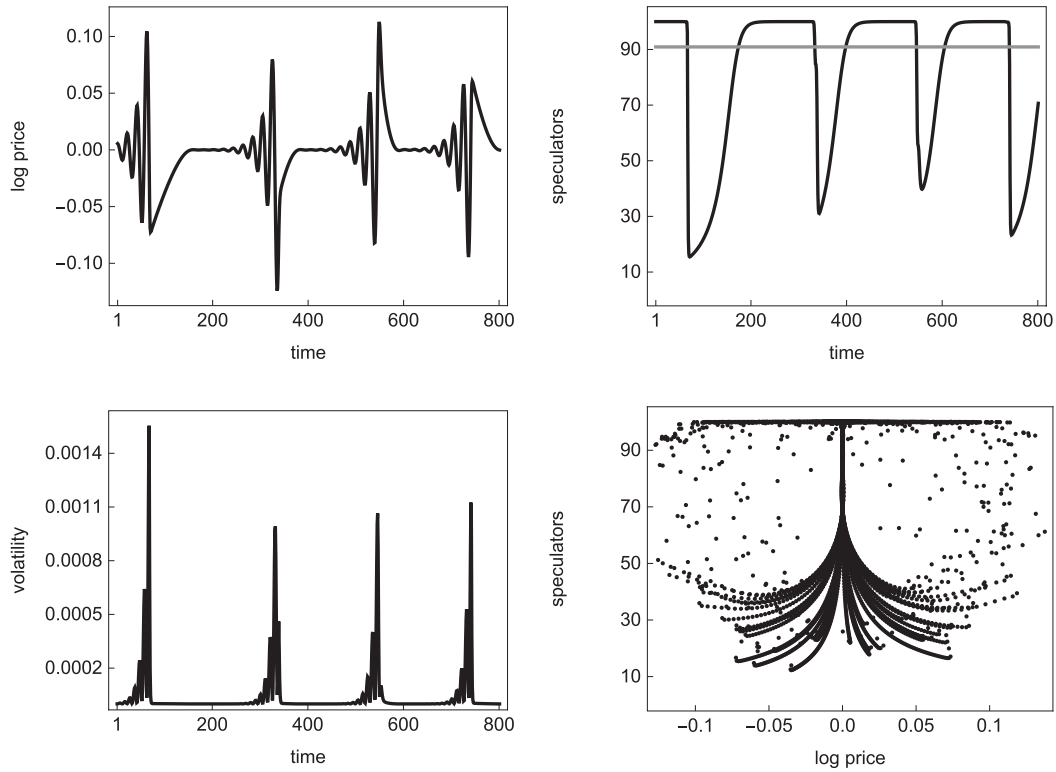


Fig. 2. Dynamics of the deterministic model for our base parameter setting. The panels show the evolution of the log price, the number of active speculators, the stock market's volatility and the number of active speculators versus the log price, respectively. The underlying parameter setting is given in Section 3.2.

3.2. Base parameter setting and functioning of the model

To be able to explain the functioning of our deterministic model, represented by the four-dimensional nonlinear map (10), we make use of the following parameter setting: $b = 0.011$, $c = 0.001$, $N = 100$, $h = 0.001$, $v = 2000$, $m = 0.25$ and $F = 0$. Recall that the scaling parameters a and λ are set equal to 1. Accordingly, we have $b > 1/N$, which implies that our fundamental steady state, i.e. $X_1^* = (F, F, 0, N)$, is unstable and that its instability is due to a Neimark–Sacker bifurcation. Fig. 2 shows a representative simulation run for 800 periods (after omitting a longer transient period). The first three panels depict the evolution of the log price, the number of active speculators and the stock market's volatility, respectively, while the fourth panel presents the number of active speculators versus the log price.

Obviously, our model is able to produce intricate dynamics, in particular, alternating periods with low and high volatility. In a nutshell, the working of our deterministic model may be summarized as follows. Note first that the gray line in the top right panel of Fig. 2 indicates the threshold for the number of active speculators for which the dynamics of our deterministic model becomes unstable, i.e. $N_c = \frac{1}{b} \approx 90.91$. If the number of active speculators is below N_c , the market is stable and prices converge slowly towards their fundamental value. Since volatility is relatively low during these periods, speculators' herding behavior dominates their risk aversion and more and more speculators enter the stock market. However, the picture changes if the number of active speculators exceeds N_c . Once $N_t > N_c$, the model dynamics becomes unstable, i.e. price fluctuations are characterized by oscillations with an increasing amplitude. As a result, volatility increases up to the point where speculators' risk aversion offsets their herding behavior. Speculators then start to exit the market and initiate a new period of relative stability. During the 800 depicted time steps, we witness four marked volatility outbursts. The strange attractor, visible in the bottom right panel, illustrates the complexity of the model dynamics.

3.3. The impact of the model's parameters on its global dynamics

To demonstrate how the global dynamics of our deterministic model depends on its parameters, we present a number of simulations in this section. Fig. 3 contains examples of how parameters b , c and N may influence the model dynamics. The first, second and third rows depict bifurcation diagrams for $0 < b < 0.020$, $0 < c < 0.004$ and $80 < N < 120$, respectively. While the left side presents their effect on log prices, the right side shows how they affect the number of active speculators. As predicted by our analytical results, we have $P^* = F = 0$ and $N^* = N = 100$ for $b < b_c = \frac{1}{N} = 0.01$. As soon as b exceeds this critical value, the fundamental steady state loses its stability and endogenous dynamics emerges. While the two top panels reveal that our model dynamics becomes unstable if technical trading is too aggressive, the second row shows that a stronger fundamental trading reduces the amplitude of price fluctuations. The two panels at the bottom also confirm our previous analytical results. The fundamental steady state loses its stability at $N = N_c = \frac{1}{b} \approx 90.91$, after which the amplitude of price dynamics and speculators' market entry and exit behavior increases with N .

In Fig. 4, we show bifurcation diagrams for $0 < h < 0.004$, $1000 < v < 3000$ and $0 < m < 1$. The left panels reveal again how log prices react to an increase in parameters h , v and m and the right panels illustrate how this affects the number of active speculators. It can be seen from the two top panels that a stronger herding behavior increases the amplitude of price fluctuations as well as fluctuations in the number of active speculators. In contrast, price dynamics is less pronounced if speculators show a stronger risk-sensitive behavior. However, the amplitude of price fluctuations also increase with m . Of course, this also causes higher fluctuations in speculators' market entry and exit behavior.

3.4. Special features: coexisting attractors and bubbles and crashes

As is well known, nonlinear dynamical systems may give rise to a number of complicated dynamic phenomena. As indicated by the (asymmetric) bifurcation diagram in the top left panel of Fig. 3, our model may also produce coexisting attractors. The left panels of Fig. 5 provide an example of this outcome. The panels show from top to bottom the evolution of the log price, the number of active speculators and the number of active speculators versus the log price, respectively. We use the base parameter setting, except that we set $b = 0.0155$ (instead of $b = 0.011$). Moreover, the dynamics is depicted for two different initial values, represented in black and red. The top left panel of Fig. 5 shows that one set of initial conditions produces a sequence of bull markets while the other set of initial conditions produces a sequence of bear markets. As it turns out, the evolution of the number of active speculators is identical for both price trajectories. The bottom left panel of Fig. 5 reveals that the bull market dynamics is intricately intertwined with the bear market dynamics. In the absence of exogenous shocks, the model generates either persistent bull or persistent bear market dynamics. However, it is clear that the addition of some exogenous noise may easily push the dynamics from one attractor to the other. The overall dynamics is then characterized by erratic switches between bull and bear market dynamics.

Interestingly, the right panels of Fig. 5 demonstrate that our model is able to generate endogenous boom-bust dynamics. Once more we use the base parameter setting but assume that $c = 0.00004$ (instead of $c = 0.001$). The top right panel of Fig. 5 demonstrates that a boom period can be followed by another boom period but also by a severe crash (and since the model is symmetric, the same is true the other way around). The right panel in the center of Fig. 5 indicates that the dynamics is again driven by speculators' market entry and exit behavior. Since speculators' mean reversion trading is now relatively weak, we do not observe repeated volatility outbursts but the emergence of pronounced and lasting bull and bear market dynamics. It is easy to check that the instability of the model's fundamental steady state is again caused by a Neimark–Sacker bifurcation. However, the strange attractor visible in the bottom right panel of Fig. 5 reveals that the model dynamics is quite complicated for the underlying parameter setting.

3.5. Alternative parameter setting

So far, we focused mainly on the Neimark–Sacker bifurcation scenario. We now turn briefly to the flip bifurcation scenario. Fig. 6 is based on an alternative parameter setting: $a = 1$, $b = 0.005$, $c = 0.0301$, $N = 100$, $h = 0.005$, $v = 10$, $m = 0.1$, $\lambda = 1$ and $F = 0$. Hence, the model's fundamental steady state is unstable due to a flip bifurcation. Since the flip bifurcation occurs at $c = c_c = 0.03$, the onset of endogenous dynamics, initially in the form of a period-two cycle and then in the form




Fig. 3. The impact of b , c and N . The first, second and third rows show bifurcation diagrams for parameters b , c and N , respectively. While the left side presents their effect on log prices, the right side depicts how they affect the number of active speculators. Parameters are as in our base parameter setting.

of more complicated dynamics, is caused by speculators' excessively aggressive fundamental trading. The first three panels of Fig. 6 show the evolution of the log price, the number of active speculators and the stock market's volatility for 800 periods, respectively. As can be seen, the model is also able to produce volatility clustering for the alternative parameter setting. The reason for this is similar to before. If the number of active speculators is low, the market is stable. Since speculators' herding behavior outweighs their risk aversion, they quickly enter the stock market. This process renders the dynamics unstable and we observe increasing (improper) oscillations. Eventually, the associated increase in stock market risk makes the stock market become unattractive. Speculators exit the stock market and there is a brief period of market stability in which the price approaches its fundamental value. Then, the process repeats itself, albeit in an intricate manner. This is also confirmed by the right panel in the center of Fig. 6, which presents the corresponding dynamics in (N_t, P_t) space.

The bottom two panels of Fig. 6 show bifurcation diagrams for parameter c . The left panel reveals how the log price reacts to increasingly aggressive fundamental trading, while the right panel shows how this affects the number of active

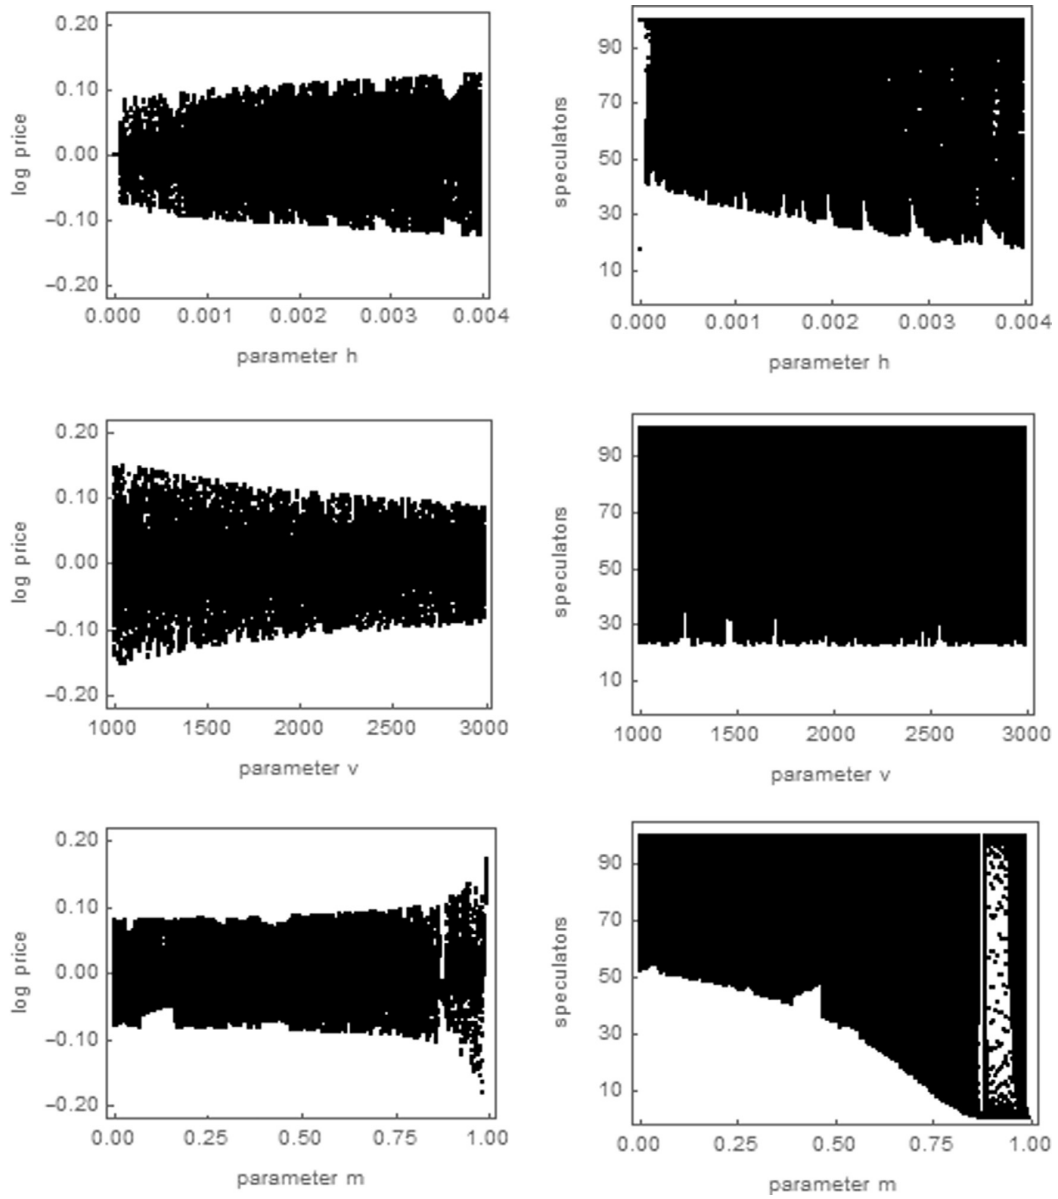


Fig. 4. The impact of h , v and m . The first, second and third rows show bifurcation diagrams for parameters h , v and m , respectively. While the left side presents their effect on log prices, the right side depicts how they affect the number of active speculators. Parameters are as in our base parameter setting.

speculators. As predicted by our analytical results, the model dynamics approaches the fundamental steady state for $c < 0.03$. However, a flip bifurcation occurs at $c = c_c = 0.03$. Within a small parameter range, the dynamics is then characterized by a period-two cycle. Afterwards, we observe the start of more complex dynamics, as already depicted in the first four panels of Fig. 6. Note furthermore that the amplitude of the price dynamics increases with parameter c . Here we have an example where excessively aggressive mean reversion trading by speculators leads to a destabilization of the stock market. As price fluctuations increase, the number of active speculators also displays more pronounced fluctuations.

4. Stochastic dynamics

In Section 3, we show that the deterministic version of our simple agent-based financial market model is – at least in a qualitative sense – able to produce bubbles and crashes, excess volatility, extreme price changes, complex price dynamics

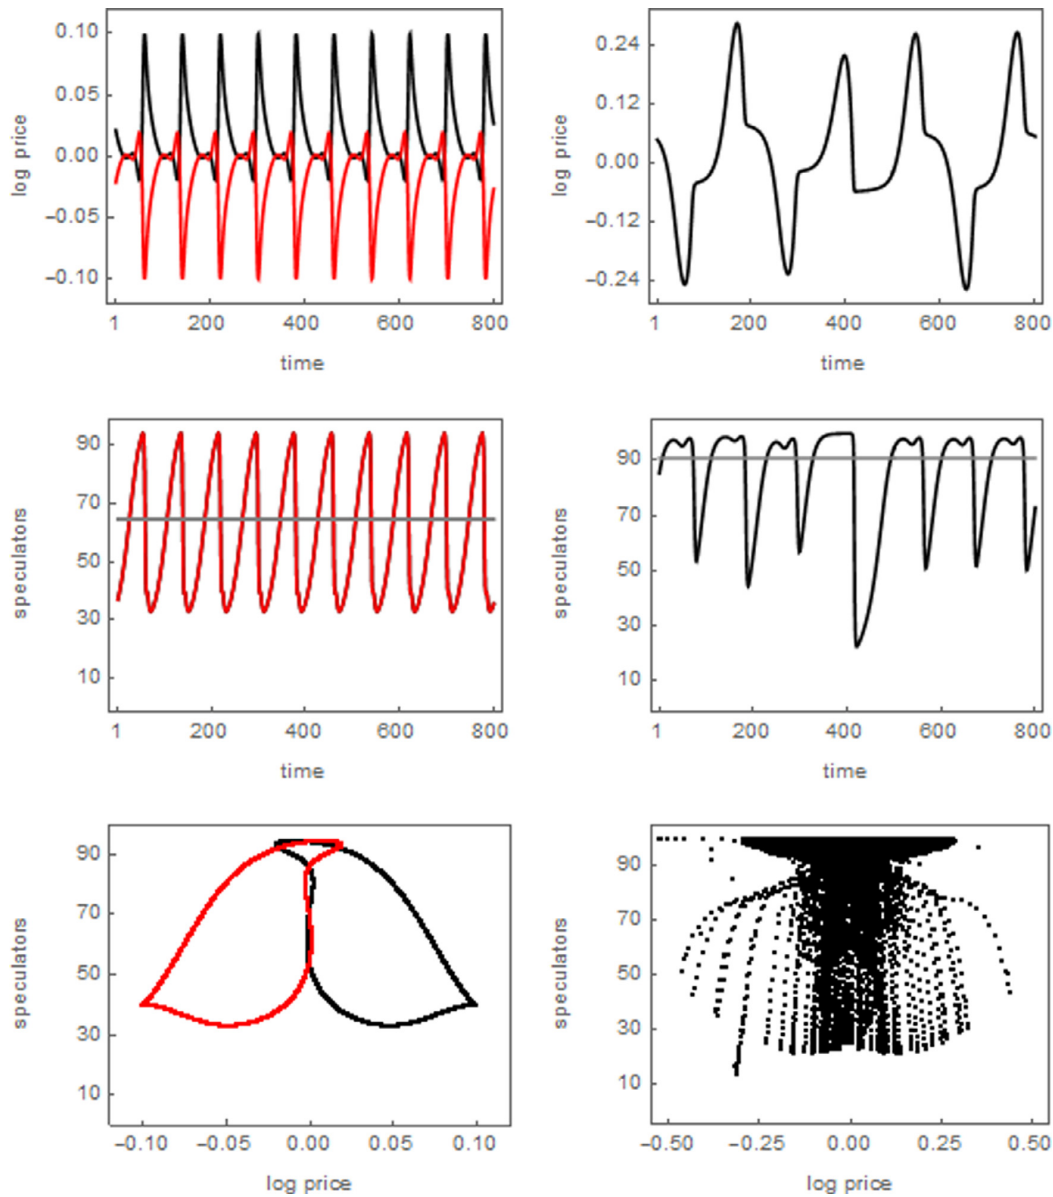


Fig. 5. Emergence of bubbles and crashes. The left panels show for two different sets of initial conditions (black and red) the evolution of log prices, the number of active speculators and the number of active speculators versus the log price for our base parameter setting, except for $b = 0.0155$. The right panels show the same for a single initial condition and our base parameter setting, except for $c = 0.00004$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and volatility clustering. In this section, we go one step further and demonstrate that the stochastic version of our model may also replicate a number of key statistical properties of actual stock markets in finer detail. In [Section 4.1](#), we first review the stylized facts of stock markets. Then, we discuss the dynamics of our stochastic model and explain its functioning in [Section 4.2](#). In [Section 4.3](#), we explore the wealth dynamics of individual speculators.

4.1. Stylized facts of stock markets

As is well known, the dynamics of stock markets is characterized by bubbles and crashes, excess volatility, fat-tailed return distributions, serially uncorrelated returns and volatility outbursts. Moreover, trading volume is persistent and cor-

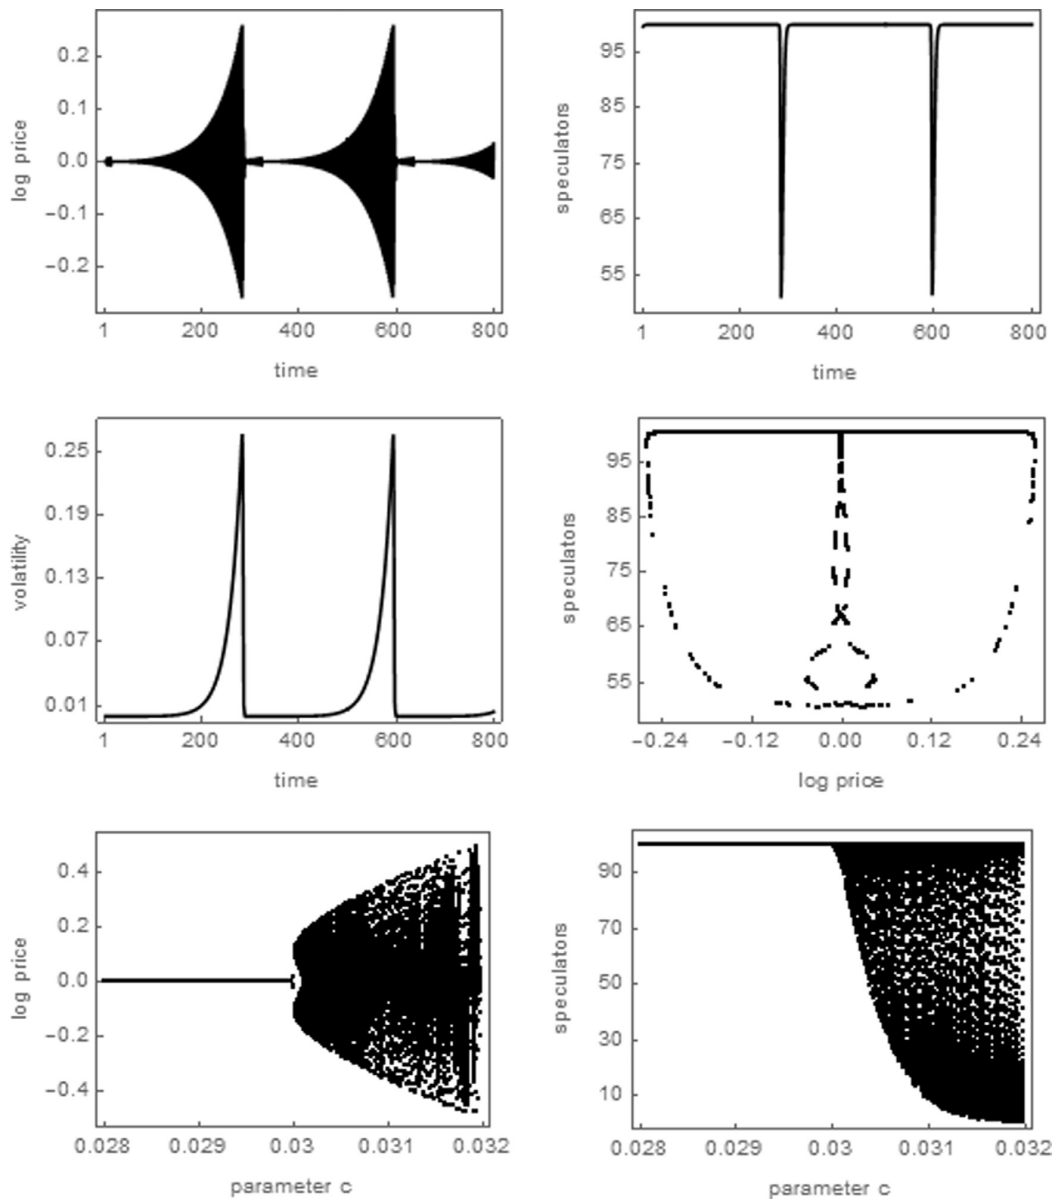


Fig. 6. Dynamics of the deterministic model for an alternative parameter setting. The panels show the evolution of the log price; the number of active speculators; the stock market's volatility; and the number of active speculators versus the log price. The last two panels show bifurcation diagrams for the log price and the number of active speculators with respect to parameter c . The underlying parameter setting is given in Section 3.5.

related with volatility. The boom-and-bust behavior of stock markets and their volatile nature is discussed thoroughly in Shiller (2015). Moreover, Mantegna and Stanley (2000), Cont (2001) and Lux and Ausloos (2002) provide excellent surveys about the statistical properties of financial markets. For illustrative reasons, we focus below on the behavior of the Dow Jones Industrial Average between 1981 and 2016, as depicted in Fig. 7. The underlying data set comes from Thomson Reuters Datastream and contains about 9000 daily observations. The top left panel of Fig. 7 shows the development of the Dow Jones Index. Despite its long-run upwards movement, a number of severe crashes can be spotted. For instance, the Dow Jones Index witnessed dramatic depreciations around 2001 and 2007. The top right panel of Fig. 7 presents the corresponding return time series. Overall, the Dow Jones Index may be regarded as quite volatile. Just to give one example, the standard deviation of the return time series is about 0.011. Furthermore, there are several larger returns visible and calm periods obviously alternate with turbulent periods.

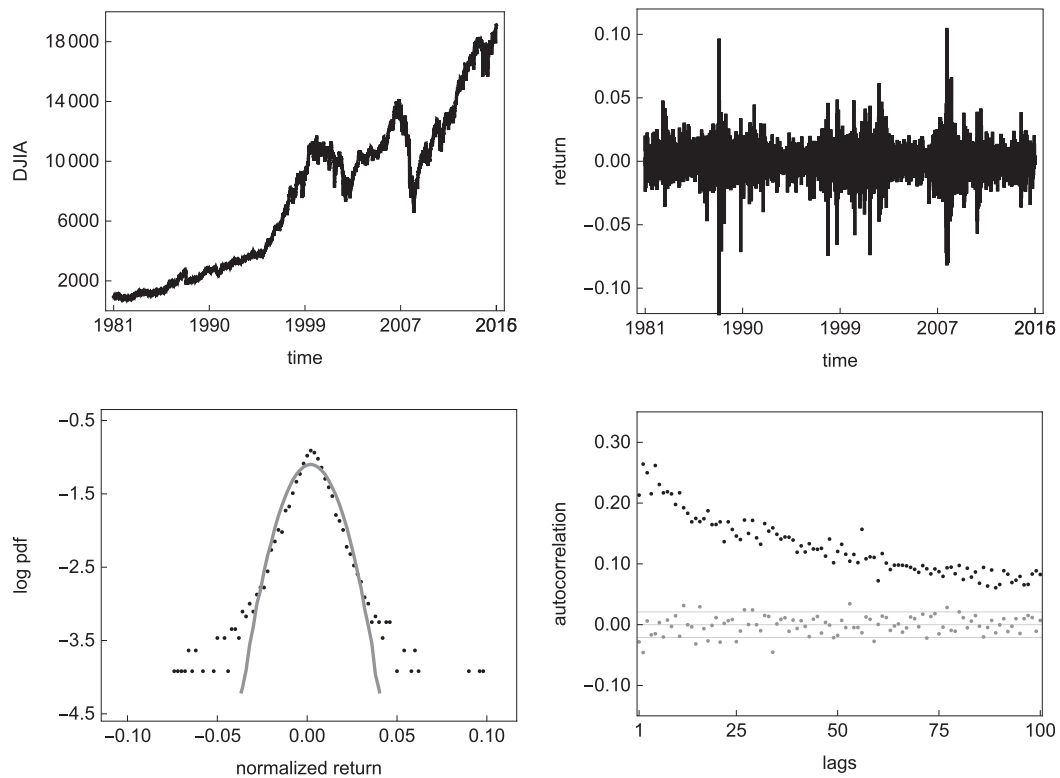


Fig. 7. Properties of the Dow Jones Industrial Average. The panels show the evolution of the Dow Jones Index between 1981 and 2016, the corresponding returns, the log probability density function of normalized empirical returns (black) and standard normally distributed returns (gray) and the autocorrelation function of raw returns (gray) together with the autocorrelation function of absolute returns (black).

The bottom right panel of [Fig. 7](#) compares the log probability density function of normalized Dow Jones Index returns (black dots) with standard normally distributed returns (gray line). Apparently, the distribution of the returns of the Dow Jones Index contains more probability mass in the center, less probability mass in the shoulders, and again more probability mass in the tails than warranted by a normal distribution with identical mean and standard deviation. Since the kurtosis of empirical returns is 42.66 and thus much larger than the kurtosis of a normal distribution, namely 3, there is clear evidence of excess kurtosis. However, the tail index provides a more reliable measure to quantify the fat-tail property of the distribution of returns ([Gopikrishnan et al., 1999](#); [Lux, 1996](#); [Lux and Alfarano, 2016](#)). Using the largest 5% of the observations, the Hill tail index estimator indicates a typical tail index of about 3.02 for this time series, suggesting that the fourth moment of the distribution of the returns does not exist. The bottom right panel of [Fig. 7](#) shows the autocorrelation coefficients of absolute returns (black dots) and raw returns (gray dots) for the first 100 lags, together with their 95% confidence bands (thin gray lines). The absence of autocorrelation of raw returns demonstrates that the evolution of the Dow Jones Index is hardly possible to predict, i.e. that its path is close to a random walk. In contrast, the autocorrelation coefficients of absolute returns are highly significant, implying a temporal persistence of volatility for more than 100 days. Finally, we remark that trading volume shows clear signs of long-memory effects and is highly correlated with volatility. Due to missing data availability, these properties are not depicted here. However, see [Brock and LeBaron \(1996\)](#), [Cont \(2001\)](#) and [Schmitt and Westerhoff \(2014\)](#) for a deeper empirical account.

4.2. Properties and functioning of the stochastic model

In the last couple of years, considerable progress has been made in estimating agent-based financial market models, see, e.g. [Alfarano et al. \(2005\)](#); [Amilon \(2008\)](#); [Boswijk et al. \(2007\)](#); [Chiarella et al. \(2014\)](#); [Hommes and in 't Veld \(2017\)](#). One powerful method to estimate such models is given by the method of simulated moments, which seeks to align a selection of empirical moments, i.e. certain summary statistics which quantify the stylized facts of financial markets, with model generated moments. Contributions in this direction include [Gilli and Winker \(2003\)](#), [Winker et al. \(2007\)](#), [Franke \(2009\)](#) and [Franke and Westerhoff \(2012\)](#). Unfortunately, the large number of parameters of our model prevents us from using this method (which requires a multi-dimensional grid search in parameter space). Instead, we rely on a more informal calibration

approach. After a tedious and time-consuming trial-and-error exercise, we can at least show that our model has some ability to match the stylized facts of stock markets.

To be precise, we use the following parameter setting to discuss the dynamics of our stochastic model: $a = 1$, $b = \beta = 0.0001$, $c = \gamma = 0.000005$, $d = \delta = 0.01$, $h = 0.00008$, $v = 130$, $m = 0.99$, $\lambda = 1$, $\sigma^n = 0.005$ and $N = 500$. Fig. 8 depicts the outcome of a typical simulation run with 9,000 observations. The top left panel of Fig. 8 displays the evolution of the stock price. While the stock price fluctuates quite erratically, there are a number of stronger price appreciations and depreciations, resembling the boom-and-bust behavior of the Dow Jones Index (since the fundamental value follows a random walk in our model, there is no long-run upwards trend in the simulated stock price dynamics). The top right panel of Fig. 8 depicts the corresponding return time series. The standard deviation of simulated returns is given by 0.011, i.e. our model matches the average volatility of the Dow Jones Index quite well. Moreover, the standard deviation of the fundamental value only amounts to 0.005. Hence, returns are roughly twice as volatile as justified by changes in the fundamental value. This panel also reveals that there are a number of larger returns as well as occasional volatility outbursts.

The bottom left panel of Fig. 8 relates the log probability density function of normalized returns (black dots) with the log probability density function of standard normally distributed returns (gray line). As can be seen, the distribution of simulated returns possesses more probability mass in the center, less probability mass in the shoulders, and again more probability mass in the tails than warranted by a normal distribution with identical mean and standard deviation. The fat-tail property is also confirmed by estimates of the kurtosis for which we obtain a value of 4.62. While this value indicates excess kurtosis, it should be noted that it is much lower than the value we observe for the Dow Jones Index. Estimates of the tail index point in the same direction. For the simulated time series, the Hill tail index estimator produces a value of 4.56, implying that the fourth moment of the distribution of returns exists. Clearly, simulated returns have more probability mass in the tails of their distribution than normally distributed returns but less than actual returns.⁵ The bottom right panel of Fig. 8 presents the autocorrelation coefficients of absolute returns (black dots) and raw returns (gray dots) for the first 100 lags. Raw returns are serially uncorrelated, i.e. also simulated prices are hardly possible to predict. The autocorrelation coefficients of absolute returns are highly significant, revealing strong evidence of volatility clustering.

Fig. 9 presents some properties of trading volume and how it relates to volatility. The panels show the evolution of trading volume, the autocorrelation function of trading volume, the return dynamics and the cross-correlation function of trading volume and absolute returns, respectively. Following Schmitt and Westerhoff (2014), we assume that all speculators trade directly with the market maker, i.e. speculators do not trade with other speculators. Trading volume can then be defined by $TV_t = \sum_{i=1}^{N_t} |D_{t,i}|$. As can be seen, trading volume is highly persistent, i.e. autocorrelation coefficients of trading volume are positive and decay rather slowly. Moreover, trading volume is positively correlated with volatility. While there is a strong contemporaneous correlation between trading volume and volatility, lagged correlations are rather low, as is the case for real markets (Brock and LeBaron, 1996).

Overall, we may thus conclude that the stochastic version of our simple agent-based financial market model is able to replicate key empirical regularities of actual stock markets. To explain its functioning in more detail, we continue with the simulation run depicted in Fig. 8 but focus our attention on a shorter time window. The four panels of Fig. 10 show from top left to bottom right the evolution of stock prices (black line) and fundamental values (gray line), the corresponding returns, the stock market's volatility and the number of active speculators between periods 6150 and 7650. During this time period, there are three pronounced volatility outbursts. Note also that volatility tends to increase with the number of active speculators. Accordingly, the functioning of our stochastic model may be understood as follows. Suppose that stock market volatility is low. In such a situation, speculators' herding behavior dominates their risk aversion. Consequently, more and more speculators enter the stock market and volatility picks up. Eventually, speculators' risk aversion offsets their herding behavior. As the number of speculators declines, the stock market becomes more stable. However, this leads directly to the next market entry wave and to another high volatility episode. Of course, higher stock market participation also drives trading volume up, i.e. stock market participation, trading volume and volatility are correlated.

It is interesting to note that the functioning of the stochastic version of our agent-based model is very similar to the functioning of its deterministic counterpart. In the deterministic setup, endogenous dynamics and volatility outbursts emerge when a model parameter crosses the Neimark–Sacker bifurcation boundary, either because speculators react too strongly to price trends or because there are too many speculators. While the calibrated parameter setting of our stochastic model implies that the fundamental steady state of the corresponding deterministic model is locally stable, the model's cyclical nature prevails. We remark that this phenomenon, i.e. realistic model dynamics for parameter settings in which the fundamental steady state of the model's deterministic skeleton is locally stable, is quite common in this line of research. As it turns out, it is the interplay of nonlinear forces and random elements that causes realistic dynamics. Nevertheless, the analytical and numerical insights we gain from studying the deterministic framework prove to be instrumental in our understanding of the much more complicated stochastic framework.

⁵ Although our model does a fairly good job of matching the stylized facts of stock markets, it produces too few extreme returns. Further experiments (available upon request) reveal that simple model extensions can alleviate this issue. In particular, our model may produce quite realistic tail indices if certain model parameters, such as speculators' reaction to fundamental shocks, are allowed to vary over time – without destroying its ability to match the other stylized facts. Since our main focus is on explaining volatility outbursts in stock markets, we abstain, for simplicity, from such model extensions.

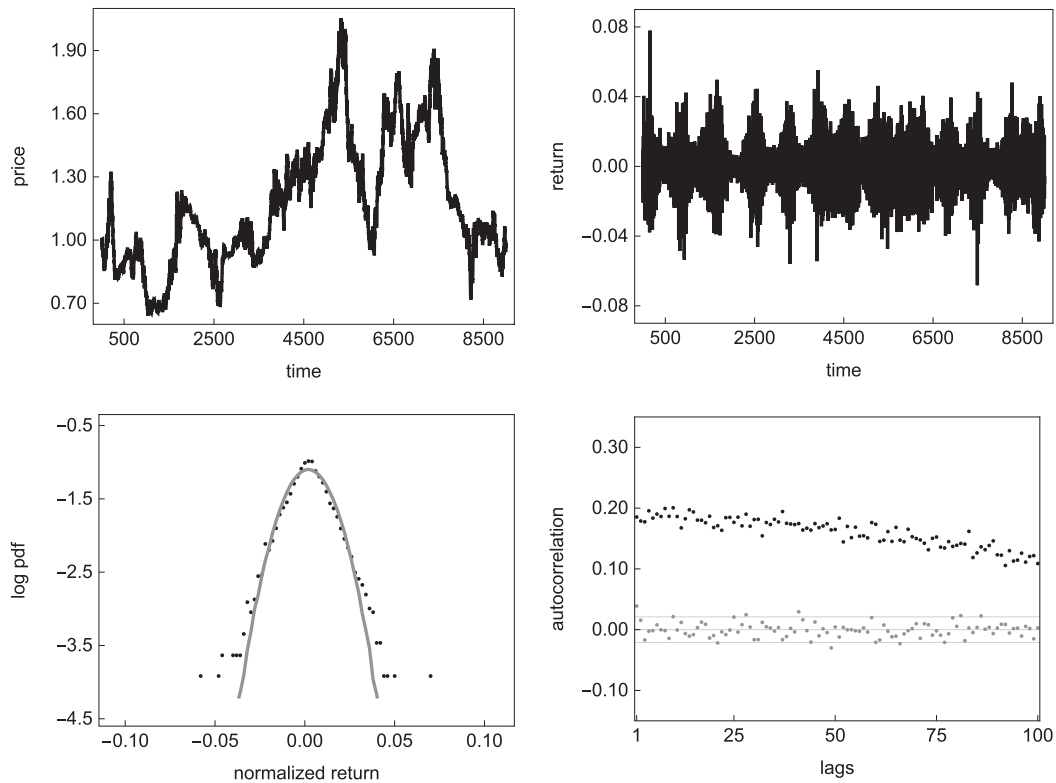


Fig. 8. Properties of the stochastic model. The panels show the evolution of the stock price for 9000 observations, the corresponding returns, the log probability density function of normalized model returns (black) and standard normally distributed returns (gray) and the autocorrelation function of raw returns (gray) together with the autocorrelation function of absolute returns (black). The underlying parameter setting is given in [Section 4.2](#).

Random elements in our model stem from speculators' probabilistic market entry decisions, from their time-varying trading rules and from changes in the fundamental value. As can be seen in the top left panel of [Fig. 10](#), stock prices and fundamental values tend to move in the same direction. However, stock prices may substantially disconnect from fundamental values. This is particularly true if the number of active speculators is rather high or rather low – the stock market then reacts too strongly or too weakly to incoming fundamental shocks. Since positive and negative fundamental shocks are equally likely, stock price changes are basically random. Moreover, in an environment in which stock prices already fluctuate quite erratically, speculators' trend extrapolation behavior does not add predictable structure to the return time series. Whether the technical part of speculators' trading rules produces a buy or sell signal is essentially equally likely. Large price changes occur if a large number of active speculators receive a strong trading signal, either because of significant price trends, pronounced misalignments or distinct fundamental shocks, or if the time-varying reaction parameters of their trading rules suggest aggressive trading. Of course, extreme returns may emerge if these forces act together, i.e. if many speculators act aggressively on heavy trading signals.

4.3. Wealth dynamics of individual speculators

Stock price changes induce diverging wealth dynamics among speculators relying on heterogeneous trading rules. In the long run, the evolutionary pressure originating from such wealth dynamics may act as a natural selection device among speculator types. While some speculator types may turn out to be successful and survive evolutionary competition, other speculator types may fail and become extinct. The so-called market selection hypothesis ([Blume and Easley, 1992](#)) therefore predicts that heterogeneity among speculators can only be a short-run phenomenon. Note that the implications of this hypothesis may be far-reaching. In particular, naive speculator types who persistently lose wealth may eventually vanish from the market, implying that stock prices would then only be subject to smart speculator types who manage to make a profit. However, [Anufriev and Dindo \(2010\)](#), [Bottazzi and Dindo \(2014\)](#) and [Bottazzi et al. \(2017\)](#) demonstrate that heterogeneity among speculators may prevail in stock markets, especially if speculators differ in their risk aversion, beliefs or sentiments. Overall, this important line of research points out that the relation between rationality and survival is rather weak and that stock prices are also influenced by boundedly rational speculators.

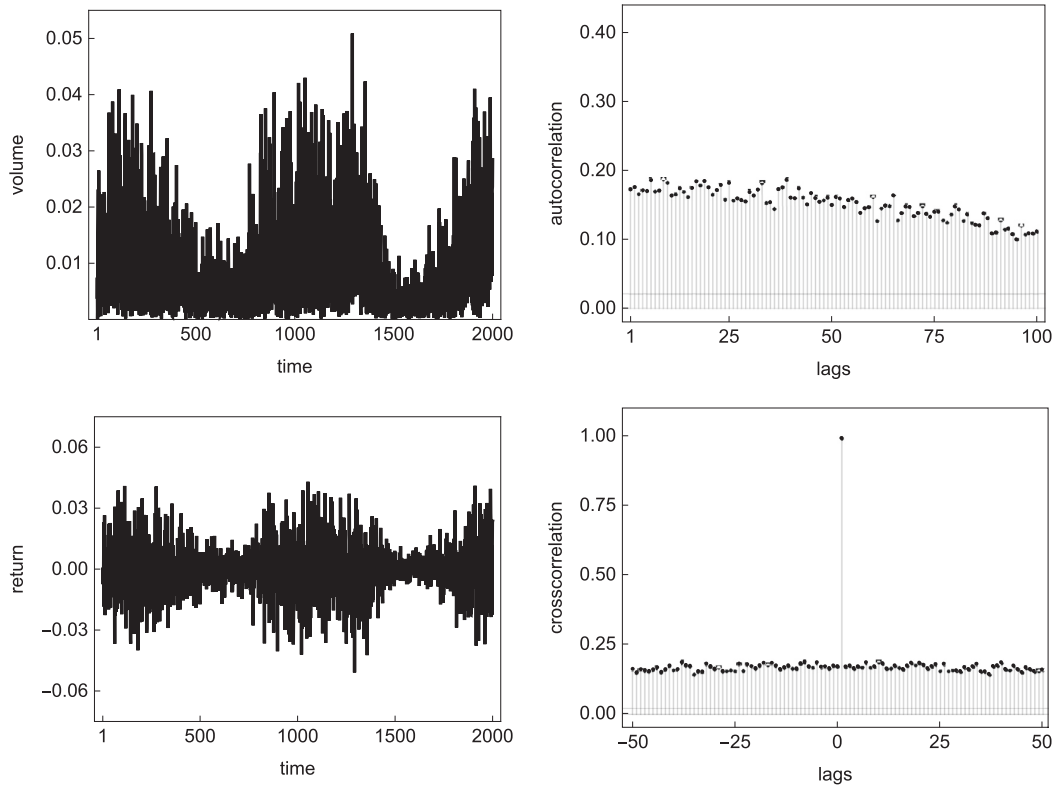


Fig. 9. The relation between trading volume and volatility. The panels show the evolution of trading volume, the autocorrelation function of trading volume, the return dynamics and the cross-correlation function of trading volume and absolute returns, respectively. The parameter setting is as in Fig. 8.

In the following, we explore how speculators' wealth evolves within our model. Fig. 11 illustrates our main results. Its top left panel shows the development of log prices for 250 trading periods, i.e. for a time span of about one year. Two things deserve our attention. First, the stock price in period 1 roughly corresponds to the stock price in period 250. Second, speculators have triggered a bubble in between. The center left panel of Fig. 11 presents the orders of two randomly selected speculators, say speculator 1 (red line) and speculator 2 (black line). Apparently, speculators' order flows differ substantially. The bottom left panel of Fig. 11 depicts (myopic) profits of the two speculators. Here we use the same time structure and definition of profits as Westerhoff and Dieci (2006) do. Accordingly, an order submitted in period $t - 2$ is filled at the price in period $t - 1$. This transaction then appears as a profit or as a loss depending on the price in period t . Formally, profits of speculator i in period t are given with $\pi_{t,i} = (\exp[P_t] - \exp[P_{t-1}])D_{t-2,i}$ (recall that P_t refers to log stock prices). As can be seen, it is difficult for speculators to beat the market. More precisely, the probability that speculator i will make a profit in period t is – due to the stock market's random walk nature – equal to the probability that he will make a loss, namely 50%.

The top right panel of Fig. 11 shows the evolution of the positions of speculators 1 and 2 while the center right panel of Fig. 11 displays the corresponding wealth dynamics. For simplicity, we abstract from trading costs, interest rates and dividend payments. Moreover, the wealth level is initially set to zero and then updated according to $\Pi_{T,i} = \exp[P_T] \sum_{t=3}^T D_{t-2,i} - \sum_{t=3}^T \exp[P_{t-1}] D_{t-2,i}$. Hence, the wealth of speculator i after T trading periods depends on the value of his position, given by

the stock price in period T multiplied by the sum of his orders, and his actual expenditures and revenues from accumulating his position, given by his executed transactions. Note that speculator 1 (red line) builds up a positive position and initially benefits from the bubble. When the bubble bursts, however, much of his wealth vanishes again. In contrast, speculator 2 enters a short position around period 75 and loses wealth. Between periods 75 and 125, speculator 2 then increases his stock position and finally depletes it again. In the end, the wealth gains and losses of both speculators are roughly equal and close to zero.

