Impact of unexpected events, shocking news, and rumors on foreign exchange market dynamics

Mark McDonald,[1,2] Omer Suleman,[1] Stacy Williams,[2] Sam Howison,[1] and Neil F. Johnson[3]
[1]*Mathematics Department, Oxford University, Oxford OX1 2EL, United Kingdom*
[2]*FX Research and Trading Group, HSBC Bank, 8 Canada Square, London E14 5HQ, United Kingdom*
[3]*Physics Department, University of Miami, Coral Gables, Florida 33146, USA*
(Received 15 January 2008; published 17 April 2008)

The dynamical response of a population of interconnected objects, when exposed to external perturbations, is of great interest to physicists working on complex systems. Here we focus on human systems, by analyzing the dynamical response of the world's financial community to various types of unexpected events—including the 9/11 terrorist attacks as they unfolded on a minute-by-minute basis. For the unfolding events of 9/11, our results show that there was a gradual collective understanding of what was happening, rather than an immediate realization. More generally, we find that for news items which are not simple economic statements—and hence whose implications for the market are not immediately obvious—there are periods of collective discovery during which opinions seem to vary in a remarkably synchronized way.

DOI: 10.1103/PhysRevE.77.046110 PACS number(s): 89.75.Fb, 89.75.Hc, 89.65.Gh

I. INTRODUCTION

Governments and planning agencies often second-guess people's collective response to important pieces of news and rumors—for example, stories concerning terrorist attacks, security risks, health scares, and economic upsets. There have even been suggestions that governments tend to bury bad news on a bad news day [1]. In the field of complex systems research, there are many physicists who have taken an interest in the dynamical properties of collections of humans—in particular, researchers such as Sornette *et al.* have carried out some fascinating analysis of peoples' collective behavior and response in real-world situations [2–9]. In Refs. [3–5] the reaction of web downloads and book sales to positive news coverage is considered, and a remarkably general result is found. In Ref. [7] the focus returns to the reaction of financial markets to news events. By comparing the reaction of market volatility to the terrorist attacks of September 11, 2001, the attempted coup in Russia during 1991, and the market crash in 1987 (Black Monday), the authors conclude that Black Monday was a purely endogenous event. In Ref. [2] concerning the Black Monday crash, log-periodic structures were found in stock index data both before and after the crash and in other smaller crashes. In Ref. [9] an Ising model is postulated which is capable of reproducing observed behavior of financial markets as well as the main features of the work described above.

Here we analyze the dynamical response of the global population of foreign exchange (FX) currency traders to various types of news. There is a vast economic literature on the effect of news on markets [10–24]. However, most of these studies focus on specific economic news announcements—in particular, the effect of having an announced economic indicator (e.g., unemployment rate) which differs from the expected value. In addition, such studies do not tend to follow the resulting temporal dynamics of the market following such announcements. By contrast, we examine the temporal response following different types of unexpected news—not just economic announcements. One of our main findings is that the dynamical response following noneconomic news shocks can differ significantly from that following market-specific news shocks. In particular, our results suggest that the former case of non-market-specific news requires a period of collective discovery as to the news' meaning, during which the population of traders struggles toward a consensus. The associated collective dynamics can be highly oscillatory as opinions fluctuate—what is remarkable, however, is that this response turns out to be well synchronized despite the underlying uncertainty as to the long-term implications of the news. By contrast, the implications of market-specific news in the latter case can be assimilated quickly at the level of each individual. Hence a consensus can be reached quickly and no collective oscillatory behavior is observed.

Prior to the September 11 terrorist attacks, economists had shown little interest in the effect of general (i.e., non-market-specific) news on markets [25]. However, interest has since burgeoned. Some of this interest has been focused on general terrorism [25–27]; but most has been focused on September 11 itself—see, for example, Refs. [28–36], which include the effects on stocks [28,29], insurance [32,33], improving the regulatory response to future shocks [31], disruption to infrastructure [30], and even effects on real estate [34]. However, the effects on the FX market have been overlooked—which is surprising given that the FX market is the largest financial market in existence, and is running perpetually. In particular, the worldwide FX markets remained trading in liquid quantities throughout the September 11 attacks. Two previous papers which did consider aspects of FX trading were Refs. [35,36]. However, Ref. [36] only looks at FX as part of a wide-ranging investigation into the effect of September 11 on many different markets, while Ref. [35] only looks at euro and U.S. dollar (EUR/USD) exchange-rate data aggregated into 10-min intervals. The main factor differentiating our work from these previous studies is the methodology that we use. While previous work focused on the actual value of prices, volatilities, or some measure of liquidity, we follow the physics tradition in focusing on the temporal behavior of correlations following such an external perturbation. In addition we are able to probe the market's response at very high frequency as a result of high-quality FX data [37,38]. Hence we avoid having to aggregate the data into typical 5- or 10-min intervals [24,35].

Like all other financial markets, currency traders are being continually bombarded by information about external events which may in turn affect their trading decisions. Recent research suggests that most everyday news items have relatively little effect on the resulting market movements [39]. But what about major news? Unlike most physical systems, such as a gas of electrons, individual humans may not all respond to external shocks in the same way. Worse still, even if a given piece of news is announced globally at a given moment, it will tend to reach different people at different times. Furthermore, their response times will tend to differ, as will the extent to which they believe the news to be true. As if this was not complicated enough, the way in which people respond will also differ in general. What is good news for one person or market may be bad for another—or just plain irrelevant. The September 11 terrorist attacks in the U.S. were unambiguously bad in terms of human lives and world affairs. But how is such news to be interpreted from the point of view of a trader dealing with a particular currency? If it is bad for the U.S. dollar, is it therefore good for the yen? After all, to sell a currency in the FX markets, you need to buy another one—which is why it is called an exchange rate. But in a moment of supposed global crisis, there may be no unambiguously safe currency to buy. So how should the unfolding of such news affect traders' decisions to buy or sell? Do they sell U.S. dollars and buy Japanese yen, or vice versa—or maybe do something entirely different? Or is there instead some more gradual process of collective learning—a sort of collective realization or discovery of the meaning and implications of such news? These are the issues that we attempt to address in this paper. The specific methodology that we use to address these issues is that of clustering [40] as explained in Sec. II. In Sec. III we describe the data sets themselves, before moving in Sec. IV to an analysis of the responses to the news. We find that similar results are seen for similar species of news. We focus on the following four case studies:

(i) the terrorist attacks of 11 September 2001 in the U.S.;

(ii) a false rumor that the Chinese currency (CNY) was about to be revalued (11 May 2005);

(iii) the real CNY revaluation (21 July 2005);

(iv) a particular piece of economic news in the form of an announced government statistic which turned out to be significantly different from that expected.

In Sec. V we establish that these results are robust to the choice of particular clustering method. Section VI provides the concluding summary.

II. CLUSTER ANALYSIS

Given that we measure the dynamical response using cluster analysis of exchange-rate correlations, it is important to establish exactly what we mean by such cluster analysis. This is the discussion that we now undertake in this section.

Cluster analysis is a field of statistics which attempts to classify objects into groups. When performing a cluster analysis, one must first choose a measure of the proximity between different objects, the dissimilarity measure, and then specify an algorithm which groups objects into clusters: the cluster algorithm. One important set of clustering algorithms are agglomerative hierarchical cluster algorithms. Performing such analyses on a set of N time series results in an indexed hierarchy, where each level of clustering is covered, beginning from N clusters of one object each, and progressing until all the objects are in one cluster containing N objects. The index details the distance at which two clusters are merged.

A frequently used measure of the similarity of two time series is their correlation,

$$\rho_{ij} = \frac{E[X_i - E(X_i)][X_j - E(X_j)]}{\sqrt{\mathrm{Var}(X_i)\mathrm{Var}(X_j)}}, \tag{1}$$

where $\mathrm{E}(X)$ and $\mathrm{Var}(X)$ denote the expectation value (i.e. mean) and variance of X respectively. The correlation in Eq. (1) is often estimated by Pearson's product-moment correlation coefficient

$$\hat{\rho}_{ij} = \frac{\sum_k (X_{ik} - \bar{X}_i)(X_{jk} - \bar{X}_j)}{\sqrt{\sum_k (X_{ik} - \bar{X}_i)^2 \sum_l (X_{jl} - \bar{X}_j)^2}}, \tag{2}$$

where $\bar{X}_i$ is the sample mean for X_i and similarly for $\bar{X}_j$ and X_j. For a set of N time series, one can form a correlation matrix $\mathbf{C}$, whose elements are $\mathbf{C}_{ij} = \hat{\rho}_{ij}$.

