Statistical identification with hidden Markov models of large order splitting strategies in an equity market

To cite this article: Gabriella Vaglica *et al* 2010 *New J. Phys.* **12** 075031

View the article online for updates and enhancements.

Statistical identification with hidden Markov models of large order splitting strategies in an equity market

Gabriella Vaglica[1], Fabrizio Lillo[1,2,3] and Rosario N Mantegna[1]

[1] Dipartimento di Fisica e Tecnologie Relative, Università di Palermo,
Viale delle Scienze, I-90128, Palermo, Italy
[2] Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA
E-mail: lillo@unipa.it

Abstract. Large trades in a financial market are usually split into smaller parts and traded incrementally over extended periods of time. We address these large trades as hidden orders. In order to identify and characterize hidden orders, we fit hidden Markov models to the time series of the sign of the tick-by-tick inventory variation of market members of the Spanish Stock Exchange. Our methodology probabilistically detects trading sequences, which are characterized by a significant majority of buy or sell transactions. We interpret these patches of sequential buying or selling transactions as proxies of the traded hidden orders. We find that the time, volume and number of transaction size distributions of these patches are fat tailed. Long patches are characterized by a large fraction of market orders and a low participation rate, while short patches have a large fraction of limit orders and a high participation rate. We observe the existence of a buy–sell asymmetry in the number, average length, average fraction of market orders and average participation rate of the detected patches. The detected asymmetry is clearly dependent on the local market trend. We also compare the hidden Markov model patches with those obtained with the segmentation method used in Vaglica *et al* (2008 *Phys. Rev.* E **77** 036110), and we conclude that the former ones can be interpreted as a partition of the latter ones.

[3] Author to whom any correspondence should be addressed.

IOP Institute of Physics **Φ** DEUTSCHE PHYSIKALISCHE GESELLSCHAFT

Contents

1. Introduction

Financial markets are complex systems where many heterogeneous agents interact. Most financial markets are continuous double auction markets where market members are allowed to trade by submitting limit and market orders. When a limit order is submitted, an order indicating the willingness to buy (or sell) a given volume of a financial asset at a given price is placed inside the order book. The order book is essentially a list of unexecuted orders waiting for a matching order. Market members can also send market orders, which result in immediate transactions at the best available price. Market members channel all orders to be processed in the market. They may act on behalf of a third party or in their own interest. Market members are heterogeneous with respect to many characteristics. Examples are their size (in volume transacted), their trading time horizons (ranging from less than a second to months) and their portfolio of clients (institutional, individuals, companies, foreign investors, etc).

A large body of the financial literature has investigated the statistical regularities observed around trades of different classes of investors, such as institutional investors [1]–[5], individual investors [3], [5]–[8] and foreign investors [4, 5]. In the investigation of the trading profile of institutional investors, the necessity that institutional investors should split large orders into many smaller transactions has been recognized [1, 6], [9]–[15]. In this paper, as was already done in previous studies [11]–[15], we address these large orders as *hidden orders*[4] to emphasize that their true size is typically not public information. The strategic reasons for the incremental execution of hidden orders were originally analyzed by Kyle [16], who developed a model for an insider trader who has information about future prices. Typically, traders of large orders attempt to keep the true size of their orders secret in order to minimize the resulting execution cost. Consider, for example, a trader who wishes to buy a large number of shares of a company because she is expecting a future price increase. She, of course, wishes to buy shares at the lowest price possible. On the other hand, as she demands a large number of shares

---

[4] Other names used in the literature are large trades, packages or meta-orders.

she will push the price up. By executing her order incrementally, she should be able to buy at least part of the order at a price not significantly affected by her action, therefore minimizing her impact on the price dynamics of the stock of interest and lowering her overall execution cost [14, 15]. Some earlier works have provided empirical evidence of statistical regularities of hidden orders [1, 6]. The empirical part of these works usually relies on specialized databases providing information about the coded identity of a set of investors. Information about the class of investors to whom each investor belongs is typically also provided to the researchers. These specialized databases are often confidential and proprietary, and therefore quite difficult to access for research purposes.

More recently, another approach for this kind of investigation has been proposed in the econophysics literature. Here, the trading dynamics of the most active market members of a large financial market have been investigated to statistically detect the strategic behavior in the execution of large orders [13], specialization in the trading strategies [17, 18], herding profile [17] and price impact of large orders broken up into many small trades [15]. These studies have successfully detected heterogeneity of market members with respect to their size, trade horizon and trading profile specialization, both at the Spanish stock market and at the London Stock Exchange. In [13], large trades broken up into many smaller orders have been statistically reconstructed and their statistical properties characterized, whereas in [15], the price impact of large orders characterized by a large fraction of market orders was found to be concave and approximately increasing as the square root of order size.

The statistical reconstruction of hidden orders is not unique. Vaglica *et al* [13] and Moro *et al* [15] adapted an algorithm [19] originally designed to investigate biomedical time series to the segmentation of the inventory profile of market members. This algorithm is quite robust against small local fluctuations, which are distinct from the overall global trend. This choice was motivated by the assumption that, at a given time, a market member could not be considered as acting only for a specific client, and therefore small deviations from the major trend should be expected at any time and taken into account by the algorithm in the process of segmentation. A corollary of this assumption is that the detected hidden order refers to the hidden order of the most important client of the investigated market member.

In the present work, we take a different perspective. Here, we propose a statistical method for the detection of hidden orders that is more sensitive to the local persistence of the trading profile of market members than the method used in [13, 15]. We achieve this goal by implementing a segmentation procedure based on hidden Markov models (HMMs). HMMs [20] were introduced in the late 1960s and have been studied since then. They are very rich and flexible in their mathematical structure and are used in several fields, ranging from speech and image recognition to the detection of biologically relevant nucleotide subsequences. HMMs [21] have also been used in econometrics and finance. Some examples are given in [22]–[26]. To the best of our knowledge, the current paper presents the first application of HMMs to the inventory time series of market members in an equity market.

In this paper, we use HMMs[5] to detect large trading orders that are split into pieces and executed incrementally. This method is used to identify periods of time when the tick-by-tick inventory variation sign is in the same underlying state according to the fitted model. We call the detected periods *patches*. Since market members act simultaneously as brokers for many clients, it may frequently be found that, in a patch, not all the transactions have the same

---

[5] In the appendix, we also consider one of the most used extension of HMMs called hidden semi-Markov models [27].

sign. In the present study, we are mainly interested in directional patches, i.e. those where the trader consistently buys or sells a large number of shares. Our working hypothesis is that each of these patches contains at least one hidden order. The patches detected are analyzed with respect to their characteristics (time duration, number of transactions and volume exchanged), their liquidity profile (relative composition of limit and market orders) and their size compared with the total volume exchanged (participation rate of the patch transactions with respect to all transactions occurring in the market during the same time interval). The analysis is carried out separately for buying and selling directional patches. By keeping these two kinds of orders separate, we are able to detect that the statistical properties of patches detected by the HMM are showing a buy–sell asymmetry that depends on the state of the market. We also compare the results obtained by a HMM with previous results obtained using the more global segmentation algorithm of [13, 15].

This paper is organized as follows. In section 2, we recall the main properties of HMMs. Section 3 discusses the data investigated and the kind of methodology used to detect patches. In section 4, we present the results obtained when extracting the most likely HMM, whereas in section 5, we investigate the statistical properties of detected patches. Section 6 discusses the asymmetry observed in buy and sell patches detected in time periods characterized by a different market state, and section 7 discusses the comparison of the HMM methodology with a methodology based on a different segmentation algorithm, recently used in [13]. Section 8 concludes the paper. In the appendix, we discuss how patches characterized by a length with a power-law distribution are reconstructed by a HMM, and we present a comparison of the results obtained by HMMs and hidden semi-Markov models.

2. Hidden Markov Models (HMMs)

The identification of hidden orders is a typical problem of regime switching with hidden states. In regime switching models, the parameters of the model change abruptly from time to time and the time series is organized in consecutive patches (regimes), each one characterized by a set of model parameters. There is a vast amount of literature on regime switching models [21], and HMMs are among those most used. Even if one considers the time series of the trading activity of a single investor, it is not immediately obvious how to identify the set of consecutive transactions that originated from a unique trading decision and the boundary between different regimes. This problem is even more serious if, as in the analysis presented in this paper, we do not have the inventory variation of an individual investor but rather we have that of a market member or a broker whose trading activity is generated by many investors. In this last case, the identification of different regimes must be performed with statistical methods, such as the HMM.

An HMM is a generalization of a Markov chain in which the states of the chain are not observable. More specifically, consider a system described by a discrete time series and where the system can be in one of N states. When it is in a given state, it emits a signal with a given probability distribution of possible observations. The transition of the system from one state to another is described by a Markov chain. It is often the case that only the observations of the signal are observable, while the underlying states are hidden. Such a system is described by an HMM.

More formally, the elements that completely characterize an HMM are [20]:

(i) N, the number of states in the model, $S = \{S_1, S_2, \ldots, S_N\}$.

(ii) M, the number of distinct observation symbols in each state. The observations correspond to the output of the signal. We denote the individual symbols as $V = \{v_1, v_2, \ldots, v_M\}$.

(iii) The state transition probability matrix $\mathbf{A} = \{a_{i,j}\}$, where

$$a_{ij} = P[q_{t+1} = S_j | q_t = S_i], \quad 1 \leqslant i \leqslant N, \quad 1 \leqslant j \leqslant N, \tag{1}$$

where q_{t+1} and q_t are the state of the system at time $t+1$ and t, respectively.