The bottom right panel of Fig. 11 extends this experiment by visualizing the wealth dynamics of 50 randomly selected speculators (alternating red and black lines). While there is a somewhat larger wealth dispersion during the height of the bubble, speculators' wealth differentials diminish after the market's mispricing reduces. Hence, our model's stock price dy-

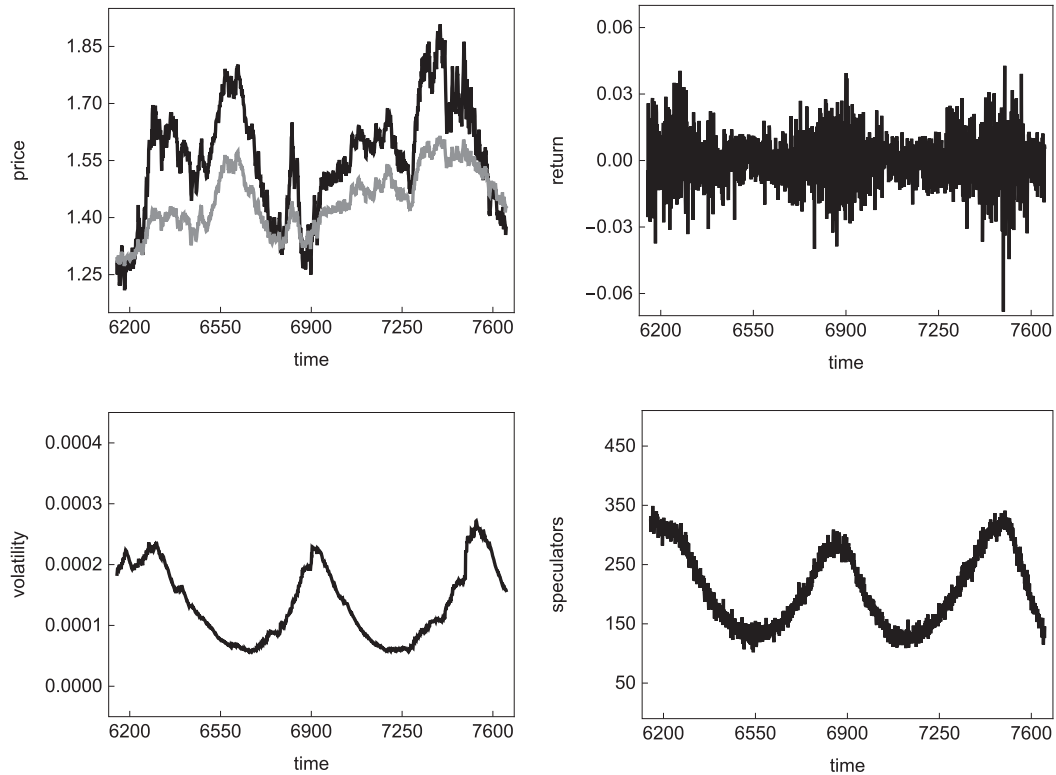


Fig. 10. Functioning of the stochastic model. The panels highlight the evolution of stock prices (black) and fundamental values (gray), the returns, the stock market's volatility and the number of active speculators between periods 6150 and 7650 of the simulation run depicted in Fig. 8.

namics does not imply that some speculators consistently lose money (and should thus eventually vanish from the market) nor that some speculators become richer and richer (and should thus dominate the market in the long run). Put differently, our analysis reveals that stock markets are difficult to beat. Some speculators may be lucky and make money; other speculators may be unlucky and lose money. However, no speculator encounters systematic profits or losses. Of course, actual speculators are not infinitely lived. In reality, speculators will eventually leave the stock market permanently, either with a higher or a lower wealth level, and newborn speculators will emerge. Taking such a perspective, the outcome depicted in Fig. 11 may be regarded as a representative snapshot of a stock market's long-run wealth dynamics.

5. Conclusions

We develop an agent-based financial market model with heterogeneous interacting speculators to explain a number of important statistical regularities of stock markets. Speculators base their orders on current price trends, the market's mispricing and new information. Speculators are heterogeneous in the sense that each of them follows his own time-varying trading rule. However, not all speculators are always active in the stock market. Two socio-economic principles govern speculators' probabilistic market entry decisions. First, speculators' market entry decisions are subject to herding behavior. The more speculators are active in the stock market, the more attractive the stock market appears to them. Second, speculators' market entry decisions depend on stock market risk. The higher the stock market risk, measured by the past volatility of the stock market, the less attractive the stock market appears to them. All orders placed by speculators are matched by a market maker who adjusts stock prices with respect to excess demand. The only extrinsic forces in our model are exogenous shocks which drive the random evolution of the fundamental value.

We use a mix of analytical, numerical and empirical tools to investigate our model. Our main result is that sporadic market entry waves may cause volatility outbursts in stock markets. To be precise, we show that a herding-induced inflow of speculators leads to rather unstable market dynamics with high volatility while a consecutive risk-driven outflow of speculators leads to more stable market dynamics with low volatility. This kind of volatility clustering is observed in the deterministic skeleton of our model, for which we provide a full steady-state and stability analysis, and in the stochastic version of our model, which we calibrate to the stylized facts of stock markets. The latter exercise demonstrates that our model is able to generate bubbles and crashes, excess volatility, fat-tailed return distributions, serially uncorrelated returns

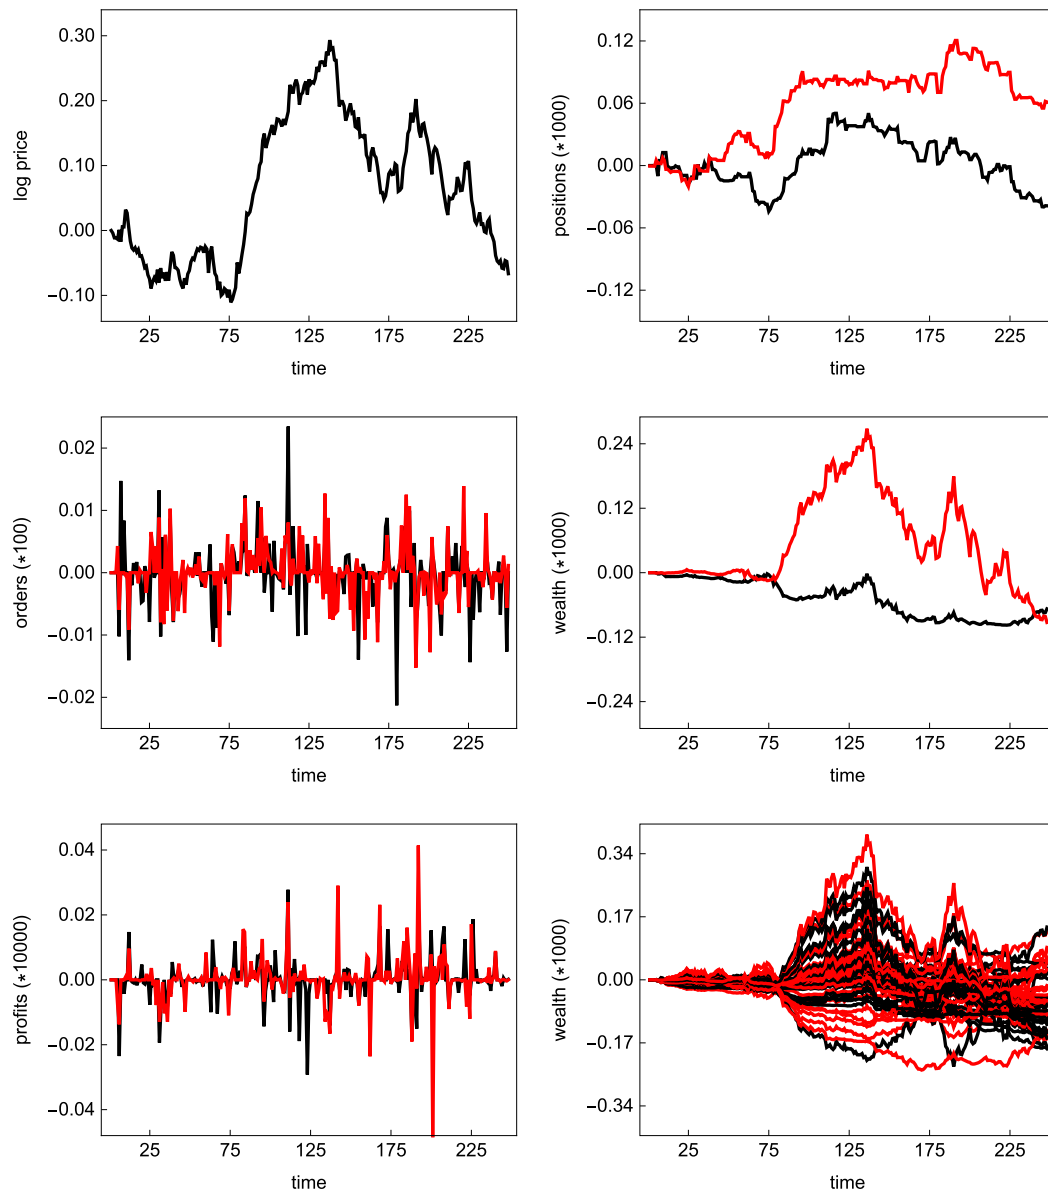


Fig. 11. Wealth dynamics of individual speculators. The panels show the evolution of log prices, the positions of two individual speculators, their orders, wealth and profits and the wealth of 50 individual speculators, respectively. The parameter setting is as in Fig. 8. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and, as already mentioned, volatility clustering. In addition, trading volume is persistent and correlated with volatility. In this sense, our model may be regarded as validated.

Our analytical results prove instrumental in our understanding of the functioning of our model. In particular, we show that the model's fundamental steady state becomes unstable once too many speculators enter the stock market. Since the instability of the fundamental steady state is due to a Neimark–Sacker bifurcation, we observe the onset of (quasiperiodic) endogenous dynamics. In this respect, it is worth mentioning that our stochastic agent-based model starts from the description of the trading behavior of a large number of individual speculators but can, after some straightforward transformations, be expressed as a four-dimensional deterministic nonlinear map. In this way, it is possible to obtain valuable analytical insights for a rather complex agent-based model. Moreover, the reduced model version greatly reduces computational efforts when it comes to a simulation-based model calibration. Of course, once the model is calibrated, one may simulate

the original agent-based framework and monitor various aspects of the behavior of individual speculators, e.g. their wealth dynamics.

We conclude our paper by pointing out a few avenues for future research. First, speculators follow a linear blend of technical and fundamental trading rules in our model. One interesting extension of our model could be to let active investors make a behavioral choice for a specific trading rule. We could then have situations with a large number of active speculators who prefer fundamental analysis or situations with a small number of active speculators who favor technical analysis – just to give two examples. Such a rule selection behavior could be modeled along the lines of Brock and Hommes (1997); Lux and Marchesi (2000) or Franke and Westerhoff (2012). Second, speculators who do not enter the stock market in our model are simply inactive. Another interesting extension of our model could be to model speculators' outside option. For instance, Dieci et al. (2018) develop a model in which speculators can invest their wealth in stock, bond and housing markets. Alternatively, one may assume that speculators switch between different stock markets. Research in that direction is surprisingly scant so far. Third, one may also use our model to conduct policy experiments (see, e.g. Jacob Leal et al., 2016 and Jacob Leal and Napoletano, 2017). Our model implies that (destabilizing) speculators increasingly enter the stock market if stock market volatility is low. This model feature may prove a real challenge for regulatory measures which seek to tame stock market fluctuations. To sum up, we hope that our paper stimulates more research in this direction. The financial crisis at the end of the noughties has not only made clear that our understanding of the dynamics of financial markets is still incomplete – it revealed, once again, how important it is to make scientific and real progress in this area.

References

- Alfarano, S., Lux, T., 2007. A noise trader model as a generator of apparent financial power laws and long memory. *Macroecon. Dyn.* 11, 80–101.
- Alfarano, S., Lux, T., Wagner, F., 2005. Estimation of agent-based models: the case of an asymmetric herding model. *Comput. Econ.* 26, 19–49.
- Alfi, V., Cristelli, M., Pietronero, L., Zaccaria, A., 2009b. Mechanisms of self-organization and finite size effects in a minimal agent based model. *J. Stat. Mech.* P03016.
- Alfi, V., Cristelli, M., Pietronero, L., Zaccaria, A., 2009c. Minimal agent based model for financial markets I. *Eur. Physical. J. B* 67, 385–397.
- Alfi, V., Pietronero, L., Zaccaria, A., 2009a. Self-organization for the stylized facts and finite-size effects in a financial-market model. *Europhys. Lett.* 86, 58003.
- Amilon, H., 2008. Estimation of an adaptive stock market model with heterogeneous agents. *J. Empir. Finance* 15, 342–362.
- Anufriev, M., Dindo, P., 2010. Wealth-driven selection in a financial market with heterogeneous agents. *J. Econ. Behav. Org.* 73, 327–358.
- Bischi, G., Gallegati, M., Gardini, L., Leombruni, R., Palestini, A., 2006. Herd behavior and nonfundamental asset price fluctuations in financial markets. *Macroecon. Dyn.* 10, 502–528.
- Bischi, G., Lamantia, F., Radi, D., 2015. An evolutionary Cournot model with limited market knowledge. *J. Econ. Behav. Org.* 116, 219–238.
- Blume, L., Easley, D., 1992. Evolution and market behavior. *J. Econ. Theory* 58, 9–40.
- Boswijk, P., Hommes, C., Manzan, S., 2007. Behavioral heterogeneity in stock prices. *J. Econ. Dyn. Control* 31, 1938–1970.
- Bottazzi, G., Dindo, P., 2014. Evolution and market behavior with endogenous investment rules. *J. Econ. Dyn. Control* 48, 121–146.
- Bottazzi, G., Dindo, P., Giachini, D., 2017. Long-run heterogeneity in an exchange economy with fixed-mix traders. *Econ. Theory*. In press.
- Brock, W., Hommes, C., 1997. A rational route to randomness. *Econometrica* 65, 1059–1095.
- Brock, W., Hommes, C., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *J. Econ. Dyn. Control* 22, 1235–1274.
- Brock, W., LeBaron, B., 1996. A dynamic structural model for stock return volatility and trading volume. *Rev. Econ. Stat.* 78, 94–110.
- Brown, J.R., Ivković, Z., Smith, P.A., 2008. Neighbors matter: causal community effects and stock market participation. *J. Finance* 63, 1509–1531.
- Chiarella, C., Dieci, R., He, X.Z., 2007. Heterogeneous expectations and speculative behavior in a dynamic multi-asset framework. *J. Economic. Behav. Org.* 62, 408–427.
- Chiarella, C., Dieci, R., He, X.Z., 2009a. Heterogeneity, market mechanisms, and asset price dynamics. In: Hens, T., Schenk-Hoppé, K.R. (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*, North-Holland, Amsterdam, pp. 277–344.
- Chiarella, C., He, X.Z., Zwickels, R., 2014. Heterogeneous expectations in asset pricing: empirical evidence from the S&P500. *J. Econ. Behav. Org.* 105, 1–16.
- Chiarella, C., Iori, G., 2002. A simulation analysis of the microstructure of double auction markets. *Quant. Finance* 2, 1–8.
- Chiarella, C., Iori, G., Perelló, J., 2009b. The impact of heterogeneous trading rules on the limit order book and order flows. *J. Econ. Dyn. Control* 33, 525–537.
- Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. *Quant. Finance* 1, 223–236.
- Cont, R., Bouchaud, J.P., 2000. Herd behavior and aggregate fluctuations in financial markets. *Macroecon. Dyn.* 4, 170–196.
- Day, R., Huang, W., 1990. Bulls, bears and market sheep. *J. Econ. Behav. Org.* 14, 299–329.
- De Grauwe, P., Dewachter, H., Embrechts, M., 1993. *Exchange Rate Theory - Chaotic Models of Foreign Exchange Markets*. Blackwell, Oxford.
- Dieci, R., Schmitt, N., Westerhoff, F., 2018. Interactions between Stock, Bond and Housing Markets. BERG Working Paper No. 133, University of Bamberg.
- Dindo, P., Tuinstra, J., 2011. A class of evolutionary models for participation games with negative feedback. *Comput. Econ.* 37, 267–300.
- Farmer, D., Joshi, S., 2002. The price dynamics of common trading strategies. *J. Econ. Behav. Org.* 49, 149–171.
- Franke, R., 2009. Applying the method of simulated moments to estimate a small agent-based asset pricing model. *J. Empir. Finance* 16, 804–815.
- Franke, R., Westerhoff, F., 2012. Structural stochastic volatility in asset pricing dynamics: estimation and model contest. *J. Econ. Dyn. Control* 36, 1193–1211.
- Gandolfo, G., 2009. *Economic Dynamics*. Springer, Heidelberg.
- Gilli, M., Winker, P., 2003. A global optimization heuristic for estimating agent based models. *Comput. Stat. Data Anal.* 42 (2003), 299–312.
- Gopikrishnan, P., Plerou, V., Amaral, L., Meyer, M., Stanley, E., 1999. Scaling of the distributions of fluctuations of financial market indices. *Phys. Rev. E* 60, 5305–5316.
- Graham, B., Dodd, D., 1951. *Security Analysis*. McGraw-Hill, New York.
- Hofbauer, J., Sigmund, K., 1988. *The Theory of Evolution and Dynamical Systems*. Cambridge University Press, Cambridge.
- Hofbauer, J., Weibull, J., 1996. Evolutionary selection against dominated strategies. *J. Econ. Theory* 71, 558–573.
- Hommes, C., in 't Veld, D., 2017. Booms, busts and behavioural heterogeneity in stock prices. *J. Econ. Dyn. Control* 80, 101–124.
- Hommes, C., Wagener, F., 2009. Complex evolutionary systems in behavioral finance. In: Hens, T., Schenk-Hoppé, K.R. (Eds.), *Handbook of financial Markets: Dynamics and Evolution*. North-Holland, Amsterdam, pp. 217–276.
- Hong, H., Kubik, J.D., Stein, J.C., 2004. Social interaction and stock-market participation. *J. Finance* 59, 137–163.
- Hong, H., Kubik, J.D., Stein, J.C., 2005. Thy neighbor's portfolio: word-of-mouth effects in the holdings and trades of money managers. *J. Finance* 60, 2801–2824.
- Iori, G., 1999. Avalanche dynamics and trading friction effects on stock market returns. *Int. J. Mod. Phys. C* 10, 1149–1162.
- Iori, G., 2000. A threshold model for stock return volatility and trading volume. *Int. J. Theor. Appl. Finance* 3, 467–472.
- Jacob Leal, S., Napoletano, M., 2017. Market stability vs. market resilience: regulatory policies experiments in an agent-based model with low- and high-frequency trading. *J. Econ. Behav. Org.*. In press.

- Jacob Leal, S., Napoletano, M., Roventini, A., Fagiolo, G., 2016. Rock around the clock: an agent-based model of low-and high-frequency trading. *J. Evolut. Econ.* 26, 49–76.
- Kirman, A., 1993. Ants, rationality, and recruitment. *Q. J. Econ.* 108, 137–156.
- Kopel, M., Lamantia, F., Szidarovszky, F., 2014. Evolutionary competition in a mixed market with socially concerned firms. *J. Econ. Dyn. Control* 48, 394–409.
- LeBaron, B., Arthur, B., Palmer, R., 1999. Time series properties of an artificial stock market. *J. Econ. Dyn. Control* 23, 1487–1516.
- LeBaron, B., Yamamoto, R., 2008. The impact of imitation on long memory in an order-driven market. *East. Econ. J.* 34, 504–517.
- Lux, T., 1996. The stable Paretian hypothesis and the frequency of large returns: an examination of major german stocks. *Appl. Financ. Econ.* 6, 463–475.
- Lux, T., 2009. Stochastic behavioural asset-pricing models and the stylized facts. In: Hens, T., Schenk-Hoppé, K.R. (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. North-Holland, Amsterdam, pp. 161–216.
- Lux, T., Alfarano, S., 2016. Financial power laws: empirical evidence, models and mechanisms. *Chaos Solitons Fract.* 88, 3–18.
- Lux, T., Ausloos, M., 2002. Market fluctuations I: scaling, multiscaling, and their possible origins. In: Bunde, A., Kropp, J., Schellnhuber, H. (Eds.), *Science of Disaster: Climate Disruptions, Heart Attacks, and Market Crashes*. Springer, Berlin, pp. 373–410.
- Lux, T., Marchesi, M., 1999. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature* 397, 498–500.
- Lux, T., Marchesi, M., 2000. Volatility clustering in financial markets: a micro-simulation of interacting agents. *Int. J. Theor. Appl. Finance* 3, 675–702.
- Mantegna, R., Stanley, E., 2000. *An Introduction to Econophysics*. Cambridge University Press, Cambridge.
- Medio, A., Lines, M., 2001. *Nonlinear Dynamics: A Primer*. Cambridge University Press, Cambridge.
- Murphy, J., 1999. *Technical Analysis of Financial Markets*. New York Institute of Finance, New York.
- Pearce, D., Roley, V., 1985. Stock prices and economic news. *J. Bus.* 58, 49–67.
- Pellizzari, P., Westerhoff, F., 2009. Some effects of transaction taxes under different microstructures. *J. Econ. Behav. Org.* 72, 850–863.
- Sandholm, W., 2015. Population games and deterministic evolutionary dynamics. In: Young, P., Zamir, S. (Eds.), *Handbook of Game Theory, Volume 4*. North-Holland, Amsterdam, pp. 703–778.
- Schmitt, N., Tuinstra, J., Westerhoff, F., 2017. Side effects of nonlinear profit taxes in an evolutionary market entry model: abrupt changes, coexisting attractors and hysteresis problem. *J. Econ. Behav. Org.* 135, 15–38.
- Schmitt, N., Westerhoff, F., 2014. Speculative behavior and the dynamics of interacting stock markets. *J. Econ. Dyn. Control* 45, 262–288.
- Schmitt, N., Westerhoff, F., 2016. Stock market participation and endogenous boom-bust dynamics. *Econ. Lett.* 148, 72–75.
- Schmitt, N., Westerhoff, F., 2017. Herding behavior and volatility clustering in financial markets. *Quant. Finance* 17, 1187–1203.
- Shiller, R., 2015. *Irrational Exuberance*. Princeton University Press, Princeton.
- Stauffer, D., de Oliveira, P., Bernardes, A., 1999. Monte Carlo simulation of volatility clustering in market model with herding. *Int. J. Theor. Appl. Finance* 2, 83–94.
- Tedeschi, G., Iori, G., Gallegati, M., 2012. Herding effects in order driven markets: the rise and fall of gurus. *J. Econ. Behav. Org.* 81, 82–96.
- Westerhoff, F., Dieci, R., 2006. The effectiveness of Keynes-Tobin transaction taxes when heterogeneous agents can trade in different markets: a behavioral finance approach. *J. Econ. Dyn. Control* 30, 293–322.
- Winker, P., Gilli, M., Jeleskovic, V., 2007. An objective function for simulation based inference on exchange rate data. *J. Econ. Interact. Coord.* 2, 125–145.

3 Heterogeneous speculators and stock market dynamics: a simple agent-based computational model

This chapter contains joint work with Noemi Schmitt and Frank Westerhoff which has already been accepted for publication in the *European Journal of Finance*, 2020: 1-20: <https://doi.org/10.1080/1351847X.2020.1832553>.¹ Frank Westerhoff conceived the original idea of the paper and supervised the project. Noemi Schmitt and Ivonne Schwartz worked out the technical details of the model and performed the numerical calculations. All three authors contributed to equal parts to the discussion of the results and the writing of the final manuscript. In the following, the available online version will be included.

¹Heterogeneous speculators and stock market dynamics: a simple agent-based computational model, Noemi Schmitt, Ivonne Schwartz, Frank Westerhoff, *European Journal of Finance*, © copyright 2020, reprinted by permission of Informa UK Limited, trading as Taylor & Francis Group, <http://www.tandfonline.com>.



Heterogeneous speculators and stock market dynamics: a simple agent-based computational model

Noemi Schmitt, Ivonne Schwartz and Frank Westerhoff

Department of Economics, University of Bamberg, Bamberg, Germany

ABSTRACT

We propose a simple agent-based computational model in which speculators' trading behavior may cause bubbles and crashes, excess volatility, serially uncorrelated returns, fat-tailed return distributions and volatility clustering, thereby replicating five important stylized facts of stock markets. Since each speculator bets on his own (technical and fundamental) trading signals, stock prices are excessively volatile and oscillate erratically around their fundamental value. However, speculators' heterogeneity occasionally vanishes, e.g. due to panic-induced herding behavior, yielding extreme returns. Lasting regimes with high volatility originate from the fact that speculators extract stronger trading signals out of past stock price movements when stock prices fluctuate strongly. Simulations furthermore suggest that circuit breakers may be an effective tool to combat financial market turbulences.

ARTICLE HISTORY

Received 10 March 2020

Accepted 28 September 2020

KEYWORDS

Stock markets; stylized facts; agent-based computational models; technical and fundamental analysis; circuit breakers; econophysics

JEL CLASSIFICATIONS

C63; D84; G15

1. Introduction

We propose a simple agent-based computational model to explain a number of important stylized facts of stock markets. In a nutshell, our model and our main results may be summarized as follows. We consider a stock market that is populated by a market maker and a given number of heterogeneous interacting speculators. The market maker adjusts stock prices with respect to the excess demand of speculators who, in turn, determine their orders by following their own individual trading signals, derived either from private market research or from applying complex (algorithmic) trading systems. Simulations reveal that speculators' trading behavior may generate bubbles and crashes, excess volatility, serially uncorrelated (log) stock price changes, fat-tailed return distributions and lasting volatility outbursts. Since speculators bet on technical and fundamental trading signals, stock prices are excessively volatile and circle in an apparently random fashion around their fundamental value. Extreme returns occur in our model due to a sporadic loss of heterogeneity. To be precise, there are short-lived periods in which speculators' behavior becomes coordinated, e.g. because they react to the same trading signals, hard-wired into their trading systems, or because they display panic-induced herding behavior, e.g. caused by sharp stock price changes. Lasting periods of high volatility occur when speculators persistently receive strong trading signals. Since many speculators infer their trading signals out of past stock price movements, the latter occurs in periods characterized by significant stock price changes. In such periods, speculators also tend to overreact to their own individual trading signals, which keeps volatility high. Our model also indicates that circuit breakers may be an effective tool to stabilize the dynamics of stock markets.

Our paper belongs to a well-developed field of literature that seeks to explain the dynamics of stock markets by taking an explicit agent-based perspective. Analytically tractable small-scale agent-based models, focusing on a few representative speculator types, have been proposed, for instance, by Zeeman (1974), Beja and Goldman (1980), Day and Huang (1990), Chiarella (1992), De Grauwe, Dewachter, and Embrechts (1993), Lux (1995), Farmer and Joshi (2002) and Chiarella and Iori (2002). More elaborated and simulation-oriented, large-scale

CONTACT Frank Westerhoff  frank.westerhoff@uni-bamberg.de

© 2020 Informa UK Limited, trading as Taylor & Francis Group

agent-based models, studying the interplay between many different and evolving speculator types, have been advanced, for instance, by Palmer et al. (1994), Arthur et al. (1997), LeBaron, Arthur, and Palmer (1999), Chen and Yeh (2001) and Raberto et al. (2001). While it is still important to better understand the forces that may create financial market havoc, current research increasingly addresses questions that revolve around input validation (Anufriev, Bao, and Tuinstra 2016; Fagiolo et al. 2017; Guerini and Moneta 2017), model estimation (Lamperti, Roventini, and Sani 2018; Platt 2020; Kukacka and Kristoufek 2020), policy applications (Stanek and Kukacka 2018; Diem, Pichler, and Thurner 2020; Schmitt, Tramontana, and Westerhoff 2020) and prediction (Demirer et al. 2019; Zhang, Sornette, and Zhang 2019; Westphal and Sornette 2020). See Delli Gatti et al. (2018), Dieci and He (2018), Iori and Porter (2018) and Lux and Zwinkels (2018) for up-to-date surveys.

Recently, Schmitt and Westerhoff (2017a, 2017b) and Schmitt (2020) started to develop rather simple agent-based computational stock market models by assuming that speculators' trading behavior can be represented at least partially by correlated random variables. For instance, Schmitt (2020) proposes an agent-based version of the asset-pricing model by Brock and Hommes (1998), keeping the correlation between speculators' random demand components constant. Nevertheless, her model produces lasting volatility outbursts when the mass of speculators switches towards destabilizing technical trading rules. Schmitt and Westerhoff (2017a) put forward an agent-based version of the asset-pricing model by Franke and Westerhoff (2012). Extreme price changes emerge within their model when the arrival of exogenous sunspots initiates a spontaneous coordination of speculators' trading behavior. Relatedly, Schmitt and Westerhoff (2017b) assume in their asset-pricing model that the correlation between speculators' trading behavior changes slowly with respect to the market's volatility. If volatility increases, speculators become afraid and follow the trading behavior of other speculators more closely. As a result, speculators' excess demand escalates, keeping volatility high. In our paper, we assume that endogenous market events may lead to a spontaneous coordination of speculators' trading behavior, and thus to extreme returns, while speculators' trading intensity depends positively on the market's volatility, an aspect that may produce lasting volatility outbursts.

Within our model, speculators' trading behavior contains a strong random component. In fact, we capture their trading behavior by a vector of multivariate normally distributed random variables to which we impose a certain minimalistic structure. Note that such a modeling strategy is quite common in certain areas of research, e.g. in econophysics. For instance, Cont and Bouchaud (2000) assume in their stock market model that the decisions of clusters of active speculators whether to buy or sell stocks are random variables with equal probabilities. See Stauffer and Penna (1998), Chang and Stauffer (1999), Stauffer and Sornette (1999), Stauffer and Jan (2000) and Iori (2002) for extensions and generalizations of this framework. Similarly, Gode and Sunder (1993, 1997), Daniels et al. (2003), Farmer, Patelli, and Zovko (2005a, 2005b) and Ladley (2012) study stock market models that are driven by zero-intelligence agents who trade randomly, subject only to their budget constraints, demonstrating that important properties of stock markets depend less on agents' strategic (rational) behavior, and more on their institutional arrangements. More recent contributions in which speculators' behavior also contains a larger random component include, for instance, Ladley et al. (2015), Xing and Ladley (2019) and Ladley (2020).

The remainder of our paper is organized as follows. In Section 2, we present a simple agent-based computational model of the stock market. In Section 3, we compare the dynamics of our approach with the behavior of actual stock markets. In Section 4, we explain the model's functioning. In Section 5, we discuss possible effects of circuit breakers. In Section 6, we conclude our paper. A number of robustness checks are presented in Appendix A.

2. A simple agent-based computational stock market model

In this section, we develop a simple agent-based computational model that aims at explaining a number of important stylized facts of stock markets. Let us start with previewing the basic setup of our approach. We consider a single stock market that is populated by a market maker and a given number of heterogeneous interacting speculators. The market maker adjusts the price of the stock with respect to speculators' order flow. Each speculator bases her orders on her own individual trading signals, derived either from private market research or from applying complex (algorithmic) trading systems. For simplicity, we model speculators' trading signals as multivariate normally distributed random variables, imposing the following minimalistic structure. First, the

means of the random variables reflect speculators' tendency to extrapolate past stock price changes and to bet on mean reversion. Second, the variances of the random variables represent speculators' trading intensities and increase in line with the stock market's volatility. Clearly, speculators infer stronger trading signals – or react more strongly to given trading signals – if the volatility of the stock market is high. The former argument is consistent with the observation that speculators derive trading signals out of past stock price movements and that the strength of these trading signals naturally grows with the stock market's volatility. The latter argument is in line with the observation that speculators tend to overreact to their trading signals in volatile periods, simply because they are agitated and thus regard their trading signals as more relevant in such times. Third, the correlation between speculators' trading signals increases if the stock market displays significant stock price patterns. This may be because speculators observe the behavior of others more strongly during periods of heightened uncertainty or because certain price patterns, such as significant reversals of stock price changes, are hard-wired into a sufficient number of speculators' complex (algorithmic) trading systems.

Let us now turn to the details of our model. We assume that a market maker adjusts the price of the stock with respect to the excess demand originating from the orders of N heterogeneous interacting speculators. As in Beja and Goldman (1980), Day and Huang (1990) and Farmer and Joshi (2002), the market maker's behavior is formalized as

$$P_{t+1} = P_t + a \sum_{i=1}^N D_{t,i}, \quad (1)$$

where P_t is the log price of the stock at time t , a is a positive price adjustment parameter, reflecting the stock market's liquidity, and $\sum_{i=1}^N D_{t,i}$ is the aggregate excess demand resulting from the individual orders $D_{t,i}$ of speculators $i = 1, 2, \dots, N$. Hence, if the sum of speculators' orders is positive (negative), the market maker increases (decreases) the log stock price.

The orders placed by speculator i depend on her own individual trading signals, derived either from private market research or by applying complex (algorithmic) trading systems. Inspired by the aforementioned line of research initiated by Gode and Sunder (1993) and Cont and Bouchaud (2000), we do not aim at formalizing speculators' trading behavior in detail. Instead, we simply represent speculator i 's order in period t by

$$D_{t,i} = \delta_{t,i} \quad (2)$$

where $\delta_t = \{\delta_{t,1}, \delta_{t,2}, \dots, \delta_{t,N}\}'$ is a vector of multivariate normally distributed random variables, i.e. $\delta_t \sim N(M_t, \Sigma_t)$. We assume for the mean vector

$$M_t = \{\mu_{t,1}, \mu_{t,2}, \dots, \mu_{t,N}\}' \quad (3)$$

and the variance-covariance matrix

$$\Sigma_t = \begin{bmatrix} \sigma_{t,1}^2 & \sigma_{t,1}\sigma_{t,2}\rho_{t,1,2} & \dots & \sigma_{t,1}\sigma_{t,N}\rho_{t,1,N} \\ \sigma_{t,2}\sigma_{t,1}\rho_{t,2,1} & \sigma_{t,2}^2 & & \vdots \\ \vdots & & \ddots & \sigma_{t,N-1}\sigma_{t,N}\rho_{t,N-1,N} \\ \sigma_{t,N}\sigma_{t,1}\rho_{t,N,1} & \dots & \sigma_{t,N}\sigma_{t,N-1}\rho_{t,N,N-1} & \sigma_{t,N}^2 \end{bmatrix} \quad (4)$$

that $\mu_t = \mu_{t,i}$, $\sigma_t^2 = \sigma_{t,i}^2$ and $\rho_t = \rho_{t,i,j}$ for $i, j = 1, 2, \dots, N$ and $i \neq j$. Despite these restrictions, each speculator submits a different order to the market maker, unless, of course, $\rho_t = 1$. In that case, all speculators submit an identical order to the market maker.

The empirical and laboratory evidence reviewed by Menkhoff and Taylor (2007) and Hommes (2011) highlights the fact that speculators rely on technical and fundamental analysis to determine their orders. The key idea behind technical analysis (Lo, Mamaysky, and Wang 2000) is that stock prices move in trends. Fundamental analysis (Graham and Dodd 1951), in contrast, postulates that stock prices display a tendency to return to

their fundamental values. Let F denote the constant log fundamental value of the stock market. We thus assume that

$$\mu_t = b(P_t - P_{t-1}) + c(F - P_t). \quad (5)$$

Note that μ_t captures the core principles of technical and fundamental analysis. The first component of (5) suggests that speculators should place a buy (sell) order if the stock market goes up (down), while the second component of (4) recommends that they sell (buy) overvalued (undervalued) stocks. The reaction parameters $b, c > 0$ determine the strength of these trading signals.

Moreover, we assume that speculators' trading intensity increases with the stock market's volatility. This assumption is supported by two arguments. First, speculators make their beliefs about future stock prices (and hence their demand) dependent on past stock price movements. If there is considerable stock price variability, then their trading signals will grow correspondingly (Murphy 1999). Second, speculators overreact to their trading signals in periods of high volatility (Manzan and Westerhoff 2005). Let us capture the stock market's volatility by

$$V_t = dV_{t-1} + (1 - d)(P_t - P_{t-1})^2, \quad (6)$$

where $0 < d < 1$ is a memory parameter. Moreover, let $\bar{V} > 0$ be a reference value for the stock market's volatility. We model the intensity of speculators' trading behavior by specifying σ_t^2 as

$$\sigma_t^2 = e^l + \frac{e^h - e^l}{1 + \exp[-e^s(V_t - \bar{V})]}. \quad (7)$$

Note that (7) represents a logistic function that is bounded between $0 < e^l < e^h$. For $V_t = \bar{V}$, speculators' trading intensity is equal to the midpoint of (7), i.e. $\sigma_t^2 = (e^l + e^h)/2$. The slope parameter $e^s > 0$ of (7) determines how sensitively σ_t^2 reacts to a change in V_t . Economically, the S-shaped function (7) implies that speculators' trading intensity increases in line with the stock market's volatility.¹

However, speculators are not isolated in their decision-making. As already observed by Keynes (1936), speculators tend to herd together in periods of heightened uncertainty. Moreover, it seems that certain price patterns are hard-wired into speculators' complex (algorithmic) trading systems. If such a price pattern emerges, speculators' trading systems generate correlated trading signals.² In reality, there may be many price/return patterns that initiate correlated actions among market participants. To keep things as simple as possible, however, we assume that the correlation of speculators' trading behavior depends on the strength of a single condition, given by

$$C_t = ((P_t - P_{t-1}) - (P_{t-1} - P_{t-2}))^2. \quad (8)$$

According to (8), C_t may take a particularly large value when a significant reversal of stock price changes occurs; say when a four percent price drop is followed by a three percent price increase. Clearly, a more developed version of our model may incorporate more than one condition. Moreover, these conditions may then evolve over time and/or contain probabilistic components.³ The correlation between speculators' trading behavior is formalized as

$$\rho_t = f^l + \frac{f^h - f^l}{1 + \exp[-f^s(C_t - \bar{C})]}, \quad (9)$$

where f^l and f^h determine the lower and upper boundary of ρ_t , with $0 \leq f^l < f^h \leq 1$, $f^s > 0$ describes the slope of (9), and $\bar{C} > 0$ marks the position of its midpoint. The greater the value of condition (8), the stronger the correlation of speculators' trading behavior. If ρ_t approaches 1, speculators' trading signals become fully correlated and, consequently, they submit identical orders. If ρ_t approaches 0, speculators' trading behavior becomes uncorrelated, implying that a substantial part of their orders cancel each other out.

In principle, we can simulate the dynamics of our simple agent-based computational stock market model by using (1) to (9). For a larger number of speculators, however, simulations soon become rather time-consuming.

Fortunately, our assumptions about speculators' trading behavior conveniently enable us to summarize their excess demand by

$$\sum_{i=1}^N D_{t,i} = N(b(P_t - P_{t-1}) + c(F - P_t)) + \sigma_t \sqrt{N + N(N-1)} \rho_t \varepsilon_t, \quad (10)$$

where $\varepsilon_t \sim N(0, 1)$. As a result, we can therefore also simulate the model's dynamics by iterating the following stochastic nonlinear dynamical system:

$$\left\{ \begin{array}{l} P_{t+1} = P_t + a \{N(b(P_t - P_{t-1}) + c(F - P_t)) + \sigma_t \sqrt{N + N(N-1)} \rho_t \varepsilon_t\} \\ \sigma_t^2 = e^l + \frac{e^h - e^l}{1 + \exp[e^s(V_t - \bar{V})]} \\ V_t = dV_{t-1} + (1-d)(P_t - P_{t-1})^2 \\ \rho_t = f^l + \frac{f^h - f^l}{1 + \exp[f^s(C_t - \bar{C})]} \\ C_t = ((P_t - P_{t-1}) - (P_{t-1} - P_{t-2}))^2 \end{array} \right. . \quad (11)$$

Note that speculators' excess demand, and, therefore, the market maker's price adjustment, increases with σ_t^2 and ρ_t , which, in turn, depend on V_t and on C_t , respectively.⁴ It might be helpful to realize that V_t changes only slowly over time, provided that the memory parameter d is not too small. As a result, speculators' trading intensity remains high during turbulent market periods, keeping volatility high. In contrast, C_t may change quickly and take larger values only for brief moments of time. In such an event, speculators' trading behavior becomes correlated and a larger price change may occur. This is exactly what we will see when we simulate our model in the next section.

3. Time series properties of actual and simulated stock markets

Before we turn to the dynamics of our model, let us briefly recap the behavior of actual stock markets. As is well known, actual stock markets are characterized by a number of prominent stylized facts, including (i) bubbles and crashes, (ii) excess volatility, (iii) fat-tailed return distributions, (iv) serially uncorrelated returns and (v) volatility clustering. See Mantegna and Stanley (2000), Cont (2001) and Lux and Ausloos (2002) for detailed reviews. In the following, we briefly visualize the dynamics of three major stock markets. The left panels of Figure 1 depict the evolution of the DAX, the NIKKEI and the DJI from 1980 to 2019. Each time series, downloaded from Refinitiv Datastream, comprises about 10,000 daily observations. Despite the long-run upward trends of the DAX and the DJI, the boom-bust nature of all three stock markets is clearly striking.⁵ The right panels of Figure 1 present the corresponding return dynamics, defined as log price changes. Obviously, actual stock markets are quite volatile. For instance, the standard deviations of the return time series of the DAX, the NIKKEI and the DJI are given by 0.013, 0.017 and 0.011, respectively. Moreover, there are a number of larger price changes. In particular, the DJI produced the largest daily loss (25.6 percent), while the NIKKEI produced the largest daily gain (13.2 percent). It is also apparent that periods of low volatility alternate with periods of high volatility.

Figure 2 documents a number of distributional and correlation properties of the DAX, the NIKKEI and the DJI, using the same color coding as in Figure 1. The top left panel of Figure 2 compares the distributions of normalized stock market returns with the distribution of standard normally distributed returns (black line). The top right panel of Figure 2 shows the same, except that we present the evidence on a log-linear scale. As can be seen, the distributions of actual stock market returns are unimodal, almost symmetric and bell-shaped. Relative to the standard normal distribution, however, the distributions of actual stock market returns possess more probability mass in the center and in the tails. This is also evident from the center left panel of Figure 2, which illustrates the cumulative distributions of normalized actual stock market returns together with the cumulative distribution of standard normally distributed returns (black line) on a log-log scale. The outer parts of the distribution of actual stock market returns can be surprisingly well fitted by a power law in the form $\text{prob}(|\text{return}| > x) \approx cx^{-\alpha}$,

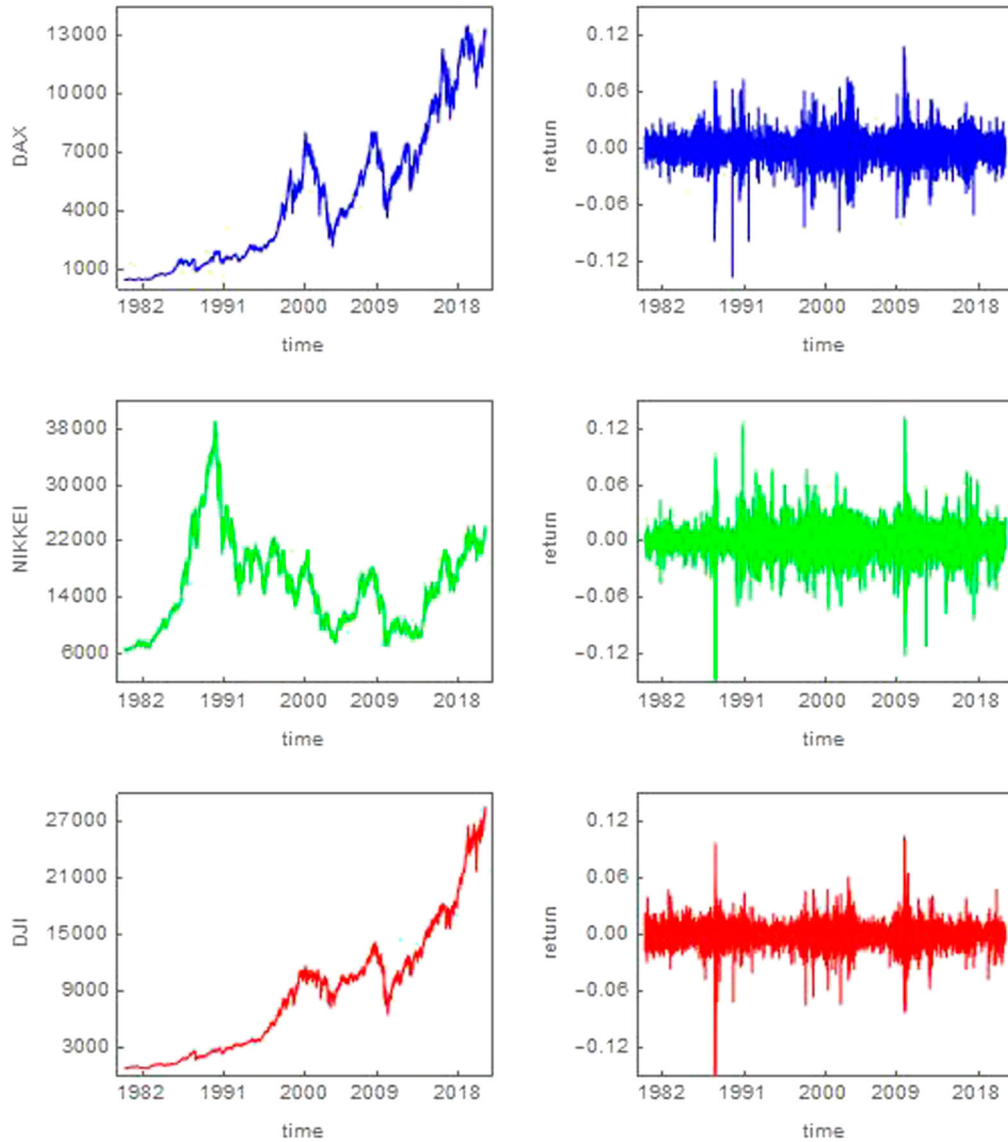


Figure 1. Time series dynamics of actual stock markets. The left panels show the evolution of the DAX, the NIKKEI and the DJI from 1980 to 2019, comprising about 10,000 daily observations. The right panels show the corresponding return dynamics.

where α is the so-called tail index. Note that a smaller tail index indicates fatter tails. In the center right panel of Figure 2, we plot the Hill tail index estimator (Hill 1975) as a function of the largest returns in percent. Using the largest 5 percent of the observations, for instance, the tail index for the DAX, the NIKKEI and the DJI is given by 3.07, 3.09 and 3.20, respectively.⁶ The bottom left panel of Figure 2 shows the autocorrelation functions of raw returns (the gray lines represent the 95 percent confidence band). As can be seen, the autocorrelation coefficients of raw returns are not significant for almost all lags, indicating that the paths of the DAX, the NIKKEI and the DJI are close to a random walk. The bottom right panel of Figure 2 reports the autocorrelation coefficients of absolute returns. Since the autocorrelation coefficients of absolute returns are significant for more than 100 lags, we can conclude that volatility outbursts are quite persistent.

