

Quantification of the High Level of Endogeneity and of Structural Regime Shifts in Commodity Markets

Vladimir Filimonov^a, David Bicchetti^b, Nicolas Maystre^{b,c}, Didier Sornette^{a,d}

^a*Dept. of Management, Technology and Economics, ETH Zürich, Zürich, Switzerland*

^b*United Nations Conference on Trade and Development, Division on Globalization and Development Strategies, Palais des Nations, 1211 Geneva 10, Switzerland*

^c*Dept. of Economics, University of Geneva*

^d*Swiss Finance Institute, c/o University of Geneva*

Abstract

We propose a novel index of short-term endogeneity (or reflexivity) derived by calibrating the Hawkes self-excited conditional Poisson model on empirical time series of trades. The Hawkes model accounts simultaneously for the co-existence and interplay between the exogenous impact of news and the endogenous mechanism by which past trading activity may influence future trading activity. Technically known in the mathematical literature on branching processes as the branching ratio, the reflexivity index is quantified for several commodity futures markets (corn, oil, soybean, sugar, and wheat) and also for a benchmark equity futures market (E-mini S&P 500). Specifically, the reflexivity index is the average ratio of the number of price moves that are due to endogenous interactions to the total number of all price changes, which also include exogenous events. We find an overall increase of the level of short-term endogeneity since the mid-2000s to October 2012, with a typical value nowadays around 0.6–0.7, implying that at least 60–70 per cent of commodity price changes are now due to self-generated activities rather than novel information. Our robustness tests show that the branching ratio provides a ‘pure’ measure of endogeneity that is independent of the rate of activity, order size, volume or volatility.

Email addresses: vfilimonov@ethz.ch (Vladimir Filimonov),
david.bicchetti@unctad.org (David Bicchetti), nicolas.maystre@unctad.org
(Nicolas Maystre), dsornette@ethz.ch (Didier Sornette)

Preprint submitted to The Journal of International Money and Finance (JIMF) March 21, 2013

We complement our analysis by relating the endogeneity dynamics of these futures markets to their price dynamics, particularly around the commodity bubble that developed since 2006 and culminated in mid-2008. While our index does not have a long-term memory, interestingly, we find that it can still provide some interesting insights when the mechanisms working at longer time scales cascade down to shorter terms.

Keywords: Commodities, endogeneity, reflexivity, branching processes, bubble, oil, regime shift, self-excitation

Disclaimer

The opinions expressed in this paper, including designation and terminology, are those of the authors and are not to be taken as the official views of the UNCTAD Secretariat or its Member States.

Highlights

- We compute the fraction of endogenous trades on highly-traded futures markets.
- Similar to equity indices, endogeneity on commodity has grown significantly during the last decade.
- Endogeneity averages 0.6–0.7 in 2012: most price changes are not due to informative news.
- Our index is independent of the rate of activity, order size, volume or volatility.

1. Introduction

Many commodity prices have experienced roller-coaster rides since the mid-2000s. These price gyrations have fueled an intense debate among academics, commodity traders and policymakers. In particular, the role of financial investors has been the subject of considerable controversies. Disagreements relate to whether these new actors have improved the price discovery process of commodities futures markets or whether they have made the process less effective and more unstable.

For the proponents of the so-called financialization of commodity markets, its benefits are at least threefold. First, it brings the futures prices of these products closer to their underlying fundamentals. Second, it provides liquidity. Third, it transfers risks to agents who are better prepared to assume it (see e.g. Stoll and Whaley (2010, 2011); Irwin and Sanders (2012) and references cited therein). In short, this process supports the efficient market hypothesis (EMH) (Samuelson, 1965; Fama, 1970, 1991). By contrast, other observers argue that financial investors can have negative effects on commodity markets (see e.g. UNCTAD (2009, 2011); Tang and Xiong (2010); Bicchetti and Maystre (2012) and references cited therein).

To contribute to this debate, we analyze the microstructure of several commodity futures markets at short time scales and provide quantitative dynamic estimates of their significant degree of reflexivity. This provides a strong counter-example to EMH, which predicts in its ideal limit that the market absorbs in full and essentially instantaneously the flow of information by faithfully reflecting it in asset prices. A corollary to this strong version of EMH states that price variations can only result from exogenous events that feed instantaneously the price determination process. In contrast, our findings show that past price changes can trigger subsequent price variations, as described qualitatively by Soros' concept of "market reflexivity" (Soros (1987)).

In this paper, we build on Filimonov and Sornette (2012) and we propose estimates about the degree of short-term endogeneity (or reflexivity) of commodity futures markets derived from the Hawkes self-excited conditional Poisson model. The Hawkes model combines in a natural and parsimonious way exogenous influences with self-excited dynamics. Indeed, it accounts simultaneously for the co-existence and interplay between the exogenous impact of news and the endogenous mechanism of trading activity where one price change triggers subsequent price changes. Thus, the Hawkes model

allows quantifying the ratio of price changes on commodity futures markets that are due to endogenous feedbacks, as opposed to exogenous news. Our research draws on concepts originally developed in the study of earthquake aftershocks (see e.g. Vere-Jones (1970); Vere-Jones and Ozaki (1982) and Ogata (1988)), which were first applied by Bowsher (2002) (published later with corrections in Bowsher (2007)) in the area of high-frequency financial data. According to the Hawkes model, each event (i.e. price change) may lead to a whole tree of offspring (i.e. subsequent price changes).

