

Segmentation algorithm for non-stationary compound Poisson processes

With an application to inventory time series of market members in a financial market

Bence Tóth^{1a}, Fabrizio Lillo^{1,2}, and J. Doyne Farmer^{1,3}

¹ Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

² Dipartimento di Fisica e Tecnologie Relative, Università di Palermo, Palermo, Italy

³ LUISS Guido Carli, Viale Pola 12, 00198, Roma, Italy

Received: date / Revised version: date

Abstract. We introduce an algorithm for the segmentation of a class of regime switching processes. The segmentation algorithm is a non parametric statistical method able to identify the regimes (patches) of the time series. The process is composed of consecutive patches of variable length, each patch being described by a stationary compound Poisson process, i.e. a Poisson process where each count is associated to a fluctuating signal. The parameters of the process are different in each patch and therefore the time series is non stationary. Our method is a generalization of the algorithm introduced by Bernaola-Galván, *et al.*, Phys. Rev. Lett., **87**, 168105 (2001). We show that the new algorithm outperforms the original one for regime switching compound Poisson processes. As an application we use the algorithm to segment the time series of the inventory of market members of the London Stock Exchange and we observe that our method finds almost three times more patches than the original one.

PACS. 02.50.Ey Stochastic processes – 05.45.Tp Time series analysis – 89.65.Gh Economics; econophysics, financial markets, business and management

1 Introduction

Many time series from natural and social phenomena exhibit non-stationarity and a proper detection and characterization of this non-stationarity is a major challenge in time series analysis. Among the possible types of non-stationarities, regime switching (or mosaic organization) plays a major role. In regime switching models, the parameters of the model change abruptly from time to time and the time series is organized in consecutive patches each one characterized by a set of model parameters. There is a vast amount of literature of regime switching models [1] and among the most used models it is worth mentioning the family of Hidden Markov Models [2]. Most models describe processes in discrete time, even if some models in continuous time (i.e point processes) exist [3].

An example of a process in which regime switching processes are applied is the case of human heartbeat interval dynamics. Studies showed [4] that the time series

of heartbeat intervals are organized in consecutive temporal segments with different local mean heart rates. Proper segmentation of the heartbeat data can bring up relevant physiological information, with the parameters possibly differentiating between healthy and ill patients. Another real example that we will consider in this paper is the inventory time series of financial market members. Market members when trading a large amount of shares tend to split their order into smaller transactions in order to limit their own impact on the market [5,6,7]. This strategic order splitting leads to long regimes when a market member is consistently buying or selling. By studying the inventory time series (the number of stocks owned at any given moment by a market member) these regimes can be identified and important information on traders' behaviour can be assessed.

Given an empirical time series which is assumed to be described by a regime switching process, the model fitting must be able to determine both the boundaries between consecutive patches and the model parameters

^a E-mail: bence@santafe.edu

in each patch. Clearly, given the probabilistic nature of the model, the boundary between two consecutive regimes is not directly observable and must be inferred from the data. A segmentation algorithm is a statistical method able to identify the different regimes (patches) of a regime switching time series.

In this paper we generalize an existing algorithm [4], originally developed to identify non-stationarity in the heart rate, as a method for segmenting regime switching compound Poisson processes. A compound Poisson process is a point process characterized by a rate (like a Poisson process) and a signal intensity distribution. In a normal Poisson process the signals are always of unit intensity. The associated counting process (i.e. neglecting the intensity of the signals) is a Poisson process and the signal intensities are independent and identically distributed random variables and are also independent of the counting process. The compound Poisson process is the simplest example of a so-called marked point process¹. Physical examples of such types of models include for example the celebrated Continuous Time Random Walk [9], which is the integral of a marked point process. Marked point processes have been applied to many different fields, ranging from earthquakes [10] to financial time series [11,12,13].

Here we are interested in segmentation of regime switching models of compound Poisson processes. We develop a null model of a time series made of consecutive patches of variable length such that in each patch the time series is described by a compound Poisson process. The parameters of the process are different in different patches and the length of the patches is described by an independent and identically distributed stochastic process with values drawn from a given distribution (in our case a lognormal distribution).

Our segmentation algorithm is a generalization of the method introduced in [4]. This is a top-down method, i.e. it first splits the whole time series in two subsets and then continues iteratively by breaking the series down to a more and more refined partition. One of the advantages of this method and of our generalization is that they are non parametric methods, i.e. they do not postulate a known distribution of the patch length and of the signal intensity (or jump size). The method of Ref. [4] has been recently

¹ A marked point process is a point process where both the counting process (rate) and the signal intensity (jump size) process are generic stochastic processes. Moreover, in general, in a marked point process the counting process and the jump process are not independent [8].

used in Refs. [6,14] for segmenting financial time series of the market members' trading activity in the Spanish Stock Exchange and in the London Stock Exchange (LSE). In this paper we will also consider the application of our algorithm to the segmentation of inventory time series of LSE market members. In order to do this in a financially reasonable way, we generalize the null model considering the possibility of long inactivity periods between two consecutive patches.