In Ref. [41], Mantegna proposed using minimum spanning trees (MSTs) as a clustering procedure to identify hierarchical structure in financial markets. The procedure proposed was to transform the matrix $\mathbf{C}$ into a distance matrix $\mathbf{D}$, where the elements of $\mathbf{D}$ are $d_{ij} = \sqrt{2(1-\hat{\rho}_{ij})}$. This distance matrix can be thought of as the adjacency matrix of a fully connected, weighted N graph, from which it is simple to construct the MST. Two algorithms commonly used to construct the MST are Kruskal's algorithm [42] and Prim's algorithm [43]. The MST has an indexed hierarchy associated with it, and it was this hierarchy that Mantegna suggested as a tool for defining a taxonomy of financial assets.

The dissimilarity measure chosen by Mantegna of $d_{ij} = \sqrt{2(1-\hat{\rho}_{ij})}$ is approximately the standardized Euclidean distance between the time series i and j [44,45]. Since there is evidence that the choice of dissimilarity measure has less effect on the clustering than the choice of clustering algorithm [46], we continue to use this measure for comparison with previous work; however, we note that for the MST any co-monotonic transformation of distances would give the same cluster structure [45,47].

The first attempt to produce a hierarchy of financial securities was in Ref. [48]; however, such techniques received little attention until Mantegna's paper [41]. Since then, there has been an explosion of papers in the area from the field of econophysics [40,41,49–60]. The MST clustering algorithm has been applied to interest rate markets [59,61], equity market indices [57], and (FX) markets [40]. In Ref. [62], the MST procedure was extended to include other links so that it was no longer a tree, but was the graph with the maximum amount of information which could be embedded in a surface. The graph which can be embedded in the surface of a sphere was termed the planar maximally filtered graph (PMFG). However, the hierarchical structure of this PMFG is identical to that of the MST and, as such, the procedure would not be of use here. The use of both the MST and the PMFG as tools for noise-undressing of correlation matrices has been investigated in Ref. [60].

Using the MST as a clustering tool is the same procedure that statisticians refer to as the single linkage clustering algorithm (SLCA) [63]. This is one of the simplest hierarchical clustering algorithms, and has been in use as a clustering algorithm since 1951 [64], although the first algorithm for calculating the MST was published in 1926 [65]. However, the simplicity of the algorithm can sometimes be detrimental. When merging two clusters, the SLCA/MST method chooses the shortest link between the two clusters and then defines the distance between all objects in the two different clusters to be this smallest distance, before merging the two clusters into a new one, indexed by this distance. As the algorithm progresses, particularly for large N, there is a tendency for the method to link objects into chains, where the objects at each end of the chain have little or nothing in common with each other. For this reason it is usual to compare the results from one clustering method with other clustering methods. If several methods which group objects in very different ways agree on the cluster structure, then one can be confident of the robustness of one's results.

One cluster algorithm which groups objects in a very different way to the SLCA is the complete linkage clustering algorithm (CLCA). Where the SLCA defines the distance between two clusters as the minimum of the pairwise distances between two objects, one in each cluster, the CLCA defines the distance as the maximum pairwise distance. If two such radically different clustering algorithms agree on the cluster structure, this is very strong evidence of robustness. While the SLCA groups objects into chains, the CLCA tends to group objects into compact groups, where unless all the objects in a group are very close together, they will not form a cluster until a very large distance. For this reason the SLCA is often termed as being *space contracting* and the CLCA as *space dilating* [45,47].

An algorithm which groups objects in a more moderate way is the average linkage cluster algorithm (ALCA). In this algorithm, the distance between two clusters being merged is the average of all the pairwise distances, one object in each cluster.

If the three algorithms above agree on the cluster structure, this is very strong evidence of the robustness of the structure being identified. *In this paper, we use all three to confirm the robustness of our results*. In addition, we use a further algorithm, Ward's Algorithm, which clusters objects by a minimizing variance procedure [66]. If this method agrees with the three linkage methods, then this provides extremely strong evidence of robustness.

All hierarchical cluster algorithms result in an indexed hierarchy. The distances of the objects in the fully connected N graph reside in a Euclidean space, and so obey the metric inequality: $d_{ij} \leq d_{ik} + d_{kj}$. However, the correct space in which to describe a hierarchy is an *ultrametric space* [47,67] in which the distances of objects under a clustering algorithm,

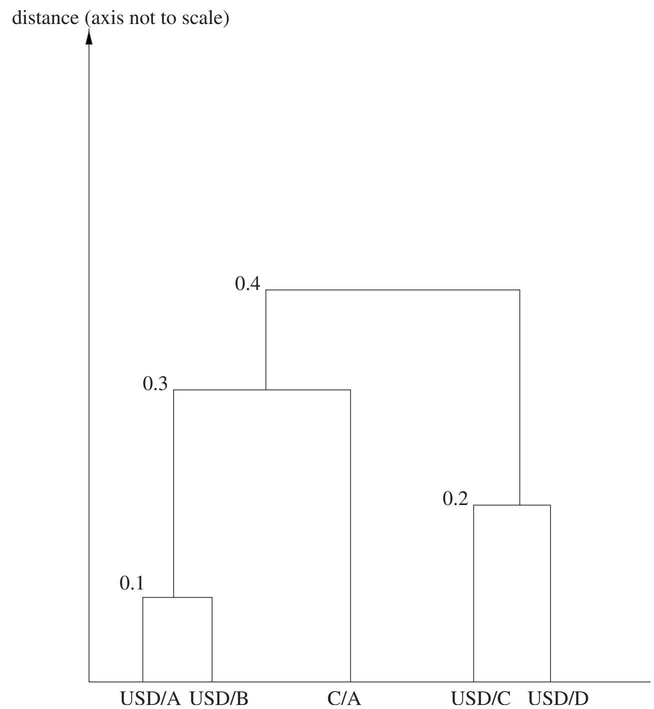

FIG. 1. A schematic section of a hierarchy to illustrate the definition of the clustering distance for a currency. In this simple example, there are five currencies: A, B, C, D, and USD (i.e., U.S. dollar). The clustering distance for USD is 0.4, since one has to travel up the hierarchy to a distance of 0.4 until all the USD-based exchange rates are included in the same cluster.

d^*_{ij}, obey the stronger ultrametric inequality: $d^*_{ij} \leq \max(d^*_{ik}, d^*_{kj})$. It is these ultrametric distances [68] which we use to define the clustering distance for a currency. We explain the definition of this clustering distance with the schematic hierarchy shown in Fig. 1.

In the FX market currencies are generally quoted against the dollar. If you know the price of two currencies against the dollar then it is trivial to calculate the price of one currency in the other. Since there is no fixed-value reference asset in terms of which they can be priced, one can only observe the relative value of two currencies, and not intrinsic value of a specific currency. However, when it comes to considering the correlation dynamics of the market then this is not appropriate. Consider the case where a strong correlation is observed between British pound/USD (GBP/USD) and EUR/USD. This could be caused by the intrinsic values of both EUR and GBP moving together in a synchronized way over time. Instead, it could be the case that the intrinsic values of these two currencies have remained basically unchanged but the intrinsic value of the USD has moved. It is only by including all possible exchange rates for a group of currencies that one is able to identify the rich and complex dynamics of the FX market. For a situation where one considers N currencies this gives $N(N-1)$ exchange rates.