We calibrate the Hawkes model using Thomson Reuters Tick History (TRTH) data on various front month commodity futures contracts, including corn, oil, soybean, sugar, and wheat, as well as on the most-traded futures contract, the E-mini S&P 500, which we use as a benchmark. Each calibration amount to estimating the parameter n , which in case of $n < 1$ equals to the ratio of price moves that are due to short-term endogenous interactions to the total number of all price changes, which also include the impact of exogenous events. In the mathematical literature of branching processes (Harris, 2002), this parameter n is usually called the “branching ratio”. We shall refer to n using both terms “reflexivity index” (which emphasizes its conceptual meaning) and “branching ratio” (which emphasizes its technical meaning).

The reflexivity index n provides a simple and illuminating characterization of markets, in particular with respect to their fragility and susceptibility to shocks. For $n < 1$, on average, the fraction $1 - n$ of price changes are due to exogenous news or surprises while the fraction n of price changes are endogenous, i.e., can be traced back to the influence of past price changes. As n approaches 1 from below, the system becomes “critical”, in the sense that its activity is mostly endogenous or self-fulfilling. More precisely, its activity becomes hyperbolically sensitive to external influences. The regime $n > 1$ corresponds to an unbounded explosion of activity nucleated by just a few external news and can only realistically occur over a finite time.

For the commodities that we have analyzed, we document the evolution of the degree of short-term endogeneity for all these markets since the second half of the 2000s, when the considered commodity exchanges moved from pit trading to full electronic platforms. Overall, we usually find average levels of short-term endogeneity above 50 per cent for all considered commodity markets with episodes well above 85 per cent. This highlights the failures of the EMH and provides evidence that price dynamics are partly driven by positive feedback mechanisms. We also discuss why higher level of endo-

geneity makes the price formation process less efficient and more prone to instability. Moreover, we show that endogeneity has increased, albeit not necessarily monotonously, over the considered periods. For instance, when trading on the oil Brent futures moved to full electronic trading in 2005, about 35 per cent of the Brent prices changes resulted from previous prices changes. In late 2008 to early 2009, this figure has increased above 75 per cent. Afterwards, it has lost some momentum but has stabilized around 60 per cent since the early 2011.

We complement this analysis on the level of reflexivity by relating the endogeneity dynamics of oil futures markets to their price dynamics, particularly around the commodity bubble that developed since 2006 and culminated in mid-2008. At first sight, our reflexivity index is not particularly well designed to capture longer-term herding mechanisms, which are responsible for bubble formation on time scales of months to years. This is because the model has been calibrated in running windows of 10 minutes, which prevents the model from capturing long term dynamics. However, we surprisingly find that the mechanisms working at longer time scales sometimes seem to cascade down to the shorter intervals and are thus detectable by our analysis. Finally, in presenting our estimates, we show that the branching ratio sometimes also exhibits abnormal increases that are concomitant with significant price swings and/or bubble developments.

Regarding the US equity futures contracts, we compute the branching ratio since late 1997. The purposes of this digression to the equity markets are twofold. First, it aims at providing a benchmark of the evolution of the reflexivity on the most traded financial derivative product since the late 1990s. In fact, the TRTH data does not allow us to compute the branching ratios prior to the introduction of full electronic trading in the mid-2000s on the considered commodity exchanges. By considering the US equity futures market, we find evidence for the growing waves of endogeneity affecting financial markets. Providing empirical evidence over a longer time scale matters because it shows that the short-term endogeneity level of highly financialized products had already increased in the early 2000s. In all likelihood, the levels of endogeneity in commodity markets in the late 1990s were smaller than in the E-mini S&P 500 futures market. This makes us believe that the already-high endogeneity level observed in commodity markets mostly in the second-half of 2000s has not been a permanent feature in the decade that preceded the availability of reliable tick data on commodity derivatives.

The rest of our paper is organized as follows. In section 2, we first present

our data and then discuss some recent technological changes that appeared on exchanges, such as the increasing influence of algorithmic trading, in particular high frequency trading (HFT) activities. Then, we introduce the Hawkes self-excited Poisson model and explain why it is well-adapted to the determination of the endogeneity level coming from discontinuous financial data at high frequencies. In section 3, we provide our estimates of the reflexivity index. We describe first the calibration process of the Hawkes model to high frequency data. Then, we present an updated analysis of the evolution of the reflexivity index for the E-mini S&P 500 futures, which was initially performed in (Filimonov and Sornette, 2012). Subsequently, based on our empirical estimates, we confirm both the relevance of the Hawkes model as an excellent data descriptor and the robustness of our estimation procedure. In the rest of this section, we present our branching ratio estimates for several highly traded commodity futures. Then, we discuss the monthly evolutions of our indices and how these relate to some key events that have influenced their development since the mid-2000s. In section 4 we conclude.

2. Data and methodology

2.1. Data

2.1.1. Nature and characteristics of the studied data set

We base our analysis on Thomson Reuters Tick History (TRTH) data. TRTH provides financial data for an extensive range of asset classes with more than 45 million unique instruments across more than 400 exchanges, based on the information transmitted by exchanges and market makers. TRTH contains historical data back to January 1996 at best. It provides granular tick as well as lower frequency pricing data, up to the microsecond level. Moreover, TRTH offers intra-day time sales or quotes, and market depth data. The database provides also over-the-counter (OTC) quotes. To our knowledge, it offers the most comprehensive pricing and reference data service, with a record of market behavior of 2 petabytes ($2 \cdot 10^{15}$ bytes) of tick data.

In this study, we limit ourselves to a few instruments. We select some of the most liquid commodity derivatives, namely futures on Brent crude oil (ICE – Europe), WTI crude oil (NYMEX), corn (CBOT), soybeans (CBOT), sugar #11 (ICE - US), wheat (CBOT) and white sugar (LIFFE). These commodity futures contracts represent the commonly used benchmarks for the world or their respective markets.