Let us now illustrate the extent to which our simple agent-based computational model can replicate the dynamics of actual stock markets (a more detailed robustness analysis is presented in Appendix A). For this

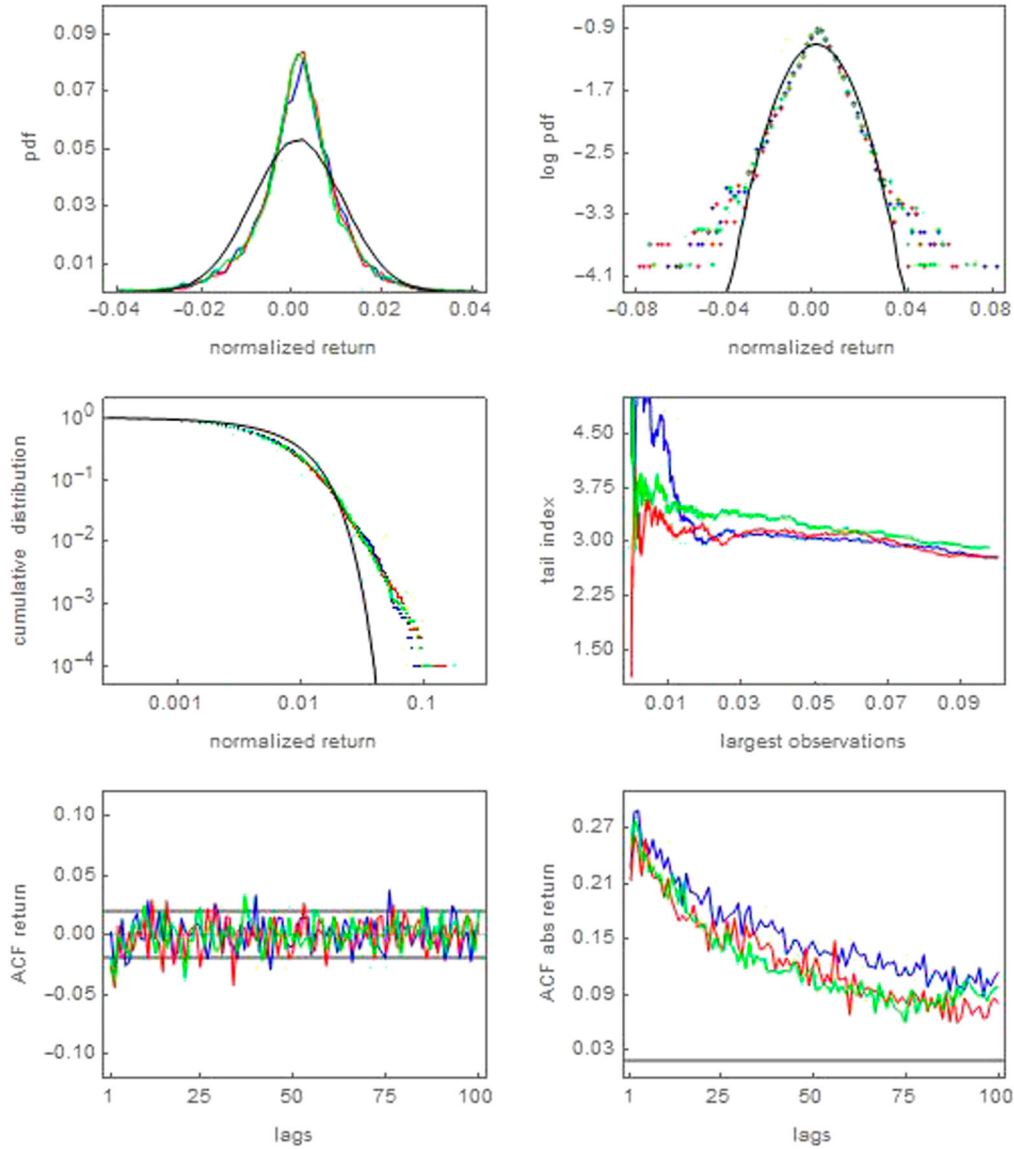


Figure 2. Distributional and correlation properties of actual stock markets. The panels show a number of distributional and correlation properties of the DAX, the NIKKEI and the DJI. The same data set and color coding as in Figure 1.

purpose, we have to determine the model's 14 parameters. In the first step, we decided to set $N = 100$, $F = 0$ and $a = 1$. Roughly speaking, parameters N and a are scaling parameters, while parameter F merely determines the level around which stock price fluctuations take place. Assuming furthermore that $e^s = f^s = \infty$ implies that the intensity of and the correlation between speculators' trading signals jumps between their lower and upper boundaries. To fix the remaining nine model parameters, we conducted a tedious trial-and-error calibration exercise. In the end, we arrived at the following parameter values: $b = 0.00005$, $c = 0.00001$, $d = 0.87$, $e^l = 0.00000055$, $e^h = 0.00000245$, $\bar{V} = 0.000125$, $f^l = 0.0006$, $f^h = 0.055$ and $\bar{C} = 0.00257$. Future work may try to estimate our model, e.g. via the method of simulated moments, as discussed by Franke and Westerhoff (2012, 2016) and Schmitt and Westerhoff (2017a, 2017b).⁷

Figures 3 and 4 portray the dynamics of three representative simulation runs. Each simulation run comprises 10,000 observations, corresponding to a time span of 40 years with 250 trading days per year. The first, second

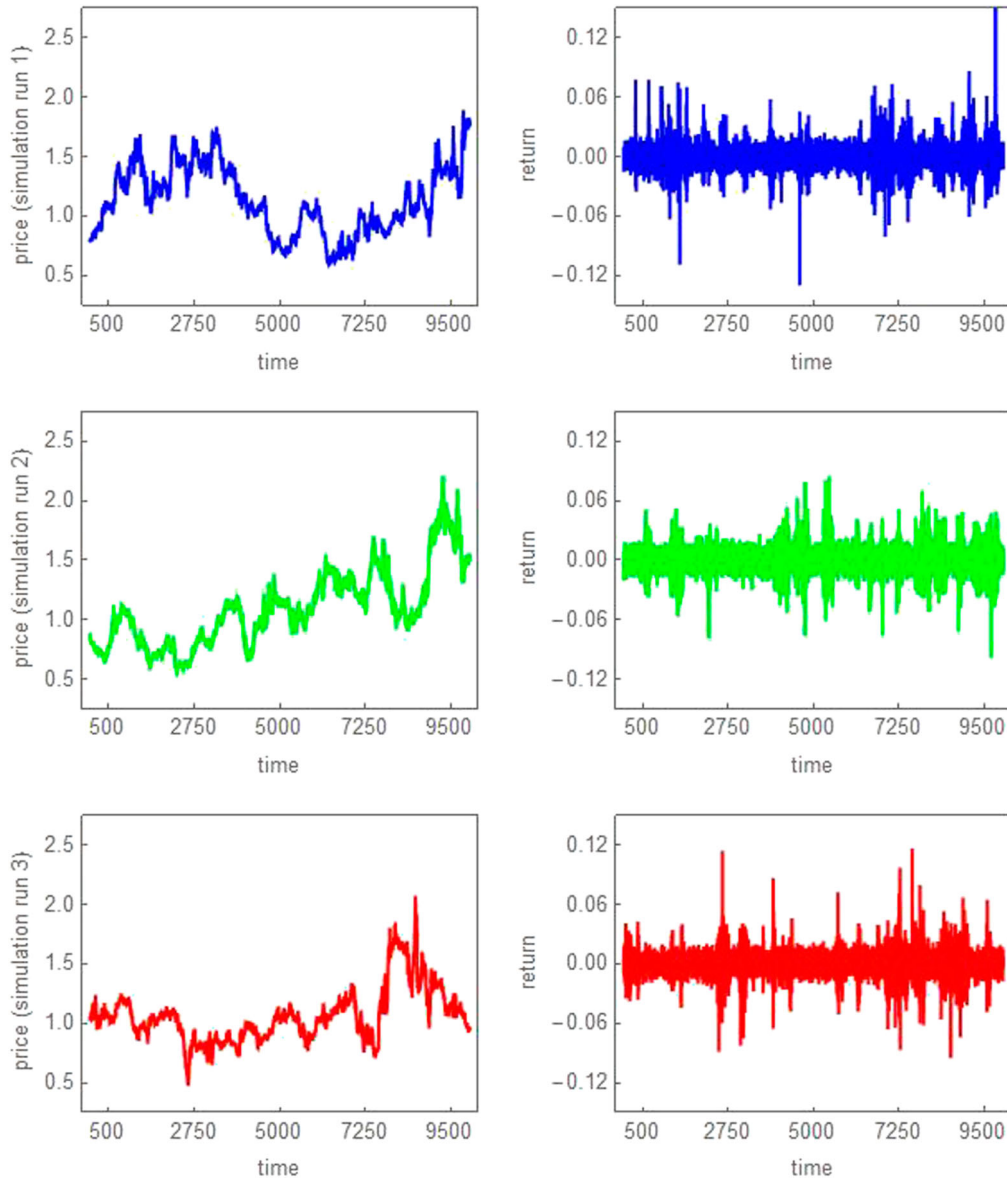


Figure 3. Time series dynamics of simulated stock markets. The left panels show the evolution of three stock market simulations, comprising 10,000 daily observations. The right panels show the corresponding return dynamics. Parameter setting as in Section 3.

and third simulation runs differ only with respect to their random seeds. For comparability reasons, we selected the same layout for Figures 3 and 4 as we did for Figures 1 and 2. The left panels of Figure 3 show the evolution of three simulated stock markets in the time domain. As can be seen, simulated stock prices oscillate around their constant fundamental value, given by $\exp[F] = 1$. The amplitude of the boom-bust dynamics suggests that simulated stock prices tend to be ‘a factor 2’ away from the fundamental value, a relation that is reported by Black (1986), Bouchaud et al. (2017) and Majewski, Ciliberti, and Bouchaud (2020) for actual stock markets, along with evidence that a self-correction of mispricing in stock markets can take several years.⁸ Note that mispricing in the simulated stock market is also quite persistent. The right panels depict the corresponding return dynamics. On average, volatility is quite high in the simulated stock markets. Although the fundamental value is constant,

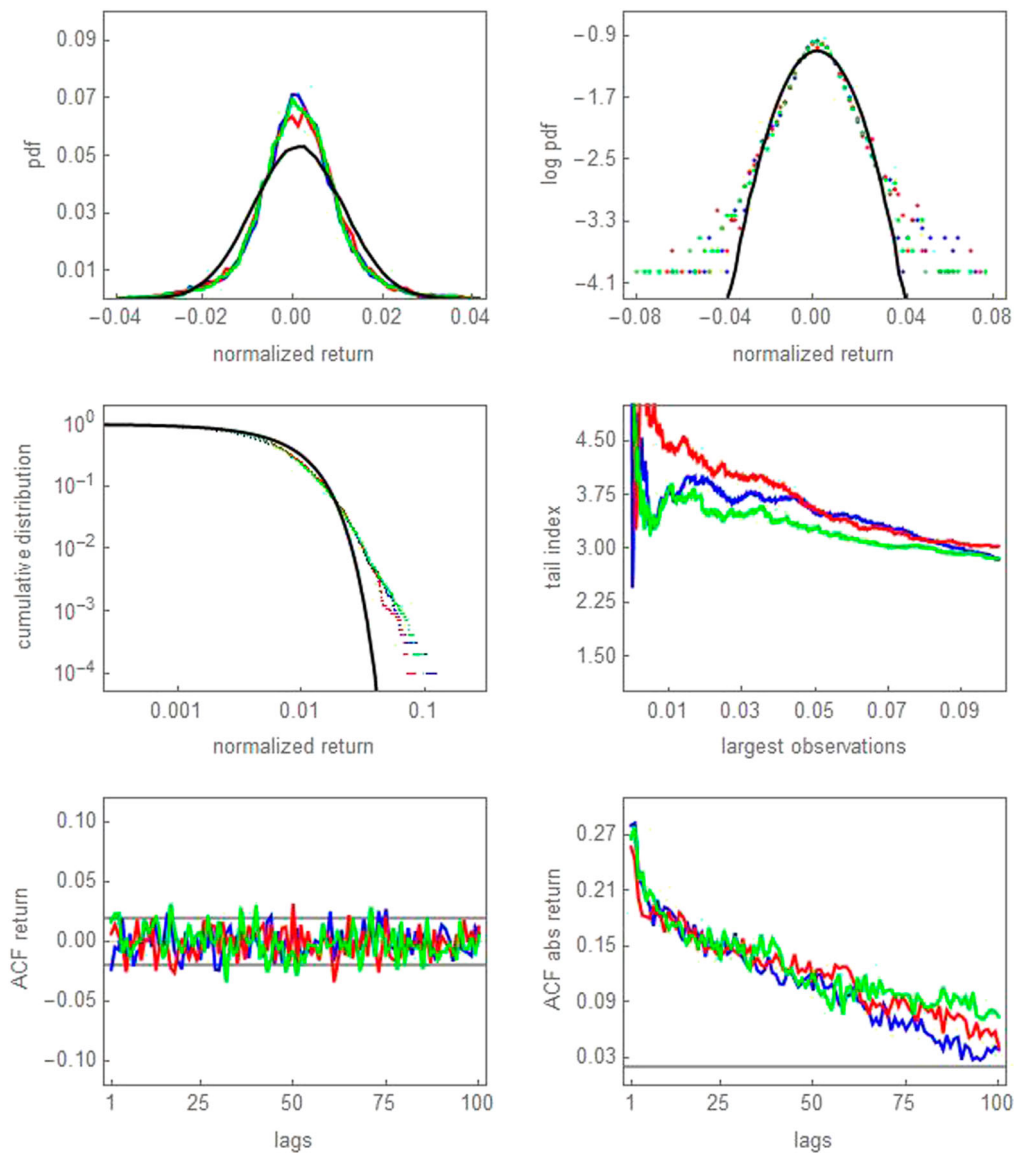


Figure 4. Distributional and correlation properties of simulated stock markets. The panels show a number of distributional and correlation properties of three representative stock market simulations. The same data set and color coding as in Figure 3.

the standard deviations of the three return time series are given by 0.0124 (top), 0.0118 (center) and 0.0117 (bottom), comparable to those reported for the DAX, the NIKKEI and the DJI. The same is true for extreme price changes, given, for instance, by 15.3 percent and -13 percent for the first simulation run.

The first three panels of Figure 4 show that the distributions of simulated stock market returns are bell-shaped, yet possess more probability mass in their tails than warranted by a normal distribution. From the center right panel of Figure 4, we can conclude that the tail indices for the three simulated time series, taking again the largest 5 percent of the returns into account, range between 3.28 and 3.53, only somewhat higher than their empirical counterparts. As revealed by the bottom right panels of Figure 4, returns hardly display any kind of serial correlation, i.e. the evolution of simulated stock markets is close to a random walk. Accordingly, it is difficult to ‘beat the market’, an important (economic) property that holds for actual and simulated stock markets.

The bottom right panel reveals that the autocorrelation coefficients of absolute returns are highly significant, up to 100 lags. Of course, the ability of our simple agent-based computational stock market model to produce volatility clustering is already apparent from its return dynamics, depicted in Figure 3.

The Monte-Carlo study presented in Appendix A1 suggests that we may indeed regard the simulation runs discussed above as representative simulation runs. Overall, we can thus conclude that our simple agent-based computational stock market model is able to match the stylized facts of stock markets in a systematic and robust manner.

4. Functioning of the model

Let us now explain the functioning of our model. Figure 5 depicts a snapshot of the dynamics of the first simulation run (750 observations, ranging from period 4351 to 5100). The left panels show the evolution of simulated stock prices and returns while the right panels show speculators' trading intensity (variance) and their coordination (correlation). Based on these panels, our model's ability to match the stylized facts of stock markets may be understood as follows:

- Bubbles and crashes: The intricate trading behavior of speculators, and in particular their reliance on technical and fundamental trading signals, creates significant bubbles and crashes. As can be seen in the top left panel of Figure 5, for instance, the stock market is overvalued up to around period 250 and then enters a significant bear market. While technical trading tends to drive stock prices away from their fundamental value, fundamental trading exercises a long-run mean reversion pressure.




Figure 5. Functioning of model. The left panels show the evolution of simulated stock prices and returns while the right panels show speculators' trading intensity (variance) and their coordination (correlation). Extract of the first simulation run, as depicted in Figure 3, ranging from period 4351 to 5100.

- Excess volatility: Since the fundamental value of the simulated stock market is constant, we have to regard all stock price changes as excessive. Clearly, once stock prices mirror their fundamental value, there is no need for further stock market adjustments. However, speculators constantly receive new trading signals, which translate into new speculative orders and prompt the market maker to quote new stock prices, as visible in the left panels of Figure 5.
- Serially uncorrelated returns: Due to speculators' heterogeneous trading behavior – each speculator obtains her own individual trading signals, either from private market research or from following complex (algorithmic) trading systems – the path of simulated stock prices closely resembles a random walk, implying that (log) price changes are serially uncorrelated.
- Fat-tailed return distributions: Occasionally, however, we observe a breakdown of speculators' heterogeneity. For instance, salient price patterns may result in panic-induced herding behavior, leading to a spontaneous synchronization of speculators' trading behavior. Moreover, certain price signals may be hard-wired into speculators' complex (algorithmic) trading system, producing coordinated buying or selling behavior. One such example occurs shortly after period 250. As evident from the bottom left panel of Figure 5, the stock market decreases by more than 12 percent. The bottom right panel of Figure 5 illustrates that this event is associated with a strong correlation between speculators' trading signals.⁹
- Volatility clustering: If volatility picks up, speculators extract stronger trading signals out of past price movements. Since this leads to more forceful trading behavior, volatility may remain high. Moreover, speculators may overreact to their trading signals in periods of heightened volatility since they are agitated and thus classify their trading signals as relatively important. Such behavior lends volatility outbursts persistency. In fact, note that in periods when speculators' trading intensity is high (top right panel of Figure 5), the variability of stock prices also tends to be high (bottom left panel of Figure 5).¹⁰

5. Circuit breakers

Understanding the functioning of stock markets is important. In particular, policymakers need to develop a sound economic knowledge of what really drives stock markets if they plan to implement new regulatory measures. Since our model is able to replicate a number of important stylized facts of stock markets, we may use it as an artificial laboratory to study the effects of regulatory policy measures. In this paper, we explore whether policymakers may stabilize the dynamics of stock markets by implementing circuit breakers.¹¹ Circuit breakers (trading halts) automatically interrupt the trading process for a given period of time when price changes are about to exceed a pre-specified limit. Policymakers hope that, by interrupting an overheated market, speculators are given time to cool off and reassess market conditions, enabling the trading process to resume in a more orderly manner after the interruption. Following the stock market crash of 1987, circuit breakers were widely implemented and are now in practice in many leading stock markets around the world. See Kim and Yang (2004) and Sifat and Mohamad (2019) for surveys.

Here we follow Westerhoff (2003, 2006, 2008) and implement circuit breakers as follows. Let parameter s stand for the maximum allowed log price change for a given trading period. Then the market maker's price adjustment rule turns into

$$P_{t+1} = \begin{cases} P_t + s \text{ if } a \sum_{i=1}^N D_{t,i} > s \\ P_t + a \sum_{i=1}^N D_{t,i} \text{ if } -s < a \sum_{i=1}^N D_{t,i} < s \\ P_t - s \text{ if } -s < a \sum_{i=1}^N D_{t,i} \end{cases} . \quad (12)$$

If policymakers set $s = 0.05$, for instance, then the market maker has to interrupt the trading process when the log price is about to either increase or decrease by more than 5 percent. The stock market reopens in the next trading period, i.e. there are no further transactions in a period when trading has been interrupted. For simplicity, we assume that all orders that have not been executed are deleted.




Figure 6. Effects of circuit breakers. The top panels show the evolution of stock prices and returns for $s = 0.05$, respectively. The bottom panels show the stock market's distortion and volatility for $0 < s < 0.1$. Blue, green and red lines are based on $\sigma^F = 0$, $\sigma^F = 0.006$ and $\sigma^F = 0.012$, respectively. Remaining parameters as in Section 3.

Figure 6 depicts a number of possible effects of circuit breakers. The top panels show the evolution of stock prices and returns for $s = 0.05$. For comparability, the simulation run is based on the same random seed as the first simulation run in Figure 3 (top panels, marked blue). First of all, circuit breakers manage to limit extreme returns to 5 percent. However, there are further important effects. The blue lines in the bottom panels of Figure 6 report the stock market's distortion, defined as $dis = 1/T \sum_{t=1}^T |P_t - F|$, and volatility, defined as $vol = 1/T \sum_{t=1}^T |P_t - P_{t-1}|$, for $0 < s < 0.1$. The sample length is set to $T = 100,000$ observations and parameter s is increased in 25 discrete steps. As circuit breakers become more restrictive, both volatility and distortion decline. In the extreme case of $s = 0$, volatility is completely eliminated. If we furthermore assume that the initial value of the stock price is identical to its fundamental value, then circuit breakers also suppress the emergence of any kind of distortion.

Let us briefly explain how circuit breakers affect the model's stock market dynamics. Obviously, circuit breakers have an immediate direct effect: if policymakers set $s = 0.05$, for instance, there will be no stock price change larger than 5 percent. Importantly, however, there are also indirect effects that amplify the direct effect. First, circuit breakers naturally reduce the strength of speculators' technical trading signals by preventing sharp stock price changes. Technically, this effect originates from Equation (5). Second, circuit breakers reduce speculators' trading intensity (variance) by reducing the stock market's volatility, as can be concluded from Equations (6) and (7). Third, circuit breakers prevent (or at least deter) speculators from displaying panic-induced herding behavior and/or from coordinating on certain salient price patterns that are hard-wired into their complex (algorithmic) trading systems, as is evident from Equations (8) and (9).

However, Fama (1989) argues that stock markets are efficient and thus warns that circuit breakers may only lead to a delayed price discovery and to a spillover of volatility. Here, volatility spillover means that a stock market that hits its upper or lower price boundary in the current trading period will experience greater volatility in the

next trading period, since the necessary price adjustment has not yet been fulfilled. Our model allows us to address this issue, at least partially, by assuming that the stock market's fundamental value is not constant, but evolves in the form of a random walk. Accordingly, we specify the stock market's log fundamental value by $F_t = F_{t-1} + n_t$, (13) where the fundamental shocks n_t that hit the stock market are normally distributed with mean zero and constant standard deviation σ^F . The blue, green and red lines depicted in the bottom lines of panels of Figure 6 are computed on the basis of $\sigma^F = 0$, $\sigma^F = 0.006$ and $\sigma^F = 0.012$. As reported in Section 3, the standard deviations of actual and simulated stock markets returns hover around 0.012. Assuming that the stock market's excess volatility is given by a factor of two (Shiller 2015), a reasonable guess for the stock market's fundamental volatility may be given by $\sigma^F = 0.006$. In order to push our analysis to the limit, we also explore the case $\sigma^F = 0.012$.

One important finding of our simulations is that circuit breakers may reduce the stock market's volatility, independently of its fundamental volatility. Another important finding of our simulations is that circuit breakers may increase the stock market's distortion if they are too restrictive. To put it differently, stock markets apparently need some price flexibility, though not a perfect price flexibility. The reason behind this outcome is that circuit breakers prevent technical and fundamental orders. If the fundamental value evolves randomly, at least some fundamental orders are needed for the stock price to be able to track its fundamental value. However, even for $\sigma^F = 0.012$, at least a mild reduction of the stock market's volatility and distortion is possible. Fundamental values are presumably less volatile than implied by $\sigma^F = 0.012$ and thus circuit breakers seem to be a useful tool for policymakers to stabilize stock markets. In this sense, our results contradict the hypothesis of a delayed price discovery process and a volatility spillover, as put forward by Fama (1989). Interestingly, the results presented in Westerhoff (2003, 2006, 2008) are quite similar to ours, despite resting on different stock market models. See also Yeh and Yang (2010, 2013) and Jacob Leal and Napoletano (2019) for more work in this direction.

6. Conclusions

Galbraith (1994), Kindleberger and Aliber (2011) and Shiller (2015) emphatically stress that the boom-bust nature of stock markets as well as their excessively volatile behavior and tendency to produce occasionally very large price changes may be quite harmful to the real economy. In this paper, we therefore develop a simple agent-based computational model that may help us to foster our understanding of the functioning of stock markets. Within our model, stock prices adjust with respect to the excess demand of speculators, who, in turn, derive their trading signals either from private market research or from applying complex (algorithmic) trading systems. Our modeling strategy is inspired by the work of Gode and Sunder (1993), Cont and Bouchaud (2000), Iori (2002) and Alfi et al. (2009) in the sense that we use a rather minimalistic approach to represent speculators' trading behavior. In particular, we formalize speculators' orders via multivariate normally distributed random variables, which allows us to acknowledge speculators' use of technical and fundamental analysis and to condition the intensity and correlation of their trading activities on the stock market's past behavior.

Despite the simplicity of our approach, simulations reveal that our model is able to mimic a number of important stylized facts of stock markets and, consequently, may be deemed to be validated. One crucial model insight is that we may regard stock markets as self-exciting systems. If volatility picks up, speculators trade more aggressively, an outcome that keeps volatility high. Moreover, certain salient price patterns may prompt complex (algorithmic) trading systems to trigger correlated trading signals or may result in panic-induced herding behavior, yielding extreme price changes. Put differently, stock markets display a life of their own and their dynamics contains a larger endogenous component that policymakers may seek to influence. In fact, simulations reveal that policymakers may stabilize the dynamics of stock markets by implementing circuit breakers.

We conclude our paper by pointing out a few avenues for future research. The simplicity of our model allows for a number of straightforward model extensions. For instance, one may try to endogenize the number of (active) speculators, e.g. by considering interactions between different stock markets. Alternatively, one may consider that the correlation of speculators' trading signals does not depend on a single, deterministic condition, but on multiple conditions, possibly time-varying and containing stochastic elements. Although our model contains a larger number of parameters, it might be interesting to try to estimate it. The method of simulated



moments seems to us to be quite appropriate for such an endeavor. We hope that our paper stimulates more work in this important and exciting research direction.

Notes

1. According to Murphy (1999), the reliability of technical trading signals increases with the trading volume of a stock market, i.e. a high trading volume indicates that the current trading signal is strong whereas a low trading volume indicates that the current trading signal is weak. Since simulations reveal that our model produces a high contemporaneous correlation between trading volume and volatility, an interesting model extension could be to condition speculators' trading intensity on the trading volume of the stock market. See Westerhoff (2006) for an example in that direction.
2. A well-known example in this respect concerns the stock market crash of October 1987, which, according to Greenwald and Stein (1991), Harris (1998) and Shiller (2015), was at least partially triggered by computer (program) trading, and could have been stopped by circuit breakers. More recent examples include the occurrence of so-called flash crashes, amplified by high-frequency traders who follow computerized trading systems. See Jacob Leal et al. (2016) and Jacob Leal and Napoletano (2019) for empirical evidence and interesting modeling approaches. Gomber and Zimmermann (2018) and Vassiliadis and Dounias (2018) provide insightful overviews of complex (algorithmic) trading systems.
3. As we will see in the next section, however, one condition may already be sufficient for our model to produce extreme price changes and, consequently, fat-tailed return distributions. We remark that we also experimented with other conditions. For instance, similar dynamics to those discussed in the next section may be observed if (8) is replaced by $C_t = gC_{t-1} + (1-g)(P_t - P_{t-1})^2$, where $0 < g < 1$ is a memory parameter. In relation to (6), however, our simulations suggest that the memory parameter has to be set to a rather low value, say $g = 0.05$, implying that coordination among market participants critically hinges on the stock market's short-run behavior. To save one parameter, we opted for specification (8). Of course, this aspect deserves more attention in future work, in particular along the lines indicated above.
4. The excess demand also increases with the number of speculators. For $a = \alpha/N$ and $N \rightarrow \infty$, however, the price adjustment equation reads $P_{t+1} = P_t + \alpha \{ (b(P_t - P_{t-1}) + c(F - P_t)) + \sigma_t \sqrt{\rho_t} \varepsilon_t \}$. Hence, it is possible to rescale our model such that its dynamics does not depend on the number of speculators. While we prefer to keep N as a model parameter, it might be worthwhile to try to endogenize the number of (active) speculators in future work. See Iori (2002), Alfi et al. (2009), Blaurock, Schmitt, and Westerhoff (2018) and Dieci, Schmitt, and Westerhoff (2018) for examples in this direction.
5. Bubbles and crashes are difficult to identify in real stock markets. However, Galbraith (1994), Kindleberger and Aliber (2011) and Shiller (2015) stress that bubbles and crashes do exist in these markets. See Schmitt and Westerhoff (2017c) and Majewski, Ciliberti, and Bouchaud (2020) for attempts on how to capture the mispricing of actual stock markets.
6. Such estimates are representative for many different financial markets, see, e.g. Gopikrishnan et al. (1999) and Plerou et al. (1999).
7. Of course, other estimation methods may also be useful, see, e.g., the work by Lamperti, Roventini, and Sani (2018), Platt (2020), Kukacka and Kristoufek (2020) and Bertschinger and Mozhorin (2020).
8. The famous "factor 2" rule by Black (1986, 533) implies that the stock "price is more than half of value and less than twice value". For our case, simulated stock prices should thus fluctuate in the interval $0.5 < F = 1 < 2$.
9. Note that a high correlation between speculators' trading behavior does not always lead to a strong stock price change. For this to be the case, speculators have to coordinate on a significant trading signal.
10. Note that speculators' trading intensity (variance) may remain high for extended periods of time, thereby producing lasting volatility outbreaks, while their coordination (correlation) spikes only occasionally, forming the base for rare but extreme returns. We discuss this aspect in more detail in Appendix A.2.
11. As pointed out by an anonymous referee, it might also be worthwhile to use our model to study the effects of margin requirements, leverage cycles and short-selling constraints. For inspiring work in this direction, see, for instance, Poledna et al. (2014), Aymanns et al. (2016) and Sng, Zhang, and Zheng (2020). Aymanns et al. (2018) and Westerhoff and Franke (2018) discuss in more detail how policymakers may use models with heterogeneous interacting agents as test beds to evaluate the effectiveness of regulatory policies.
12. Our model is also able to generate extreme returns for lower values of $\sigma_t^2 = 0.000003$. However, we found that this number matches actual tail indices quite well.

Acknowledgement

We would like to thank two anonymous referees and Giulia Iori for their valuable feedback. This research was carried out in the Bamberg Doctoral Research Group on Behavioral Macroeconomics (BaGBeM) supported by the Hans-Böckler Foundation.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

- Alfi, V., M. Cristelli, L. Pietronero, and A. Zaccaria. 2009. "Minimal Agent Based Model for Financial Markets I." *The European Physical Journal B* 67: 385–397.
- Anufriev, M., T. Bao, and J. Tuinstra. 2016. "Microfoundations for Switching Behavior in Heterogeneous Agent Models: an Experiment." *Journal of Economic Behavior and Organization* 129: 74–99.
- Arthur, B., J. Holland, B. LeBaron, R. Palmer, and P. Tayler. 1997. "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market." In *The Economy as an Evolving Complex System II*, edited by B. Arthur, S. Durlauf, and D. Lane, 15–44. Addison-Wesley: Addison-Wesley, Reading.
- Aymanns, C., F. Caccioli, D. Farmer, and V. Tan. 2016. "Taming the Basel Leverage Cycle." *Journal of Financial Stability* 27: 263–277.
- Aymanns, C., D. Farmer, A. Kleinnijenhuis, and T. Wetzler. 2018. "Models of Financial Stability and Their Application in Stress Tests." In *Handbook of Computational Economics: Heterogeneous Agent Modeling*, edited by C. Hommes, and B. LeBaron, 329–391. Amsterdam: North-Holland.
- Beja, A., and M. Goldman. 1980. "On the Dynamic Behaviour of Prices in Disequilibrium." *The Journal of Finance* 35: 235–248.
- Bertschinger, N., and I. Mozhorin. 2020. "Bayesian Estimation and Likelihood-Based Comparison of Agent-Based Volatility Models." *Journal of Economic Interaction and Coordination*. doi:10.1007/s11403-020-00289-z.
- Black, F. 1986. "Noise." *The Journal of Finance* 41: 528–543.
- Blaurock, I., N. Schmitt, and F. Westerhoff. 2018. "Market Entry Waves and Volatility Out-Bursts in Stock Markets." *Journal of Economic Behavior and Organization* 153: 19–37.
- Bouchaud, J.-P., S. Ciliberti, Y. Lemperiere, A. Majewski, P. Seager, and K. Sin Ronia. 2017. Black was Right: Price is Within a Factor 2 of Value. SSRN <https://ssrn.com/abstract=3070850>.
- Brock, W., and C. Hommes. 1998. "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model." *Journal of Economic Dynamics and Control* 22: 1235–1274.
- Chang, I., and D. Stauffer. 1999. "Fundamental Judgement in Cont-Bouchaud Herding Model of Financial Fluctuations." *Physica A: Statistical Mechanics and its Applications* 264: 294–298.
- Chen, S.-H., and C.-H. Yeh. 2001. "Evolving Traders and the Business School with Genetic Programming: A New Architecture of the Agent-Based Artificial Stock Market." *Journal of Economic Dynamics and Control* 25: 363–393.
- Chiarella, C. 1992. "The Dynamics of Speculative Behavior." *Annals of Operations Research* 37: 101–123.
- Chiarella, C., and G. Iori. 2002. "A Simulation Analysis of the Microstructure of Double Auction Markets." *Quantitative Finance* 2: 346–353.
- Cont, R. 2001. "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues." *Quantitative Finance* 1: 223–236.
- Cont, R., and J.-P. Bouchaud. 2000. "Herd Behavior and Aggregate Fluctuations in Financial Markets." *Macroeconomic Dynamics* 4: 170–196.
- Daniels, M., D. Farmer, L. Gillemot, G. Iori, and E. Smith. 2003. "Quantitative Model of Price Diffusion and Market Friction Based on Trading as a Mechanistic Random Process." *Physical Review Letters* 90: 108102.
- Day, R., and W. Huang. 1990. "Bulls, Bears and Market Sheep." *Journal of Economic Behavior and Organization* 14: 299–329.
- De Grauwe, P., H. Dewachter, and M. Embrechts. 1993. *Exchange Rate Theory: Chaotic Models of Foreign Exchange Markets*. Oxford: Blackwell.
- Delli Gatti, D., G. Fagiolo, M. Gallegati, M. Richiardi, and A. Russo. 2018. *Agent-Based Models in Economics: A Toolkit*. Cambridge: Cambridge University Press.
- Demirer, R., G. Demos, R. Gupta, and D. Sornette. 2019. "On the Predictability of Stock Market Bubbles: Evidence from LPPLS Confidence Multi-Scale Indicators." *Quantitative Finance* 19: 843–858.
- Dieci, R., and X.-Z. He. 2018. "Heterogeneous Agent Models in Finance." In *Handbook of Computational Economics: Heterogeneous Agent Modeling*, edited by C. Hommes, and B. LeBaron, 257–328. Amsterdam: North-Holland.
- Dieci, R., N. Schmitt, and F. Westerhoff. 2018. "Interactions between Stock, Bond and Housing Markets." *Journal of Economic Dynamics and Control* 91: 43–70.
- Diem, C., A. Pichler, and S. Thurner. 2020. "What is the Minimal Systemic Risk in Financial Exposure Networks?" *Journal of Economic Dynamics and Control* 116: 103900.
- Fagiolo, G., M. Guerini, F. Lamperti, A. Moneta, and A. Roventini. 2017. "Validation of Agent-Based Models in Economics and Finance." In *Computer Simulation Validation*, edited by C. Beisbart, and N. Saam, 763–787. Berlin: Springer.
- Fama, E. 1989. "Perspectives on October 1987, or, What Did we Learn from the Crash?" In *Black Monday and the Future of Financial Markets*, edited by R. Kampuis, R. Kormendi, and J. Watson, 71–82. Homewood: Irwin.
- Farmer, D., L. Gillemot, G. Iori, S. Krishnamurthy, E. Smith, and M. Daniels. 2005b. "A Random Order Placement Model of Price Formation in the Continuous Double Auction." In *The Economy as an Evolving Complex System, III: Current Perspectives and Future Directions*, edited by L. Blume, and S. Durlauf, 133–173. Oxford: Oxford University Press.
- Farmer, D., and S. Joshi. 2002. "The Price Dynamics of Common Trading Strategies." *Journal of Economic Behavior and Organization* 49: 149–171.
- Farmer, D., P. Patelli, and I. Zovko. 2005a. "The Predictive Power of Zero Intelligence in Financial Markets." *Proceedings of the National Academy of Sciences* 102: 2254–2259.
- Franke, R., and F. Westerhoff. 2012. "Structural Stochastic Volatility in Asset Pricing Dynamics: Estimation and Model Contest." *Journal of Economic Dynamics and Control* 36: 1193–1211.

- Franke, R., and F. Westerhoff. 2016. "Why a Simple Herding Model May Generate the Stylized Facts of Daily Returns: Explanation and Estimation." *Journal of Economic Interaction and Coordination* 11: 1–34.
- Galbraith, J. K. 1994. *A Short History of Financial Suphoria*. London: Penguin Books.
- Gode, D., and S. Sunder. 1993. "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality." *Journal of Political Economy* 101: 119–137.
- Gode, D., and S. Sunder. 1997. "What Makes Markets Allocationally Efficient?" *The Quarterly Journal of Economics* 112: 603–630.
- Gomber, P., and K. Zimmermann. 2018. "Algorithmic Trading in Practice." In *The Oxford Handbook on Computational Economics and Finance*, edited by S.-H. Chen, M. Kaboudan, and Y.-R. Du, 311–332. Oxford: Oxford University Press.
- Gopikrishnan, P., V. Plerou, L. Amaral, M. Meyer, and E. Stanley. 1999. "Scaling of the Distributions of Fluctuations of Financial Market Indices." *Physical Review E* 60: 5305–5316.
- Graham, B., and D. Dodd. 1951. *Security Analysis*. New York: McGraw-Hill.
- Greenwald, B., and J. Stein. 1991. "Transactional Risk, Market Crashes, and the Role of Circuit Breakers." *The Journal of Business* 64: 443–462.
- Guerini, M., and A. Moneta. 2017. "A Method for Agent-Based Models Validation." *Journal of Economic Dynamics and Control* 82: 125–141.
- Harris, L. 1998. "Circuit Breakers and Program Trading Limits: What Have we Learned?" In *Brookings-Wharton Papers on Financial Services*, edited by R. Litan, and A. Santomero, 17–64. Washington: Brookings Institution Press.
- Hill, B. 1975. "A Simple General Approach to Inference About the Tail of a Distribution." *The Annals of Statistics* 3: 1163–1174.
- Hommes, C. 2011. "The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab." *Journal of Economic Dynamics and Control* 35: 1–24.
- Iori, G. 2002. "A Microsimulation of Traders Activity in the Stock Market: The Role of Heterogeneity, Agents' Interactions and Trade Frictions." *Journal of Economic Behavior and Organization* 49: 269–285.
- Iori, G., and J. Porter. 2018. "Agent-based Modeling for Financial Markets." In *The Oxford Handbook of Computational Economics and Finance*, edited by S.-H. Chen, M. Kaboudan, and Y.-R. Du, 635–666. Oxford: Oxford University Press.
- Jacob Leal, S., and M. Napoletano. 2019. "Market Stability vs. Market Resilience: Regulatory Policies Experiments in an Agent-Based Model with Low- and High-Frequency Trading." *Journal of Economic Behavior and Organization* 157: 15–41.
- Jacob Leal, S., M. Napoletano, A. Roventini, and G. Fagiolo. 2016. "Rock Around the Clock: An Agent-Based Model of Low- and High-Frequency Trading." *Journal of Evolutionary Economics* 26: 49–76.
- Keynes, J. M. 1936. *The General Theory of Employment, Interest, and Money*. New York: Harcourt, Brace and Company.
- Kim, Y., and J. Yang. 2004. "What Makes Circuit Breakers Attractive to Financial Markets? A Survey." *Financial Markets, Institutions and Instruments* 13: 109–146.
- Kindleberger, C., and R. Aliber. 2011. *Manias, Panics, and Crashes: A History of Financial Crises*. New Jersey: Wiley.
- Kukacka, J., and L. Kristoufek. 2020. "Do 'Complex' Financial Models Really Lead to Complex Dynamics? Agent-Based Models and Multifractality." *Journal of Economic Dynamics and Control* 113: 103855.
- Ladley, D. 2012. "Zero Intelligence in Economics and Finance." *The Knowledge Engineering Review* 27: 273–286.
- Ladley, D. 2020. "The High Frequency Trade Off between Speed and Sophistication." *Journal of Economic Dynamics and Control* 116: 103912.
- Ladley, D., T. Lensberg, J. Palczewski, and K. R. Schenk-Hoppe. 2015. "Fragmentation and Stability of Markets." *Journal of Economic Behavior and Organization* 119: 466–481.
- Lamperti, F., A. Roventini, and A. Sani. 2018. "Agent-Based Model Calibration Using Machine Learning Surrogates." *Journal of Economic Dynamics and Control* 90: 366–389.
- LeBaron, B., B. Arthur, and R. Palmer. 1999. "Time Series Properties of an Artificial Stock Market." *Journal of Economic Dynamics and Control* 23: 1487–1516.
- Lo, A., H. Mamaysky, and J. Wang. 2000. "Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation." *The Journal of Finance* 55: 1705–1765.
- Lux, T. 1995. "Herd Behaviour, Bubbles and Crashes." *The Economic Journal* 105: 881–896.
- Lux, T., and M. Ausloos. 2002. "Market Fluctuations I: Scaling, Multiscaling, and Their Possible Origins." In *Science of Disaster: Climate Disruptions, Heart Attacks, and Market Crashes*, edited by A. Bunde, J. Kropp, and H. Schellnhuber, 373–410. Berlin: Springer.
- Lux, T., and R. Zwinkels. 2018. "Empirical Validation of Agent-Based Models." In *Handbook of Computational Economics: Heterogeneous Agent Modeling*, edited by C. Hommes, and B. LeBaron, 437–482. Amsterdam: North-Holland.
- Majewski, A., S. Ciliberti, and J.-P. Bouchaud. 2020. "Co-Existence of Trend and Value in Financial Markets: Estimating an Extended Chiarella Model." *Journal of Economic Dynamics and Control* 112: 103791.
- Mantegna, R., and E. Stanley. 2000. *An Introduction to Econophysics*. Cambridge: Cambridge University Press.
- Manzan, S., and F. Westerhoff. 2005. "Representativeness of News and Exchange Rate Dynamics." *Journal of Economic Dynamics and Control* 29: 677–689.
- Menkhoff, L., and M. Taylor. 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: New York Institute of Finance.
- Palmer, R., B. Arthur, J. Holland, B. LeBaron, and P. Tayler. 1994. "Artificial Economic Life: A Simple Model of a Stock Market." *Physica D: Nonlinear Phenomena* 75: 264–274.

- Platt, D. 2020. "A Comparison of Economic Agent-Based Model Calibration Methods." *Journal of Economic Dynamics and Control* 113: 103859.
- Plerou, V., P. Gopikrishnan, L. Amaral, M. Meyer, and E. Stanley. 1999. "Scaling of the Distribution of Price Fluctuations of Individual Companies." *Physical Review E* 60: 6519–6529.
- Poledna, S., S. Thurner, D. Farmer, and J. Geanakoplos. 2014. "Leverage-Induced Systemic Risk Under Basle II and Other Credit Risk Policies." *Journal of Banking and Finance* 42: 199–212.
- Raberto, M., S. Cincotti, S. Focardi, and M. Marchesi. 2001. "Agent-Based Simulation of a Financial Market." *Physica A: Statistical Mechanics and its Applications* 299: 319–327.
- Schmitt, N. 2020. "Heterogeneous Expectations and Asset Price Dynamics." *Macroeconomic Dynamics*. doi:10.1017/S1365100519000774.
- Schmitt, N., F. Tramontana, and F. Westerhoff. 2020. "Nonlinear Asset-Price Dynamics and Stabilization Policies." *Nonlinear Dynamics*. doi:10.1007/s11071-020-05828-8.
- Schmitt, N., and F. Westerhoff. 2017a. "Heterogeneity, Spontaneous Coordination and Extreme Events Within Large-Scale and Small-Scale Agent-Based Financial Market Models." *Journal of Evolutionary Economics* 27: 1041–1070.
- Schmitt, N., and F. Westerhoff. 2017b. "Herding Behaviour and Volatility Clustering in Financial Markets." *Quantitative Finance* 17: 1187–1203.
- Schmitt, N., and F. Westerhoff. 2017c. "On the Bimodality of the Distribution of the S&P 500's Distortion: Empirical Evidence and Theoretical Explanations." *Journal of Economic Dynamics and Control* 80: 34–53.
- Shiller, R. 2015. *Irrational Exuberance*. Princeton: Princeton University Press.
- Sifat, I., and A. Mohamad. 2019. "Circuit Breakers As Market Stability Levers: A Survey of Research, Praxis, and Challenges." *International Journal of Finance and Economics* 24: 1130–1169.
- Sng, H., Y. Zhang, and H. Zheng. 2020. "Margin Trade, Short Sales and Financial Stability." *Journal of Economic Interaction and Coordination* 15: 673–702.
- Stanek, F., and J. Kukacka. 2018. "The Impact of the Tobin Tax in a Heterogeneous Agent Model of the Foreign Exchange Market." *Computational Economics* 51: 865–892.
- Stauffer, D., and N. Jan. 2000. "Sharp Peaks in the Percolation Model for the Stock Markets." *Physica A: Statistical Mechanics and its Applications* 277: 215–219.
- Stauffer, D., and T. Penna. 1998. "Crossover in the Cont-Bouchaud Percolation Model for Market Fluctuations." *Physica A: Statistical Mechanics and its Applications* 256: 284–290.
- Stauffer, D., and D. Sornette. 1999. "Self-Organized Percolation Model for Stock Market Fluctuation." *Physica A: Statistical Mechanics and its Applications* 271: 496–506.
- Vassiliadis, V., and G. Dounias. 2018. "Algorithmic Trading Based on Biologically Inspired Algorithms." In *The Oxford Handbook on Computational Economics and Finance*, edited by S.-H. Chen, M. Kaboudan, and Y.-R. Du, 295–310. Oxford: Oxford University Press.
- Westerhoff, F. 2003. "Speculative Markets and the Effectiveness of Price Limits." *Journal of Economic Dynamics and Control* 28: 493–508.
- Westerhoff, F. 2006. "Technical Analysis Based on Price-Volume Signals and the Power of Trading Breaks." *International Journal of Theoretical and Applied Finance* 9: 227–244.
- Westerhoff, F. 2008. "The Use of Agent-Based Financial Market Models to Test the Effectiveness of Regulatory Policies." *Jahrbücher für Nationalökonomie und Statistik (Journal of Economics and Statistics)* 228: 195–227.
- Westerhoff, F., and R. Franke. 2018. "Agent-Based Models for Policy Analysis: Two Illustrative Examples." In *The Oxford Handbook of Computational Economics and Finance*, edited by S.-H. Chen, M. Kaboudan, and Y.-R. Du, 520–558. Oxford: Oxford University Press.
- Westphal, R., and D. Sornette. 2020. "Market Impact and Performance of Arbitrageurs of Financial Bubbles in an Agent-Based Model." *Journal of Economic Behavior and Organization* 171: 1–23.
- Xing, G., and D. Ladley. 2019. Noise Trading and Market Stability. SSRN <https://ssrn.com/abstract=3415141>.
- Yeh, C.-H., and C.-Y. Yang. 2010. "Examining the Effectiveness of Price Limits in an Artificial Stock Market." *Journal of Economic Dynamics and Control* 34: 2089–2108.
- Yeh, C.-H., and C.-Y. Yang. 2013. "Do Price Limits Hurt the Market?" *Journal of Economic Interaction and Coordination* 8: 125–153.
- Zeeman, E. C. 1974. "On the Unstable Behaviour of Stock Exchanges." *Journal of Mathematical Economics* 1: 39–49.
- Zhang, Q., D. Sornette, and H. Zhang. 2019. "Anticipating Critical Transitions of the Housing Market: New Evidence from China." *The European Journal of Finance* 25: 1251–1276.