(iv) The observation symbol probability distribution in state j, $\mathbf{B} = \{b_j(k)\}$, called emission probabilities, where

$$b_j(k) = P[v_k \quad \text{at} \quad t | q_t = S_j], \quad 1 \leqslant j \leqslant N, \quad 1 \leqslant k \leqslant M. \tag{2}$$

(v) The initial state distribution $\pi = \{\pi_i\}$, where

$$\pi_i = P[q_1 = S_i], \quad 1 \leqslant i \leqslant N. \tag{3}$$

The specification of the parameters N and M and the specification of the three probability measures $\mathbf{A}$, $\mathbf{B}$ and π completely define a given HMM (that we address simply as $\lambda(A, B, \pi)$, as in [20]).

Given an empirical time series of observed symbols $O = \{v_{t_1}, v_{t_2}, \ldots\}$, one can ask for the best HMM able to explain it. To this end, one has to compute $P(O|\lambda)$, the probability of observations' time series O given a specific HMM λ. Then one finds the optimal HMM parameters by maximizing this probability. This maximum likelihood estimation can be done either through an iterative procedure (the Baum–Welch reestimation procedure) or by using gradient techniques [20].

After having found the best model λ capable of fitting the data, one can ask for the 'optimal' sequence of hidden states $Q = \{S_{t_1}, S_{t_2}, \ldots\}$ that best explains the observation sequence O given the model λ. While the solution of the first problem (i.e. to find the best model λ) is unique, for the second problem there are several optimality criteria. For example, one possibility is to ask for the sequence of states Q that are individually most likely (this optimality criterion maximizes the expected number of correct individual states), or one can ask for the state sequence that maximizes the expected number of correct pairs of states or triplets, and so on. One can also ask for the single best state sequence. An algorithm exists to solve this last problem. This is the Viterbi algorithm [28, 29]. In the following, we will search for the best sequence of states that are individually most likely.

HMMs are parametric models and therefore assume a given structure of the data. One of the most important structures is the Markovian assumption of the transition between states. As shown in the appendix, in a HMM the distribution of patch length is exponential. Therefore, the HMM poses a serious constraint to the shape of the distribution of patch length. The distribution of real hidden order size is not known, but there are indications that it might be characterized by long tails. Recently, by using a non-parametric segmentation algorithm, it has been suggested that the tail of the distribution of hidden order size is described by a power law [13]. Thus, it seems that fitting the inventory variation with a HMM might pose an unrealistic structure to the patch size distribution. To verify the nature and degree of constraints imposed by our methodology to the detection of patches characterized by a long tailed distribution, we have performed extensive numerical simulations, which are discussed in the appendix. As detailed

there, the patches detected with the HMM are quite close to the ones obtained with a more computationally intensive method, called the hidden semi-Markov model (HSMM), which in numerical simulations and in real data is detecting patch distributions having approximately the same distribution as the one detected with the HMM. Therefore, it seems that the use of the more sophisticated (and probably more appropriate) HSMM does not change significantly the conclusions we draw by using the HMM. Given the intense computational work needed to apply the HSMM to the whole dataset and the requirement of a long time series for the input sequence of this method, in the present study we have chosen to use a HMM.

3. Data and methods

Our database of the electronic market SIBE (Sistema de Interconexión Bursátil Electrónico) allows us to follow each on-book transaction performed by all the market members registered at Spanish Stock Exchange (BME, Bolsas y Mercados Españoles). In 2004, the BME stood eighth in the world in market capitalization. We consider market members' transactions for four highly capitalized stocks, specifically Banco Bilbao Vizcaya Argentaria (BBVA), Repsol (REP), Banco Santander Central Hispano (SAN) and Telefónica (TEF). Our database covers the time period 2001–2006. However, in most of the following analysis we will restrict our attention to the period 2003–2006. The reason for this choice is that, in 2001–2002, we observe a behavior that is markedly different from the one observed in 2003–2006. Since we show the results pooled across different years, we restrict our attention to the years 2003–2006, which are more homogeneous. However, in section 6, we discuss a buy–sell asymmetry providing a justification for the observation that some results change when we investigate the years 2001–2002.

At BME, market members are local and foreign credit entities and investment firms, which are members of the stock exchange and are the only firms entitled to trade. In fact, orders to buy and sell are entered into the market only through members of the stock market. Approximately 75% of them are major financial institutions and approximately 25% are established security dealers. Both kinds of members may trade on their own behalf and/or on behalf of other individuals and institutions that are not members of the market. It is important to stress that market members are not necessarily quoted companies (stocks) but rather are the only institutions entitled to trade stocks directly. In this paper, we consider only the most active market members defined by the criteria that each market member makes at least 1000 transactions per year and is active at least 200 days per year in each of the six investigated years. The number of investigated market members is 43 for BBVA, 37 for REP, 46 for SAN and 52 for TEF. These market members perform approximately 80% of the on-book transactions and are responsible for approximately 78% of the on-book volume (see table 1 for the precise values for each investigated stock).

We fit a HMM to the time series of the sign of the transactions rather than to the signed volume, as in [13, 15]. The reason for this choice is that, in this way, we can fit a HMM without making any parametric assumption about the transaction volume distribution. The sign of a transaction is +1 if the investigated member buys and −1 if she sells. In the fit, performed by using an R software package named *hmm.discnp*, we set the number of hidden states of the HMM as $N = 3$, keeping in mind that they could correspond to a buy state, a sell state and a state without a well-defined trading direction. As we will see below, in most cases the fitted emission probabilities for these states confirm this interpretation. Therefore, we fit each time series of the sign of the transactions of a market member with a HMM with $N = 3$ and $M = 2$. Since

IOP Institute of Physics **Φ** DEUTSCHE PHYSIKALISCHE GESELLSCHAFT

the detected patches are rather short, to detect possible slow time dependencies of the model parameters, we decided to fit the HMM and to apply the state reconstruction algorithm to each year separately. Since we cannot present the results of all the $((43 + 37 + 46 + 52) \times 4 = 712)$ fitted models, in the following we will show the results obtained for the four stocks and four years lumped together. As an example of the single asset behavior, we also show the result for one specific year and stock, which is TEF in 2004.

4. Fitting of HMMs

We use the expected maximization algorithm to perform the maximum likelihood fit of a HMM to the time series of transaction sign for each market member. Specifically, for each market member trading a stock in a particular year, we extract the most likely HMM. Since the observed variable is binary (± 1), the emission probabilities are characterized by assigning, for example, the buy probability ($+1$) for each HMM state. We label the three states as states 1, 2 and 3. In the three states, the average (across market members) buy emission probability is $(0.92 \pm 0.07, 0.50 \pm 0.14, 0.075 \pm 0.05)$ for TEF in 2004, while for the pooled sample it is $(0.95 \pm 0.05, 0.51 \pm 0.17, 0.06 \pm 0.05)$, where errors are standard deviations. These values show that state 1 is characterized by a buy probability very close to one, while state 3 is characterized by a buy probability very close to zero.

Transactions in state 1 are therefore preferentially buy transactions, whereas in state 3 they are preferentially sell transactions. State 2 is characterized by an approximately balanced number of buy and sell transactions. Therefore, we call state 1 *buy state* and state 3 *sell state*. These two states are collectively termed *directional states*. In contrast, the buy probability in state 2 is centered at $1/2$ and we call state 2 *neutral state*, i.e. a state describing a prolonged activity of the market member without a specified direction to buy or sell. It is worth noting that the standard deviation of the buy probability in state 2 is quite large. This shows that, for some market members, state 2 is associated with a buy probability significantly different from $1/2$. In other words, these market members lack a proper neutral state. We do not have an interpretation of this fact. It is finally worth noting that this effect is more evident in 2001 and 2002.

We then consider the transition probability matrices $\mathbf{A}$ of the fitted HMM for each market member trading a stock in one particular year. The average (across market members) transition probability matrix for TEF in 2004 is

$$\mathbf{A}_{\text{TEF04}} = \begin{pmatrix} 0.89 \pm 0.06 & 0.07 \pm 0.05 & 0.04 \pm 0.05 \\ 0.06 \pm 0.04 & 0.87 \pm 0.08 & 0.07 \pm 0.07 \\ 0.04 \pm 0.05 & 0.06 \pm 0.04 & 0.90 \pm 0.07 \end{pmatrix}, \tag{4}$$

while for the whole sample of four stocks and four years it is

$$\mathbf{A}_{\text{pool}} = \begin{pmatrix} 0.89 \pm 0.07 & 0.06 \pm 0.05 & 0.04 \pm 0.06 \\ 0.07 \pm 0.05 & 0.85 \pm 0.07 & 0.08 \pm 0.06 \\ 0.04 \pm 0.06 & 0.07 \pm 0.05 & 0.89 \pm 0.08 \end{pmatrix}, \tag{5}$$

where, again, errors are standard deviations. The analysis of the transition probability matrix shows that the diagonal elements, i.e. the conditional probabilities that the market member remains in the same state, are quite large, indicating the presence of long periods during which the market member persists in the same state. Moreover, the off diagonal elements are on average very similar, indicating that, when the system is in a given state, it has a high probability of remaining in its state, but, if it changes state, it has almost equal probability of jumping