Table 1 summarizes the main characteristics of each future contract. Each derivative contract has an underlying physical asset described in the “Specification” column and reaches maturity on specific dates, which we refer to as “Contract month”. Several futures contracts referring to the same underlying asset are traded in parallel during the trading sessions but are differentiated by their maturity dates. The front months for each future contracts usually have the greatest liquidity. For each considered commodity, TRTH provides a so-called continuous futures contract by taking the front month and switching to the next contract at expiration date.

The continuous Front month futures contract (which has suffix “c1” in the TRTH notation) usually exhibit greater trading volumes than other futures contracts with subsequent maturities. However, as the *expiration date* approach, this is no longer true as financial investors switch to the next maturity contract to avoid delivery. The peak of these rollover processes has traditionally preceded the expiration date by about one week. For instance, E-mini S&P 500 futures contracts are traded on a quarterly basis, expire on the third Friday of March, June, September and December. However, in our observations, the number of trade on the Front month futures contract become smaller than the ones on the next maturity contract eight days before the expiration dates, i.e. on the second Thursday of each of these months. Hence, the liquidity (measured in volume) of the expiring contract is switched from the expiring contract to the next quarter maturity at these *rollover dates*. For Brent and WTI futures, monthly settlements results in rollover dates closer to the expiration dates as most traders typically roll their positions two days before the expiration. In order to be consistent in our analysis of different assets, we have excluded periods between the rollover and the expiration dates from our analysis. For the Corn, Wheat and Soybean futures contracts, we could not clearly identify rollover dates because, in some instances, the Second month contracts are more heavily traded than the Front month. For these contracts we have excluded five trading days before the expiration.

As can be seen from table 1, different exchanges moved the trading activities from pit trading to full electronic platforms at different times. Brent crude oil, which was originally traded on the open outcry International Petroleum Exchange (IPE) in London, was the first oil contract that fully switched in 2005 to the electronic platform of Intercontinental Exchange (ICE) based in London. However, the white sugar traded in Europe at LIFFE moved already in 2000 on a full electronic platform. As discussed below, we

base our analysis on the so-called mid-quote price, which averages the best bid price and the best ask price. Due to the specifics of the open outcry pits, which were mostly driven by designated (official) market makers, the quotes on the pre-electronic exchanges could not serve as a reliable source of information because time was not registered precisely. For this reason, in our analysis, we have considered only time periods starting from the date of the complete switch to full electronic trading for each contract (indicated in table 1). For Sugar #11, we have excluded the period prior to March 2, 2008 even though electronic trading started on January 12, 2007 at ICE US. In fact, pit trading had continued to exist in parallel at NYBOT and, unfortunately, because of NYBOT feed limitations, the timestamps of quotes as well as volumes and settlement values are not entirely reliable before the decommissioning of NYBOT pit trading on February 28, 2008.

Table 2 summarizes the number of annual transactions and volumes for each considered derivative contract¹. The emergence of full electronic trading in the course of the 2000s on the considered commodity exchanges marks the beginning of an increase in the amount of transactions. At the same time, the dynamics of volume (presented at the monthly scale also in fig. 1) exhibits very moderate growth. As a result, the average volume per transaction (A-VPT) decreased significantly between 2005 and 2012 for all considered contracts: starting at an average of 5–40 contracts per transaction in 2005, this figure declined below 3 contracts per transaction in 2012 for all considered commodities. The more striking dynamics are observed on the median (M-VPT) and other quantiles of the distribution of the volume per transaction. No later than 2009, M-VPT settles at 1 for all the considered commodities, which means that at least 50 per cent of all the transactions involve only one contract. Similarly, the 90%-quantile of volumes per transaction (Q_{90} -VPT) have remained equal or below 7 contracts per transaction from 2009 onwards for commodities. The VPT decline primarily reflects the increase of HFT on commodity futures markets, whose typical strategies imply ultra-fast market-making with only a few contracts per limit order.

The beginning of HFT can probably be traced back to 1998 when the U.S. Securities and Exchange Commission (SEC) authorized electronic exchanges.

¹Due to feed limitations, TRTH does not contain reliable information of trading volumes for Soybean and Sugar #11 contracts prior introduction of electronic trading. Therefore for these contracts we presented data starting 2006 and 2007 respectively.

However, in the early 2000s, HFT, defined as the high speed component of algorithmic trading, was quite rare and accounted for less than 10 per cent of all equity orders. In subsequent years, its importance grew rapidly (Duhigg, 2009). In 2009, the proportion of high frequency trading in US markets was estimated as more than 60 per cent by the TABB Group (Sussman et al., 2009) and the Aite Group (2009), with less conservative early estimates of the TABB Group of the order of 73% (see Iati (2009)).

Although reliable estimates of algo trading activities on commodity markets are not systematically available, Reuters quoted the chief executive officer of the CME Group, saying that 45 per cent of volume exchanged on the NYMEX—a commodity futures exchange owned and operated by its group—was computer driven (Sheppard, March 3, 2011). In light of the figures presented in Table 2, this probably represents a conservative estimate. The one-time released report by the CFTC on “Large Trader Net Position Change” reveals the dominating role of day traders in volatile commodity markets (CFTC, 2011). In some instance, like WTI crude oil, almost 95 percent of trading volume is generated by day trading², which suggest that long-term bets have little effect on the commodity volatility (Meyer, July 5, 2011).