Figure 1 shows a section of a hypothetical hierarchy resulting from using a hierarchical clustering method on the exchange rates from five currencies: A, B, C, D, and USD (i.e., U.S. dollar). The clustering distance for a currency is the minimum distance at which one can partition the hierar-

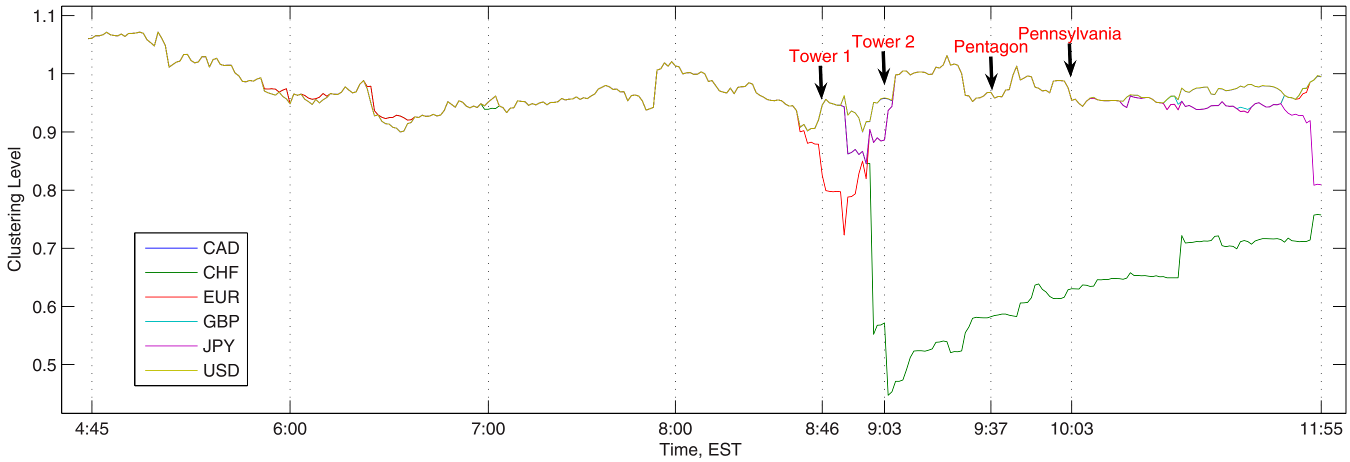

FIG. 2. (Color online) The clustering distances for September 11, 2001.

chy and get all the base rates for that currency in the same cluster. Let us consider what the clustering distance is for USD for the case shown in Fig. 1. In this example, if one partitions the hierarchy at 0.1, one has a cluster of (USD/A, USD/B), one cluster of (USD/C), and one cluster of (USD/D), so 0.1 is not the clustering distance. If one partitions the hierarchy at 0.2, there are two clusters containing USD based exchange rates: (USD/A,USD/B) and (USD/C,USD/D), so 0.2 cannot be the clustering distance. Partitioning the hierarchy at 0.3 still results in two clusters containing USD based exchange rates, one cluster of (USD/A,USD/B,C/A) and another cluster of (USD/C,USD/D). However, if one were to partition the hierarchy at 0.4, there would be one cluster containing (USD/A,USD/B,C/A,USD/C,USD/D). All the USD based exchange rates are in the same cluster—hence the clustering distance for USD in this rather simplistic example is 0.4. Note that as a result of the distance definition, a smaller distance corresponds to a stronger correlation.

III. DATA

The data used for this paper were provided by HSBC Bank. We use the last tick of each minute to form 1-min prices for all exchange rates between the following currencies: Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), euro (EUR), British pound (GBP), Japanese yen (JPY), and U.S. dollar (USD) [69]. This gives rise to 42 time series, from which we then form 1-min log returns.

When performing regressions or calculating correlations, it is necessary to ensure that each time series in the data is *autocovariance stationary* [71,72]. To this end, augmented Dickey-Fuller (ADF) tests (of various orders) [71,73,74] were performed on the returns. The null hypothesis for this test is that the data are nonstationary. However, when ADF tests were performed on the return data, this null hypothesis is strongly rejected—evidence that the returns are autocovariance stationary.

Missing values were deleted prior to calculating correlations. To preserve the positive-definite property required of correlation matrices, if a time step was deleted from one exchange rate, that time step was deleted from all currency pairs under investigation.

IV. RESULTS

We now turn to discuss each one of our four case studies in detail. To investigate the dynamics of the clustering distance in each case, we use a window of length $T=120$ time steps which we then move through the data one time step at a time. There are $N=42$ time series, so $T \geq N$, as required. For each window the clustering distance for each currency was calculated, as described above.

A. 11 September 2001

The results for 11 September 2001 are shown in Fig. 2. The equity markets closed as a result of the terrorist attacks, but the (FX) markets remained open thereby allowing us to investigate the effect of the unfolding news on global markets. Figure 2 shows the power of this technique, with the market impact by currency clearly visible.

The first feature to note in Fig. 2 is a contagionlike effect whereby clustering gets successively transferred between currencies. Prior to the first tower being hit at 8:46 a.m., there is no significant clustering in any currency, although the Euro seems to be slightly "in play" [40]. As soon as the first tower is hit at 8:46 a.m., a small amount of additional clustering is seen in the Euro.

Remarkably, the CHF then undergoes a dramatic clustering *prior to the second tower being hit*. In other words, even though the second tower has not yet been hit—and hence the hitting of the first tower can still be regarded as some form of awful accident—traders have converged on the notion that something strange is happening. As a result, the CHF comes into play, taking center stage in trading and hence giving rise to a very small clustering level as shown in Fig. 2. At the same time, there seems to be a small split in opinion in that the JPY is very slightly in play as well. As soon as the second tower is hit, and people therefore realize that this is more than just two very unlikely accidents, the CHF undergoes another dramatic reduction in clustering level. In essence, activity in the CHF is then driving the market.

Just as remarkable is the *absence* of any additional clustering features associated with the subsequent attacks, i.e., the Pentagon at 9:37 a.m. and the crash in Pennsylvania at

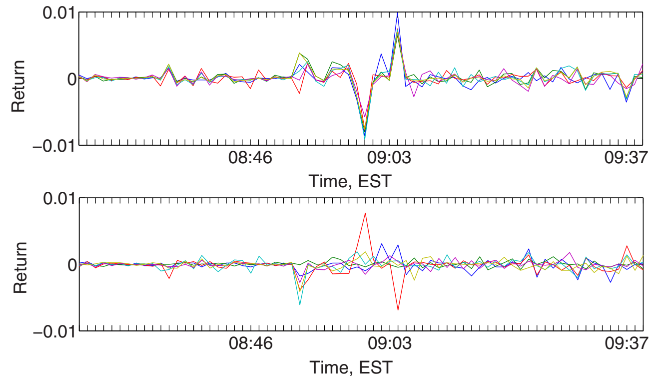

FIG. 3. (Color online) The returns for 11 September 2001, over a period of time which includes the first three attacks. Top panel: returns from the CHF-based exchange rates. Bottom panel: returns from the USD-based exchange rates. The horizontal axis is the same for both panels.

10:03 a.m. It is as though the market had already decided that the news is as bad as it can get. Hence these two otherwise extremely dramatic events—which on any other day would surely have driven the market into mayhem—then had no visible effect. It is as if they had not happened.

It is interesting that it is the CHF which is the clustered currency in Fig. 2 and not the USD. One might think that the reason is trivial—CHF may be acting as a "safe haven" currency which people buy in order to try to protect themselves in uncertain times. However, a quick look at the actual returns shows that this cannot be the complete answer. In particular, Fig. 3 shows the returns from a 1-h window for 11 September 2001. This 1-h period includes the first three attacks. The top panel shows the returns for the CHF-based exchange rates; the bottom panel shows the USD-based exchange rate returns over the same period. The horizontal time axis is the same in both cases. It can be seen that before the first attack, the market is quiet. The first large return comes a short time after the first plane hits the World Trade Center. In other words, *the market takes time to absorb the news, assimilate its meaning and validity, interpret its importance, and hence react to it by trading*. It can be seen that the reactions of the CHF exchange rates become *synchronized* for a short time, and are certainly more synchronized than the USD exchange rates. This is an important feature, since there is obviously no invisible hand or central controller coordinating such a collective dynamic and such an event is unlikely to occur by accident.

Most importantly, it is not simply the case that people are buying CHF and selling everything else. To a first approximation, increased demand for CHF (i.e., an increased number of buy orders) will increase the value of CHF/other exchange rates while reduced demand for CHF (i.e., an increased number of sell orders) will decrease them. We can see from Fig. 3 that there is a spontaneous synchronization of buy *and* sell decisions which is occurring in oscillatory fashion—in other words, *synchronized* oscillatory fluctuations emerge from an otherwise uncoordinated system. For this reason, the rationale that people are simply seeking to buy CHF cannot be right. Instead, there appears to be a rich process of collective change of mind.

Even though this period of CHF-based synchronization is not long, the magnitude of the returns has a large effect on the correlation—hence the drop seen in the clustering distance. To illustrate precisely how clustered the CHF-based exchange rates are at that time, Fig. 4 shows the hierarchy resulting from using the SLCA on a 120 time step window of data ending at 9:05 EST on 11 September 2001. The CHF cluster is picked out in bold, red lines [75]. It is evident that all the CHF based rates are clustered together. The clustering distance of 0.4532 corresponds to a correlation of 0.897.

