J Econ Interact Coord (2016) 11:1–34
DOI 10.1007/s11403-014-0140-6

REGULAR ARTICLE

Why a simple herding model may generate the stylized facts of daily returns: explanation and estimation

Reiner Franke · Frank Westerhoff

Received: 22 April 2013 / Accepted: 2 August 2014 / Published online: 20 August 2014

Abstract The paper proposes an elementary agent-based asset pricing model that, invoking the two trader types of fundamentalists and chartists, comprises four features: (i) price determination by excess demand; (ii) a herding mechanism that gives rise to a macroscopic adjustment equation for the market fractions of the two groups; (iii) a rush towards fundamentalism when the price misalignment becomes too large; and (iv) a stronger noise component in the demand per chartist trader than in the demand per fundamentalist trader, which implies a structural stochastic volatility in the returns. Combining analytical and numerical methods, the interaction between these elements is studied in the phase plane of the price and a majority index. In addition, the model is estimated by the method of simulated moments, where the choice of the moments reflects the basic stylized facts of the daily returns of a stock market index. A (parametric) bootstrap procedure serves to set up an econometric test to evaluate the model's goodness-of-fit, which proves to be highly satisfactory. The bootstrap also makes sure that the estimated structural parameters are well identified.

Keywords Structural stochastic volatility · Method of simulated moments · Autocorrelation pattern · Fat tails · Bootstrapped p values

Helpful comments by two anonymous referees and an associate editor are gratefully acknowledged. With respect to a (more limited) precursor of this paper, we also wish to express our thanks to a referee of another journal for his or her careful remarks.

R. Franke (✉)
University of Kiel, Kiel, Germany
e-mail: franke@uni-bremen.de

F. Westerhoff
University of Bamberg, Bamberg, Germany

JEL Classification D84 · G12 · G14 · G15

1 Introduction

In the last two decades numerous models with heterogeneous interacting agents and simple heuristic trading strategies have been designed that in this way seek to contribute to an explanation of the behaviour of financial markets.[1] Guided by questionnaire evidence (Menkhoff and Taylor 2007), this literature focusses on the behaviour of fundamental and technical traders.[2] The latter, also called chartists, employ trading methods that attempt to extract buying and selling signals from past price movements (Murphy 1999). By contrast, fundamentalists bet on a reduction in the current mispricing with respect to some fundamental value of the asset (see already Graham and Dodd 1951).

Small models with extremely simple versions of these two strategies have proven to be quite successful in generating dynamic phenomena that share central characteristics with the time series from real financial markets, such as fat tails in the return distributions, volatility clustering and long memory effects. Two features are particularly useful in this respect. First, a device that permits the agents to switch between fundamentalist and technical trading, so that the market fractions of the two groups endogenously vary over time. Second, the concept of *structural stochastic volatility* (SSV henceforth). By this, we mean a random term that is added to the deterministic "core demand" of each of the two strategies, which is supposed to capture some of the real-life heterogeneity within the groups. Given that the two noise terms may differ in their variance, the variations of the market fractions will induce variations in the overall noise level of the asset demand, which then can carry over to the price dynamics.

Several models with these features have been put forward and (partly) also successfully estimated by Franke (2010) and Franke and Westerhoff (2011, 2012a, b). The present paper reconsiders a model of this origin that emphasizes a herding mechanism. Here we wish to provide an in-depth inquiry into its dynamic properties, which takes place in the phase plane of a majority index and the asset price. Integrating analytical and numerical methods, this framework allows us to study the conditions of a stochastic switching between a tranquil fundamentalist regime of relatively long duration and a more volatile chartist regime of shorter duration. In this way, we are able to go beyond the mere observation of a simulation outcome and obtain a better understanding of why the model performs so effectively.

We also take up the issue of estimating this model once again, albeit with two new aspects. First, the computation of the weighting matrix for the objective function is based on an alternative bootstrap procedure, which we have not seen applied before and which we believe is superior to the block bootstrap used in previous work. Apart

[1] For recent surveys of this burgeoning field of research, see Chiarella et al. (2009), Hommes (2006), Hommes and Wagener (2009), LeBaron (2006), Lux (2009a) and Westerhoff (2009), among others.

[2] Other evidence is based on laboratory experiments; see, e.g., Heemeijer et al. (2009) or Hommes et al. (2007).

from this improvement, we wish to make sure that the resulting parameter estimates are nevertheless robust. Second, complementary to the measures of a model's goodness-of-fit discussed in other contributions, we propose the concept of a more straightforward p value. This statistic is derived from a large number of re-estimations of the model which, in particular, give us a distribution of the minimized values of the objective function under the null hypothesis that the model is true. The model fails to be outright rejected if this p value exceeds the five per cent level; and the higher it is, the better the fit.

The estimation approach itself, which proves most suitable for our purpose of reproducing the aforesaid stylized facts, is the method of simulated moments (MSM). "Moments" refers to the time series of one or several variables and means certain selected summary statistics computed from them, the empirical values of which the model-generated moments should try to match. In our case, the latter have no analytical expressions but must be simulated. Hence the estimation searches for the parameter values of a model that minimize the distance between the empirical and simulated moments, where the distance is defined by a quadratic loss function (specified by the weighting matrix mentioned above). In the present context, the moments that we choose will reflect what is considered to be the most important stylized facts of the daily stock returns from the S&P 500 stock market index, in particular, volatility clustering and fat tails. After all, this is what the evaluation of the models in the literature usually centres around. It thus also goes without saying that the MSM estimation approach may equally be applied to other financial market models of a similar complexity.[3]

The remainder of the paper is organized as follows. The model is introduced in the next section. In Sect. 3 its dynamic properties are studied in the phase plane, first in a deterministic and then in the full stochastic setting. Section 4 briefly recapitulates the MSM approach, carries out the estimation on the empirical moments and then applies the econometric testing of the model's goodness-of-fit. At the same time these computations provide us with the confidence intervals of the estimated parameters. Section 5 concludes. Several appendices are added for the discussion of finer details. Appendix 1 summarizes the value added of the present paper *vis-à-vis* previous work. Appendix 2 contains a few remarks on the technical treatment of our herding mechanism in the earlier literature. The mathematical proofs of two propositions in the main text are relegated to Appendices 3, 4 and 5 collect some estimation details.

2 Formulation of the model

2.1 Excess demand and price adjustments

We consider a financial market for a risky asset on which the price changes are determined by excess demand. The market is populated by two types of speculative traders, fundamentalists and chartists. Fundamentalists have long time horizons and base their

[3] The choice of MSM does not rule out that other estimation methods may be tried as well. For a brief summary of the comparative advantages of MSM, see Franke (2009, pp. 804f). In our opinion, its main merits are the high transparency in the evaluation of a model's goodness-of-fit, and the relatively low computational cost.

demand on the differences between the current price and the fundamental value. Even though they might expect the gap between the two prices to widen in the immediate future, they do not trade on the likeliness of this event and rather choose to place their bets on an eventual rapprochement. Chartists, on the other hand, have a short-term perspective and bet on the most recent price movements, buying (selling) if prices have been rising (falling). However, the agents are allowed to switch from one strategy to the other, where their choice is governed by a herding mechanism combined with an evaluation of the most recent price levels.

Let us start with the demand for the asset.[4] We join numerous examples in the literature and, in the first step, postulate two extremely simple deterministic rules. These rules govern what we may call the core demand in each group. For the fundamentalists, this demand is inversely related to the deviations of the (log) price p_t from its fundamental value $p^\star$, where we treat the latter as an exogenously given constant (for simplicity and to show that no random walk behaviour of the fundamental value is required to obtain the stylized facts). On the other hand, the core demand of the group of chartists is hypothesized to be proportional to the returns they have just observed, i.e. $(p_t - p_{t-1})$, where as already indicated, the time unit may be thought of as one day.

A crucial feature of our models is that we add a noise term to *each* of these demand components (and not just their sum). The two terms are meant to reflect a certain within-group heterogeneity, which we do not wish to describe in full detail. Since the many individual digressions from the simple rules as well as their composition in each group will more or less accidentally fluctuate from period to period, it is a natural short-cut to have this heterogeneity represented by two independent and normally distributed random variables ε_t^f and ε_t^c for the fundamentalists and chartists, respectively.[5] Combining the deterministic and stochastic elements, the net demands of an average fundamentalist and chartist trader for the asset in period t are supposed to be given by

$$d_t^f = \phi\,(p^\star - p_t) + \varepsilon_t^f \qquad \varepsilon_t^f \sim N(0, \sigma_f^2) \quad \phi > 0 \tag{1}$$

$$d_t^c = \chi\,(p_t - p_{t-1}) + \varepsilon_t^c \quad \varepsilon_t^c \sim N(0, \sigma_c^2) \quad \chi \geq 0 \tag{2}$$

where here and in the following Greek symbols denote constant and nonnegative parameters. Total demand (normalized by the population size) results from multiplying d_t^f and d_t^c by the market fractions of the two groups.

[4] To be exact, by demand we mean the orders (positive or negative) per trading period, not the desired positions of the agents.

[5] For example, individual and presently active traders with a fundamentalist strategy may adopt different values for their fundamental price, they react with different intensities to their trading signal, or they experiment with more complex trading rules which may also be continuously subjected to further modifications. Similarly so for the chartists, which explains the independence of ε_t^f and ε_t^c. In short, the two noise variables can be conceived of as a most convenient short-cut of certain aspects that are more specifically (but to some extent also more arbitrarily) dealt with in models with hundreds or thousands of different agents that one would have to keep track of over time (see Farmer and Joshi 2002; LeBaron 2006).

It is an intricate matter to judge whether or not the stochastic noise may "dominate" the deterministic terms in (1) and (2). More specifically, it may be observed that a higher signal-to-noise ratio within the fundamental rule (1) implies a stronger mean-reversion, which would eventually lead to (counterfactual) negative autocorrelations in the raw returns. On the other hand, a higher signal-to-noise ratio within the chartist rule (2) will bring about more pronounced bubbles and thus positive autocorrelations in the returns (which would equally be counterfactual). We will leave it to the data to decide about the levels of these ratios and, in particular, whether the coefficients ϕ and χ are significantly different from zero. In this regard, it may be noted that $\chi = 0$ would turn the chartists into pure noise traders. Even the additional assumption of a zero variance $\sigma_c^2 = 0$ would make sense; under these circumstances 'chartism' is tantamount to not trading at all. In other words, the agents would choose between fundamentalist strategies and complete inactivity.[6]

Concerning the market fractions of fundamentalism and chartism, it will be convenient below to fix the population size at $2N$. Then, with n_t^f and n_t^c being the number of fundamentalists and chartists, define $x_t := (n_t^f - n_t^c)/2N$ as the majority index of the fundamentalists. By construction, x_t is contained between -1 (all traders are chartists) and $+1$ (all traders are fundamentalists). Expressing the population shares of the two groups in terms of this index yields[7]

$$n_t^f/2N = (1 + x_t)/2, \quad n_t^c/2N = (1 - x_t)/2 \tag{3}$$

Total (normalized) excess demand, which is thus given by $(1+x_t)\, d_t^f/2+(1-x_t)\, d_t^c/2$, will generally not balance. A market maker is assumed to absorb any excess of supply, and to serve any excess of demand from his inventory. He reacts to this disequilibrium by changing the price for the next period, where we make use of the derivation of the market impact function in Farmer and Joshi (2002, p. 152f), according to which the market maker adjusts the price with a factor $\mu > 0$ in the direction of excess demand.[8] The coefficient μ is inversely related to market liquidity, or market depth. Following common practice in models that do not further discuss the microstructure of the market, it is treated as a fixed parameter. In sum, the equation determining the price for the next period $t + 1$ may be written as

$$p_{t+1} = p_t + \frac{\mu}{2}\left[(1 + x_t)\, \phi\, (p^\star - p_t) \; + \; (1 - x_t)\, \chi\, (p_t - p_{t-1}) \; + \; \varepsilon_t \right] \tag{4}$$

[6] In actual fact, $\chi = \sigma_c = 0$ results from an estimation of the USD–DEM exchange rate; see Franke and Westerhoff (2011, Section 7). The situation for $\phi = 0$ and, possibly, $\sigma_f = 0$ would be formally analogous. In this case, however, the price dynamics would no longer be anchored on the fundamental value.

[7] To see this, define $n_t = (n_t^f - n_t^c)/2 = x_t N$, write the identity $n_t^f + n_t^c = 2N$ as $n_t^f/2 = N - n_t^c/2$ and add $n_t^f/2$ on both sides of this equation. This yields $n_t^f = N + n_t$ and, after division by $2N$, the first part of Eq. (3). The derivation of the second part is analogous.