The paper is organized as follows. In Section 2 we discuss our null-models of time series. In Section 3 we introduce our segmentation methods and present simulation results. We present empirical results for the inventory of market members at the LSE in Section 4 and finish the paper with conclusions in Section 5.

2 Null models

We present two variants of the null model of time series that we expect to segment. The first is a pure regime switching compound Poisson process while in the second there may be long inactivity periods between two consecutive active patches. This second model is designed to better describe the financial time series investigated in Section 4.

In the pure regime switching compound Poisson process, the patches with different Poisson rate of activity follow each other. We consider the simplest case in which the jump size can be ± 1 . Since jump sizes are independently and identically distributed random variables, the jump sizes in a patch are statistically characterized by one probability (say, the probability of a jump size equal to $+1$). The sign associated with the largest probability is the dominant sign of the patch. Moreover we assume that patches have length T_{act} distributed according to a lognormal distribution,

$$T_{act} = C_{act} e^{N(1,1)}, \quad (1)$$

where C_{act} determines the characteristic length of the patches, $N(1,1)$ is a normal distribution with mean $\mu = 1$ and $\sigma = 1$. The mean and standard deviation of T_{act} are

$$\langle T_{act} \rangle = e^{3/2} C_{act} \simeq 4.48 C_{act} \quad (2)$$

$$\Delta T_{act} = \sqrt{(e^1 - 1)e^3} C_{act} \simeq 5.87 C_{act} \quad (3)$$

In order to generate a realization of the process, for each patch we choose a patch length T_{act} , a dominant sign (± 1

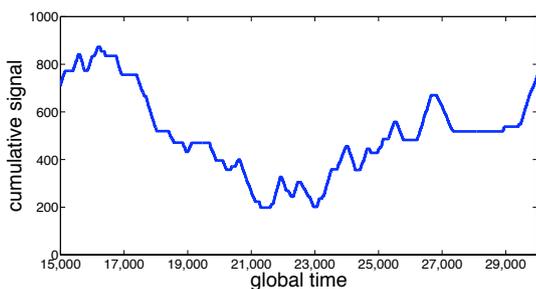


Fig. 1. Snapshot of the integral of a simulated time series of a regime switching compound Poisson process with inactivity periods, with parameters $C_{act} = 50$, $C_{inact} = 30$, $\delta = 0.2$ and $\eta = 0$.

randomly), a noise level $0 \leq \eta \leq 1$, and a rate α . The dominant sign and the noise level determine the jump process. With probability $1 - \eta$ the jump has size equal to the dominant sign and with probability η it has the opposite sign. The rate α determines the counting process, i.e. the probability per unit time that an event occurs is α . This parameter is drawn in each patch from a uniform distribution in the interval $[0.5 - \delta, 0.5 + \delta]$, where $0 \leq \delta \leq 0.5$ is the dispersion of the rate. For each patch, we generate a compound Poisson time series with the chosen parameters and of length T_{act} .

In the second type of null model we also add inactive patches, i.e. patches when no jumps are present between each pair of consecutive active patches. The active patches are generated in the same manner as before. Inactive patches have length T_{inact} distributed according to a lognormal distribution:

$$T_{inact} = C_{inact} e^{N(1,1)}, \quad (4)$$

where, as before, C_{inact} determines the characteristic length of the patches. Figure 1 shows a snapshot of the integral of the time series. In our simulations we varied the ratio $\langle T_{inact} \rangle / \langle T_{act} \rangle = C_{act} / C_{inact}$ between the mean length of active and inactive patches, sweeping the ratio from 0.02 (almost negligible length of inactivity) to 1 (equal mean length of active and inactive patches).

We simulate the process in discrete time with steps of unit time length. Moreover we consider two different times for the process. The first is the *global time* and is the (coarse grained) real time. The second is the *local time* and in this coarse grained time we consider only the times when an event occurs. In other words, in local time we discard all the times when there are no events and we end up with a time series containing only ± 1 . This second time discards all the information related to the point pro-

cess nature of the time series and preserves only the jump (event) process.

3 Segmentation method

3.1 Segmenting regime switching compound Poisson processes

As we have mentioned above, our segmentation algorithm is based on the algorithm in Ref. [4]. This algorithm works as follows. We move a sliding pointer over the series and for each position of the pointer we measure the mean of signal in the subset to the left and to the right of the pointer. We search for the position of the pointer for which the t statistic is maximal (t_{max}). To measure the significance of the difference between the two subsets we compute a modified t statistic. We make a cut if its significance exceeds a given threshold that we set to 99% as in Ref. [6]. After the cut we continue the method recursively on the newly created subsequences, until no further cut can be made. It is important to note that before a new cut is accepted we compute the modified t value between the new segments and their neighbours and check if both values exceed the above significance level.