B. CNY revaluation—rumor

The 11 September analysis above shows that the SLCA clustering distance can identify a currency dominating a group of others in the market, even for 1-min returns. In this section we consider a more market-specific shock. At 08:22 GMT, Wednesday 11 May 2005, a rumor began to propagate through the market that the Chinese government would imminently remove the peg between the Chinese currency (the Chinese Yuan Renminbi, CNY) and the USD [76,77]. It turned out that the rumor was false, but this was not known at the time. The credibility of the rumor, and hence its contagion period, survived for approximately 30 min. Toward the end of this period, it gradually became evident that the rumor was untrue.

This example is of particular interest for two reasons. First, the shock is specific to the FX market, yet external to the currencies included in our analysis. Hence it does not represent news which is unambiguously "good" or "bad" for a given currency—nor are the subsequent actions which a trader should take immediately obvious. Second, the real revaluation of the CNY happened only a couple of months later. Hence it is almost as if the same shock happened twice. Comparing the difference between the reaction of the market in both situations therefore enables us to investigate how the market processes such information. In other words, does the FX market collectively go ahead and repeat its earlier behavior or do something completely different? A physical system would, in the absence of hysteresis effects, respond in a similar way if subjected to similar conditions—but it is not obvious whether human-based complex systems behave in such a way. This study therefore provides a unique experiment in a real-world system.

Prior to the rumor, the market was actually expecting some form of official announcement about revaluing the CNY in the ensuing days, weeks, or months. In other words, the market was "susceptible" to such a rumor. The Chinese government had been under pressure for some time to revalue its currency, which other countries felt was fixed at an unrealistically low price. However, the knock-on effects of such a revaluation were not known beforehand. Some assumed that the announcement would have a large impact on other Asian currencies. Others argued that since the news had been expected for some time, all possible effects on the Asian currencies would have already been priced into the market.

Figure 5 shows the clustering levels for the day before, and the day of, this false rumor (10–11 May 2005). It is

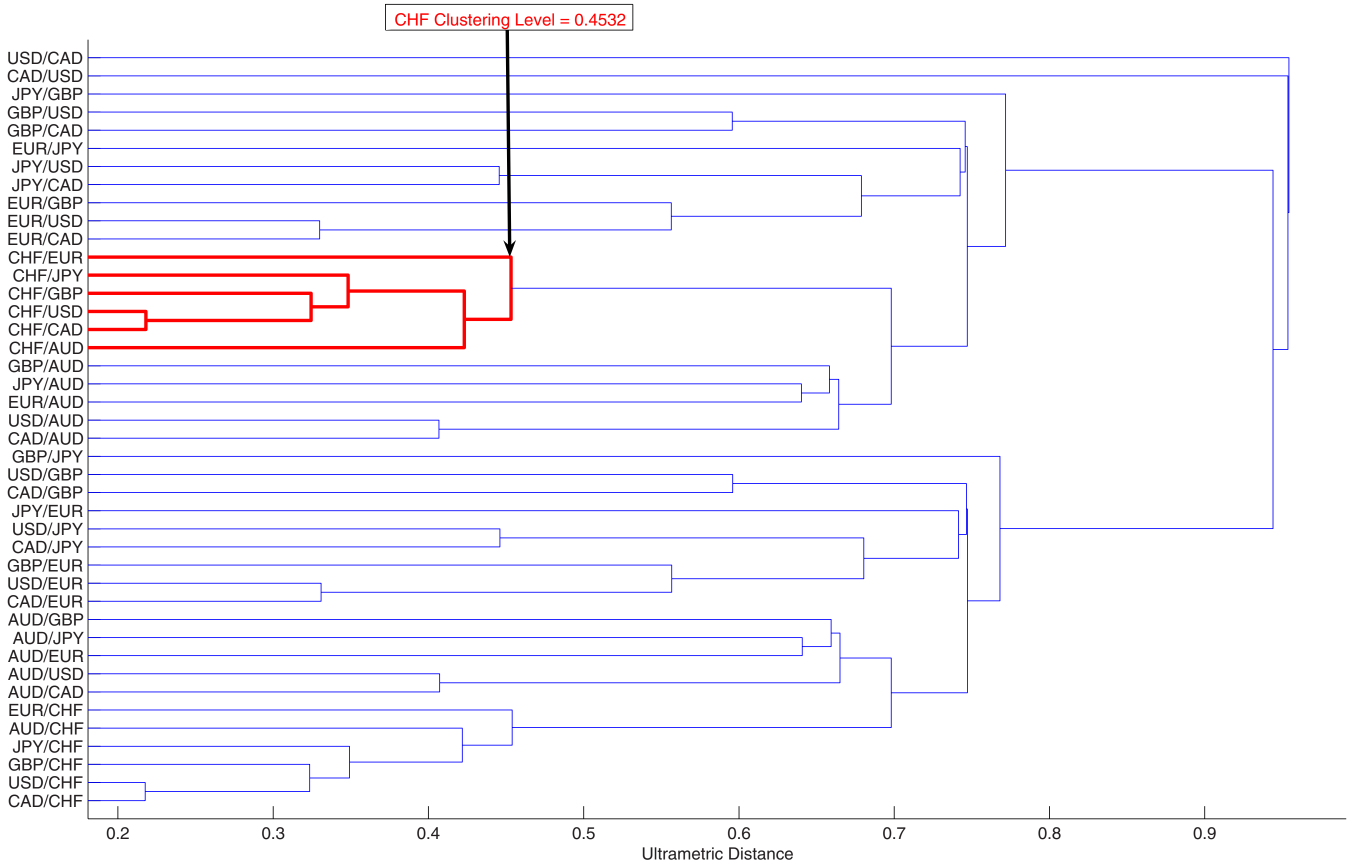

FIG. 4. (Color online) The hierarchy resulting from using the SLCA on a 120 time step window ending at 9:05 EST on September 11, 2001.

evident that there is a strong dip in the clustering distance for JPY. There are two important points to note. Since we look at 1-min data, the first data point which could be affected by news released in the minute of 08:22 is the return at 08:23. The JPY cluster level responds from the very first window which incorporates this. In other words, the market collectively converged to a consensus concerning which currency should be traded—despite the fact that there was no invisible hand or central controller to conduct the coordination process. Second, the observed clustering lasts for approximately 30 min longer than the width of one window, indicating that strong clustering is seen for the entire time that the rumor is

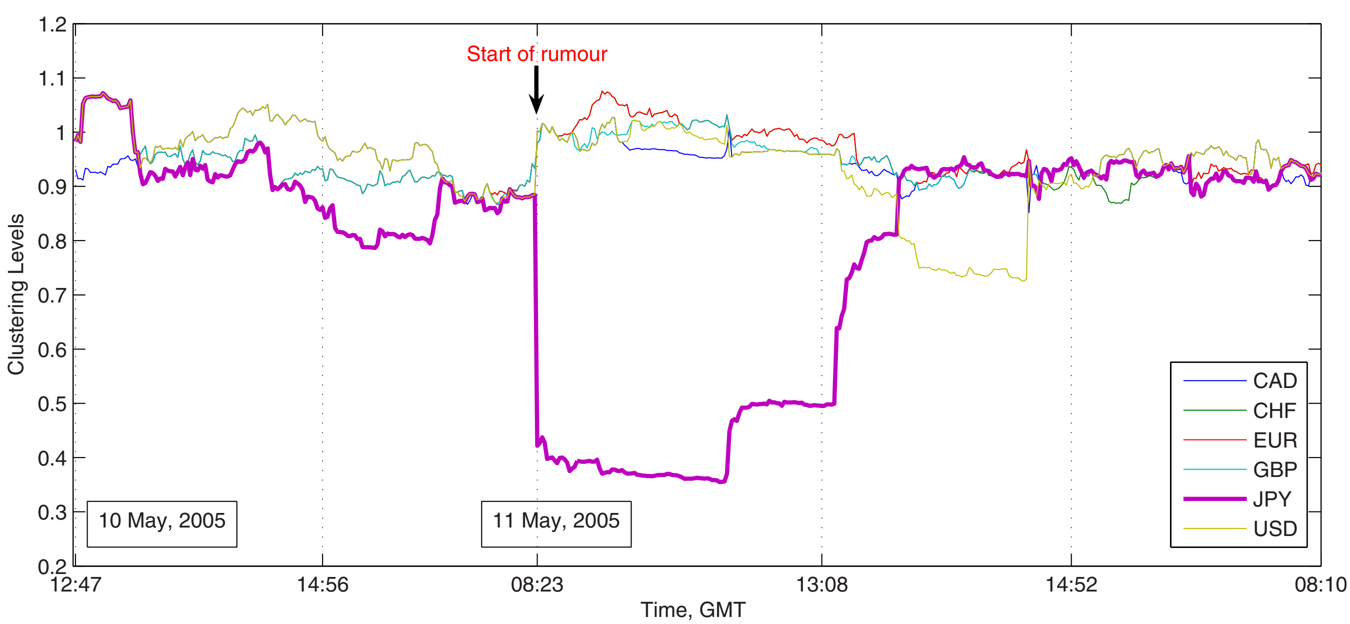

FIG. 5. (Color online) The clustering distances for the day before, and the day of, the false rumor concerning CNY revaluation.