[8] As usual in this kind of framework, any other feedbacks when his inventory continues to deviate from some target are ignored, which (in a stochastic model) is clearly an inconsistency. It could be removed by adding the risk aversion concept of the market maker (and also the other agents) studied in Franke and Asada (2009). We forgo this option to avoid blurring the central mechanisms of the model.

$$\varepsilon_t \sim N(0, \sigma_t^2), \quad \sigma_t^2 = [(1 + x_t)^2 \sigma_f^2 + (1 - x_t)^2 \sigma_c^2] / 2 \tag{5}$$

Equation (5) is derived from the fact that the sum of the two normal distributions in (1) and (2), which are to be multiplied by the market fractions $(1 \pm x_t)/2$, is again normally distributed, with mean zero and the variance being equal to the sum of the two single variances. Obviously, if σ_f^2 and σ_c^2 are different, σ_t^2 will change with the changes in the majority index x_t. The time-varying variance σ_t^2 will, in fact, be a key feature of the model. While this stochastic volatility component might be akin to a GARCH-type of modelling, we stress that it is not just a handy technical device but emerges from a structural (though parsimonious) modelling approach. The random components introduced in the formulation of the group-specific demand may therefore be said to give rise to *structural stochastic volatility* (SSV) in the returns (i.e. the log differences in prices).[9]

Before continuing, a general feature is worth pointing out. First, in a pure chartist regime, $x_t \equiv -1$, the two-dimensional price process is easily seen to have a zero and a unit root. Second, in a pure fundamentalist regime, $x_t \equiv 1$, the root of the one-dimensional price dynamics is $1 - \mu\phi$, where in estimations the product $\mu\phi$ turns out to be around 0.01 or less. Hence there is broad scope for persistent price misalignment, which is certainly a good general selling point for the model.

2.2 Evolution of the market fractions

The model is completed by setting up the motions of the majority index x_t. In light of earlier presentations in the literature (e.g. Weidlich and Haag 1983; Lux 1995), we wish to emphasize that x_t is the index actually prevailing in period t (and not some expected value; see the discussion in Appendix 2). The index is predetermined in each period, and only changes from one period to the next.[10]

The law governing the adjustments of x_t rests on the supposition that in period t all fundamentalists, whose population share is $(1 + x_t)/2$, have the same transition probability π_t^{fc} to convert to chartism, and all chartists, whose population share is $(1 - x_t)/2$, have the same probability π_t^{cf} to convert to fundamentalism. If the number of agents is sufficiently large, the intrinsic noise from different realizations when the individual agents apply their random mechanism can be neglected. So the changes in the groups are given directly by their size multiplied by the transition probabilities.

[9] Randomized demand functions of heterogeneous traders were also considered in Westerhoff and Dieci (2006) and Westerhoff (2008). The idea as such may be traced back to Westerhoff (2003). However, the implied feature of stochastic volatility and its scope for matching certain stylized facts of (daily) returns was not fully elaborated there. More on the particular effects of SSV can be learned from the investigation in Franke (2010), where this principle of heterogeneous noise was incorporated into two other model types.

[10] This is different from the discrete choice approach, which is a constituent part of the Brock–Hommes (1998) model variety. There, the population shares of the agents—and not their rates of change—are directly a function of the state variables of the model. However, introducing an asynchronous updating of strategies in the latter, it becomes essentially the same as the Weidlich–Haag–Lux approach (see the discussion in Franke 2013).

Accordingly, the population share of the fundamentalists decreases by $\pi_t^{fc}\,(1+x_t)/2$ due to the fundamentalists leaving this group, and it increases by $\pi_t^{cf}\,(1-x_t)/2$ because of the chartists who newly join this group. As a net effect, the following deterministic adjustment equation for x_t is obtained,[11]

$$x_{t+1} = x_t + (1-x_t)\,\pi_t^{cf} \;-\; (1+x_t)\,\pi_t^{fc} \tag{6}$$

As indicated by the time subscripts, the two transition probabilities are not constant. The effects determining their changes over time are summarized in a switching index $s = s_t$. An increase in s_t is supposed to increase the probability that a chartist becomes a fundamentalist, and to decrease the probability that a fundamentalist becomes a chartist. Assuming that the relative changes of π_t^{cf} and π_t^{fc} in response to the changes in s_t are linear and symmetrical, the specification of the transition probabilities reads (where 'exp' is the exponential function),[12]

$$\pi_t^{cf} = \pi^{cf}(s_t) = \nu\,\exp(s_t), \qquad \pi_t^{fc} = \pi^{fc}(s_t) = \nu\,\exp(-s_t) \tag{7}$$

Certainly, (7) ensures positive values of the probabilities. They also remain below unity if the switching index is bounded and ν is sufficiently low.[13]

A special feature of (7) is $\pi_t^{cf} = \pi_t^{fc} = \nu > 0$ in a situation $s_t = 0$. Hence even in the absence of active feedback forces in the switching index, or when the different feedback variables behind s_t neutralize each other, the individual agents will still change their strategy with a positive probability. These reversals, which can occur in either direction, are ascribed to idiosyncratic circumstances. Although they appear as purely random from a macroscopic point of view, in the aggregate they will only cancel out in a balanced state when $x_t = 0$. For nonzero values of the switching index, on the other hand, the coefficient ν measures the general responsiveness of the transition probabilities to the socio-economic aspects summarized in s_t. So ν may be generally characterized as a *flexibility parameter* (Weidlich and Haag 1983, p. 41).

The switching index itself is specified as follows,

$$s_t = s(x_t,\, p_t) := \alpha_o + a_x\, x_t + \alpha_m \cdot (p_t - p_t^{\star})^2 \tag{8}$$

The coefficient α_o can be interpreted as a *predisposition parameter*, since in a state where the other effects in (8) cancel out, a positive α_o gives rise to a probability π_t^{cf}

[11] In contrast to the more elaborate treatment in Lux (1995, 1997), this reasoning, which can also be found in Lux (1998, p. 149), is sufficient for an infinite population. A rigorous mathematical argument that begins with a finite population size and the intrinsic noise it implies is spelled out in Franke (2008a; 2008b).

[12] The precise hypothesis is $d\pi_t^{cf}/\pi_t^{cf} = \alpha\, ds_t$ and $d\pi_t^{fc}/\pi_t^{fc} = -\alpha\, ds_t$ for some constant α, which may be unity without loss of generality (since s_t may be arbitrarily scaled). Integrating these relationships with an integration constant ν yields (7).

[13] As it depends on the other parameters in the model, it is *a priori* not clear what "sufficiently low" would exactly mean. When for the numerical simulations we had to settle down on a specific positive value of ν, we checked that indeed the upper-bound of unity for π_t^{cf}, π_t^{fc} was never reached.

of switching from chartism to fundamentalism that exceeds $\nu = \nu \cdot \exp(0)$, while the reverse probability π_t^{cf} is less than ν (and vice versa for $\alpha_o < 0$).

The second term on the right-hand side of (8) captures the idea of herding. The greater the number of traders who are already fundamentalists (i.e. the higher x_t), the higher the probability that the remaining chartists will also convert to fundamentalism (and vice versa, since $x_t < 0$ if chartists are in the majority). In addition, it will be seen in the analysis below that suitable values of α_x, which may be called a *herding parameter*, can give rise to one, two or three equilibrium points of the deterministic skeleton of the model.[14]

With $\alpha_m > 0$, the third term in (8) measures the influence of misalignment, or distortion. The idea behind it also has some empirical support. It states that when the price is further away from its fundamental value, "professionals tend more and more to anticipate" its "mean-reversion towards equilibrium" (Menkhoff et al. 2009, p. 251). In our context, this means that the probability of becoming a fundamentalist rises. The underlying expectations should actually be self-fulfilling and should constitute a stabilizing mechanism, by virtue of the negative feedback in the core demand (1) of the fundamentalists.

To sum up, the two central dynamic equations of the model are (i) the price adjustments (4), (5) with the structural stochastic volatility component σ_t^2, and (ii) the changes in the majority index x_t described in (6)–(8), which basically represent a herding dynamics curbed by a control for strong price misalignment. The pivotal point of the model is that the time-varying population shares from the mechanism in (ii) feed back on the variance σ_t^2 in (i) and may therefore lead to variations in price volatility.

3 How the model functions

3.1 The deterministic skeleton

Although the structural stochastic volatility in the form of the time-varying variance in (5) is essential to the model's desired properties, it is useful to analyze the deterministic skeleton in order to understand how the model works. To this end, we first study the number of equilibrium points and their location as two of the parameters in the switching index (8) are varied. Subsequently, the nature of the resulting dynamics is sketched in phase diagrams in the (x_t, p_t)-plane. The discussion does not deal with all of the phenomena that are *a priori* possible. Instead, we concentrate on the cases that lead, step by step, to the scenario that will generate the stochastic trajectories with the desired properties.

To begin with the deterministic equilibrium points, it is clear from the market maker Eq. (4) that the price is at rest if and only if it coincides with the fundamental value $p^\star$. On the other hand, as it is typical for models employing the switching mechanism (6),

[14] There are several stories about the ways in which x_t influences the transition probabilities. If the individual agents base their switching decision on the publicly available knowledge of the current majority index, these observations might also involve some noise. We disregard this option for simplicity.

(7), the majority index can attain multiple equilibrium values. The cases of interest to us are collected in a separate proposition. Its proof is given in Appendix 3.

Proposition 1 *A stationary point of the deterministic skeleton of the dynamic system formulated in Sect. 2 is constituted by a price* $p = p^\star$*, while the following cases can be distinguished for the majority index* x*:*

(a) If the herding parameter satisfies $0 < \alpha_x < 1$*, then there exists a unique interior equilibrium value* x^o *of the majority index.*
(b) If the herding parameter exceeds unity and the predisposition parameter is zero, $\alpha_x > 1$ *and* $\alpha_o = 0$*, then there exist three equilibrium values* x^{cd}*,* x^o*,* x^{fd} *of the majority index, with* $-1 < x^{cd} < x^o < x^{fd} < 1$*. This configuration is maintained if* α_o *is moderately lowered below zero (or increased above zero).*
(c) If for given $\alpha_x > 1$ *the predisposition parameter* α_o *is sufficiently negative, then again a unique interior equilibrium value* x^{cd} *of the majority index exists, which is closer to* -1 *than the value of* x^{cd} *brought about by* $\alpha_o = 0$*.*

Clearly, the superscript *cd* for the majority index indicates a distribution of trading rules where the chartists dominate, and *fd* represents one where fundamentalism is dominant.[15] Often multiple equilibria configurations, such as that in part (b), are a good basis for interesting dynamic phenomena; in particular, because the outer equilibria typically prove to be attracting and can thus be said to describe 'bubble equilibria', i.e. a persistently bullish or bearish market, respectively (a characteristic example of this is analyzed in Lux 1995). In the present model, however, it is part (c) of the proposition with its dominance of chartist traders that will turn out to be the most promising situation for our purpose, i.e., for generating volatility clustering in the stochastic model further below.

In the next step of the analysis we turn to the deterministic motions of the market fractions of traders. We need to know in which regions of the state space the majority index rises or falls. As is easily seen from (6) to (8), the change in x depends only on the contemporaneous values of x itself and the price. Hence the movements of the majority index can be conveniently sketched in the (projection onto the) phase plane for the variables (x_t, p_t). The basic information for this is given by the isoclines $\Delta x_{t+1} = x_{t+1} - x_t = 0$, that is, the geometric locus of all pairs (x_t, p_t) on which (6)–(8) would temporarily cause x_t to come to a halt. The description of the isoclines and whether x_t increases/decreases above or below them in the plane makes use of the following function $g(\cdot)$ of the majority index,

$$g(x) := \alpha_o + \alpha_x\, x - \frac{1}{2} \ln \left[\frac{1+x}{1-x} \right] \tag{9}$$

The analytical conditions on the combinations of (x_t, p_t) under which x_t rises or falls are summarized by the next proposition. Its proof can again be found in Appendix 3.

[15] Symmetrically to point (c) in the proposition, a sufficiently positive predisposition parameter α_o would establish a unique equilibrium value of $x = x^{fd}$ where fundamentalism takes over. As has just been stated, this situation will be of no concern to us.