IOP Institute of Physics Φ DEUTSCHE PHYSIKALISCHE GESELLSCHAFT
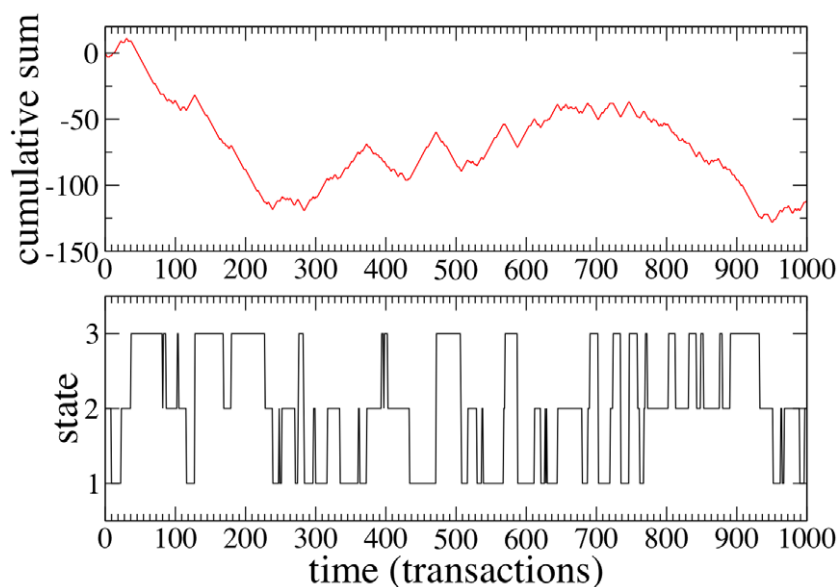
Figure 1. Top panel: cumulative sum of the transaction sign of a sequence of 1000 transactions of the market member Interdin Bolsa, Barcelona (code 9409) when trading TEF. Clear trends or ramps are seen. Bottom panel: corresponding most probable state according to the fitted HMM. States 1, 2 and 3 are the buy, neutral and sell state, respectively.

to either of the other two states. Finally, the fitted transition probability matrices are quite consistent across the four stocks and the four years.

5. Detecting patches and their statistical properties

After the fitting of the HMM to the transaction sign time series, we calculate the most probable hidden state underlying each transaction for all the investigated market members. We call these intervals of transactions, patches of sequential buying or selling, and we interpret them as proxies of the traded hidden orders. To give an example of this reconstruction, in the top panel of figure 1, we show the cumulative sum of a sequence of 1000 transaction signs made by market member Interdin Bolsa, Barcelona (code 9409) when trading TEF. The cumulative sum of the transaction sign shows the presence of ramps corresponding to time periods when the market member was preferentially buying or selling. The identified HMM states in the bottom panel confirm our interpretation of the model states.

5.1. Summary statistics of HMM patches

We summarize some statistical properties of the detected HMM patches in table 1. We first observe that the number of patches is very large, being more than 10^5 per year and per stock. The unconditional mean length of the directional patches is around ten transactions, but the standard deviation is significantly larger, suggesting a non-exponential length distribution. For neutral patches, the mean length is a bit smaller but the standard deviation is again much larger than the mean. The distributional properties of the length of the patches will be investigated

Table 1. Summary statistics of the patches detected in the time series of transaction signs of active market members. Data refer to the four investigated stocks for the period 2003–2006. Line 1 gives the total number of transactions for the considered period. Line 2 gives the number of active market members selected. Lines 3 and 4 give the percentage of on-book transactions and on-book volume in which at least one of the investigated market members is involved, respectively. Lines 5 and 6 give the total number of patches in directional states and in the neutral state, respectively. Lines 7 and 8 are the total number of patches with at least ten transactions in the directional and neutral states, respectively. Lines 9–12 give the mean and standard deviation (sd) of patch length (in transactions) obtained by considering directional and neutral patches of all the selected market members. Lines 13–16 give the mean and standard deviation of patch length (in transactions) obtained by considering directional and neutral patches with at least ten transactions of all the selected market members.

		BBVA	REP	SAN	TEF
1	Total number of transactions	4 413 563	3 289 239	4 914 738	5 940 651
2	Number of active market members	43	37	46	52
3	% of transactions of active market members	77.1	75.1	82.0	83.9
4	% of volume (shares) of active market members	72.5	74.1	79.7	83.8
5	Total number of directional patches	319 057	228 647	301 691	478 175
6	Total number of neutral patches	192 308	137 252	192 236	267 160
7	Total number of directional patches ($N \geqslant 10$)	87 758	66 737	107 073	133 208
8	Total number of neutral patches ($N \geqslant 10$)	42 792	30 656	56 807	72 005
9	Mean length of directional patches	10.51	11.12	12.55	9.73
10	sd of length of directional patches	31.11	22.14	30.78	19.70
11	Mean length of neutral patches	7.41	7.09	9.78	8.62
12	sd of length of neutral patches	13.23	9.57	20.54	14.22
13	Mean length of directional patches ($N \geqslant 10$)	27.42	28.99	26.91	24.01
14	sd of length of directional patches ($N \geqslant 10$)	55.77	34.86	48.37	33.11
15	Mean length of neutral patches ($N \geqslant 10$)	20.97	19.34	24.27	22.20
16	sd of length of neutral patches ($N \geqslant 10$)	23.02	14.00	33.37	21.95

in more detail in section 5.3. Around 30% of the patches are longer than ten transactions, and this subsample will be investigated in depth in the following. This threshold was also used in [13, 15]. The mean length of these long patches is between 20 and 30 transactions. The typical number of patches of a market member in a year is of the order of 1000.

Our results are based on the investigation of the time series of transactions signs of each market member. In this way, we discard any information concerning the volume (in shares or Euros) of the transactions. Since volume is a highly fluctuating quantity [12, 31, 32], it is important to verify that there is consistency between the interpretation of states 1 and 3 as directional states and the sign of the net volume exchanged in these states. Each patch is characterized by the total volume purchased V_{buy} and the total volume sold V_{sell}. We also define $V_{tot} = V_{buy} + V_{sell}$ and we use Euros to measure volume. Analogously, we also define $N_{tot} = N_{buy} + N_{sell}$, where N_{buy} is the number of buy transactions and N_{sell} is the number of sell transactions present in the patch.
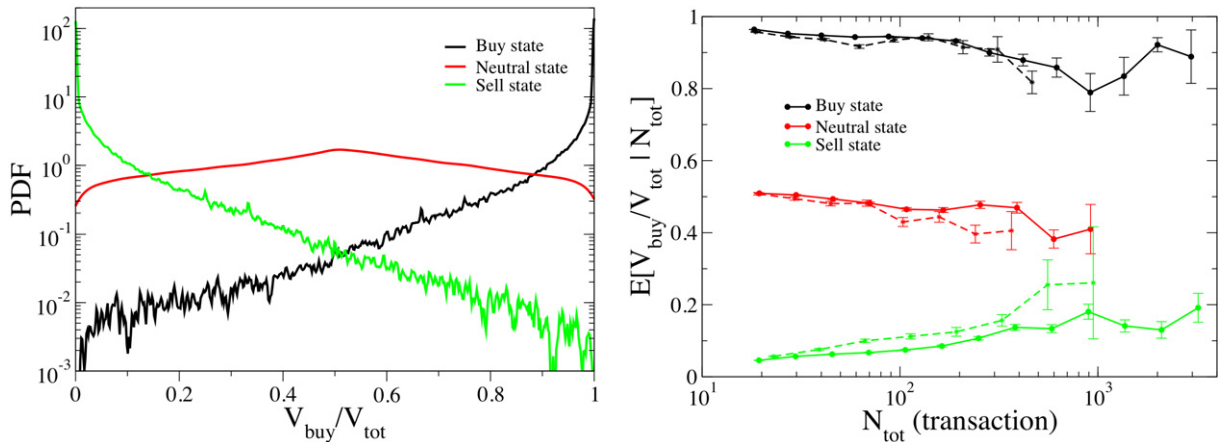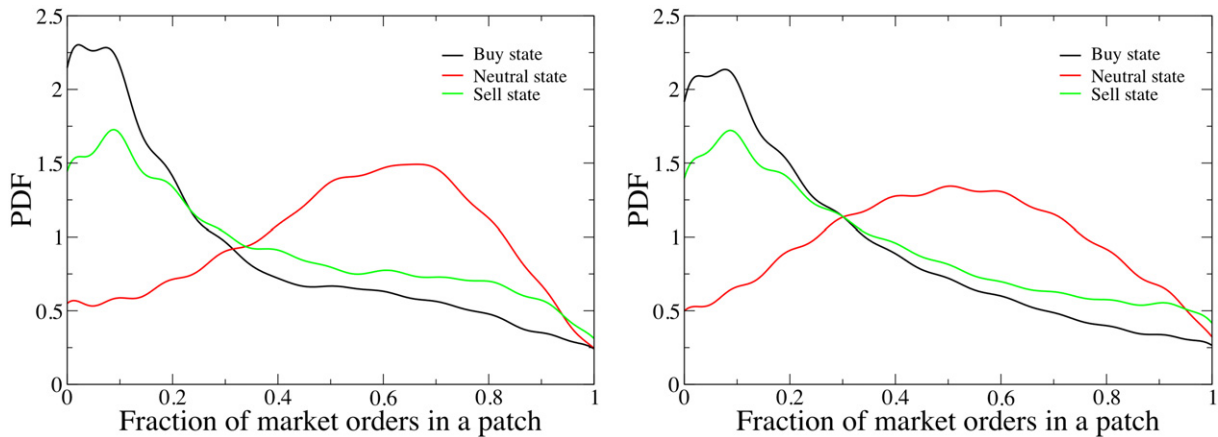
Figure 2. Left panel: PDF of the ratio V_{buy}/V_{tot} for the patches in the buy (black), neutral (red) and sell (green) states. We consider only patches with at least ten transactions for the pooled sample. Right panel: expected value of the ratio V_{buy}/V_{tot} conditional to the patch length N (in transactions) for the patches of TEF in 2004 (dashed lines) and for the pooled sample (solid lines).

