Statistical analysis of financial returns for a multiagent order book model of asset trading

Tobias Preis,* Sebastian Golke, Wolfgang Paul, and Johannes J. Schneider
Institute of Physics, Johannes Gutenberg University of Mainz, Staudinger Weg 7, 55099 Mainz, Germany
(Received 21 December 2006; revised manuscript received 26 April 2007; published 20 July 2007)

We recently introduced a realistic order book model [T. Preis *et al.*, Europhys. Lett. **75**, 510 (2006)] which is able to generate the stylized facts of financial markets. We analyze this model in detail, explain the consequences of the use of different groups of traders, and focus on the foundation of a nontrivial Hurst exponent based on the introduction of a market trend. Our order book model supports the theoretical argument that a nontrivial Hurst exponent implies not necessarily long-term correlations. A coupling of the order placement depth to the market trend can produce fat tails, which can be described by a truncated Lévy distribution.

DOI: 10.1103/PhysRevE.76.016108 PACS number(s): 89.65.Gh, 05.10.Ln, 02.50.Ey

I. INTRODUCTION

Already some decades ago Osborne [1,2] and Mandelbrot [3,4] started to apply methods and models from statistical physics and complexity theory to problems arising in the description of the behavior of financial markets and thus became founders of a new area of research nowadays often called econophysics. They discovered that the price movements at the financial markets do not show Gaussian fluctuations [5–7], which is still one of the basic assumptions in the economic sciences. The modeling of financial time series by a diffusive stochastic process dates back to the Ph.D. thesis of Bachelier [8] in 1900. Geometric Brownian motion [1,9,10], a variant of Brownian motion, which was independently investigated by Bachelier and Einstein [11], has become the standard mathematical model, and the theory of fair option prices [12,13] with its huge impact on financial market is based on it.

However, financial market time series indicate a more complex statistical behavior, which contains non-Gaussian characteristics which are, e.g., also found in critical fluctuations of physical systems.

The work of Mandelbrot was based on data records of very limited length at that time. In the course of the progress in information technology, which was accompanied by constantly growing computing resources, trading processes were adapted to the computational infrastructure at the international financial markets and full electronic exchanges were created. Thus, nowadays an impressive quantity of historical financial market time series is available. Using this constantly extending data basis, Mandelbrot's results could be confirmed and the question arises as to what extent established assumptions in economics are not sufficient.

Noticing this discrepancy between theory and financial-economical reality, physicists started again to examine the complex "multiparticle systems" of financial markets in the 1990s with a large variety of physical methods and to develop simple models for their description [6,14]. An early very simple agent-based model [14], which is based on the reaction diffusion process $A+B\rightarrow 0$ and which contains the elements of imitation and feedback, was already able to reproduce some nontrivial characteristics of real financial markets, like a nontrivial Hurst exponent and a non-Gaussian price change distribution, in which large price changes are more probable than predicted by the Gaussian distribution.

In the last years, physicists have intensified their efforts to understand the process of price formation in financial markets in more detail, since simple models were not able to reproduce the behavior of financial markets completely. In this context, Maslov and Mills [15,16] published an alternative model, in which two different types of orders, the limit order and the market order, were introduced. Another model, into which some additional rules are integrated, such as, for example, a Poisson process for the deletion of limit orders, was suggested by Challet and Stinchcombe [17]. Further agent-based market models [17–24] were published targeted at reproducing and interpreting empirical stylized facts. Especially one work [25,26] has to be mentioned, in which a statistical model of the continuous double auction is examined analytically and numerically.

We recently suggested an agent-based order book model [27] which can be used to obtain a "mechanistic understanding" of the price formation process leading to the so-called stylized facts observed in financial markets. Here we present a detailed analysis of this model and show that the components leading to so important properties as non-Gaussian return distributions or a persistent price dynamics on intermediate time scales (Hurst exponent $H>1/2$) can be identified and realized individually.

Section II will present the definition of the basic model. In Sec. III we analyze the parameter space of this model to identify meaningful regions in parameter space. In Sec. IV we then discuss several augmentations of the model leading to a nontrivial Hurst exponent and a non-Gaussian return distribution. Finally Sec. V gives our conclusions and an outlook.

II. DEFINITION OF THE ORDER BOOK MODEL

In this section, the order book model in its basic form is defined. It is inspired by the model for the continuous double auction introduced in [25,26]. Our aim is to accurately reproduce the structure and the mechanisms of an order book at real financial markets, as shown in Fig. 1. In our simulations, we limit ourselves to only one order book in which one

*mail@preis.co.uk

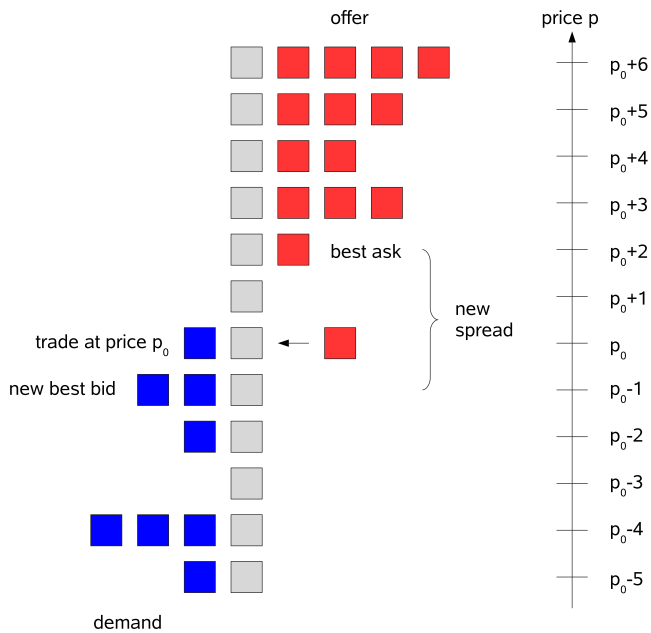

FIG. 1. (Color online) Structure of the order book: the limit buy ($p \leq p_0$) and limit sell orders ($p > p_0$) are added chronologically to the appropriate discrete price level p. Here a matching algorithm with price time priority is implemented. In this example, a new sell order is placed at price p_0, at which there is already a buy order stored in the order book. The two orders at price p_0, which are on the demand and offer site of the order book, are executed against each other. Thus, a trade is performed at price p_0, which is then called the last traded price. The spread s, which is the difference between the best ask and best bid prices, was two ticks before the arrival of the executable sell order at price level p_0 and increases to three ticks after this trade.

individual asset is traded. This asset can be, e.g., a share, a loan, or a derivative product. From the various types of orders which can be found at real financial markets we use only the two most important types: namely, limit orders and market orders. The order book model contains two different types of agents, liquidity providers and liquidity takers, which differ in the types of orders they are permitted to submit.

On the one hand, N_A liquidity providers transmit only limit orders. In the case of a limit sell order, an agent offers an asset for sale at a given limit price or any price better for this agent. Analogously, a limit buy order indicates a demand for buying a traded asset and will be executed at a given limit price or a lower price. Let p_a be the so-called best ask, which is the lowest price level for which at least one limit sell order in the order book exists, and analogously p_b the so-called best bid, being the highest price level for which at least one limit buy order is stored in the order book. In our model, limit orders are placed adjusted around the midpoint

$$p_m = \frac{p_a + p_b}{2}, \tag{1}$$

with a rate α; i.e., αN_A new limit orders are inserted in the order book per time step. We denote q_{provider} to be the probability with which a limit order, which is to be placed, is a limit buy order; thus, with probability $1-q_{\text{provider}}$, the limit order to be placed is a limit sell order. The liquidity provider, which one could identify as so-called market maker, supplies the order book liquidity in this way. The aim of these market participants is to use the nonzero spread $s=p_a-p_b$ for earning money: they intend to sell an asset at price p_a or higher and then to buy it back at price p_b or lower, thus having earned at least the spread s (if it remained constant between the sale and the purchase of the asset). Of course, they can analogously try to make money the other way around by first buying at price p_b or lower and then selling at price p_a or higher. Like in real financial markets, we allow agents to sell assets even if they do not possess them and thus to perform so-called short sales.

On the other hand, N_A liquidity takers transmit only market orders with a rate μ; i.e., μN_A market orders are inserted per time step. (Thus, there are altogether $2N_A$ agents in the system.) A market order is immediately executed after arriving in the order book: a market sell order is executed at price p_b, a market buy order at price p_a. The market order of a liquidity taker is a market buy order with probability q_{taker} and a market sell order with probability $1-q_{\text{taker}}$. In the first approach of our order book model, we simply use

$$q_{\text{provider}} = q_{\text{taker}} = \frac{1}{2}. \tag{2}$$

Thus, limit orders and market orders are produced symmetrically around the midpoint.

Limit orders, which are stored in the order book, can expire or can be deleted. The removal of a limit order is realized in the way that each inserted order is deleted with probability δ per time unit.