Table 1 Numerical benchmark parameters (rounded)

ϕ	0.198	Aggressiveness of fundamentalists in the market
χ	2.263	Aggressiveness of chartists
σ_f	0.782	Noise in fundamentalist demand
σ_c	1.851	Noise in chartist demand
μ	0.010	Market impact factor of demand
$p^\star$	0.000	Log of fundamental value
ν	0.050	Flexibility parameter in the population dynamics
α_o	−0.155	Predisposition parameter in the switching index
α_x	1.299	Herding parameter in the switching index
α_m	12.648	Misalignment parameter

Proposition 2 *(a) Suppose the majority index in a period t brings about $g(x_t) = 0$. Then $x_{t+1} > x_t$ if at the same time $p_t \neq p^\star$, and $x_{t+1} = x_t$ if p_t equals the fundamental value.*

(b) The case $g(x_t) > 0$ implies $x_{t+1} > x_t$, irrespective of the current level of the price.

(c) Suppose $g(x_t) < 0$. Then $x_{t+1} > x_t$ if either

$$p_t > p^\star + \sqrt{-g(x_t)/\alpha_m} \qquad or \qquad p_t < p^\star - \sqrt{-g(x_t)/\alpha_m}.$$

Furthermore, $x_{t+1} = x_t$ if equality prevails in these relationships, and $x_{t+1} < x_t$ if the inequality signs are reversed.

The geometric locus of the isocline $\Delta x_{t+1} = 0$ is therefore given by the equality relationship in Proposition 2(c). Deducing the properties of $g(\cdot)$ and the square root function from a general mathematical analysis would be possible but rather cumbersome and not very illustrative. On the other hand, a few numerical examples are sufficiently informative about the number of equilibria, the shape of the isocline in the phase plane, and the cases of different branches that may have to be distinguished (in the latter case we may also use the plural, isoclines). As can be seen from Proposition 2, the isocline depends on the three parameters α_o, α_x, α_m in the switching function only. For a plot of some typical trajectories, however, the other reaction coefficients are required as well. Table 1 presents a benchmark parameter scenario for this investigation. Including the standard deviations for the noise terms, it actually anticipates the result of the estimation in Sect. 4, where the underlying time unit is one day. Of course, the values $p^\star = 0$ and $\mu = 0.010$ are just a matter of scaling, and for the present analysis of the deterministic model we put $\sigma_f = \sigma_c = 0$.

Regarding the role of the coefficients α_o, α_x, α_m, let us first consider the herding parameter α_x. This is best done by abstracting from a possible predisposition towards chartism or fundamentalism. So for the moment being we set α_o equal to zero, adopt the other parameter values (except σ_f, σ_c) from Table 1, and plot the equilibria and isoclines from Propositions 1 and 2 for selected values of α_x in Fig. 1.

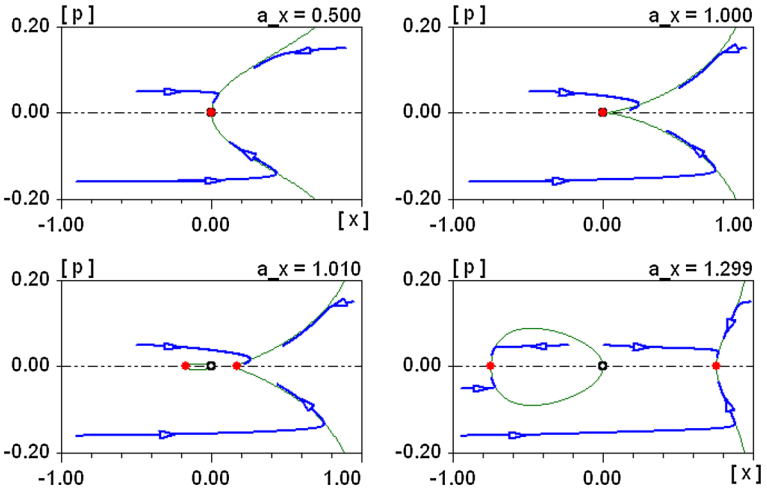

Fig. 1 Phase diagrams of the deterministic skeleton under *ceteris paribus* variations of the herding parameter α_x. *Note*: $\alpha_o = \sigma_f = \sigma_c = 0$, other parameters from Table 1. *Thin* (*green*) *solid lines* are the isoclines $\Delta x_{t+1} = 0$

The upper-left panel shows the outcome for a relatively low level of herding, $\alpha_x = 0.50$. Here, as stated in Proposition 1(a), we have a unique equilibrium $(x^\star, p^\star)$, which by virtue of $\alpha_o = 0$ is given by $(x^\star, p^\star) = (0, 0)$ and which is globally attracting.[16] The isocline $\Delta x_{t+1} = 0$, the thin (green) solid line, divides the plane into two regions in which x_t increases and decreases, respectively. The sample trajectories, the bold (blue) lines, indicate that left to the isocline the majority index rises so fast relative to the price that the motion is almost horizontal. We emphasize that such a phase of temporarily strong herding in the convergence process is a universal phenomenon in the model; we find it for practically all parameter combinations that are of any relevance. On the price side, the main reason for it is the relatively low value of ϕ in comparison with χ, which limits the mean-reverting tendencies from the fundamentalist strategy. But once again, this only applies in a part of the phase plane.

It may also be noted in the first panel that near the equilibrium the upper branch of the isocline is a concave function, and it eventually becomes convex for x sufficiently high (of course, the lower branch is symmetric to this). A *ceteris paribus* increase of α_x shifts the point of inflection closer and closer to the equilibrium, until at $\alpha_x = 1.00$ the entire branch is a convex function. This case is illustrated in the upper-right panel in Fig. 1.

According to Proposition 1(b), a qualitative change occurs when now α_x rises above unity. In this way the previously stable equilibrium $(x^\star, p^\star) = (0, 0)$ becomes unstable (indicated by the empty dot) and two new equilibria arise symmetrically to its left and its right, which are locally stable (indicated by the filled dots). The lower-left panel in Fig. 1 illustrates the situation for the minimal increase up to $\alpha_x = 1.010$. The isocline $\Delta x_{t+1} = 0$ from the upper two panels remains qualitatively the same, except that it

[16] A mathematical proof is omitted.

is no longer anchored in $(x^\star, p^\star) = (0, 0)$ but in the equilibrium $(x^{fd}, p^\star) \geq (0, 0)$ (and there is no symmetry in that no isocline runs through the opposite equilibrium). The fundamentalist equilibrium attracts the great majority of all motions, while the basin of attraction of its chartist counterpart $(x^{cd}, p^\star)$ is so small that none of our three sample trajectories happens to converge to it.

Since $(x^{cd}, p^\star)$ is so close to the inner equilibrium $(0, p^\star)$, we cannot see what happens in the small region between the two. This dynamics becomes clear when α_x is increased to our benchmark value $\alpha_x = 1.299$ in the lower-right panel, the reaction being that the (unstable) inner equilibrium stays put and the other two (stable) equilibria shift to the outside. In this way a second part of the isocline grows and forms a lens between x^{cd} and $x^\star = 0$, within which the market converges to the chartist equilibrium (the lens is already present, though hardly visible, in the lower-left panel).

We may so far distinguish between weak and strong herding; weak herding is constituted by $\alpha_x \leq 1$ in the upper two panels in Fig. 1 with their unique equilibrium at $(x^\star, p^\star) = (0, 0)$, and strong herding prevails for $\alpha_x > 1$ in the lower two panels. The latter brings about two additional bubble equilibria, where the higher α_x, the larger the corresponding majority of fundamentalists or chartists, and the broader the scope for convergence towards $(x^{cd}, p^\star)$ by increasing the lens just mentioned.

Interestingly, the estimation suggests strong herding. However, besides fixing the herding coefficient at $\alpha_x = 1.299$, it also advises us to decrease the predisposition parameter α_o below zero. Let us see in Fig. 2 what this means for the isoclines and equilibria.

To begin with, the top-left panel reproduces the situation $\alpha_x = 1.299$ and $\alpha_o = 0.000$ from Fig. 1. The effect of a moderate decrease in α_o to $\alpha_o = -0.10$, which represents a moderate predisposition towards chartism, is that the $\Delta x_{t+1} = 0$ isoclines in the left and right half of the plane move towards each other; see the top-right panel in Fig. 2. In particular, the inner equilibrium is no longer fixed but moves to the right, too. Nevertheless, the trajectories remain largely unaffected. It requires a stronger bias towards chartism (a stronger fall of α_o) for the system to undergo a structural change, such that in line with Proposition 1(c) the fundamentalist equilibrium disappears. Geometrically, when α_o further declines below -0.10, the two equilibria $(x^o, p^\star)$ and $(x^{fd}, p^\star)$ first collapse into a single point and then dissolve, so that the separate two original isoclines are now connected. This has happened in the middle-left panel, where α_o attains the value of the benchmark scenario from Table 1, $\alpha_o = -0.155$.

Here the chartist equilibrium $(x^{cd}, p^\star)$ is not only unique but also globally stable. This derives from the fact that the price increases (decreases) if $p_t < p^\star$ (if $p_t > p^\star$); that the majority index x_t decreases if the system is inside the region bounded by the upper and lower branch of the isocline; and that eventually every trajectory will enter this region (which can also be algebraically verified). Moreover, as already observed in Fig. 1, farther away from the isocline the price reactions are so slow relative to the strategy changes that the motions of (x_t, p_t) trace out almost horizontal lines.

The trajectory starting in the lower-left corner of the middle-left panel illustrates the stabilizing force of the misalignment component in the switching mechanism (represented by the parameter α_m in (8)). Due to the strong initial misalignment, the market first moves straight into the fundamentalist region. However, there is no more fundamentalist equilibrium towards which it could converge or around which it could

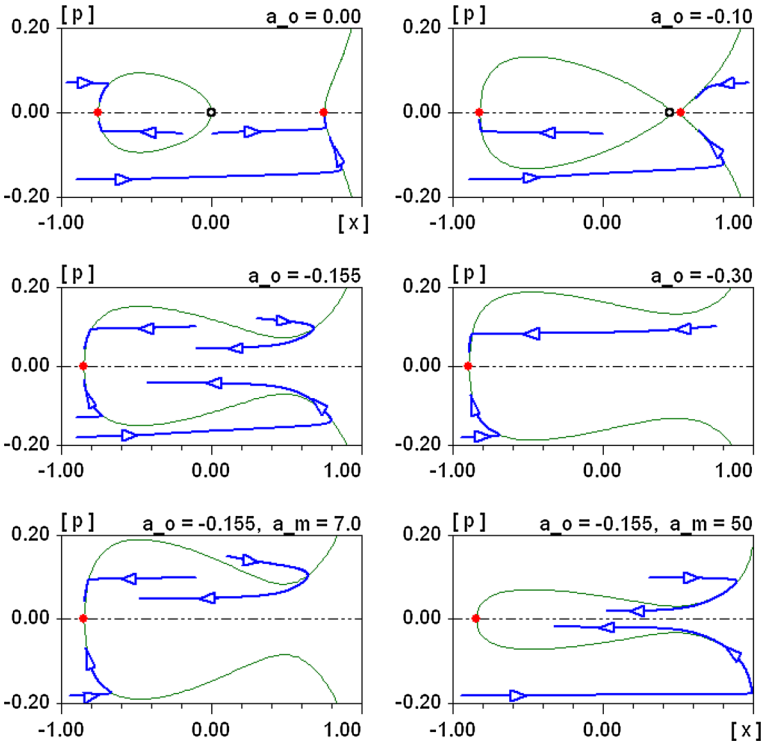

Fig. 2 Phase diagrams of the deterministic skeleton under *ceteris paribus* variations of α_o and α_m. *Note*: $\sigma_f = \sigma_c = 0$, other parameters (in particular, $\alpha_x = 1.299$) from Table 1

fluctuate. Hence, sooner or later such a trajectory must return to the chartist region. On this path, the switches in strategy will again be relatively fast once the trajectory disconnects from the isocline in the local maximum (minimum) in the lower (upper) half of the phase plane. Now the price misalignment is of secondary importance, and the herding mechanism reinforced by the predisposition effect (the behavioural bias towards chartism) re-establishes a chartist regime.

The main features of the $\Delta x_{t+1} = 0$ isocline are maintained under the parameter variations considered in the remaining three panels of Fig. 2. As shown in the middle-right panel, it makes good sense that a stronger predisposition towards chartism (a further *ceteris paribus* decrease in α_o) enlarges the region where convergence takes the form of a declining x_t, i.e. where the market fraction of the chartists steadily increases. Likewise, a weaker or stronger influence of price misalignment (lower or higher values of the coefficient α_m in the lower two panels, with α_o reset to -0.155)) widen or narrow, respectively, this region in the phase space with its dominance of the herding mechanism.