2.1.2. *Mid-quote price as informative proxy*

The choice of a *proxy* for the price movements at high frequency (minute, second and sub-second time scales) matters and depends on the particular application. At any given moment t , one may distinguish three different prices: (i) *the last transaction price* $p_{tr}(t)$, at which the previous transaction was executed, (ii) *the best ask price* $a(t)$ and (iii) *the best bid price* $b(t)$ at which market participants may immediately correspondingly buy and sell an asset. Best bid and best ask prices are usually aggregated in the so-called *mid-quote price*, which averages the two: $p_m(t) = (a(t) + b(t))/2$ (see fig. 2). The bid and ask prices reflect demand and supply of the liquidity providers, respectively. The transaction price reflects actions of liquidity takers. And mid-quote price changes result from actions of all market participants, both liquidity providers and takers. The transactions are triggered when a market order arrives. In case of a buy market order, the transaction is executed at

²defined as trades in and out of the market that are performed within a given day and whose positions do not roll over to the next day or longer

the best ask price, while a sell market order triggers a transaction at the best bid price. Since the sequence of order arrivals is stochastic with the sign of order being a random variable, the last transaction price will jump from best bid to best ask price and back even without changes in the balance between supply and demand. This stochastic behavior, which is called “bid-ask bounce”, represents a kind of “noise source” to the price.

The idea that the last transaction price in high frequency financial data is a poor proxy of the unobservable asset’s value, because it is subjected to the additive “microstructure noise”, is a well established concept in the market microstructure literature (see for instance, Ait-Sahalia et al. (2005), and the concept of “noise traders” by F. Black (1986)). In contrast to the last transaction price, the mid-quote price is free from the bid-ask bounce and changes only when the balance between supply (liquidity providers) and demand (liquidity takers) is modified. Therefore, the mid-quote price can be argued to be a better proxy for the asset value, given the available information (Hasbrouck, 1991; Engle, 2000). In the “price impact” literature dedicated to the question of the price response to an execution of a single (or series of) market order(s), the mid-quote price has become the “default measure” to monitor price movements (see, for instance, the extensive review in Bouchaud et al. (2009)). In the present study, we stick to the mid-quote price as the best proxy for market movements as a whole. However, we disregard the direction of the price movements, considering the so-called *point process* of the timestamp of events — i.e. mid-quote price changes — as represented by red squares in Fig. 2.

2.2. The Hawkes self-excited model and endogeneous mechanisms of price formation

2.2.1. Definition of the self-excited point process (Hawkes) model

The typical null hypothesis in modeling point processes³ is the so-called *Poisson process* in which events occur independently of one another with a constant average arrival rate λ . Having no correlation structure, the Poisson point process cannot describe the wide range of empirical stylized facts of real order flows, such as (i) clustering of order arrivals, (ii) long memory in

³Without going into precise mathematical definitions, point processes are special types of random processes, for which the realization consists of isolated events and the modeled variable is the timestamp (and coordinate as well as marks if applicable) of each event.

inter-trade intervals (Ivanov et al., 2004; Jiang et al., 2009), (iii) slower-than-exponential decay of the distribution of inter-trade intervals (Ivanov et al., 2004; Eisler and Kertész, 2006; Politi and Scalas, 2008), (iv) long memory of the signs of successive trades (Bouchaud et al., 2009), and (v) multi-fractal scaling of inter-trade intervals (Jiang et al., 2009; Oswiecimka et al., 2005; Perelló et al., 2008). Traditionally, two large classes of self-excited point processes have been used to account at least partially for these stylized facts that are characteristic of high frequency price data. The first one is the so-called Autoregressive Conditional Durations (ACD) model (Engle and Russell, 1997, 1998) and its extensions, which describes the inter-event durations with a GARCH-type equation. The second one is more parsimonious and flexible and is called the self-excited Hawkes model (Hawkes, 1971), which was first applied to high frequency financial data in the working paper (Bowsher, 2002) (published later with corrections as (Bowsher, 2007)). Nowadays, the Hawkes point process has become the “gold standard” of self-excited models to describe discontinuous financial data. It has a wide range of applications going from modeling high frequency order flows (Hewlett, 2006; Bauwens and Hautsch, 2009) and the construction process of the order book (Large, 2007; Toke, 2011; Cont, 2011), to modeling extreme events clustering at daily and hourly scales (Embrechts et al., 2011), estimating Value-at-Risk (Chavez-Demoulin et al., 2005) or modeling correlated default times in a portfolio of firms (Errais et al., 2010; Azizpour et al., 2011).

The Hawkes point process can be regarded as the generalization of the non-homogeneous Poisson process, whose intensity $\lambda(t)$ (defined such that $\lambda(t)dt$ is the expected value of the number of events in the time interval $[t, t + dt)$) not only depends on time t but also on the history of the process. Within the Hawkes model, the intensity of a process is conditional on history and has the form

$$\lambda_t(t) = \mu(t) + \sum_{t_i < t} h(t - t_i), \quad (1)$$

where t_i are the timestamps of the events of the process, $\mu(t)$ is a *background intensity* that accounts for *exogenous* events (not dependent on history) and $h(t)$ is a *memory kernel function* that weights how much past events influence the generation of future events and thus controls the amplitude of the *endogenous* feedback mechanism. Traditionally, the memory kernel is assumed

to be exponentially decaying in time:

$$h(t) = \alpha \exp\left(-\frac{t}{\tau}\right) \quad (2)$$

It is parametrized by two variables $\alpha > 0$ and $\tau > 0$. Below, we will validate the choice of the kernel with a goodness-of-fit analysis. Apart from the good agreement with the data, the choice of a short-memory exponential kernel reflects the main target of our analysis, namely — the impact of the short-term speculative mechanisms of reflexivity, that are operated on the scales of minutes and less (see section section 3.1.2 for the discussions).