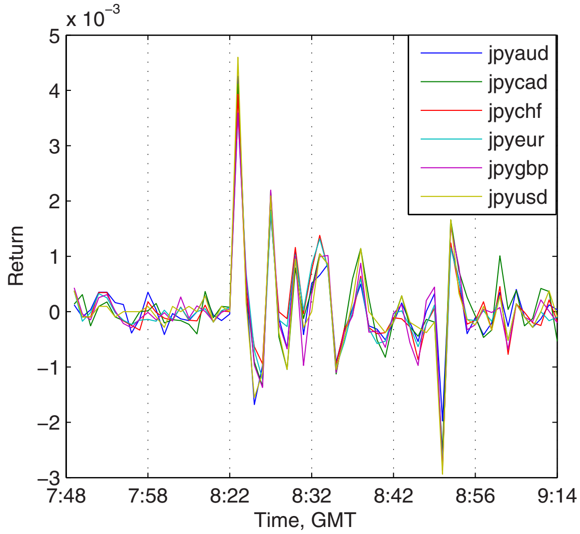

FIG. 6. (Color online) The JPY-based exchange rate returns for the time including the false CNY revaluation rumor. In the inset, jpyaud denotes the exchange rate between JPY and AUD, etc.

prevalent in the market. Note that for such a strong clustering effect to be seen as soon as the window includes the first time step from the rumor, the returns must be very large in order to induce such a rapid change in the correlation.

Figure 6 shows the returns for the JPY-based exchange rates for the period of time which includes the false CNY revaluation rumor. The degree of synchronization between the different rates once the rumor emerges is remarkably strong—even more so than for 11 September. The buying and selling of JPY is dominating any activity in the rest of the currencies included in the analysis.

C. CNY revaluation—actual event

Figure 7 shows the clustering distance for the day before, and the day of, the real CNY revaluation (20–21 July 2005). The similarities with Fig. 5 for the rumor are striking. The market appears to be noisier before the release of the news, which is probably a result of some failed tube bombings in London earlier that same morning. However, the clustering profile seen in the JPY is very similar in both the rumor and the real revaluation. Again, the clustering is seen to last for longer than one window length, indicating genuine clustering and not simply one large outlying point. Figure 5 shows that the clustering distance for JPY returns to its original level once the rumor becomes discredited, whereas Fig. 7 shows that the clustering distance for JPY remains lower for a long time afterwards. In other words, similar patterns are observed for the rumor and real cases, but the real case evolves to a modified state.

Figure 8 shows the JPY-based exchange rate returns for the time period which includes the real CNY revaluation. As with the false rumor, the extent to which the JPY-based exchange rates become synchronized once the news is released is striking.

D. Surprising economic news

Looking again at Fig. 5, and in particular the behavior later in the day of 11 May, there is a feature which at first seems puzzling. There is a dip in the USD cluster distance which lasts for a short time (it lies within the interval 13:08–14:52). This can be understood more easily by referring to Fig. 9, in which we show the clustering distance results from 11 May 2005 using the smallest possible window size ($T=N=42$). The time marked "X" is the last time for which the window includes the return at 08:23 GMT. The clustering in JPY continues beyond this time, implying that the JPY cluster is the result of a systematic clustering of JPY-based exchange rates. The time marked "Y" is the last time for which the window contains the result at 12:31 GMT (which is where the USD clustering starts). It can be seen that as soon

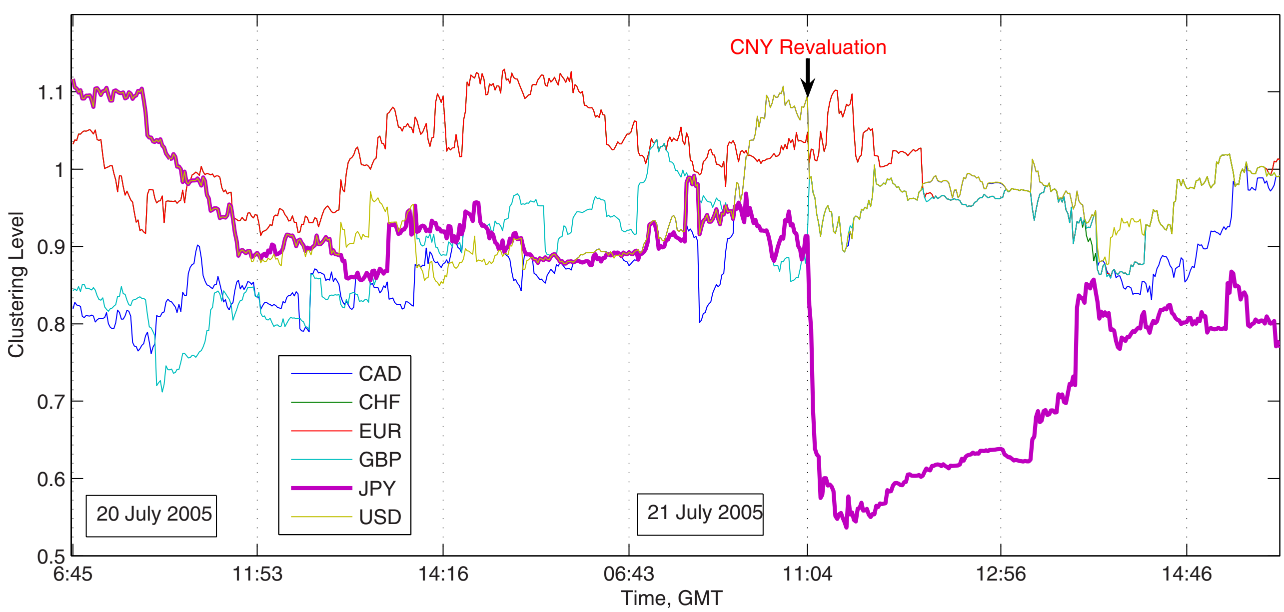

FIG. 7. (Color online) The clustering distances for the day before, and the day of, the real CNY revaluation.

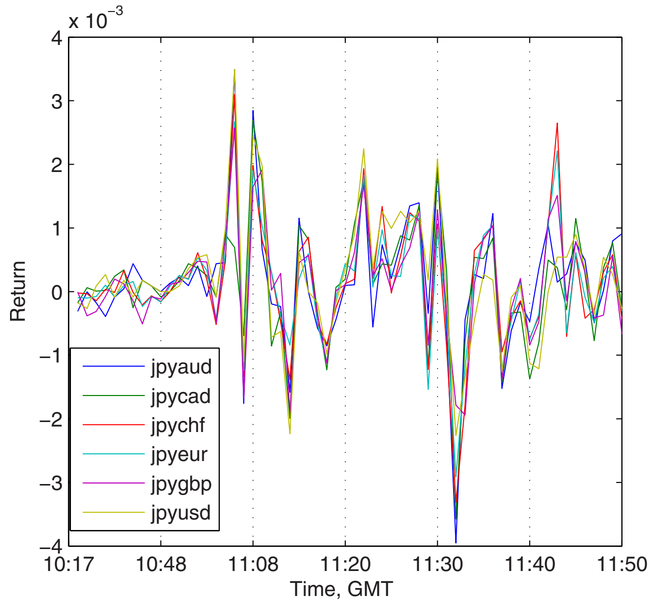

FIG. 8. (Color online) The JPY-based exchange rate returns for the time period which includes the real CNY revaluation.