For each patch we compute the ratio V_{buy}/V_{tot}, and, in order to avoid discretization errors due to a small number of transactions, we consider only patches with at least ten transactions. The average ratio for these patches is 0.97 ± 0.09, 0.5 ± 0.2 and 0.04 ± 0.09 for patches in state buy, neutral and sell, respectively. In the left panel of figure 2, we show the probability density function (PDF) of the ratio for the patches in the three states. The figure shows that the buy and sell state display a distribution of the ratio V_{buy}/V_{tot} that is very narrowly peaked close to 1 and zero, respectively. On the other hand, this ratio is more broadly distributed for the neutral state. This result confirms the interpretation of states 1 and 3 as directional states, not only in number of transactions but also in volume. The left panel of figure 2 considers all the patches with at least ten transactions. The relation between our interpretation of the states and the ratio V_{buy}/V_{tot} holds also if one measures the ratio conditional on the patch length. The right panel of figure 2 shows the mean ratio V_{buy}/V_{tot} conditional on the patch length measured in total number of transactions. Long directional states tend to be slightly less pure, i.e. roughly 10% of the volume in a long buy patch is associated with sell transactions. However, this fluctuation is relatively small, and we conclude that our interpretation of states 1 and 3 as buy and sell states, respectively, remains valid, even when one considers the volume exchanged.

5.2. Liquidity and participation rate of HMM patches

Our working hypothesis is that HMM patches are related to the strategic behavior of traders. To provide empirical support to our hypothesis, it is important to investigate the properties of the patches concerning the liquidity properties and the participation rate of the patch. Each transaction in an electronic double auction market happens between a limit order and a market order. The use of limit and market orders usually indicates different trading strategies. Market orders are placed by more impatient traders who want their orders to be executed as soon as possible, while traders using limit orders are more patient and are willing to wait for a
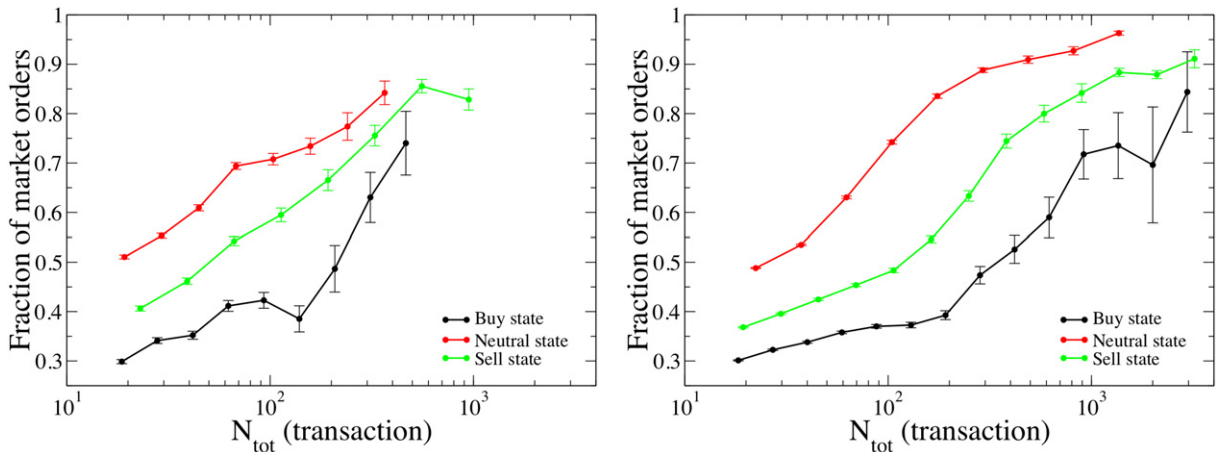
Figure 3. PDF of the fraction of market orders of patches in the buy (black), neutral (red) and sell (green) states for TEF in 2004 (left panel) and for the pooled sample (right panel). We consider only patches with at least 10 transactions.

transaction in order to get a better price. In our database, we do not have direct information about who initiated a transaction, and therefore we use the Lee and Ready algorithm [33] to infer whether the initiator of the transaction, i.e. the agent who placed the market order, was the buyer or the seller.

We investigate the fraction of market orders in the detected HMM patches. Figure 3 shows the probability density function of the fraction of market orders in a patch for the three states detected by the HMM. As before, to avoid discretization problems, we consider only patches made of at least ten transactions. The patches in state 2 show a fraction of market orders that is broadly distributed and the density function is roughly centered on 50%. In contrast, directional patches show a clear preference for a low fraction of market orders. Moreover, the figure shows that there is a difference between patches in buy and sell states. Buy patches tend to have a smaller proportion of market orders than sell patches. Patches in 2001–2002 show the opposite behavior.

The above result considers all the patches independent of their length. To investigate the dependency of the fraction of market orders on the patch length, we show the mean fraction of market orders conditional on the length of the patch in figure 4. The figure shows that, for a given length, neutral patches are richer in market orders than patches in the sell state, which in turn are richer in market orders than patches in the buy state, confirming the unconditional result shown in figure 3. Again, the patches in 2001–2002 display the opposite buy–sell behavior. A possible explanation for the difference between the neutral and directional states is market impact. It has recently been shown that a hidden order made of market orders typically has a positive impact, while an *executed* hidden order made of limit orders typically has a negative[6] market impact [15]. Thus, if a trader wants to execute a buy or a sell hidden order, the use of limit orders will lower the impact when compared to the use of market order (of course, the use of limit orders does not give certainty of execution because the price could move in the opposite direction). Therefore, the low fraction of market orders in directional patches could be the consequence of a strategic choice of minimizing price impact. The asymmetry of behavior

---

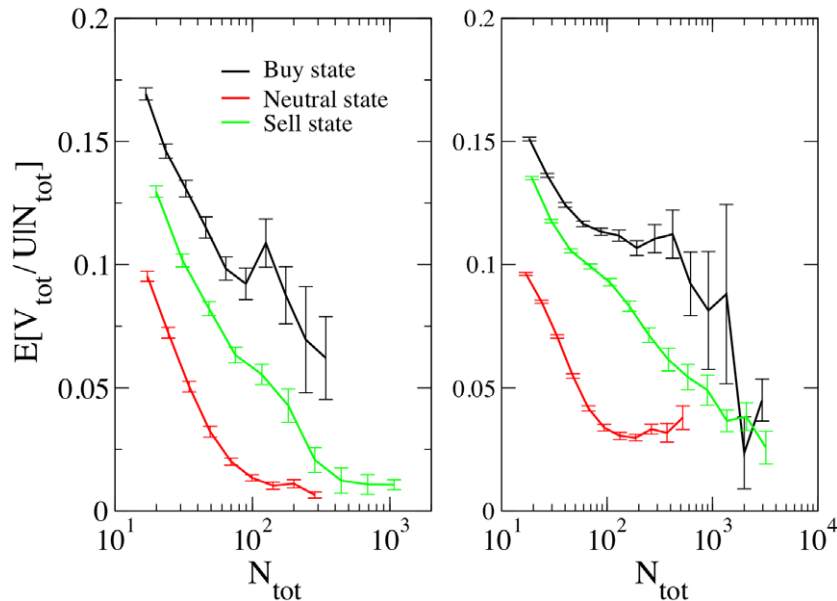[6] See [15] for a discussion of this apparent paradox.

Figure 4. Average value of the fraction of market orders in a patch in the buy (black), neutral (red) and sell (green) states conditional to the patch length for TEF in 2004 (left panel) and for the pooled sample (right panel). Error bars are standard errors.
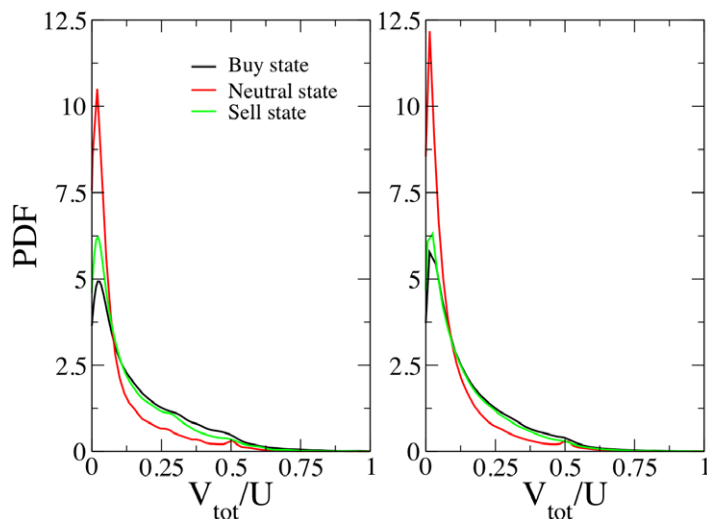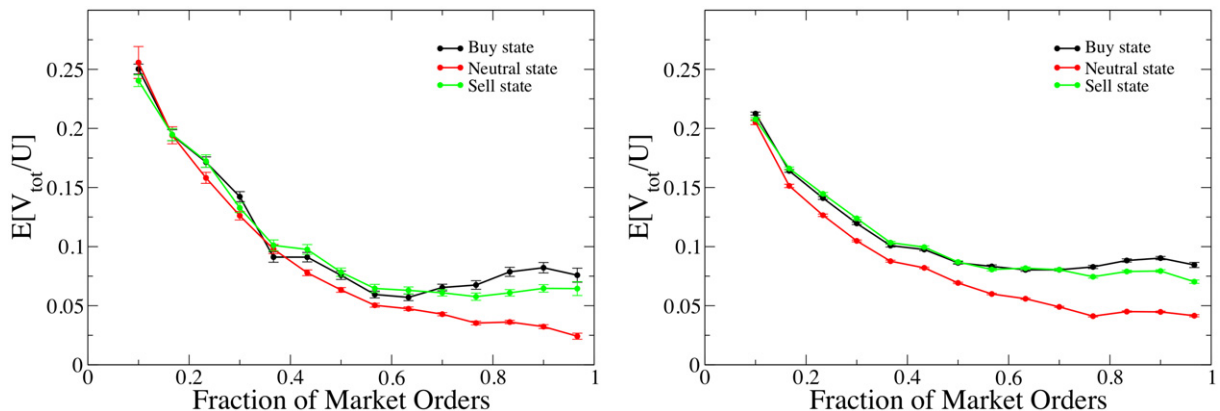

Figure 5. PDF of the participation rate of patches in the buy (black), neutral (red) and sell (green) states for TEF in 2004 (left panel) and the pooled sample (right panel).

between buy and sell patches will be discussed in more detail in section 6. Finally, figure 4 shows that the fraction of market orders in a patch increases with patch length.