One direct method for segmenting a compound Poisson process is to consider the time series in local time and run the algorithm of Ref. [4]. We will refer to this method as the *local time t-test* throughout the paper. This approach has been used in Ref. [6] for segmenting inventory time series in financial markets (see also Section 4). However this approach is unable to identify changes in rate due to a regime shift because all the information on the rate is lost when one considers the compound Poisson process in local time. For this reason we introduce a new method for segmenting compound Poisson processes that we call the *global time t-test*. The idea is to apply the algorithm to the time series in global time, i.e. our series will be composed by many zeroes (in the coarse grained time intervals, when no event occurs) and ± 1 (when one event occurs). Naturally the power of the segmentation method depends considerably on the noise level η and on the variance of the rate of Poisson processes inside different patches. In order to study this aspect we made simulations for different values of the dispersion of the rate of the Poisson process in different patches. As specified in the previous section, the rate is drawn from a uniform distribution in the interval $[0.5 - \delta, 0.5 + \delta]$ and we vary δ . After generating the time series we run the local time t-test and the

global time t-test. We compare the segmentation made by each method with the true segmentation. All the results we show are the average of 100 simulation runs and each run is made of 100 patches.

For the assessment of the segmentation methods we compute the Jaccard index [15] between the detected segmentation and the true segmentation. We denote the number of point pairs that are in the same patch both in the true time series and the segmented one by M_{11} . The number of point pairs that are in the same patch in the true time series but not in the segmented one is M_{10} , and similarly the number of pairs being in different patches in the true time series but in the same patch in the segmented one is M_{01} . The Jaccard index is defined as

$$J = \frac{M_{11}}{M_{10} + M_{01} + M_{11}}. \quad (5)$$

A perfect segmentation has $J = 1$. In our case there are two possible ways of computing the Jaccard index. The first is to compute the Jaccard index in local time, i.e. by restricting the comparison to the time series where we have removed all the zeroes. We denote this measure by J_{Local} . The second way takes into account all time steps (even when there are no counts) and we denote it by J_{Global} . J_{Local} tells us what fraction of the signals are segmented well, while J_{Global} tells us, what fraction of global time is segmented well.

We studied both J_{Local} and J_{Global} for growing dispersion of the Poisson rates and different noise levels. Figure 2 shows J_{Local} for the local time t-test and the global time t-test. We also included the results that come from entirely random segmentation of the time series, applying the same number of cuts as the local time t-test. In this figure we set $\eta = 0$. We can clearly see that for low values of δ the two methods perform roughly the same, while with growing δ the global time t-test outperforms the local time t-test by a large margin. The same result is seen in Figure 3 where we show J_{Global} .

Figure 4 shows the local Jaccard index J_{Local} for a model with $\delta = 0.5$ and for different values of η . Clearly, J_{Local} decreases as the noise level increases, but for all values of $\eta \lesssim 0.4$ the global time t-test outperforms the local time t-test segmentation algorithm, and for $0.4 \lesssim \eta \leq 0.5$ the two methods perform the same. A similar result is observed also with the global Jaccard index, as can be seen in Figure 5.

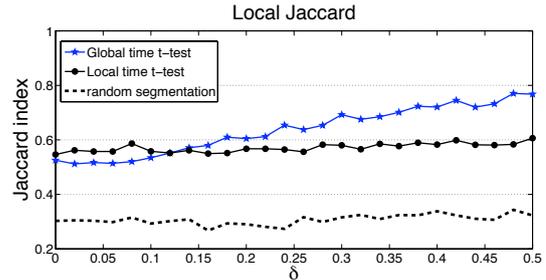


Fig. 2. (Color online) The local Jaccard index J_{Local} as a function of the dispersion δ of the rate α for the local time t-test (black circles), the global time t-test (blue stars). The black dashed line shows the Jaccard index for random segmentation. Here we set $\eta = 0$.

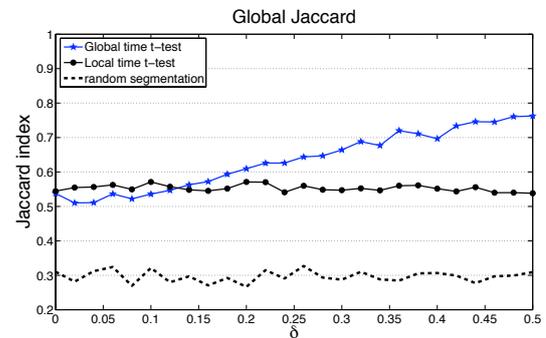


Fig. 3. (Color online) The global Jaccard index J_{Global} as a function of the dispersion δ of the rate α for the local time t-test (black circles), the global time t-test (blue stars). The black dashed line shows the Jaccard index for random segmentation. Here we set $\eta = 0$.