Appendix A. Robustness analysis

The robustness analysis we carry out in this appendix consists of two parts. In Appendix A1, we first conduct a Monte Carlo study to demonstrate that the simulation runs presented in the main body of our paper may in fact be deemed as representative simulation runs. In Appendix A2, we then conduct a sensitivity analysis to explain in more detail how certain building blocks of our model may affect its dynamics.

Appendix A1. Monte Carlo study

Our Monte Carlo study rests on 5000 simulation runs with 10,000 observations each, generated with the parameter setting introduced in Section 3 and different random seeds. Based on these simulations, Figure A1 shows probability density functions for volatility, distortion, the tail index at the 5 percent level, the autocorrelation coefficient of raw returns at lag 1 and the autocorrelation coefficients of absolute returns at lag 5 and at lag 95, respectively. Using these summary statistics (moments), we seek to capture a number of important stylized facts of stock markets, as discussed in Sections 3 and 4.

Our measure of volatility is defined as in Section 5, i.e. $vol = 1/T \sum_{t=1}^T |P_t - P_{t-1}|$. As can be seen from the top left panel of Figure A1, our volatility estimates hover around a value of about 0.0082. Further computations reveal that 90 percent of the volatility estimates are located in the range 0.0080 and 0.0096. To put these numbers into perspective, note that the volatility estimate for the DAX, at 0.0093, fits nicely into this interval. The top right panel of Figure A1 portrays the probability density function of the simulated stock markets' distortion, defined as $dis = 1/T \sum_{t=1}^T |P_t - F|$. Apparently, the average mispricing of simulated stock




Figure A1. Monte Carlo study. The panels show probability density functions for volatility, distortion, the tail index, the autocorrelation coefficient of raw returns at lag 1 and the autocorrelation coefficients of absolute returns at lag 5 and lag 95, respectively, based on 5000 simulation runs with 10,000 observations each. Parameter setting as in Section 3.

markets is usually above 10 percent, and can easily increase to as much as 30 percent or more. While we cannot compute the distortion for the time series discussed in Section 3, we remark that Schmitt and Westerhoff (2017b) report that the distortion of the S&P500 between 1871 and 2015 was about 30 percent.

The center left panel of Figure A1 depicts the probability density function of our estimates of the tail index (at the 5 percent level). The median estimate is 3.55, while the 90 percent confidence interval ranges from 3.2 to 3.88. According to Lux and Ausloos (2002), the tail indices for most financial market data scatter between 3 and 4. In this sense, our simple agent-based computational model is able to replicate the fat-tail property of stock market returns (though Gopikrishnan et al. 1999 and Plerou et al. 1999 stress that the tail indices of major stock markets are somewhat closer to 3). The center right panel of Figure A1 reveals that the estimated autocorrelation coefficients of raw returns at lag 1 are near zero. To be more precise, 90 percent of the estimated autocorrelation

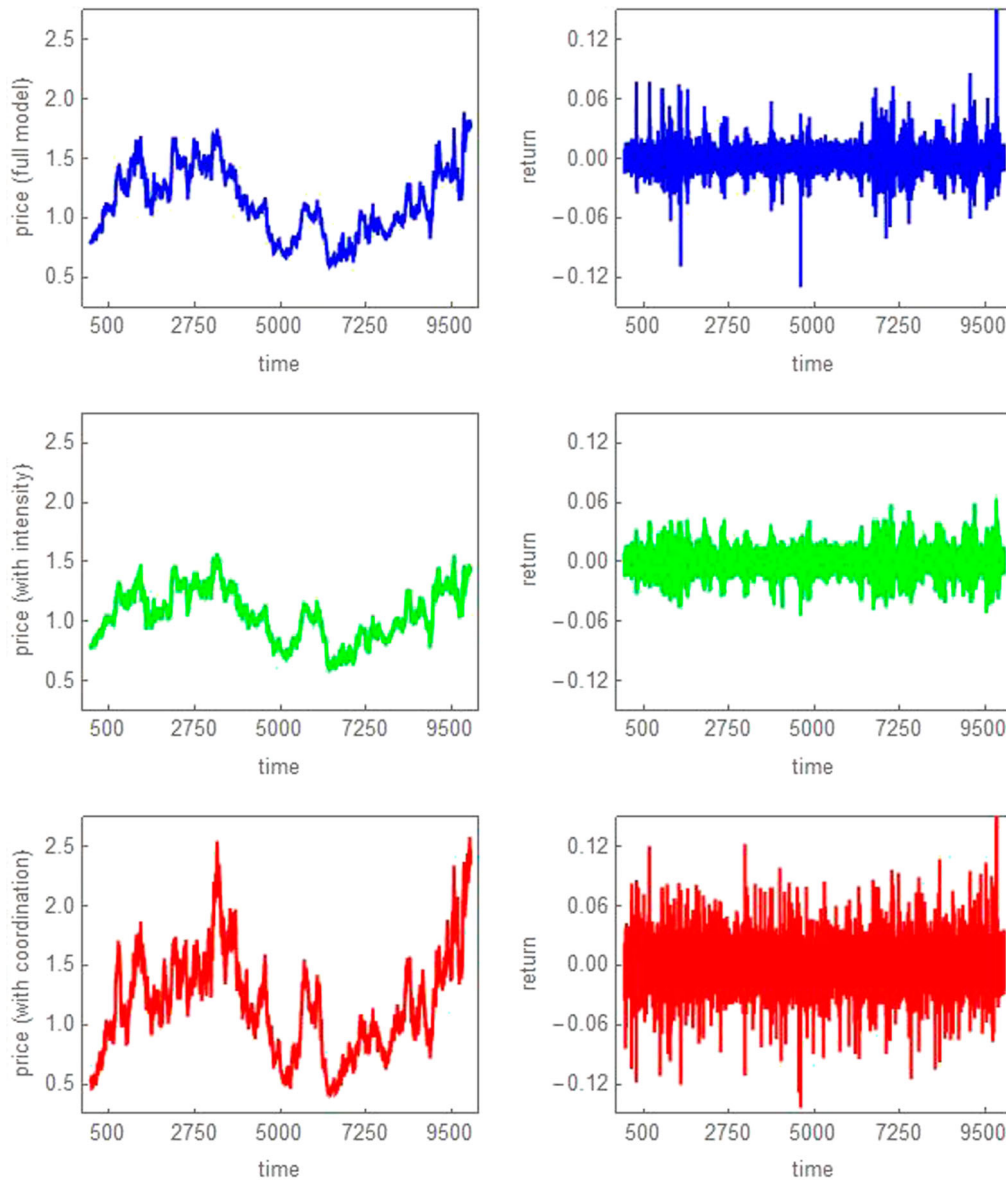


Figure A2. Sensitivity analysis. The left and right panels show simulated stock prices and return dynamics for 10,000 time steps, respectively. Top: full model. Center: model without coordination mechanism, i.e. $\rho_t = 0.0006$. Bottom: model without trading intensity mechanism, i.e. $\sigma_t^2 = 0.000003$. Remaining parameters as in Section 3.



coefficients fall into the interval -0.03 and 0.03 , with a median estimate of about 0.01 , implying that the paths of simulated stock prices are indeed close to random walks.

The bottom two panels show probability density functions for the autocorrelation coefficients of absolute returns at lag 5 and lag 95, respectively, demonstrating our model's ability to generate volatility clustering and long memory effects. For instance, 90 percent of the estimated autocorrelation coefficients of absolute returns at lag 5 are between 0.7 and 0.22 , while more than 95 percent of the estimated autocorrelation coefficients of absolute returns at lag 95 are still larger than 0.04 . Hence, the volatility outbursts produced by our model are quite persistent.

Appendix A2. Sensitivity analysis

Finally, we outline how certain building blocks of our model may affect its dynamics. For convenience, the top panels of Figure A2 show simulated stock prices and the corresponding return dynamics for the full model, using the same parameter setting and random seed as in the top panels of Figure 3. As demonstrated above, our model is able to replicate key stylized facts of stock markets, in particular fat-tailed return distributions and volatility clustering. The center panels of Figure A2 report the dynamics of our model when its coordination mechanism is switched off (achieved by setting $\rho_t = f^l = 0.0006$). Apparently, the model is still able to produce lasting volatility outbursts, yet its ability to generate extreme returns diminishes. The bottom panels of Figure A2 present the dynamics of the model when the trading intensity mechanism is switched off (we now fix $\sigma_t^2 = 0.000003$).¹² Obviously, our model is able to generate extreme returns, yet its ability to produce lasting volatility outbursts is basically gone. More precisely, the 90 percent confidence intervals of simulated autocorrelation coefficients of absolute returns at lag 5 and at lag 95 are 0.16 and 0.21 and 0.03 and 0.09 for the model without the coordination mechanism and 0.00 and 0.04 and -0.02 and 0.02 without the trading intensity mechanism. Moreover, the 90 percent confidence intervals for the tail index (at the 5 percent level) are 3.73 and 4.59 for the model without the coordination mechanism and 3.26 and 3.78 without the trading intensity mechanism. To conclude, our simple agent-based computational model is only able to match the stylized facts of stock markets when the coordination and the trading intensity mechanism act together.

4 Estimation of agent-based models: Testing and applying a simulated joint moment approach

This chapter contains single-authored work. In the following, the working paper version will be included.

Estimation of agent-based models

Testing and applying a simulated joint moment approach

Ivonne Schwartz*

University of Bamberg, Department of Economics, Germany

June 18, 2022

Abstract

The growing literature on the estimation of agent-based models offers new sophisticated techniques which often demand immense computational resources. In this paper, we advocate for estimation methods that provide intuitive understanding, straightforward implementation and focus on the feasibility of empirical applications. We investigate the performance of a simulated joint moment estimator first proposed by Franke and Westerhoff (2012). Applying it on a simple agent-based model, we thoroughly explore the behavior of the response surface of the objective function and test the estimator's ability to correctly recover model parameters. We run large Monte Carlo studies, enhanced by machine-learning surrogates, to study the estimator's properties and report common pitfalls. Key findings indicate that biased parameter estimates come from strong correlations between model parameters.

Keywords

Estimation; agent-based models; model validation; method of simulated moments.

JEL Classification

C15; C52; C63; G15.

* ivonne.schwartz@mailbox.org This research was carried out in the Bamberg Doctoral Research Group on Behavioral Macroeconomics (BaGBeM) and received funding from the Hans-Böckler-Stiftung (PK 045).

1 Introduction

With the global meltdown of 2007-2009, our limited understanding of the workings of financial markets became apparent and traditional economic models have increasingly faced criticism because of their flawed assumptions. Consequently, alternative approaches emerged focusing more on the behavioral dynamics of financial markets. Among them is the rapidly growing literature on agent-based (computational) models (ABM). ABM are abandoning the traditional paradigm of the representative agent framework by looking at the interactions of boundedly rational speculators who are heterogeneous in their expectations and behavior.¹ These models have proven to help understand the dynamics of financial markets since many of them are able to replicate an important number of stylized facts. The literature is rich and reviews can be found in Hommes (2006), LeBaron (2006), Chiarella et al. (2009), Lux (2009), Delli Gatti et al. (2018), Dieci and He (2018) and Iroi and Porter (2018) among others.

Despite their proven potential for better explanation of financial crises, the acceptance of agent-based financial market models is not yet widespread even in the economics profession. To increase acceptance, it is crucial that we have models that are both powerful yet simple and can be econometrically tested and verified. In order to create policy-oriented models that can compete with more traditional models, we have to bring our models to the data. However, given the specific properties of ABM, this is not a trivial task. In fact, most ABM are characterized by high-dimensionality, inherent non-linearity, stochasticity and lack of closed-form expressions. Hence, standard estimation techniques (e.g. maximum likelihood, ordinary least squares, generalized method of moments) are often applicable only partially if at all.² Investigating the micro-macro relationships and causal mechanisms requires different methods and techniques. Simulation-based techniques are among the preferred choices, yet they vary much in their details. Following pioneering contributions by McFadden (1989), Pakes and Pollard (1989), Lee and Ingram (1991) and Duffie and Singleton (1993), the simulated method of moments (SMM) has made its path towards the estimation of economic ABM, see e.g. seminal works by Gilli and Winker (2003), Winker et al. (2007), Franke (2009) and Franke and Westerhoff (2011, 2012, 2016).³ In general, the goal of any SMM is

¹This family of behavioral models is often also referred to as heterogeneous agent models (HAM) or agent-based computational economics (ACE). We will here stick to ABM.

²Exceptions are for example given by Alfarano et al. 2005, 2006, 2007 (maximum likelihood), Boswijk et al. 2007 (nonlinear least square) and Amilon 2008 (efficient method of moments).

³This method is also known under the name of 'method of simulated moments' (MSM).

to find parameter values that minimize the distance between a set of empirical and simulated summary statistics, i.e. moments. Besides such a SMM standard version with a quadratic loss function, Franke and Westerhoff (2012) propose the concept of a joint moment coverage ratio (JMCR). This JMCR aims at maximizing the fraction of simulation runs for which all simulated moments jointly fall into the 95 percent confidence intervals of their empirical counterparts. Franke and Westerhoff (2012) report that the joint moment coverage ratio’s informational value differs from that of the traditional SMM leading to different parameter estimates. Since this approach is not overly technical, very transparent and offers an intuitive understanding of the estimates’ validity, it has often been used for the calibration of stock market ABM (see e.g. Schmitt and Westerhoff 2017a,b or Schmitt 2021).

The first goal of this paper is to contribute to the blooming research on the empirical validation of agent-based financial market models by providing a systematic numerical study of the performance and properties of a simulated joint moment estimator as guidance for future research. Given that numerical studies are extremely expensive in terms of computation time, the literature is still missing an in-depth study of the joint moment estimator. In fact, the SJM has been used mostly for calibration exercises so far. We will hence test the estimator’s ability to recover model parameters consistently and efficiently in an estimation setting. To explore the behavior of the objective function’s response surface, we will run a large number of Monte Carlo simulations. Being aware of the computational burden of such a Monte Carlo study, we will use a machine-learning meta-model as proposed in Lamperti et al. (2018).⁴ We will train a surrogate meta-model with a limited number of true (cost-intensive) model evaluations. By doing so, we have been able to efficiently explore the parameter space at low cost. Since Lamperti et al. (2018) only consider the p-value of a Kolmogorov-Smirnov test as a validation criterion, we will, thus, take their approach one step further by coupling the use of machine-learning surrogates with a simulated moment estimator. This allows us to test the potential of AI-enhanced surrogates in a broader scaled applied estimation routine.

To undergo such a study we need an agent-based financial market model that is rather simple to keep the computational burden manageable and yet sophisticated enough to produce realistic dynamics. We found the model of Schmitt et al.

⁴In recent years, machine-learning has been a hot topic and it also made its way towards the field of economics. See e.g. Varian (2014) for an introductory discussion on big data, machine learning and potential in economics.

(2020) which is henceforth referred to as SSW quite appealing for this purpose. They propose a small-scale behavioral asset-pricing model (including nine relevant parameters) that is able to explain some important stylized facts of financial markets. Assuming that all speculators follow their own individual trading signals, it turns out that stock prices are excessively volatile and fluctuate erratically around their fundamental value. Schmitt et al. (2020) then show that during short periods of high uncertainty, speculators start to panic and coordinate their behavior. This panic-induced herding behavior leads to a sporadic loss of heterogeneity and gives rise to extreme market events. Periods of high volatility are long lasting since speculators persistently receive stronger trading signals out of past stock price movements. Despite the model's simple nature, simulations in Schmitt et al. (2020) indicate that the SSW model is indeed able to replicate a number of important stylized facts of financial markets. Yet, in their paper, the authors only provide a rough calibration of the model. Any true empirical validation is missing so far. Therefore, our study aims at closing this gap and attempts at a thorough empirical analysis of the SSW model. Hence, we make a second contribution by estimating the model of Schmitt et al. (2020) for a total of five empirical financial time series.

Our paper also relates to the work by Grazzini et al. (2012a,b) who study the performance of an SMM estimator for a simple innovation diffusion model and show that parameter estimates of small sample populations are biased due to the non-linear nature of moments. They report that biases vanish with increasing sample size. In a similar vein is the work by Jang (2015). He uses SMM estimation for the model of Alfarano and Lux (2007) and finds the objective function to be non-smooth, having multiple local optima and large flat areas. This hampers the consistent and efficient estimation of model parameters. Following this line of research, Chen and Lux (2018) study the efficiency of SMM estimation for the model by Alfarano et al. (2008). Their main findings include the arbitrariness of results due to different starting conditions. They, hence, propose the use of an extensive grid search. They further mention the limited general usefulness of SMM estimation for ABM using univariate financial time series because of the limited number of available moments and the very nature of stylized facts of financial markets. More macroeconomic-related SMM studies can be found in Grazzini and Richiardi (2015), Kukacka et al. (2019) and Franke (2018). Since the research on empirical validation and estimation of economic ABM is blooming, other approaches have emerged as well. They include information-theoretic similarity-based measures as given by Lamperti (2015) and Barde (2016). Guerini and Moneta (2017) use

causal inference through SVAR regressions. Most recently, a number of Bayesian alternatives have emerged, too (see Grazzini et al. 2017, Platt 2019, Bertschinger and Mozzhorin 2020, Delli Gatti and Grazzini 2020 and Lux 2020a). Yet, these sophisticated techniques often suffer from immense computational burden. In contrast, this work presents an estimation methodology which can be realized with the computing resources of a state-of-the-art personal computer.⁵ So far, there is no real consensus on the empirical validation of ABM. The status quo is discussed in Fagiolo et al. (2019), Delli Gatti et al. (2018), Lux and Zwinkels (2018).⁶

The remainder of this paper is organized as follows. Section 2 briefly describes a simple agent-based financial market model of Schmitt et al. (2020) which will be used to explain the econometric challenges of estimating such a model. Section 3 proposes a simulated joint moment estimator. In Section 4, we present results of several Monte Carlo simulation runs to study the performance of the joint moment estimator. Section 5 includes an empirical application of the previously proposed estimation approach on a large set of financial data. We conclude this paper in Section 6.

2 A simple agent-based financial market model

A study of estimation techniques is not feasible without any actual (economic) model under consideration. Hence, in this section, we will introduce a simple stochastic agent-based financial market model of Schmitt et al. (2020) that aims at replicating the stylized facts of financial markets. This model considers a single speculative stock market in which a market maker adjusts stock prices with respect to the speculators' order flow. Speculators have different beliefs about future prices and base their orders on their own trading signals, either derived from private market research or from applying complex trading systems. In calm periods, i.e. when volatility is low, speculators act rather independently. Their trading signals, which they derive out of past price movements are also rather weak. What happens in more turbulent times? If volatility is high, speculators receive strong trading signals which results in long-lasting periods of high volatility, i.e. volatility clustering. In these times of great uncertainty, speculators also tend to herd more strongly. Heterogeneity breaks down and the herding-induced coordination of

⁵All experiments in this paper have been performed using Julia 1.1.0 on a standard desktop computer with an eight core 3.00 Ghz AMD Ryzen 7 1700 CPU, 32 gigabytes of memory and an ASUS GeForce GTX Phoenix OC 6 GB GPU.

⁶For a good overview on empirically estimated ABM, see Chen et al. (2012), which is expanded in Kukacka and Barunik (2017), and Ellen and Verschoor (2018).

speculators leads to extreme returns.

Let us now turn to a detailed description of the model. A market maker adjusts stock prices with respect to the sum of the orders of N heterogeneous speculators. Inspired by Day and Huang (1990), the market maker's behavior is thus given by

$$P_{t+1} = P_t + \alpha \sum_{i=1}^N D_{t,i}, \quad (1)$$

where P_t is the log price of the risky asset at time t , α is a positive price adjustment parameter reflecting the market's liquidity, and $\sum_{i=1}^N D_{t,i}$ is the aggregate excess demand resulting from the individual orders $D_{t,i}$ of speculators $i = 1, 2, \dots, N$.⁷ Hence, if the sum of total orders is positive (negative), the market maker increases (decreases) the price of the risky asset.

The order of speculator i depends on her or his own trading signal, which is a complex mixture of e.g. private market research, rumors, or the application of (algorithmic) trading systems. Thus, all agents have different expectations about future prices and their expectations are correlated to some extent. For simplicity, a single order $D_{t,i}$ of speculator i at time t is represented by

$$D_{t,i} = \delta_{t,i}, \quad (2)$$

where $\delta_t = \{\delta_{t,1}, \delta_{t,2}, \dots, \delta_{t,N}\}'$ is a multivariate normally distributed random variable vector with, i.e. $\delta_t \sim \mathcal{N}(\mu_t, \Sigma_t)$.

The mean vector is given as $\mu_t = \{\mu_{t,1}, \mu_{t,2}, \dots, \mu_{t,N}\}'$ with $\mu_t = \mu_{t,i}$ and incorporates empirical and laboratory insights according to which financial speculators rely on fundamental and technical trading rules (Menkhoff and Taylor 2007, Hommes 2011). It is thus assumed that

$$\mu_t = \beta^c(P_t - P_{t-1}) + \beta^f(F - P_t), \quad (3)$$

where μ_t is a linear function to capture the core principles of technical and fundamental analysis. The first component of (3) implies that speculators rely on the most recent price trend. If prices go up (down), they receive a positive (negative) trading signal. The second component of (3) suggests that speculators sell (buy) overvalued (undervalued) assets with F being the stock's fundamental value. The

⁷The use of a market maker allows us to avoid further sources of uncertainty when analyzing the root causes of biases in the parameter estimates. For example, Platt and Gebbie (2018) found parameter degeneracies for agent-based order book models suggesting that a stylized fact-centric validation approach is insufficient for intraday models.

strength of trading signals is determined by the reaction parameters β^c and β^f .

Next, let us now turn to the symmetric variance-covariance matrix of the multivariate normally distributed random variable δ_t which is given by

$$\Sigma_t = \begin{bmatrix} \sigma_{t,1}^2 & \sigma_{t,1}\sigma_{t,2}\rho_{t,1,2} & \cdots & \sigma_{t,1}\sigma_{t,N}\rho_{t,1,N} \\ & \sigma_{t,2}^2 & \cdots & \vdots \\ & & \ddots & \sigma_{t,N-1}\sigma_{t,N}\rho_{t,N-1,N} \\ & & & \sigma_{t,N}^2 \end{bmatrix}, \quad (4)$$

with $\sigma_t^2 = \sigma_{t,i}^2$ and the time-varying correlation coefficients $\rho_t = \rho_{t,i,j}$ for $i, j = 1, 2, \dots, N$ and $i \neq j$. It is assumed that the intensity of trading signals increases with the market's volatility. This reflects two behavioral aspects. First, speculators are known to make up their beliefs about future prices (and hence their demand) (at least partially) dependent upon past and current price movements. If there is greater variability in the market, then their trading signals will grow correspondingly. Second, speculators are known to be biased in times of uncertainty. Shiller (2015) argues that they tend to become prone to hysteria. Hence, speculators overreact to their signals in more volatile periods and regard them as highly relevant. Based on these insights, the market's volatility is defined as a smoothed measure of past and current volatility, i.e.

$$V_t = \lambda V_{t-1} + (1 - \lambda)(P_t - P_{t-1})^2, \quad (5)$$

where $0 \leq \lambda < 1$ is a memory parameter governing the relative strength of past and current volatility. If λ approximates 0, speculators are becoming more and more myopic, focusing only on the most recent price movements. For $\lambda = 1$, speculators have a behavioral bias towards past volatility solely. The intensity of speculators' trading behavior is then modeled as

$$\sigma_t^2 = \gamma^l + \frac{\gamma^h - \gamma^l}{1 + \exp[-\gamma^s(V_t - \gamma^m)]}. \quad (6)$$

Note that (6) represents a logistic function with $0 < \gamma^l < \gamma^h$. Considering $\gamma^m > 0$ as a reference value for the stock market's volatility, speculators' trading intensity is equal to the midpoint of (6) if $V_t = \gamma^m$. Slope parameter $\gamma^s > 0$ determines sensitivity of reactions to changes.

As supported by empirics (Welch 2000) and experiments (Hommes et al. 2005), financial speculators are not isolated in their decision-making. Instead, they are subject to herding behavior and observe the decisions of their peers quite carefully

in times of heightened uncertainty as already argued by Keynes (1936). For simplicity, herding behavior is dependent upon current regime shifts in the market, given by

$$C_t = ((P_t - P_{t-1}) - (P_{t-1} - P_{t-2}))^2. \quad (7)$$

According to (7), the functional value of C_t increases when price reversals occur. The correlation between speculators' trading behavior is then formalized as

$$\rho_t = \phi^l + \frac{\phi^h - \phi^l}{1 + \exp[-\phi^s (C_t - \phi^m)]}. \quad (8)$$

In analogy to (6), ρ_t represents again a logistic function with $0 \leq \phi^l < \phi^h \leq 1$, where ϕ^l and ϕ^h determine the lower and upper boundary of ρ_t . The slope of ρ_t is defined by parameter $\phi^s > 0$ and $\phi^m > 0$ marks its midpoint. The greater the current price movements as captured by (7), the stronger is the speculators' tendency to coordinate their behavior. Thus, the strength of herding, i.e. the correlation coefficients, increases with (7) due to the sigmoid shape of (8). Note that if ρ_t approaches 1, trading signals become fully correlated and, hence, are identical among all market participants. If ρ_t approaches 0, speculators' beliefs are uncorrelated.

Given our assumptions, the evolution of the log price is governed by the following law of motion

$$P_{t+1} = P_t + \alpha \left\{ N (\beta^c (P_t - P_{t-1}) + \beta^f (F - P_t)) + \sigma_t \sqrt{N + N(N-1)\rho_t} \cdot \epsilon_t \right\}, \quad (9)$$

with $\epsilon_t \sim \mathcal{N}(0, 1)$. For the simulation of (9), we face a total of 14 parameters. We set $N = 100$, $\alpha = 1$ and $F = 0$. Parameters a and N are scale parameters while parameter F determines the stock's (constant) fundamental value around which prices will fluctuate. Further, we assume that $\gamma^s = \phi^s = \infty$ implying that the trading signals' intensity and correlation jump discontinuously between lower and upper boundaries of (6) and (8), respectively. For the remaining parameters, we define $\boldsymbol{\theta}$ as a 9×1 parameter vector from the nine-dimensional parametric space Θ with $\boldsymbol{\theta} := (\lambda, \gamma^h, \gamma^l, \gamma^m, \phi^l, \phi^h, \phi^m, \beta^c, \beta^f)$. The main goal of this estimation study will be to find empirically validated estimates for the parameter vector $\boldsymbol{\theta}$. Regarding its dimensionality, it may be deemed to be representative for most small scale ABMs.⁸

⁸Related studies mostly use prototype-like models with only 3 to 5 parameters (e.g. Franke 2008, Jang 2015 or Chen and Lux 2018). In a recent study by Delli Gatti and Grazzini (2020), they consider a medium-sized macro ABM with 11 parameters.

3 Simulated Method of Moments

In this section, we will define a simulated joint moment estimator as first proposed by Franke and Westerhoff (2012). In general, the goal of any moment-based estimation approach is to find parameter values that minimize the distance between a set of empirical and simulated moments. Those moments can be any summary statistic which puts the choice of moments at the very core of every SMM estimation. Since the moment conditions are not predetermined by default, we face a broad range of possibilities. That means, the simulated method of moments requires the modeler to be clear about her goals. This arbitrariness is often stressed as the main disadvantage of SMM. In fact, it is less a question of correct choice than of transparency, reasonableness and thoughtfulness. Given the model that is studied in this paper, we try to explain the emergence of some of the well-documented stylized facts of financial markets.⁹ These include bubbles and crashes, excess volatility, fat-tailed return distributions, absence of autocorrelations in the raw returns, volatility clustering and long memory. See left panels of Figure 9 to 13 in Appendix A which illustrate these universal properties showing the dynamics of the S&P500 stock market index, the German DAX, the Japanese Nikkei, the price of gold and the USD-EUR exchange rate. A set of reasonable summary statistics, i.e. moments, will incorporate quantitative measurements of these features as summarized in Table 1. Hence, our first moment is the volatility of returns defined as the mean of absolute returns. The second and the third moment are the Hill tail estimates at 2.5 and 5.0 %, respectively. Next, we take the autocorrelation of raw returns at lag 1 as the fourth moment to show the absence of serial correlation in returns. The final five moments are the autocorrelations of absolute returns at lags 3, 6, 12, 25, 50 and 100 and take the long memory effect into account. Section 4.2 includes a sensitivity study of our set of moments showing that it is robust and reasonable.

Based on the previously defined moments, we will now define an objective function to be optimized. Basically, the objective function for a standard SMM estimation would be formulated to minimize the distance between the empirical and the simulated moments. This is typically done by a quadratic loss function including a positive definite weighting matrix to account for the differences in the sampling variability of the moments. Franke and Westerhoff (2012) have shown that minimizing such an objective function does not necessarily lead to the best

⁹Surveys about the statistical properties of financial markets can be found in Mantegna and Stanley (2000), Cont (2001), Lux and Ausloos (2002) and Lux (2009).

Moment	Statistic
m_1	$E(r_t)$
m_2	$\alpha_{2.5}$
m_3	$\alpha_{5.0}$
m_4	$E(r_t r_{t-1})$
m_5	$E(r_t r_{t-3})$
m_6	$E(r_t r_{t-6})$
m_7	$E(r_t r_{t-12})$
m_8	$E(r_t r_{t-25})$
m_9	$E(r_t r_{t-50})$
m_{10}	$E(r_t r_{t-100})$

Table 1: Moments of interest. We define a set of ten moments including the returns' volatility, the Hill tail indices $\alpha_{2.5}$ and $\alpha_{5.0}$, the autocorrelation coefficient of raw returns at lag 1 and the autocorrelation coefficients of absolute returns for lags 3, 6, 12, 25, 50 and 100.

model in terms of joint moment matching. They introduce the joint moment coverage ratio which aims at maximizing the fraction of simulation runs for which all simulated moments jointly fall into the 95 percent confidence intervals of their empirical counterparts. This estimator offers an intuitive interpretation of the parameter estimates' validity, i.e. how often can the model and the real-world not be told apart, which makes the SJM an appealing choice. While research on the properties of the standard SMM estimator is abundant, the SJM has been used mostly for calibration exercises so far. We will try to close this gap by providing a coherent study of the performance and properties of a simulated joint moment estimator as guidance for future research.

Based on the work by Franke and Westerhoff (2012), we will, therefore, now define a simulated joint moment estimator. We consider a large number of simulation runs whereof each single realization of our economic model is given as

$$Y = \{y_t(\boldsymbol{\theta}, s)\}_{t=1}^T, \quad (10)$$

where Y represents the simulated return series over T time steps. Initial conditions for y_t are given by the parameter vector $\boldsymbol{\theta}$ and random seed s . Given the artificial data from our economic model, we can compute the set of moments m_k as

$$m_k(y(\boldsymbol{\theta}, T; s)). \quad (11)$$

Next, we conduct the test statistic

$$z_k(y; m_k) = \left(\frac{\hat{m}_k - m_{k,0}}{\sigma_{k,0}} \right), \quad (12)$$

with $m_{k,0}$ and $\sigma_{k,0}$ being the (pseudo-) true k th moment value and its corresponding standard error and $\hat{\mathbf{m}}_k$ being the simulated moment. We now define an indicator function

$$f(z; \dots) := \begin{cases} 1 & \text{if } z_k \leq z_c, \forall k \\ 0 & \text{otherwise} \end{cases}, \quad (13)$$

which is only true if all (ten) moments are (jointly) smaller than a critical value z_c . Under asymptotic theory, a 95% confidence interval amounts to a value of 1.96 for z_c . Since we are interested in the number of positive realizations, i.e. joint moment matches, the total loss function maximizes the sum of the joint moment events. Our objective function is, hence, given by

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta} \in \Theta} \sum_{MC} f(z; \dots). \quad (14)$$

To reduce sample variability, the SJM estimator $\hat{\boldsymbol{\theta}}$ is run over a large number of simulations MC . The strong penalty in (13) can become problematic when searching on a broader parameter space. One might end up getting zero information. Franke and Westerhoff (2012) introduced for that reason a piecewise linear function. Experimenting with similar adjustments, however, gave us no better results than the proposed loss function here.

3.1 Response Surface

To get a first idea of the objective function's behavior, we do some preliminary sensitivity analyses. This allows us to get a better understanding of potential hurdles during estimation and to define an appropriate parameter space for our upcoming simulation study.

For the remaining nine parameters, we set the pseudo-true parameter vector $\boldsymbol{\theta}_0 = \{0.81, 0.35, 0.058, 14.5, 0.001, 0.065, 0.0032, -0.1, 10\}$. The parameter set is chosen in proximity to actual empirical estimates to satisfy that the generated data resembles the statistical properties of financial returns. Note that parameters $\gamma^h, \gamma^l, \gamma^m, \beta^c$ and β^f are always multiplied by 10^5 for better readability.

In order to illustrate the behavior of the response surface of our objective function as given in (14), we generate a two-dimensional grid of varying parameter pairs of $\lambda \in [0.795, 0.825]$ and $\gamma^h \in [0.32, 0.38]$. The grid size contains 961 equidistant parameter combinations of the two parameters. We run simulations over combinations of $T = 20\,000$ and $T = 80\,000$ for 400 and 5 000 Monte Carlo repetitions. The results of the response surfaces of the SJM estimator are shown in




Figure 1: Objective function's response surface. The panels reveal the asymptotic behavior of the objective function for different time length T and different number of repeated simulation runs MC over varying values for parameters λ and γ^h . All other parameters are fixed. Black lines indicate the (pseudo-) true parameter values for λ and γ^h and red dots visualize the highest peak of the response surface.

Figure 1. Black lines indicate the (pseudo-) true parameter values for λ and γ^h and red dots visualize the highest peak of the response surface. Apparently, the objective's function response surface is rugged, yet globally concave and differs among the four different settings. Increasing both time length T and the number of Monte Carlo simulations MC makes the surface less rugged and more smooth. We clearly see that the SJM estimator struggles with identifying parameter λ , while parameter γ^h is almost always correctly estimated. Only in a somewhat ideal world where we could run 5 000 simulations over time length $T = 80\,000$ the red dot coincides with the intersection of both black lines.¹⁰ However, this

¹⁰Running a complete estimation, i.e. over all nine parameters with 5 000 Monte Carlo simula-

comes at the price of less information available, i.e. the large dark red-shaded area. Let us explain this in more detail. The SJM estimator appears to be a much stronger criterion compared to other traditional SMM estimators which focus on simply minimizing the distance between the simulated and empirical moments. That often leads to parametrizations for which some moments are fit very well while others are far away from their empirical counterparts. A joint moment estimator aims at avoiding this drawback by considering all moments simultaneously. This way, the informational value of the SJM estimator is limited, though, which hampers estimation. Let us consider, for example, any parameter combination θ_1 that is not even close to the true parameter vector θ_0 . The corresponding SJM would be 0 giving us no information at all. A small change in the parameter values may probably still give us a SJM of 0. Hence, the response surface in this area of the parameter space is completely flat leaving us unable to detect any direction for optimization. This is especially true for large values of T . The reason is that for generating pseudo-empirical data, the length of T is directly linked to the moments' variance. The longer T is, the smaller tends to be the moments' variance (given a high enough number of Monte Carlo repetitions to ensure convergence). Consequently, time length T restricts the size of the areas of the parameter space that carry valuable information as shown in the panels of Figure 1.

This is very important to keep in mind for the estimation routine and so far the biggest drawback of the here presented SJM estimator. What does this mean for the design of our estimation study? We have to find an optimal trade-off between empirical and simulated time length T and the number of Monte Carlo runs MC .




Figure 2: Box Plots for $T_{emp} = 20.000$ (left) and $T_{emp} = 80.000$ (right). The Figures show Monte Carlo studies with 100 repetitions for parameter $\lambda = 0.81$ and different settings of $R = \frac{T_{sim}}{T_{emp}}$.

Given that the available length of real financial data naturally limits the choice for T , we run a second sensitivity analysis. In fact, we tested if we can reduce some

tions and over $T = 80\,000$ observations, using the same computational resources and setup as in this paper would take us approximately 250 days.

of the estimator’s variability by increasing the simulated time length only (and keeping the empirical time length constant). We define R as the ratio of simulated time length per empirical time length. In related studies, e.g. Chen and Lux (2018), R is often chosen to be larger than 1. Therefore, we run a small simulation study with 100 independent repetitions for varying values of parameter λ under ceteris paribus. MC is set to 400. Figure 2 depicts the relationship between simulated and (pseudo-) empirical time lengths. We see that the variability is not decreasing but increasing for higher values of R . Also, estimates are biased and not centered around the (pseudo-) true value. The reason that higher values of R have no positive effect here, might be related to findings by Secchi and Seri (2017). They discuss the potential risks of under- and overpower of ABM in terms of statistical power analysis. When overpowered, differences in effects may become so small that the model might disregard them. This is exactly what we see here. Given our test statistic as proposed earlier in this section, one way to solve this issue is to set the critical value z_c significantly lower. This means reducing the significance level. Figure 2 shows that this would reduce the estimator’s variance. However, this comes again at the cost of less available information which is the reason why we did not include lower values of z_c in our numerical study.

4 Monte Carlo Study

In this paper, we aim at developing a robust estimation routine, giving a better understanding of possible obstacles, e.g. sources of biases, and providing a sound base for an empirical application. Before we will apply the proposed SJM estimator on real-world data, we test its performance in a controlled laboratory environment in order to investigate the consistency and efficiency properties over small samples from simulated data of our model. Section 4.1 introduces a general setup of the estimation routine. Section 4.2 conducts an extensive Monte Carlo study and reports the corresponding results.

4.1 Estimation Routine

Conducting a massive estimation study is a computationally challenging attempt. It is well known that ABM are computationally demanding in terms of simulation time per run, their validation and estimation. More details on the computational costs of running and estimating ABM is given by Fagiolo et al. (2019). Recall that our objective function as given in (14) possesses no analytical expression. This

limits the number of possible optimization techniques at hand by large. Further, as we have shown previously, the response surface of the loss function tends to be non-smooth. This makes it even harder to find the right tools to investigate the emergent behavior of an ABM. In such cases, optimization heuristics are often used (see e.g. Gilli and Winker 2009) and the Nelder-Mead algorithm has been an especially popular choice. However, it is known for its strong variability of results depending on its starting values. To mitigate any bias, one would have to run the algorithm several times with different initial values leading to increased computational costs.

If one is lucky enough to have access to powerful cluster infrastructures, or willing enough to pay for cloud computing services (e.g. Amazon Web Services or Google Computing Cloud), one might be less concerned with computational limitations. Yet, we strongly believe in academic equity and advocate for inclusive and fully accessible research. Hence, we here provide an estimation routine that is easily doable on any up-to-date home computing device. To make that happen, we reduce the computational burden with the help of a meta-model.

Meta-models (also known as surrogates) can be used as a computationally cheap proxy since they approximate the mostly non-linear relationship between input parameters and model output. They, thus, reduce the computational burden. This meta-modeling approach has first been introduced into the field of economic ABM by Salle and Yıldızoğlu (2014), followed by Dosi et al. (2018a, b, c).

The statistical tool, that these studies use, is called Kriging and it facilitates sensitivity analyses over a large parameter space. However, Kriging struggles to correctly approximate the non-smooth response surface of most ABM. Therefore, Lamperti et al. (2018) propose the use of machine-learning surrogates as a more efficient alternative. Machine-learning, as an applied form of artificial intelligence, learns autonomously from given data and is able to make predictions based on that data.¹¹ In their study, they introduce a (semi-) supervised regression/classification approach. The basic idea is to map a user-defined fitness measure and a specific parameter set of the underlying ABM. Using a budgeted online approach, the surrogate’s performance improves successively over multiple rounds while correctly identifying parametrizations of interest. We will adopt this approach which is described in the following. A corresponding pseudo code is given in Figure 3. The source code to run and replicate the numerical experiments in this section can be found online: <https://github.com/IvonneSchwartz/SurrogateEnhanced-Mo>

¹¹See Van der Hoog (2019) for a broad discussion of potential uses of machine learning surrogates to agent-based modeling in the field of economics.

```

for i=1 to 100, do
  RandomSeed 1:  $seed_1(i)$ 
  Pseudo empirical data  $T_{emp}$ 
  RandomSeed 2:  $seed_2(i)$ 
  Simulated data  $T_{sim}$  (Grid 2500)
  for j=1 to 4, do
    RandomSeed 3:  $seed_3(i)$ 
    Train the machine-learning surrogate (CatBoost)
    Make predictions over a large pool of parameter combinations
    400 best predictions: Simulate data  $T_{sim}^n$  with  $seed_2(i)$ 
  end for
end for

```

Figure 3: Pseudo code of our estimation routine.

mentEstimation. The specific design of our estimation algorithm is mainly due to the computational limitations we faced. Our goal was to find a feasible balance between the computational burden and the significance of our results. To get any insightful results at all, we set the number of repetitions to 100 per setting as the bare minimum. A first random seed is used to generate the pseudo empirical data. We choose two different time length; $T = 20\,000$ and $T = 80\,000$. Using a second random seed, we run and evaluate an initial grid over 2 500 parameter combinations. This is done for two reasons.

First, given the behaviour of the objective function’s response surface as discussed in Section 3.1 with its flat valleys and multiple local optima, any attempt to converge to a potential global optimum will critically depend on initial conditions. Different initial conditions may lead to drastic differences in parameter estimates. This has also been reported by Grammig and Schaub (2015) and Chen and Lux (2018). The usual work-around is to start with a rough grid search to provide a first mapping of the objective function’s behavior. Such an initial guidance prior to the optimization is crucial here for another reason. To efficiently explore the parameter space with the help of machine-learning surrogates, we first need to train the machine-learning model. In order to get an acceptable prediction performance, the training sample size should not be too low. We stick to the results by Lamperti et al. (2018) who report a satisfactory explanatory performance by training the surrogate on 2 500 points (i.e. parameter combinations). This works great here too, but it is important to mention that the initial sample size most probably depends on the complexity of the economic model under consideration. Given the dimensionality of our model, equi-distant grid points, as

in Chen and Lux (2018), are computationally not feasible here. Instead, we used Quasi Monte Carlo sampling based on Sobol sequences to get a set of 2 500 parameter combinations that are evenly distributed over the nine-dimensional parameter space of interest. Sobol sequences generate a uniform distribution in probability space which are more evenly distributed compared to random sequences. Alternative sampling techniques include Monte Carlo sampling based on pseudo-random numbers or Latin Hypercube sampling as used in Salle and Yıldızoğlu (2014). Yet, Kucherenko et al. (2015) have shown that Sobol sequences possess space-filling properties which are superior to the other aforementioned sampling techniques.

We, thus, start with a preliminary manual search over a restricted subset of the parameter space as given in Table 2. Parameter boundaries are chosen such that (i) (pseudo-) true values lay well within the parameter space and (ii) pre-grid search finds at least one positive output, i.e. at least one joint moment event. Otherwise, we would search in the dark.¹² However, this is not an easy task given the high dimensionality of the ABM and computational costs will probably outrun the performance benefits. Therefore we decided to manually calibrate the restricted subset.

The training of the surrogate model is now done repeatedly over multiple rounds. This way it is successively improved to make more accurate predictions after each training round.¹³ With another set of random seeds, we train the machine-learning surrogate using our previously simulated grid data. For the meta-model learning, we resort to CatBoost (Prokhorenkova et al. 2018). CatBoost is a gradient boosting toolkit that is freely available and open-source.¹⁴ One of CatBoost’s advantages is its little required training time and little required expertise for parameter tuning. After experimenting with other algorithms, e.g. the XGBoost library as proposed in Lamperti et al. (2018), we found that the prediction time is much faster for CatBoost. Although CatBoost’s default performance is already high, we initially optimized the surrogate’s hyperparameters to foster predictive performance further. In contrast to Lamperti et al. (2018), this is not done automated. Since the nature of our economic model and the estimator don’t change throughout our study, we use the same set of optimized hyperparameters for all our experiments. To further improve CatBoost’s predictive performance, we made

¹²We thank an anonymous referee who remarked that, ideally, this should be done by a search algorithm. We takes this as an open issue for future research.

¹³We have experimented with different options of splitting up our total evaluation budget and found the presented one to result in the most precise parameter estimates given our limited resources.

¹⁴The Github repository can be found here: <https://github.com/catboost>. For a thorough introduction into boosting, see Schapire and Freund (2012).

use of feature engineering and included additional input features. This is done by creating new features out of already existing ones in order to improve model performance. Feature engineering aims at highlighting import patterns within the data, i.e. interactions between features. What we did is try to highlight the relationship between the variables of the (economic) model’s variance function and correlation function, respectively. We hence included features considering the ratio/quotient of the upper and lower bound parameters of those functions. Plus the quotient of the upper and lower bound and the threshold parameter as well. We also experimented with different type of normalizing our input data, but we could not detect any improvements from that.¹⁵ After successfully training the




Figure 4: Surrogate behavior and performance. The figure reveals the performance of the surrogate for a simulation exercise run over 10 000 parameter combinations. The black line (belonging to the right y-axis) shows the median value of the mean-squared error (MSE) together with its confidence band for increasing training sample size. The solid and dashed red lines (belonging to the left y-axis) indicate the computation time in seconds for the surrogate and the true ABM, respectively.

surrogate, we can finally make predictions over a very large number of parameter combinations. In fact, we set up a pool containing 40 million parameter combinations. Given the predictions, we take the best 400 (in terms of the predicted SJM estimator) and simulate those via the true ABM. Adding the newly simulated data to the initial data set drives the estimation towards parameterizations that match our desired properties since the machine-learning surrogate becomes successively smarter. In total, we repeat this procedure four times. Obviously, our algorithm can be easily adjusted at several points to scale it up (or down) given the available

¹⁵See Zheng and Casari (2018) for a useful guide on applied feature engineering.

computational resources.