The other parameter characterizing the strategy behind hidden order placement is the participation rate. We define the participation rate of a patch as V_{tot}/U, where V_{tot} is the total volume exchanged by the patch measured in Euros and U is the total volume, also measured in Euros, exchanged in the traded stock during the execution of the patch. In figure 5, we show the PDF of the participation rate for the three states. Again, we find a different behavior for the neutral and for the directional states. The neutral patches have a smaller participation rate than the directional patches. This result also holds when conditioning on the order length. Figure 6

Figure 6. Average value of the participation rate conditional to the length of the patch in the buy (black), neutral (red) and sell (green) states for TEF in 2004 (left panel) and the pooled sample (right panel). Error bars are standard errors.

Figure 7. Average value of the participation rate conditional to the fraction of market orders of patches in the buy (black), neutral (red) and sell (green) states for TEF 2004 (left panel) and the pooled sample (right panel). Error bars are standard errors.

shows the mean value of the participation rate conditional on the patch length. We see that, for a given patch length, buy patches have on average a larger participation rate than sell patches, which in turn have a larger participation rate than neutral patches. As before, in 2001–2002, the buy–sell asymmetry of behavior of directional patches is inverted. The participation rate decreases with the length of the patch, and this is expected, given that a large order is typically executed with a smaller participation rate in order to avoid detection by other traders and in order to decrease the impact cost. In figure 7, we consider the mean participation rate conditional on the fraction of market orders in a patch. There we find that this is a decreasing function, in

contrast with what is found in [15], where patches detected with the more global segmentation algorithm were investigated.

In summary, we conclude that large directional patches are made mainly of market orders and have a small participation rate, while small directional patches have a larger fraction of limit orders and a larger participation rate. The buy–sell asymmetry conforms to this pattern. In fact, in the period 2003–2006, buy patches have a smaller fraction of market orders and a larger participation rate than sell patches. In section 6, we show that this asymmetry may be associated with the trend of the price of the traded stock.

5.3. Statistical properties of HMM patch length size

Three variables characterize the size of patches: T, N_{tot} and V_{tot}. The time T measured in seconds is the time elapsed from the first to the last transaction of the patch. We use trading time, i.e. we remove night-time, holidays and in general any time when the market is closed. The total number of transactions composing the patch is given by $N_{tot} = N_{buy} + N_{sell}$, and the total volume traded in the patch is $V_{tot} = V_{buy} + V_{sell}$. These variables are different from the one used in [13], where only the most frequent type of transactions (buy or sell) in a directional patch were considered in the definition of N and V. However, as we have seen above, in HMM directional patches the most frequent type of transaction is responsible for almost all the volume exchanged in the patch, and, in fact, we observe very little difference in the distributions obtained by considering all the transactions or only the ones of most frequent type. The motivation for the slight change in variables is that our present choice allows us to investigate the distributional properties both of the directional and of the neutral patches with the same set of variables.

Figure 8 shows the cumulative distribution of T, N_{tot} and V_{tot} for the four considered stocks pooling together the four different years. We pool together the directional patches because, in this case, we find a negligible difference of behavior between buy and sell patches. This is true also when we investigate the 2001–2002 period. We observe that, for neutral patches, the distributions of N_{tot} and V_{tot} have a thinner tail than the corresponding directional patch distributions. For T, the distributions for the three states are very similar. The cumulative distribution shows that approximately 90% of the detected patches are characterized by a time duration of less than an hour. Therefore, the largest fraction of detected patches seems to be proxies of intraday hidden orders. This might indicate a preference of the traders to close daily the position taken, or might be related to a specific aspect of our approach, which is monitoring market members and not single investors. In fact, it is realistic to hypothesize that market members are most of the time performing multiple, although heterogeneous, strategies, and the interference of multiple strategies might end up in a local fragmentation of trading strategies planned over longer time horizons by a limited number of large investors. Fragmentation might be revealed by the HMM approach.

However, in spite of this limitation, from figure 8 we can see that, also in the case of a HMM, the patch size has a fat tail distribution. We performed a Jarque–Bera test of the hypothesis that the above distributions are lognormal and in all cases we reject the null hypothesis at a 0.01 confidence level. It is difficult to assess if the tail of the distributions is power law, given the small extension of the tail. In order to compare the tail of the distributions with those obtained in [13], we estimate the tail exponent by using the Hill estimation. Table 2 gives the Hill estimation for the tail exponent for all the analyzed sets and variables when the estimation is applied to the top 5% percentile.
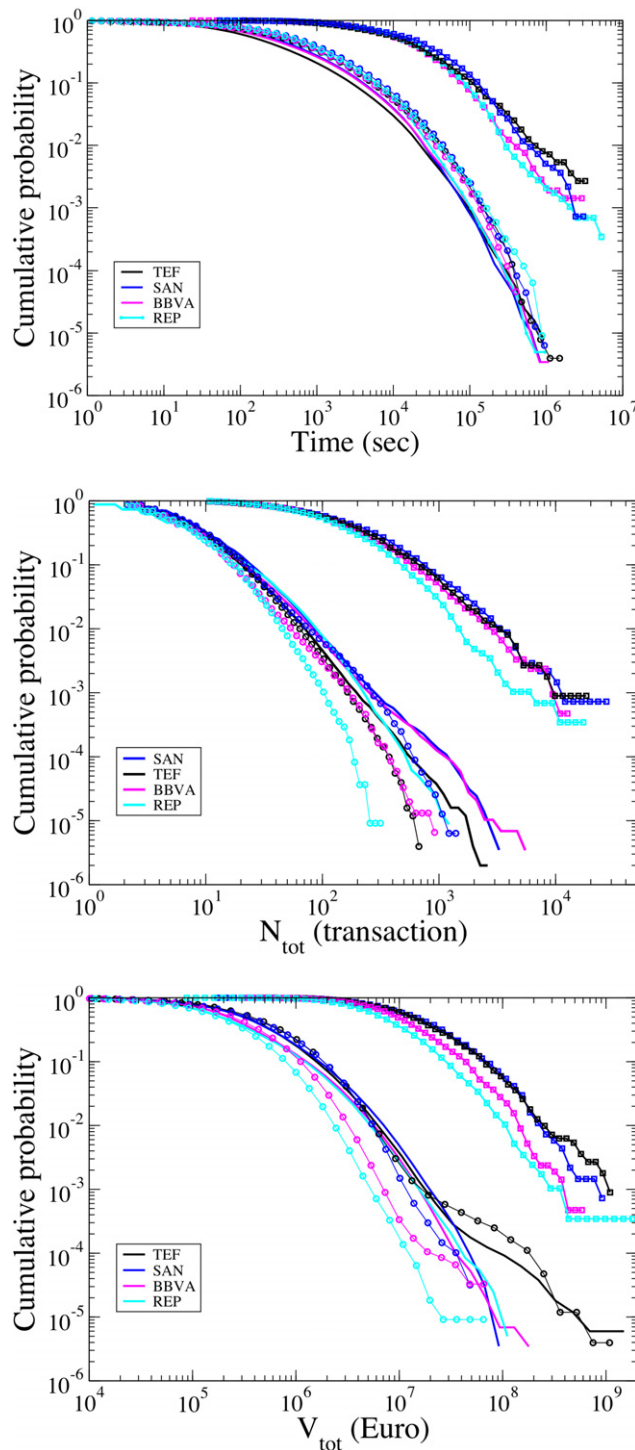
Figure 8. Top panel: distribution of time interval T elapsed between the first and last transactions in an HMM patch. Middle panel: distribution of the total number of transactions N_{tot} in an HMM patch. Bottom panel: distribution of the total volume V_{tot} exchanged in an HMM patch. The solid lines refer to the directional patches and the circles to the neutral patches. The squares show the distribution of the T, N_{tot} and V_{tot} variables for patches detected with the algorithm of [13]. The investigated years are 2003–2006.

Table 2. Tail exponent of the cumulative distributions of T, N_{tot} and V_{tot} of HMM patches. The exponent is estimated with the Hill estimation applied to the top 5% quantile. The top part of the table refers to directional patches (buy and sell state), and the bottom part refers to the neutral state. The errors identify a 95% confidence interval.