As there are overall $2N_A$ agents in the system, each Monte Carlo step (MCS) consists of $2N_A$ moves, in which one agent is randomly selected and can perform one action. If the agent is a liquidity provider, then this agent submits a limit order with probability α. Subsequently, independently of whether it came to an order placement or not, each limit order of this liquidity provider is deleted with probability δ. On the other hand, if the randomly selected agent is a liquidity taker, then this agent places a market order with probability μ which is immediately executed.

In the investigations of the order book model shown here, we only consider order volumes of 1; i.e., only one asset can be offered or demanded with a single order. The matching algorithm works according to a price time priority; i.e., first those limit orders at best ask and best bid are executed, respectively. If there is more than one order at a given price level, then the orders are executed in the same chronology as they were inserted in the order book.

A. Liquidity providers vs liquidity takers

On the basis of this model definition, first an unrealistic independent identically distributed (IID) order placement within two intervals (for the buy and for the sell orders, respectively), which have both a width of p_{int}, is assumed, according to which every liquidity provider enters his limit buy orders to each price level in the interval of $[p_a-1-p_{\rm int};p_a-1]$ with the same probability. Accordingly, limit sell orders are transmitted uniformly distributed in the interval of $[p_b+1;p_b+1+p_{\rm int}]$. With these settings, we were able to reproduce the results of [25,26]. In [25,26] an appropriate microscopic-dynamic-statistical model for the continuous double auction is examined with analytic approximations under the assumption of an IID order flow and the limit $p_{\rm int}\to\infty$.

Already with this comparatively simple realization of the order book model, the profit loss distribution of the agents can be analyzed. In the order book model, the main distinction between the agents is that they are either liquidity providers or liquidity takers, which is reflected in the types of orders they are permitted to perform. Thus, first we want to investigate how the nature of a trader influences the temporal development of his account balance.

Let $\kappa_i(t)$ be the account balance of agent i at time t. Each agent i possesses no money at the beginning of the simulation at $t=t_0$, such that $\kappa_i(t_0)=0$ applies to $i=1,2,\ldots,2N_A$. In order not to restrict their trades, we provide an unlimited credit line to each agent, free of charge.

Furthermore, each agent buys and sells assets over time. The number of assets which agent i possesses at time t is given by $\pi_i(t)$. Also here we set $\pi_i(t_0)=0$ for all $1\leqslant i\leqslant 2N_A$, such that each agent possesses zero assets at the beginning. Note that we allow also negative values of $\pi_i(t)$, as an agent can sell an asset he does not possess.

Of course, when buying an asset, agent i has to pay the price of this asset, such that $\pi_i(t)$ is incremented by 1 with this trade at time t, but $\kappa_i(t)$ is decreased by $p(t)$, with $p(t)$ being the transaction price of this trade, which then becomes the last traded price. Analogously, when selling an asset, $\pi_i(t)$ is decremented by 1 and $\kappa_i(t)$ is increased by $p(t)$.

The overall wealth $\gamma_i(t)$ of agent i at time t thus consists both of the account balance $\kappa_i(t)$ and the number of assets $\pi_i(t)$:

$$\gamma_i(t)=\kappa_i(t)+\pi_i(t)p(t). \tag{3}$$

The change of the wealth of agent i between time t_0 and t is given by

$$\Delta\gamma_i=\gamma_i(t)-\gamma_i(t_0)=\gamma_i(t), \tag{4}$$

as $\gamma_i(t_0)=0$.

When simulating this order book model with Monte Carlo techniques, we always find a significant differentiation between the wealth distributions of the liquidity providers and of the liquidity takers, respectively. The group of liquidity takers is systematically disadvantaged in relation to the group of liquidity providers. Although it is possible that some liquidity takers obtain a positive trading result, nevertheless a separation of the two groups arises, because liquidity takers have to pay additionally the spread $s=p_a-p_b$ to liquidity providers if opening or closing a position in the traded asset, i.e., if either first buying and then reselling or if first selling and then rebuying an asset.

The wealth values of liquidity takers and liquidity providers drift apart linearly in time as shown in Fig. 2.

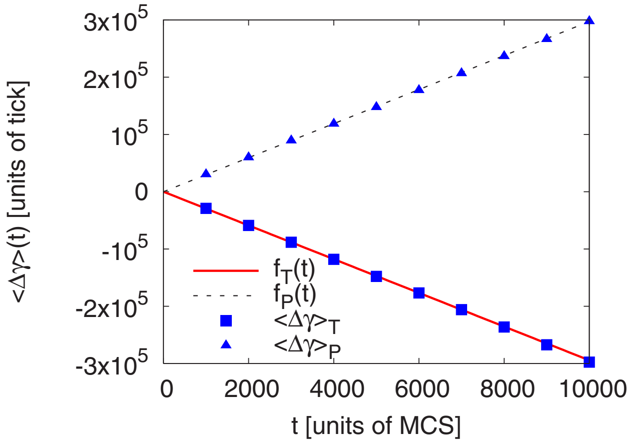

FIG. 2. (Color online) Liquidity providers vs liquidity takers: Averaged wealth shift of the liquidity providers and the liquidity takers, respectively, as a function of t for $\alpha=0.15$, $\mu=0.1$, $\delta=0.025$, and $N_A=250$. These dependences can be fitted with the functions $f_{\rm T}(t)=a_{\rm T}t$ and $f_{\rm P}(t)=a_{\rm P}t$ with the parameters $a_{\rm T}\approx-29.45$and $a_{\rm P}\approx 29.75$. The difference between the absolute values of these two gradients is due to statistical inaccuracies occurring during the determination of the averages of the distributions.

The identical absolute values of the gradients $a_{\rm P}$ and $a_{\rm T}$ are related to the spread s: an individual liquidity taker loses on average $|a_{\rm T}|$ ticks per MCS. That is the average gain of a liquidity provider per MCS, whose wealth is increased averagely by $|a_{\rm P}|$ ticks per MCS. Thus, on average $N_A|a_{\rm T}|$ ticks are transferred from the group of liquidity takers to the group of liquidity providers per MCS. Since $N_A\mu$ market orders are submitted to the order book and therefore $N_A\mu$ trades take place, the average wealth transfer per transaction is given through

$$\Gamma=\frac{N_A\langle|a_i|\rangle}{N_A\mu}=\frac{\langle|a_i|\rangle}{\mu}, \tag{5}$$

with $\langle|a_i|\rangle=(|a_{\rm P}|+|a_{\rm T}|)/2$ being the averaged absolute value of the gradients. For the parameters $\alpha=0.15$, $\mu=0.1$, $\delta=0.025$, $p_{\rm int}=2000$, and $N_A=250$ used in the simulations for the results shown in Fig. 2, the averaged wealth transfer per transaction can be determined to $\Gamma=296$ ticks. This is approximately half of the averaged spread s, which was determined in the simulation to be $\langle s\rangle=606$ ticks:

$$\Gamma=\frac{\langle s\rangle}{2}. \tag{6}$$

The factor of $1/2$ is to be attributed to the fact that a liquidity provider needs two transactions for earning the complete spread s—he has to buy once and to sell once. Therefore the agent earns only $\langle s\rangle/2$ on average per transaction.

If comparing these results with the situation given in real markets, it has to be mentioned that in fact liquidity takers are disadvantaged financially compared to liquidity providers. Our distinction in the order book model between liquidity providers and liquidity takers reflects the two different types of orders which are used: limit orders and market orders. In real financial markets there is no strict distinction between these two groups of traders. In general, each market participant can send both types of orders to the electronic order book, so that the idealized situation of Fig. 2 will be difficult to verify in real markets by looking at the wealth evolution of individuals.

B. Exponential order placement depth

An IID order placement depth, as used in the previous section, is not in agreement with the conditions which are found at real financial markets. In contrast to the uniform cumulative order volume generated by the IID order flow of [25,26], the order book depth of real markets can be described by a log-normal distribution [16].

To take this into account we replace the IID limit order placement in the fixed interval p_{int} around the midpoint p_m by an exponentially distributed order placement depth.

For placing a limit order i, the limit price p_i^l is determined for a limit buy order with

$$p_i^l = p_a - 1 - \eta \tag{7}$$

and for a limit sell order according to

$$p_i^l = p_b + 1 + \eta, \tag{8}$$

whereby η is an exponentially distributed integer random number created by

$$\eta = \lfloor -\lambda_0 \ln(x) \rfloor, \tag{9}$$

with x being a uniformly distributed random number in the interval $[0;1)$ and $\lfloor z \rfloor$ denoting the integer part of z. Thus, the transmittal of limit orders minimizes a potentially existing spread. Also the situation is avoided that a limit order becomes instantaneously executable at p_a or p_b, such that a limit order degenerates to a market order.