In sum, the three parameters $\alpha_x, \alpha_o, \alpha_m$ fulfil the following tasks: a sufficiently strong herding α_x brings the two bubble equilibria into existence; a sufficiently strong predisposition towards chartism (α_o sufficiently negative) lets the fundamentalist as well as the inner equilibrium disappear; and the aversion α_m against the risk of price

misalignment governs the curvature of the Δx_{t+1} isocline. The latter becomes important for fine-tuning the volatility clustering in the stochastic dynamics in the next subsection. This extension will also qualify the significance of the remaining, globally stable chartist equilibrium in the deterministic setting; there will still be sufficient scope for a temporary fundamentalist regime.

3.2 The stochastic dynamics

Let us now study the full model that includes the daily random perturbations to the price. The numerical parameters are those from Table 1. On the basis of the deterministic dynamics in the middle-left panel of Fig. 2, a first and immediate idea might be that not many interesting things can happen here since the market will eventually settle down in a region around the unique and globally stable chartist equilibrium. While the general noise σ_t^2 in the system would perhaps be high, the variations of the resulting volatility of the returns would be rather limited, leaving not much room for long memory effects or a non-normal distribution of the returns. This reasoning, however, does not take into account that a sequence of the random shocks ε_t in (4) may cause the system to jump across the $\Delta x_{t+1} = 0$ isocline. If this happens at a stage where x_t has declined towards the chartist equilibrium value and the noise level σ_t^2 from (5) has increased accordingly, the motion would be reversed towards fundamentalism and σ_t^2 may even systematically decline again for a while.

In order to check whether events of this type might be able to lead to significant clusters of low and high volatility, the model has to be simulated. The first three panels in Fig. 3 present a sample run over 6,867 days. These roughly 27 years cover the same time span as the empirical returns from the S&P 500 stock market index, which is plotted in the bottom panel.[17]

The top panel in the figure illustrates the model-generated fluctuations of the (log) price around the fundamental value $p^\star = 0$. They clearly reproduce the informal stylized fact of fairly long and irregular swings with a considerable amplitude. The second panel displays the corresponding composition of the traders in the form of the market share of chartists, $n_t^c/2N = (1 - x_t)/2$ as stated in (3). It shows that the market is ruled by the fundamentalists most of the time. Every now and then, however, a relatively rapid motion to a chartist regime is observed. Normally these regimes do not last very long, although there are exceptions where chartists are in the majority for even more than one year (roughly 300 days from $t = 3{,}450$ onward). The conditions for these features to occur will become clearer from the discussion of Fig. 4.

Comparing the upper two panels in Fig. 3, it can be seen that fundamentalists take over in the presence of stronger mispricing, and chartists only gain ground when the price returns to the fundamental benchmark. This phenomenon is easily explained by the term $\alpha_m\,(p_t - p_t^\star)^2$ in the switching index s_t in (8), higher values of which increase the probability that the agents convert to fundamentalism rather than to chartism. In

[17] Reckoning 250 days per year. Specifically, the empirical sample period is January 1980 to March 2007 (just before the financial crisis began to unfold).

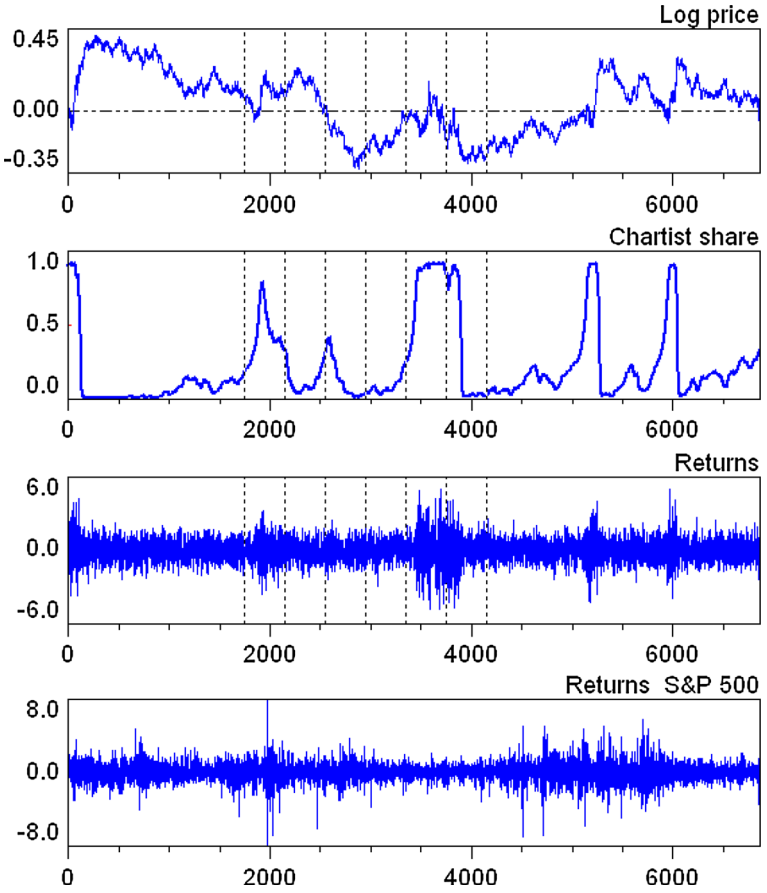

Fig. 3 Sample run of the model and empirical daily returns. *Note*: Numerical parameters from Table 1. *Vertical dotted lines* indicate the subperiods shown in Fig. 4

combination with the other parameters, $\alpha_m \approx 12$ is high enough for this mechanism to become effective.

The third panel in Fig. 3 demonstrates the implications of the irregular regime switches for the returns r_t, which are specified in percentage points,

$$r_t := 100 \cdot (p_t - p_{t-1}) \tag{10}$$

Owing to the greater variability in chartist demand *vis-à-vis* fundamentalist demand, $\sigma_c^2 > \sigma_f^2$ in (1), (2) or (4), (5), respectively, the noise level in the returns during a chartist regime exceeds the level in a fundamentalist regime. Since the fundamentalists dominate the market over longer periods of time, it looks as if a certain "normal" noise in the returns is occasionally interrupted by outbursts of increased volatility. In other words, the pattern in the evolution of the simulated returns can indeed be characterized as volatility clustering.

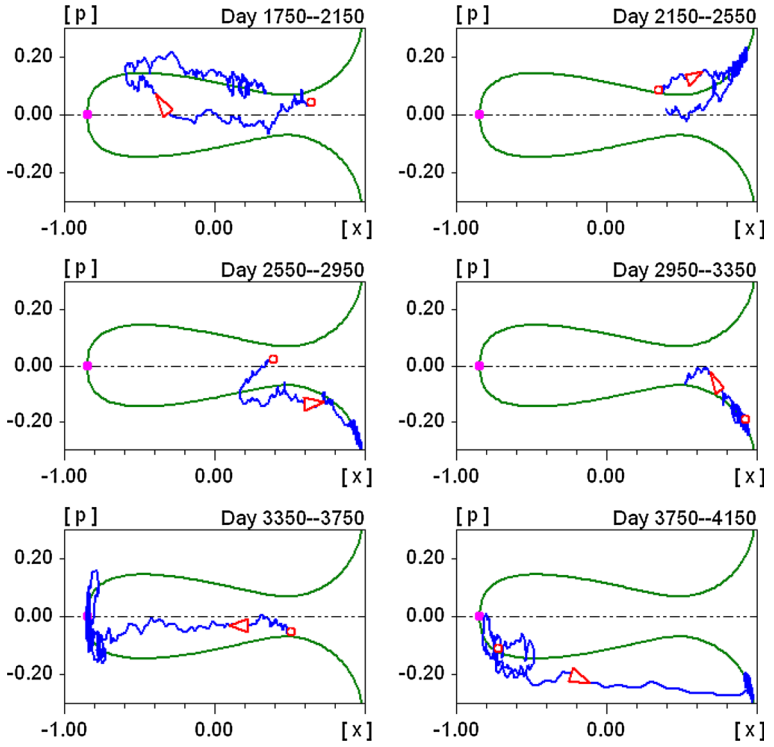

Fig. 4 Subperiods of sample run from Fig. 3 in the phase plane. *Note*: As indicated by the (*red*) *empty dots*, panel 1 (*top-left*) starts from $(x, p) = (0.64, 0.036)$, panel 2 (*top-right*) from $(0.34, 0.086)$, panel 3 (*middle-left*) from $(0.39, 0.018)$, panel 4 from $(0.92, -0.205)$, panel 5 from $(0.51, -0.056)$, and panel 6 from $(-0.73, -0.119)$

The bottom panel in the diagram displays the daily returns from the S&P 500 over the same time horizon. A comparison with the third panel shows that the qualitative pattern of the alternation of periods of tranquillity and volatility in the returns is similar for the simulated and empirical series. Also the quantitative outbursts are comparable in size (note that the two panels do not have the same scale). Differences can be seen in the band width of the returns in the periods of relative tranquillity. While the noise level is then constant in the simulated series, the empirical series exhibits certain changes from the first, say, 1,800 days of the sample to the period between $t = 3{,}000$ and $t = 4{,}000$, where the band becomes narrower, and from there to the end of the series, where the band again widens somewhat. Obviously, a simple model cannot easily endogenize these more refined 'regime shifts', if they were found to be significant at all.

To obtain a better understanding of what we observe in the time series diagrams, let us follow the dynamic evolution of the market over six consecutive subperiods in the phase diagrams of Fig. 4. These periods are indicated by the vertical dotted lines in Fig. 3. The $\Delta x_{t+1} = 0$ isocline is reproduced from the middle-left panel in Fig. 2, but the vertical price axis now covers a wider range.

The discussion of Fig. 4 begins at $t = 1{,}750$, when the system is at $(x_t, p_t) = (0.64, 0.036)$ and the chartist share amounts to 18 per cent. The system remains in the inner region bounded by the two branches of the $\Delta x_{t+1} = 0$ isocline and, in herding towards the chartist equilibrium, hovers around the fundamental value for more than one hundred days. Then the shocks start to shift the market to the upper isocline. Eventually, after 8.5 months at $t = 1{,}927$, the market crosses it—at a time when the market fraction of the chartists has risen to almost 80 per cent. From then on, the trajectory (essentially) stays above the isocline for the next few hundred days, and the misalignment mechanism in the switching index leads the market back to a fundamentalist regime. Note that it nevertheless takes a while until the chartist share falls again below values of, say, 20 or 10 per cent.

The second panel in Fig. 4 sets in at $t = 2{,}150$; its starting point at $(x_t, p_t) = (0.34, 0.086)$ is the final point in the first panel. From here, the system moves up the isocline, and after about half of the second subperiod it returns into the inner region. In this episode, the speculators' herding towards fundamentalism was first reinforced by the misalignment term, while with the ensuing stabilization, i.e. reduction in the mispricing, the fundamentalist regime eased off somewhat. In fact, at the end, around $t = 2{,}550$, the system is close to the situation where it had started from in the first panel. The third (middle-left) panel, however, shows that this time the dynamics leaves the inner region much earlier and downwards across the lower isocline, from which time on the price remains below the fundamental value. Consequently, the dynamics re-enters a pronounced fundamentalist regime. At the end of the third and for most of the fourth subperiod, it crawls up and down the outer lower branch of the isocline in the lower-right corner of the two panels.

At the end of the fourth subperiod, from approximately $t = 3{,}290$ on, the system continues to stay in the inner region, where we also find the starting point of the fifth subperiod. Although it is close to the boundary, it does not cross it once again. Instead, within 120 days until $t = 3{,}470$, the system relatively quickly builds up a chartist majority. Since strong shocks happen to be absent then, the deterministic stability of the chartist equilibrium continues to work out and the chartist share stays between 85 and 92 per cent. Correspondingly, at this stage the market fluctuates up and down the steep part of the Δx_{t+1}-isocline. At the end of the fifth and the beginning of the sixth subperiod, the trajectory moves slightly to the right in the phase diagrams, then for a short while returns to the chartist equilibrium, until finally the shocks drive the price so low that the market rushes towards the fundamentalist regime in the lower-right corner in the sixth phase diagram.

To summarize this discussion, the deterministic structure of the model with, in particular, the three coefficients $\alpha_x, \alpha_o, \alpha_m$ establishes the nonlinear $\Delta x_{t+1} = 0$ isocline, which serves to see in which subregions of the phase space the market share of the chartists systematically increases and decreases. The random forces are, however, strong enough to lead the dynamics towards and across the isocline. On the other hand, they are not strong enough to let the market permanently fluctuate back and forth near this geometric locus. Occasionally, the deterministic core of the model becomes dominant, that is, the market remains on one side of the isocline for a longer time, implying that it changes from a more or less fundamentalist regime to a chartist regime, or vice versa.