	BBVA	REP	SAN	TEF
ζ_T	1.33 ± 0.02	1.50 ± 0.03	1.40 ± 0.02	1.22 ± 0.01
$\zeta_{N_{tot}}$	1.95 ± 0.03	1.93 ± 0.03	1.84 ± 0.03	1.92 ± 0.02
$\zeta_{V_{tot}}$	1.73 ± 0.03	1.72 ± 0.03	1.75 ± 0.02	1.82 ± 0.02
ζ_T	1.31 ± 0.03	1.34 ± 0.03	1.36 ± 0.03	1.26 ± 0.02
$\zeta_{N_{tot}}$	2.07 ± 0.04	2.48 ± 0.06	1.87 ± 0.03	2.19 ± 0.03
$\zeta_{V_{tot}}$	2.36 ± 0.04	2.31 ± 0.05	2.38 ± 0.05	2.15 ± 0.03

Table 3. Tail exponent of the cumulative distributions of T, N_{tot} and V_{tot} of patches. The exponent is estimated by the Hill estimator applied to the top 5% quantile of the directional patches detected in the signed volume time series by using the segmentation algorithm of [13]. The errors identify a 95% confidence interval.

	BBVA	REP	SAN	TEF
ζ_T	1.53 ± 0.25	1.73 ± 0.24	1.54 ± 0.32	1.23 ± 0.28
$\zeta_{N_{tot}}$	1.64 ± 0.27	2.10 ± 0.30	1.57 ± 0.32	1.58 ± 0.36
$\zeta_{V_{tot}}$	1.91 ± 0.31	1.83 ± 0.26	2.20 ± 0.45	1.79 ± 0.41

We also compare the size distribution of the HMM patches with those found by applying the segmentation method used in [13] to the signed volume time series. We therefore run the segmentation algorithm to the inventory time series and we consider as size variables T, N_{tot} and V_{tot}. The figure shows that these patches are much longer than the one identified by the HMM. Table 3 shows the Hill estimation of the tail exponent of the directional patches identified by the segmentation algorithm on the top 5% quantile. By comparing the exponents in this table and in table 2, we observe that the values are not very different, suggesting a relation between the small HMM patches and the larger patches detected by the segmentation algorithm. In section 7, we will make this comparison more quantitative.

6. The buy–sell asymmetry

We have seen in the previous sections that several properties of HMM patches show a significant behavior asymmetry between buy and sell patches. This may be due to two possible reasons. The first possible explanation is that there is an intrinsic asymmetry between the buy and sell strategies of hidden order splitting. This may be due to behavioral biases, margin calls, short selling restriction, etc. The second possible explanation is related to the presence of a different behavior in the buy–sell order splitting strategy that is *conditional* on the market state: for example, the fact that the price is trending up or down. In this second case, there is a buy–sell

**IOP** Institute of Physics Φ DEUTSCHE PHYSIKALISCHE GESELLSCHAFT

asymmetry when the market is, for example, trending up, but the asymmetry is reverted when the market trends down.

In order to investigate which hypothesis explains the data better, we consider Telefonica stock and we divide our six-year period (2001–2006) into 72 one-month sub-periods. For each of them we compute the ratio $x = \langle r \rangle / \sigma$ between the average daily price return in the month and the volatility σ, here computed as the standard deviation of daily returns in the month. This ratio gives a measure of the price trend in the month. Then we consider the HMM directional patches detected in the considered month with at least ten transactions, and we compare the statistical properties of the buy and sell patches as a function of x. The top part of figure 9 shows the results. Specifically, panel (a) shows the difference in the number of buy and sell patches as a function of x. Panel (b) shows the difference in the average length measured in transactions between buy and sell patches as a function of x. Panel (c) shows the difference in the average fraction of market orders between buy and sell patches as a function of x. Panel (d) shows the difference in the average participation rate between buy and sell patches as a function of x. In each panel, we observe a highly significant linear relation between the plotted variables. The correlation coefficient is -0.42, -0.42, -0.61 and 0.64, respectively. This result indicates that, when the price increases, there are more sell patches, which are typically longer, richer of market order and with a lower participation rate when compared with buy patches. The opposite happens when the price decreases. The intercept of the regressions is significantly different from zero for the number and the fraction of market orders, suggesting also an overall buy–sell asymmetry, even in the absence of a price trend ($x = 0$). However, the dependence of the four quantities on x is much stronger, indicating that the second hypothesis stated above is probably the most relevant one in explaining the buy–sell asymmetry. Finally, for the neutral patches, we do not observe any significant dependence of the statistical properties as a function of x. A full explanation of these asymmetries is beyond the scope of this paper but will be considered in a future paper.

These asymmetries help us to explain why we investigate the market in 2003–2006 and leave out the period 2001–2002. The bottom part of figure 9 shows the closing price of the four stocks in the period 2001–2006. We see that the market was in a bear phase (i.e. decreasing price) in 2001–2002, while it was in a bullish phase (i.e. increasing price) in 2003–2006. By investigating each year separately, we find that some properties of the HMM patches are similar in all years, while others have a different behavior in 2001–2002 and in 2003–2006. The main differences we find in 2001–2002 are as follows.

- Sell patches typically have a smaller fraction of market orders than buy patches, i.e. contrary to what was observed in 2003–2006, where sell patches have a higher fraction of market orders than sell patches (figures 3 and 4).

- Sell patches have a higher participation rate than buy patches, i.e. contrary to what was observed in 2003–2006, where sell patches have a smaller participation rate than buy patches (figures 5 and 6).

These differences are consistent with the dependence of the buy–sell asymmetry on the price trend investigated above on the monthly scale. We observe that this explanation seems to work even when one considers longer timescales, such as the yearly timescale. While this explains our choice of investigating only the four-year period 2003–2006, a full understanding of the interdependence between long scale price trends and short scale HMM buy–sell asymmetries requires further investigation.
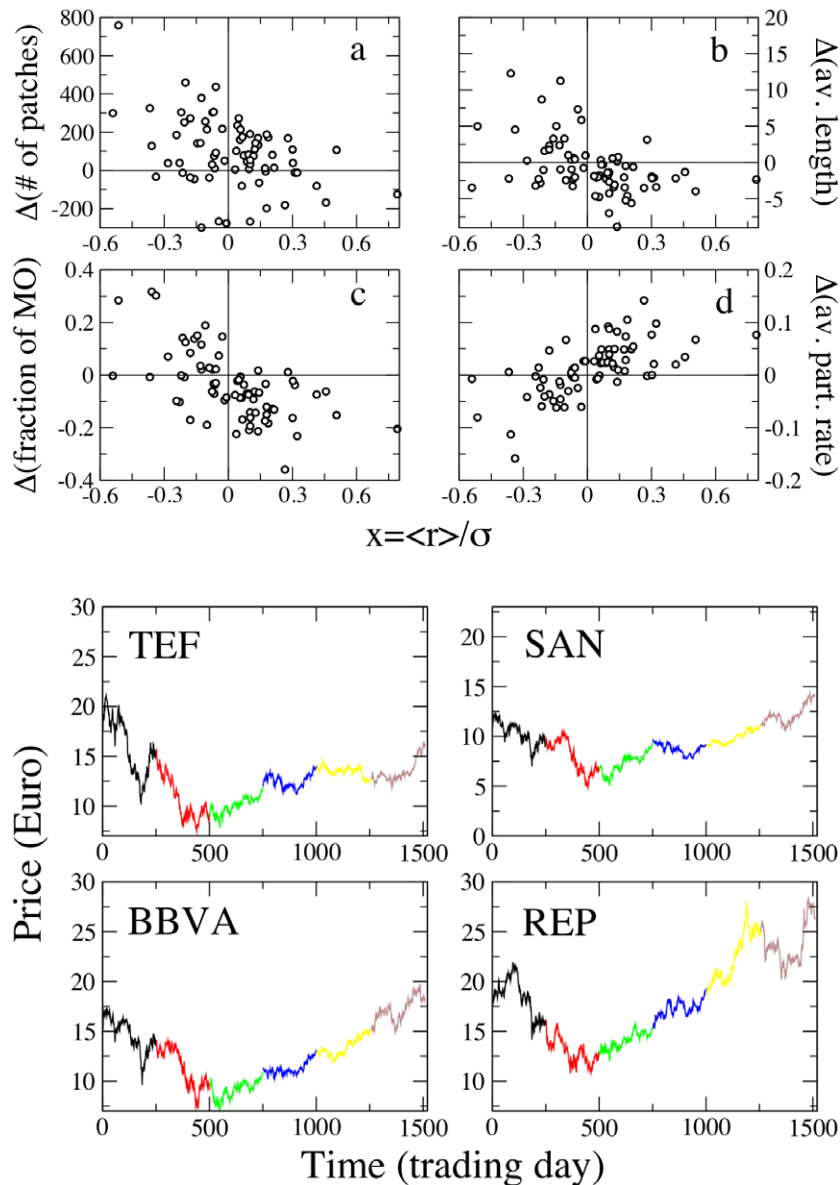
Figure 9. Top four panels: buy–sell asymmetry measures as a function of the ratio $x = \langle r \rangle / \sigma$ for directional patches with at least ten transactions. The symbol $\Delta(.)$ indicates the difference in the quantity in parentheses between buy and sell patches. Specifically, (a) shows the difference in the number of buy and sell patches, (b) shows the difference in the average length (in transactions) between buy and sell patches, (c) shows the difference in the average fraction of market orders between buy and sell patches and finally (d) shows the difference in the average participation rate between buy and sell patches. Bottom four panels: closing price time series for the four investigated stocks in the period 2001–2006. Each color refers to a different calendar year.
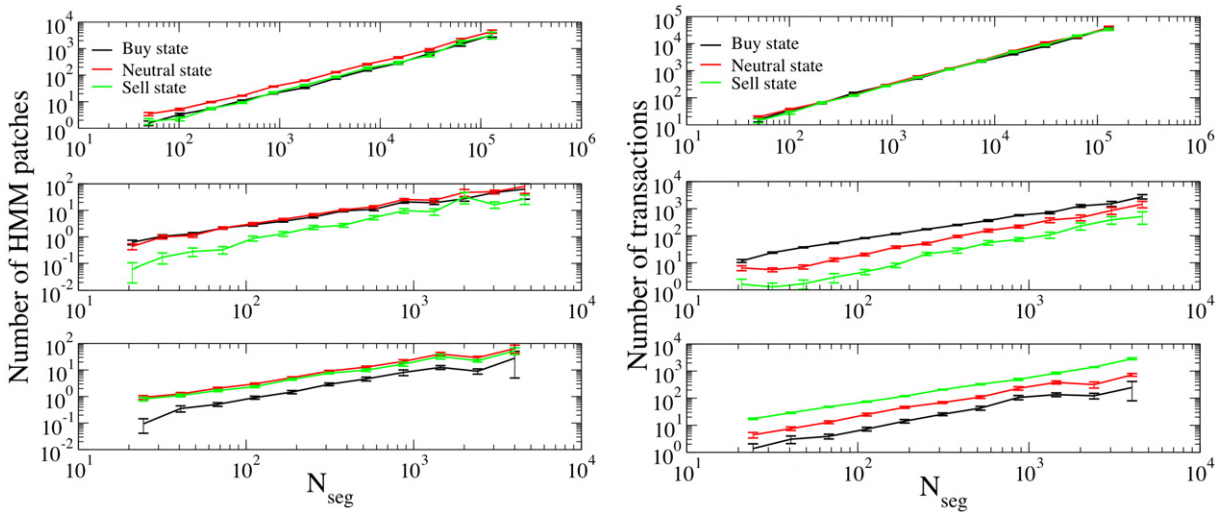
Figure 10. Left panel: average number of HMM patches in the three HMM states found inside a segment patch of type neutral (top), buy (middle) and sell (bottom) as a function of N_{seg}, which is the length of the segment patch. Right panel: average number of transactions in HMM patches in the three HMM states found inside a segment patch of type neutral (top), buy (middle) and sell (bottom) as a function of N_{seg}.