Using this exponentially distributed order placement depth, a log-normal-distributed order book depth is achieved, as shown exemplarily in Fig. 3. This implementation of the order book model is from now on regarded as the basic version of the order book model, for which we will show the most important time series characteristics. In Fig. 4, first an exemplary price sequence of 10^6 MCS is shown. The autocorrelation of the corresponding time series $\delta p(t) = p(t+1) - p(t)$ of the price changes shows the same behavior as can be observed in real financial market data. A significant negative autocorrelation exists for time lag $\delta\tau = 1$. Thus, a positive price change succeeds a negative price change with large probability and vice versa. The autocorrelation vanishes for $\delta\tau > 1$ [27].

Figure 5 shows the Hurst exponent $H(\Delta\tau)$ for different agent numbers N_A. Generally the Hurst exponent $H(q)$ is calculated by the relationship

$$\langle |p(t+\Delta\tau) - p(t)|^q \rangle^{1/q} \propto \Delta\tau^{H(q)}, \tag{10}$$

as defined, for example, in [28]. If not mentioned otherwise, we use $H(\Delta\tau) \equiv H(\Delta\tau, q=2)$. For comparison, the constant Hurst exponent $H = 1/2$ of the random walk is given in Fig. 5 additionally. On short time scales, the price process indicates an antipersistent behavior, which is due to the order book structure. On long time scales, the process converges towards a diffusive regime.

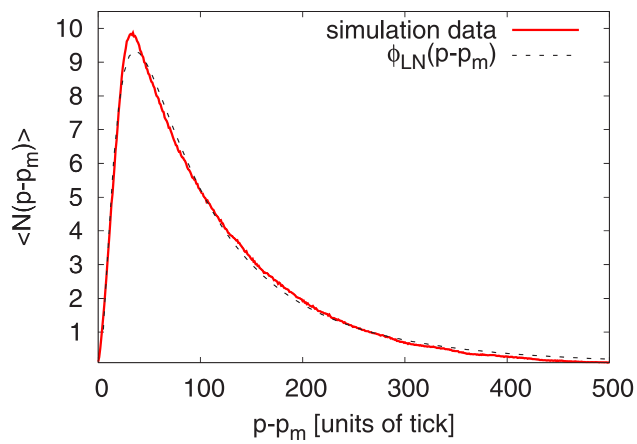

FIG. 3. (Color online) Exponential order placement depth: with the parameter selection $\alpha=0.15$, $\mu=0.025$, $\delta=0.025$, $\lambda_0=100$, $N_A=250$, and $q_{\mathrm{provider}}=q_{\mathrm{taker}}=0.5$, one gets the order book depth shown here. The results were averaged over 10^4 MCS. A fit to the log-normal distribution function $\phi_{\mathrm{LN}}(p-p_m)=a/x\exp\{-[\ln(p-p_m)-b]^2/c\}$ leads to the parameters $a=527\pm(1)$, $b=4.483\pm(4\times10^{-3})$, and $c=1.79\pm(10^{-2})$.

The antipersistent price behavior on short time scales can be found in actual financial time series [29] and is a consequence of the "mechanics" of the order book. At a constant order influx an executed market order of any kind automatically increases the probability that the next transaction price will be anticorrelated with the previous one.

The price change distributions shown in Fig. 6 exhibit no fat tails, but can rather well be approximated by a Gaussian distribution. Deviations from the Gaussian distribution are found for large price changes, where the Gaussian distribution overestimates the probability for these price changes. Note that the introduction of the exponentially distributed order placement depth influences only the order book depth,

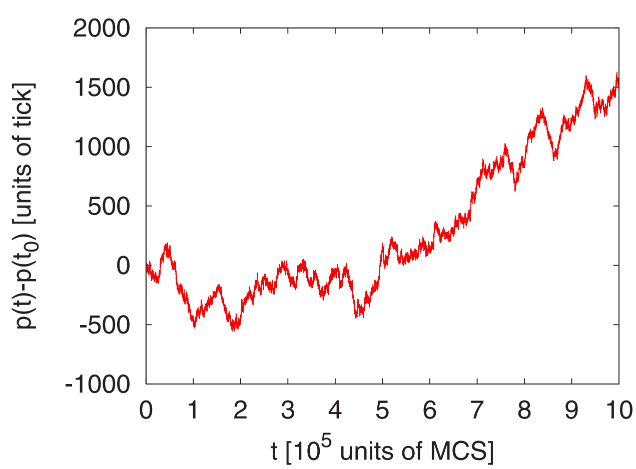

FIG. 4. (Color online) Exponential order placement depth: for $\alpha=0.15$, $\mu=0.025$, $\delta=0.025$, $\lambda_0=100$, $N_A=250$, and $q_{\mathrm{provider}}=q_{\mathrm{taker}}=0.5$, an exemplary price sequence of 10^6 MCS is shown.

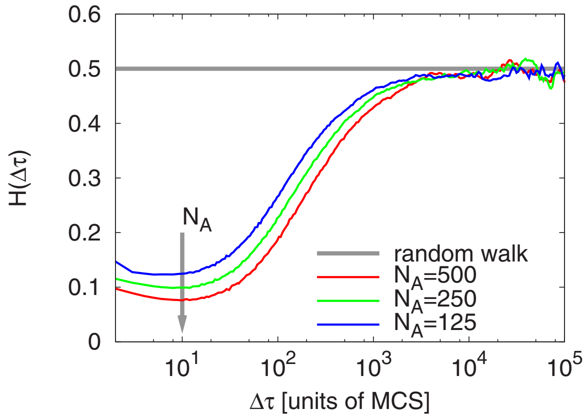

FIG. 5. (Color online) Exponential order placement depth: for the same parameters as in Fig. 4, the Hurst exponent is shown here for different numbers of agents N_A, averaged over 50 simulation runs each. The arrow points in the direction of larger values of the parameter N_A.

but the Hurst exponent and the price change distributions are not influenced in a qualitative way in comparison with results for the IID order placement approach.

III. PARAMETER SPACE

In order to get well founded information about a reasonable parameter selection for future extensions of this order book model, the parameter space of this basic version of the order book model has to be analyzed in detail. As the order book depth has to be considered as a vital criterion for the stability of the order book in the Monte Carlo simulations, it is examined quantitatively in dependence of the two most important parameters λ_0 and μ. For the other parameters, we

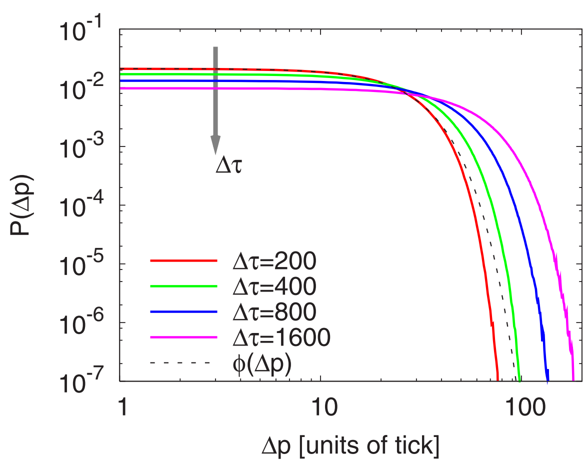

FIG. 6. (Color online) Exponential order placement depth: distributions of price changes for different $\Delta\tau$ for the same parameter values as in Fig. 4. Again the results were averaged over 50 simulation runs each. For $\Delta\tau$=200, the fit parameters for the Gaussian distribution function $\phi(\Delta p)=a\exp(-b\Delta p^2)$ are given by $a=2.122\times10^{-2}\pm(4\times10^{-5})$ and $b=1.388\times10^{-3}\pm(10^{-6})$. The arrow points in the direction of larger values of the parameter $\Delta\tau$.

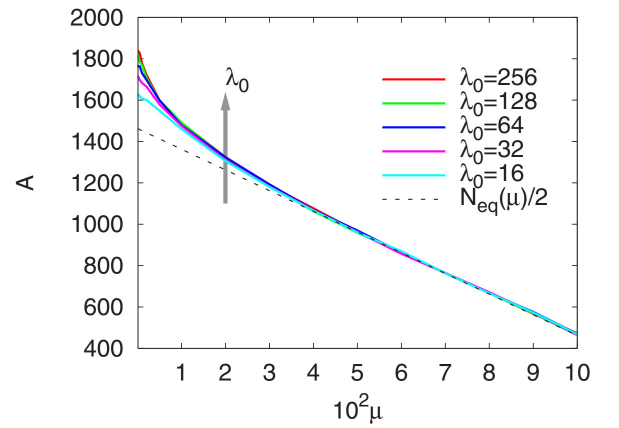

FIG. 7. (Color online) Parameter space: Fit parameter A as a function of the market order rate μ for different values of the order placement depth λ_0. The arrow points in the direction of larger values of the parameter λ_0.

use the values N_A=500, α=0.15, and δ=0.025 within this section. After a transient time of few thousands MCS the order book depth is stationary and log-normally distributed according to

$$P_{\mathrm{LN}}(x)=A\frac{1}{Sx\sqrt{2\pi}}\exp\left(-\frac{(\ln x-M)^2}{2S^2}\right), \qquad (11)$$

with the parameters A, S^2, and M. In our simulation runs, we wait 22 500 MCS, as the order book equilibrates within a few thousand MCS and then take the average over the following 2500 MCS.