On the whole, the present numerical scenario renders these mechanisms so effective that we obtain the volatility clustering of the temporary chartist markets demonstrated in Fig. 3. We may furthermore expect that this pattern of the returns gives rise to a non-normal distribution or fat tails, respectively. This is certainly a qualitatively satisfactory result. In the next section, we must make sure that the usual summary statistics describing these phenomena also match their empirical counterparts in a quantitatively satisfactory manner.

4 Estimation of the model

This section is devoted to a rigorous estimation of the model by the method of simulated moments (MSM). The first subsection begins with a recapitulation of the MSM approach, explaining its minimization of the quadratic distance between certain model-generated and empirical summary statistics, i.e. "moments". Subsequently, two specific problems will be addressed: (a) the determination of the weighting matrix for the moments by a more suitable (nonparametric) bootstrap procedure than the usual block bootstrap; (b) the sample variability in the model's stochastic simulations, which we propose to straighten out by the concept of a "representative estimation".

A second subsection introduces another (parametric) bootstrap. It generates as many artificial moments as we want in order to re-estimate the model on them. From the frequency distribution of the thus minimized values of the objective function, a measure will then be derived (actually a p value) that can serve for an overall evaluation of the model's goodness-of-fit. At the same time, a (marginal) distribution for each of the re-estimated parameters is obtained, by which we can assess the precision of the original estimation. The subsection is concluded with a brief discussion on how to assess the "dominance" of the stochastic noise as it is implied by the estimation.

4.1 The method of simulated moments

The model has been designed to explain—at least partially—the most important stylized facts of financial markets.[18] Referring to the price changes at daily intervals, we aim to check the four features that have received the most attention in the literature on agent-based models. These are the absence of autocorrelations in the raw returns, fat tails in their frequency distributions, volatility clustering, and long memory (see Chen et al. 2012).[19] For the quantitative analysis, we measure these features by a number of summary statistics or, synonymously, moments. The first moment is the volatility of the returns, which we define as the mean value of the absolute returns $v_t = |r_t|$ (here and in the autocorrelations below it makes no great difference whether one works with the absolute or squared returns). Reproducing it is basically a matter of scaling, and in the first instance it should have a bearing on the admissible general noise level in

[18] Detailed descriptions of the statistical properties of asset prices can be found in Cont (2001), Lux and Ausloos (2002), or Lux (2009b).

[19] Generally, one might also include a negative skewness of stock returns. Stylized small-scale asset pricing models, such as the present one, do not, however, provide for any asymmetry in this respect.

the model, as it is brought about by the two variances σ_f^2 and σ_c^2. The second moment is the first-order autocorrelation of the raw returns. The requirement that it be close to zero should balance the reaction intensities of the chartists and fundamentalists in the form of the parameters χ and ϕ (as χ is conducive to positive and ϕ to negative autocorrelations). On the other hand, we checked that if this moment is matched, the autocorrelations at the longer lags will practically all vanish, too. Because of this lack of additional information, it suffices to make use of only one moment of the raw returns.

Next, in order to capture the long memory effects, we invoke the autocorrelation function (ACF) of the absolute returns v_t up to a lag of 100 days. As the ACF slowly decays without becoming insignificant at the long lags, we have an entire profile to match. We view it as being sufficiently well represented by the six coefficients for the lags $\tau = 1, 5, 10, 25, 50, 100$. The influence of accidental outliers that may occur here is reduced by using the centred three-lag averages.[20] Lastly, the fat tail property is measured by the well-known Hill estimator of the tail index of the absolute returns, where the tail is conveniently specified as the upper 5 per cent. Thus, on the whole, we evaluate the performance of the model on the basis of nine moments, which we collect in a (column) vector $m = (m_1, \ldots, m_9)'$ (the prime denotes transposition).

It has already been indicated that the simulated moments from the model should be as close as possible to the empirical moments that we compute for the daily returns of the S&P 500 stock market index. To make the informal summary of "fairly close" more precise in a formal estimation procedure, it is only natural for us to employ the method of simulated moments (MSM). To this end, an objective function, or loss function, has to be set up that defines a distance between two moment vectors. It is given by a quadratic function, which is characterized by a weighting matrix $W \in \mathbb{R}^{9\times 9}$ (to be specified shortly). Considering the general situation where a moment vector $m \in \mathbb{R}^9$ is to be compared to another set of reference moments $m^{ref} \in \mathbb{R}^9$, the function reads,

$$J = J(m, m^{ref}) := (m - m^{ref})' \, W \, (m - m^{ref}) \tag{11}$$

The weighting matrix takes the sampling variability of the moments into account. The basic idea is that the higher the sampling variability of a given moment i, the larger the differences between m_i and m_i^{ref} that can still be deemed insignificant. The loss function can account for such a higher tolerance by correspondingly smaller diagonal elements w_{ii}. In addition, matrix W should provide for possible correlations between the single moments. These two tasks are fulfilled by specifying the weighting matrix as the inverse of an estimated variance-covariance matrix $\widehat{\Sigma}$ of the moments,

$$W = \widehat{\Sigma}^{-1} \tag{12}$$

An obvious, since asymptotically optimal, choice for W would be the inverse of a Newey-West estimator of the long-run covariance matrix of the empirical moments

[20] That is, at lag τ the mean of the three autocorrelation coefficients for $\tau - 1, \tau, \tau + 1$ is computed, except for $\tau = 1$, where it is the average of the first and second coefficient. It may also be noted that volatility clustering, which describes the tendency of large changes in the asset price to be followed by large changes, and small changes to be followed by small changes, is closely related to these long-term dependencies between the returns.

(see, e.g., Lee and Ingram 1991, p. 202, or the application of MSM in Franke 2009, Sect. 2.2). Optimality, however, does not necessarily carry over to small samples.[21] We therefore choose a bootstrap procedure to construct, from the empirical observations of length T, additional samples of the same size and derive the covariances in $\widehat{\Sigma}$ from them. We nevertheless depart from the block bootstraps that have been used in Winker et al. (2007) or Franke and Westerhoff (2011, 2012b), since the original long-range dependence in the return series is interrupted every time two non-adjacent blocks are pasted. The fact that our estimation is concerned with summary statistics and not the one-period ahead predictions of a time series allows us to sample the single days t_k ($k = 1, \ldots, T$) and, associated with each of them, the history of the past few lags required to calculate term t_k in the formula for the lagged autocorrelations. Avoiding thus the join-point problem, this alternative seems more trustworthy than a block bootstrap (see Appendix 4 for details).

The bootstrap gives us a collection of $b = 1, \ldots, B$ values for each of the nine moments, where $B = 5{,}000$ is large enough (indices b may be identified with the random seed for the sequence of the (pseudo-)random numbers that set up the single bootstrap samples). Letting $m^b = (m_1^b, \ldots, m_9^b)'$ be the corresponding moment vectors and computing the vector of their mean values $\overline{m} := (1/B) \sum_b m^b$, the bootstrap estimate of the moment covariance matrix $\widehat{\Sigma}$ in (12) is given by

$$\widehat{\Sigma} = \frac{1}{B} \sum_{b=1}^{B} (m^b - \overline{m})(m^b - \overline{m})' \tag{13}$$

We are now ready to turn to the estimation problem.[22] With respect to $T = 6{,}866$, the length of the empirical sample of the returns, denote the moments computed from it by m_T^{emp}. Let θ be the vector of the model parameters to be estimated. While they are generally contained in a certain set, beginning with possible nonnegativity constraints, we can omit an explicit reference to it since no estimated values or their confidence intervals will have any problem in this respect. MSM, then, means finding a parameter vector θ such that the simulated moments to which it gives rise minimize the loss function.

To limit the variability in the stochastic simulations, their sample size, designated S, should be appreciably larger than the number of the empirical observations T, where $S/T = 10$ is a common proportion (S is the effective simulation size, after discarding the first few hundred days to rule out any transient effects). Furthermore, the comparability of different trials of θ requires them to have the same random number sequence underlying.[23] The latter are determined by a random seed, which we generally identify by an integer number, such as $a = 1, 2, \ldots$, let us say. Thus, the moment vector obtained by simulating the model with a parameter vector θ over

[21] To reduce the thus arising bias, even the identity matrix could be a superior weighting matrix; see Altonji and Segal (1996).

[22] We checked that the weighting matrix resulting from our bootstrap procedure is indeed positive definite.

[23] For the normally distributed ε_t with variance σ_t^2 in (4), (5), this means, more precisely, that for each simulation run at time t the same random number $\tilde{\varepsilon}_t$ is drawn from the standard normal distribution $N(0, 1)$ and ε_t is set as $\sigma_t \tilde{\varepsilon}_t$.

S periods on the basis of a random seed a is denoted as $m^a(\theta; S)$. The parameter estimates based on this random seed a read $\widehat{\theta}^a$, and are the solution of the following minimization problem,[24]

$$\widehat{\theta}^a = \arg\min_{\theta} \; J[m^a(\theta; S), m_T^{emp}], \quad S = 10 \cdot T \tag{14}$$

The fundamental value $p^\star$ and the market impact factor μ are two parameters in the model that just serve scaling purposes. We exogenously fix them at $p^\star = 0$ and $\mu = 0.010$. The flexibility parameter ν approximately scales the switching index s_t (this would be exact if $\exp(\cdot)$ were a linear function). Given the interpretation of ν in the remark on Eq. (7) as an 'autonomous' switching probability, its value should be distinctly below unity. Here we choose $\nu = 0.050$, which says that in the hypothetical absence of predisposition and any other influences, an agent would on average change his strategy every 20 days, i.e. every month.[25] On the whole, there are thus seven parameters left to estimate.

Although it might seem that a simulation over $S = 68{,}660$ days generates a large sample to base the moments on, the variability arising from such different samples still turns out to be considerable. Hence it would not be pertinent to pick out an arbitrary random seed and present the corresponding results. This way, we may simply be lucky or unlucky and obtain a particularly good or bad match. Therefore, when for a succinct estimation summary we will have to settle down on a specific parameter set, the loss J it produces should be more or less 'representative', in the sense of an expected value.

To this end, it seems most appropriate to carry out a great number of estimations and choose the one with an average loss. In detail, 1,000 estimations will suffice. We then select the parameter set $\widehat{\theta}$, the associated loss of which is the median value of the entire distribution of the 1,000 estimated losses. This outcome may be viewed as our "representative" estimation, an idea that is apparently new in the literature. Formally, with reference to (14),

$$\begin{aligned} \widehat{\theta} &= \widehat{\theta}^{\tilde{a}}, \qquad \text{where } \tilde{a} \text{ is such that } \widehat{J}^{\tilde{a}} \text{ is the median of } \{\widehat{J}^a\}_{a=1}^{1000}, \text{ and} \\ \widehat{J}^a &= J[m^a(\widehat{\theta}^a; S), m_T^{emp}], \qquad a = 1, \ldots, 1{,}000 \end{aligned} \tag{15}$$

The parameter vector $\widehat{\theta}$ resulting from this battery of estimations has already been reported in Table 1. For convenience, it is reproduced in the first row of Table 2. The corresponding minimized loss amounts to 7.28,[26]

$$\widehat{J} := J[\, m^{\tilde{a}}(\widehat{\theta}; S), m_T^{emp} \,] = 7.28 \tag{16}$$

[24] We use the Nelder-Mead simplex search algorithm (see Press 1986, pp. 289–293) and restart it upon convergence several times until no further noteworthy improvement in the minimization occurs.

[25] Admittedly, the value $\nu = 0.57$ in Franke and Westerhoff (2011) is psychologically not very convincing.

[26] This value can be slightly reduced to $\widehat{J} = 6.98$ by treating ν as a free parameter, too. We then get a higher value $\nu = 0.067$ which, however, is something that we had sought to avoid. Besides, given the random seed $\tilde{a}$, a marginal improvement, $J = 7.16$, can also be obtained by a lower value of the flexibility parameter, $\nu = 0.033$.

4.2 Evaluation of the estimation results

As such, the figure in Eq. (16) is not very informative. To put it into perspective, whether it indicates a good or a bad overall match of the moments, we make use of another bootstrap procedure. It is a parametric bootstrap, which means we work with the null hypothesis that there is a parameter vector θ^o for which the model is a true description of the aspects of the stock market summarized by our moments. In other words, the moments simulated with θ^o over an horizon $S = 10 \cdot T$ are assumed to be drawn from the same distribution as the data in the real world. Naturally, the true parameter vector θ^o is proxied by the estimated, "representative" vector $\widehat{\theta}$ from (15).[27]

Nevertheless, the null hypothesis allows us to produce as many returns series of an empirical length T and artificial moment vectors as we like—and to re-estimate the model on them. In this way, we obtain an entire distribution of minimized losses, to which we can then compare our benchmark value $\widehat{J}$ from (16). If the null applies and the empirical moments, too, could therefore have been generated by the model, $\widehat{J}$ should be in the range of that loss distribution. Conversely, the null has to be rejected, and it must be concluded that the model is definitely incompatible with the data at a 5 % significance level, if $\widehat{J}$ exceeds the 95 % quantile of the distribution.