2.2.2. Branching ratio and level of endogeneity/reflexivity

For our purposes, the Hawkes model presents two interesting properties. First, the model clearly isolates the external influences on the system, $\mu(t)$, from the internal feedback mechanisms, $h(t)$, in the conditional intensity $\lambda_t(t)$. Second, the linear structure $\lambda_t(t)$ of the Hawkes model allows one to map it exactly onto a so-called *branching process* (Daley and Vere-Jones, 2008). This mapping introduces a key parameter called the *branching ratio* n , which we define more precisely below. As described in (Filimonov and Sornette, 2012), the branching ratio n quantifies the degree of self-excitation (or reflexivity) occurring in the system.

In the language of branching processes (Harris, 2002), all events belong to one of two classes: *immigrants* (zero-order events) and *descendants* (first-, second-, and higher order events). The exogenous immigration — described by the background intensity $\mu(t)$ — triggers clusters of descendants⁴. Namely, every zero-order event (mother) can trigger one or more first-order events (daughters), each of whom becoming mother-event in turn can trigger several daughters (second-order events or grand-daughters) and so on over many generations (see Fig. 3). All first-, second-, and higher order events form the cluster of aftershocks of the main event as a result of the self-excited (endogenous) generating mechanism of the system.

Applying the Hawkes process to interpret high frequency price dynamics, each event can be either exogenous or endogenous to the system. In the first case, its external origin could be interpreted as due to idiosyncratic news,

⁴In other words, $\mu(t)$ is the frequency at which exogenous events impact the system, and the share of exogenous events are measured by $1 - n$, described later in the paper.

which are not anticipated by investors and surprise them, thus providing genuine new pieces of information forcing them to reassess their investments and react. We refer to these events as “*fundamental*” events and, in the language of branching processes, as zero-order events (see Fig. 3). In contrast, endogenous events are triggered by preceding price changes as the result of internal feedback mechanisms. The existence of triggered events embodies the mechanism of self-excitation of the system onto itself, i.e., the influence of past price changes on future price changes. Self-excitation and endogeneity are related to the concept of *reflexivity*, which has extensive roots in philosophy and sociology, and has more recently been advocated by Soros (1987) to provide a useful framework to understand financial markets and beyond.

We can propose the following non-exhaustive list of mechanisms that can trigger endogeneity/reflexivity.

- *Technical analysis, including algorithmic and HFT trading*: market participants send orders that are not based on changes in economic fundamentals but on technical analysis of price and volumes movements.
- *Behavioral mechanisms and herding*: changes in fundamentals triggers an avalanche of new orders based on momentum or market sentiments that cause prices to over/undershoot.
- *Optimal portfolio execution*: in order to minimize market impact when buying/selling large numbers of shares, orders are split into smaller orders that are traded incrementally;
- *Hedging strategies* increase cross-excitation and cross-correlation between markets. When combined with portfolio execution issues, hedging strategies amplify self-excitation, as for instance in insurance portfolio (Kyle and Obizhaeva, 2012).
- *Margin and leveraged trading* occurs when price changes above or below a certain limit against the initial position triggers a *margin call*, which, if not addressed, results into the automatic liquidation of the leveraged position, exacerbating price movements through a domino effect.
- complex orders such as *stop-loss orders*, and so on.

- Finally, a combination of any of the above mechanisms could increase market reflexivity

All these mechanisms create clusters of first-, second- and higher-order price changes. To describe such cluster structure, the theory of branching processes introduces the *branching ratio* (Harris, 2002), denoted as n . This key parameter corresponds to the average number of daughter events of first generation per mother event. When the branching ratio is small ($n \ll 1$), the dynamic is stationary and mostly driven by the external (exogenous) uncorrelated immigrants in the system, as most of clusters contains only one or few events. When the average number n of daughters per mother increases, clustering rises and the self-excitation mechanisms play an increasingly important role in the system's dynamics. When the branching ratio is close to one ($n \lesssim 1$), the external stimulation of the system by zero-order (news-driven) events is strongly dominated by the enormous growth in the clustering of endogenous events. Finally, when the branching ratio is above one ($n > 1$), implying that each price change triggers on average more than one future price change, the dynamic becomes non-stationary. With finite probability, the system explodes in an infinite number of events without need for a permanent supply of fundamental (exogenous) events. In the theory of branching processes (Harris, 2002), these three regimes are called: (i) *sub-critical* ($n < 1$), (ii) *critical* ($n = 1$) and (iii) *super-critical* or explosive ($n > 1$).

In the sub-critical regime ($n < 1$), in the case of a rate of truly informative (or exogenous) news $\mu(t)$ which is constant ($\mu(t) = \mu = \text{const}$), the branching ratio n can be shown to be exactly equal to the average fraction of endogenous events (i.e., due to past price changes) within the whole population of events (Helmstetter and Sornette, 2003; Filimonov and Sornette, 2012). In other words, the branching ratio is equal to the fraction of events that are triggered due to the internal feedback mechanisms described above. This can be seen as follows. The total average activity (average number of trades per unit time) is $\mu + \mu n + \mu n^2 + \mu n^3 + \dots = \mu/(1 - n)$, reflecting the cascade of triggering over the successive generations (mother \rightarrow daughter \rightarrow grand-daughter \rightarrow ...). Subtracting the rate μ of immigrants, we get the rate $\mu/(1 - n) - \mu = \mu n/(1 - n)$ of triggered events of all generations. Their ratio to the total activity $\mu/(1 - n)$ is indeed n . To repeat, while by definition n is the average number of triggered events of first generation per exogenous event, it is also the average fraction of all triggered events. The number $1 - n$

is therefore the fraction of price changes that are due directly to exogenous surprising news. This means that, as n tends to 1, most of the activity becomes endogenous (or reflexive) and the total observed activity diverges. As a quantitative illustration, for $n = 0.8$, the observed number of trades is 5 times the number of trades that would exist if each trade was only reacting directly to an incoming unanticipated news.