In Fig. 7, the influence of the parameter μ on the scaling factor A is illustrated for various values of λ_0. A is the prefactor in Eq. (11), determining the size of the area under the curve of the log-normal distribution, as the integral over Eq. (11) gives the value A. We find that A decreases linearly for medium and large values of μ, independently of the value of λ_0. Only for small values of μ can deviations from this linear behavior be found. This result can be analytically explained if we consider the order book having already reached its equilibrium: in its stationary state, the order book depths on the bid and on the ask side are identical. Thus, the scaling factor A, as it is the area under the curve of the log-normal distribution, corresponds to half the total number of orders stored in the order book. The total number of limit orders at time $t+1$ can be described recursively by the order rates α, δ, and μ; the number of agents, N_A; and the number $N(t)$ of limit orders at time t by

$$N(t+1)=N(t)+\alpha N_A-(N(t)+\alpha N_A)\delta-\mu N_A. \qquad (12)$$

In equilibrium, one arrives at

$$\frac{N_{\mathrm{eq}}}{N_A}=\alpha\left(\frac{1}{\delta}-1\right)-\frac{\mu}{\delta} \qquad (13)$$

for the number of limit orders per liquidity provider. If defining an effective limit order rate $\alpha^*=\alpha(1-\delta)$, this relation results in

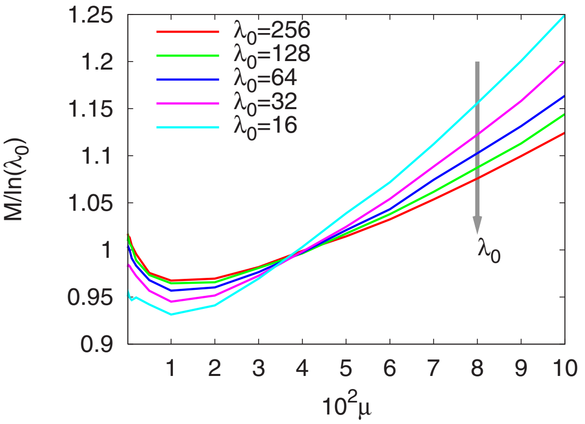

FIG. 8. (Color online) Parameter space: fit parameter M as a function of the market order rate μ for different values of the order placement depth λ_0. The arrow points in the direction of larger values of the parameter λ_0.

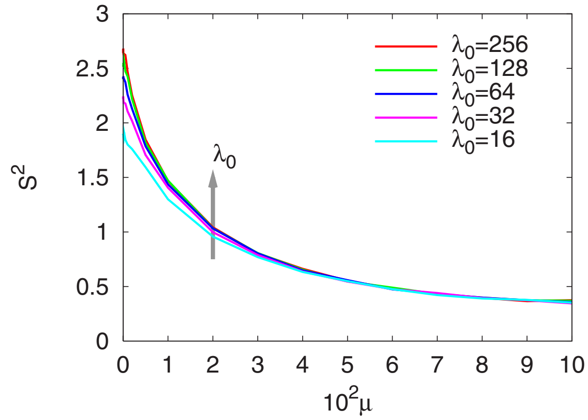

FIG. 9. (Color online) Parameter space: fit parameter S^2 as a function of the market order rate μ for different values of the order placement depth λ_0. The arrow points in the direction of larger values of the parameter λ_0.

$$\frac{N_{\rm eq}}{N_A}\delta = \alpha^* - \mu. \tag{14}$$

A stable order book is therefore achieved on average if the conditions $\alpha^* > \mu$ and $\delta > 0$ are fulfilled.

Inserting the values we used in the simulations in this equation, we get $N_{\rm eq}/2 = 1462.5 - 10^4\mu$. This relation is drawn as a dashed line in Fig. 7. We find that the theoretical considerations and the simulation results coincide for medium and large values of μ. The deviations for small μ result from the fact that the order book depth does not approach a log-normal distribution anymore, such that a fit to the log-normal distribution function is no longer valid in this case.

As shown in Fig. 8, also the center point M of the log-normal distribution is affected by a modification of the parameters λ_0 and μ. The larger the market order rate μ, the more limit orders are removed from the bid and the ask side of the order book by transactions. The more limit orders, however, are removed from the inside of the distributions at the bid and the ask side, the more M departs from the midpoint p_m. As can be clearly seen in Fig. 8, this behavior occurs only for market order rates $\mu \gtrsim 10^{-2}$. An increase of the exponentially distributed order placement depth leads to a larger value of M.

For completion, the squared variance S^2 of the log-normal distribution is shown in Fig. 9 as a function of μ for different values of λ_0. The parameter S^2 decreases with increasing μ, this decrease being rather independent of the value of λ_0 for large values of μ.

Significant deviations from the log-normal behavior we discussed so far are found for parameter combinations with small values for either λ_0 or μ: for very small μ, the distribution of the order book depth resembles the exponentially distributed order placement depth, as the market orders are not sufficient to remove the limit orders at best ask and best bid, respectively. For small λ_0, limit orders are placed closely around the midpoint, There are so many limit orders at best bid and best ask, respectively, that there is a limiting value for μ (depending on λ_0) under which the market orders are not able to shift p_b and p_a anymore. In both cases, when considering the extreme situations of vanishing μ and λ_0, the order book freezes in a way that p_a and p_b become constant and $s = p_a - p_b = 1$. The price then jumps between p_a and p_b.

Another limiting situation is given by large market order rates μ. The more market orders are submitted, the larger becomes the probability that one side of the order book is completely cleared and trading stops. The total number of limit orders, $N(t)$, which are stored in the order book at time t, is crucial for the stability of the simulation. Independently of the value chosen for μ, fluctuations in the selection of agents can empty the bid or ask side of the order book, if the order book depth is too small. This can result in a price crash or a price explosion. Especially, at the beginning of the simulation, only limit orders have to be placed in a preopening phase which we choose to last 10 MCS.

IV. AUGMENTATION OF THE MODEL

A. Deterministic perturbation

After the investigation of the parameter space of the basic version of the order book model, extensions of this model shall be considered. The basic model is only able to reproduce the antipersistent price behavior on short time scales as well as the diffusive price behavior on long time scales, which are both also found at real financial markets. However, the persistent behavior on medium time scales is not reproduced in this basic version which models a stationary market. Moreover, there are no fat-tailed price change distributions found in the basic variant. So far, we always considered a symmetry between the buy and sell probabilities, both for the group of liquidity providers and for the group of liquidity takers, as $q_{\rm provider} = q_{\rm taker} \equiv 1/2$. This is an appropriate assumption presuming a stationary behavior of a financial market. Such a stationary behavior is, however, not compatible with financial-economical conditions. Real order rates indicate asymmetries. In a bull market, an increased buy probability can be found and one can measure an increased sell probability in a bear market. This applies not only for long-

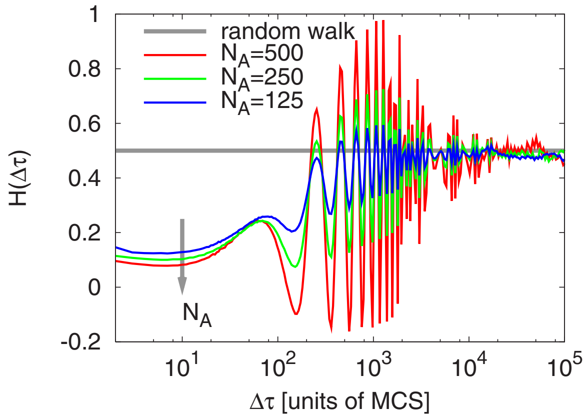

FIG. 10. (Color online) Deterministic perturbation realized by a sawtooth modulation of q_{taker}: Hurst exponent $H(\Delta\tau)$ for different agent numbers N_A as a function of $\Delta\tau$. The parameter values used are $\alpha=0.15$, $\mu=0.025$, $\delta=0.025$, and $\lambda_0=100$. The arrow points in the direction of larger values of the parameter N_A.