Before running a batch of experiments, let us first consider the behavior of the machine-learning surrogate to check if it is truly able to approximate our ABM. To do so, we test the predictive performance of the surrogate for increasing size of the training sample. The experiment is based on a set of 10 000 parameter combinations run over 100 independent experiments with different random seeds for the surrogate. Predictive performance is measured in terms of the mean-squared error (MSE) which is the squared deviations of the predicted estimator's value and the true value (that comes from running/simulating and evaluating the true ABM). Figure 4 shows the results.

As expected, the MSE is decreasing for an increasing training sample size. Given that we are using a real-valued setting (regression problem instead of classification), the MSE is already satisfyingly low for small sizes of the training sample. The red dashed line shows the increase in computation time for running the true ABM. The timing results for the surrogate include average training and prediction time and remain steady around two seconds for all considered training sample sizes. In contrast to the growing computational demand of the ABM, we see no increase for the surrogate. Depending on the available computational resources, we need to decide for a certain trade-off between computational cost and performance. As described earlier, we decided for a training size of 2 500 parameter combinations.

4.2 Simulation study

Let us now turn to our simulation study where we investigate the performance of our proposed estimation approach. Therefore, we start this subsection with a detailed description of our setup for the numerical experiments. Afterwards, we will report and discuss the results of our Monte Carlo simulations.

We start our estimation study with first considering a reduced subversion of the economic model including only four parameters. This first experiment lets us investigate if our approach and set of moments is basically suitable to estimate the chosen economic model. Therefore, we switch off the coordination mechanism of the model, as given in (8), by setting $\rho_t = \phi^l = 0.0006$. For the remaining pseudo-true parameter set $\boldsymbol{\theta} = \{\lambda, \gamma^h, \gamma^l, \gamma^m, \phi^l, \beta^c, \beta^f\}$, we choose the (pseudo-) true parameter vector $\boldsymbol{\theta}_0 = \{0.81, 0.35, 0.058, 14.5, 0.0006, -0.1, 10\}$, whereof we include only the first four parameters for the estimation exercise. Results are shown in the first part of Table 2 and look very promising. All four parameters are

precisely estimated with their sample means being very close to their (pseudo-) true values. Standard errors of the estimates are fairly low, too. We can conclude that our approach is, indeed, able to correctly identify model parameters.


	λ	γ^h	γ^l	γ^m	ϕ^h	ϕ^m	ϕ^l	β^c	β^f	tsec
θ_0	0.81 [0.795,0.825]	0.35 [0.22,0.48]	0.058 [0.045,0.071]	14.5 [13,16]	-	-	0.0006	-0.1	10	
<hr/>										
$T = 20\,000$										
	0.808	0.346	0.057	14.324						
	0.002	0.006	0.001	0.242						
	0.003	0.007	0.001	0.298						
<hr/>										
θ_0	0.81 [0.795,0.825]	0.35 [0.22,0.48]	0.058 [0.045,0.071]	14.5 [13,16]	0.065 [0.05,0.08]	0.0032 [0.0022,0.0042]	0.001 [0.0001,0.01]	-0.1 [-1,-0.01]	10 [1,100]	
<hr/>										
$T = 20\,000$										
MEAN	0.8125	0.2661	0.0468	14.4186	0.0689	0.0026	0.0038	-0.405	12.8029	855.87
FSSE	0.0044	0.0115	0.0018	0.5125	0.0053	0.0002	0.0007	0.2382	4.541	9.56
RMSE	0.0050	0.0847	0.0114	0.5166	0.0065	0.0007	0.0029	0.3863	5.3187	
$T = 80\,000$										
MEAN	0.8125	0.2746	0.0477	13.8439	0.0654	0.0025	0.0030	-0.3746	12.3685	2177.75
FSSE	0.0039	0.0173	0.0034	0.4858	0.0063	0.0001	0.0007	0.2367	6.2806	13.48
RMSE	0.0046	0.0773	0.0109	0.8150	0.0062	0.0007	0.0021	0.3617	6.6834	

Table 2: The table shows the means, finite sample standard errors (FSSE) and root-mean squared errors (RMSE) of the parameter estimates of 100 Monte Carlo repetitions per scenario. MC is fixed at 400. The last column reports the average computation time in seconds needed for the execution of one single estimation.


Next, we go one step further and take the full version of the model with its nine parameters into consideration. For the pseudo-true parameter set $\theta = \{\lambda, \gamma^h, \gamma^l, \gamma^m, \phi^l, \phi^h, \phi^m, \beta^c, \beta^f\}$, we choose the (pseudo-) true parameter vector $\theta_0 = \{0.81, 0.35, 0.058, 14.5, 0.001, 0.065, 0.0032, -0.1, 10\}$. A summary of the included parameter ranges is found in Table 2. Parameter intervals are chosen based on preliminary sensitivity analyses over the global parameter space and accordingly restricted given the computational resources at hand. We have further chosen parameter intervals in a way that the pseudo-true values are perfectly centered within (except for β^c and β^f). Note that, for this test scenario, we restrict the reaction parameter of chartists β^c to only allow for contrarian behavior.

We run the Monte Carlo simulations over two different time length $T = \{20\,000, 80\,000\}$. Experimenting with smaller time lengths led to high sampling variations which is why we chose 20 000 observations as our baseline setting. By including an increased sample size of 80 000 observations, we aim at obtaining valuable insights about the asymptotic tendencies of the estimator. Table 2 reveals the corresponding results of our Monte Carlo study where we report the means, standard errors and root-mean squared errors of the sample estimates.

Apparently, some parameters are easier to identify than others. Results show that parameters λ , γ^m , ϕ^h and β^f are very precisely identified with their mean estimates being quite close to their (pseudo-) true values. The previously well



(a) $T = 20\,000$



(b) $T = 80\,000$

Figure 5: Densities of parameter estimates for experiment 1. The panels show smoothed kernel densities for parameter estimates with red lines indicating the (pseudo-) true values and dotted black lines visualizing the 95% confidence intervals of the sample estimates.

identified parameters γ^h and γ^l are now far off their true values. Parameters ϕ^m and ϕ^l are troublesome as well. There also seems to be no improvement in the precision for these parameters with increasing time length. It seems that the estimator converges to different optima here. This might be due to the fact that these parameters belong to one of the s-shaped functions. In fact, parameter γ^h limits the upper bound and parameter γ^l the lower bound of the function for the variance of the trading signals, respectively. So it seems reasonable that there exists a multitude of parameter combinations which are equally good. A similar pattern can be observed for the other s-shaped function, i.e. the trading signals' correlation function. Figure 5 reveals additional insights about our numerical experiment. The panels show the parameter estimates' smoothed kernel densities for $T = 20\,000$ and $T = 80\,000$. Red lines indicate the (pseudo-) true values while dotted black lines visualize the 95% confidence intervals of the sample estimates. Figure 5a confirms multitudes of maxima since it shows multimodal distributions for most of the parameters. Increasing the time length T from 20 000 to 80 000, the

multimodality and distortions becomes less pronounced. In total, for five (six) out of nine parameters, the (pseudo-) true values are well within in the 95% confidence intervals given $T = 20\,000$ ($T = 80\,000$). Yet, increasing T does not necessarily improve the FSSE and RMSE for all parameters. This suggests convergence slower than \sqrt{T} as expected by asymptotic theory. Possible explanations may be found in Table 3 where we report correlations of the sample estimates. Inspecting the sample correlation matrix, we see that biases in the estimates can most probably be explained by high correlations between model parameters. We find especially high values for the parameter pairs defining the intensity and correlation of speculators' trading behavior: (γ^l, γ^h) , (ϕ^h, ϕ^m) and (ϕ^l, ϕ^h) . They are also high for the cross-functional pairs (γ^h, ϕ^l) and (γ^m, ϕ^l) . This suggests that a possible lack of identification is mainly due to the specifics of our underlying economic model. More precisely, this means that the S-shaped functional form of the two core model mechanisms (see Equation (6) and (8)) are complicating our estimation endeavor.

In the last column of Table 2, we report average timing results which refer to the average computation time in seconds needed for the execution of one single estimation run. Note that the increase in computation time from $T = 20\,000$ to $80\,000$ is less than proportional given the fact that we efficiently run all available cores of our CPU in parallel. How can we quantify these numbers in terms of performance improvement. Running the proposed routine, we need a total of 4 100 cost-intensive ABM evaluations. Using the same budget to evaluate a simple Sobol sample (without using any sophisticated estimation technique) leads to a significantly lower SJM estimator. More precisely, our approach gives us a qualitative performance increase of at least 25%. To get the same qualitative results, the Sobol Sequence has to cover a minimum of 200 000 data points. A full estimation run, i.e. including 100 Monte Carlo repetitions, over such a sample size would take around 12 days for $T = 20\,000$. This equals a reduction in computation time of 90%.

With regard to the empirical application, we run a second set of experiments over a slightly different and less restricted parameter space (details can be found in Table 2). For this setting, we allow for both contrarian and trend-following behavior. Note also that the (pseudo-) true parameter values are not necessarily centered within in the used parameter space anymore. We were suspecting that results would get worse with a broader parameter space. But that did not happen. However, results look differently which we report in Table 4. Compared to our first setting, mean estimates have significantly improved for parameters γ^h and γ^l which are now fairly close to their true values. The same holds for

	λ	γ^h	γ^l	γ^m	ϕ^h	ϕ^m	ϕ^l	β^c	β^f
λ	1.0	-0.40	0.24	0.68	0.21	0.12	0.45	0.08	-0.06
γ^h		1.0	0.73	-0.36	-0.58	-0.10	-0.86	-0.29	-0.24
γ^l			1.0	0.17	-0.56	-0.16	-0.58	-0.22	-0.23
γ^m				1.0	0.55	0.54	0.70	0.20	-0.09
ϕ^h					1.0	0.84	0.82	0.41	0.01
ϕ^m						1.0	0.50	0.31	-0.14
ϕ^l							1.0	0.32	0.04
β^c								1.0	0.16
β^f									1.0

Table 3: Correlation matrix of parameter estimates.

parameter ϕ^m . However, this comes at the cost of higher distortions for other parameters, including λ and γ^m which have been estimated precisely before. Figure 6 reveals the densities of the parameter estimates for the second setting. It shows again skewed multimodal distributions for most of the parameter estimates. For $T = 20\,000$, they are far from being normally distributed. For $T = 80\,000$, this becomes a bit less pronounced which is a common pattern between the first and second experiment. Further increasing T might eventually give us any true convergence behavior. Note that any statement about the asymptotic properties of the estimator should be taken cautiously given that the considered sample size is not long enough. It is known that the choice of moments effects the efficiency and precision of parameter estimates (Chen and Lux 2018). The reason is that the informational value differs among moments. Given that the choice of moments is also often criticized as being arbitrary, we aim at offering a high degree of transparency here. Therefore, we run a small sensitivity analysis for $T = 20\,000$ where we tested the impact of different moment sets on the estimator’s performance.¹⁶ Results are reported in Table 4:

- M6 includes the following six moments: volatility, Hill tail index at 5%, autocorrelation of raw returns at lag 1 and the autocorrelations at lags 3, 12 and 50. M6 might be regarded as the minimum of set of useful moments.
- M7 uses all moments of M6 plus the Hill tail index at 2.5%.
- M9 includes all moments except for the Hill tail index at 2.5%.

At first glance, the overall quality of the results did not worsen. Yet, inspecting the RMSE shows that results for M6 and M9 fall behind the standard setting


¹⁶Alternatively, Creel and Kristensen (2016) propose a cross-validation method for the selection of appropriate statistics for Approximate Bayesian Computing. They also suggest their approach for SMM-related estimation.

including all ten moments. Interestingly, M7 performs pretty well if not slightly better than M10. What we can learn from this is that the chosen set of moments is quite robust. Cutting off single moments does not deteriorate the results at large. We have also experimented with an extended set of moments adding the autocorrelations of squared returns. Results are also reported in Table 4 under M16. We find minor improvements for some of the parameter estimates which come at the expense of increased RMSE for other model parameters plus increased computational overhead of roughly ten percent. Overall, our baseline set of ten moments seems to be robust since the extended set of 16 moments could not foster performance significantly.


	λ	γ^h	γ^l	γ^m	ϕ^h	ϕ^m	ϕ^l	β^c	β^f	tsec
θ_0	0.81 [0.78,0.92]	0.35 [0.25,0.46]	0.058 [0.046,0.076]	14.5 [13.5,18]	0.065 [0.003,0.1]	0.0032 [0.0023,0.0040]	0.001 [0,0.003]	-0.1 [-10,10]	10 [1,50]	
$T = 20\,000$										
M6										
MEAN	0.817	0.355	0.066	16.699	0.054	0.0029	0.0014	3.047	19.51	
FSSE	0.007	0.025	0.003	0.391	0.006	0.0004	0.0005	3.295	9.595	
RMSE	0.01	0.025	0.009	2.234	0.012	0.0004	0.0006	4.546	13.479	
M7										
MEAN	0.817	0.344	0.063	16.622	0.063	0.0032	0.0018	3.323	19.479	
FSSE	0.005	0.012	0.002	0.325	0.007	0.0004	0.0001	3.491	8.096	
RMSE	0.008	0.014	0.005	2.146	0.007	0.0004	0.0008	4.877	12.441	
M9										
MEAN	0.821	0.34	0.065	16.637	0.054	0.0028	0.0017	3.904	21.764	
FSSE	0.005	0.016	0.002	0.323	0.007	35.433	0.0003	3.219	9.608	
RMSE	0.012	0.019	0.007	2.161	0.013	53.739	0.0007	5.129	15.163	
M10										
MEAN	0.818	0.341	0.063	16.553	0.063	0.0031	0.0019	4.000	21.339	857.80
FSSE	0.005	0.014	0.002	0.36	0.008	0.0004	0.0002	4.208	9.723	6.94
RMSE	0.01	0.016	0.005	2.084	0.008	0.0004	0.0009	5.861	14.907	
M16										
MEAN	0.811	0.343	0.058	15.904	0.074	0.0040	0.0020	1.240	15.409	948.73
FSSE	0.009	0.014	0.004	0.771	0.006	0.0003	0.0003	5.012	10.862	7.62
RMSE	0.009	0.015	0.004	1.600	0.011	0.0004	0.0011	5.165	12.089	
$T = 80\,000$										
M10										
MEAN	0.820	0.34	0.061	16.569	0.075	0.0034	0.0021	5.586	21.235	2179.10
FSSE	0.005	0.009	0.002	0.559	0.003	0.0001	0.0002	3.34	7.540	14.73
RMSE	0.011	0.016	0.004	2.143	0.01	0.0003	0.0012	6.589	13.510	

Table 4: The table shows the means, finite sample standard errors (FSSE) and root-mean squared errors (RMSE) of the parameter estimates of 100 Monte Carlo repetitions per scenario. MC is fixed at 400. The last column reports the average computation time in seconds needed for the execution of one single estimation.

Summing it up here, we find five out of nine parameters fairly well matched with the true values lying within the 95% confidence intervals of the sample estimates. For all other parameters, we observe small to strong biases. Our results suggest that estimates tend to depend quite strongly on initial conditions. We have shown that varying the parameter space gives us quite different results. This is a well-



(a) $T = 20\,000$



(b) $T = 80\,000$

Figure 6: Densities of parameter estimates for experiment 2. The panels show smoothed kernel densities for parameter estimates with red lines indicating the (pseudo-) true values and dotted black lines visualizing the 95% confidence intervals of the sample estimates.

known and common problem in the estimation of ABM and is mostly rooted in the struggle that search algorithms face when dealing with non-globally convex surfaces. Gilli and Winker (2003) suggest to include a stochastic acceptance rule in order to avoid being trapped in local optima. However, Platt and Gebbie (2018) find that the Nelder-Mead paired with a threshold accepting heuristic can still converge towards local minima. Here, it seems that the machine-learning surrogate is prone to this as well. Our workaround is to combine results from various estimation runs based on different settings. Figure 7 shows smoothed kernel densities of combined parameter estimates coming from both experiments for $T = 20\,000$. This seems a natural and effortless task to gain a better understanding of the estimator's performance for the considered economic model. In fact, we can boost the estimator's efficiency given that now eight out of nine (pseudo-) true values fall well within the 95% confidence intervals of their corresponding sample estimates. We, further, see that bimodality is even more apparent here. This again supports our hypothesis that it may most probably be explained by the specifics of the

economic model. Recall the S-shaped functions of Equations (6) and (8) for the intensity and correlation of speculators' trading behavior, respectively. Given their functional form, we have to deal with possible multiple local optima. Convergence to a global optimum can, thus, not be expected.

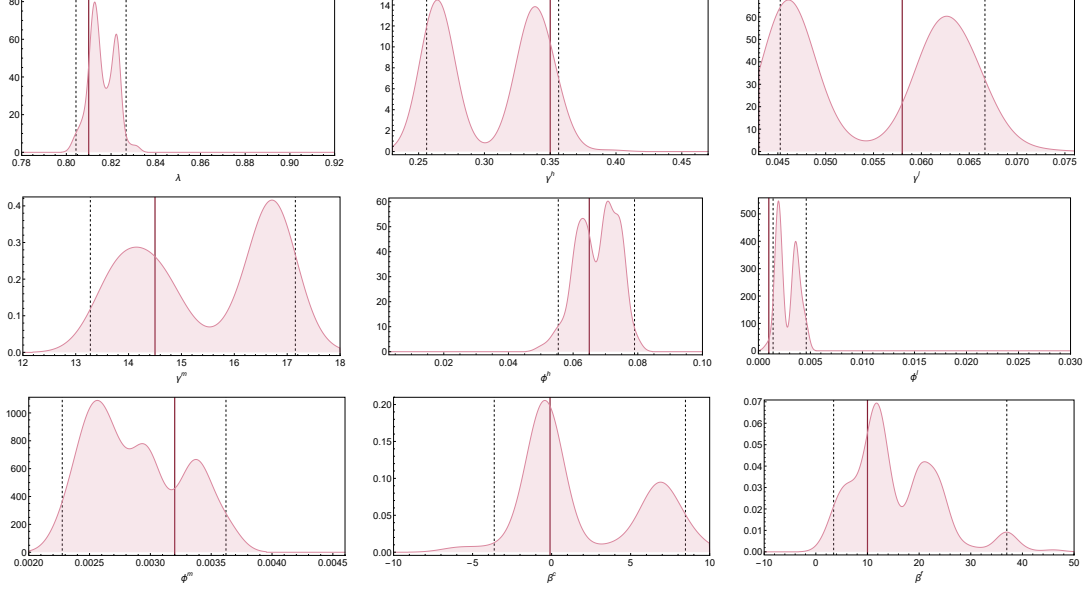


Figure 7: Combined densities of (theoretical) parameter estimates. The panels show smoothed kernel densities for parameter estimates with red lines indicating the (pseudo-) true values and dotted black lines visualizing the 95% confidence intervals of the sample estimates coming from both settings for $T = 20\,000$.

5 Empirical Application

In this section, we proceed with an empirical application and apply our approach to estimate the model for several financial time series. Since Schmitt et al. (2020) only run calibration exercises for the model, we will provide a thorough empirical validation with which we are able to investigate the model in a far better way. Using different asset classes, including stock market indices, commodities and foreign exchange rates, will allow us to get valuable insights. We use three stock market indices which are the German DAX, the S&P500 and the Japanese Nikkei225. As commodities, we chose the price of gold. To account for the specifics of foreign exchange markets, we also include the U.S. dollar to euro (USD/EUR). All five time series cover 25 years from January 1, 1993 to December 31, 2017 with the total number of daily observations ranging between 6146 to 6421. In table

6, we report for each empirical dataset all ten moments together with their standard deviations. Since the standard deviations of the empirical moments need to be approximated, we resort to do this via bootstrapping.¹⁷ Following Winker et al. (2007), we used a block bootstrap to account for the long-range dependence in the empirical time series when computing the frequency distributions for the volatility and the tail indices. We, hence, subdivided the original time series into blocks of 250 daily observations and constructed 5 000 block bootstrapped time series from random draws with replacement. For the frequency distributions of the lagged autocorrelations, we consider the reasoning of Franke and Westerhoff (2016) that the long-range dependence in the return series gets distorted when two non-adjacent blocks are put consecutively. To avoid a bias in these statistics, we sampled single observations together with their accordingly lagged past data points as they suggest.¹⁸ For the empirical estimation, we use the same setup from the previous numerical exercise reported in 4.2. For all five empirical time series, we run a total of 100 independent estimations over $T = 20\,000$.¹⁹ The results of our empirical study are reported in Table 5 and visually depicted in Figure 8. Based on our findings in the previous section, we are well aware that part of the parameter estimates are most likely biased to some degree. Using the same initial conditions for the estimation of all empirical time series, the probable biases are supposedly consistent. Hence, we are still able to get some insights when comparing the empirical estimates. A first inspection of the parameter estimates shows that the values are quite similar for the SP500, Nikkei, Gold and DAX. This is not much of a big surprise when looking at the empirical moments (see Appendix B) which are in a close range for these four time series. Looking at the estimates for reaction parameters of chartists β^c and fundamentalists β^f , we recognize strong fundamental forces. In case of the first four series, fundamentalists tend to react around six times more strongly than chartists to given trading signals. Further, we found contrarian behavior on average for the SP500 and Nikkei index paired with very high mean estimates for the reaction parameter of chartists β^f . In contrast, we find trend-following behavior for Gold, the DAX index and the USD-EUR ex-

¹⁷Alternatively, the Newey-West estimator of the covariance-matrix of the empirical moments might be used to obtain confidence intervals. Yet, those intervals might be less narrow and hence not optimal.

¹⁸If an estimated moment from the empirical time series happened to differ significantly from its computed median value, we increased the number of bootstrap repetitions up to 50 000 until the moment's frequency distribution centered around its empirical observation.

¹⁹Note that the simulated time length is significantly longer than the empirical time series. As stated earlier, running simulations on a smaller time length would result in high sampling variations and thus leading to strongly biased parameter estimates.

	SP500	Nikkei	Gold	DAX	USD-EUR		USD-EUR
λ	0.8108	0.7939	0.8908	0.7992	0.8723	λ	0.9092
$\in [0.78, 0.92]$	(0.0064)	(0.0087)	(0.0031)	(0.0116)	(0.0073)	$\in [0.83, 0.925]$	(0.0028)
γ^h	0.3759	0.3695	0.2911	0.3999	0.2788	γ^h	0.1766
$\in [0.25, 0.46]$	(0.0136)	(0.0239)	(0.0084)	(0.0101)	(0.0208)	$\in [0.15, 0.3]$	(0.0031)
γ^l	0.0515	0.0640	0.0718	0.0518	0.0539	γ^l	0.0533
$\in [0.046, 0.076]$	(0.0012)	(0.0021)	(0.0022)	(0.0044)	(0.0012)	$\in [0.04, 0.058]$	(0.0006)
γ^m	16.4974	15.8042	17.5011	16.4373	15.3356	γ^m	10.7996
$\in [13.5, 18]$	(0.4641)	(0.4538)	(0.4251)	(0.6546)	(0.6641)	$\in [10, 15]$	(0.1462)
ϕ^h	0.063	0.0662	0.0567	0.0584	0.0085	ϕ^h	-
$\in [0.003, 0.1]$	(0.0051)	(0.0070)	(0.0061)	(0.0059)	(0.0049)		
ϕ^l	0.0021	0.0019	0.0019	0.0025	0.0008	ϕ^l	0.0003
$\in [0.0, 0.003]$	(0.0002)	(0.0004)	(0.0003)	(0.0005)	(0.0003)	$\in [0.0, 0.003]$	(0.0002)
ϕ^m	0.0034	0.0036	0.0028	0.0030	0.0033	ϕ^m	-
$\in [0.0023, 0.0040]$	(0.0003)	(0.0003)	(0.0002)	(0.0004)	(0.0004)		
β^c	-9.1625	-5.9052	2.7243	4.1345	4.4039	β_c	7.7633
$\in [-10, 10]$	(0.5034)	(3.1288)	(5.0827)	(4.5304)	(3.4059)	$\in [-10, 10]$	(1.7912)
β^f	47.575	35.7463	16.4968	26.4697	6.9985	β_f	9.1051
$\in [1, 50]$	(1.534)	(10.5134)	(9.6083)	(11.7108)	(5.5135)	$\in [1, 50]$	(5.6502)
SJM	0.192	0.099	0.247	0.226	0.0594		0.2102
	(0.0076)	(0.0075)	(0.0177)	(0.0256)	(0.0149)		(0.0105)

Table 5: Simulation results of the empirical application. The table shows the mean of the parameter estimates. The used parameter intervals are given in the first and seventh column. For better readability, parameter estimates for γ^h , γ^l , γ^m , β^c and β^f are multiplied by 10^5 . The final row shows the mean value of the SJM estimator with estimations based on 5 000 repetitions over time length T set equal to the corresponding empirical time series. Standard errors are reported in parentheses.

change rate. Speaking of Gold, it is interesting to see that the memory parameter λ is especially high. This is in line with other research, e.g. Baur and Glover (2014), who found long-term oriented trend-following behavior for the price of gold.

With regard to the second main mechanism of the economic model, we found the herding parameter ϕ^h to be significant and similar for almost all time series except for the USD-EUR exchange rate. In fact, the estimates for the USD-EUR exchange rate differ significantly for most parameters from those of the other series. First, we find overall low values for the trading intensity mechanism defined by parameters γ^h , γ^m and γ^l together with a comparatively high memory parameter λ . Second, the estimates for the parameters of the coordination/herding mechanism reveal statistical insignificance. Finally, we observe reaction parameters for chartists β^c and fundamentalists β^f which are almost equal in strength. Given these results, we hence decided for an slightly adjusted second estimation run where we cut off strong herding effects and allowed only mild yet persistent coordination behavior

among speculators ruled by parameter ϕ^l . See the last two columns of Table 5 for the adjusted parameter space and the corresponding results. Estimations of this second run reveal that speculators' trading intensity is rather weak. Speculators tend to be long-term oriented given a high value for the memory parameter λ . Since the time series of the USD-EUR exchange rate barely shows extreme events (highest daily return is around 4 percent, see Appendix A for a sample simulation run), speculators' herding behavior appears weak, too. Reaction parameters are again of comparable strength.

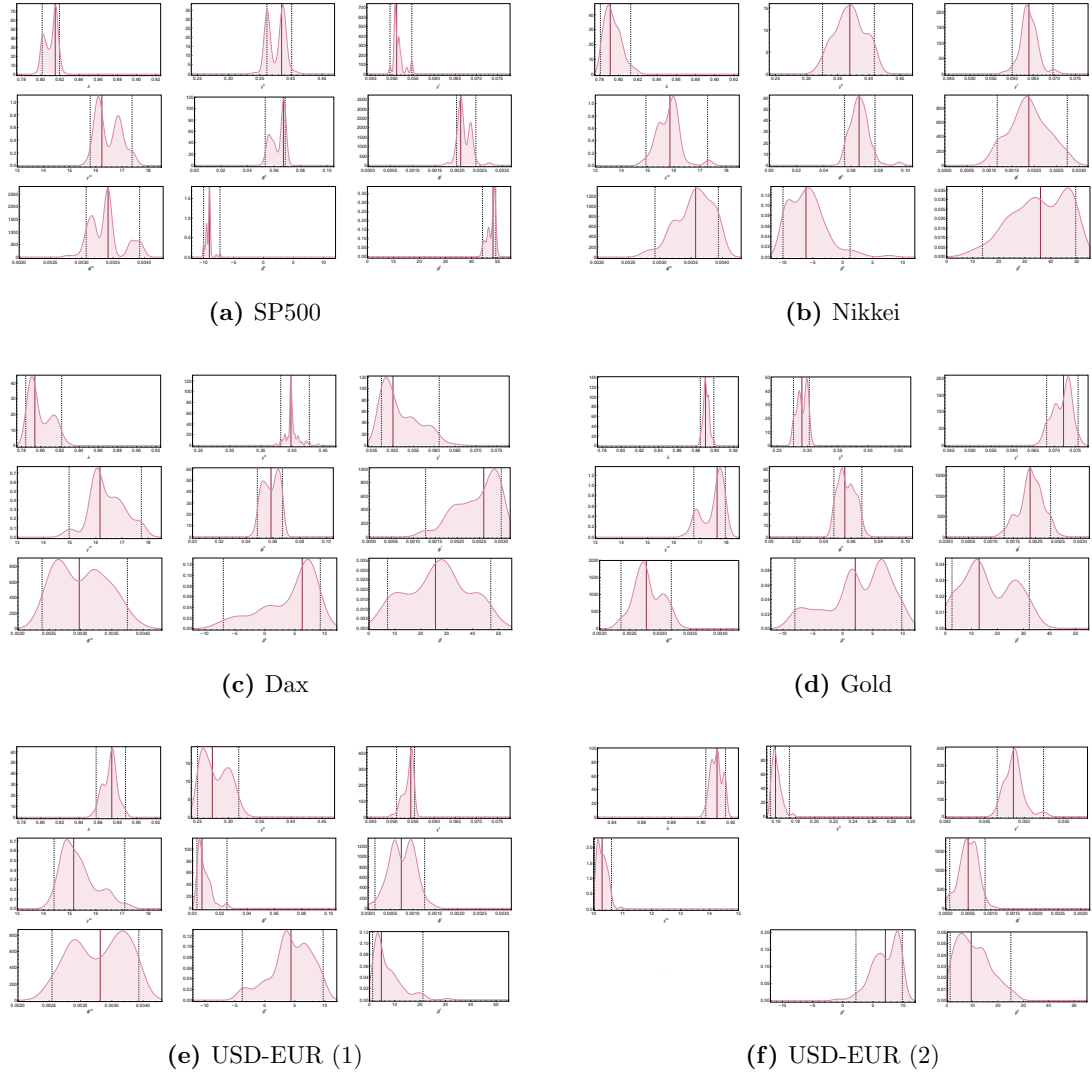


Figure 8: Densities of empirical parameter estimates. The panels show smoothed kernel densities for parameter estimates with red lines indicating the mean values and dotted black lines visualizing the 95% confidence intervals of the sample estimates.

Figure 8 shows the smoothed kernel densities for parameter estimates. Red lines indicate the median values and solid block lines mark the 95% confidence

intervals of the sample estimates. Given the small sample sizes we are using here, it is no surprise that parameter estimates suffer from right-hand and left-hand skewness. Most densities also show at least some degree of multimodality (mostly bimodality). The bimodality of the parameter estimates tends to be more pronounced for the SP500 and Gold, is less pronounced for DAX and USD-EUR and least pronounced for the Nikkei. Overall, we find the multimodality to be less distinct than in the purely numerical studies before suggesting that the quality and quantity of empirical data is sufficient enough. It is also interesting to see that the standard errors are similar in size for most of the parameters. Only for the reaction coefficients of chartists and fundamentalists, we find some differences. While parameter estimates of β^c and β^f are narrowly distributed for the SP500, the densities for Nikkei, Gold and Dax show much greater variability. Hence, it might make sense to pre-calibrate them to achieve more efficiency during estimations for these three time series.

Let us next assess the validity of the empirical parameter estimates. The SJM's great advantage is that it comes with an inherent goodness-of-fit measure. Therefore, we calculate the fraction of simulation runs for which all simulated moments jointly fall into their empirical counterparts. This is a simple, intuitive and computationally cheap way. Recall that we used an extended time length with $T = 20\,000$ for the empirical estimation. To get unbiased estimates of the SJM estimator, we re-evaluated the results for the same time length as the empirical time series (ranging between 6146 and 6421 daily observations) over 5 000 repetitions. The last row of Table 5 reports the mean estimates for the joint-moment measure which ranges from almost 10% to 25% with the Nikkei performing worst and Gold performing best. That means for the Nikkei that for every tenth simulation run all ten moments jointly fall into their empirical counterparts. And for Gold it is even every fourth simulation run. Given the simplicity of the economic model and the known struggle to estimate certain model parameters, we find the overall performance of the model to be fairly satisfactory. Sample simulation runs vis-a-vis their empirical counterparts can be found in Appendix A.

6 Conclusion

In this paper, we study the performance of a simulated joint moment estimator as first proposed by Franke and Westerhoff (2012) which has been used (mostly) for calibration so far. This estimator belongs to the bigger group of simulated moment estimators which is rich but varies in details. We aim at contributing to this line of

research by providing a comprehensive and thorough estimation study of the SJM. The economic model under scrutiny here is a simple agent-based computational stock market model by Schmitt et al. (2020) which is able to replicate a number of important stylized facts of stock markets.

In a first bunch of sensitivity analyses, we reveal that the estimator’s objective function is non-smooth and has partially flat response surfaces. We further show that the response surface is strongly dependent on both simulated and empirical time length, plus the number of simulation runs. Hence, stochasticity causes identification problems here. To overcome these, we can increase the time lengths and number of simulation runs leading to a massive computational burden. A machine-learning surrogate helps to reduce the computational costs and to maximize the objective function. A (relatively massive) initial grid search provides a first rough mapping of the objective function and is used to feed the machine learning meta-model. The proposed methodology offers an easily scalable estimation framework which can be run on any standard home computing device. In fact, we find that our approach leads to a reduction in computation time of at least 90%. There is, hence, no need for a High-Performance Computing Cluster.

Running a massive Monte Carlo study, we test if our estimation framework is able to precisely estimate model parameters. We show that most parameters of the SSW model are correctly identified. Results also suggest biases in the estimates and convergence behavior for just some model parameters in small sample populations. To be precise, parameter estimates show bi-modal densities. This is a common and persistent pattern across model parameters and experimental settings. The reason lies in the functional form of the theoretic model with its S-shaped functions for the intensity and correlation of speculators’ trading leading to high correlation between model parameters. Considering our economic model as naturally misspecified (as is every model), it is essential to understand the relation between model and estimation routine. This allows to detect possible biases and carefully interpret meaningful results. Given the dimensionality of the economic model, the number of model parameters and the limited computational resources, we find the precision of the estimates quite satisfying.

The joint moment estimator has often been criticized for being somewhat arbitrary and offering limited validity and comparability. It is indeed quite sensitive to a lot of assumptions, e.g. time length T , the number of Monte Carlo repetitions MC or the set of chosen moments. Hence, comparisons across different studies might not be feasible. What this work offers is a focus on the distributional properties of the estimator instead of a single point estimate. This way, evaluations

of the performance of the model are much more insightful. This is also the first study to use this kind of estimator on various empirical time series. Indeed, we have demonstrated that the model by Schmitt et al. (2020) is valid for all five datasets. This compensates, at least to some degree, for the missing comparability across studies.

We have further shown that the machine-learning surrogates are prone to convergence to local optima too. This is most probably due to the specific choice and size of the pool of parameter combinations which we identify as the main weak spot of this approach. Due to limited memory capacity, the size of the pool cannot be set arbitrarily large.

Our work offers possible new research directions. Inspired by the recent work of Lux (2020b) and Platt (2020) who test the performance of different estimation routines, we find that a similar study could be a worthwhile task. In this regard, we would like to test an upscaled version of our framework given more computational resources and benchmark it against other estimation approaches. Moreover, the broader applicability of our estimation framework opens avenues for future research. For instance, it may be interesting to check and see if it is still useful and computationally manageable in larger and more complex settings of e.g. macroeconomic models. One possible direction for future work could also be to use a mixture of surrogates including classification considering CatBoost’s superiority in handling categorical features. Since our approach is somewhat related to Approximate Bayesian Computation, we also see potential for a methodological marriage here.

References

- [1] 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.
- [2] Alfarano, S., Lux, T. and Wagner, F. (2006): Estimation of a simple agent-based model of financial markets: An application to Australian stock and foreign exchange data. *Physica A: Statistical Mechanis and its Applications*, 370, 38-42.
- [3] Alfarano, S. and Lux, T. (2007): A noise trader model as a generator of apparent financial power laws and long memory. *Macroeconomic Dynamics*, 11, 80-101.
- [4] Amilon, H. (2008): Estimation of an adaptive stock market model with heterogeneous agents. *Journal of Empirical Finance*, 15, 342-362.
- [5] Barde, S. (2016): A practical, accurate, information criterion for nth order markov processes. *Computational Economics*, 1-44.
- [6] Barde, S. and van der Hoog, S. (2017): An empirical validation protocol for large-scale agent-based models. *Studies in Economics*, 1712, School of Economics, University of Kent.
- [7] Bargigli, L., Riccetti, L., Russo, A., and Gallegati, M. (2018): Network calibration and metamodeling of a financial accelerator agent based model. *Journal of Economic Interaction and Coordination*.
- [8] Baur, D. and Glover, K.J. (2014): Heterogeneous expectations in the gold market: Specification and estimation. *Journal of Economic Dynamics and Control*, 40, 116-133.
- [9] Bertschinger, N. and Mozzhorin, I. (2020): Bayesian estimation and likelihood-based comparison of agent-based volatility models. *Journal of Economic Interaction and Coordination*.
- [10] Boswijk, P., Hommes, C. and Manzan, S. (2007): Behavioral heterogeneity in stock prices. *Journal of Economic Dynamics and Control*, 31, 1938-1970.
- [11] Chen, S.-H., Chang, C.-L. and Du, Y.-R. (2012): Agent-based economic models and econometrics. *The Knowledge Engineering Review*. 27, 187–219.

- [12] Chen, Z. and Lux, T. (2018): Estimation of Sentiment Effects in Financial Markets: A Simulated Method of Moments Approach. *Computational Economics*, 52, 711–744.
- [13] Chiarella, C., Dieci, R., He, X.-Z. (2009a): Heterogeneity, market mechanisms, and asset price dynamics. In: Hens, T., Schenk-Hoppé, K.R. (eds.): *Handbook of financial markets: dynamics and evolution*. North-Holland, Amsterdam, 277-344.
- [14] Cont, R. (2001): Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1, 223-236.
- [15] Creel, M. and Kristensen, D. (2016): On selection of statistics for approximate Bayesian computing (or the method of simulated moments). *Computational Statistics & Data Analysis*, 100, 99-114.
- [16] Day, R. and Huang, W. (1990): Bulls, bears and market sheep. *Journal of Economic Behavior and Organization*, 14, 299-329.
- [17] Delli Gatti, D., Fagiolo, G., Gallegati, M., Richiardi M. and Russo, A. (2018): *Agent-based models in economics: a toolkit*. Cambridge University Press, Cambridge.
- [18] Dieci, R. and He, X.-Z., (2018): Heterogeneous agent models in finance. In: Hommes, C., LeBaron, B. (Eds.), *Handbook of Computational Economics*, Vol. 4, *Heterogeneous Agent Modeling*. North-Holland, 257–328.
- [19] Dosi, D., Pereira, M.C., Roventini, A. and Virgillito, M.E. (2018a): The effects of labour market reforms upon unemployment and income inequalities: an agent-based model. *Socio-Economic Review*, 16, 687–720.
- [20] Dosi, G., Pereira, M.C. and Virgillito, M.E. (2018b): On the robustness of the fat-tailed distribution of firm growth rates: a global sensitivity analysis. *Journal of Economic Interaction and Coordination*, 13, 173–193.
- [21] Dosi, G., Pereira, M.C., Roventini, A. and Virgillito, M.E. (2018c): Causes and consequences of hysteresis: aggregate demand, productivity, and employment. *Industrial and Corporate Change*, 27, 1015–1044.
- [22] Duffie, D. and Singleton, K.J. (1993): Simulated moments estimation of Markov models of asset prices. *Econometrica*, 61, 929–52.

- [23] Ellen, S. and Verschoor, W.F.C. (2018): Heterogeneous Beliefs and Asset Price Dynamics: A Survey of Recent Evidence. In: Jawadi, F. (eds.): Uncertainty, Expectations and Asset Price Dynamics. Dynamic Modeling and Econometrics in Economics and Finance, 24, Springer, Cham.
- [24] Fabretti, A. (2013): On the problem of calibrating an agent based model for financial markets. *Journal of Economic Interaction and Coordination*, 8, 277–293.
- [25] Fagiolo, G., Guerini, M., Lamperti, F., Moneta, A. and Roventini, A. (2019): Validation of Agent-Based Models in Economics and Finance. In: Beisbart C., Saam N. (eds.): Computer Simulation Validation. Simulation Foundations, Methods and Applications. Springer, Cham.
- [26] Franke, R. (2009): Applying the method of simulated moments to estimate a small agent-based asset pricing model. *Journal of Empirical Finance*, 16, 804–815.
- [27] Franke, R. and F. Westerhoff (2011): Estimation of a structural stochastic volatility model of asset pricing. *Computational Economics*, 38, 53–83.
- [28] Franke, R. and Westerhoff, F. (2012): Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *Journal of Economic Dynamics and Control*, 36, 1193–1211.
- [29] Franke, R. and Westerhoff, F. (2016): Why a simple herding model may generate the stylized facts of daily returns: explanation and estimation. *Journal of Economic Interaction and Coordination*, 11, 1–34.
- [30] Gilli, M. and Winker, P. (2003): A global optimization heuristic for estimating agent based models. *Computational Statistics & Data Analysis*, 42(3), 299–312.
- [31] Gilli, M. and Winker, P. (2009): Heuristic Optimization Methods in Econometrics. In: Belsley, D.A. and Kontoghiorghes, E.J. (eds.): *Handbook of Computational Econometrics*, 81–119, Wiley, New York.
- [32] Grammig, J. and Schaub, E.M. (2015): Give Me Strong Moments and Time: Combining GMM and SMM to Estimate the Long-Run Risk Asset Pricing Model. CFS Working Paper, 479.
- [33] Grazzini J., Richiardi, M. and Sella, L. (2012): Small sample bias in MSM estimation of agent-based models. In: Teglio A., Alfarano S., Camacho-Cuena,

- E. and Ginés-Vilar, M. (eds.): Managing Market Complexity. Lecture Notes in Economics and Mathematical Systems, 662. Springer, Berlin, Heidelberg.
- [34] Grazzini, J. and Richiardi, M. (2015): Estimation of ergodic agent-based models by simulated minimum distance. *Journal of Economic Dynamics and Control*, 51, 148-165.
- [35] Guerini, M. and Moneta, A. (2017): A method for agent-based models validation. *Journal of Economic Dynamics and Control*, 82, 125-141.
- [36] Hill, B. (1975): A simple general approach to inference about the tail of a distribution. *Annals of Statistics*, 3, 1163-1174.
- [37] Hommes, C. (2006): Heterogeneous agent models in economics and finance. In: Tesfatsion, L. and Judd, K. (eds.): *Handbook of computational economics: agent-based computational economics*. North-Holland, Amsterdam, 1109-1186.
- [38] Iori, G. and Porter, J. (2018): Agent based Modelling for Financial Markets. In: Chen, S.-H., Kaboudan, M. and Du, Y.-R. (eds.): *The Oxford Handbook on Computational Economics and Finance*, OUP Oxford University Press.
- [39] Jang, T.-S. (2015): Identification of Social Interaction Effects in Financial Data. *Computational Economics*, 45, 207-238.
- [40] Kucherenko, S., Albrecht, D. and Saltelli, A. (2015): Exploring multi-dimensional spaces: A comparison of Latin hypercube and quasi Monte Carlo sampling techniques. arXiv preprint arXiv:1505.02350.
- [41] Kukacka, J. and Barunik, J. (2017): Estimation of financial agent-based models with simulated maximum likelihood. *Journal of Economic Dynamics and Control*, 85, 21-45.
- [42] Lamperti, F. (2018a): An information theoretic criterion for empirical validation of simulation models. *Econometrics and Statistics*, 5, 83-106.
- [43] Lamperti, F. (2018b): Empirical validation of simulated models through the GSL-div: an illustrative application. *Journal of Economic Interaction and Co-ordination*, 13, 143-171.
- [44] Lamperti, F., Roventini, A. and Sani, A. (2018): Agent-based model calibration using machine learning surrogates. *Journal of Economic Dynamics and Control*, 90, 366-389.

- [45] LeBaron, B. (2006): Agent-based computational finance. In: Tesfatsion, L. and Judd, K. (Eds.): Handbook of computational economics: agent-based computational economics. North-Holland, Amsterdam, 1187-1233.
- [46] Lee, B. and Ingram, B. (1991): Simulation estimation of time series models. Journal of Econometrics, 47, 197-205.
- [47] Lux, T. (2009): Stochastic behavioural asset-pricing models and the stylized facts. In: Hens, T. and Schenk-Hoppé, K.R. (eds.): Handbook of Financial Markets: Dynamics and Evolution. North-Holland, Amsterdam, 161-216.
- [48] 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.
- [49] Lux, T. (2018): Estimation of agent-based models using sequential Monte Carlo methods. Journal of Economic Dynamics and Control, 91, 391-408.
- [50] Lux, T. (2020a): Bayesian estimation of agent-based models via adaptive particle Markov chain Monte Carlo. Economics Working Paper, Christian-Albrechts-University of Kiel.
- [51] Lux, T. (2020b): Inference for Nonlinear State Space Models: A Comparison of Different Methods applied to Markov-Switching Multifractal Models. Econometrics and Statistics, in press.
- [52] Lux, T. and Zwinkels, R. (2018): Empirical validation of agent-based models. In: Hommes, C. and LeBaron, B. (eds.): Handbook of Computational Economics, Vol. 4, Heterogeneous Agent Modeling. North-Holland, Amsterdam, 437-488.
- [53] Mantegna, R. and Stanley, E. (2000): An introduction to econophysics. Cambridge University Press, Cambridge.
- [54] Marks, R. (2013): Validation and model selection: three similarity measures compared. Complexity Economics, 2, 41-61.
- [55] McFadden, D. (1989): A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. Econometrica, 57, 995-1026.

- [56] Pakes, A. and Pollard, D. (1989): Simulation and the Asymptotics of Optimization Estimators. *Econometrica*, 57, 1027-1057.
- [57] Platt, D. and Gebbie, T. (2018): Can agent-based models probe market microstructure? *Physica A: Statistical Mechanics and its Applications*, 503, 1092-1106.
- [58] Platt, D. (2022): Bayesian estimation of economic simulation models using neural networks. *Computational Economics*, 59(2): 599-650.
- [59] Platt, D. (2020): A comparison of economic agent-based model calibration methods. *Journal of Economic Dynamics and Control*, 113: 103859.
- [60] Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A.V. and Gulin, A. (2018): CatBoost: unbiased boosting with categorical features. *Advances in Neural Information Processing Systems*, 31, 6638-6648.
- [61] Salle, I. and Yildizoglu, M. (2014): Efficient sampling and meta-modeling for computational economic models. *Computational Economics*, 44, 507-536.
- [62] Schmitt, N. (2021): Heterogeneous expectations and asset price dynamics. *Macroeconomic Dynamics*, Vol. 25, 1538–1568.
- [63] Schmitt, N., Schwartz, I. and Westerhoff, F. (2020): Heterogeneous speculators and stock market dynamics: a simple agent-based computational model. *The European Journal of Finance*, forthcoming.
- [64] Secchi, D. and Seri, R. (2017): Controlling for false negatives in agent-based models: a review of power analysis in organizational research. *Computational and Mathematical Organization Theory*, 23, 94–121.
- [65] Shiller, R. (2015): *Irrational exuberance*. Princeton University Press, Princeton.
- [66] ten Broeke, G., van Voorn, G. and Ligtenberg, A. (2016): Which sensitivity analysis method should i use for my agent-based model. *Journal of Artificial Societies and Social Simulation*, 19.
- [67] van der Hoog, S. (2019) *Surrogate Modelling in (and of) Agent-Based Models: A Prospectus*. *Computational Economics*, 53, 1245–1263.
- [68] Varian, H.R. (2014): Big Data: New Tricks for Econometrics. *Journal of Economic Perspectives*, 28, 3-28.