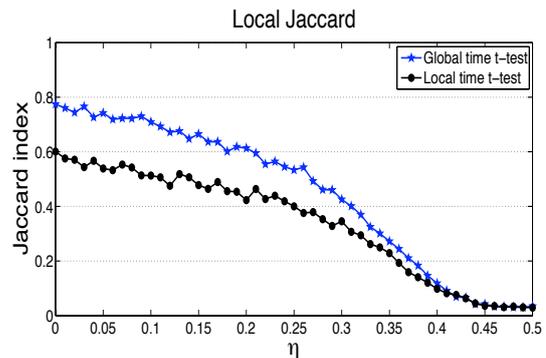


Fig. 4. (Color online) The local Jaccard index J_{Local} as a function of the noise η for the local time t-test (black circles), the global time t-test (blue stars). Here we set $\delta = 0.5$.

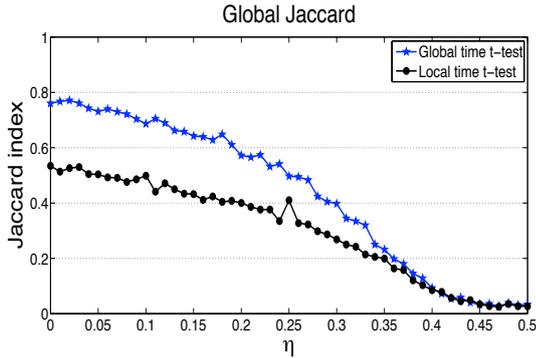


Fig. 5. (Color online) The global Jaccard index J_{Global} as a function of the noise η for the local time t-test (black circles), the global time t-test (blue stars). Here we set $\delta = 0.5$.

3.2 Segmenting regime switching compound Poisson processes with inactivity periods

In the case of regime switching compound Poisson processes with inactivity periods we modify the above algorithm. The segmentation algorithm, called the *composite test*, has two modules. First we run the global time t-test as described above. After the global time t-test, there is a second module that we call the *rate test*: we check if the patches found by the global time t-test are consistent with the Poisson assumption. For each patch we estimate α as the inverse of the average inter-event time in the patch. Then for each patch, we search for the longest period of inactivity between two events and we test whether it is consistent with Poisson waiting times. Given a Poisson process with rate α , the probability of having at least one waiting time longer than W is

$$\mathbb{P}(\omega_{max} > W) = 1 - (1 - e^{-\alpha W})^N, \quad (6)$$

where N is the number of jumps. Given the longest observed waiting time ω_{obs} , we make additional cuts if the probability of having a waiting time $\omega_{max} \geq \omega_{obs}$ is less than a given threshold q i.e.,

$$\omega_{obs} > \frac{\log(1 - (1 - q)^{1/N})}{-\alpha}. \quad (7)$$

In this paper we use $q = 0.01$. If Eq. 7 holds, we make two additional cuts, one at the beginning and one at the end of the inactivity period and then we continue the process recursively on the new sequences.

To present the power of our composite test we show how it works for simulated time series of regime switching compound Poisson process and inactivity periods as described in the previous section. For activity periods we

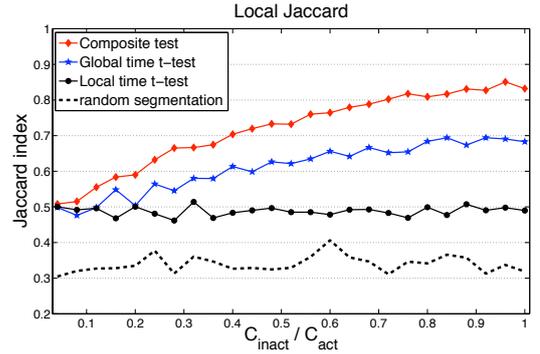


Fig. 6. (Color online) The local Jaccard index J_{Local} for the local time t-test (black circles), the global time t-test (blue stars) and the composite test (red diamonds). The black dashed line shows the Jaccard index for random segmentation.

draw the Poisson rate randomly from the interval $\alpha \in [1/15, 1/5]$. We consider here the noiseless case ($\eta = 0$), but we observe similar results for all values of η . After generating the time series we test the local time t-test segmentation method and the two methods introduced here, the global time t-test and the composite test. We compare the segmentations obtained by each method with the generated segments by computing the Jaccard index in local time and in global time. As before, all results we show are the average of 100 simulation runs.

Figure 6 shows the results for J_{Local} . We present the 3 methods, the local time t-test (black circles), the global time t-test (blue stars) and the composite test (red diamonds). With the dashed line we also show the Jaccard index obtained by an entirely random segmentation, applying the same number of cuts as in the local time t-test. It can be clearly seen that the curve of the composite test is well above the other curves. For very short inactivity periods the 3 methods perform roughly same, while for growing length of inactivity periods the performance of the composite method (red diamonds) becomes better quite fast. The local time t-test (black circles) performs uniformly when changing the ratio of active and inactive patch length.