In detail, take the estimated parameter vector $\widehat{\theta}$, consider $c = 1, \ldots, 1{,}000$ different random seeds, simulate the model over the empirical time horizon for each of them, compute the moments $m^c(\widehat{\theta}; T)$ from these series, and then re-estimate the model on the latter.[28] These MSM estimations are carried out on the basis of different random seeds $d = 1, \ldots, 1{,}000$, one such d for each artificial sample $m^c(\widehat{\theta}; T)$. This procedure provides us with a distribution of estimated parameters $\widehat{\theta}^d$ and their losses $\widehat{J}^d$,

$$\widehat{\theta}^d = \arg\min_{\theta} J[m^d(\theta; S), m^c(\widehat{\theta}; T)], \qquad (c, d) = 1, \ldots, 1{,}000 \tag{17}$$

$$\widehat{J}^d = J[m^d(\widehat{\theta}^d; S), m^c(\widehat{\theta}; T)] \tag{18}$$

where, with a slight slip in notation, the pairs (c, d) are also referred to by the integers $1, \ldots, 1{,}000$. The critical value for our test of the model's goodness-of-fit is the 95 % quantile of the loss distribution $\{\widehat{J}^d\}_{d=1}^{1000}$, which results as $J_{0.95} = 13.23$. Since $\widehat{J}$ from (16) falls short of it we fail to reject the null hypothesis, even by a wide margin as it seems.

We can take a small step further than the reject-or-not decision and put forward a quantitative evaluation of the model. This is readily done by deriving a p value from

[27] In Franke and Westerhoff (2011, 2012b), a nonparametric bootstrap was employed. There we also discussed statistical measures that could characterize the matching of the single moments.

[28] To perfectly imitate the original estimation, one would also have to take into account that different return series r_t^c (in obvious notation) give rise to different weighting matrices in the loss function. Unfortunately, this would mean carrying out an extra bootstrap for each of the 1,000 artificial samples. We refrain from this additional computational effort and employ the original weighting matrix W from (12), (13) for all of the re-estimations.

Table 2 Estimation results (rounded)

	ϕ	χ	σ_f	σ_c	α_o	α_x	α_m	p
Est.	0.198	2.263	0.782	1.851	−0.155	1.299	12.65	17.3
Lower	0.145	1.621	0.737	1.531	−0.194	1.265	7.97	**8.7**
Upper	0.239	2.571	0.837	2.119	−0.132	1.498	15.36	**32.6**

Exogenously fixed are $\mu = 0.010$, $p^\star = 0$, $\nu = 0.050$. The first row is the 'representative' estimation (15), with the p value from (19) (all p values in per cent). The two bottom rows indicate the 95 % confidence intervals for the distributions of $\widehat{\theta}^d$ in (17) and p^a in (20); the Hall percentile intervals for the former (as explained in Appendix 5) and the standard percentile intervals for the latter. Bold face figures summarize the overall model evaluation

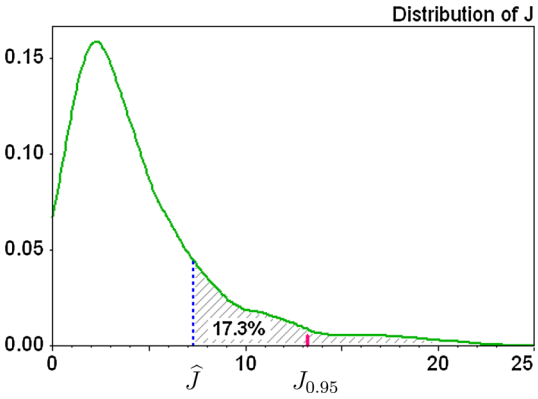

Fig. 5 Distribution $\{\widehat{J}^d\}$ from (18), its 95 % quantile $J_{0.95}$, and the estimated $\widehat{J}$ from (16)

the loss distribution $\{\widehat{J}^d\}$.[29] With respect to the estimated loss in (16), it is given by

$$p\text{value} = \text{solution of } \left\{(1-p) \text{ quantile of } \{\widehat{J}^d\} = \widehat{J}\right\} \tag{19}$$

This statistic says that if $\widehat{J}$ were employed as a benchmark for model rejection, then p is the error rate of falsely rejecting the null hypothesis that the model is true. Thus, if the p value exceeds the 5 % level, it gives us an impression of the width of the margin by which we fail to reject the null. Incidentally, it is also a particularly useful measure if there are several models to compare. As reported by the last entry in the first row of Table 2, we compute a p value of 17.3 % for the present model. Figure 5 illustrates the concept with the additional information about the 95 % quantile of the loss distribution $\{\widehat{J}^d\}$.[30]

While the 17.3 % error rate evaluates the model's goodness-of-fit as it emerges from our representative estimation, the same concept can be applied to the other losses $\widehat{J}^a$ from the original estimations on the empirical moments in (15). In this way, we also

[29] Concerning symbol p, there should be no confusion with the log prices p_t, which by now will have disappeared from the scene.

[30] The density functions in this and the next diagram are estimated using the Epanechnikov kernel; see Davidson and MacKinnon (2004, pp. 678–683) for the computational details.

obtain an entire distribution $\{p^a\}$ of p values,

$$p^a = \text{solution of } \left\{(1-p) \text{ quantile of } \{\widehat{J}^d\} = \widehat{J}^a\right\}, \qquad a = 1, \ldots, 1{,}000 \quad (20)$$

A 95 % standard percentile interval gives us a reliable range over which, owing to the small-sample variability in the simulations for the MSM estimations, the p values can vary; the upper and lower boundary are reported in the last column of Table 2. In particular, the 2.5 % quantile of $\{p^a\}$, $p = 8.7\,\%$, is a very conservative measure of the model's ability to generate the desired stylized facts. Still, even that value exceeds the critical 5 % level.[31] How much the range of the p values in (20) could be narrowed by adopting a larger simulation size S might be left for future research.[32]

In concluding our investigation of the model's general goodness-of-fit, it may be recalled that the positive evaluation at which we arrived is conditional on the specific choice of the moments the model is desired to match. Certainly, if more and qualitatively different moments were added to the present list, for which (at least intentionally) the model was not designed, the p values will dwindle and eventually lead to a rejection.

In a last step, we wish to assess the precision of our representative parameter vector $\widehat{\theta}$ in (15). Standard errors for its components can be derived from the diagonal elements of the covariance matrix of the parameters as it results from the asymptotic econometric theory.[33] However, due to the considerable small-sample variability in our estimations (as evidenced by the relatively wide range of p values), this approach may perhaps not be wholly credible. On the other hand, we already have a distribution of 1,000 parameters from our bootstrap procedures, namely, the distribution $\{\widehat{\theta}^d\}$ that we obtain from the re-estimations in (17) under the null hypothesis of a true model.[34] They readily provide us with confidence intervals for the single parameters.

Figure 6 shows the frequency distributions of the seven single components $\widehat{\theta}_i^d$, where the shaded area indicates the probability mass of the standard percentile confidence intervals, the lower and upper bounds of which are given by the 2.5 and 97.5 % quantiles. It is immediately apparent that all of the parameters are well identified.[35] We can therefore say that the numerical specification of the model rests on solid grounds.

In finer detail, it has to be taken into account that, although the standard percentile confidence intervals in Fig. 6 are a straightforward specification, they may not have the desired coverage probability. This is, for instance, the case with the distributions of χ or α_x, for which one may infer that the estimates from (15) are biased. This feature suggests that the bootstrap distribution of these parameters will be asymptotically

[31] In fact, among the 1,000 estimations there is only one case where p^a is slightly below 5 %.

[32] Since presently a set of 1,000 estimations on an average personal computer takes between 27 and 31 hours, an increase in S would require a parallel computing device.

[33] See Lee and Ingram (1991, p. 202).

[34] The estimates $\{\widehat{\theta}^a\}$ in (15) only take the sample variability in the simulations into account but not the variability arising from different realizations of the data generation process.

[35] On the basis of a number of explorations, we are confident that the intervals continue to be bounded and so the conclusion remains valid if ν is also treated as a free parameter.

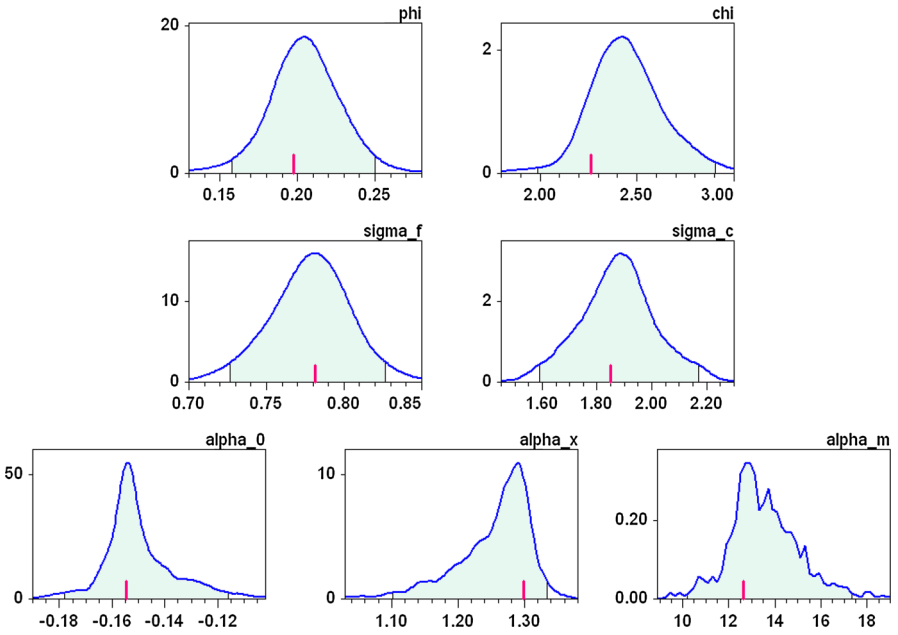

Fig. 6 Distributions of parameter re-estimates $\widehat{\theta}^d$ from (17). *Note*: The *shaded areas* represent the standard 95 % confidence intervals. The *short vertical bars* (in *red*) indicate the benchmark estimates $\widehat{\theta}_i$ from (15)

centred around the pseudo-true value plus a bias term, which would imply that the intervals shown are the 95 % confidence interval for the latter quantity. Thus, they may have a grossly distorted range as a confidence interval for the pseudo-true parameter value.[36] An alternative that solves the problem is Hall's percentile confidence interval (see Appendix 5). This is the reason why the lower and upper boundaries that we report in Table 2 are based on this device. The Hall intervals for χ and α_x, in particular, are seen here to be fairly different from the intervals in Fig. 6. The feature of a limited range of the intervals is, of course, maintained.

At the end of the evaluation of the structural parameters we should not conceal a problem that some perceptive readers might have with the estimation's overall noise level brought about by σ_f and σ_c. If one compares the in some sense typical order of magnitude in the price changes as they are caused by the deterministic and by the stochastic forces, the latter are found to strongly dominate the former. Is this more than an interesting observation and even a sufficient reason to discard the model altogether?[37]

[36] Even though the model may be misspecified, a pseudo-true parameter vector θ^o is a well-defined concept. If m^o is the expected moment vector of the true model of the stock market, θ^o satisfies $J[m(\theta^o), m^o] \leq J[m(\theta), m^o]$ for all admissible θ, where $m(\theta) = lim_{S\to\infty} E[m^a(\theta; S)]$ (assuming ergodicity, the expected values converge to the same limit for all random number sequences). This definition corresponds to that in Hnatkovska et al. (2011, p. 6), where the expected moments of the model can be analytically computed.

[37] A conclusion that actually was strongly insinuated to us by some severe readers of a previous version of the paper.

We would like to make three points on this issue. To begin with, we know of no model of similar complexity in the literature that would fare better in this respect. In fact, the present paper is the first pointing out this problem at all. Second, the word 'dominate' does not mean that the stochastic forces are also more important than the deterministic forces. The analysis in Sect. 3 has clearly shown that it is just the permanent interaction of the two types of forces that generate the stylized facts; one of them is absolutely useless without the other. This is more directly supported by the re-estimations of the deterministic and stochastic parameters ϕ, χ and σ_f, σ_c, respectively, all of which are definitely bounded away from zero.