- [69] Windrum, P., Fagiolo, G. and Moneta A. (2007): Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, 10.
- [70] 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.
- [71] Winker, P. and Maringer, D. (2009): The convergence of estimators based on heuristics: theory and application to a GARCH model. *Computational Statistics*, 24, 533–550.
- [72] Zheng, A. and Casari, A. (2018): *Feature Engineering for Machine Learning: Principles and Techniques for Data Scientists*. O'Reilly Media, Sebastopol.

A Sample simulation runs

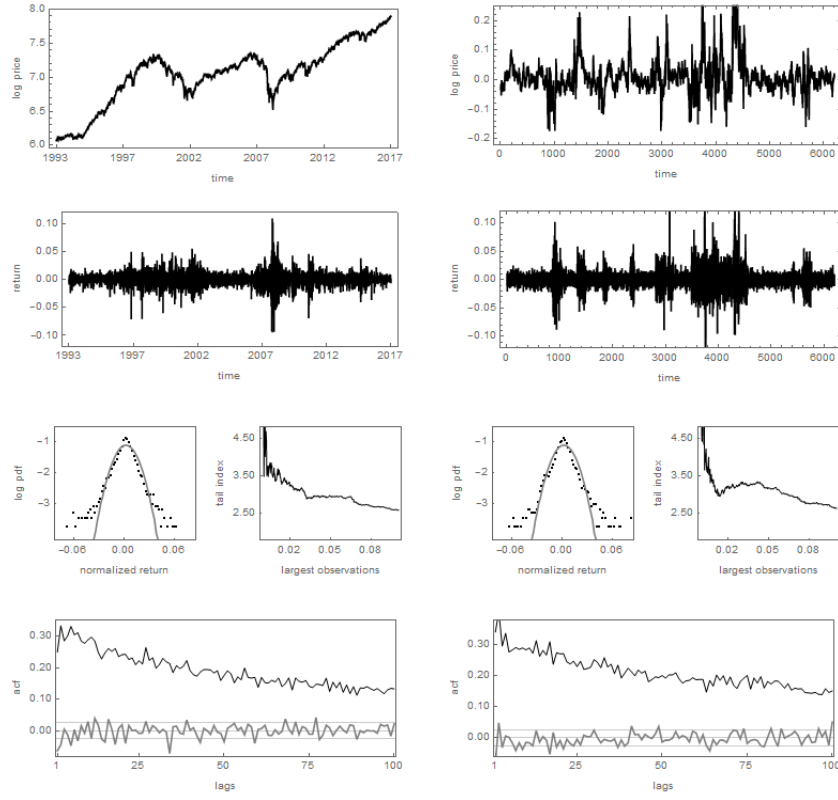


Figure 9: Dynamics of the SP500 between 1993 and 2017 (left panel) and dynamics of a sample simulation run (right panel) using the following parameter set: $\lambda = 0.82$, $\gamma^h = 0.36$, $\gamma^l = 0.051$, $\gamma^m = 16.1$, $\phi^l = 0.002$, $\phi^h = 0.067$, $\phi^m = 0.0034$, $\beta^c = -0.0001$, $\beta^f = 0.0005$. The panels show from top to bottom the price evolution, the log returns, the log probability density functions of normalized returns (black) and standard normally distributed returns (gray), the Hill tail index estimator and the autocorrelation functions of raw (gray) and absolute returns (black), respectively.

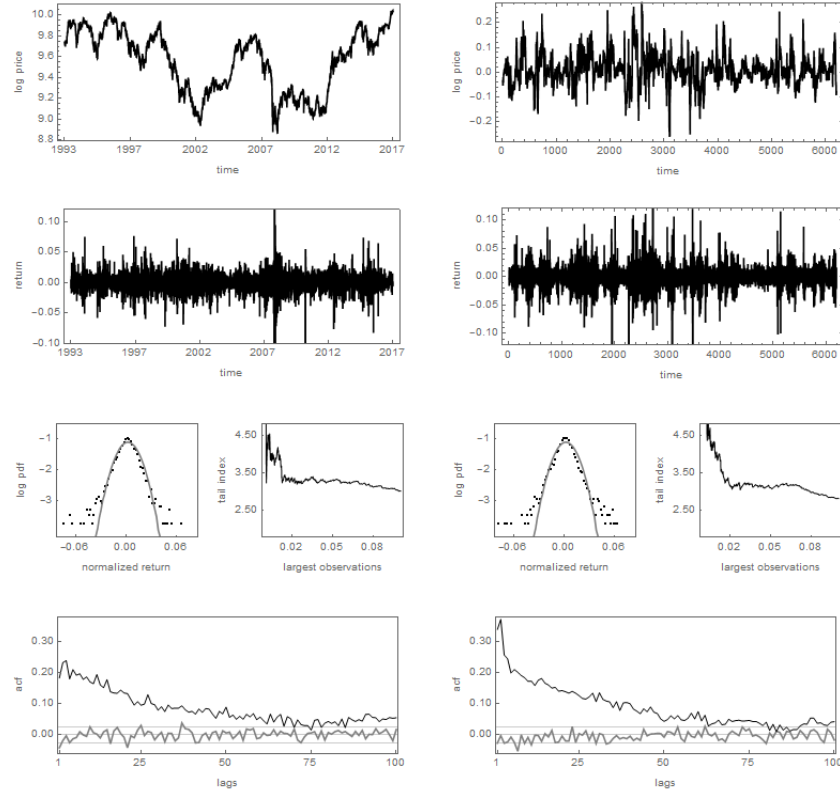


Figure 10: Dynamics of the Nikkei225 between 1993 and 2017 (left panel) and dynamics of a sample simulation run (right panel) using the following parameter set: $\lambda = 0.80$, $\gamma^h = 0.34$, $\gamma^l = 0.064$, $\gamma^m = 16.2$, $\phi^l = 0.002$, $\phi^h = 0.073$, $\phi^m = 0.0036$, $\beta^c = 0.000024$, $\beta^f = 0.00047$. The panels show from top to bottom the price evolution, the log returns, the log probability density functions of normalized returns (black) and standard normally distributed returns (gray), the Hill tail index estimator and the autocorrelation functions of raw (gray) and absolute returns (black), respectively.

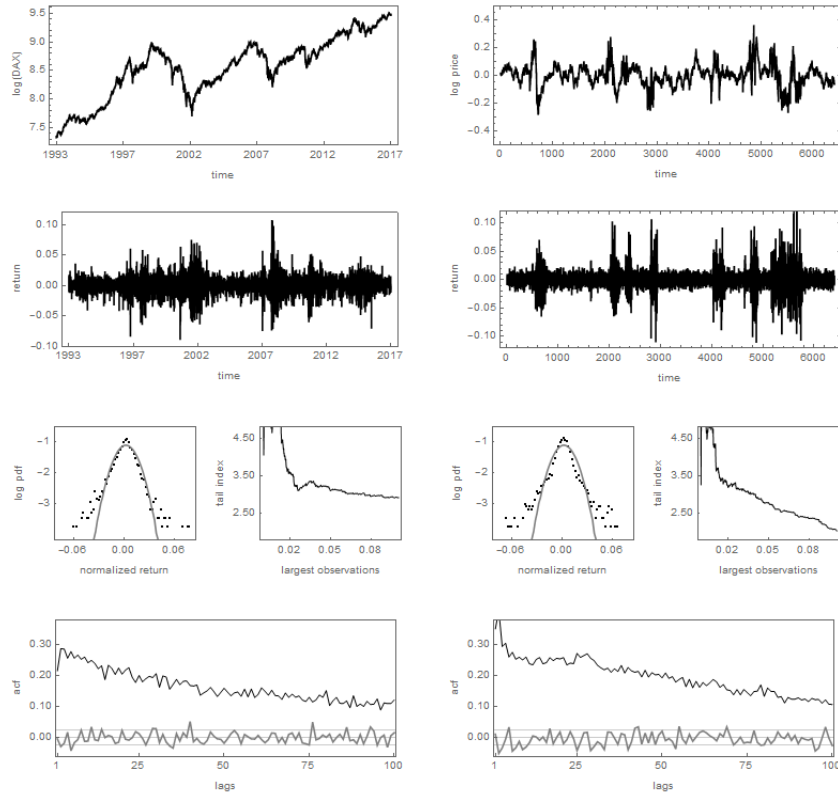


Figure 11: Dynamics of the German DAX between 1993 and 2017 (left panel) and dynamics of a sample simulation run (right panel) using the following parameter set: $\lambda = 0.82$, $\gamma^h = 0.39$, $\gamma^l = 0.061$, $\gamma^m = 17.7$, $\phi^l = 0.002$, $\phi^h = 0.05$, $\phi^m = 0.0024$, $\beta^c = -0.000016$, $\beta^f = 0.00016$. The panels show from top to bottom the price evolution, the log returns, the log probability density functions of normalized returns (black) and standard normally distributed returns (gray), the Hill tail index estimator and the autocorrelation functions of raw (gray) and absolute returns (black), respectively.

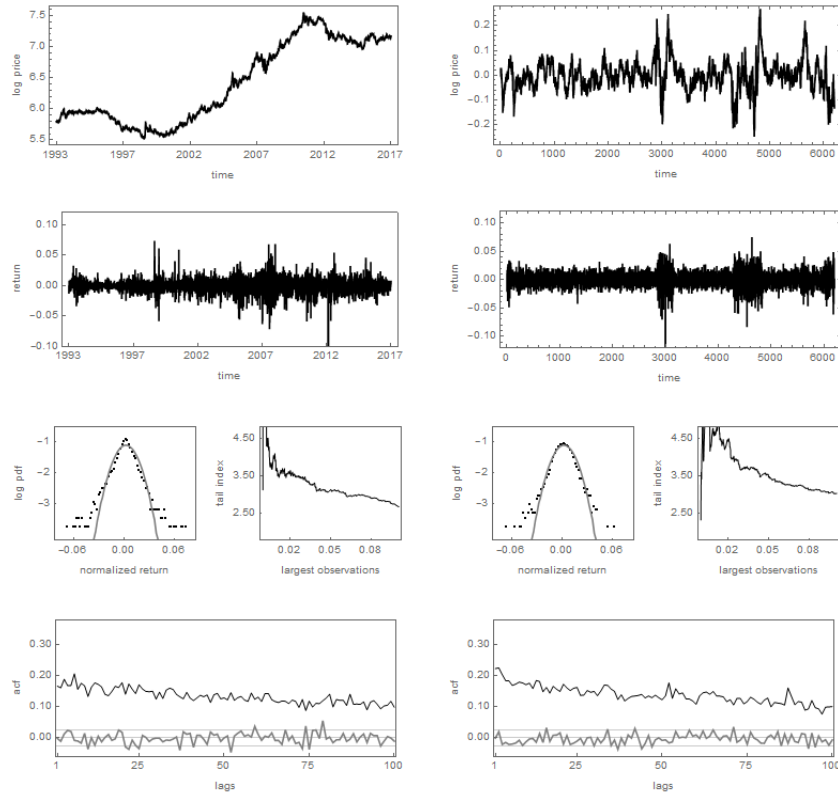


Figure 12: Dynamics of the gold price between 1993 and 2017 (left panel) and dynamics of a sample simulation run (right panel) using the following parameter set: $\lambda = 0.90$, $\gamma^h = 0.27$, $\gamma^l = 0.074$, $\gamma^m = 17.1$, $\phi^l = 0.002$, $\phi^h = 0.049$, $\phi^m = 0.0024$, $\beta^c = 0.00006$, $\beta^f = 0.00024$. The panels show from top to bottom the price evolution, the log returns, the log probability density functions of normalized returns (black) and standard normally distributed returns (gray), the Hill tail index estimator and the autocorrelation functions of raw (gray) and absolute returns (black), respectively.

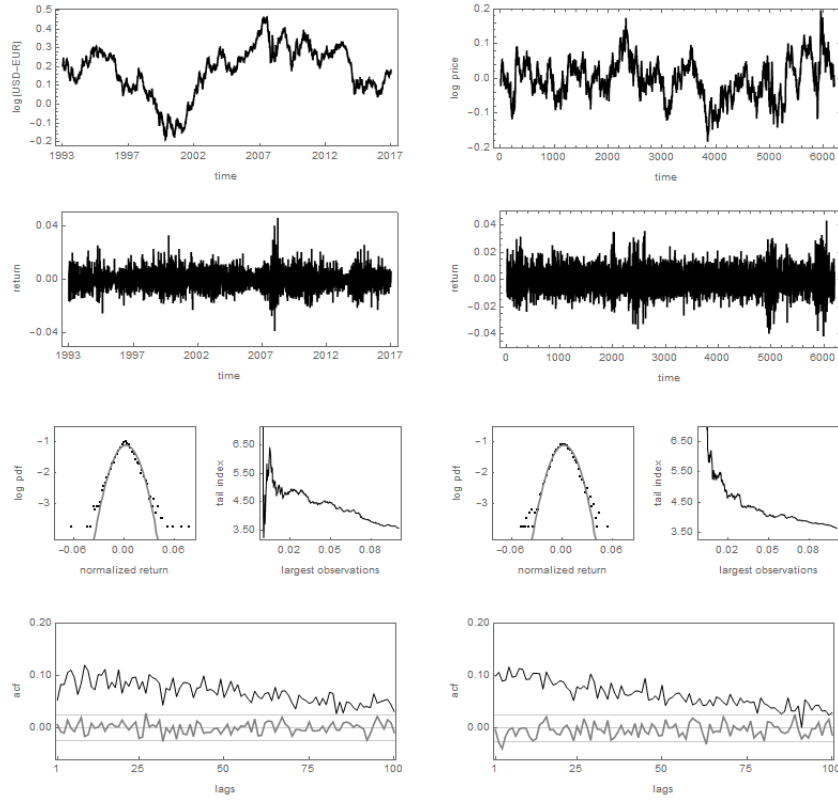


Figure 13: Dynamics of the USD-EUR exchange rate between 1993 and 2017 (left panel) and dynamics of a sample simulation run (right panel) using the following parameter set: $\lambda = 0.91$, $\gamma^h = 0.17$, $\gamma^l = 0.053$, $\gamma^m = 10.5$, $\phi^l = 0.0001$, $\beta^c = 0.0001$, $\beta^f = 0.0001$. The panels show from top to bottom the price evolution, the log returns, the log probability density functions of normalized returns (black) and standard normally distributed returns (gray), the Hill tail index estimator and the autocorrelation functions of raw (gray) and absolute returns (black), respectively.

B Empirical moments and their bootstrapped standard deviations

	SP500	Nikkei	Gold	DAX	USD-EUR
$E(r_t)$	1.136 (0.109)	1.487 (0.098)	1.022 (0.085)	1.435 (0.103)	0.603 (0.03)
$\alpha_{2.5}$	3.124 (0.192)	3.235 (0.227)	3.464 (0.228)	3.232 (0.219)	4.877 (0.272)
$\alpha_{5.0}$	2.954 (0.137)	3.317 (0.17)	3.143 (0.139)	3.104 (0.148)	4.466 (0.248)
$E(r_t r_{t-1})$	-0.062 (0.023)	-0.042 (0.022)	-0.002 (0.02)	-0.004 (0.019)	0.007 (0.014)
$E(r_t r_{t-3})$	0.285 (0.025)	0.24 (0.031)	0.19 (0.021)	0.288 (0.023)	0.084 (0.014)
$E(r_t r_{t-6})$	0.305 (0.032)	0.193 (0.028)	0.207 (0.021)	0.258 (0.02)	0.097 (0.016)
$E(r_t r_{t-12})$	0.285 (0.031)	0.195 (0.027)	0.178 (0.018)	0.241 (0.023)	0.073 (0.015)
$E(r_t r_{t-25})$	0.223 (0.024)	0.096 (0.017)	0.172 (0.019)	0.202 (0.018)	0.074 (0.015)
$E(r_t r_{t-50})$	0.176 (0.019)	0.09 (0.015)	0.129 (0.016)	0.161 (0.018)	0.059 (0.014)
$E(r_t r_{t-100})$	0.133 (0.018)	0.055 (0.016)	0.099 (0.016)	0.123 (0.016)	0.032 (0.013)

Table 6: Empirical moments of the S&P500, the Japanese Nikkei, the German DAX, Gold and the USD-EUR exchange rate. The first line contains estimates of the returns standard deviation V , the tail indexes at 2.5 and 5.0%, the autocorrelation coefficients of raw returns for lag 1, and the autocorrelation coefficients of absolute returns for lags 3, 6, 12, 25, 50, 100. The second line shows the standard deviations of the summary statistics that are reported in the first line. Standard deviations have been computed from bootstrapped distributions of the moments. See Section 5 for further details.

5 Time is limited on the road to asymptopia: An analysis of the ergodic properties of moment functions in the validation of financial agent-based models

This chapter contains joint work with Mark Kirstein. Ivonne Schwartz took the lead in designing the study, carried out all numerical simulations, performed all data analyses, interpreted results, designed figures and wrote Sections 2.5, 3 and 4. Mark Kirstein developed the theoretical formalism of the study and discussed its numerical implications. He is responsible for writing Sections 2.1 - 2.4 and designing the corresponding figures. Ivonne Schwartz and Mark Kirstein were involved in equal parts in structuring and writing Section 1. Both authors read, commented and revised the manuscript. In the following, the working paper version will be included.

Time is limited on the road to asymptopia*

An analysis of the ergodic properties of moment functions in the validation of financial agent-based models

IVONNE SCHWARTZ⁺¹ AND MARK KIRSTEIN^{‡2}

¹*University of Bamberg, Department of Economics, Germany*

²*Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany*

29th September 2022

Abstract

One challenge in the estimation of financial market agent-based models (FABMs) is to infer reliable insights using numerical simulations validated by only a single observed time series. Ergodicity (besides stationarity) is a strong precondition for any estimation, however it has not been systematically explored and is often simply presumed. For finite-sample lengths and limited computational resources empirical estimation always takes place in pre-asymptopia. Thus broken ergodicity must be considered the rule, but it remains largely unclear how to deal with the remaining uncertainty in non-ergodic observables. Here we show how an understanding of the ergodic properties of moment functions can help to improve the estimation of (F)ABMs. We run Monte Carlo experiments and study the convergence behaviour of moment functions of two prototype models. We find infeasibly-long convergence times for most. Choosing an efficient mix of ensemble size and simulated time length guided our estimation and might help in general.

Keywords Broken Ergodicity, Simulated Method of Moments, Validation, Calibration, Agent-based Models

JEL Classification C61 · D01 · D81 · D9

*For valuable feedback we thank participants of the Ergodicity Economics 2022 Conference and the virtual seminar on *Quasi-Ergodic Measures* jointly organized by Zachary Adams (MPI MiS) and Maximilian Engel (FU Berlin).

⁺ ✉ ivonne.schwartz@mailbox.org Funded by the Hans-Böckler-Stiftung (PK045) in the Bamberg Doctoral Research Group on Behavioral Macroeconomics (BaGBeM).

[‡] ✉ mark.kirstein@mis.mpg.de

1 Introduction

Financial crashes, global pandemics or disruptive innovations all qualify as formative events during which economic systems reach a new order. On an aggregate level such events can trigger the transition from an old equilibrium to a new equilibrium. A rapid sequence of such formative events can prevent the economy, financial markets and observables from ever reaching their long-time equilibria and thus keeping them steadily out-of-equilibrium. In an effort to use a new methodology to understand such complex emergent behaviors and answer urging policy questions¹ a community of researchers adopted financial agent-based models (FABMs) that deliberately allow for more complexity than traditional models.² The scientific challenge is to work with FABMs that are capable to endogenously reproduce known statistical regularities and that remain statistically and computationally tractable for estimation exercises.³ The ever increasing amounts of data and computational resources make FABMs a promising candidate to ultimately join the set of models which drive informed policy decisions.⁴

The development of a consensual validation and estimation protocol for (F)ABMs is important for scientific standards of reproducibility. However, the estimation and validation of FABMs remains challenging and despite a growing body of research there exists no consensus in the literature, yet.⁵ Having a standardized protocol after all would address main methodological critiques raised against agent-based modeling.⁶

Here we follow a direction outlined in GAFFEO et al. (2007, p. 96) of a ‘descriptive output validation’, where we focus on simulated moment estimators which have been studied already intensively for FABMs.⁷ Apart from the simulated method of moments (SMM) used here, different estimation approaches have recently emerged in the literature including likelihood-based methods, information-theoretic similarity-based measures and more recently a number of Bayesian estimations of FABMs appeared,

¹BOUCHAUD 2008; FARMER and FOLEY 2009; LUX and WESTERHOFF 2009.

²TESFATSION and JUDD 2006; HOMMES and LeBARON 2018.

³The most important stylized facts of financial time series include phenomena like excess volatility, fat-tailed return distributions, absence of autocorrelations in the raw returns but long memory in absolute returns and volatility clustering. Regarding the study of financial markets, there exists a large strand of literature which provides models that explain the stylized facts well. Surveys about the statistical properties of financial markets can be found in CONT (2001), LUX and AUSLOOS (2002), HOMMES (2006), LeBARON (2006) and LUX (2009).

⁴WESTERHOFF and FRANKE 2018.

⁵RICHIARDI et al. (2006) is one attempt in this direction, DELLI GATTI, FAGIOLO et al. (2018) is a helpful textbook. See FAGIOLO, GUERINI et al. (2019) and LUX and ZWINKELS (2018) for recent surveys.

⁶LEOMBRUNI and RICHIARDI 2005; RICHIARDI et al. 2006; FAGIOLO and ROVENTINI 2017.

⁷AMILON 2008; GILLI and WINKER 2003; WINKER, GILLI et al. 2007; FRANKE 2009; SCHMITT 2021; SCHWARTZ 2022.

too.⁸

Reliable estimation rests on sufficient convergence speed given finite-sample lengths and limited computational resources. Surprisingly, the development of better tools and protocols in the validation of (F)ABM has often overlooked the ergodicity issue so far. Our contribution enters exactly at this issue. Already the seminal contribution by CONT (2001, p. 225) emphasised ergodicity in the process of estimation of parameters in models of financial market:

One needs an ergodic property which ensures that the time average of a quantity converges to its expectation. *Ergodicity* is typically satisfied by IID observations but it is *not obvious* – *in fact it may be very hard to prove or disprove* – for processes with complicated dependence properties such as the ones observed in asset returns. (emphasis added)

Admittedly, ergodicity is an elusive property and the fact that no ready-to-use econometric test for ergodicity exists stands in the way of a comprehensive treatment. It may come as less of a surprise then that only few publications identify ergodicity as an important assumption at all.⁹ Regardless, ergodicity often lurks in the background as an implicit regularity condition.¹⁰ As mentioned by CONT (2001) a proof of ergodicity of the relevant observables is not trivial, because it depends on the convergence speed and the respective time scales are often large. The danger for modellers is to run into unnoticed trouble during the process of estimation and validation since knowledge about the estimator's convergence behaviour is instrumental.

In this paper, we aim to contribute to a growing body of research on the validation and estimation of FABMs. In particular, we add useful steps towards a coherent study of the effects of broken ergodicity. We do not propose any new estimation technique, but rather enhance the understanding and applicability of existing ones with regard to the ergodicity issue. Our main contributions are the following. First, there exists a nucleus of publications studying ergodicity in the estimation of (F)ABMs especially by JAKOB GRAZZINI and MATTEO RICHIARDI, which started with the study of models much

⁸KUKACKA and BARUNIK 2017; LAMPERTI 2018; BARDE 2016; GRAZZINI, RICHIARDI and TSIONAS 2017; PLATT 2021; BERTSCHINGER and MOZZHORIN 2021; DELLI GATTI and GRAZZINI 2020; LUX 2018, 2021a,b.

⁹To convey a rough idea, the four volumes of the *Handbook of Computational Economics* only casually mention 'ergodic'; esp. in the volumes most relevant, i.e. Vol. 2 on *Agent-based Computational Finance* (TESFATSION and JUDD 2006) and Vol. 4 on *Heterogeneous Agent Modeling* (HOMMES and LEBARON 2018), 'ergodic' is mentioned in Ch. 22 & 32 and resp. in Ch. 8 & 14, however with no relation to the validation of ABMs.

¹⁰SOLTYK and CHAN 2021, p. 19; FRANKE and WESTERHOFF 2016, p. 25, footnote 36; LEBARON 2021, p. 91, footnote 8.

simpler than the FABMs we study here.¹¹ We analyse broken ergodicity in the SMM estimation of two well established FABMs, the ALFARANO et al. (2008) (ALW model) and the FRANKE and WESTERHOFF (2012) (FW model), which reproduce a much larger set of stylized facts of financial markets. We perform Monte Carlo simulations and use the simulated method of moments (SMM) approach for the estimation of both ABMs.¹²

Second, we use the terminology of time averages and ensemble averages which allows to express many challenges in the estimation process as depending on taking different types of averages. For instance, we explicitly analyse the convergence behaviour over time vis-a-vis the convergence behaviour over replications. Doing so, we show how to reduce the uncertainty coming from broken ergodicity by analysing the estimation for different choices of ensemble size and simulation time length under limited computational resources, for which most estimators remain in a pre-asymptotic regime.

A brief survey of related work puts our contribution in context of the existing literature. Our work is closely related to a series of publications initiated by GRAZZINI and RICHIARDI who in GRAZZINI (2012) and GRAZZINI and RICHIARDI (2015) introduced a useful terminology and also use the SMM approach. However, they analyse the ergodic properties of much simpler models than those FABMs that we study here. As stated by GRAZZINI and RICHIARDI (2015, p. 154), ergodicity assures consistent estimates even in the adjustment phase, *i.e.* when the system has not yet fully reached the limiting distribution, because the properties are invariant across different realizations of the stochastic process. This assumes that certain regularities are ‘stable enough across different replications of the model’ to be identified. In general, GRAZZINI (2012) and GRAZZINI and RICHIARDI (2015) use non-parametric statistical tests to test for ergodicity (and stationarity). GRAZZINI (2012) applies a Wald-Wolfowitz runs test to check whether the observed and simulated data come from the same population. This can also be done using a Kolmogorov-Smirnov test.¹³ These tests can offer a first starting point. However, DOSI et al. (2018) find conflicting results using these econometric tests and thus conclude that they are not reliable to detect (non-)ergodicity. Another approach based on robust statistics is given by WINKER, GILLI et al. (2007) who study situations when a typical distribution or time series is better represented by the median instead of the mean. The divergence of the median from the mean is a typical symptom of

¹¹GRAZZINI 2011; GRAZZINI and RICHIARDI 2013, 2015.

¹²Standard references for the SMM are MCFADDEN (1989), PAKES and POLLARD (1989), LEE and INGRAM (1991) and DUFFIE and SINGLETON (1993). A textbook treatment is in DAVIDSON and MACKINNON (2003, Chap. 9.6).

¹³For ergodic Markov processes, the tests by DOMOWITZ and EL-GAMAL (1993, 2001) are applicable.

broken ergodicity as the mean or ensemble average is prone to outliers and because of that tracks the *exceptional* individual whereas the median is resistant to extremes and approximates the time-averaged behaviour and thus the behavior of *typical* individuals.

Furthermore, sufficient asymptotic convergence might be beyond reach for all practical purposes. BOUCHAUD and FARMER (2021) recently termed these situations *quasi non-ergodicity*, i.e. convergence times are astronomically-long such that for most practical purposes the estimation takes place in a pre-asymptotic regime where ergodicity is notoriously hard to achieve. Similarly DESSERTAINE et al. (2022) draw another analogy from physics and refer to non-ergodic effects of estimation in pre-asymptopia as violations of the *adiabatic assumption*, ‘in the sense that the time needed for the system to reach equilibrium is much shorter than the time over which the environment changes, so out-of-equilibrium effects can be neglected’. The model data-generating process (mDGP) of a (F)ABM is usually not analytically solvable let alone available in closed-form, which makes it even harder to choose the number of simulation runs and length without any prior knowledge about the speed of convergence. Our study helps to reduce this uncertainty.

The remainder of the paper is organised as follows. In Section 2 we provide a basic understanding of ergodicity for the purpose of moment selection. It turns out that the simulated method of moments like any estimation approach closely relies on ergodicity somewhere along the process. Section 2 shows where exactly ergodicity enters. In Section 3 we investigate how (broken) ergodicity influences the validation of FABMs. In doing so, we run a bunch of numerical experiments on two financial market models and present the results of our simulations. Section 4 contains a summary and concludes the paper.

2 Ergodicity and moment selection

This section starts with a brief discussion of ergodicity for the purpose to guide the selection of proper moments. In short, proper moments must be ergodic to be invariants of the model and not artefacts of particular simulated trajectories. With the help of an example of (geometric) Brownian motion we show that it is possible and necessary to derive ergodic observables from non-ergodic processes. These insights about ergodicity then guide the process of moment selection in Subsec. 2.3. The moments enter the method of simulated moments which is briefly described in Subsec. 2.5.

2.1 Ergodicity

Ergodic theory can be understood as a mathematical theory that studies the asymptotic behaviour of averages. Results from this mathematical field become relevant in an economics context whenever an economic observable is modelled as a stochastic process. In our context of financial markets, the asset price as our economic observable of interest is modelled as a random variable in a static context or as a stochastic process if we are interested in the time evolution as well. Figure 1 provides an intuitive understanding of the core aspects of ergodicity and shows five sample time series.




Figure 1: **Two Averaging Procedures.** Averaging over all realizations at a fixed time, *e.g.* at t^* , yields the ensemble average at that time (blue shade): $\langle X_{\bullet, t^*} \rangle = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N X_{x_n, t^*}$. Averaging over time of a single realization of the stochastic processes, *e.g.* the sample path associated with initial condition x_3 , yields the time average (red shade): $\bar{X}_{x_3, \bullet} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} X_{x_3, t}$ (adapted from KIRSTEIN 2019, p. 68).

From a modelling standpoint we think of the five time series as five realizations of a

stochastic process $\{X_t\}_{t \in \mathbb{T}}$ with five different initial conditions x_1, \dots, x_5 .¹⁴ For FABMs, the raw observables will be time series of asset prices or returns initialized by the series' first observations. We think of the observed data as being generated by some unknown real-world data-generating stochastic process (rwDGP), which remains a purely mental construction, of course. The FABM is the model data-generating process (mDGP), from which we are able to sample. Validation and estimation in econometrics can then be understood as a test of the match between the mDGP and the rwDGP with respect to some metric. In SMM, this metric consists of a vector of moments. The selection of moments is discussed in Subsec. 2.3.

With the help of ergodicity and Fig. 1, we distinguish between two types of the physical identity of an average, ensemble average (blue shade) and time average (red shade). Ensemble and time averages perform different aggregations on different instances of the observables. A time average computes an average along a single time series (or a single realization of a stochastic process) which is depicted in Fig. 1 by the red shade tracking the time series $X_t(x_3)$. An ensemble average computes an average over different realizations of the random variable (or stochastic process) at a single point in time which is depicted in Fig. 1 by the blue shade covering realizations at time t^* . In econometric terminology time averages correspond to longitudinal averages and ensemble averages to cross-sectional averages. Intuitively, a time average gives an aggregate statement about a single system over time, an ensemble average gives an aggregate statement about many systems at an instance in time.¹⁵

Ergodicity is then a mathematical property of an observable that describes the conditions under which both averages exist and coincide. A common definition of ergodicity establishes some type of convergence of longitudinal or time series averages on the LHS towards the cross-sectional average on the RHS:

$$\underbrace{\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T f(x_t)}_{\text{average along a single infinite time series}} \xrightarrow{\text{a. s.}} \underbrace{\int_{\Omega} f(x) P(x) dx}_{\text{average over population}} = \langle X \rangle, \quad (2.1)$$

where f is an arbitrary function of the observables. This handy notation allows to cover the case of analysing the raw observables as well as some transformation of them, in the former case f is simply the identity mapping, $f(x) = x$. In our context, the function

¹⁴For simplicity we will omit the initial conditions in the notation if they are not relevant.

¹⁵A key insight that has recently become better understood is to distinguish not only between the physical identity of an average (ensemble/time) but to further abstract and distinguish the physical identity from the mathematical identity or mathematical functional form of an average (arithmetic/geometric/harmonic/etc.) (KIRSTEIN 2019).

f potentially represents some summary statistic of the data, *i.e.* a moment function which we discuss below in more detail. The LHS of Eq. (2.1) yields a time average in the long-time limit. It is important to note that the RHS is also the result of a limit, it yields the ensemble average in the large-ensemble limit. The latter is rarely made explicit and instead the RHS is often presented just as *the* expectation value, *i.e.* $\langle X \rangle = E[X]$.

2.2 Identifying ergodic observables

In this section we briefly discuss the possibility to extract ergodic observables from non-ergodic processes. This corresponds to extracting moment functions from some rwDGP of which we don't need to know whether it is ergodic or not as long as we select an ergodic moment function. The goal is to identify invariants of a model which is a key goal in science.¹⁶ Invariants are adequate quantities when describing and comparing models of real-world phenomena. Identifying these invariants and package them in moment functions will be instrumental for the selection of proper moments discussed in the subsequent section.

Ergodicity as a property of mathematical models characterises such possible invariants in the specific sense of Eq. (2.1), *i.e.* an ergodic observable converges (almost always) towards a unique invariant limit distribution.¹⁷ Most models of real-world phenomena are *per se* non-ergodic processes, because some dynamic aspect (like growth) is relevant.¹⁸ For illustration purposes let us look at the standard model of asset price evolution which is that of (geometric) Brownian motion (GBM).

As such the model of GBM is clearly non-ergodic. Yet it is possible to derive ergodic observables from this non-ergodic process. Quite generally, a stochastic process is a model of an uncertain economic observable that can vary in time and over realizations. A simple stochastic process is the Wiener process W_t . In finance the Wiener process is used as a model of the price evolution, *i.e.* the repeated measurements of the price of an asset.¹⁹ The Wiener process W_t is non-ergodic, as its time average behaves as $\overline{W_t} \sim \mathcal{N}(0, t/3)$ for $t \rightarrow \infty$, and thus does not converge, because the distribution has a time-dependent variance, *i.e.* the distribution is broadening over time. Therefore, the *time-dependent* time average cannot almost always be equal to its *time-independent*

¹⁶PETERSEN 1996.

¹⁷In a sense, an ergodic observable ensures that more data is better for estimation and validation purposes, because more observations are likely to lead to better convergence.

¹⁸Let us emphasise that ergodicity is not a property of reality but of the mathematical model.

¹⁹The Wiener process is sometimes also called *Brownian motion* because it was initially conceived to model the real-world phenomenon of Brownian motion, *i.e.* the jittery motion of pollen suspended in fluids.

ensemble average $\langle W_t \rangle = 0$, if so then only occasionally (by chance).

Again, we are interested in identifying ergodic observables or summary statistics of a process. In the unlikely case that the raw observable is ergodic we are lucky. More generally some transformations might be necessary to arrive at an ergodic observable. Typical derived observables are increments, squared increments (sometimes referred to as squared displacements), correlations, multiples, returns or growth rates. If we analyse the squared displacement from the origin of a Wiener process as our observable of interest, $D_T^2 = (W_T - W_0)^2$, then this observable is ergodic, $\langle D_T^2 \rangle = \overline{D_T^2}$.²⁰ The increments of a Brownian motion converge towards a unique invariant time-independent distribution, and are in this sense an ergodic observable. The increments of a geometric Brownian motion, however, are not converging but remain level-dependent. For GBM it is the exponential growth rate that converges to an ergodic invariant distribution instead. This little example is meant to show that it is possible to derive ergodic observables from non-ergodic processes. In fact, we use the ergodic observables to characterize non-ergodic processes and thus in a similar way like in this little example use ergodic moments to characterize the mDGP of a (F)ABM.

It is the ergodic property that makes a summary statistic a suitable moment for the validation of (F)ABMs. However, the high degrees of freedom in (F)ABMs often prevent the existence of a closed-form solution, which would simplify the direct analysis of the stochastic data-generating process and to analytically derive the limit distributions of observables. Therefore, we have to rely on numerical simulations to study the convergence behavior and ergodic properties of moment functions. With this in mind, we now turn to the process of moment selection.

2.3 Selection of moment functions

The term *moment* is central in all moment-based estimations as the generalized method of moments (GMM) and the simulated method of moments (SMM) which we use in this paper. However, the word *moment* has different meanings in different fields. Let us therefore briefly clarify how we use it here. Whenever we refer to *moments*, we use this term in the narrow sense of their ordinary meaning in stochastics as properties of a probability distribution \mathcal{P} in Eq. (2.2). In the context of SMM estimation of (F)ABMs, the literature refers to moments in a wider sense as some summary statistic. To avoid confusion we will thus refer to them as *moment functions* in this paper.

²⁰METZLER et al. 2014.

2.3.1 Moments

In stochastics a moment m of order $q = 0, 1, 2, 3, \dots$; $q \in \mathbb{N}$ of a distribution is computed as the mean of the q -th power of its realizations, $m_q = \int_{-\infty}^{\infty} x^q P(x) dx = E[X^q]$. It is often convenient and therefore desirable to work with probability density functions such as P , but the density of a measure doesn't always exist, thus we can express moments using only the probability measure \mathcal{P} instead,

$$m_q = E[X^q] = \int_{-\infty}^{\infty} x^q d\mathcal{P}(x). \quad (2.2)$$

As the mDGP for most FABMs is not readily available in closed form, the moment functions can't simply be calculated analytically. Thus we need to simulate the data first and then compute the moments numerically. Inference from a given sequence of moments to a distribution is a classical problem in mathematics known as the moment problem. The estimation task of a (F)ABM is an inverse form of the classical moment problem. Generally, moment problems are only solvable if long sequences of moments are available, which is usually not the case for distributions of observables of financial markets because they turn out to be heavy-tailed.²¹ Higher-order moments of heavy-tailed distributions (HTDs) quickly vanish and in extreme cases even the first moment doesn't exist.²² It seems therefore unlikely that solutions to classical moment problems are of immediate help for our estimation task. However, it provides an additional motivation to extend the selection of moment functions in SMM beyond the ordinary moments.²³

Empirical analyses of economic and financial data at least since the 1960s²⁴ have repeatedly confirmed that many observables are not very well described by the family of GAUSSIAN distributions but instead follow HTDs, often with a power law tail. Since the increasing magnitude of an extreme event raised to the q -th power, x^q , can't be offset by its decreasing probability, $d\mathcal{P}(x)$, the integral in Eq. (2.2) diverges and higher-order moments cease to exist for HTDs. For instance, the calculation of the second central moment, the variance, is numerically always possible but does not converge for larger samples. Associated with the (existence of the) variance is a typical scale of dispersion.

²¹A classic text on the moment problem is AKHIEZER (1965), more recent research reveals the solvability of the HAMBURGER moment problem depends on the non-negative extendability of the HANKEL matrix, see BOLOTNIKOV (1996, 2008), CHEN and HU (1998) and DYUKAREV et al. (2008).

²²RACHEV 2003.

²³SOLTYK and CHAN (2021) use solutions to classical moment problems for modelling time-varying higher-order conditional moments of financial time series.

²⁴MANDELBROT (1960, 1963) and FAMA (1963) are prominent references that find heavy-tailed price increments (returns).

If the distribution has no second moment, then such the variations are so extreme that they lack a typical scale. Obviously, non-existing moments would not be good candidates for moment functions. An invariant of such volatile behavior might be the scaling exponent of the increments or returns. In fact, it is now understood that power laws are invariants of scaling relationships which again are imprints of complex systems and their otherwise out-of-equilibrium dynamics.

2.3.2 Moment functions

Moments ‘provide a natural source of information about the parameters, other functions of the data may also be useful’ (GREENE 2018, p. 491). A natural motivation to move beyond moments to moment functions is to not only use static properties of distributions but also dynamic properties of the underlying processes, *e.g.* correlations are not a property of a distribution but of observations at different times, locations or both. We refer to these ‘other functions of the data’ as *moment functions*. Our notation proves convenient here, as it is possible to interpret the function f in Eq. (2.1) simply as some moment function $f(x) = m(x)$. Eventually, our vector of moments \mathbf{m} can contain ordinary moments in the sense of Eq. (2.2) as well as moment functions.

A hierarchy of convergences in ergodic theorems In order to emphasise how strong a statement about an ergodic property actually is, it is worthwhile to make explicit a hierarchy of convergences contained in ergodic theorems like Eq. (2.1). For an ergodic theorem to hold:

1. almost all individual time series averages need to converge at all,
2. almost all individual time series averages need to converge to some (arbitrary) identical value,
3. and this value must be a particular one, namely coincide with the expectation value.

Let us comment on this hierarchy. On the first two points, most raw time series do not converge at all to a single value and are thus not stationary in this sense, *e.g.* all growth processes, security prices, business cycle dynamics or exchange rates. Some transformation like first or higher differencing is necessary to impose stationarity and thus convergence.²⁵ However, such transformations do not necessarily impose

²⁵DELLI GATTI, FAGIOLO et al. 2018, p. 145.

ergodicity, as they operate only on a single time series. On the third point, the expectation value is a particular average, *i.e.* the probability-weighted arithmetic mean of all realizations of a random variable. Thus its mathematical identity is an arithmetic mean and its physical identity is an ensemble average. Note that in general both time and ensemble averages do not coincide and hence non-ergodicity or broken ergodicity is our default condition. In fact, ergodicity only holds for very special cases or carefully chosen transformations of the observables.

This convergence hierarchy also applies to moment functions. Already FRANKE (2009, p. 805) and RUGE-MURCIA (2012, p. 918) note that many empirical moment functions are computed as time averages of some function of the (time series) data, $1/T \sum_{t=1}^T m(x_t)$. We only take the obvious next step to demand the ergodicity of moment functions in the selection process such that they converge towards a unique value or distribution and capture invariants of the mDGP. So far it is only tacitly assumed in the literature that moment functions are ergodic but rarely explicitly analyzed, which implies that the calculation for (almost) all time series samples converges (almost) always towards a particular and identical value (namely the expectation value). For FABMs that offer no closed-form analytical solution, the convergence of moment functions can only be assessed numerically which will be done in Subsec. 3.2.2. Thus, the hierarchy of convergences is directly informative in the context of the validation of ABMs. Firstly, the time averages of the moment functions need to converge at all and additionally to some value, $1/T \sum_{t=1}^T m(x_t) \rightarrow \bar{m}$. Corresponding to the third aspect in the hierarchy, this convergence needs to go towards the expected value we observed empirically, $\bar{m} = 1/T \sum_{t=1}^T m(x_t) \rightarrow \langle m \rangle$. Eventually, proper moment functions have to rely be ergodic invariants of the mDGP.

Ergodic theory implies a further highly relevant aspect to the validation of FABMs which is its asymptotic nature. Roughly speaking, ergodic theorems live in the land of asymptopia, while reality and the estimation of (F)ABMs take place in pre-asymptopia with finite samples. For all practical purposes real-world data will always be limited even if the available computational resources are constantly expanding. If there are only finite observations in the time series of the FABMs in pre-asymptopia then the time scales of the convergences will become crucially important sooner or later. Any uncertainty about the LHS in Eq. (2.1) vanishes only in the large-time limit.²⁶ In pre-asymptopia all quantities are computed under the constraint of finite computational resources. Thus we operate with samples finite in time and ensemble size, $\lim_{t \rightarrow T}, T \ll$

²⁶But there is also some uncertainty about the true expectation value on the RHS in Eq. (2.1), which only vanishes in the large-ensemble limit.

∞ and $\lim_{n \rightarrow N} N \ll \infty$. Given any limited budget of computational resources, the uncertainties in computing moment functions of our FABM depend on such limits and behave differently and might never vanish sufficiently fast.

2.4 Monte Carlo cube

So far we have discussed (i) ergodicity in general, (ii) the importance of the ergodic property, (iii) assumptions hidden in the ergodic theorem about specific convergences that reappear in the selection of suitable moments of FABMs and (iv) the fact that validation of FABMs always takes place in pre-asymptopia where ergodicity is generically broken due to slow convergences. In this section we focus on the different dimensions where convergences take place when performing Monte Carlo simulations and how to visualize them. This section prepares for a better understanding of the effects of broken ergodicity along the different dimensions as discussed in Sec. 3. It will turn out to be instructive for the choice of how to split a limited number of total observations over (i) the simulation length, (ii) the number of simulation runs and (iii) the number of Monte Carlo repetitions to improve validation efforts.

We model observables as realizations of a random variable X .²⁷ Repeated observations or realizations x of the same random variable at regular intervals from time 0 to time T form a time series of our observable, $\{x_0, x_1, x_2, \dots, x_T\}$, also denoted by $\{x_t\}_0^T$. Within this stochastic model a specific observed time series can therefore be interpreted as a realization of finite length of a stochastic process $\{x_t\}_{t \in \mathbb{T}}$, where \mathbb{T} denotes the time domain of the time index.²⁸ A stochastic process, $\{x_t\}_{t \in \mathbb{T}}$, is then a family of all infinite time series.

Our notation establishes the following relation between the real world (which is observable) and our model. Nature or the financial markets are thought of as a rwDGP which is unknown and of which we often observe only exactly one unique time series. On the other hand we have an economic model of a financial market – a FABM in our case – which is our model data-generating stochastic process (mDGP), from which we can generate an ensemble of (at least in principle) arbitrarily many N time series of (finite) length T . At least two facts distinguish two time series generated by the same stochastic process. First, the two time series (or realizations of the stochastic process)

²⁷Throughout this paper we denote random variables by capital letters and their realizations by the respective lower case.

²⁸The flexible notation of the time domain as \mathbb{T} allows to easily adopt different time domains. For a continuous time model the time parameter t of the stochastic process is $\mathbb{T} = \mathbb{R}$ or $\{x_t\}_{t \in \mathbb{R}}$. For a model in discrete time $\mathbb{T} = \mathbb{N}$ which yields the stochastic process $\{x_t\}_{t \in \mathbb{N}}$.

differ in their random initial conditions or random seeds.²⁹ The dependence on some initial conditions is indicated in Fig. 1 by the argument in the function, *e.g.* $X_t(x_1)$ signifies a process X_t that started in x_1 . Second, randomness in the noise realizations – there are simply different realizations of the random variable that appear over time or are generated and used in the simulation.