Figure 7 shows the results for J_{Global} . Again we see that for very short inactivity periods the 3 methods perform roughly the same, while for longer inactivity patches the composite method (red diamonds) significantly outperforms the other methods. The global time t-test performs roughly constant as C_{inact}/C_{act} changes, while the performance of the local time t-test declines as the relative length of inactivity periods increases.

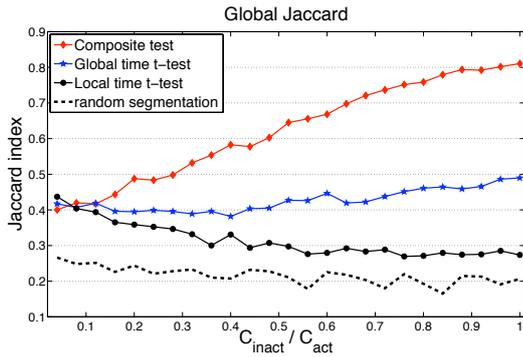


Fig. 7. (Color online) The global Jaccard index J_{Global} for the local time t-test (black circles), the global time t-test (blue stars) and the composite test (red diamonds). The black dashed line shows the Jaccard index for random segmentation.

From both plots we conclude that our proposed composite test outperforms the direct application of the original test to time series which can be described by regime switching compound Poisson processes (with or without inactivity periods).

4 Empirical analysis of financial data

As already mentioned in the introduction, a real example where the proposed segmentation methods can be very useful is the analysis of financial inventory time series and the determination of hidden orders.

It has recently been shown that the signed market order flow, i.e. the time series of the sign of the market orders initiating the trades is a long memory process [16,17]. It has been proposed that the long memory property is due to the practice of order splitting [5]. When a large investor decides to trade a large volume, it is quite unlikely that she places one large order in the market. What typically happens is that large trading orders are split into pieces and executed incrementally. We call these large orders *hidden orders*. For strategic reasons traders attempt to keep the true size of their orders secret in order to minimize transaction cost and trade the order at a more favorable price (for a review of this problem, see [18]). Empirical evidence of a widespread practice of order splitting has been given in the empirical results of Refs [19,20,21,6,14]. More recently it has been shown more directly that order splitting is the main cause of the long memory of order flow [7].

From an empirical point of view the detection of hidden orders is difficult because market participants do not

reveal their trading intentions and strategy. It has been recently proposed to detect hidden orders in a statistical way from the analysis of the time series of the inventory of market members [6]. Inventory variation is the signed transaction volume of a given market member, where conventionally a buy (sell) transaction has a positive (negative) volume equal in absolute value to the number of shares traded. Beside trading for themselves, market members act often as brokers or dealers. Therefore they are often not the real initiators of hidden orders, but the assumption of this approach is that it is unlikely that more than one large hidden order is given simultaneously to one broker. Under this assumption a large hidden order should be visible in the inventory time series of the broker as a random walk with a very strong drift. This approach has been proved useful and applied to the determination of the market impact of hidden orders [14]. The segmentation method used in [6,14] is the local time t-test of [4] applied to the time series of inventory variation of each broker *in local time*. Now, if two consecutive buy hidden orders with the same typical volume per trades (largely determined by the liquidity present in the market) are placed by the same broker with a time difference between the end of the first and the beginning of the second one that is short compared to the length of the two hidden orders, a local time t-test segmentation method will identify them as a single hidden order, because the algorithm has no way of separating them. For this reason, the main claim of this paper is that a global time t-test segmentation method should work better in identifying hidden orders.

One possible financial interpretation of the null model with inactivity periods described in Section 3.2 is the following. Brokers typically use for hidden order splitting a VWAP (Volume Weighted Average Price) algorithm [22], which is a simple trading protocol in which the trader splits the order in pieces and trades it incrementally in order to keep approximately constant the ratio, α_V , between the traded volume by him and the total volume simultaneously traded on that stock in the market. We assume that a VWAP strategy is well represented by a Poisson process with a rate related to the participation rate of the order. The noise in the inventory takes into account the fact that the broker is acting on behalf of many small customers while working the large order. Finally, the inactivity period takes into account the possibility that the broker trades with a bursting activity corresponding to

hidden orders and separated by inactivity periods, as recently shown empirically [7].

The database we study is the on-book (SETS) market of the London Stock Exchange (LSE), for the period of January 2002 through December 2004. The data set contains all orders placed together with the participant code of the market member placing them. A member can both act as a broker, i.e., handling trades for other institutions that are not members of the market, and may trade for its own account. Thus a single membership code may lump together trades from many different institutions. We call patches the segments that the statistical methods identify. Patches with a well-defined direction (mostly buying or mostly selling) can be understood as hidden orders. In the following we present results from empirical financial data applying both the local time t-test (as in Vaglica et al. [6]) and our composite test for segmenting the time series.