2.2.3. Implications of high endogeneity for the inefficiency and possible instability of price discovery

A priori, the excess endogenous trading described by the Hawkes process could be interpreted as reflecting the tatonnement process of convergence of the price towards the fundamental price, and not taken as a diagnostic of a potential source of inefficiency or instability. This interpretation must be tampered by taking into consideration the following facts. First, it can be shown that the convergence time is also proportional to $1/(1 - n)$, which increases without bounds as n increases towards 1. In the mathematical literature on bifurcations (Sornette, 2006; Scheffer, 2009), this is referred to as “critical slowing down”: it takes more and more time for the system to adjust to new immigrants, due to the larger and larger number of triggered descendants and the longer and longer sequences of generations. This means that the convergence process to any true price becomes longer and longer, in other words, less and less efficient. Rather than agreeing rapidly on the “correct” price after the arrival of some unanticipated news, the traders trade longer and longer as $n \rightarrow 1$, not knowing on what price to settle. As the branching ratio n increases from, say 0.2 to 0.8, not only the activity increases by a factor of 4 but the convergence time to the true price is multiplied also by this factor 4. This supports the interpretation that, as n increases, endogeneity makes the market less efficient. Moreover, not only the rate of price changes and the convergence time diverges proportionally to $1/(1 - n)$, but the variance of the event rate also diverge as $n \rightarrow 1$. The susceptibility to external shocks diverges similarly. All these singular behaviors (in the mathematical sense) point to a growing instability of the system as the branching ratio increases.

2.2.4. Estimation of the branching ratio (reflexivity index)

Two different methodologies can be used to compute the branching ratio. The first one involves reverse-engineering the clusters by reconstructing ensembles of scenarios for the top structure in Fig. 3 from the known bottom

timeline and calculating the ensemble average number of direct descendants from any given event. This can be done using *stochastic declustering* procedures (see Zhuang et al. (2002); Marsan and Lengline (2008)). This approach is sophisticated and time consuming, since it provides in principle the full reconstructed history of the generation process. Moreover, its precision deteriorates when the memory τ of the kernel $h(t)$ given by expression (2) increases (Sornette and Utkin, 2009). If interested only in the determination of the branching ratio n , it is simpler to estimate the parameters α and τ of the kernel $h(t)$ (2) by maximum likelihood and use the relation

$$n = \int_0^{\infty} h(t)dt = \alpha\tau. \quad (3)$$

Indeed, the Maximum Likelihood Estimation (MLE) method benefits from the fact that the log-likelihood function is known in closed form for Hawkes processes (see Ogata (1978); Ozaki (1979) for an analytical expression of the likelihood function). Then, the MLE method provides a statistical estimation of α and τ , and therefore of n , and in addition of μ .

The standard quantification of the goodness-of-fit of the data by the Hawkes process uses *residual analysis* (Ogata, 1988), which consists in studying the *residual process*, defined as the nonparametric transformation of the initial series of the event time stamps t_i into

$$\xi_i = \int_0^{t_i} \hat{\lambda}_t(t)dt = \mu t_i + \alpha \sum_{t_j < t_i} \exp\left(-\frac{t_i - t_j}{\tau}\right), \quad (4)$$

where $\hat{\lambda}_t(t)$ is the conditional intensity of the Hawkes process (1) estimated with the maximum likelihood method. As it was shown by Papangelou (1972), under the null hypothesis that the data has been generated by the Hawkes process (1) with kernel (2), the residual process ξ_i should be Poisson (memoryless) with unit intensity. The goodness-of-fit can then be verified both by (i) visual cusum plot or Q-Q plot analysis and (ii) rigorous statistical tests, such as independence tests applied to the sequence of ξ_i and/or tests of the exponential distribution of the transformed inter-event times $\xi_i - \xi_{i-1}$, which amounts to testing the uniform distribution of the random variables $U_i = 1 - \exp[-(\xi_i - \xi_{i-1})]$ in the interval $[0, 1]$.

3. Calibration of the Hawkes model to high frequency data

3.1. Description of the calibration methodology

3.1.1. Trade-off in choosing the relevant time-windows

The present study is based on the analysis of the reflexivity of the mid-quote price movements. As discussed above, the branching ratio n estimated with the Hawkes model (1)-(2) provides a direct quantification of the degree of reflexivity, but only under three important assumptions: (i) stationarity of the underlying process, (ii) sub-criticality of the regime ($n < 1$), and (iii) constant parameters (μ, n, τ) of the model. As pointed out earlier, trading activity has been in general increasing over the analyzed period, thus showing a non-stationary behavior. In addition, during any given day, the trading activity is very low outside trading hours (see table. 1, please note that due to international venue these hours are different from the so-called Regular Trading Hours), and exhibits strong intraday seasonality during the active trading hours, being on average almost twice as large at the beginning and at the end of the trading session in comparison with lunch time. Under these circumstances, to ensure that the assumptions of stationarity and constant parameters are approximately met, one needs to consider the smallest possible intervals that are still compatible with reasonable stable statistical estimations. However, smaller time intervals imply smaller numbers of events for the estimation and thus decreased robustness. More importantly, the size of the time window limits the memory of the endogenous process that can be recovered from the estimation procedure to about the size of the considered window. In other words, considering time intervals of just a few minutes prevents capturing memory effects that may develop over time scale of hours and longer.