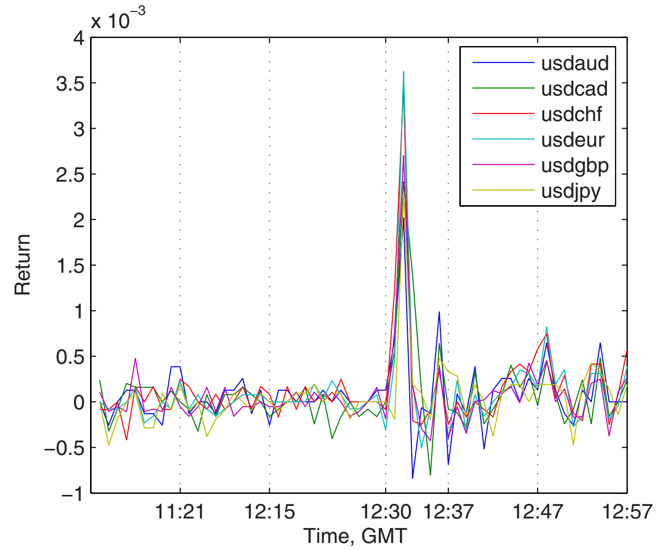

FIG. 10. (Color online) The returns for the USD-based exchange rates for a time period including the Trade Balance news release on 12 May 2005.

as the window does not include the time step at 12:31, there is no clustering. Hence the "clustering" of USD in this situation is simply caused by the behavior at one time step. But what is this behavior due to?

It turns out that at 12:30 GMT on 11 May 2005, the Census Bureau of the U.S. Department of Commerce made a very surprising announcement about the U.S. Trade Balance. This caused a very rapid change in the value of the USD against all other currencies. The returns for the USD-based exchange rates around the time of the announcement are shown in Fig. 10, showing that between 12:30 and 12:31 GMT, the value of the USD had moved against all the other currencies shown in the analysis [78]. The directional price change for all these exchange rates occurred in less than 1 min, whereas it appears that the volatility of the exchange rates is larger following the announcement. This is in agreement with papers researching the impact of economic announcements on financial markets [10].

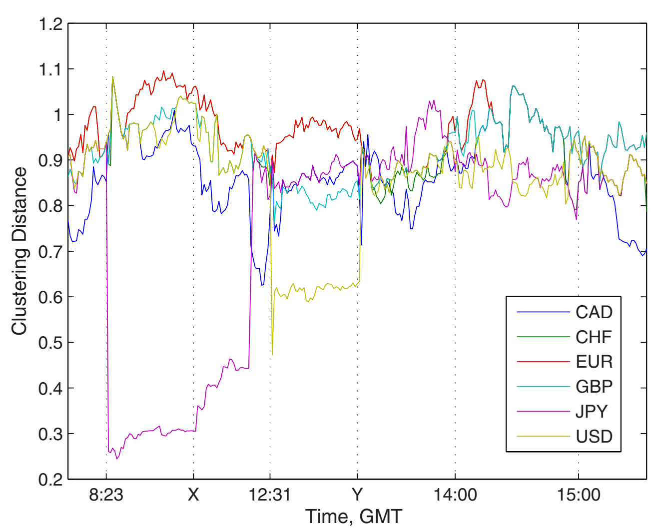

FIG. 9. (Color online) The clustering distances from using the shortest possible window size (42 returns) for the day of 11 May 2005.

V. ROBUSTNESS OF RESULTS

The results detailed above are all for the SLCA/MST clustering method. As we mentioned, there are potential problems with using this method if the data structure is not appropriate. In this section we repeat the analysis for three different clustering algorithms: ALCA, CLCA, and Ward. Note that the Ward algorithm is constrained to use the Euclidean distances between time series, and not the standardized Euclidean distance. Thus the distances seen in the Ward algorithm graphs are not the same distances as in the other graphs.

Figure 11 shows the results for 11 September. The top panel is the ALCA distance, the middle panel is the CLCA distance, and the bottom panel is the Ward distance. There are several points of interest in this figure. First, the initial clustering of EUR, which then becomes dominated by the CHF cluster, is evident in all three panels. Second, the CLCA distances (middle panel) show that prior to the attacks the cluster distance is consistently close to 2.0 for many currencies, for most of the time, corresponding to perfect anticorrelation and hence $\hat{\rho}_{ij}=-1.0$. In general this occurs for the returns of two inverse exchange rates (for example USD/CHF and CHF/USD). For the cluster distance for a currency to be 2.0 it means that by the time all the exchange rates with that currency as the base are in the same cluster, there are two inverse exchange rates in the cluster. It is obvious that

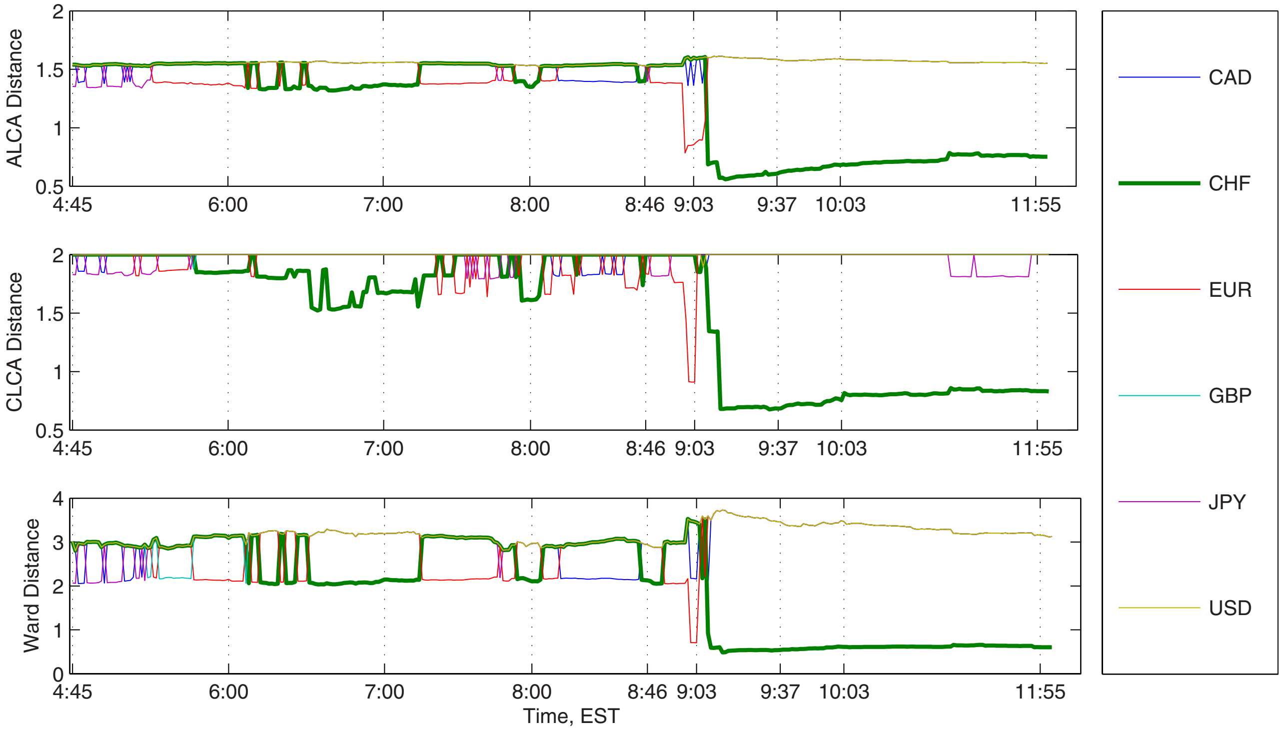

FIG. 11. (Color online) 11 September cluster distances using three alternative clustering algorithms.

one cannot consider two inverse exchange rates to move together, so this level of clustering must be interpreted as zero level of clustering. This allows us to distinguish real clustering from noise.

Figures 12 and 13 show the equivalent figures for the case of the false CNY revaluation rumor and the actual revaluation, respectively. Similar effects can be seen in these figures also. It can be seen that the clustering observed as a result of the news can be distinguished from the case of zero clustering. In addition, it can be seen that the clustering levels from the CLCA and Ward algorithms are quite noisy, but still show the same clustering for JPY.