We study the inventory variation of a market member in *global transaction time*. This means that for each transaction, in which the studied participant is not involved we set his inventory variation to 0, and when he makes a transaction we put $+v$ for a buy trade and $-v$ for a sell trade. Here v is the volume of the transaction in number of shares. In this way, the time series of the inventory variation contains all information on the participation rate of the agent. We present here results on the data of AstraZeneca (AZN), a highly liquid stock of the LSE. We study the 50 most active market members for this period. We consider patches (hidden orders) with a well defined direction. For this reason, as in Refs. [6,14], we apply some filters to the patches found by the methods. Specifically, we define directed patches as those where at least 75% of the volume of transactions inside the patch have the same direction. Furthermore we only study patches that are constituted of at least 10 transactions. It is also important that in case of financial data, when applying the rate test, we only check for inactivity periods that are at least 50 transactions long. In real time, this corresponds on average to 20 minutes of trading.

As we can expect the composite test identifies more directional patches than the local time t-test. In case of AZN the original algorithm determines 3702 directed patches, while the composite test finds 10613. This difference suggests that the local time t-test segmentation method tends to stick several patches together, not taking into account changes in the rate and the presence of inactivity periods.

Table 1. Summary statistics of the local time t-test and the composite test applied to financial data. The number of directed patches found (num), the average number of transactions made in a directed patch (\bar{N}), the average length of a directed patch in global transaction time (\bar{T}), the average participation rate in directed patches ($\bar{\alpha}_V$), and the average fraction of market orders in directed patches (\bar{f}_{MO}).

measure	Local time t-test	Composite test
num	3702	10613
\bar{N}	171.6	64.6
\bar{T}	1267	286
$\bar{\alpha}_V$	0.15	0.23
\bar{f}_{MO}	0.48	0.46

This problem is taken into account by our composite test. This way we expect a better resolution of the patches and less intrinsic error in the detection.

As characterizing statistics of the directed patches, we study the number of transactions made inside the patches, N , the length of directed patches in global transaction time, T , the participation rate in the patch, α_V , and the fraction of transactions made via market orders in the patch (f_{MO}). In Table 1 we summarize some basic statistics of the results, showing the number of directed patches found with each method and the average of the above measures computed over all directed patches found.

Figure 8 shows the cumulative distribution of the number of transactions N in the directed patches. We can see that the distribution of patch length found using the composite test has a smaller variance than the one found with the local time t-test. In both cases we find fat tailed distributions. The average number of transactions constituting a directed patch in case of the local time t-test is 171.6, while in case of the composite test it is 64.6. Concerning the length of directed patches in global transaction time we find a very large difference (see Table 1): in case

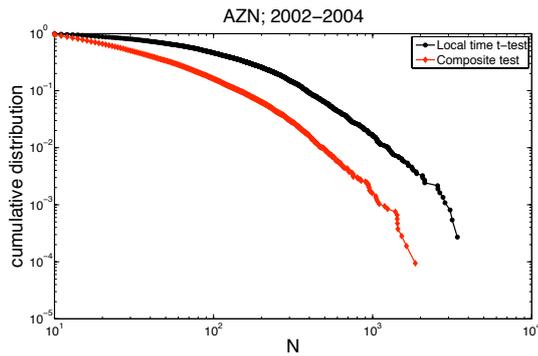


Fig. 8. (Color online) Cumulative distribution of the patch length for directed patches found by the local time t-test (black circles) and the composite test (red diamonds).

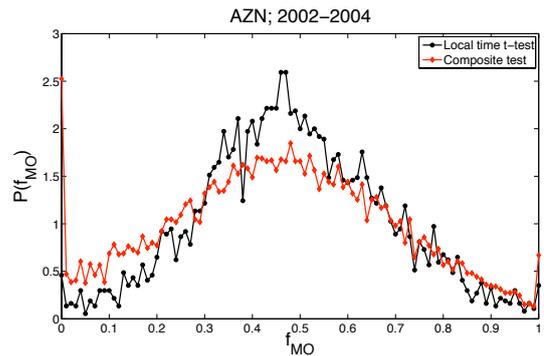


Fig. 10. (Color online) The distribution of the ratio of market orders for directed patches found by the local time t-test (black circles) and the composite test (red diamonds).

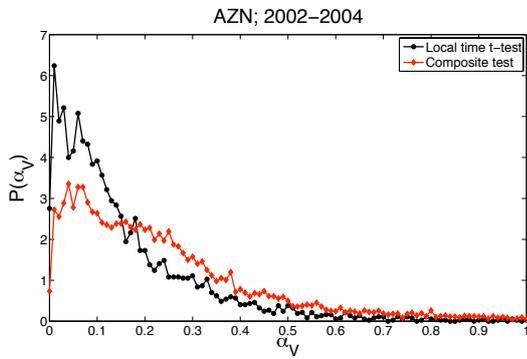


Fig. 9. (Color online) Distribution of the participation rate for directed patches found by the local time t-test (black circles) and the composite test (red diamonds).

of the local time t-test $\bar{T} = 1267$ transactions, which is almost one trading day, while for the composite test $\bar{T} = 286$ transactions.