Throughout the paper, the distinction between the time dimension and the ensemble dimension will be crucial. We denote time by T and denote the size of the ensemble by N . A Monte Carlo simulation experiment is then fully determined by the

1. simulated time length T , which yields a $(1 \times T)$ -matrix or row vector of size T ;
2. an ensemble of N different generated time series, which yields a $(N \times T)$ -matrix;
3. number of M Monte Carlo realizations, which yields M different $(N \times T)$ -matrices.

Figure 2 contains a visualization of the three dimensions of simulations of (F)ABMs. Commonly, this third dimension of the Monte Carlo runs is understood as belonging to the ensemble dimension as only the random seeds might be different, but additional numerical effects appear that make a distinction between the two dimensions possible which are discussed in more detail in Sec. 3.³⁰ The convergence behavior in the numerical estimation crucially depends on the ensemble size N , the time length T and the number of Monte Carlo simulations M in a non-intuitive and non-linear way. In principle different convergence behaviours in these three dimensions have not been investigated before. Our analysis is thus contributing to the joint community efforts of improving the validation of (F)ABMs.

2.5 Simulated method of moments

Let us now turn to the simulated method of moment estimation. This estimation method is studied for different kind of models such as dynamic stochastic general equilibrium (DSGE) models³¹ as well as (F)ABMs.³² In general SMM is more robust to

²⁹DELLI GATTI, FAGIOLO et al. 2018, pp. 39–40.

³⁰Related two-dimensional visualisations in the context of tests for ergodicity can be found in GUERINI and MONETA (2017, Fig. 3, p. 132). GRAZZINI and RICHIARDI (2015, p. 155) refer to the three-dimensional mental model of a cube and refer to it as the ‘replications’ dimension. Similar reasoning about correct alignment of research question and statistical identification led to the idiographic paradigm in behavioral psychology and physiology and similar visualisations (MolenaarCampbell2009; NeumannEtAl2022).

³¹RUGE-MURCIA 2012, 2013.

³²FRANKE 2009; FAGIOLO and ROVENTINI 2017; FAGIOLO, GUERINI et al. 2019.

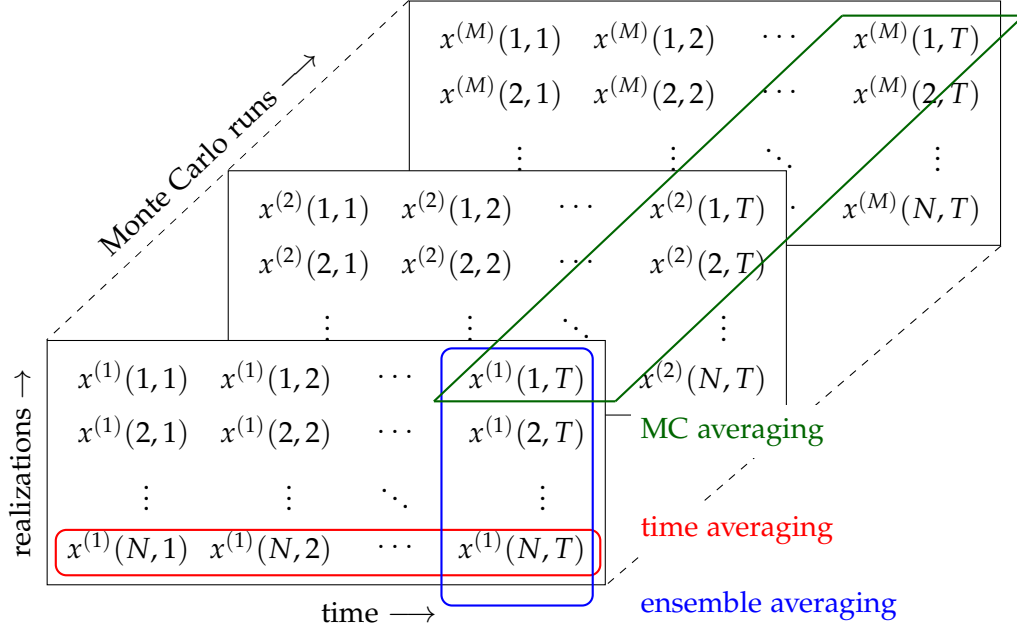


Figure 2: **Monte Carlo Cube.** Every column depicts an ensemble, *i.e.* a collection of N realizations of the random variable x at some instance of time. Every row of a matrix depicts a time series, *i.e.* realizations of the random variable x in a particular realization given by the first argument over time t which is denoted by the second argument. Every matrix of such rows and columns depicts the outcome of a single Monte Carlo run denoted by the superscript number in parenthesis. Three possibly different dimension along which averaging and thus convergence can take place, the ensemble dimension (blue), the time dimension (red) and a third dimension along different Monte Carlo runs (green).

misspecification than for example maximum likelihood methods³³ and performs better for large-scale models.³⁴ SMM belongs to the broader class of simulated minimum distance methods whose goal is to identify the unknown parameters of the model, which generate the least distance between the simulated moment functions and the empirical observed moment functions.

Let us now introduce a simulated method of moment estimator. Therefore, we consider a FABM with unknown parameter vector $\theta \in \Theta$, which we ultimately want to estimate from the parameter space Θ . In many empirical studies it is the primary goal to fit parameter values of the mDGP *e.g.* by matching empirical and simulated data moment functions. In our context, the goal is to assess the quality of the estimation approach itself, *i.e.* how well does the estimation approach identify some known ‘true’ parameter values θ_0 . If no empirical data is used for the evaluation of the estimation approach but data simulated from an ABM – like in our study – then it is a common practice to use known parameter values from the literature as the benchmark of the estimation that match the moment functions well and refer to them as ‘true’ values.³⁵ Since SMM belongs to the broader class of simulated minimum distance methods, a distance function d measures the difference between empirical and simulated data given some parameters θ .³⁶ Thus, we compute the distance between the two moment function vectors for the simulated and empirical moment functions $d = m^{\text{emp}}(x_t) - m^{\text{sim}}(x_t|\theta)$ given some parameters θ . In the following, we explicitly consider two different types of averaging over simulated observables. First, simulated moment functions can be computed as time averages over one realization of observables x_t of time length T :

$$m^{\text{sim}} = \frac{1}{T} \sum_{t=1}^T m(x_t) . \quad (2.3)$$

Or second, they can also be captured as averages over an ensemble of size N of time averages of length T :

$$m^{\text{sim}} = \frac{1}{NT} \sum_{n=1}^N \sum_{t=1}^T m(x_t) . \quad (2.4)$$

The objective function or criterion function for a given a set of model parameters θ

³³RUGE-MURCIA 2007.

³⁴PLATT 2020.

³⁵They are sometimes also referred to as ‘pseudo-true’ values as they are derived from synthetic/simulated data. See also Subsubsec. 3.2.1.

³⁶To be very precise, the empirical data is produced given the true parameters θ_0 , which we want the estimation to recover reliably.

aggregates then the distances

$$J(\boldsymbol{\theta}) = \mathbf{d}(\boldsymbol{\theta})' \mathbf{W} \mathbf{d}(\boldsymbol{\theta}) , \quad (2.5)$$

for a given weighting matrix \mathbf{W} which is positive semi-definite. The estimator $\hat{\boldsymbol{\theta}}$ yields that vector of model parameters for which the objective function is minimized:

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) . \quad (2.6)$$

Under standard regularity conditions, the distance function \mathbf{d} is assumed to be stationary and ergodic resulting in an asymptotically consistent estimator.³⁷ For most (F)ABMs, however, a corresponding SMM estimator might not have such properties. In fact, we will show that in pre-asymptopia the following inequality holds due to non-commutativity of the limits in ensemble size, time and also of the number of MC runs:

$$\frac{1}{NT} \sum_{n=1}^N \sum_{t=1}^T m(x_t) \neq \frac{1}{NT} \sum_{t=1}^T \sum_{n=1}^N m(x_t) . \quad (2.7)$$

Put simply, under broken ergodicity and/or in pre-asymptopia the order matters in which the limits are taken.³⁸ We thus study the effect of an efficient allocation of a limited budget of observations and how to get an estimator with improved properties.

The efficiency of the SMM estimator is affected by the design of the weighting matrix \mathbf{W} . The optimal choice is given by a weighting matrix with the smallest asymptotic covariance for the estimator. One popular choice would be the use of the Newey-West estimator. In our numerical study, we are interested in the uncertainty that can be solely associated with broken ergodicity. Therefore, we will consider the inverse of the long-run covariance matrix of the true data as the optimal weighting matrix.

To summarize this section, we have discussed how the ergodic property plays a crucial role in moment selection. As we will see next, convergences along the three different dimensions might behave differently due to broken ergodicity in pre-asymptopia. How exactly shows the evaluation of the simulation experiments in the next section.

3 Simulation experiments

To conduct our study of (broken) ergodicity in the convergence of moment functions, we take two established financial agent-based models: ALFARANO et al. (2008) (ALW

³⁷LEE and INGRAM 1991; DUFFIE and SINGLETON 1993.

³⁸GRAZZINI 2011, eq. 2 on p. 6.

model) and FRANKE and WESTERHOFF (2012) (FW model). They are both based on a herding mechanism that has its roots in KIRMAN (1993). Both models are commonly used for estimation exercises in the community.³⁹ They replicate most of the stylized facts mentioned above and require comparatively little computational resources despite their prototype-nature.

3.1 Agent-based financial market models

In this subsection we briefly discuss the key mechanisms underlying the ALW model and the FW model which are the basis of our numerical experiments hereafter.

3.1.1 ALW model

The first model by ALFARANO et al. (2008) incorporates a behavioural herding mechanism based on KIRMAN (1993) with precedent analysis in ALFARANO et al. (2005).⁴⁰

The ALW model assumes two types of financial speculators, fundamentalist and chartist traders. Fundamentalists' excess demand is given by $D_f = N_f V_f (p_t^* - p_t)$ where V_f is the average demand of N_f fundamental speculators. The fundamental value p_t^* is assumed to follow a random walk $p_t^* = p_{t-1}^* + \sigma_f \cdot \epsilon_{f,t}$ with $\epsilon_f \sim \mathcal{N}(0, 1)$. The excess demand of chartist traders is given by $D_c = N_c V_c x_t$. Chartists are in one of two opinion states, either optimistic or pessimistic. A sentiment index x_t is defined as $x_t = 2n_t/N_c - 1$ with n_t optimistic traders at time t of a total number of N_c chartists. They are assumed to change their sentiment based on the (extensive) transition rates $\pi^+ = a + bn$ and $\pi^- = a + b(N - n)$, where parameter a indicates idiosyncratic switches and b measures the herding intensity. The resulting sentiment dynamics is approximated by the following Langevin equation with drift component $A(x) = -2ax$ and diffusion term $D(x) = 2b(1 - x_t^2) + 4a/N$ which gives

$$dx_t = A(x) dt + D(x) dW_t = -2ax_t dt + \sqrt{2b(1 - x_t^2)} dW_t. \quad (3.1)$$

Eq. (3.1) can be discretized with $\Delta t = 1$:

$$\Delta x = x_{t+\Delta t} - x_t = -2ax_t + \sqrt{2b(1 - x_t^2)} \epsilon_t \quad (3.2)$$

³⁹CHEN and LUX 2018; LUX 2018; KUKACKA and KRISTOUFEK 2020; BERTSCHINGER and MOZZHORIN 2021.

⁴⁰For a simple variant of the herding model by KIRMAN (1993) with local interaction between the agents (or equivalently extensive transition rates) a closed-form exists for which analytical solutions of the time-variation of moments and some related quantities of interest exist can be computed (see ALFARANO et al. 2005 and esp. ALFARANO et al. 2008, Section 4).

with $\epsilon \sim \mathcal{N}(0, 1)$. The distribution of the sentiment index is known to be bimodal for $a < b$ and unimodal for $a > b$.

Price dynamics are governed by a standard Walrasian adjustment mechanism, depending on total excess demand of both trading groups:

$$\begin{aligned} p_{t+1} &= p_t + \beta (D_{f,t} + D_{c,t}) \\ &= p_t + \beta \left[N_{f,t} V_{f,t} (p_t^* - p_t) + N_{c,t} V_{c,t} x_t \right] , \end{aligned} \quad (3.3)$$

where β is the assumed price adjustment speed. With instantaneous market clearing, *i.e.* $\beta \rightarrow \infty$, and setting $N_c V_c / N_f V_f = 1$, we get the following evolution of returns:

$$\begin{aligned} r_{t+1} &= p_{t+1} - p_t \\ &= \sigma_f \epsilon_t + (x_t - x_{t-1}) . \end{aligned} \quad (3.4)$$

Finally, the parameter vector to be estimated for the ALW model contains three items $\theta^{\text{ALW}} = (a, b, \sigma_f)^\top$. In accordance with the literature we use for all our simulation experiments the following true model parameters $\theta_0 = (0.3, 1.4, 30)^\top$.⁴¹

3.1.2 FW model

FRANKE and WESTERHOFF (2012) propose an entire model zoo for which they run a model contest. Here, we will only consider their top performing DCA-HPM model version (discrete choice with herding, predisposition and price misalignment). FRANKE and WESTERHOFF (2012) apply a market maker model that considers as well two types of speculators: chartist and fundamentalists, whose fractions are denoted by n_c and n_f , respectively. The evolution of the log prices is then determined by

$$p_{t+1} = p_t + \mu (n_{f,t} D_{f,t} + n_{c,t} D_{c,t}) , \quad (3.5)$$

with μ reflecting the speed of price adjustment. Excess demand of chartists D_c and fundamentalists D_f is given by

$$\begin{aligned} D_{f,t} &= \phi (p^* - p_t) + \epsilon_{f,t} \\ D_{c,t} &= \chi (p_t - p_{t-1}) + \epsilon_{c,t} , \end{aligned} \quad (3.6)$$

⁴¹See GHONGHADZE and LUX (2016) and CHEN and LUX (2018). The (pseudo-)true parameter setting for the ALW model that is used in this paper represents the bimodal case of the underlying sentiment index given that $b > a$. Since CHEN and LUX (2018) have shown that the bimodal case gives the best fit to empirical data, we will only focus on this model scenario. Note that for better readability parameters of the ALW model are always multiplied by 10^3 throughout the paper.

with $\epsilon_f \sim \mathcal{N}(0, \sigma_f^2)$ and $\epsilon_c \sim \mathcal{N}(0, \sigma_c^2)$. Parameters χ and ϕ are always positive and indicate the strength of reaction. The switching between both trading strategies is governed by a fitness function measuring the attractiveness a_t of fundamentalism over chartism:

$$a_t = \underbrace{\alpha_n (n_{f,t} - n_{c,t})}_{\text{herding intensity}} + \alpha_0 + \underbrace{\alpha_p (p_t - p^*)^2}_{\text{misalignment}}, \quad (3.7)$$

where α_0 represents a constant idiosyncratic predisposition for one of the two trading strategies, α_n relates to herding intensity and α_p accounts for price misalignments from the fundamental value p^* . Note that while α_n and α_p are always strictly positive, parameter α_0 might be negative as well. The current market shares of fundamentalists $n_{f,t}$ and chartists $n_{c,t}$ at time t are then updated according to the following discrete choice approach:

$$\begin{aligned} n_{f,t} &= \frac{1}{1 + \exp(-\beta a_{t-1})} \\ n_{c,t} &= 1 - n_{f,t} \end{aligned} \quad (3.8)$$

As in FRANKE and WESTERHOFF (2012, p. 1199) we set the intensity of choice parameter to $\beta = 1$, the speed of price adjustment parameter to $\mu = 0.01$ and the fundamental price to $p^* = 0$. Finally, the parameter vector to be estimated for the FW model contains a total of seven parameters $\theta^{\text{FW}} = (\phi, \chi, \alpha_0, \alpha_n, \alpha_p, \sigma_f, \sigma_c)^\top$. In accordance with FRANKE and WESTERHOFF (2012) we use for all our simulation experiments the following true model parameters $\theta_0 = (0.12, 1.5, -0.336, 1.839, 19.671, 0.708, 2.147)^\top$.

To sum up, the ALW model by ALFARANO et al. (2008) is a small-scale model which can be reduced to only three estimation parameters $\theta^{\text{ALW}} = (a, b, \sigma_f)^\top$. The FW model by FRANKE and WESTERHOFF (2012) is more complex with a total of seven parameters $\theta^{\text{FW}} = (\phi, \chi, \alpha_0, \alpha_n, \alpha_p, \sigma_f, \sigma_c)^\top$. Sample simulation runs for both model dynamics can be found in appendix A.

3.2 Moment functions

The following subsection builds on Sec. 2 and explains the motivation behind the choice of our set of moment functions. As explained in Subsec. 3.1, we focus on two agent-based asset pricing models which replicate many of the stylized facts of financial markets. The vector of moment functions \mathbf{m} should contain a set of reasonable summary statistics that (partly) capture these stylized facts expressed as observables

and thus measurable statistical quantities. A necessary condition for the choice of moment functions is given by the order condition, *i.e.* for estimations with more than one model parameter the number of moment functions K needs to be greater or equal than the number of model parameters or the cardinality of the parameter vector θ . Thus, the order condition provides a lower bound for the number of moment conditions. Theoretically, the (full) rank condition is a sufficient condition for identification assuring that only moment functions without linear dependence are included. While the order condition is easy to meet, the rank condition is barely testable for most (F)ABMs given their non-linearity and the lack of analytical closed-form expressions.

Besides these restrictions the number and choice of moment conditions is not strictly limited which may render it arbitrary. However, this is not a weakness *per se*, because SMM is designed to capture complex patterns which likely escape any single moment condition. Thus, if empirical data show complex patterns that are hard to squeeze into a single metric, it can be necessary to wrap the patterns in more than a single moment condition. The stylized facts of financial markets show such complex patterns like bubbles and crashes, excess volatility, heavy-tailed return distributions, absence of autocorrelation in raw returns or slow decay in volatility. Especially the statistical pattern of slow non-linear decay in autocorrelation over many lags requires more than one moment condition. This justifies the comparatively large size of our vector of moment functions with $K = 18$ listed in Table 1. For the SMM approach this is a common size of the moment vector, see *e.g.* CHEN and LUX (2018, p. 722) who analyse the ALW model with up to 15 moment functions or RUGE-MURCIA (2012, p. 930) who uses 16 moment functions for the estimation of a macroeconomic DSGE model.

As the first two moment functions we use the mean of absolute returns and the variance of the return series. As mentioned in Sec. 2, ordinary moments in the sense of Eq. (2.2) for HTDs do only exist for orders lower than the tail index. For example, if a return series has a tail index of $\alpha = 3$, only the first and second moment exist. This means that sample estimates of the third or higher moments will never converge with increasing sample size. Hence, such moments might not be suitable as moment functions for estimation purposes. We yet consider the kurtosis of the return series as one moment function and additionally include the Hill tail index estimates at 2.5 % and 5.0 %. The autocorrelation of raw returns at lag 1 checks for the absence of serial correlation in returns. Finally, in order to capture the slow decay of volatility we include autocorrelations of absolute and squared returns for the lags 1, 5, 10, 25, 50 and 100.

Table 1 offers an extensive yet non-comprehensive list of possible summary statistics considered for the estimation of financial market models. We abstain from modelling

Table 1: **Moment functions.** We use a set of eighteen moment functions to cover the most important stylized facts. Therefore we include the returns' volatility defined as the mean value of absolute returns, unconditional second and fourth return moments, Hill estimators of the power law tail index for the top 2.5% and top 5%, $\alpha_{2.5}$ and $\alpha_{5.0}$, the first order autocorrelation coefficient of raw returns and the autocorrelation coefficients of absolute and squared returns for lags 1, 5, 10, 25, 50 and 100.

Moment function	Notation	Statistic
expectation value	m_1	$E[r_t]$
variance	m_2	$E[r_t^2]$
kurtosis	m_3	$E[r_t^4]$
power law tail exponent top 2.5%	m_4	$\alpha_{2.5}$
power law tail exponent top 5%	m_5	$\alpha_{5.0}$
AC raw returns lag 1	m_6	$E[r_t r_{t-1}]$
AC absolute returns lag 1	m_7	$E[r_t r_{t-1}]$
AC squared returns lag 1	m_8	$E[r_t^2 r_{t-1}^2]$
AC absolute returns lag 5	m_9	$E[r_t r_{t-5}]$
AC squared returns lag 5	m_{10}	$E[r_t^2 r_{t-5}^2]$
AC absolute returns lag 10	m_{11}	$E[r_t r_{t-10}]$
AC squared returns lag 10	m_{12}	$E[r_t^2 r_{t-10}^2]$
AC absolute returns lag 25	m_{13}	$E[r_t r_{t-25}]$
AC squared returns lag 25	m_{14}	$E[r_t^2 r_{t-25}^2]$
AC absolute returns lag 50	m_{15}	$E[r_t r_{t-50}]$
AC squared returns lag 50	m_{16}	$E[r_t^2 r_{t-50}^2]$
AC absolute returns lag 100	m_{17}	$E[r_t r_{t-100}]$
AC squared returns lag 100	m_{18}	$E[r_t^2 r_{t-100}^2]$

moment functions as further derived processes, like a GARCH(1,1), firstly because of the additional computational cost associated with the estimation of the GARCH parameters. Secondly and more importantly, they introduce additional uncertainty leading to possible biases in the estimates.

Table 2: True (theoretical) values of moment functions. The table lists mean value, variance and p-value of Kolmogorov-Smirnov statistic for the ALW and FW model for all 18 moment functions listed in Tab. 1 over $T_1 = 10000$, $T_2 = 100000$, $T_3 = 1000000$ and $T_4 = 1000000$. Simulations are run over $M = 5000$ repetitions.

	m_1		m_2		m_3		m_4		m_5		m_6		m_7		m_8		m_9											
	mean	var	p-value	mean	var	p-value	mean	var	p-value	mean	var	p-value	mean	var	p-value	mean	var	p-value										
ALW	T1	0.0332	0.0	0.0365	0.0019	0.0	0.0153	0.8199	0.0184	0.0038	4.7256	0.0543	0.0083	-0.0	0.0001	0.7456	0.0993	0.0002	0.8175	0.0953	0.0002	0.6847	0.0973	0.0002	0.7464			
	T2	0.0332	0.0	0.668	0.0019	0.0	0.6258	0.8529	0.0019	0.7856	5.6635	0.0159	0.7591	4.6839	0.0057	0.7714	-0.0002	0.0	0.467	0.1034	0.0	0.9827	0.0983	0.0	0.3522			
	T3	0.0332	0.0	0.4123	0.0019	0.0	0.7276	0.8563	0.0002	0.9226	5.6597	0.0016	0.8183	4.6804	0.0006	0.2509	-0.0002	0.0	0.9348	0.1038	0.0	0.8736	0.0986	0.0	0.979			
	T4	0.0332	0.0	0.7663	0.0019	0.0	0.6951	0.8563	0.0	0.6985	5.6595	0.0002	0.9216	4.6799	0.0001	0.7177	-0.0002	0.0	0.9835	0.1038	0.0	0.9149	0.0986	0.0	0.834			
ALW	T1	0.0935	0.0003	0.7791	0.0951	0.0002	0.5688	0.0915	0.0002	0.2882	0.0884	0.0002	0.9002	0.0851	0.0002	0.2923	0.0778	0.0002	0.8222	0.0749	0.0002	0.4867	0.0607	0.0002	0.9124	0.0586	0.0002	0.2626
	T2	0.0966	0.0	0.768	0.0991	0.0	0.9564	0.0944	0.0	0.5326	0.0924	0.0	0.9785	0.0882	0.0	0.9196	0.0824	0.0	0.6354	0.0787	0.0	0.8017	0.0657	0.0	0.5687			
	T3	0.0969	0.0	0.786	0.0996	0.0	0.9963	0.0947	0.0	0.9349	0.0929	0.0	0.547	0.0886	0.0	0.4198	0.083	0.0	0.8863	0.0793	0.0	0.4013	0.0662	0.0	0.9985			
	T4	0.0969	0.0	0.4051	0.0996	0.0	0.6847	0.0947	0.0	0.7574	0.093	0.0	0.947	0.0886	0.0	0.8556	0.083	0.0	0.9503	0.0793	0.0	0.7092	0.0662	0.0	0.7422			
FW	T1	0.0071	0.0	0.0584	0.0001	0.0	0.0249	2.3495	0.0781	0.0021	4.3731	0.1543	0.0001	3.5976	0.0401	0.0	0.0073	0.0002	0.2485	0.1881	0.0004	0.0	0.1796	0.0004	0.5184	0.1839	0.0004	0.0
	T2	0.0071	0.0	0.2857	0.0001	0.0	0.3108	2.4246	0.0074	0.9997	4.3516	0.0175	0.139	3.5312	0.0039	0.0056	0.0072	0.0	0.4614	0.1937	0.0	0.1256	0.1824	0.0	0.7128	0.1893	0.0	0.1071
	T3	0.0071	0.0	0.9359	0.0001	0.0	0.5381	2.4302	0.0007	0.6599	4.3502	0.0018	0.5474	3.5254	0.0004	0.8748	0.0072	0.0	0.6108	0.1942	0.0	0.1158	0.1826	0.0	0.8993	0.1898	0.0	0.7911
	T4	0.0071	0.0	0.9987	0.0001	0.0	0.9397	2.4314	0.0001	0.9388	4.3505	0.0002	0.4357	3.5249	0.0	0.9837	0.0072	0.0	0.5315	0.1943	0.0	0.6216	0.1826	0.0	0.9461	0.1899	0.0	0.4562
FW	T1	0.1742	0.0004	0.3067	0.1764	0.0004	0.0	0.1661	0.0004	0.0672	0.1497	0.0005	0.0	0.1388	0.0004	0.0853	0.1112	0.0005	0.0029	0.1013	0.0004	0.8188	0.0641	0.0005	0.8457	0.057	0.0005	0.5279
	T2	0.1766	0.0	0.4711	0.1819	0.0	0.1044	0.1689	0.0	0.9349	0.1553	0.0	0.1422	0.1423	0.0	0.7186	0.117	0.0	0.9478	0.1058	0.0	0.8696	0.0698	0.0001	0.3663	0.0619	0.0001	0.0873
	T3	0.177	0.0	0.6711	0.1824	0.0	0.3943	0.1691	0.0	0.9735	0.1559	0.0	0.7974	0.1428	0.0	0.9909	0.1176	0.0	0.788	0.1063	0.0	0.9977	0.0704	0.0	0.8684	0.0624	0.0	0.8808
	T4	0.177	0.0	0.9552	0.1825	0.0	0.9178	0.1692	0.0	0.9023	0.156	0.0	0.4667	0.1428	0.0	0.8939	0.1177	0.0	0.5447	0.1064	0.0	0.6114	0.0705	0.0	0.3633	0.0625	0.0	0.7818

3.2.1 True moment vector

In order to apply any moment estimator, it is important that simulated moments (asymptotically) equal their true values. This means that true values exist and convergence happens with an increasing number of observations (either over time T and/or the ensemble N). Otherwise, meaningful inference is not possible. Hence, the very first step of our analysis is to numerically check if the moment functions do exist, *i.e.* if they quickly converge to identical values for some finite number of observations. Since the mDGP for most (F)ABMs is not available in closed form, the ensemble average $\langle X \rangle$ is not readily at hand. Living in a pre-asymptotic numerical simulation, sufficient convergence to the theoretical average needs to take place rather quickly for finite observations. Recall that we talked about the conditions for the existence of higher moments in the previous section. To find the moments' true values, we run simulations over different time lengths T with $T_1 = 10\,000$, $T_2 = 100\,000$, $T_3 = 1\,000\,000$ and $T_4 = 10\,000\,000$. We run each scenario over $M = 5000$ replications. Results are presented in Tab. 2. For each moment function, we report mean, variance and p -value of a Kolmogorov-Smirnov test statistic.

We observe convergence behaviour for all moment functions for both the ALW model and FW model for an increased number of observations. In fact, the variance decreases consistently for all considered moment functions and models. Given that the variance is low for all considered moment functions, we can assume that the long-run estimates are informative for the estimation. Therefore, we will consider the estimates of the longest time length T_4 to compute the inverse of the long-run covariance matrix as the optimal weighting matrix in our SMM estimator. High p -values for the KS-test also suggest normally distributed samples of the moment estimates. Given that the true values exist and converge, we can move on to the next step which is to test if the chosen moment functions are suited for identification of the model parameters.

3.2.2 Pre-asymptotic properties of moment functions

In this subsection we are concerned with the pre-asymptotic properties of our moment functions and their identification power. We want to check if the mapping between model parameters and model output in terms of moment functions is unique. Only a one-to-one mapping works as a sufficient condition (*i.e.* full rank condition) for identification. For the one parameter case, the identification issue is quite obvious. A sufficient condition for identification would be a strictly monotone relationship between model parameter and moment function. For higher dimensional models, this is less

trivial. In fact, as stated earlier, the rank condition is barely directly testable for ABMs.

Nevertheless, running sensitivity analyses allows us to study how strongly moment functions react towards changes in model parameters.⁴² Therefore, we break down the model's dimensionality and study how sensitive our moment functions are to variations of the model parameters.

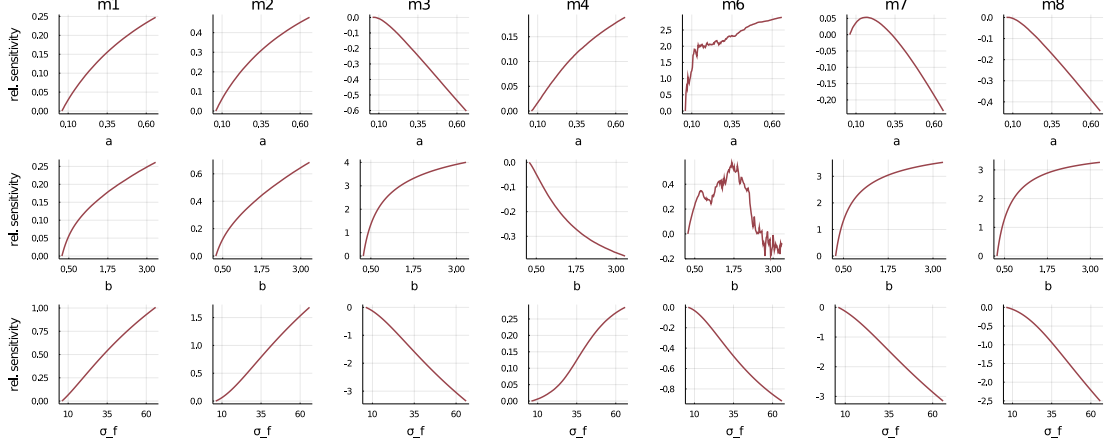


Figure 3: **Sensitivity of the moment functions to model parameters in the ALW model.** Model parameters are varied on horizontal axis, moment response is on the vertical axis. The top panel shows selected moment function responses for parameter a , the middle panel for parameter b and the bottom panel for parameter σ_f . Simulations are run over $M = 100$ Monte Carlo replications with $T = 20\,000$ and $N = 50$.

For the ALW model, we sample 151 equidistant points over a large range of parameter variations: $a \in [0.06, 0.66]$, $b \in [0.28, 3.28]$ and $\sigma_f \in [6.0, 66.0]$. We plot the relative responsiveness of moments m_1 , m_2 , m_3 , m_4 , m_6 , m_7 and m_8 for the ALW model in Fig. 3.⁴³ We find non-linear relationships for all moment-parameter pairs. Moment function m_6 (autocorrelation of raw returns at lag 1) shows highly non-smooth behaviour towards changes in model parameters a and b of the ALW model. Regarding the strength of reactions, moment functions react most sensitive to parameter σ_f . This is in

⁴²Since not all moment functions are equally informative, another way to identify relevant moment functions is in terms of statistical efficiency. GALANT and TAUCHEN (1996) suggest using scores of an auxiliary model as moment functions in a generalized method of moment (GMM) estimator in order to reduce loss of efficiency due to uninformative moments. Hence, by searching for the statistically most informative moments, such an approach tries to reach the efficiency of maximum-likelihood. Auxiliary models are considered to be especially attractive when the number of moment functions is rather limited.

⁴³Moment functions m_9 to m_{18} for the autocorrelations of absolute and squared returns at higher lags show qualitatively the same functional behaviour as m_7 and m_8 , respectively.

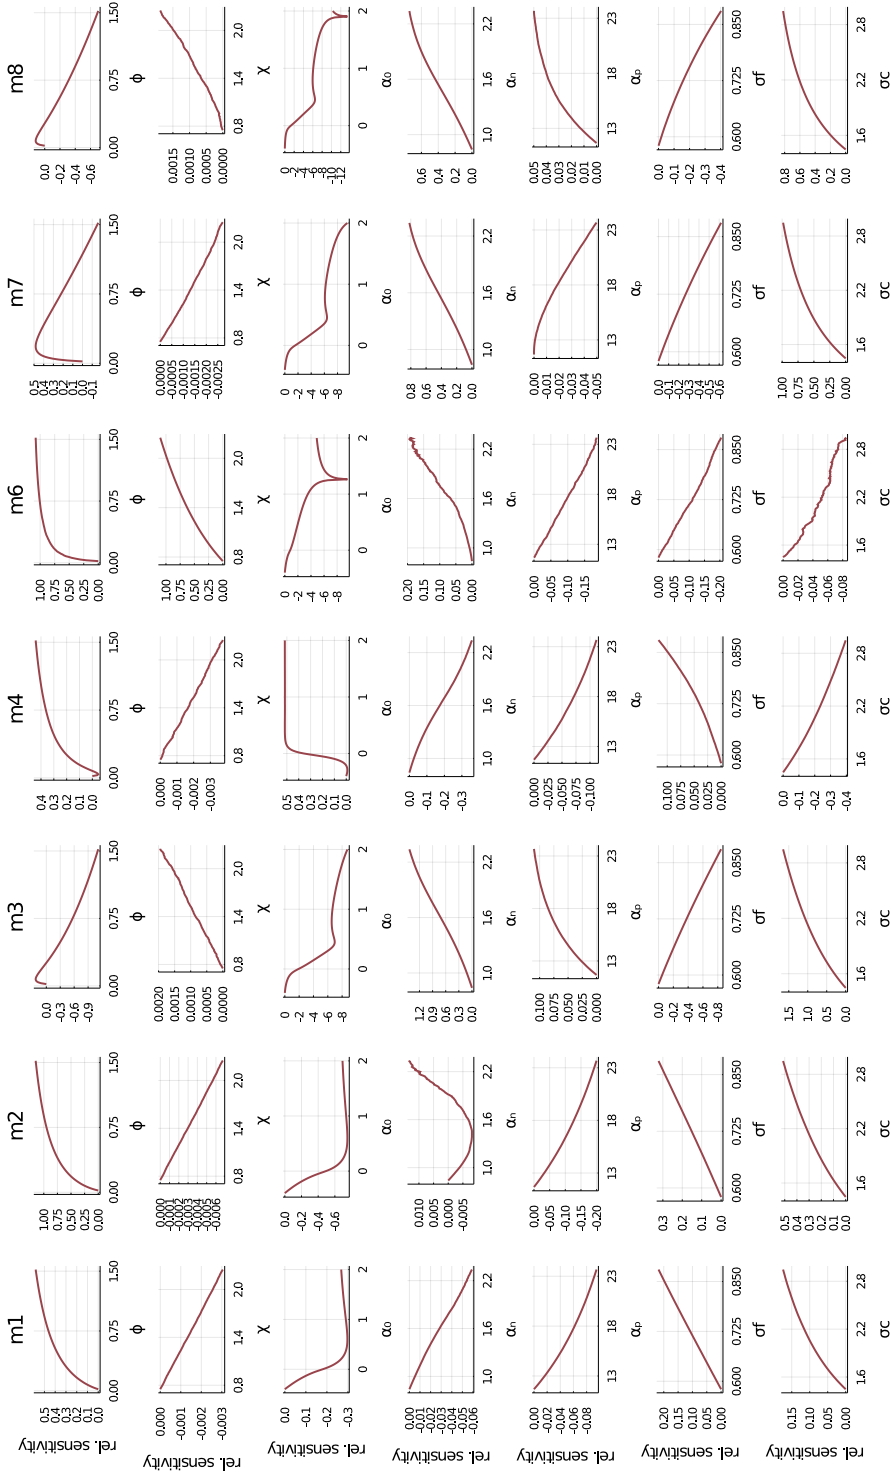


Figure 4: **Sensitivity of the moments to model parameters in the FW model.** Model parameters are varied on horizontal axis, moment response is on the vertical axis. The panels show selected moment responses for the seven parameters of the FW model: $\phi, \chi, \alpha_0, \alpha_n, \alpha_p, \sigma_f$ and σ_c . Simulations are run over $M = 100$ Monte Carlo replications with $T = 20000$ and $N = 50$.

line with earlier findings by CHEN and LUX (2018). In fact, the moment functions are least sensitive towards varying values of parameter a .

For the FW model, we also run simulation for 151 equidistant points over a broad range of parameter variations: $\phi \in [0.02, 1.52]$, $\chi \in [0.75, 2.25]$, $\alpha_0 \in [-0.4, 2.0]$, $\alpha_n \in [0.839, 2.339]$, $\alpha_p \in [11.671, 23.671]$, $\sigma_f \in [0.58, 0.88]$ and $\sigma_c \in [1.447, 2.947]$. Results are shown in Fig. 4. We find again highly non-linear functional relationships for almost all moment-parameter combinations. We further see parameter α_0 as probably hard to identify correctly given the wild and untamed nature of the moment function responses. Additionally, moment functions react comparatively insensitive towards parameter χ suggesting that it might trigger identification problems, too.

To conclude, sensitivity analyses reveal that most functional relationships are non-linear which might result in (small) biases of the estimates. As mentioned in GRAZZINI, RICHIARDI and SELLA (2012) and GRAZZINI and RICHIARDI (2015), if a moment function happens to be non-linear, there will be a small bias with the direction of the bias depending on moment functions' derivatives. Thereby concave (convex) moment functions lead to an upward (downward) bias. Such biases can be reduced in different ways: (i) given the analytical expression of the moment function, however, this is unknown for most ABMs, (ii) through monotonic transformations and (iii) through increased number of observations. The latter brings us to the focus of our study.

Recall that biases decrease for increasing number of observations since the simulated moment functions converge to their true theoretical value. Hence, the next step is to take a closer look at the convergence speed of the moment functions. The goal here is twofold. First, we want to check if an ensemble N of simulation runs converges to the exact same value as the corresponding long-run realization, *i.e.* how strong the effect of broken ergodicity is. Second, we aim at evaluating the speed of convergence. For this experiment, we run simulations over $M = 5000$ repetitions for two different scenarios: (i) $T = 400\,000$, $N = 1$ and (ii) $T = 40\,000$, $N = 10$. In Fig. 5 we plot again the convergence behaviour for the moment functions $m_1, m_2, m_3, m_4, m_6, m_7$ and m_8 defined as relative deviations from their true values, *i.e.* their long-run estimates reported in Subsubsec. 3.2.1. The top and lower panels of Fig. 5a and Fig. 5b show results for scenario (i) and (ii), respectively. The left (right) panels refer to the selected moment function including m_6 (excluding m_6). We find for both the ALW model and FW model that moment functions m_1, m_2 and m_4 closely converge towards their true values, while m_3, m_7 and m_8 show more or less pronounced and persistent biases. Regarding the speed of convergence, we observe long transition phases for scenario (i) and comparatively fast adjustments for scenario (ii). Yet, deviations are more

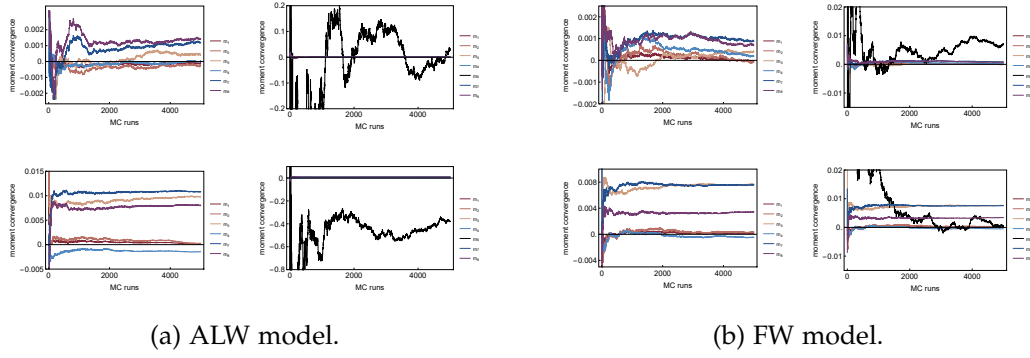


Figure 5: **Convergence behaviour for selected moment functions.** Plots reveal the relative deviations of moment functions from their true values for the ALW model in Fig. 5a and for the FW model in Fig. 5b. The top and lower panels show results for scenario (i) and (ii), respectively. The horizontal axis is representing the number of Monte Carlo runs ($M = 5000$).

severe here, too. Looking at the right panels including moment function m_6 which is the autocorrelation of raw returns at lag 1, we find very volatile and probably non-convergent behaviour for both models. This suggests that moment m_6 might not be suited for identification of the given model outcomes.

3.3 Properties of the objective function

Next, we study the properties of the SMM's objective function with explicit regard to the impact of broken ergodicity. Let us start with investigating the objective function's response surface. The aim is to check for discontinuities, (non-)smoothness and flat valleys, which create problems during the optimization as they impede finding global optima during the estimation. Such plots provide further help in gaining an understanding of the mapping between model parameters and the objective function.

Table 3: **Simulation Scenarios.** Allocation of limited budget of observations of $T \times N \times M = 80\,000\,000$ observations over different simulation dimensions.

Scenario	Time only		$T \times N$ mix			
	T_{short}	T_{long} (a)	(b)	(c)	(d)	(e)
Time Length T	40k	400k	40k	20k	10k	40k
Ensemble Size N	1	1	10	20	40	20
Monte Carlo Runs M	200	200	200	200	200	200

3.3.1 Response surface

For this analysis, we will focus only on the most problematic model parameters in terms of identification power. For the ALW model, we have identified both herding parameters a and b as troublesome. We run simulation scenarios summarized in Tab. 3 with a constant observation budget of 8 000 000 throughout all experiments, *i.e.* one short time series realization $T_{\text{short}} = 40\,000$, a long time series with $T_{\text{long}} = 400\,000$ and over an ensemble with $N = 10$ and $T = 40\,000$. The parameter grid consists of 41 equidistant points between $[0.06, 0.66]$ for parameter a and between $[0.28, 3.1]$ for parameter b . The moment conditions are based on all 18 moment functions of Tab. 1. Figure 6 shows the corresponding results for three different sample simulation runs using different random seeds. The solid lines mark the true value of the model parameters $(a, b) = (0.0003, 0.0014)$. The intersection of the two lines build the theoretical optimum of the objective function. Note that we plot the response surface of the inverse of the objective function to better visualize the global optimum. Ideally, the highest peak on the response surface coincides with the intersection of the true parameter values.

The response surface of the objective function is clearly non-smooth and contains multiple local optima. Depending on the choice of random seeds, the fitness of the objective function differs greatly. While parameter b is correctly identified in all three simulation samples, estimates of parameter a happen to be far (left plot) to slightly off (middle plot) compared with the true value of $a = 0.0003$. One might be even lucky enough to identify both parameters correctly (right plot).

Increasing the number of observations from $T = 40\,000$ to $T = 400\,000$, we find a smoother surface as well as less variations as expected. However, even for this higher sample size, local optima arise (see the middle plot) depending again on the choice (or luck) of random seeds. Comparing this with an ensemble of $N = 10$ simulation runs

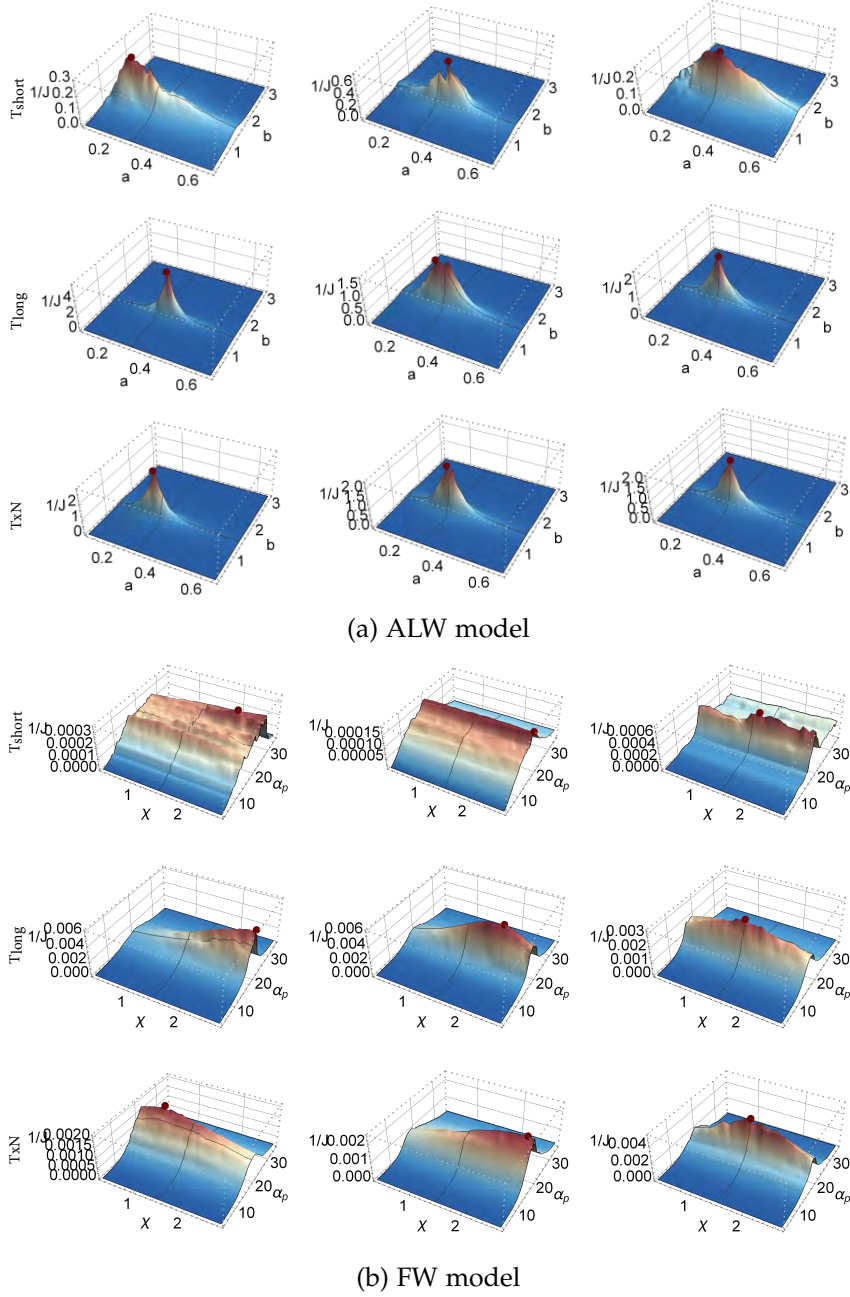


Figure 6: **3D surface plots for critical model parameters.** Contour plots of the inverse objective function J^{-1} for varying values for parameter pairs (a, b) for the ALW model in Fig. 6a and (χ, α_p) for the FW model in Panel Fig. 6b. Sample simulations are run over different settings: $T_{\text{short}} = 40\,000$, $T_{\text{long}} = 400\,000$ and $T \times N$ with $T = 40\,000$ and $N = 10$. Black lines indicate the (pseudo-)true parameter values for selected model parameters and red dots denote the highest peak of the inverse of the objective function's response surface.

for $T = 40\,000$, we qualitatively observe the same patterns.