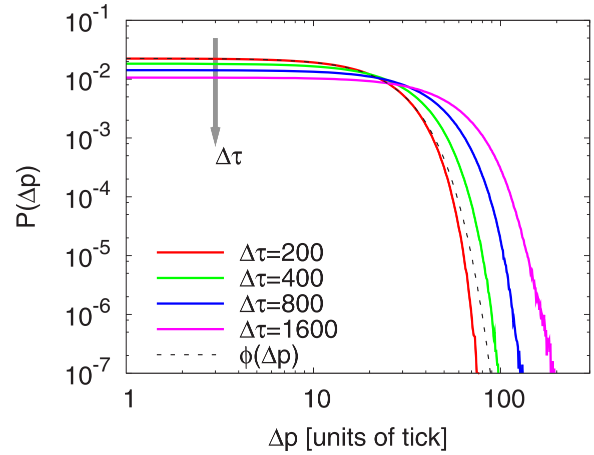

FIG. 11. (Color online) Deterministic perturbation realized by a sawtooth modulation of q_{taker}: Return distributions for different values of $\Delta\tau$ and for $N_A=500$. For $\Delta\tau=200$, the function $\phi(\Delta p) = a\exp(-b\Delta p^2)$ can be numerically fitted with $a=2.251\times10^{-2}\pm(3\times10^{-5})$ and $b=1.573\times10^{-3}\pm(5\times10^{-6})$. The arrow points in the direction of larger values of the parameter $\Delta\tau$.

time movements, but also for short-time trends on intraday time scales. Therefore, the symmetry $q_{\text{provider}}=1-q_{\text{provider}}=1/2$ and $q_{\text{taker}}=1-q_{\text{taker}}=1/2$ used so far will now be broken in a way that q_{provider} stays at its constant value of 1/2, whereas q_{taker} shall be changed in time but still have an average value of 1/2. One could say that by the modulation of q_{taker}, the market is "deflected" from its stationary state.

For the practical realization of such a market deflection, first of all a deterministic symmetry disturbance shall be investigated. As a simple ansatz, q_{taker} shall be varied by a sawtooth modulation

$$q_{\text{taker}} = \begin{cases} \dfrac{1}{2} + t\Delta S & \text{for } 0 \leqslant t \leqslant S/\Delta S, \\ \dfrac{1}{2} + 2S - t\Delta S & \text{for } S/\Delta S \leqslant t \leqslant 3S/\Delta S, \\ \dfrac{1}{2} - 4S + t\Delta S & \text{for } 3S/\Delta S \leqslant t \leqslant 4S/\Delta S, \end{cases} \tag{15}$$

which is periodically repeated and for which we use the amplitude value $S=1/20$ and the step size $\Delta S=1/1000$, by which the variable q_{taker} is changed after each MCS time step denoted by t. This sawtooth modulation has a period of $4S/\Delta S$ and thus q_{taker} returns to the value of 1/2 after every $\tau_r=2S/\Delta S=100$ MCS. We choose a sawtooth modulation because it exhibits a constant residence distribution.

Using such a sawtooth modulation for the implementation of an asymmetry in the order flow, price time series with new properties are achieved. An averaged Hurst exponent is shown exemplarily in Fig. 10. Each simulation lasted 10^6 MCS; the average was taken of 50 simulation runs. The deterministic modulation of q_{taker} with a period of 200 MCS is reflected in the mean-square displacement, leading to quasiperiodic oscillations of the Hurst exponent $H(\Delta\tau)$ at medium time scales. The Hurst exponent oscillates a few times with a period of approximately 200 MCS before it shows a diffusive behavior for the price development on long time scales. The more agents are trading in the order book, the larger are the amplitudes of the oscillations. Such periodic oscillations of the Hurst exponent have nothing in common with the behavior the Hurst exponent is showing for data from real financial markets.

Figure 11 shows the corresponding distributions of the price changes. Like for the basic version of the order book model with a symmetric order flow, no fat-tailed price change distributions are found here either. Instead a fit to a Gaussian distribution works rather well, except that it again overestimates the probability for large price changes. Concluding, we can state that with this approach using a deterministic perturbation we were not successful in producing a more realistic price behavior, but we were able to change the behavior of the Hurst exponent especially on medium time scales with this periodic modulation.

B. Stochastic perturbations

As any deterministic perturbation with a discrete return time spectrum will be reflected in oscillations of the Hurst exponent, we focus now on stochastic perturbations with continuous return time distributions. As a straightforward approach, we change the value of q_{taker}, which shall keep an average value of 1/2, in a bounded random walk and again keep $q_{\text{provider}}=1/2$ constant. This approach was already introduced and analyzed in [27]. There one gets an antipersistent price behavior on short time scales, a persistent price behavior on medium time scales, and a diffusive behavior on long time scales. However, the maximum values of $H(\Delta\tau)$ found for the medium time scales are too large; one gets even values up to 0.9, but according to [3,4], only a maximum value of $\approx$0.6 may be found. This is also the case for other financial time series; see, e.g., the results in [30] for foreign exchange time series. The price change distributions of the bounded random walk approach show a bimodal shape, which is in contrast to the behavior of price changes in time series of real financial markets. This wrong behavior is caused by the constant residence distribution of a bounded random walk. As q_{taker} returns slowly from the extreme areas of the modulation to the mean value of 1/2, strong trend phases are created, leading in turn to the bimodal shape.

From these insights, one arrives straightforwardly at using a feedback random walk, i.e., a random walk with increased probability for returning to the mean value. This approach, as shall be mentioned already here, is able to produce a nontrivial Hurst exponent, which is comparable to those of time series found at real financial markets, and this approach generates an almost Gaussian price change distribution.

This feedback random walk, which is again only applied to q_{taker}, whereas q_{provider} stays constant at 1/2, works as follows: at the beginning of the Monte Carlo simulation, q_{taker} starts at the mean value of 1/2. The variable q_{taker} is incremented and decremented by a value of ΔS after each MCS. But in contrast to a standard random walk, the probability for returning to the average value of 1/2 is given by $1/2+|q_{\text{taker}}(t)-1/2|$ and thus the probability for departing from the mean value is given by $1/2-|q_{\text{taker}}(t)-1/2|$. This feedback random walk has the tendency to return back to its average value $\langle q_{\text{taker}}\rangle=1/2$ more often compared to the bounded random walk, which can be easily understood: the expectation value of $q_{\text{taker}}(t+1)$ is given by

$$\langle q_{\text{taker}}(t+1)\rangle = [q_{\text{taker}}(t)-\Delta S]q_{\text{taker}}(t) + [q_{\text{taker}}(t)+\Delta S][1-q_{\text{taker}}(t)], \qquad (16)$$

both for $q_{\text{taker}}(t)\geqslant 1/2$ and for $q_{\text{taker}}(t)\leqslant 1/2$. Thus, we generally have

$$\langle q_{\text{taker}}(t+1)\rangle - q_{\text{taker}}(t) = \Delta S[1-2q_{\text{taker}}(t)]\begin{cases} >0 & \text{if } q_{\text{taker}}(t)<1/2, \\ <0 & \text{if } q_{\text{taker}}(t)>1/2, \end{cases} \qquad (17)$$

such that this feedback random walk on average approaches its expectation value 1/2. The stochastic process of the feedback random walk is characterized by a continuous return time spectrum, qualitatively comparable to that of the bounded random walk. However, the residence distribution of the probability q_{taker} exhibits an almost Gaussian shape in contrast to the bounded random walk, for which it is uniform. For the simulation results shown here, we again used $\Delta S=1/1000$.

Figure 12 shows the behavior of the Hurst exponent $H(\Delta\tau)$, which was averaged over 50 simulation runs lasting 10^6 MCS each. Again one finds an antipersistent behavior on short time scales, a persistent behavior on medium time scales, and a diffusive regime on long time scales. The maximum of the Hurst exponent increases with increasing agent number N_A. If comparing the results shown here with measurements of the Hurst exponent of financial time series achieved in real financial markets, we find that agent numbers in the range $150\leqslant N_A\leqslant 500$ are best able to reproduce a realistic maximum value of the Hurst exponent.

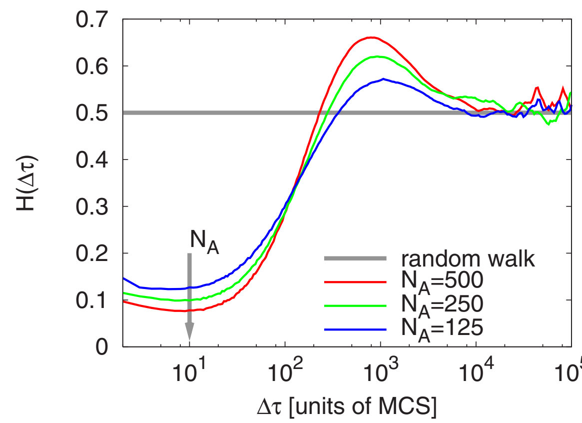

FIG. 12. (Color online) Stochastic perturbation realized by the feedback random walk modulation of q_{taker}: Hurst exponent $H(\Delta\tau)$ for different agent numbers N_A as a function of $\Delta\tau$. The parameters used are $\alpha=0.15$, $\mu=0.025$, $\delta=0.025$, and $\lambda_0=100$. The arrow points in the direction of larger values of the parameter N_A.

In Fig. 13, price change distributions for this approach are shown. In contrast to the bounded random walk approach, we obtain here more reasonable, almost Gaussian-shaped price change distributions.

However, also this extension of the order book model fails in producing fat-tailed price change distributions. Note that identical results can be achieved for a temporal modulation of q_{provider} and a constant $q_{\text{taker}}=1/2$ and if q_{provider} and q_{taker} are independently of each other changed in time by a feedback random walk. After the following section, in which the relationship between the Hurst exponent and the autocorrelation is analyzed, a further extension will be introduced to create fat tails.