Figure 9 shows the distribution of the rate (participation rate) for the directed patches found by the two methods. The distribution for the directed patches found by the local time t-test is peaked close to zero, with an average of 0.15, while the distribution of those found by the composite test have a maximum at a higher value and decaying slower, having an average of 0.23. A simple reason for this difference might be the fact that the composite test finds shorter patches on average and in case of shorter patches it is more likely to find high participation rate than for long patches (it is very hard to trade with a high participation rate for long periods). Also, if the patches detected with the local time t-test are composed by several shorter patches detected with the composite test and interspersed by inactivity period, it is clear why the rate is larger in the new than in the local time t-test.

Figure 10 shows the distribution of f_{MO} for the two methods. Both distributions seem to be symmetric roughly around $f_{MO} = 0.5$. However we find that in case of the patches found by the local time t-test, there is a higher weight of the distribution in the middle, and lower values at the extremes. In case of the composite test, the weight in the middle is lower and there are large values for the extreme cases, i.e. $f_{MO} = 0$ or $f_{MO} = 1$. Especially, for $f_{MO} = 0$ we find the highest value of the distribution. Again these differences are probably due to the fact that the composite test finds shorter patches. In case of shorter patches, it is more probable that the participant sticks to one type of order (in our case limit orders).

4.1 Comparing the segmentations

An important question is to know the differences between the segmentations found by the two methods. To understand if the composite test is a “refinement” of the local time t-test (finding roughly the same patches and cutting them further into pieces) or if it makes an entirely different segmentation, we study the distance of the cuts made by the local time t-test from the nearest cut made by the composite test. We find that this distance is on average one fourth of the distance that would be found if the two cuts were made independently and on average 90% of the distances are lower than what we can expect in the independent case. This means that usually there is a cut made by the local time t-test not far from the cuts made by the composite test. We also measured the “nesting” of the two methods. Specifically, we measure the fraction of patches found by the composite test, that are entirely contained in a patch found by the local time t-test. For this nesting

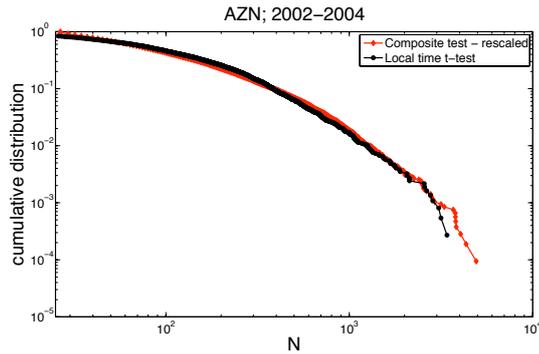


Fig. 11. (Color online) Cumulative distribution of the patch lengths for directed patches found by the local time t-test (black circles) and the rescaled cumulative distribution of the patch lengths found by the composite test (red diamonds).

coefficient for AZN we find a value of 0.87 ± 0.09 . This high value suggests that the composite test can partially be seen as a refinement of the local time t-test. On Figure 11 we plot the cumulative distribution of N for the two methods (similarly to Figure 8), but rescaling the N values of the composite test by the ratio of \bar{N} : $171.6/64.6 \approx 2.65$ (see Table 1). The two distributions are similar, suggesting again that the composite test is a homogeneous refinement of the local time t-test.

We study the similarity of the two segmentations by computing the Jaccard index of the two patch structures. For the 50 studied brokerage codes, we find a global Jaccard index of $J = 0.24 \pm 0.21$. This average is interestingly not far from the value that we expect if the composite test were a homogeneous refinement of the local time t-test. In fact the Jaccard index between a partition and its refinement where each patch in the first partition is divided in M patches of equal length, is $1/M$. In our case this would lead to $\bar{T}_{composite}/\bar{T}_{localtime} = 286/1267 \approx 0.22$.

5 Conclusions

We presented two related methods for segmenting time series that show non-stationarity in both time (rate) and intensity. Specifically our segmentation method has been developed and tested on regime switching compound Poisson processes. In this type of processes the events occur as in a Poisson process and each event is associated with a signal described by an independent and identically distributed stochastic variable. Moreover the time series is organized in patches where the parameters of the model are constant while they are different in different patches.

We have shown that our algorithms perform quite well in segmenting the simulated time series in a wide range of possible parameter distributions.

Even if the algorithm has been designed and tested for regime switching compound Poisson processes, it can be used for more general regime switching marked point processes. This is due to the non parametric nature of the segmentation method. Therefore we expect that the method performs quite well in segmenting regime switching time series, where in each patch (regime) the signal is a point process with an intensity associated to each event. Examples of time series described by marked point processes can be taken from a wide range of phenomena observed, for example, in physiological (neuron spikes), geophysical (earthquakes), astrophysical (solar flares), and socioeconomic (financial transactions) systems.