7. Comparison of segmentation algorithms

In this section, we compare the patches detected with the HMMs and the ones detected with the algorithm of [13] in more detail. As we have described in section 6, the two types of patches describe regimes at very different scales. HMM patches are typically short, while the patches of [13] (hereafter referred to as *segment patches*) are very long. In order to compare the two segmentations, we divide the segment patches into three groups of type buy, sell and neutral, according to the criterion used in [13]. We therefore compute the number of HMM patches in the three HMM states contained in a segment patch of a given type. The left panel of figure 10 shows the average number of HMM patches in the three HMM states for segment patches of the three types as a function of N_{seg}, which is the number of transactions present in the segment patch. Neutral segment patches contain roughly the same number of neutral HMM patches as the sum of the numbers of buy and sell HMM patches. The middle and bottom panels of the left part of figure 10 show that a buy (sell) segment patch contains typically an equal number of buy (sell) and neutral HMM patches and a much smaller number of sell (buy) HMM patches. Therefore, directional segment patches are typically a mixture of neutral and directional (with the same sign) HMM patches.

The comparison is even clearer if one considers the number of transactions in a segment patch, which are assigned by the HMM to a specific type of state (see the right panel of figure 10). Neutral segment patches have, on average, an equal number of transactions in the three HMM states (see the top right panel). Taken together with the top left panel, this result indicates that neutral segment patches are composed of relatively short neutral HMM patches and an equal mixture of longer buy and sell HMM patches. By contrast, buy (sell) segment patches are composed of a large fraction of transactions in the buy (sell) HMM state, a smaller

IOP Institute of Physics ⬥ DEUTSCHE PHYSIKALISCHE GESELLSCHAFT

fraction of transactions in the HMM neutral state and a much smaller fraction of transactions in the opposite directional HMM state. We therefore conclude that directional segment patches are mainly composed of an equal mixture of short neutral HMM patches and long HMM directional patches with the same direction as the segment patches.

8. Conclusions

In this paper, we have shown that HMMs can be used to detect patches in the series of transaction signs performed by market members. We interpret the directional patches as possible hidden orders submitted by the market members. By analyzing the statistical properties of these patches, we observe that the time distribution of these patches is fat tailed. Unconditionally, long patches are characterized by a large fraction of market orders and a low participation rate, while short patches have a large fraction of limit orders and a high participation rate. We also observe an interesting buy–sell asymmetry, which depends on the trend of the stock price. When the price goes up, we detect more sell patches, which are longer, with more market orders and less aggressive compared with the buy patches detected in the same period. When the price decreases, the opposite behavior is observed. These results may be considered a contribution to the identification of stylized facts in the strategic behavior of the investors.

The comparison of the HMM patches and those obtained with the segmentation algorithm used in [13] can shed light on their interpretation. The two segmentation methods perform a coarse graining of the inventory time series on two different time scales. The HMM is able to detect short patches with a well-defined drift, while the segmentation algorithm detects large-scale patches composed of many smaller HMM directional patches with the same sign, interspersed with even smaller HMM neutral patches. A possible interpretation is that the segment patches represent the trading decision of the portfolio managers that are executed on time scales that can go up to several days, while HMM patches represent the trader's executions of usually intraday smaller packages, which might be part of the larger portfolio manager orders.

There are many interesting points raised by the present study that might deserve future investigations. Firstly, it is important to build a stronger relation between the stylized facts observed in HMM patches and the possible strategies adopted by the investors. Secondly, the buy–sell asymmetry deserves a more in-depth analysis to understand (i) how it depends on the considered time scale and (ii) a possible causal relation between the asymmetry and the price trend. Specifically, is the price going up or down because of the asymmetry or is the order flow adapting to the price drift? Thirdly, it is interesting, both from a theoretical and from an applied point of view, to compute the market impact of the HMM patches, in a similar way to what has been done in [15] for patches detected with the segmentation algorithm. Finally, in our study, we have pooled together all the market members without considering possible differences in the order splitting strategies across market participants. An understanding of the heterogeneity and mutual relationship in the order splitting strategies of market participants might result in a better appreciation of the fascinating process of price formation.

Appendix. HMMs, HSMMs and power-law distributions

In this paper, we have used HMMs to model hidden orders. However, HMMs are parametric models and therefore assume a given structure of the data. A known property of HMMs is that patch length is exponentially distributed. In fact, if a system is described by an HMM, which is initially in a given state S_i, the probability $p_i(\ell)$ that the system stays in that state for exactly ℓ steps, i.e. that

$$O = \{q_1 = S_i, q_2 = S_i, q_3 = S_i, q_4 = S_i, \ldots, q_\ell = S_i, q_{\ell+1} \neq S_i\}, \tag{A.1}$$

is

$$p_i(\ell) = (a_{ii})^{\ell-1}(1 - a_{ii}). \tag{A.2}$$

Since we do not have any prior information about the length distribution, we test numerically whether HMMs are able to detect non-exponentially distributed patch lengths. More specifically, we perform two in-depth simulation studies. In the first, we simulate an artificial time series with patch length taken from a Pareto (i.e. power law) distribution. Then we fit an HMM and we reconstruct the hidden states of the process by looking for the best sequence of states that are individually most likely. The second simulation study involves considering a generalization of HMMs, namely the HSMM in which the patch length distribution is not fixed by the model but is fitted from the data. Both studies show that our procedure is able to detect non-exponential distributed patch lengths.

### A.1. Reconstruction of power-law-distributed patches

Here, we show the results of a simulation study of the reconstruction with a HMM of the hidden states of an artificial time series generated with power-law-distributed patch length. Specifically, we simulate a time series composed of patches of variable length and characterized by different compositions of the symbols +1 and −1. The length of each patch is extracted from a Pareto distribution $P(\ell) \propto \ell^{-\mu}$. We assign to each patch a dominant sign (+1 or −1), and then for each patch we simulate a time series of length ℓ as a binary time series with a bias in favor of the dominant sign. Consecutive patches have alternate dominant signs. In the following, we show the results obtained for a surrogate time series composed of 5000 patches with length extracted from a Pareto distribution with exponent $\mu = 2$. In figure A.1, we compare the distribution of the simulated patch length with the distribution of the HMM reconstructed patch length. Figure A.1 shows that the tail behavior of the two distributions is quite similar, apart from a global factor, which is due to the fact that the HMM is unable to detect very short patches due to statistical fluctuations. This result suggests that HMMs are able to reproduce the power-law behavior of the distributional properties of the patches in a time series, despite the exponential structure of the patch length expected in a HMM.