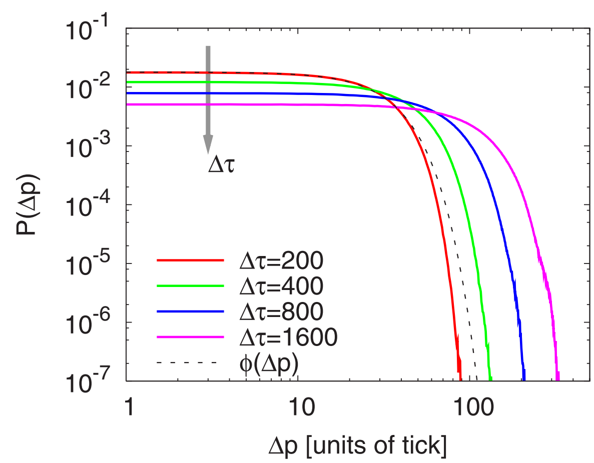

FIG. 13. (Color online) Stochastic perturbation realized by the feedback random walk modulation of q_{taker}: price change distributions for different values of $\Delta\tau$ and for $N_A=500$. The other parameters are the same as in Fig. 12. Fitting the function $\phi(\Delta p) = a\exp(-b\Delta p^2)$ to the data for $\Delta\tau=200$, one gets $a=1.79\times 10^{-2}\pm(10^{-4})$ and $b=9.9\times 10^{-4}\pm(10^{-5})$. The arrow points in the direction of larger values of the parameter $\Delta\tau$.

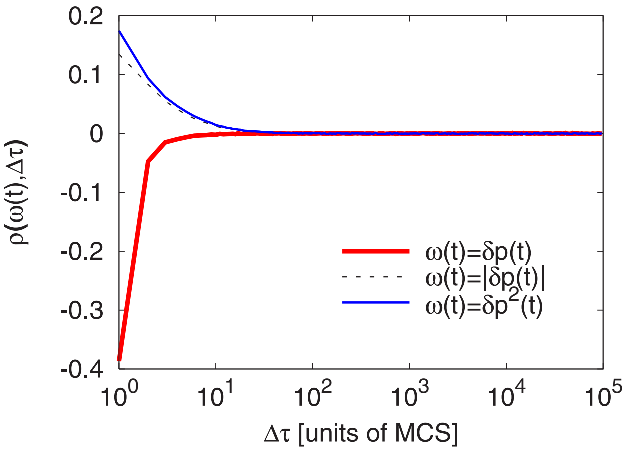

FIG. 14. (Color online) Autocorrelation $\varrho(\delta p(t),\Delta\tau)$ of price changes, autocorrelation $\varrho(|\delta p(t)|,\Delta\tau)$ of absolute price changes, and autocorrelation $\varrho(\delta p(t)^2,\Delta\tau)$ of squared price changes with $\delta p(t)=p(t+1)-p(t)$. The parameters are the same as for the results in Fig. 12 with $N_A=125$.

C. Hurst exponent and autocorrelation

The Hurst exponent is often used for the characterization of stochastic processes. Often a connection to autocorrelations is drawn, which describes memory effects within stochastic processes. In the literature, it is widely assumed that a Hurst exponent of $H\neq 1/2$ implies long-time correlations, but recent theoretical work [31,32] shows that this is not necessarily true. A persistent behavior with $H=1/2$ also occurs for Markov processes (i.e., processes without memory) in the case nonstationary increments. This is what we employed in the last section. This result affects also the interpretation of the Hurst exponent of financial market time series. In this context, the Hurst exponent is used in order to measure the efficiency of a market. The Hurst exponent $H=1/2$ of the random walk corresponds to an efficient market. However, this criterion alone is not sufficient for the determination of the efficiency according to the results of [31,32]. From a measurement of the Hurst exponent alone, the existence of a long-time memory cannot be derived or the existence of an efficient financial market be deduced. Instead, an additional investigation of autocorrelations is necessary.

Figure 14 shows the autocorrelation of the price change, of the absolute price change, and of the quadratic price change time series created by the same parameter set as was used for the corresponding Hurst exponent shown in Fig. 12 for $N_A=125$, but instead of averaging over 50 simulations, here only one simulation was performed whose calculation time was increased to 10^7 MCS in order to improve the statistics of the autocorrelations. The autocorrelation $\varrho(\omega(t),\tau)$ of a time-dependent function $\omega(t)$ is given by

$$\varrho(\omega(t),\tau)=\frac{\langle\omega(t+\tau)\omega(t)\rangle-\langle\omega(t)\rangle^2}{\langle\omega(t)^2\rangle-\langle\omega(t)\rangle^2} \tag{18}$$

in the stationary case. We find that a nontrivial Hurst exponent does not coincide with long-time correlations, as the autocorrelation functions converge rather fast to zero. Although the price time series show on short time scales an antipersistent and on medium time scales a persistent price behavior, no nonvanishing autocorrelation $\varrho(\delta p(t),\delta\tau)$ with $\delta p(t)=p(t+1)-p(t)$ can be recognized for $\Delta\tau>10$. The autocorrelation functions for the quadratic price change and for the absolute price change, which are also shown in Fig. 14, are positive and converge roughly exponentially towards zero. This result can be interpreted as volatility clustering on short time scales, which is also a stylized empirical fact of financial data records and analyzed in detail in [33–35]. The analysis in [34], for example, shows for data sets from the New York Stock Exchange that volatility correlations are power laws on time scales from 1 day to 1 year and that the exponent is not unique.

These results of the order book model presented here are in good agreement with the results in [31,32], showing that a Hurst exponent $H>1/2$ implies not necessarily long-time correlations.

D. Dynamic order placement depth

In the previous approaches, a constant order placement depth λ_0 was used. According to the remarks in [36], one can expect an equilibrium on real markets between the effective costs of a market order and of a limit order. If the spread is large, the submission of a limit order is advantageous. In this case, the execution of a limit sell order at best ask or the execution of a limit buy order at best bid is connected with a smaller risk than in the situation of a small spread. The risk consists of the establishment of a market trend, which is directed against the position entered by the limit order, leading to a loss. However, if a small risk exists, also other liquidity providers are ready to place orders around a smaller spread. The spread decreases down to a level at which the risk and thus the effective costs of a market order and of a limit order are comparable [36].

From the above discussion it follows that the liquidity providers can reduce their risk exposure by adapting their limit order placement depth to the prevailing market conditions. In trendless market phases, in which no large price fluctuations are to be expected, liquidity providers place their limit orders close to the midpoint, in order to be able to participate in small price movements. But if the volatility increases, which can be, e.g., recognized in strong trend phases, the risk of the liquidity providers to possess positions which are orientated against the prevailing market trend increases. In these market phases, also the probability decreases to close such a position on the opposite side of the order book by a limit order without loss. Therefore, it is an obvious consequence that liquidity providers adapt their characteristic order placement depth to changing conditions. The market risk is reduced by an enlargement of this characteristic order placement depth in times of a trend. In the order book model the strength of a trend is given by the deviation of the market order influx from the symmetric case $q_{\text{taker}}=1/2$. We therefore replace the constant order placement depth λ_0 by

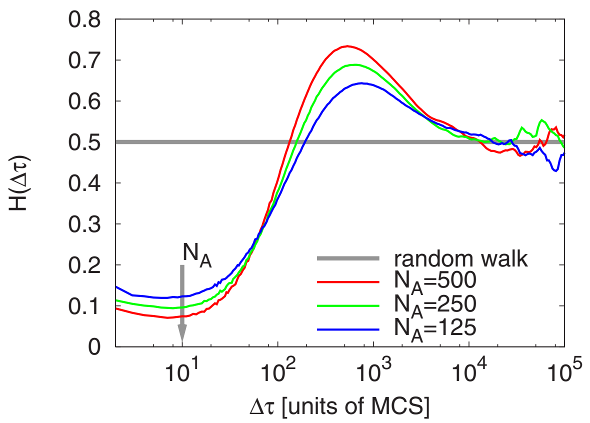

FIG. 15. (Color online) Dynamic order placement depth: Hurst exponent $H(\Delta\tau)$ for different values of N_A with $\lambda_0=100$, $C_\lambda=10$, $\alpha=0.15$, $\mu=0.025$, $\delta=0.025$, and $\Delta S=0.001$. The arrow points in the direction of larger values of the parameter N_A.

$$\lambda(t) = \lambda_0\left(1 + \frac{\left|q_{\text{taker}}(t) - \frac{1}{2}\right|}{\sqrt{\left\langle\left[q_{\text{taker}}(t) - \frac{1}{2}\right]^2\right\rangle}} C_\lambda\right). \tag{19}$$

For $C_\lambda=0$, this further extension of the order book model corresponds to the variant of the order book model with static order placement parameter. The averaged value $\langle(q_{\text{taker}}(t)-\frac{1}{2})^2\rangle$ is determined in a separate Monte Carlo simulation lasting 10^6 MCS before the main simulation starts.