The third point is that we should not be too surprised about the relatively high noise levels since after all the model is still very simple. It is rather by no means obvious that just the two noise sources in the agents' demand are already so efficient, if they are suitably scaled. Our perspective is that we can take them as a point of departure and that it is now time to build more structure into them. It is presently an open question for us whether this could be satisfactorily done with furthermore a few groups of agents or whether we would have to introduce a great number of individual agents with more differentiated strategies (using individual thresholds to become active on the market, for example). In other words, we understand the observation of the 'dominating' noise not as a vice but as a challenge to enter a new stage of agent-based modelling.

5 Conclusion

In the recent past increased efforts have been made to create small-scale agent-based models that are able to reproduce the stylized facts of financial markets, especially regarding the volatility clustering and fat tails of the daily returns. In previous work, we put forward the concept of structural stochastic volatility which, despite its parsimony, appeared to be fairly successful in this respect. Generally, it consists of two components. First, the core excess demand of two groups of speculative traders, to each of which a random term is added that is said to reflect the heterogeneity within the groups. Second, a mechanism that governs endogenous switches of the agents between the two strategies. If the noise terms differ in their variance, the variations of the two market fractions will induce variations in the overall noise level of the asset demand, and thus in the returns.

In this paper, a version of this modelling device with fundamentalist and chartist traders was reconsidered where the switching mechanism incorporates three socio-economic principles: herding, a certain predisposition towards chartism, and a propensity to withdraw from chartism as the gap between prices and the fundamental value widens. Beyond a mere observation of the model's ability to mimic the statistical regularities that we find in the empirical daily returns, a deeper understanding of these phenomena was obtained by an analysis of the dynamics in the phase plane of the asset price p_t and a strategy majority index x_t.

The key elements in this investigation are the isoclines of the majority index, i.e. the geometric locus where temporarily, in the deterministic part of the model, $\Delta x_{t+1} = 0$. Our analysis highlighted the fact that it is the synthesis of the deterministic and stochastic components that make the model work. The deterministic part would be

nothing without the random forces, and the latter would remain ineffective without an appropriate shape of the nonlinear $\Delta x_{t+1} = 0$ isoclines, which can be brought about by a skillful combination of the behavioural parameters in the switching function.

While these parameters are essentially responsible for the qualitative volatility clustering effects, the other parameters take care of the quantitative effects. The precise numerical values were obtained by a formal econometric estimation. As the 'stylized facts' are readily described by a set of summary statistics, or 'moments', our method of choice is the method of simulated moments (MSM), which seeks for values of the structural coefficients such that the simulated moments of the model come as close as possible to their empirical counterparts.

In addition to finding suitable parameters, we advanced the concept of a p value for the model's overall goodness-of-fit (conditional on the chosen moments, of course). Treating the estimated model as the true data generation process, simulating samples of artificial moments from it, and then re-estimating the model on them, this p value is the original estimation's error rate of falsely rejecting the null hypothesis. It should be higher than five per cent, and the higher it is, the better the fit. Moreover, by estimating the model with MSM on the empirical moments a great number of times, we took account of the problem of small-sample variability in the model simulations. In this way, we were able to compute an entire distribution of p values, one for each of these re-estimations, and finally set up a confidence interval for them. Thus we arrived at an upper and lower boundary for the p values of 32.6 and 8.7 %, respectively, which is the paper's main message to summarize the model's performance.

On the whole, besides another application of MSM as a powerful estimation approach, this paper proposed a further rigorous and simulation-based econometric test to quantify the goodness-of-fit of an asset pricing model. We believe that the aforementioned figures can be considered a success and present a challenge to other models of similar complexity. Regarding the analytical underpinnings of the present model's dynamic properties, the switching mechanism of which is based on the transition probability approach, it may be worthwhile to attempt a similar analysis for its "twin" model, which is based on the discrete choice approach and fared so well in the model contest discussed in Franke and Westerhoff (2012b). In this sense, the paper is more of a stimulus for further research than a final once-and-for-all result, where we have not yet mentioned the challenge for richer modelling discussed at the end of the previous section.

Appendix 1: Value added of the present contribution

Given that two papers of ours (Franke and Westerhoff 2011, 2012b; abbreviated 'FW'), deal with the same model, one may wonder about the current paper's original features.[38] Although they are mentioned when they are introduced in the text, it may be convenient to collect them in a self-contained overview.

[38] Incidentally, large parts of this paper were written before we even started with FW (2011, 2012b). Unfortunately, after its completion it took a rather thorny path through the journal landscape.

- FW (2011, 2012b) are mainly concerned with estimation issues. The geometric analysis in Sect. 3 about the functioning of the model, which allows us to discuss the precise nature of the interplay between the random shocks and the model's deterministic skeleton with its specific nonlinear isoclines, is entirely original material. Actually, as far as we know, there is no comparable in-depth analysis in the literature on agent-based asset pricing models.
- While in FW (2011) the flexibility parameter ν was estimated at the not very convincing value $\nu = 0.57$, it is here (and in FW 2012b) exogenously set at $\nu = 0.050$, which [as pointed out in the paragraph following Eq. (14)] makes more psychological sense. The estimates of the other parameters are somewhat affected by this change, but not too much. This robustness is not obvious and perhaps even a bit surprising.
- The weighting matrix W in the loss function J is given by the inverse of the covariance matrix $\widehat{\Sigma}$. In the other two papers the latter was obtained from a block bootstrap, which is plagued with the so-called joint-point problem. As a remedy we now propose what we call a history sampling bootstrap [mentioned in the paragraph following Eq. (12)]. Despite the straightforward idea on which it is based, we know of no precursor in the bootstrap literature.
- A general problem of the method of simulated moments is that the estimation rests on a given sequence of random numbers; another sequence—another solution of the loss minimization in eq. (14). Here we introduce the concept of a "representative" estimation (see the brief discussion of eq. (15)), which is a new idea.
- A common feature of this and the other two papers is that we generate a frequency distribution of values $\{J^b\}_{b=1}^B$ of the loss function and take its 95 % quantile as a yardstick to evaluate the previously estimated loss $\widehat{J}$. It also serves to determine a p value for the model. In FW (2011, 2012b), such a reference distribution was obtained from a nonparametric bootstrap, which only makes use of the empirical data.[39] In the present paper, a parametric bootstrap is employed, which involves the model and its previous "representative" parameter estimates. This is now pointed out in footnote 27.

Appendix 2: A note on the nature of variable x in the literature

The role of the majority index x_t in an adjustment equation such as (6) may seem slightly unclear in some of the literature, so that the concepts involved here may not always have been fully understood.[40] In early publications, the equation was only formulated after the transition probabilities were utilized to set up the so-called Master equation. From this point of view, the stochastic process is characterized not by the actual values of x_t and some other state variables, but by entire probability distributions

[39] In fact, we are now aware that we did not fully exploit the information that could be obtained from a nonparametric bootstrap, where the critical point is a recentring of the loss function [cf. the discussion in Franke (2012), Sect. 2.2]. The parametric bootstrap in the present paper does not suffer from this kind of problem.

[40] The present authors do not exempt themselves from this.

of them, which are furthermore subject to change over time. The adjustment equation, which is a deterministic equation, is referred to here as "an approximative mean value equation for the original stochastic system", whose analysis "is sufficient to determine the most probable development from any initial state." Neglecting the other aspect of the probability distributions can technically be justified "by the convenient assumption of a sharply peaked *initial* distribution" (Lux 1995, p. 885; emphasis in the original).

Two questions arise from these presentations. (1) As the probability distribution varies over time, is it ensured that it remains so sharply peaked?[41] (2) Equilibrium (i.e. time-invariant) probability distributions that have a bimodal density function are of particular interest. This implies that over longer periods of time a sample trajectory fluctuates around some (low) value of the majority index, then eventually switches over into the neighbourhood of another (high) value of x, fluctuates around it for another period of time, until it switches back into a neighbourhood of the first value, etc. Since the probability distribution does not change during all this, its mean value does not change either. The specific value it attains would indeed be some constant in an intermediate range between the two more extreme values. In this situation, the assumption of peakedness is violated, although the stochastic process itself is in its (unique) equilibrium. The expected value would only provide misleading information about what is actually going on between the agents.

The ambiguities can be resolved by deriving the so-called Langevin equation for x_t. Although it looks similar to Eq. (6), x_t is here not an approximative mean value but the actual value of the majority index in a sample trajectory. This equation can be viewed as a stochastic adjustment rule for x_t. In general, it includes an additive noise term with a variance that decreases with the population size. It moreover becomes the deterministic Eq. (6), i.e. the variance tends to zero, as the population size becomes infinitely large.

For more information about the historical background of the transition probability approach as well as a rigorous derivation of Eq. (6) in a stochastic and the present deterministic version, see Franke (2008a, b).

Appendix 3: Mathematical proofs

Proof of Proposition 1 To determine the equilibrium value(s) of the majority index, it proves useful to resort to the definition of the hyperbolic sine and cosine (sinh and cosh). This allows us to rewrite (6) and (7) as $\Delta x_{t+1} := x_{t+1} - x_t = 2\nu \{ [\exp(s_t) - \exp(-s_t)]/2 - x_t [\exp(s_t) + \exp(-s_t)]/2 \} = 2\nu [\sinh(s_t) - x_t \cosh(s_t)]$. With $\tanh = \sinh / \cosh$ for the hyperbolic tangent, we then get

$$\Delta x_{t+1} = x_{t+1} - x_t = 2\nu \{\tanh[s(x_t, p_t)] - x_t\} \cosh[s(x_t, p_t)] \qquad (21)$$

Since cosh is an everywhere positive function, $\Delta x_{t+1} = 0$ if and only if the term in curly brackets vanishes. Hence, taking $p = p^\star$ in the switching index (8) into

[41] For a specific system, this question is answered by an explicit (elaborate) mathematical analysis in Lux (1997, Sects. 4.1 and 4.2).

account, any equilibrium value of x has to satisfy the relationship $\tanh(\alpha_o + \alpha_x x) = x$. Applying the inverse function $\operatorname{arctanh}(\cdot)$ to both sides of this equation and using the identity $\operatorname{arctanh}(x) = (1/2)\ \ln[(1+x)/(1-x)]$, the equilibrium condition for the majority index can be reformulated as

$$g(x) := \alpha_x\, x - \frac{1}{2}\ \ln\left[\frac{1+x}{1-x}\right] + \alpha_o = 0 \tag{22}$$

To locate the roots of the function $g(\cdot)$, note that it tends to $+\infty$ as x approaches -1 from the right, and to $-\infty$ as x approaches $+1$ from the left. In addition, the derivative is computed as $g'(x) = \alpha_x - 1/(1-x^2)$. If, as in part (a) of the proposition, α_x is contained between zero and unity, $g'(x)$ is negative over the entire domain. Hence a unique equilibrium value x^o exists in this case.[42]

Consider next $\alpha_x > 1$ together with a zero intercept $\alpha_o = 0$ in the switching index. One equilibrium value satisfying (22) is then given by $x^o = 0$, in which $g(\cdot)$ is now upward sloping. Equating the derivative to zero, it is seen that $g(\cdot)$ has exactly one local minimum between -1 and x^o, in which g is negative, and (symmetrical to it) exactly one local maximum between x^o and $+1$, in which g is positive. From the limiting behaviour of the function for $x \to \pm 1$, we thus infer the existence of exactly two additional outer equilibria; one between -1 and x^o and the other between x^o and $+1$. This proves part (b) of the proposition.

As for part (c), fix $\alpha_x > 1$ and, starting from zero, let the predisposition parameter α_o decrease. Obviously, this shifts the function $g(\cdot)$ downwards. As a consequence, x^o and x^{fd} move towards each other, x^o as the interior and x^{fd} as the outer-right point of intersection of $g(\cdot)$ with the zero line. Eventually, as the downward shift of α_o continues, the local maximum of $g(\cdot)$ will be zero. When this occurs, x^o and x^{fd} collapse into one single point of intersection. Subsequently, if α_o decreases further, they disappear. Under these circumstances, x^{cd} remains as the only equilibrium point, where the shifting procedure has moved it monotonically to the left all the time. This observation completes the proof. □

Proof of Proposition 2 Given a pair (x_t, p_t), we have $\Delta x_{t+1} \geq 0$ if and only if the term in curly brackets in (21) is nonnegative, or $\tanh[\alpha_o + \alpha_x x_t + \alpha_d (p_t - p^\star)^2] \geq x_t$. Applying the strictly increasing arctanh function on both sides of the inequality and using the abovementioned identity for $\operatorname{arctanh}(x_t)$ as well as the definition of the function $g(\cdot)$, this relationship is equivalent to $g(x_t) \geq -\alpha_d\,(p_t - p^\star)^2$. It is certainly fulfilled if $g(x_t) > 0$ or, in the case $g(x_t) = 0$, if $p_t \neq p^\star$.