A.2. Comparison of HMMs and HSMMs

The problem of the exponential distribution of the patch length is well known in the HMM literature. The most used extension of HMMs to cope with the problem of exponential distribution of patch length is the HSMM [27]. In a HSMM, the diagonal elements of the transition probability are set to zero and the model is specified by assigning also the patch length
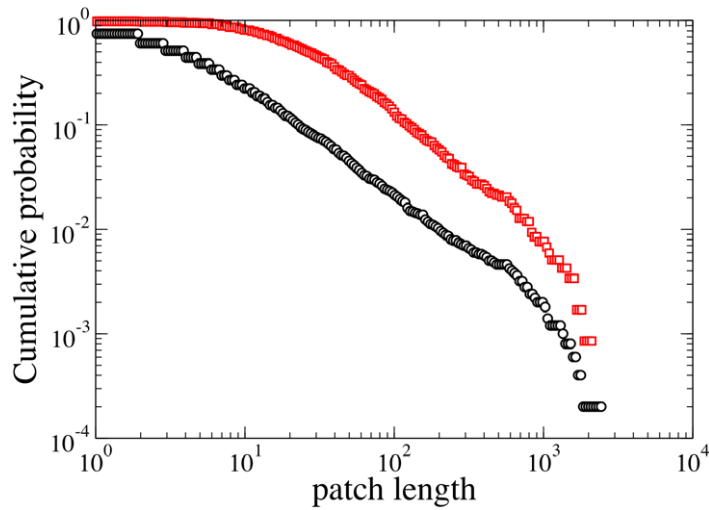
Figure A.1. Black circles: cumulative probability of the patch length simulated from a Pareto distribution of tail exponent $\mu = 2$. Red squares: cumulative probability of the patch length detected by the HMM.
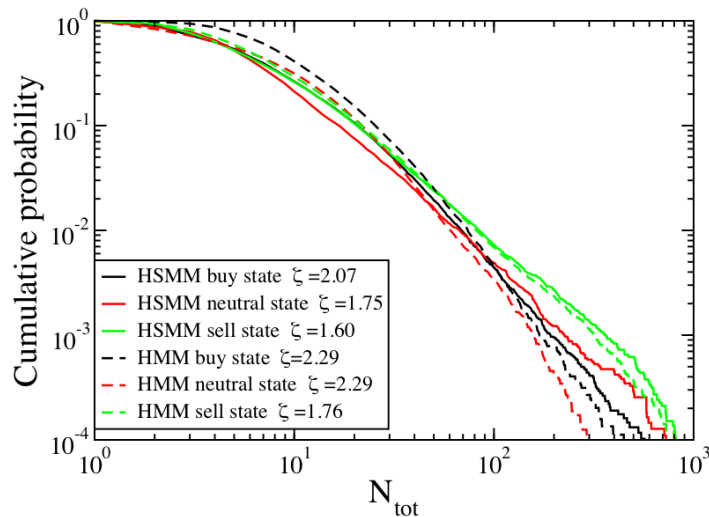
Figure A.2. Cumulative probability of the patch length for the three states obtained by the HSMM (solid line) and HMM (dashed lines). In the legend, there are also shown the respective values of tail exponent ζ performed by Hill estimation on the top 5% quantile.

distribution (sometimes called sojourn time distribution or runlength distribution) associated with each state and defined as

$$d_j(\ell) = P(S_{t+\ell+1} \neq j, S_{t+\ell} = j, \ldots, S_{t+2} = j | S_{t+1} = j, S_t \neq j) 1 \leqslant j \leqslant N. \tag{A.3}$$

In other words, in a HSMM, the patch length distribution of each hidden state is assigned explicitly in the specification of the model and is not therefore necessarily exponential. There exist algorithms that allow us to fit a HSMM with a given sojourn time distribution to a time

series. However, if one has no prior information about the functional form of the sojourn time distribution, it is possible to fit a non-parametric HSMM to a time series. In this case, the distributions of sojourn time are fitted from the data. HSMMs with non-parametric state sojourn time distributions were first proposed by Ferguson in 1980 [27] in the field of speech recognition, and an efficient algorithm for the model fit has been recently proposed [30]. In the HSMM investigation, we use a non-parametric fitting, i.e. we fit also the distribution of sojourn time. This estimation requires a long computational time and can be only applied to a long time series.

For the estimation of the model parameters of the HSMM and the subsequent reconstruction of the underlying state sequence, we use the *R* library *hsmm* [34]. The algorithms for maximum likelihood parameter estimation in this library are based on a method introduced in [30]. The non-parametric estimation of the sojourn time distribution requires a large number of observations. For this reason, we choose to analyze only a time series with at least 20 000 records. Accordingly, we select 28 market members trading TEF in the 2006. The mean buy probability computed across market members in the three hidden states of the fitted HSMM is 0.05 ± 0.05, 0.53 ± 0.16 and 0.97 ± 0.05, which supports the interpretation of the three states as sell state, neutral state and buy state, respectively. The mean transition probability matrix is

$$
\begin{pmatrix}
0 & 0.3 \pm 0.2 & 0.7 \pm 0.2 \\
0.4 \pm 0.2 & 0 & 0.6 \pm 0.2 \\
0.6 \pm 0.1 & 0.4 \pm 0.1 & 0
\end{pmatrix}. \tag{A.4}
$$

After the model parameter estimation, we compute the hidden state sequence for each of the 28 inventory variation sign time series. In figure A.2, we compare the cumulative probability of patch length obtained by the HMM and HSMM. For the directional states, the distribution of patch length obtained by the two methods is quite similar. The most pronounced difference is observed in the tail of the distribution of the length of neutral patches. We also estimate the tail exponents with the Hill estimation on the 5% quantile, and we notice similar tail exponent values for HMM and HSMM patches.

The parameter estimations for the HSMM require much more computational resources than the HMM. The fitting and reconstruction of the time series for one year of transactions of 28 participants with the HSMM required about 6 days of computing time compared with a few minutes for the HMM. The complexity of the algorithm is linear in the sequence length, and we estimated that the modeling of 6 years of transactions for the participants trading Telefonica stock would have required roughly one month of computations. Even though the HSMM is a more realistic choice than the HMM in modeling inventory time series, the small differences in the distributions observed in figure A.2 convinced us to use the HMM to model our data.

References

[1] Chan L K C and Lakonishok J 1995 *J. Finance* **50** 1147
[2] Grinblatt M, Titman S and Wermers R 1995 Momentum investment strategies, portfolio performance and herding: a study of mutual fund behavior *Am. Econ. Rev.* **85**(5) 1088–105
[3] Nofsinger J R and Sias R W 1999 *J. Finance* **54** 2263
[4] Choe H, Kho B-C and Stulz R M 1999 *J. Finance Econ.* **54** 227
[5] Grinblatt M and Keloharju M 2000 *J. Finance Econ.* **55** 43
[6] Gallagher D R and Looi A 2006 *Account. Finance* **46** 125

[7] Barber B M, Lee Y, Liu Y and Odean T 2008 *Rev. Fin. Stud.* **22** 609

[8] Grinblatt M and Keloharju M 2009 *J. Finance* **64** 549

[9] Torre N 1997 *BARRA Market Impact Model Handbook* (Berkeley, CA: BARRA)

[10] Almgren R, Thum C, Hauptmann H L and Li H 2005 *Risk* July 2005 p 57

[11] Lillo F and Farmer J D 2004 The long memory of the efficient market *Studies in Nonlinear Dynamics and Econometrics* **8**(3) http://www.bepress.com/snde/vol8/iss3/art1

[12] Lillo F, Mike S and Farmer J D 2005 *Phys. Rev.* E **7106** 066122

[13] Vaglica G, Lillo F, Moro E and Mantegna R N 2008 *Phys. Rev.* E **77** 036110

[14] Bouchaud J P, Farmer J D and Lillo F 2009 *Handbook of Financial Markets: Dynamics and Evolution* vol 57 ed T Hens and K R Schenk-Hoppé (Amasterdam: North-Holland)

[15] Moro E, Vicente J, Moyano L G, Gerig A, Farmer J D, Vaglica G, Lillo F and Mantegna R N 2009 *Phys. Rev.* E **80** 066102

[16] Kyle A S 1985 *Econometrica* **53** 1315

[17] Lillo F, Moro E, Vaglica G and Mantegna R N 2008 *New J. Phys.* **10** 043019

[18] Zovko I I and Farmer J D 2007 *Complexity, Metastability and Nonextensivity: An International Conference (AIP Conf. Proc.)* (Berlin: Springer) arXiv.org:0709.3261v1

[19] Bernaola-Galvan P, Ivanov P Ch, Amaral L A N and Stanley H E 2001 *Phys. Rev. Lett.* **87** 168105

[20] Rabiner L R 1989 *Proc. IEEE* **77** 257

[21] Hamilton J D 2008 Regime-switching models *Palgrave Dictionary of Economics* 2nd edn ed S Durlauf and L Blume (New York: Palgrave Mcmillan)

[22] Goldfeld S M and Quandt R E 1973 *J. Econometrics* **1** 3

[23] Tyssedal J S and Tjøstheim D 1988 *Appl. Stat.* **37** 353

[24] Hamilton J D 1989 *Econometrica* **57** 357

[25] Rydén T, Teräsvirta T and Åsbrink S 1998 *J. Appl. Econ.* **13** 217

[26] Diebeld F X and Inoue A 2001 *J. Econometrics* **105** 131

[27] Ferguson J D 1980 *Proc. Symp. on the Applications of Hidden Markov Models to Text and Speech* ed J D Ferguson (Princeton, NJ) p 143

[28] Viterbi A J 1967 *IEEE Trans. Inf. Theory* **13** 260

[29] Forney G D 1973 *Proc. IEEE* **61** 268

[30] Guedon Y 2003 *J. Comput. Graph. Stat.* **12** 604

[31] Gopikrishnan P, Plerou V, Gabaix X and Stanley H E 2000 *Phys. Rev.* E **62** R4493

[32] Eisler Z and Kertesz J 2006 *Eur. Phys. J.* B **51** 145

[33] Lee C M C and Ready M J 1991 *J. Finance* **46** 733

[34] Bulla J, Bulla I and Nenadic O 2010 *Comput. Stat. Data Anal.* **54** 611