VI. CONCLUSION

To summarize, we have shown that the clustering distance is a useful tool for investigating the collective human re-

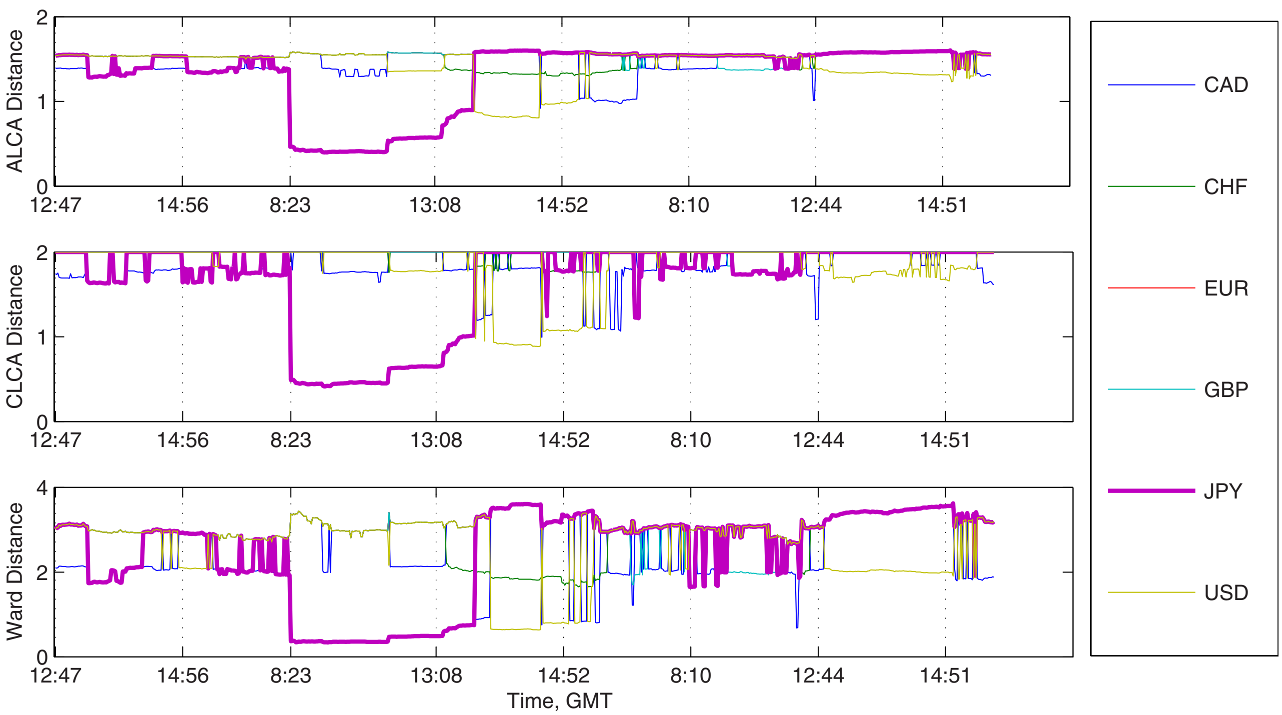

FIG. 12. (Color online) The clustering distances for the days of 10–11 May 2005, using three alternative clustering algorithms.

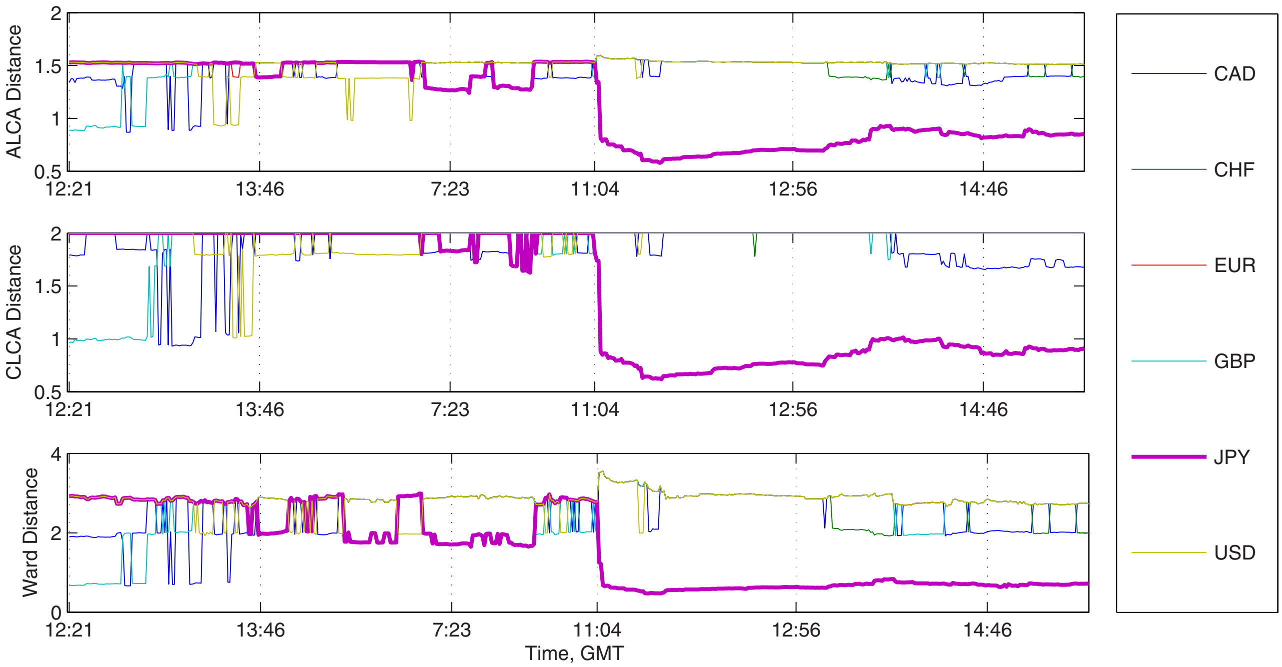

FIG. 13. (Color online) The clustering distances for the days of 20–21 May 2005, using three alternative clustering algorithms.

sponse to news in the FX market. In particular we have shown that there is a markedly different response seen when the news is genuinely unexpected, as opposed to a scheduled economic announcement that has a value far from what was expected by the market. We believe that this tool may eventually prove very useful for classifying news events.

We have also confirmed the robustness of the results from the MST method applied to the FX data. This is particularly useful since it justifies the use of the intuitive graphical representation provided by the MST. Indeed, the clustering distances from the MST are less noisy than for the other methods investigated. Section V illustrates the importance of comparing results with alternative clustering methods, so as to be able to distinguish real clustering from noise.

ACKNOWLEDGMENTS

We would like to thank Mark Austin and David Pavitt for helpful discussions.

[1] See http://tinyurl.com/3e4cbx for a BBC report of recent allegations of this practice in the U.K.
[2] D. Sornette, A. Johansen, and J.-P. Bouchaud, J. Phys. I **6**, 167 (1996).
[3] A. Johansen and D. Sornette, Physica A **276**, 338 (2000).
[4] D. Sornette, F. Deschatres, T. Gilbert, and Y. Ageon, Phys. Rev. Lett. **93**, 228701 (2004).
[5] F. Deschatres and D. Sornette, Phys. Rev. E **72**, 016112 (2005).
[6] B. M. Roehner, D. Sornette, and J. V. Andersen, Int. J. Mod. Phys. C **15**, 809 (2004).
[7] D. Sornette, Y. Malevergne, and J.-F. Muzy, in *Proceedings of the 2nd Nikkei Symposium on Econophysics*, edited by H. Takayasu (Springer-Verlag, New York, 2004).
[8] D. Sornette, in *Extreme Events in Nature and Society*, edited by S. Albeverio, V. Jentsch, and H. Kantz (Springer, Heidelberg, 2005), Chap. 5.
[9] W. Zhou and D. Sornette, Physica A **361**, 297 (2006).
[10] L. H. Ederington and J. H. Lee, J. Finance **48**, 1161 (1993).
[11] G. W. Schwert, J. Finance **36**, 15 (1981).
[12] D. K. Pearce and V. Roley, Technical Report No. 1296, National Bureau of Economic Research, Inc., 1985. Available at http://ideas.repec.org/p/nbr/nberwo/1296.html
[13] G. Hardouvelis, J. Econ. Bus. **39**, 131 (1987).
[14] C. R. Harvey and R. D. Huang, Rev. Financ. Stud. **4**, 543 (1991).
[15] G. McQueen and V. Roley, Rev. Financ. Stud. **6**, 683 (1993).
[16] See http://ideas.repec.org/a/fip/fedlrv/y1989imarp34-46.html
[17] T. Cook and S. Korn, Technical Report No. 5, Federal Reserve Bank of Richmond, 1991.
[18] U. Savanayana, T. Schneeweis, and J. Yau, *Recent Developments in International Banking and Finance* (Blackwell, Cambridge, England, 1991), Vol. 6.
[19] M. D. Evans and R. K. Lyons, Technical Report No. 9433, National Bureau of Economic Research, Inc., 2003.
[20] T. G. Andersen and T. Bollerslev, J. Finance **53**, 219 (1998).
[21] M. J. Fleming and E. M. Remolona, Technical Report 9706, Federal Reserve Bank of New York, 1997.
[22] P. Balduzzi, E. J. Elton, and T. C. Green, J. Financ. Quant. Anal. **36**, 523 (2001).
[23] N. Hautsch and D. Hess, Eur. Finance Rev. **6**, 133 (2002).