During a pre-exploratory experiment for the FW model, we identified the most problematic parameter pairs (χ, α_0) and (χ, α_p) . The latter pair is visualized in Fig. 6.⁴⁴ In the following, we set the (pseudo-)true values $(\chi, \alpha_p) = (1.5, 19.671)$ and grid points result from 41 equidistant parameter variations over $[0.3, 2.7]$ for parameter χ and over $[3.9342, 3.9342]$ for parameter α_p . Let us first note that we find qualitatively the same results as for the ALW model. Yet, the contour plots look rather different for the FW model. We observe that the surface is less elevated with peaks at much lower levels than we see for the ALW model. Especially for the T_{short} -setting, we find very flat hilltops for parameter χ and even flat plateaus for both parameters. Increasing the number of observations helps to smoothen the surface. However, identification issues leading to biases in the estimates might be more pronounced here which is most probably due to the higher dimensionality and complexity of the FW model.

While these contour plots of the objective function reveal a partially non-smooth surface, they allow no assessment of the estimator's further distributional properties. To do so, we need to include the third dimension of our simulation cube shown in Fig. 2 which is referred to the Monte Carlo repetitions. Hence, we run the simulations repeatedly over $M = 200$ and present graphical results in Figure 7. The left panel shows boxplots for the three different simulation settings. First, we notice that for the ALW model the estimator is consistent and overall able to detect the true value of parameter a on average. Unsurprisingly, variations are comparatively larger for the T_{short} -setting. Regarding the other two settings, variations differ only slightly with almost unbiased estimates for the T_{long} -setting and a small downward bias in the mix-setting. The paired histogram plot on the right panel in Fig. 7 confirms this observation. While for the T_{long} -simulation the histogram is quite symmetric, the histogram for the mix-simulation is skewed towards lower values of a .

For the FW model results differ slightly. We find small biases of the estimates for the T_{short} and T_{long} -cases while the $T \times N$ -setting is able to perfectly identify the true value of parameter χ on average. This suggest that the ensemble setting might lead to more consistent and efficient estimates. However, we are cautious to not overinterpret our results here given that we have included only two varying parameters while keeping the rest fixed.

⁴⁴Graphical results for the objective function's response surface of the FW model are almost identical for parameter pairs (χ, α_0) and (χ, α_p) .

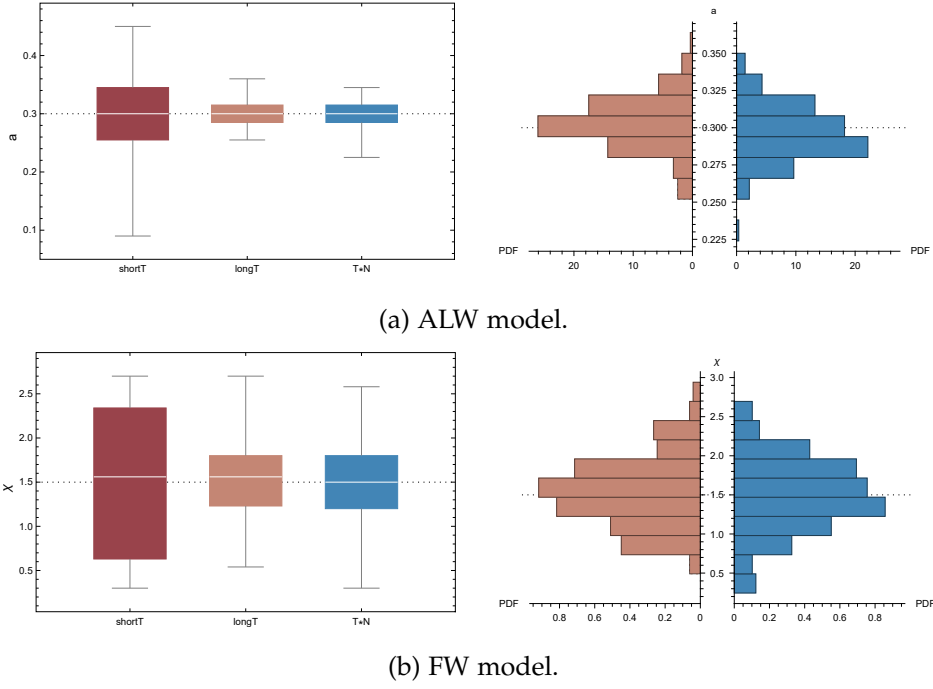


Figure 7: **Box plots and paired histograms for critical model parameters.** Upper panel Fig. 7a for the ALW model and lower panel Fig. 7b for the FW model. Left figures show box plots of the estimates for parameter a for the ALW model and parameter χ for the FW model. The settings are the following: $T_{\text{short}} = 40\,000$, $T_{\text{long}} = 400\,000$ and $T \times N$ with $T = 40\,000$, $N = 10$, whereby the latter two have the same budget of observations. Simulations are run over $M = 200$ Monte Carlo repetitions.

3.3.2 Numerical estimation experiment

This brings us directly to the next experiment where we include all model parameters for estimation. In the following, we will run a small estimation exercise over a Sobol-sequenced sample of 2000 parameter combinations. The included parameter space is given by $\pm 25\%$ of the true values which results in the following parameter ranges: $a \in [0.225, 0.375]$, $b \in [1.05, 1.75]$, $\sigma_f \in [22.5, 37.5]$ for the ALW model and $\phi \in [0.09, 0.15]$, $\chi \in [1.125, 1.875]$, $\alpha_0 \in [-0.252, -0.42]$, $\alpha_n \in [1.379, 2.299]$, $\alpha_p \in [14.753, 24.589]$, $\sigma_i \in [0.531, 0.885]$, $\sigma_c \in [1.61, 2.684]$ for the FW model.⁴⁵ The number of Monte Carlo runs is set again to $M = 200$. Since we are interested in the degree of uncertainty that solely

⁴⁵The defined parameter ranges provide broad parameter variations while keeping execution times of the simulation runs feasible.

comes from broken ergodicity, we abstain from running a full estimation exercise given that the choice of the optimization algorithm would highly influence and distort our results.⁴⁶ We report graphical results in Fig. 8 and the corresponding mean estimates together with their standard deviations and root-mean squared errors in Tab. 4. Looking at Fig. 8, the left vertical axis shows the value of the objective function J . Settings (a) and (b) result in the lowest values of J . For the ALW model, it seems that both (a) and (b) are compatible in terms of minimized objective function. Mean sample estimates in Tab. 4 confirm these findings. Decreasing the length of T while increasing the ensemble size N raises the level of uncertainty in the estimates, pushing mean estimates slowly away from their true values. Accordingly, standard deviations and root-mean squared errors tend to increase, too.

Regarding the FW model and its comparatively higher dimensionality, results are qualitatively similar yet the loss in efficiency for increasing ensemble sizes N is more pronounced here. This loss in efficiency becomes particularly obvious for setting (d) with $T = 10\,000$ and $N = 40$. In fact, the number of observations per realized time series is too short here resulting in uninformative samples.

Summing up, we find different values of the objective function J for different allocations of the total budget of observations. While the differences might be negligible for longer time lengths T , they significantly deteriorate estimates for shorter time lengths due to broken ergodicity. Recall the ergodic theorem in Eq. (2.7) which is effective here.

However, investigating only the finite sample properties of the parameter estimates neglects the computational resources needed to run these simulations. Therefore, we add the execution times as black stars, belonging to the right vertical axis, to Fig. 8. Setting (a) is by far the most computationally demanding simulation. The three $T \times N$ -settings (b), (c) and (d) have similar execution times which reduce the benchmark setting (a) by a factor of three for both the ALW model and FW model. The reason is that single simulations over T_{long} are memory intense and inefficient. In fact, repeated simulations over different random seeds are so called *embarrassingly parallel* tasks and can, thus, be easily distributed over multiple processing units which is done here.⁴⁷ We finally run a fifth scenario (e) for which we take the best performing $T \times N$ -setting (b)

⁴⁶WINKER and MARINGER (2009) provide a study of the joint convergence of an estimator and a heuristic optimization algorithm. They compute the necessary number of Monte Carlo repetitions of the optimization routine to derive robust estimates.

⁴⁷All numerical simulations in this paper have been performed using Julia 1.0.5 on a standard desktop computer with an eight core 3.00 Ghz AMD Ryzen 7 1700 CPU and 32GB of memory. We made sure to run all simulation settings most efficiently to keep the different scenarios comparable.

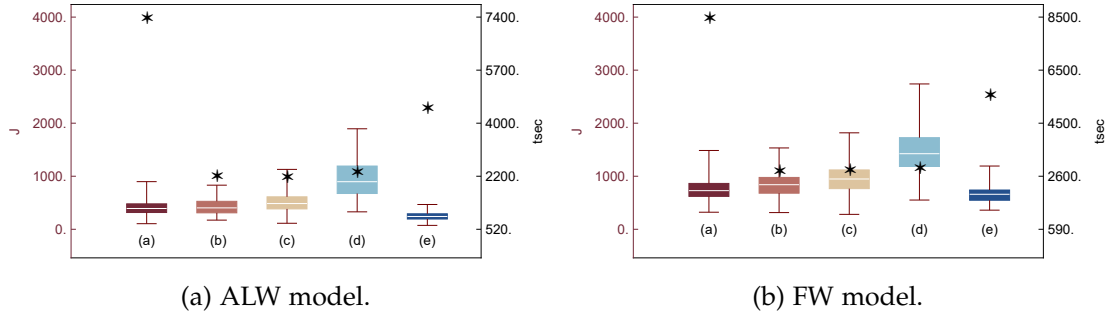


Figure 8: **Box plots and timing results of estimation exercise.** Box plots of objective function J (left vertical axes) together with star points representing the required computational time to run the simulations (right vertical axes). Left panel Fig. 8a shows results for the ALW model. Right panel Fig. 8b shows results for the FW model. The scenarios allocate differently the total budget of 400 000 observations in the following ways: (a) $T = 400\,000, N = 1$; (b) $T = 40\,000, N = 10$; (c) $T = 20\,000, N = 20$; (d) $T = 10\,000, N = 40$ and (e) $T = 40\,000, N = 20$. Simulations are run over $M = 200$ Monte Carlo repetitions.

with $T = 40\,000$ but double the amount of ensemble size $N = 20$. This setting is able to easily outperform all other settings, even scenario (a) while taking only two third of setting (a)'s computational time. We find this to be a consistent pattern for both FABMs considered here.

With this analysis, we are not only aiming for an adequate mixture of $T \times N$ from the perspective of (broken) ergodicity but also considering the associated computational costs. We conclude that for a limited budget of computational resources, using a mix of $T \times N$ is favourable over a single long simulation run with the same number of observations.

Table 4: **Estimation results.** The table lists mean estimates, their standard deviations and root-mean squared errors of the estimated parameters for the scenarios (a)-(e) given in Tab. 3. We also report J -values in the last column of the ALW model and FW model, respectively. Results are based on $M = 200$ repetitions.

		ALW model				FW model							
		a	b	σ_f	J	ϕ	χ	α_0	α_n	α_p	σ_f	σ_c	J
θ_0		0.3	1.4	30	0	0.12	1.5	-0.336	1.839	19.671	0.708	2.147	0
(a)	mean	0.298	1.4	30.003	400.09	0.123	1.801	-0.287	1.871	19.153	0.702	2.19	747.07
	STD	0.03	0.019	0.219	125.26	0.016	0.108	0.024	0.135	1.253	0.009	0.11	193.05
	RMSE	0.03	0.019	0.218	0	0.016	0.32	0.055	0.138	1.353	0.011	0.118	0
(b)	mean	0.292	1.405	30.017	423.84	0.123	1.785	-0.284	1.88	18.905	0.702	2.194	829.72
	STD	0.026	0.022	0.207	149.98	0.015	0.147	0.019	0.131	0.984	0.01	0.119	222.86
	RMSE	0.027	0.022	0.208	0	0.015	0.321	0.056	0.137	1.245	0.012	0.128	0
(c)	mean	0.288	1.412	30.086	504.92	0.121	1.696	-0.296	1.925	18.564	0.707	2.153	963.66
	STD	0.03	0.021	0.258	179.94	0.017	0.233	0.032	0.14	1.184	0.013	0.126	262.44
	RMSE	0.033	0.024	0.272	0	0.017	0.304	0.051	0.164	1.619	0.013	0.126	0
(d)	mean	0.287	1.426	30.111	948.87	0.123	1.579	-0.306	1.982	18.212	0.714	2.109	1481.87
	STD	0.028	0.024	0.282	344.38	0.016	0.249	0.038	0.116	1.111	0.013	0.112	386
	RMSE	0.031	0.035	0.302	0	0.016	0.261	0.048	0.184	1.832	0.015	0.118	0
(e)	mean	0.29	1.403	30.033	244.35	0.128	1.832	-0.277	1.897	18.634	0.702	2.179	622.62
	STD	0.019	0.013	0.162	70.73	0.009	0.084	0.01	0.068	0.49	0.005	0.065	156.9
	RMSE	0.022	0.013	0.165	0	0.012	0.342	0.06	0.089	1.147	0.008	0.073	0

3.3.3 Robustness of results

How robust are these numerical results especially regarding variations of the moment function set? In the following, we try to answer this question by running sensitivity analyses for varying sizes K of the moment vector \mathbf{m} . For this, we randomly sample 17 different moment sets per number of moment vector size and evaluate them for the numerical estimation data of setting (a). Fitness is defined as the euclidean distance between the (pseudo-)true and estimated parameter vector. We decide to investigate two different scenarios here. For the first scenario we sample from the complete range of 18 moment functions as listed in Tab. 1. For the second scenario we exclude m_6 known as the autocorrelation of raw returns at lag 1. The reason is that we have seen for example in Fig. 3 problematic to non-convergent behaviour for moment function m_6 . We suspect improved results in terms of fitness given that m_6 offers very limited identification power. We present graphical results in Fig. 9. The two figures show boxplots for the ALW model on the left and the FW model on the right. Unsurprisingly, results improve for an increasing number of included moment functions. More interestingly

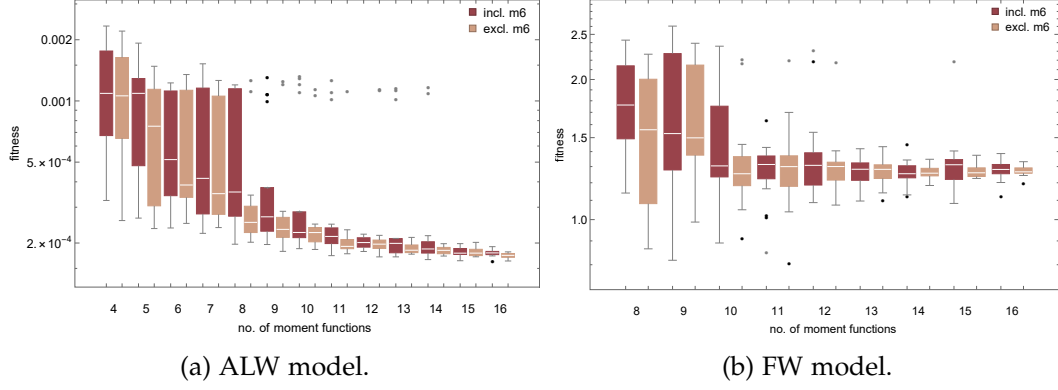


Figure 9: **Sensitivity analysis for varying number of moment functions.** Boxplots show improved fitness for increased moment vector size with fitness defined as the euclidean distance between the (pseudo-)true and estimated parameter vector. We consider two different settings for the sampling of moment functions: dark red boxplots including moment function m_6 and orange boxplots excluding m_6 . Left panel Fig. 9a shows results for the ALW model. Right panel Fig. 9b shows results for the FW model. Note that we use a logarithmic scale for the vertical axis to better visualize the fitness.

and rather unexpectedly, the improvement does not happen linearly. While variations in the fitness are quite large for smaller moment vector sizes, we observe a sudden drop for moment vector sizes of nine (ALW model) and ten (FW model), respectively. After that drop results are pretty robust for further increasing number of moments. This is especially true for the FW model while we see slight improvements for the ALW model. Comparing both scenarios (with and without m_6 , graphically represented by dark red and orange) we find consistently better fitness for moment sample sets excluding the autocorrelation of raw returns at lag 1. Additionally, estimates related to this scenario tend to produce less outliers (see the grey and black dots).

This suggest that re-evaluating our estimation results from Tab. 4 for an adjusted moment vector should lead to significantly better results. Therefore, we take our previous results as presented in Tab. 4 and recompute them for a reduced moment vector excluding m_6 . We relabel these results J_{18} since they have been actually estimated for a complete set of 18 moments. Now we can compare them with the J_{17} results that we estimated based on 17 moment functions. We report the differences between J_{18} and J_{17} values for all five settings in Tab. 5. As we anticipated, we find clearly improved J -values as a common and consistent pattern. Furthermore, one-tailed t -tests reveal statistical significance for all results except for setting (d).

Table 5: *J* results for adjusted moment vector. For the scenarios (a) – (e) given in Tab. 3. Significance levels of a one-tailed *t*-test are represented by *, ** and *** for 10%, 5% and 1%, respectively.

	$J_{18} - J_{17}$				
	(a)	(b)	(c)	(d)	(e)
ALW model	23.84***	26.24***	23.15**	23.25	12.17***
FW model	27.42**	34.02**	29.18*	24.8	17.99**

4 Conclusion

(F)ABMs are often conceived as black boxes lacking a sound and well-behaved mathematical model. This has caused considerable resistance in the (economics) community to accept computer simulations as proper research methods. The development of estimation and validation tools that are particularly suited for the special properties of most (F)ABMs is a very active field of research. This paper aims to contribute to this line of research by explicitly considering the uncertainty coming from broken ergodicity.

We systematically study the properties and convergence behaviour of a SMM estimator and its individual moment functions. We have seen that assuring a one-to-one mapping between model output in terms of moment functions and model parameters is not trivial. Most moment functions are indeed non-linear and non-monotone. This can lead to biased estimations. More important is however the relative responsiveness of moment functions towards changes in the model parameters, *i.e.* how strongly they react. Parameters for which we find weak responsiveness of moment functions tend to be the most troublesome during estimation. We have further shown that not all moment functions which are considered to be a popular choice for the estimation of univariate asset pricing models are actually suited. As it turns out, the convergence of the autocorrelation of raw returns at lag 1 is highly erratic making it an unfavourable candidate statistic. We show that estimation results improve significantly when leaving out this summary statistic. For all other moment functions we find robust results for the moment matching fitness. Therefore, our analysis confirms our choice of moment functions for (univariate) asset price models and thereby offers a response towards the arbitrariness critique of SMM.

Estimation methods like SMM are perfect tools for perfect models. Living in an ever-changing world, it should be no design fault to create models that possess non-

ergodic properties. Even if regularity conditions are not fully met, models and tools may still be useful. Being aware of broken ergodicity is the key here. We further suggest researchers to not blindly follow previous studies based on other models. In fact, every (F)ABM behaves differently as we have seen here. Therefore, we run a bunch of Monte Carlo analyses to learn about the FABMs' sensitivity and responsiveness towards changes in model parameters and simulation settings. We have shown that for a given computational time budget an adequate mixture of ensemble size N and time length T is better suited than the same number of observations in one realization. We can conclude that due to careful selection of moment functions and consideration of ergodicity, we can improve the objective function's potential to correctly identify model parameters for a limited budget of computational resources.

So far, the energy demand of simulation models seems to play a minor role. Access to institute-owned high performance computing clusters and the recent global increase in cloud computing may have made the necessity for energy-efficient simulations obsolete. Yet, the opposite is true given the looming threat of climate change and a likely exponential increase in demand for simulations like FABMs especially for policy purposes in the near future. A holistic approach to fight global warming will include many different aspects of our lives. This also includes an efficient and mindful use of resources when it comes to research. Our work provides steps into this direction.

We have limited our focus in this paper on the use of an SMM estimator. The scale of our work can be extended to including other estimation methods as well. We specifically aim to apply our approach to likelihood-based methods. Since computing the likelihood function is expensive in terms of computational costs, applying our insights might improve cost efficiency there too.

References

- AKHIEZER, NAUM I. (1965). *The Classical Moment Problem. and some related questions in analysis*. Oliver & Boyd, Edinburgh London ()
- ALFARANO, SIMONE, THOMAS LUX and FRIEDRICH WAGNER (2005). 'Estimation of agent-based models: the case of an asymmetric herding model'. *Computational Economics* 26 (1), 19–49. DOI: [10.1007/s10614-005-6415-1](https://doi.org/10.1007/s10614-005-6415-1) ()
- ALFARANO, SIMONE, THOMAS LUX and FRIEDRICH WAGNER (2008). 'Time variation of higher moments in a financial market with heterogeneous agents: An analytical approach'. *Journal of Economic Dynamics & Control* 32 (1), 101–136. DOI: [10.1016/j.jedc.2006.12.014](https://doi.org/10.1016/j.jedc.2006.12.014) ()
- AMILON, HENRIK (2008). 'Estimation of an adaptive stock market model with heterogeneous agents'. *Journal of Empirical Finance* 15.2, 342–362. DOI: [10.1016/j.jempfin.2006.06.007](https://doi.org/10.1016/j.jempfin.2006.06.007) ()
- BARDE, SYLVAIN (2016). 'Direct comparison of agent-based models of herding in financial markets'. *Journal of Economic Dynamics & Control* 73, 329–353. ISSN: 0165-1889. DOI: [10.1016/j.jedc.2016.10.005](https://doi.org/10.1016/j.jedc.2016.10.005) ()
- BERTSCHINGER, NILS and IURI MOZZHORIN (2021). 'Bayesian estimation and likelihood-based comparison of agent-based volatility models'. *Journal of Economic Interaction and Coordination* 16 (1), 173–210. DOI: [10.1007/s11403-020-00289-z](https://doi.org/10.1007/s11403-020-00289-z) ()
- BOLOTNIKOV, VLADIMIR (1996). 'On degenerate Hamburger moment problem and extensions of nonnegative Hankel block matrices'. *Integral Equations and Operator Theory* 25 (3), 253–276. DOI: [10.1007/bf01262294](https://doi.org/10.1007/bf01262294) ()
- BOLOTNIKOV, VLADIMIR (2008). 'On degenerate Hamburger moment problem and extensions of positive semidefinite Hankel block matrices'. The corrected version of the paper BOLOTNIKOV (1996). arXiv: [0812.4567](https://arxiv.org/abs/0812.4567) ()
- BOUCHAUD, JEAN-PHILIPPE (2008). 'Economics needs a scientific revolution'. *Nature* 455 (7217), 1181. DOI: [10.1038/4551181a](https://doi.org/10.1038/4551181a) ()
- BOUCHAUD, JEAN-PHILIPPE and ROGER FARMER (2021). 'Self-Fulfilling Prophecies, Quasi Non-Ergodicity and Wealth Inequality'. arXiv: [2012.09445](https://arxiv.org/abs/2012.09445) [[econ.TH](https://arxiv.org/archive/econ)] ()
- CHEN, GONG-NING and YONG-JIAN HU (1998). 'The truncated Hamburger matrix moment problems in the nondegenerate and degenerate cases, and matrix continued fractions'. *Linear Algebra and its Applications* 277 (1-3), 199–236. DOI: [10.1016/s0024-3795\(97\)10076-3](https://doi.org/10.1016/s0024-3795(97)10076-3) ()

- CHEN, ZHENXI and THOMAS LUX (2018). 'Estimation of Sentiment Effects in Financial Markets: A Simulated Method of Moments Approach'. *Computational Economics* 52 (3), 711–744. DOI: [10.1007/s10614-016-9638-4](https://doi.org/10.1007/s10614-016-9638-4) ()
- CONT, RAMA (2001). 'Empirical properties of asset returns: stylized facts and statistical issues'. *Quantitative Finance* 1 (2), 223–236. DOI: [10.1088/1469-7688/1/2/304](https://doi.org/10.1088/1469-7688/1/2/304) ()
- DAVIDSON, RUSSELL and JAMES G. MACKINNON (2003). *Econometric Theory and Methods* ()
- DELLI GATTI, DOMENICO, GIORGIO FAGIOLO, MAURO GALLEGATI, MATTEO RICHIARDI and ALBERTO RUSSO, eds. (2018). *Agent-based Models in Economics. A Toolkit*. Cambridge University Press. DOI: [10.1017/9781108227278](https://doi.org/10.1017/9781108227278) ()
- DELLI GATTI, DOMENICO and JAKOB GRAZZINI (2020). 'Rising to the challenge: Bayesian estimation and forecasting techniques for macroeconomic Agent Based Models'. *Journal of Economic Behavior & Organization* 178, 875–902. DOI: [10.1016/j.jebo.2020.07.023](https://doi.org/10.1016/j.jebo.2020.07.023) ()
- DESSERTAINE, THÉO, JOSÉ MORAN, MICHAEL BENZAQUEN and JEAN-PHILIPPE BOUCHAUD (2022). 'Out-of-equilibrium dynamics and excess volatility in firm networks'. *Journal of Economic Dynamics and Control* 138, 104362. DOI: [10.1016/j.jedc.2022.104362](https://doi.org/10.1016/j.jedc.2022.104362) ()
- DOMOWITZ, IAN and MAHMOUD A. EL-GAMAL (1993). 'A Consistent Test of Stationary-Ergodicity'. *Econometric Theory* 9 (04), 589–601. DOI: [10.1017/s0266466600007994](https://doi.org/10.1017/s0266466600007994) ()
- DOMOWITZ, IAN and MAHMOUD A. EL-GAMAL (2001). 'A consistent nonparametric test of ergodicity for time series with applications'. *Journal of Econometrics* 102 (2), 365–398. DOI: [10.1016/s0304-4076\(01\)00058-6](https://doi.org/10.1016/s0304-4076(01)00058-6) ()
- DOSI, GIOVANNI, M. C. PEREIRA, ANDREA ROVENTINI and M. E. VIRGILLITO (2018). 'Causes and consequences of hysteresis: aggregate demand, productivity, and employment'. *Industrial and Corporate Change* 27 (6), 1015–1044. DOI: [10.1093/icc/dty010](https://doi.org/10.1093/icc/dty010) ()
- DUFFIE, DARRELL and KENNETH J. SINGLETON (1993). 'Simulated Moments Estimation of Markov Models of Asset Prices'. *Econometrica* 61 (4), 929–952. DOI: [10.2307/2951768](https://doi.org/10.2307/2951768) ()
- DYUKAREV, YURIY M., BERND FRITZSCHE, BERND KIRSTEIN, CONRAD MÄDLER and HELGE C. THIELE (2008). 'On Distinguished Solutions of Truncated Matricial Hamburger Moment Problems'. *Complex Analysis and Operator Theory* 3 (4), 759–834. DOI: [10.1007/s11785-008-0061-2](https://doi.org/10.1007/s11785-008-0061-2) ()
- FAGIOLO, GIORGIO, MATTIA GUERINI, FRANCESCO LAMPERTI, ALESSIO MONETA and ANDREA ROVENTINI (2019). 'Validation of Agent-Based Models in Economics and Fin-

- ance'. *Computer Simulation Validation. Fundamental Concepts, Methodological Frameworks, and Philosophical Perspectives*. Ed. by CLAUS BEISBART and NICOLE J. SAAM. Simulation Foundations, Methods and Applications. Springer International Publishing, 763–787. DOI: [10.1007/978-3-319-70766-2_31](https://doi.org/10.1007/978-3-319-70766-2_31) ()
- FAGIOLO, GIORGIO and ANDREA ROVENTINI (2017). 'Macroeconomic Policy in DSGE and Agent-Based Models Redux: New Developments and Challenges Ahead'. *Journal of Artificial Societies and Social Simulation* 20 (1). DOI: [10.18564/jasss.3280](https://doi.org/10.18564/jasss.3280) ()
- FAMA, EUGENE F. (1963). 'Mandelbrot and the stable Paretian hypothesis'. *Journal of Business* 36 (4), 420–429. JSTOR: [2350971](https://www.jstor.org/stable/2350971) ()
- FARMER, J. DOYNE and DUNCAN FOLEY (2009). 'The economy needs agent-based modeling'. *Nature* 460 (7256), 685–686. DOI: [10.1038/460685a](https://doi.org/10.1038/460685a) ()
- FRANKE, REINER (2009). 'Applying the method of simulated moments to estimate a small agent-based asset pricing model'. *Journal of Empirical Finance* 16 (5), 804–815. DOI: [10.1016/j.jempfin.2009.06.006](https://doi.org/10.1016/j.jempfin.2009.06.006) ()
- FRANKE, REINER and FRANK WESTERHOFF (2012). 'Structural stochastic volatility in asset pricing dynamics: Estimation and model contest'. *Journal of Economic Dynamics & Control* 36 (8), 1193–1211. DOI: [10.1016/j.jedc.2011.10.004](https://doi.org/10.1016/j.jedc.2011.10.004) ()
- FRANKE, REINER and FRANK WESTERHOFF (2016). 'Why a simple herding model may generate the stylized facts of daily returns: explanation and estimation'. *Journal of Economic Interaction and Coordination* 11 (1), 1–34. DOI: [10.1007/s11403-014-0140-6](https://doi.org/10.1007/s11403-014-0140-6) ()
- GAFFEO, EDOARDO, MICHELE CATALANO, FABIO CLEMENTI, DOMENICO DELLI GATTI, MAURO GALLEGATI and ALBERTO RUSSO (2007). 'Reflections on modern macroeconomics: Can we travel along a safer road?' *Physica A: Statistical Mechanics and its Applications* 382 (1), 89–97. DOI: [10.1016/j.physa.2007.02.011](https://doi.org/10.1016/j.physa.2007.02.011) ()
- GALLANT, A RONALD and GEORGE TAUCHEN (1996). 'Which moments to match?' *Econometric Theory* 12 (4), 657–681. DOI: [10.1017/S0266466600006976](https://doi.org/10.1017/S0266466600006976) ()
- GHONGHADZE, JABA and THOMAS LUX (2016). 'Bringing an elementary agent-based model to the data: Estimation via GMM and an application to forecasting of asset price volatility'. *Journal of Empirical Finance* 37, 1–19. DOI: [10.1016/j.jempfin.2016.02.002](https://doi.org/10.1016/j.jempfin.2016.02.002) ()
- GILLI, MANFRED and PETER WINKER (2003). 'A global optimization heuristic for estimating agent based models'. *Computational Statistics & Data Analysis* 42 (3), 299–312. DOI: [10.1016/S0167-9473\(02\)00214-1](https://doi.org/10.1016/S0167-9473(02)00214-1) ()
- GRAZZINI, JAKOB (2011). 'Consistent Estimation of Agent Based Models'. 110. URL: http://www.laboratoriorevelli.it/_pdf/wp110.pdf ()

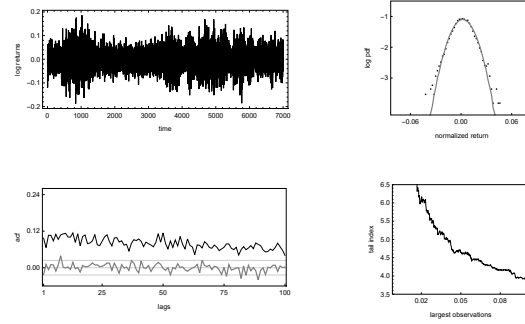
- GRAZZINI, JAKOB (2012). 'Analysis of the Emergent Properties: Stationarity and Ergodicity'. *Journal of Artificial Societies and Social Simulation* 15 (2). DOI: [10.18564/jasss.1929](https://doi.org/10.18564/jasss.1929) ()
- GRAZZINI, JAKOB and MATTEO RICHIARDI (2013). 'Consistent Estimation of Agent Based Models by Simulated Minimum Distance'. 130. URL: http://www.laboratoriorevelli.it/_pdf/wp130.pdf ()
- GRAZZINI, JAKOB and MATTEO RICHIARDI (2015). 'Estimation of ergodic agent-based models by simulated minimum distance'. *Journal of Economic Dynamics & Control* 51, 148–165. DOI: [10.1016/j.jedc.2014.10.006](https://doi.org/10.1016/j.jedc.2014.10.006) ()
- GRAZZINI, JAKOB, MATTEO RICHIARDI and LISA SELLA (2012). 'Small sample bias in MSM estimation of agent-based models'. *Managing Market Complexity. The Approach of Artificial Economics*. Ed. by ANDREA TEGLIO, SIMONE ALFARANO, EVA CAMACHO-CUENA and MIGUEL GINÉS-VILAR. Lecture Notes in Economics and Mathematical Systems. Springer Berlin Heidelberg, 237–247. DOI: [10.1007/978-3-642-31301-1_19](https://doi.org/10.1007/978-3-642-31301-1_19) ()
- GRAZZINI, JAKOB, MATTEO RICHIARDI and MIKE TSIONAS (2017). 'Bayesian estimation of agent-based models'. *Journal of Economic Dynamics and Control* 77, 26–47. DOI: [10.1016/j.jedc.2017.01.014](https://doi.org/10.1016/j.jedc.2017.01.014) ()
- GREENE, WILLIAM H. (2018). *Econometric Analysis*. 8th ed. Pearson, New York ()
- GUERINI, MATTIA and ALESSIO MONETA (2017). 'A method for agent-based models validation'. *Journal of Economic Dynamics & Control* 82, 125–141. DOI: [10.1016/j.jedc.2017.06.001](https://doi.org/10.1016/j.jedc.2017.06.001) ()
- HOMMES, CARS and BLAKE LeBARON, eds. (2018). *Handbook of Computational Economics. Heterogeneous Agent Modeling*. Vol. 4. Elsevier ()
- HOMMES, CARS H. (2006). 'Heterogeneous Agent Models in Economics and Finance'. *Handbook of Computational Economics*. Ed. by LEIGH TEFATSION and KENNETH L. JUDD. Vol. 2. Handbooks in Economics. North Holland Elsevier, Amsterdam Oxford (UK). Chap. 23, 1109–1186. DOI: [10.1016/s1574-0021\(05\)02023-x](https://doi.org/10.1016/s1574-0021(05)02023-x) ()
- KIRMAN, ALAN (1993). 'Ants, rationality, and recruitment'. *The Quarterly Journal of Economics* 108 (1), 137–156. DOI: [10.2307/2118498](https://doi.org/10.2307/2118498) ()
- KIRSTEIN, MARK (2019). 'The Ergodicity Problem in Economics. The Advent of Ergodicity Economics'. PhD thesis. TU Dresden. URL: bit.ly/mkphdr1 ()
- KUKACKA, JIRI and JOZEF BARUNIK (2017). 'Estimation of financial agent-based models with simulated maximum likelihood'. *Journal of Economic Dynamics and Control* 85, 21–45. DOI: [10.1016/j.jedc.2017.09.006](https://doi.org/10.1016/j.jedc.2017.09.006) ()

- KUKACKA, JIRI and LADISLAV KRISTOUFEK (2020). 'Do 'complex' financial models really lead to complex dynamics? Agent-based models and multifractality'. *Journal of Economic Dynamics and Control* 113, 103855. DOI: [10.1016/j.jedc.2020.103855](https://doi.org/10.1016/j.jedc.2020.103855) ()
- LAMPERTI, FRANCESCO (2018). 'An information theoretic criterion for empirical validation of simulation models'. *Econometrics and Statistics* 5, 83–106. DOI: [10.1016/j.ecosta.2017.01.006](https://doi.org/10.1016/j.ecosta.2017.01.006) ()
- LEBARON, BLAKE (2006). 'Agent-based Computational Finance'. *Handbook of Computational Economics*. Ed. by LEIGH TEFATSION and KENNETH L. JUDD. Vol. 2. Handbooks in Economics 13. North Holland Elsevier, Amsterdam Oxford (UK). Chap. 24, 1187–1233. DOI: [10.1016/S1574-0021\(05\)02024-1](https://doi.org/10.1016/S1574-0021(05)02024-1) ()
- LEBARON, BLAKE (2021). 'Microconsistency in Simple Empirical Agent-Based Financial Models'. *Computational Economics* 58 (1), 83–101. DOI: [10.1007/s10614-019-09917-8](https://doi.org/10.1007/s10614-019-09917-8) ()
- LEE, BONG-SOO and BETH FISHER INGRAM (1991). 'Simulation Estimation of Time Series Models'. *Journal of Econometrics* (2-3), 197–205. DOI: [10.1016/0304-4076\(91\)90098-X](https://doi.org/10.1016/0304-4076(91)90098-X) ()
- LEOMBRUNI, ROBERTO and MATTEO RICHIARDI (2005). 'Why are economists sceptical about agent-based simulations?' *Physica A: Statistical Mechanics and its Applications* 355 (1), 103–109. DOI: [10.1016/j.physa.2005.02.072](https://doi.org/10.1016/j.physa.2005.02.072) ()
- LUX, THOMAS (2009). 'Stochastic Behavioral Asset-Pricing Models and the Stylized Facts'. *Handbook of Financial Markets: Dynamics and Evolution*. Ed. by THORSTEN HENS and KLAUS REINER SCHENK-HOPPÉ. Elsevier, 161–215. DOI: [10.1016/B978-012374258-2.50007-5](https://doi.org/10.1016/B978-012374258-2.50007-5) ()
- LUX, THOMAS (2018). 'Estimation of agent-based models using sequential Monte Carlo methods'. *Journal of Economic Dynamics & Control* 91, 391–408. DOI: [10.1016/j.jedc.2018.01.021](https://doi.org/10.1016/j.jedc.2018.01.021) ()
- LUX, THOMAS (2021a). 'Approximate Bayesian Inference for ABMs in Economics: A case study'. *mimeo* ()
- LUX, THOMAS (2021b). 'Bayesian Estimation of Agent-Based Models via Adaptive Particle Markov Chain Monte Carlo'. *Computational Economics*. DOI: [10.1007/s10614-021-10155-0](https://doi.org/10.1007/s10614-021-10155-0) ()
- LUX, THOMAS and MARCEL AUSLOOS (2002). 'Market Fluctuations I: Scaling, Multiscaling, and Their Possible Origins'. *The Science of Disasters. Climate Disruptions, Heart Attacks, and Market Crashes*. Ed. by ARMIN BUNDE, JÜRGEN KROPP and HANS JOACHIM SCHELLNHUBER. Springer Berlin Heidelberg. Chap. 13, 372–409. DOI: [10.1007/978-3-642-56257-0_13](https://doi.org/10.1007/978-3-642-56257-0_13) ()

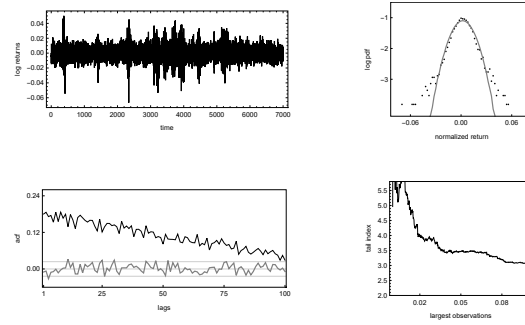
- LUX, THOMAS and FRANK WESTERHOFF (2009). 'Economics crisis'. *Nature Physics* 5 (1), 2–3. DOI: [10.1038/nphys1163](https://doi.org/10.1038/nphys1163) ()
- LUX, THOMAS and REMCO C.J. ZWINKELS (2018). 'Empirical Validation of Agent-Based Models'. *Handbook of Computational Economics. Heterogeneous Agent Modeling*. Ed. by CARS HOMMES and BLAKE LEBARON. Vol. 4. Elsevier, 437–488. DOI: [10.1016/bs.hescom.2018.02.003](https://doi.org/10.1016/bs.hescom.2018.02.003) ()
- MANDELBROT, BENOÎT B. (1960). 'The Pareto-Lévy Law and the Distribution of Income'. *International Economic Review* 1 (2), 79–106. DOI: [10.2307/2525289](https://doi.org/10.2307/2525289) ()
- MANDELBROT, BENOÎT B. (1963). 'The Variation of Certain Speculative Prices'. *Journal of Business* 36 (4), 394–419. eprint: [2351623](https://doi.org/10.2307/2351623) ()
- McFADDEN, DANIEL (1989). 'A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration'. *Econometrica* 57 (5), 703–708. DOI: [10.2307/1913621](https://doi.org/10.2307/1913621) ()
- METZLER, RALF, JAE-HYUNG JEON, ANDREY G. CHERSTVY and ELI BARKAI (2014). 'Anomalous diffusion models and their properties: non-stationarity, non-ergodicity, and ageing at the centenary of single particle tracking'. *Physical Chemistry Chemical Physics* 16 (44), 24128–24164. DOI: [10.1039/c4cp03465a](https://doi.org/10.1039/c4cp03465a) ()
- PAKES, ARIEL and DAVID POLLARD (1989). 'Simulation and the asymptotics of optimization estimators'. *Econometrica* 57 (5), 1027–1057. DOI: [10.2307/1913622](https://doi.org/10.2307/1913622) ()
- PETERSEN, KARL (1996). 'Ergodic Theorems and the Basis of Science'. *Synthese* 108 (2), 171–183. DOI: [10.1007/bf00413496](https://doi.org/10.1007/bf00413496) ()
- PLATT, DONOVAN (2020). 'A comparison of economic agent-based model calibration methods'. *Journal of Economic Dynamics and Control* 113, 103859. DOI: [10.1016/j.jedc.2020.103859](https://doi.org/10.1016/j.jedc.2020.103859) ()
- PLATT, DONOVAN (2021). 'Bayesian estimation of economic simulation models using neural networks'. *Computational Economics*, 1–52. DOI: [10.1007/s10614-021-10095-9](https://doi.org/10.1007/s10614-021-10095-9) ()
- RACHEV, SVETLOZAR T., ed. (2003). *Handbook of Heavy Tailed Distributions in Finance*. North-Holland, Amsterdam. DOI: [10.1016/b978-0-444-50896-6.x5000-6](https://doi.org/10.1016/b978-0-444-50896-6.x5000-6) ()
- RICHIARDI, MATTEO, ROBERTO LEOMBRUNI, NICOLE SAAM and MICHELE SONNESSA (2006). 'A Common Protocol for Agent-Based Social Simulation'. *Journal of Artificial Societies and Social Simulation* 9.1. DOI: <https://www.jasss.org/9/1/15.html> ()
- RUGE-MURCIA, FRANCISCO J. (2007). 'Methods to estimate dynamic stochastic general equilibrium models'. *Journal of Economic Dynamics & Control* 31 (8), 2599–2636. DOI: [10.1016/j.jedc.2006.09.005](https://doi.org/10.1016/j.jedc.2006.09.005) ()

- RUGE-MURCIA, FRANCISCO J. (2012). 'Estimating nonlinear DSGE models by the simulated method of moments: With an application to business cycles'. *Journal of Economic Dynamics & Control* 36 (6), 914–938. DOI: [10.1016/j.jedc.2012.01.008](https://doi.org/10.1016/j.jedc.2012.01.008) ()
- RUGE-MURCIA, FRANCISCO J. (2013). 'Generalized Method of Moments estimation of DSGE models'. *Handbook of Research Methods and Applications in Empirical Macroeconomics*. Ed. by NIGAR HASHIMZADE and MICHAEL A. THORNTON. Edward Elgar Publishing. Chap. 20, 464–485. DOI: [10.4337/9780857931023.00028](https://doi.org/10.4337/9780857931023.00028) ()
- SCHMITT, NOEMI (2021). 'Heterogeneous expectations and asset price dynamics'. *Macroeconomic Dynamics* 25.6, 1538–1568. DOI: [10.1017/S1365100519000774](https://doi.org/10.1017/S1365100519000774) ()
- SCHWARTZ, IVONNE (2022). 'Estimation of agent-based models: Testing and applying a simulated joint moment approach' ()
- SOLTYK, SYLVIA J. and FELIX CHAN (2021). 'Modeling time-varying higher-order conditional moments: A survey'. *Journal of Economic Surveys*. DOI: [10.1111/joes.12481](https://doi.org/10.1111/joes.12481) ()
- TESFATSION, LEIGH and KENNETH L. JUDD, eds. (2006). *Handbook of Computational Economics. Agent-based Computational Finance*. Vol. 2. North Holland Elsevier, Amsterdam Oxford. DOI: [10.1016/S1574-0021\(05\)02024-1](https://doi.org/10.1016/S1574-0021(05)02024-1) ()
- WESTERHOFF, FRANK and REINER FRANKE (2018). 'Agent-Based Models for Economic Policy Design. Two Illustrative Examples'. *The Oxford Handbook of Computational Economics and Finance*. Ed. by SHU-HENG CHEN, MAK KABOUDAN and YE-RONG DU. Oxford University Press. Chap. 18, 520–558. DOI: [10.1093/oxfordhb/9780199844371.013.40](https://doi.org/10.1093/oxfordhb/9780199844371.013.40) ()
- WINKER, PETER, MANFRED GILLI and VAHIDIN JELESKOVIC (2007). 'An objective function for simulation based inference on exchange rate data'. *Journal of Economic Interaction and Coordination* 2 (2), 125–145. DOI: [10.1007/s11403-007-0020-4](https://doi.org/10.1007/s11403-007-0020-4) ()
- WINKER, PETER and DIETMAR MARINGER (2009). 'The convergence of estimators based on heuristics: theory and application to a GARCH model'. *Computational Statistics* 24.3, 533–550. DOI: [10.1007/s00180-008-0145-5](https://doi.org/10.1007/s00180-008-0145-5) ()

A Sample simulation runs of model dynamics



(a) ALW model.



(b) FW model.

Figure 10: **Model dynamics of sample simulation runs.** Upper panel Fig. 10a shows dynamics for the ALW model using the following parameter set: $a = 0.3$, $b = 1.4$, $\sigma_f = 30$. The lower panel Fig. 10b reveals dynamics for the FW model using the following parameter set: $\phi = 0.12$, $\chi = 1.5$, $\alpha_0 = -0.336$, $\alpha_n = 1.839$, $\alpha_p = 19.671$, $\sigma_f = 0.708$, $\sigma_c = 2.147$. The panels show from top left clockwise to bottom left the evolution of log returns, the log probability density functions of normalized returns (black) and standard normally distributed returns (gray), the Hill tail index estimator and the autocorrelation functions of raw (gray) and absolute returns (black), respectively.