If $g(x_t) < 0$, we can multiply the inequality by -1, which reverses the inequality sign, and then take the square root on both sides. This yields the condition $p_t - p^\star \geq \sqrt{-g(x_t)/\alpha_d}$ if $p_t > p^\star$ and $p_t - p^\star \leq -\sqrt{-g(x_t)/\alpha_d}$ if $p_t < p^\star$. The remaining statements in part *(c)* are obvious. □

[42] Incidentally, the argument remains the same if $\alpha_x \leq 0$, although we would then have the opposite of herding.

Appendix 4: The history-sampling procedure in bootstrapping the empirical moments

Bootstrapping the empirical autocorrelations of r_t and $v_t = |r_t|$ requires a second thought. As a representative example, consider the hth-order autocorrelation of v_t ($h \in \mathbb{N}$), which for a sample of size T reads,

$$\rho_v(h) = (1/T) \sum_{t=1+h}^{T} (v_t - \bar{v})\,(v_{t-h} - \bar{v}) \,/\, s_v^2,$$

$$\text{where} \qquad \bar{v} = (1/T) \sum_{t=1}^{T} v_t \,, \qquad s_v^2 = (1/T) \sum_{t=1}^{T} (v_t - \bar{v})^2$$

With a view to the bootstrap procedure to be specified shortly, it is convenient to define the set of time indices

$$I^o = \{1, 2, \ldots, T\}$$

and rewrite the autocorrelation as

$$\rho_v^{emp}(h) = (1/T) \sum_{t \in I^o} (v_t - \bar{v})\,(v_{t-h} - \bar{v}) \,/\, s_v^2 \quad (\text{putting } v_{t-h} = \bar{v}^b \quad \text{if } \; t - h \leq 0)$$

(the superscript '*emp*' has been added for greater clarity.)

Bootstrapping summary statistics that involve lagged values of the dynamic variables is often carried out as a block bootstrap of the time series data. For longer lags h, however, this is not an entirely satisfactory procedure because the independence between the randomly selected single blocks cannot reproduce the dependence structure of the original sample, a phenomenon known as the join-point problem. In addition, the variability of various moments may thus be increased (cf. Andrews 2004, p. 674).

While these are serious problems in likelihood or dynamic regression estimations,[43] they can be circumvented in the present moment matching approach. To put up a bootstrap sample b, we need not form a new series of consecutive data points and compute the moments from them, but can sample directly from the time indices: alternatively to I^o, they give us a new set I^b on which we can base the same calculations as above (of course, the same index set I^b for each of the moments, with and without lags). Accordingly, a bootstrap sample in our approach is constituted by T random draws with replacement from the set I^o (each time index having the same probability $1/T$). Repeating this B times, we have $b = 1, \ldots, B$ index sets

$$I^b = \{t_1^b, t_2^b, \ldots, t_T^b\}$$

[43] For which Andrews (2004) proposes the concept of a block–block bootstrap.

from which, analogously to the empirical magnitudes, we can subsequently obtain the bootstrapped moments

$$\rho_v^b(h) = (1/T) \sum_{t \in I^b} (v_t - \bar{v}^b)\,(v_{t-h} - \bar{v}^b) \,/\, (s_v^2)^b, \quad b = 1, \ldots, B; \tag{23}$$

$$\text{where} \quad v_{t-h} = \bar{v}^b \text{ if } t \leq h, \quad \bar{v}^b = (1/T) \sum_{t \in I^b} v_t, \quad (s_v^2)^b = (1/T) \sum_{t \in I^b} (v_t - \bar{v}^b)^2$$

It might be noted that, while in an empirical autocorrelation $\rho_v^{emp}(h)$ exactly h of the T terms in the sum vanish, there may be more or less such zero terms in a bootstrapped autocorrelation $\rho_v^b(h)$. Given the large sample we have, however, this effect will be negligible.

The statistics computed according to (23) are the components of the moment vectors m^b from which subsequently the covariance matrix $\widehat{\Sigma}$ in (13) is made up.

Appendix 5: Hall's percentile confidence interval

Let a collection $\{\, \widehat{\theta}^b : b = 1, \ldots, B \,\}$ of parameter re-estimates be given. With respect to a significance level $\alpha = 0.05$, let $\widehat{\theta}_{i,L}$ be such that only a fraction $\alpha/2$ of all the bootstrap estimates $\widehat{\theta}_i^b$ are less than this value, and likewise let $\widehat{\theta}_{i,H}$ be the value that is exceeded by only $\alpha/2$ of the bootstrap estimates. The standard percentile confidence interval is then given by

$$CI_S(\theta_i) = [\, \widehat{\theta}_{i,L},\, \widehat{\theta}_{i,H} \,] \tag{24}$$

(where the index S indicates that (24) is regarded as the standard method). To fix the problem that $CI_S(\theta_i)$ will not have the desired coverage probability in the presence of a bias, Hall's percentile confidence interval is proposed (see Hall 1992, Chapter 3). With respect to the original estimate $\widehat{\theta}_i$ on the empirical moments, it is defined as

$$CI_H(\theta_i) = [\, 2\widehat{\theta}_i - \widehat{\theta}_{i,H},\, 2\widehat{\theta}_i - \widehat{\theta}_{i,L} \,] \tag{25}$$

(presupposing that it does not contain inadmissible values of the parameter). Letting θ_i^o be the pseudo-true parameter value, this specification is based on the idea that the bootstrap distribution $(\widehat{\theta}_i^b - \widehat{\theta}_i)$ approximates the distribution $(\widehat{\theta}_i - \theta_i^o)$. This implies that $\text{Prob}(\widehat{\theta}_{i,L} - \widehat{\theta}_i < \widehat{\theta}_i - \theta_i^o < \widehat{\theta}_{i,H} - \widehat{\theta}_i) \approx \text{Prob}(\widehat{\theta}_{i,L} - \widehat{\theta}_i < \widehat{\theta}_i^b - \widehat{\theta}_i < \widehat{\theta}_{i,H} - \widehat{\theta}_i) = 1 - \alpha$, and the first probability expression is easily seen to be equal to $\text{Prob}(2\widehat{\theta}_i - \widehat{\theta}_{i,H} < \theta_i^o < 2\widehat{\theta}_i - \widehat{\theta}_{i,L}) = \text{Prob}(\theta_i^o \in CI_H(\theta_i))$. Hence Hall's percentile method (25) is asymptotically correct.

References

Altonji J, Segal L (1996) Small-sample bias in GMM estimation of covariance structures. J Bus Econ Stat 14:353–366

Andrews DWK (2004) The block-block bootstrap: improved asymptotic refinements. Econometrica 72:673–700
Brock W, Hommes C (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. J Econ Dyn Control 22:1235–1274
Chen S-H, Chang C-L, Du Y-R (2012) Agent-based economic models and econometrics. Knowl Eng Rev 27:187–219
Chiarella C, Dieci R, He X-Z (2009) Heterogeneity, market mechanisms, and asset price dynamics. In: Hens T, Schenk-Hoppé KR (eds) Handbook of financial markets: dynamics and evolution. North-Holland, Amsterdam, pp 277–344
Cont R (2001) Empirical properties of asset returns: stylized facts and statistical issues. Quant Financ 1:223–236
Davidson R, MacKinnon JG (2004) Econometric theory and methods. Oxford University Press, Oxford
Farmer JD, Joshi S (2002) The price dynamics of common trading strategies. J Econ Behav Organ 49:149–171
Franke R (2008a) Microfounded animal spirits and Goodwinian income distribution dynamics. In: Flaschel P, Landesmann M (eds) Effective demand, income distribution and growth. Research in memory of the work of Richard M. Goodwin. Routledge, London, pp 372–398
Franke R (2008b) Estimation of a microfounded herding model on German survey expectations. Interv Eur J Econ Econ Policies 5:301–328
Franke R (2009) Applying the method of simulated moments to estimate a small agent-based asset pricing model. J Empir Financ 16:804–815
Franke R (2010) On the specification of noise in two agent-based asset pricing models. J Econ Dyn Control 34:1140–1152
Franke R (2012) How much backward-looking is the New-Keynesian three-equations model? Evidence from moment matching estimations. Working paper, University of Kiel
Franke R (2013) Aggregate sentiment dynamics: a canonical modelling approach and its pleasant nonlinearities. Working paper, University of Kiel
Franke R, Asada T (2009) Incorporating positions into asset pricing models with order-based strategies. J Econ Interact Coord 3:201–227
Franke R, Westerhoff F (2011) Estimation of a structural stochastic volatility model of asset pricing. Comput Econ 38:53–83
Franke R, Westerhoff F (2012a) Converse trading strategies, intrinsic noise and the stylized facts of financial markets. Quant Financ 12:425–436
Franke R, Westerhoff F (2012b) Structural stochastic volatility in asset pricing dynamics: estimation and model contest. J Econ Dyn Control 36:1193–1211
Graham B, Dodd D (1951) Security analysis. McGraw Hill, New York
Hall P (1992) The bootstrap and edgeworth expansion. Springer, New York
Heemeijer P, Hommes CH, Sonnemans J, Tuinstra J (2009). Price stability and volatility in markets with positive and negative expectations feedback: an experimental investigation. J Econ Dyn Control (forthcoming)
Hnatkovska V, Marmer V, Tang Y (2011) Comparison of misspecified calibrated models: the minimum distance approach. University of British Columbia and Bowdoin College, Brunswick, mimeo
Hommes C (2006) Heterogeneous agent models in economics and finance. In: Tesfatsion L, Judd K (eds) Handbook of computational economics, vol 2: agent-based computational economics. North-Holland, Amsterdam, pp 1107–1186
Hommes CH, Sonnemans J, Tuinstra J, van de Velden H (2007) Learning in cobweb experiments. Macroecon Dyn 11(S1):8–33
Hommes C, Wagener F (2009) Complex evolutionary systems in behavioral finance. In: Hens T, Schenk-Hoppé KR (eds) Handbook of financial markets: dynamics and evolution. North-Holland, Amsterdam, pp 217–276
LeBaron B (2006) Agent-based computational finance. In: Tesfatsion L, Judd K (eds) Handbook of computational economics, vol 2: agent-based computational economics. North-Holland, Amsterdam, pp 1187–1233
Lee B-S, Ingram BF (1991) Simulation estimation of time series models. J Econom 47:197–205
Lux T (1995) Herd behaviour, bubbles and crashes. Econ J 105:881–889
Lux T (1997) Time variation of second moments from a noise trader/infection model. J Econ Dyn Control 22:1–38

Lux T (1998) The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. J Econ Behav Organ 33:143–165
Lux T (2009a) Applications of statistical physics in finance and economics. In: Rosser JB Jr (ed) Handbook of research on complexity. Edward Elgar, Cheltenham, pp 213–258
Lux T (2009b) Stochastic behavioural asset-pricing models and the stylized facts. In: Hens T, Schenk-Hoppé KR (eds) Handbook of financial markets: dynamics and evolution. North-Holland, Amsterdam, pp 161–216
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
Menkhoff L, Taylor M (2007) The obstinate passion of foreign exchange professionals: technical analysis. J Econ Lit 45:936–972
Menkhoff L, Rebitzky RR, Schröder M (2009) Heterogeneity in exchange rate expectations: evidence on the chartist-fundamentalist approach. J Econ Behav Organ 70:241–252
Murphy J (1999) Technical analysis of financial markets. New York Institute of Finance, New York
Press WH et al (1986) Numerical recipes: the art of scientific computing. Cambridge University Press, Cambridge
Weidlich W, Haag G (1983) Concepts and models of a quantitative sociology: the dynamics of interacting populations. Springer, Berlin
Westerhoff F (2003) Expectations driven distortions in the foreign exchange market. J Econ Behav Organ 51:389–412
Westerhoff F (2008) The use of agent-based financial market models to test the effectiveness of regulatory policies. Jahrbücher für Nationalökonomie und Statistik (J Econ Stat) 228:195–227
Westerhoff F (2009) Exchange rate dynamics: a nonlinear survey. In: Rosser JB Jr (ed) Handbook of research on complexity. Edward Elgar, Cheltenham, pp 287–325
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