3.1.2. The choice of intervals for analysis

In the present work, we choose a trade-off by considering time intervals of 10 minutes, such that parameters of the Hawkes model (1)-(2) can be considered approximately constant. At the same time, these intervals are wide enough to capture a significant part of the endogenous memory of the system: indeed, using the exponential kernel, our estimations give a characteristic memory time much smaller than minutes. The number of mid-quote price changes amounts typically to more than 100–200 events over a typical 10 minute window, which is sufficient to perform a reliable calibration. In the most active periods it can reach up to several tens of thousands. As

shown in (Filimonov and Sornette, 2012), the model calibration is robust to the choice of the interval: an increase of the window size to 20 or 30 minutes does not result in significant changes of the estimated parameters.

However, ten minutes intervals are not long enough to capture the memory of long-term herding mechanisms, which are responsible for bubble formation on time scales of months to years (Sornette, 2003). For instance, the recent oil bubble started approximately in 2004 and developed over several years until its burst in summer of 2008. Its detection with the Hawkes process calibrated in running windows of 10 minutes can only be done if the mechanisms working at large time scales somehow cascade down to the minute time scale, so that the branching ratio exhibits abnormal increase concomitant with the development of the bubble. We shall see that this indeed happened, for instance during the oil bubble culminating in July 2008 (see the results discussion on oil in section 3.4). Similarly, optimal portfolio executions take from tens of minutes to hours and (for extremely large portfolios) days). Therefore, execution of large orders has a minor impact on the quantification of the reflexivity index done here at the 10 minute time scale. The time that clearing houses give retail clients to react to margin call is also typically one day. This mechanism of reflexivity is thus also negligible at our time scale of investigation. Hedging usually also involve longer time scales. In the list of the sources of reflexivity presented in section 2.2, mostly (i) short-term human reflexivity, (ii) algorithmic and HF trading strategies, (iii) herding in algorithmic strategies and (iv) complex stop-loss and other orders are operating at the time scales of 10 minutes or less. As a consequence, only these mechanism can be captured by the model.

The combination of the use of small time windows of 10 minutes together with the choice of the short-memory exponential kernel (2) allows us to investigate the short-term speculative mechanisms of reflexivity. As we shall see, the short-term reflexivity shows interesting patterns that can be attributed to the changes in algorithmic and HF trading activity over the time of our analysis. We will see that self-excitation at short time scales has been growing steadily in most of the commodity markets in 2005–2009 (and for sugar markets even until the end of analysis at 2012). Moreover, Filimonov and Sornette (2012) suggests the usage of this short-term reflexivity for the forecasting of HF instabilities of markets such as “flash-crashes”. In order to account for long-term behavioral mechanisms described in section 2.2.2, one indeed needs considering much longer time windows up to several months and, as a consequence, power-law kernel (2) that can account

for long-memory effects. Such analysis performed in Hardiman et al. (2013) suggests that this long-term self-excitation of financial market has always been much higher as reflected by the long-term branching ratio being around the critical value of 1.

To analyze the short-term reflexivity of the market, we used the maximum likelihood estimator (Ogata, 1978; Ozaki, 1979) to calibrate the Hawkes model (1)–(2) in time windows of 10 minutes spanning every day from 2005 to 2012 with one minute time step (from 1997 to 2012 for the E-Mini S&P 500 futures contracts). We excluded the days when trading was closed before the end of hours of active trading (table 1) or with daily volume less than the 5% quantile of daily volumes for each given year.

Finally, we need to acknowledge that even such small time windows could be susceptible to some non-stationarity effects. In particular major macroeconomic news announcements (such as FOMC rate decision or EIA weekly report) clearly results in abrupt change in the dynamics of trading. In particular, it is no longer possible to assume that the background intensity $\mu(t)$ is constant over the period, if an important announcement falls within the window of analysis. For this reason, we have excluded the 10-minutes windows that contain announcement of FOMC (eight times per year) and EIA (weekly). However, since our analysis is based on monthly averages, the presence in our estimations of some small number of “outliers” due to news announcements do not change the overall statistics. The tight quantile ranges around the monthly average that we observe for all of commodities support this hypothesis.

3.1.3. Dealing with the TRTH uncertainties of timestamp recording

Despite the fact that ticks in Thomson Reuters Tick History (TRTH) are stamped with microsecond resolution, a rather large number of quote changes have identical timestamps. In the most recent years, we can observe up to several hundreds quote changes during active trading hours for the same timestamp. The origin of this phenomenon lies in the nature of the data feed from the exchange, which is obtained by the FAST/FIX protocol. The protocol bundles multiple updates of multiple instruments within a single message by an algorithm designed by the exchange. Then, the package is sent to the Thomson Reuters collection system, and TRTH timestamps relate to the time when the messages reach the collection system, but not to the time when the transactions were actually executed and recorded by the exchange. Since the exchange time, coded in the FAST/FIX protocol, is stamped with

a resolution of seconds, the actual time of any tick is uncertain within a range that is larger than or equal to the time between two consecutive FAST/FIX packages. This range varies from tens of milliseconds in recent years to several hundreds of milliseconds or even seconds in early 2000–2005.

An additional source of uncertainty to the timestamps is introduced by the latency of the message traveling from the exchange to Thomson Reuters collection system and by the overhead brought by processing the FAST/FIX protocol on both sides. However, both factors introduce a shift to the timestamp, which is constant when the latency does not fluctuate. Such constant time shift would not change the analysis and could be omitted in principle. In reality however, both actors may vary in time, but the order of magnitude of these variations is much smaller than the time between consecutive packages and thus could be neglected. The typical time for package processing is of the order of tens microseconds. The latency is usually also of the order of tens of milliseconds. This suggests a rough estimate of their variations, typically of the order of milliseconds, which is much smaller than the uncertainty introduced by the bundling of updates to a single message.