The results of coupling the order placement depth λ to the prevailing trend are shown in Figs. 15 and 16. The average was taken of 50 simulation runs lasting 10^6 MCS each. With

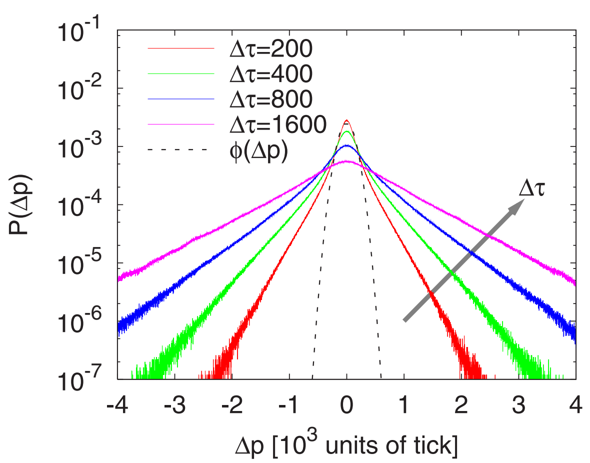

FIG. 16. (Color online) Dynamic order placement depth: distributions of price changes for the same parameter values as in Fig. 15. One clearly finds fat tails. A Gaussian approximation $\phi(\Delta p) = a\exp(-b\Delta p^2)$ for $\Delta\tau=200$ with the parameters $a=2.54\times10^{-3}\pm(10^{-5})$ and $b=2.8\times10^{-5}\pm(2\times10^{-7})$ strongly underestimates the probability for large price changes. The arrow points in the direction of larger values of the parameter $\Delta\tau$.

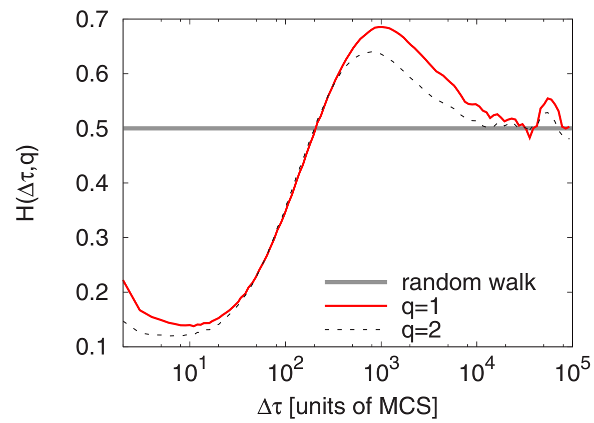

FIG. 17. (Color online) Dynamic order placement depth: Hurst exponent $H(\Delta\tau,q)$ for different values of q with $\lambda_0=100$, $C_\lambda=10$, $\alpha=0.15$, $\mu=0.025$, $\delta=0.025$, $\Delta S=0.001$, and $N=125$ for a single simulation over 5×10^7 MCS.

this additional extension of our order book model, which already includes the feedback random walk approach as first extension, it is now possible to produce not only a persistent Hurst exponent for medium time scales but also fat-tailed price change distributions.

In Fig. 17 the Hurst exponent $H(\Delta\tau)\equiv H(\Delta\tau,q=2)$ is shown in comparison with $H(\Delta\tau,q=1)$, as defined in Eq. (10). The scaling exponent for absolute price changes $H(\Delta\tau,q=1)$ exhibits a larger persistent behavior on medium time scales and the antipersistence on short time scales is smaller, consistent with earlier findings [34,35].

A widely discussed problem in physics is the origin of the fat-tailed price distributions generated by the complex system financial markets. Often the truncated Lévy distribution [6,7,37] is considered as an approximation of fat-tailed price change distributions found at real financial markets. A Lévy-stable distribution is scale invariant and exhibits an infinite variance [6]. The truncated Lévy distribution (TLD) [38,39] has finite variance and shows scaling behavior in a large, but finite interval. However, also the possibility of power-law tails is mentioned at great length in the physics community [6]. In [40] it is shown for the S&P 500 index that the price change distributions for time lags $\Delta\tau\leq4$ days are consistent with a power-law behavior with an exponent $\alpha\approx3$, outside the stable Lévy regime $(0<\alpha_L<2)$. For larger time lags a slow convergence to Gaussian behavior was found.

One possibility to analyze if a process generates true Lévy distributions or not, is given by shuffling the short term returns of the time series [40,41]. If after shuffling the longer term returns maintain the same power-law exponent, then the stochastic process generates true Lévy distributions, as a stable Lévy process is invariant under folding. However, if after shuffling of the short-term returns the longer term returns appear Gaussian, then the distribution tails are not as "fat" as Lévy fat tails as a result of the central limit theorem. However, this argument assumes that the return distributions on short time scales are independent from each other. In our case one can find the results of Fig. 16 after shuffling returns on the time horizon of one MCS in Fig. 18. It is obvious that the return distributions show a convergence to a Gaussian

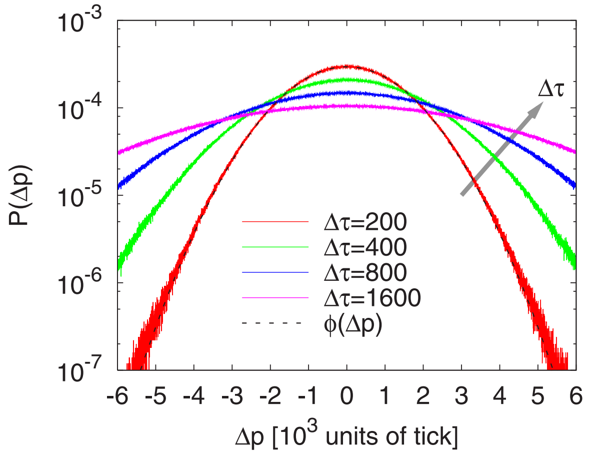

FIG. 18. (Color online) Dynamic order placement depth: Distributions of shuffled price changes for the same parameter values as in Fig. 16. After shuffling the returns on the time horizon of one MCS the longer term aggregate returns appear Gaussian. For $\Delta\tau=200$, the function $\phi(\Delta p)=a\exp(-b\Delta p^2)$ can be numerically fitted with $a=2.9505\times10^{-4}\pm(4\times10^{-8})$ and $b=2.7419\times10^{-7}\pm(9\times10^{-11})$. The arrow points in the direction of larger values of the parameter $\Delta\tau$.

behavior. This shows that the fat tails are indeed a result of the dynamic order entry depth, which is coupled with the strongness of the prevailing market trend. So we find a qualitatively similar result of data shuffling as for independent returns, however for a different reason.

In the following, a TLD as described by Koponen [39] will be used to fit the fat-tailed price changing distributions of Fig. 16. It features a smooth exponential cutoff and its characteristic function is given through

$$\Lambda_{\alpha_L,c_1,l}(f_n)=\exp\left(c_0-c_1\frac{(f_n^2+1/l^2)^{\alpha_L/2}}{\cos(\pi\alpha_L/2)}\cos[\alpha_L\arctan(l|f_n|)]\right), \tag{20}$$

with the scaling factors

$$c_0=\frac{l^{-\alpha_L}}{\cos(\pi\alpha_L/2)} \tag{21}$$

and c_1. As only the characteristic function of the TLD is given in an analytic form, a discrete Fourier transformation of the price change distributions of our order book model is necessary, in order to fit $\Lambda_{\alpha_L,c_1,l}(f_n)$ to the simulation data.

TABLE I. Fit parameters α_L, c_1, and l of the TLF distribution for $C_\lambda=1$ in dependence on $\Delta\tau$: the errors given originate only from the process of fitting.

$\Delta\tau$	α_L	c_1	l
200	1.606±0.002	4271±40	3635±52
400	1.477±0.003	3799±52	5433±72
800	1.426±0.004	5542±113	8510±124
1600	1.490±0.005	16901±460	12073±210

TABLE II. Corresponding to Table I, the fit parameters are given for $C_\lambda=5$.

$\Delta\tau$	α_L	c_1	l
200	1.232±0.002	1872±23	20903±252
400	1.181±0.002	2209±32	39040±497
800	1.148±0.003	3215±73	70650±1293
1600	1.223±0.005	11058±398	75158±1559

The values of the fit parameters are documented in Tables I–III for some exemplary values of C_λ. Comparing to the Lévy exponent α_L measured for real financial data time series, which takes values in the range of $\approx1.4-1.5$ [38,42], $C_\lambda=1$ seems to be the best approximation to real market behavior. In Fig. 19, the Lévy exponent α_L is shown as a function of $\Delta\tau$ for different values of C_λ. Again one clearly finds that α_L stays in the correct interval for $C_\lambda\approx1$, whereas a larger value of C_λ leads to too small values and a smaller one to too large values.