[24] T. G. Andersen, T. Bollerslev, F. X. Diebold, and C. Vega, Working Paper No. 11312, National Bureau of Economic Research, Inc., 2003. Available at http://www.nber.org/papers/w11312
[25] G. A. Karolyi and R. Martell, *Terrorism and the stock market*. Available at http://ssrn.com/abstract=823465
[26] A. Abadie and J. Gardeazabal, Am. Econ. Rev. **93**, 113 (2003).
[27] A. H. Chen and T. F. Siems, Eur. J. Polit. Econ. **20**, 349 (2004).
[28] D. A. Carter and B. J. Simkins, Working Paper No. OK 740784011, Oklahoma State University, Stillwater, 2002.
[29] M. T. Hon, J. Strauss, and S. K. Yong, J. Financ. Res. **27**, 95 (2004).
[30] J. M. Lacker, J. Monet. Econ. **51**, 935 (2004).
[31] R. B. Johnston and O. M. Nedekescu, Working Paper No. WP/05/06, International Monetary Fund, 2006. Available at www.imf.org/external/pubs/ft/wp/2005/wp0560.pdf
[32] J. R. Brown, D. J. Cummins, C. M. Lewis, and R. Wei, J. Monet. Econ. **51**, 861 (2004).
[33] D. J. Cummins and C. M. Lewis, J. Risk Uncertain. **26**, 153 (2003).
[34] J. Kallberg, C. Liu, and P. Pasquariello, working paper, University of Michigan, 2005. Available at http://ssrn.com/abstract=716721
[35] A. Mende, *Effect of 09/11 on the USD/EUR foreign exchange market*. Available at http://ideas.repec.org/p/han/dpaper/dp-312.html
[36] *How has September 11 influenced the global economy?* World economic outlook, International Monetary Fund, December 2001. Available at http://www.imf.org/external/pubs/ft/weo/2001/03/pdf/chapter2.pdf
[37] The last authoritative survey of the FX market was performed by the Bank for International Settlements in April 2007 [38]. In this survey the turnover of the global FX market was estimated to be 3.2 trillion U.S. dollars per day.
[38] *Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity*, Bank for International Settlements, September 2007. Available from http://www.bis.org/publ/rpfx07.pdf
[39] For a comprehensive quantitative overview of the world's foreign exchange markets, see M. M. Dacorogna, R. Gencay, U. Muller, R. B. Olsen, and O. V. Pictet, *An Introduction to High-frequency Finance* (Academic Press, San Diego, 2001), and references within.
[40] M. McDonald, O. Suleman, S. Williams, S. Howison, and N. F. Johnson, Phys. Rev. E **72**, 046106 (2005).
[41] R. N. Mantegna, Eur. Phys. J. B **11**, 193 (1999).
[42] J. B. Kruskal, Proc. Am. Math. Soc. **7**, 48 (1956).
[43] R. C. Prim, Bell Syst. Tech. J. **36**, 1389 (1957).
[44] R. N. Mantegna and H. E. Stanley, *An Introduction to Econophysics: Correlations and Complexity in Finance* (Cambridge University Press, Cambridge, England, 2000). See also the account of correlations in J. P. Bouchaud and M. Potters, *Theory of Financial Risk and Derivative Pricing* (Cambridge University Press, Cambridge, England, 2003), and also the thorough discussion of different measures of dependence given in Y. Malevergne and D. Sornette, *Extreme Financial Risks (from Dependence to Risk Management)* (Springer, Heidelberg, 2006).
[45] L. Kaufman and P. J. Rousseeuw, *Finding Groups in Data: An Introduction to Cluster Analysis* (Wiley-Interscience, New York, 1990).
[46] P. Hirst, M.S. thesis, Oxford Centre for Industrial and Applied Mathematics, Oxford University, 2003.
[47] P. H. Sneath and R. R. Sokal, *Numerical Taxonomy: The Principles and Practice of Numerical Classification* (Freeman, San Francisco, 1973).
[48] B. F. King, J. Bus. **39**, 139 (1966).
[49] G. Bonanno, G. Caldarelli, F. Lillo, and R. N. Mantegna, Phys. Rev. E **68**, 046130 (2003).
[50] G. Bonanno, G. Caldarelli, F. Lillo, S. Micciche, N. Vandewalle, and R. N. Mantegna, Eur. Phys. J. B **38**, 363 (2004).
[51] J.-P. Onnela, A. Chakraborti, K. Kaski, and J. Kertesz, Eur. Phys. J. B **30**, 285 (2002).
[52] S. Micciche, G. Bonanno, F. Lillo, and R. N. Mantegna, Physica A **324**, 66 (2003).
[53] J.-P. Onnela, A. Chakraborti, K. Kaski, and J. Kertesz, Physica A **324**, 247 (2003).
[54] J.-P. Onnela, A. Chakraborti, K. Kaski, and J. Kertesz, Phys. Scr., T **T106**, 48 (2003).
[55] J.-P. Onnela, A. Chakraborti, K. Kaski, J. Kertesz, and A. Kanto, Phys. Rev. E **68**, 056110 (2003).
[56] P. Ormerod and C. Mounfield, Physica A **280**, 497 (2000).
[57] G. Bonanno, N. Vandewalle, and R. N. Mantegna, Phys. Rev. E **62**, R7615 (2000).
[58] G. Bonanno, F. Lillo, and R. N. Mantegna, Quant. Finance **1**, 96 (2001).
[59] T. Di Matteo, T. Aste, and R. N. Mantegna, Physica A **339**, 181 (2004).
[60] V. Tola, F. Lillo, M. Gallegati, and R. N. Mantegna, e-print arXiv:physics/0507006; See also the following more recent papers by R. N. Mantegna and co-workers: M. Tumminello, T. Di Matteo, T. Aste, and R. N. Mantegna, Eur. Phys. J. B **55**, 209 (2007); M. Tumminello; C. Coronnello, F. Lillo, S. Micciche; R. N. Mantegna, Int. J. Bifurcation Chaos Appl. Sci. Eng. **17**, 2319 (2007).
[61] M. Bernaschi, L. Grilli, and D. Vergni, Physica A **308**, 381 (2001).
[62] M. Tumminello, T. Aste, T. Di Matteo, and R. N. Mantegna, Proc. Natl. Acad. Sci. U.S.A. **102**, 10421 (2005).
[63] J. C. Gower and J. S. Ross, J. Appl. Stat. **18**, 54 (1969).
[64] K. Florek, J. Lukaszewicz, J. Perkal, H. Steinhaus, and S. Zubrzycki, Colloq. Math. **2**, 282 (1951).
[65] O. Boruvka, Moravske Priodovedecke Spolecnosti **3**, 37 (1926).
[66] J. H. Ward, J. Am. Stat. Assoc. **58**, 236 (1963).
[67] R. Rammal, G. Toulouse, and M. A. Virasoro, Rev. Mod. Phys. **58**, 765 (1986).
[68] These are sometimes referred to as cophenetic distances.
[69] These currency codes are standard in the financial industry. They have been established by the International Organization for Standardization (ISO) and are known as ISO 4217 [70].
[70] See http://www.iso.org/iso/en/prods-services/popstds/currencycodeslist.html
[71] C. Alexander, *Market Models: A Guide to Financial Data Analysis* (John Wiley, London, 2001).
[72] C. W. J. Granger and P. Newbold, J. Econ. **2**, 111 (1974).
[73] D. A. Dickey and W. A. Fuller, J. Am. Stat. Assoc. **74**, 427 (1979).
[74] J. G. MacKinnon, in *Long-Run Economic Relationships: Read-*

ings in Cointegration, edited by R. F. Engle and C. W. J. Granger (Oxford University Press, Oxford, 1991), Chap. 13.

[75] Upon closer inspection, one can see the equivalent cluster of *X*/CHF rates in the hierarchy.

[76] See http://www.bloomberg.com/apps/news?pid=10000080&sid=al3eF8ebU3qw&refer=asia, for a news story detailing this.

[77] A pegged exchange rate (also known as a fixed exchange rate) is one where the value of a currency is matched to the value of either (i) another currency, (ii) a basket of other currencies, or (iii) another measure of value (for example, the price of gold). As the value of the reference changes, so too does the value of pegged currency.

[78] The exchange rate USD/*X* is the number of units of *X* that one USD can buy. Hence if USD/*X* decreases, USD is less valuable with respect to currency *X*.