We also presented an application of our segmentation algorithm to financial time series. Specifically, we consider the segmentation of the time series describing the inventory of market members. The method looks for periods of time when a market member consistently buys or sells at a roughly constant rate. We postulate that these periods correspond to hidden orders, i.e. large orders that are split and traded incrementally over a long period of time. Hidden orders are an important element for a proper characterization of order flow, which in turn is one of the main ingredient needed to understand the price formation process. Given the privacy issues related to the trading activity of market investors, the detection of hidden orders at a financial market scale can be performed only with the aid of statistical methods. Some of the recent efforts in this direction include the methods in Refs. [6,14,23,24]. The method we introduced in this paper is the first to consider the segmentation in global transaction time, i.e. it is the only one that takes into account the trade rate of a given market member. By contrast the other existing methods work in local time, i.e. discard all the information related to the trading rate. When we compare the statistics of the patches found by our algorithm with those found by the algorithm of Refs. [4,6] we find almost three times more patches. This suggests a better resolution in the detection of hidden orders. Comparing the results of the two methods we find evidence that the segmentation of the composite test can be understood as a refinement of the patches found by the local time t-test.

Acknowledgments

B.T and F.L acknowledge Gabriella Vaglica for useful discussions. F.L. acknowledges financial support from the PRIN project 2007TKLTSR “Computational markets design and agent-based models of trading behavior”. J.D.F. and B.T. acknowledge support from NSF grant HSD-0624351. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. J.D. Hamilton, *Regime-Switching Models*, in Palgrave Dictionary of Economics 2nd ed., edited by Steven Durlauf and Lawrence Blume, (Palgrave MacMillan Ltd., 2008)
2. L. R. Rabiner, *Proceedings of the IEEE*, **77**, 257-286 (1989)
3. S. Chib, *Journal of Econometrics* **86** 221-241 (1998)
4. P. Bernaola-Galván, P. Ch. Ivanov, L. A. N. Amaral, H. E. Stanley, *Phys. Rev. Lett.*, **87**, 168105 (2001)
5. F. Lillo, Sz. Mike, and J.D. Farmer. *Phys. Rev. E*, **71** 066122 (2005)
6. G. Vaglica, F. Lillo, E. Moro, R. N. Mantegna, *Phys. Rev. E*, **77**, 036110 (2008).
7. B. Tóth, I. Palit, F. Lillo, J.D. Farmer (in preparation).
8. G. Last and A. Brandt, *Marked Point Processes on the Real Line: The Dynamic Approach*, (Springer, 1995).
9. E.W. Montroll, G.H. Weiss, *Journal of Mathematical Physics* **6** 167-181 (1965)
10. L. Holden, S. Sannan, H. Bungum, *Natural Hazards and Earth System Sciences* **3**: 95-101(2003)
11. Engle, R, Russell, J. *Econometrica*, **66**, 1127-1162 (1998)
12. E. Scalas, R. Gorenflo, F. Mainardi, *Physica A* **284**, 376 (2000)
13. M. Montero, J. Perelló, J. Masoliver, F. Lillo, S. Miccichè, R.N. Mantegna, *Phys. Rev. E* **72**, 056101 (2005).
14. E. Moro, J. Vicente, L. G. Moyano, A. Gerig, J. D. Farmer, G. Vaglica, F. Lillo, R. N. Mantegna, preprint: <http://arxiv.org/abs/0908.0202> (2009). Accepted in *Phys. Rev. E*.
15. P. Jaccard, *Bulletin del la Société Vaudoise des Sciences Naturelles* **37**, 547-579 (1901)
16. F. Lillo and J. D. Farmer, *Studies in Nonlinear Dynamics & Econometrics*, **8**, 1, (2004)
17. J.-P. Bouchaud, Y. Gefen, M. Potters, M. Wyart, *Quantitative Finance*, **4** 176-190 (2004)
18. J. P. Bouchaud, J. D. Farmer, F. Lillo, *How markets slowly digest changes in supply and demand*, in *Handbook of Financial Markets: Dynamics and Evolution*, T. Hens and K.R. Schenk-Hoppé (Editors), 57 (North-Holland, 2009)
19. L. K. C. Chan, J. Lakonishok, *J. Financ.* **50**, 1147, (1995)
20. R. Almgren, C. Thum, H. L. Hauptmann, H. Li, *Risk*, July 2005, 57 (2005)
21. D. R. Gallagher, A. Looi, *Account. Finance* **46**, 125 (2006)
22. R. Kissell and M. Glantz. *Optimal Trading Strategies: Quantitative Approaches for Managing Market Impact and Trading Risk*. American Management Association, (2003).
23. A.N. Gerig, *A theory for market impact: How order flow affects stock price*. Ph.D. Thesis. University of Illinois (2007).
24. G. Vaglica, F. Lillo, R. N. Mantegna, *Identification and characterization with hidden Markov models of large orders splitting strategies in a financial market*, (submitted).