Looking closely at Fig. 19, we furthermore find an interesting relation between the Lévy exponent α_L and the parameter C_λ for not too long time lags $\Delta\tau$: α_L depends on C_λ via a power law according to $\alpha_L=\xi\, C_\lambda^{-0.15}$, with the prefactor ξ only depending on the time lag.

The coupling of the order placement depth to the prevailing trend thus leads to fat-tailed price change distributions. This property is, however, independent of the persistence of the price time series on medium time scales, as we already got this persistence by imposing nonstationary increments of the price process. On the other hand, one can also get fat tails without $H>1/2$ for medium time scales. We achieved this scenario by determining $\lambda(t)$ according to a mean reverting random walk as in Eq. (19) but using always symmetric order placement behavior.

V. CONCLUSION AND OUTLOOK

We examined the order book model, which we introduced in [27] as a multiagent system for the modeling of financial markets, in more detail. According to a bottom-up approach, we started out with a simple model variant, whose key feature is the distinction between liquidity providers and liquidity takers. We showed that liquidity providers have a systematic advantage by the possibility of transmitting limit orders, compared to the group of the liquidity takers, who are only allowed to use market orders. Moreover, this simple variant

TABLE III. Corresponding to Table I, the fit parameters are given for $C_\lambda=10$.

$\Delta\tau$	α_L	c_1	l
200	1.101±0.002	1427±20	83954±1376
400	1.046±0.003	1583±29	298864±14000
800	1.039±0.004	2773±83	591232±52480
1600	1.143±0.006	12696±642	265875±10860

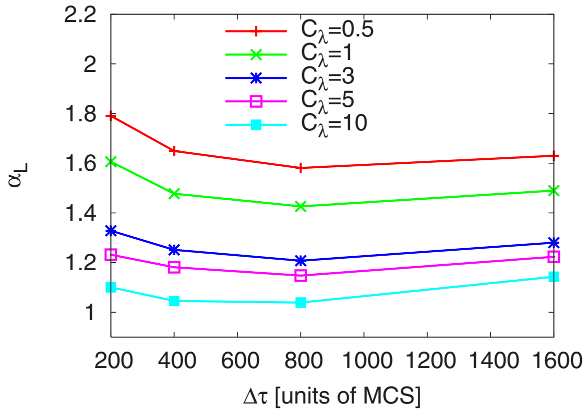

FIG. 19. (Color online) Dynamic order placement depth λ: Lévy exponent α_L as a function of $\Delta\tau$ for different C_λ.

is able to reproduce the results of [25,26]. But it does not correspond to order books at real financial markets at all, because it exhibits a constant nonzero order book depth, even far away from the midpoint. The introduction of an exponentially distributed order placement depth of limit orders creates a log-normal-distributed depth of the order book. This model variant was then regarded as the basic version of the order book model. Its price time series possesses an antipersistent price behavior on short time scales which is due to the order book structure. On medium and long time scales the Hurst exponent converges towards a diffusive regime. The price change distributions exhibit an almost Gaussian shape. This basic version of the order book model, which is characterized by a symmetry created by identical buy and sell probabilities, describes a stationary market. However, when one additionally introduces a symmetry disturbance, the order book model is displaced from its stationary state. This extension is implemented by a temporal modulation of the buy probability q_{taker} of the liquidity takers or the buy probability q_{provider} of the liquidity providers. Qualitatively identical results are achieved if both probabilities are modulated independently of each other. Employing a feedback random walk to introduce micro market trends into the market, one additionally obtains a persistent price behavior on medium time scales. However, no fat tails can be reproduced with such a symmetry-breaking extension of the order book model. When one furthermore couples the characteristic order placement depth to the prevailing market trend, widened price change distributions are achieved, with so-called fat tails. A truncated Lévy distribution can be fitted to these price changes. Thus, with these extensions of our order book model, we could demonstrate that the generation of a nontrivial Hurst exponent is independent of the generation of fat tails. This disproves the implication which can be often found in the literature that a persistent price behavior corresponds to non-Gaussian price changes. Furthermore, we are able to support the statement in [31,32] that $H > 1/2$ implies not necessarily long time correlations. We plan to examine further characteristics of our order book model in the future and to investigate whether other mechanisms than those extensions described and examined in this paper can lead to fat tails or a nontrivial Hurst exponent.

[1] M. F. M. Osborne, Oper. Res. **7**, 145 (1959).
[2] M. F. M. Osborne, *The Stock Market and Finance From a Physicist's Viewpoint* (Crossgar Press, Minneapolis, 1995).
[3] B. Mandelbrot, J. Business **36**, 394 (1963).
[4] B. Mandelbrot, J. Business **37**, 393 (1964).
[5] J.-P. Bouchaud and M. Potters, *Theory of Financial Risks—From Statistical Physics to Risk Management* (Cambrigde University Press, Cambridge, England, 2000).
[6] R. N. Mantegna and H. E. Stanley, *An Introduction to Econophysics—Correlations and Complexity in Finance* (Cambrigde University Press, Cambridge, England, 2000).
[7] W. Paul and J. Baschnagel, *Stochastic Processes: From Physics to Finance* (Springer, Heidelberg, 2000).
[8] L. Bachelier, Ph.D. thesis, Faculté des Sciences de Paris, 1900.
[9] K. Itô, Proc. Imp. Acad. (Tokyo) **20**, 519 (1944).
[10] P. A. Samuelson, Ind. Management Rev. **6**, 13 (1965).
[11] A. Einstein, Ann. Phys. **17**, 549 (1905).
[12] F. Black and M. Scholes, J. Polit. Econ. **81**, 637 (1973).
[13] R. Merton, Bell J. Econ. Manage. Sci. **4**, 141 (1973).
[14] P. Bak, M. Paczuski, and M. Shubik, Physica A **246**, 430 (1997).
[15] S. Maslov, Physica A **278**, 571 (2000).
[16] S. Maslov and M. Mills, Physica A **299**, 234 (2001).
[17] D. Challet and R. Stinchcombe, Quant. Finance **3**, 155 (2003).
[18] R. Cont and J.-P. Bouchaud, Macroecon. Dyn. **4**, 170 (2000).
[19] D. Stauffer and T. J. P. Penna, Physica A **256**, 284 (1998).
[20] D. Stauffer and D. Sornette, Physica A **271**, 496 (1999).
[21] T. Lux and M. Marchesi, Nature (London) **397**, 498 (1999).
[22] I. Giardina, J.-P. Bouchaud, and M. Mézard, Physica A **299**, 28 (2001).
[23] A. Krawiecki, J. A. Holyst, and D. Helbing, Phys. Rev. Lett. **89**, 158701 (2002).
[24] D. Challet, M. Marsili, and Y.-C. Zhang, Physica A **294**, 514 (2001).
[25] M. G. Daniels, J. D. Farmer, L. Gillemot, G. Iori, and E. Smith, Phys. Rev. Lett. **90**, 108102 (2003).
[26] E. Smith, J. D. Farmer, L. Gillemot, and S. Krishnamurthy, Quant. Finance **3**, 481 (2003).
[27] T. Preis, S. Golke, W. Paul, and J. J. Schneider, Europhys. Lett. **75**, 510 (2006).
[28] G. A. Darbellay and D. Wuertz, Physica A **287**, 429 (2000).
[29] T. Preis, J. J. Schneider, and W. Paul (unpublished).
[30] M. Ausloos, Physica A **285**, 48 (2000).
[31] K. E. Bassler, G. H. Gunaratne, and J. L. McCauley, Physica A **369**, 343 (2006).
[32] J. L. McCauley, G. H. Gunaratne, and K. E. Bassler, e-print arXiv:cond-mat/0609671.
[33] R. Cont, Quant. Finance **1**, 223 (2001).
[34] M. Pasquini and M. Serva, Econ. Lett. **65**, 275 (1999).
[35] M. Serva, U. L. Fulco, I. M. Gléria, M. L. Lyra, F. Petronia,

and G. M. Viswanathan, Physica A **363**, 393 (2006).
[36] M. Wyart, J.-P. Bouchaud, J. Kockelkoren, M. Potters, and M. Vettorazzo, e-print arXiv:physics/0603084 (2006).
[37] J. Voit, *The Statistical Mechanics of Financial Markets* (Springer, Berlin, 2003).
[38] R. N. Mantegna and H. E. Stanley, Phys. Rev. Lett. **73**, 2946 (1994).
[39] I. Koponen, Phys. Rev. E **52**, 1197 (1995).
[40] P. Gopikrishnan, V. Plerou, L. A. Nunes Amaral, M. Meyer, and H. E. Stanley, Phys. Rev. E **60**, 5305 (1999).
[41] G. M. Viswanathan, U. L. Fulco, M. L. Lyra, and M. Serva, Physica A **329**, 273 (2003).
[42] R. N. Mantegna and H. E. Stanley, Nature (London) **376**, 46 (1995).