Two possibilities can be considered to deal with the uncertainty in the timestamps resulting from FAST/FIX protocol. One is to consider only the timestamps provided by the exchange (with resolution of seconds) as a reliable source of data. The other is to use enriched millisecond timestamps of TRTH, while accounting for the uncertainty due to bundling updates in FAST/FIX packages. In this paper, for the sake of caution, we follow the second option by relying on the non-zero difference of timestamps between consecutive transactions or updates of quotes as the proxy for the uncertainty in arrival times. Table 3A and B provide respectively the annual average and median uncertainties of the event timestamps for our different considered instruments. Starting in the range of 200–300 milliseconds in mid 2000s, the average timestamp uncertainties have decreased progressively over the years. In 2012, the average and median durations between two consecutive FAST/FIX packages that were recorded with different timestamps by TRTH reached a relatively low range of 103–242 milliseconds and 22–135 milliseconds, respectively.

In order to make the data compatible with the Hawkes model for which the probability of having multiple events with identical timestamps is equal to zero, we follow the methodology developed by Filimonov and Sornette (2012). Specifically, we randomly redistribute the TRTH timestamps around their recorded values within an interval of duration Δ . In doing so, we

implicitly assume that each event occurring within the interval of width Δ is independent of all the others within the same interval (but not between different intervals). This processing step tends to lead to underestimations of the endogeneity levels presented below.

The intuition, that Δ should be chosen to be of the order of the typical duration between consecutive FAST/FIX packages, has been validated by numerical tests. Similarly to (Filimonov and Sornette, 2012), we have verified the procedure on synthetic time series obtained by numerical synthesis of the Hawkes process (1) with parameters (μ, n, τ) close to the calibrated values of the real data. The results of such synthetic tests and comparisons with the estimation on real data has revealed quantitative limitations of the proposed method. In particular, the distortion of the distribution of inter-event times becomes significant and, as a consequence, the estimation of the parameters of the Hawkes process becomes unreliable when Δ is chosen to be significantly smaller than the typical waiting time between consecutive FAST/FIX packages. As a rule of thumb, Δ should be chosen to be more than the median duration (table 3B). As seen from table 3B, for most years, the median timestamp uncertainty is of the order or below 100 milliseconds, while its average value is of the order or below 200 milliseconds. This suggests that a reasonable value for Δ is 200 milliseconds. In order to check the robustness of the method applied to real data, we have also used the values $\Delta = 50$ milliseconds, $\Delta = 100$ milliseconds and $\Delta = 300$ milliseconds, where the last value corresponds to the upper bound of the average uncertainty (table 3A). As an extreme case, we have also considered $\Delta = 1$ second, which corresponds to the resolution of the exchange time.

3.1.4. Testing the goodness-of-fit of our calibration

The goodness-of-fit tests are essential to quantify the agreement between the model and the data. As goodness-of-fit tests, we have used residual analysis, described in section 2.2. In a nutshell, after performing a calibration of the Hawkes model (1)–(2), we performed the non-parametric transform (4) to obtain the residual process and then obtained the transformed inter-event intervals $U_i = 1 - \exp(\xi_{i-1} - \xi_i)$. Under the null hypothesis that the data is generated by the Hawkes model, these transformed inter-event intervals U_i should be iid uniformly distributed in the interval $[0, 1]$. We have used the Kolmogorov-Smirnov test in order to test the uniformity of the distribution of U_i 's.

Each 10-minute data interval is characterized by the estimated parameters

(μ, n, τ) and a p -value of the goodness-of-fit. We reject the null hypothesis if the p -value is smaller than the confidence level of 0.05. Table 4 presents the fraction of time-windows for which the null hypothesis was rejected, for a set of future contracts, at different years and for different values of Δ . We also exclude from the analysis the years prior to the introduction of electronic trading for each analyzed commodity (see table 1), due to the weak reliability of the corresponding quote timestamps when the open pit still existed. As discussed above, for Sugar #11 we have additionally excluded the time period before March 6, 2008, when electronic trading on the ICE platform coexisted with pit trading at NYBOT. Those years are marked with dash lines (—). The reliability of timestamp (Δ) should be comparable with the typical inter-package times. We have excluded from the analysis time periods where the median inter-package times (table 3B) was larger than Δ . Those time intervals are marked in the table 4 with stars (***) .

As seen from the tables 4A to 4E, the quality-of-fit is usually good. In total, for all the analyzed commodities over the years of 2005–2012 for 10-minutes intervals and with $\Delta = 200$ milliseconds (msec), we reject 452'514 estimations out of a total of 3'332'016 estimations, corresponding to a rejection rate of 13.6 per cent. For smaller Δ 's, the agreement of the Hawkes model with the data worsens: for $\Delta = 100$ msec, we reject 576'532 out of 2'605'129 estimations, a rejection rate of 22.1 per cent; for $\Delta = 50$ msec, we reject 434'662 out of 1'165'761 estimations, a rejection rate of 37.3 per cent. By contrast, for $\Delta = 300$ msec, we reject 303'420 out of 3'375'079 estimations, a rejection rate of 9.0 per cent; for $\Delta = 1$ second, we reject only 70'126 out of 3'375'079 estimations, a rejection rate of 2.1 per cent. Due to the strong distortion of the statistics of inter-event intervals occurring for small Δ 's and the generally poor agreement of the Hawkes model with the data (as quantified by the Kolmogorov-Smirnov test), we have not presented results for $\Delta < 100$ msec. However, we must acknowledge that the results obtained for $\Delta = 50$ msec agree within the confidence intervals with those obtained for larger Δ 's. Despite very good agreement between model and data for $\Delta = 1$ second, we will see that use of such large Δ (which lasts 3–10 times more than the typical waiting time between packages) results in a significant overestimation of the reflexivity index n .

The good results of our tests support our use of the exponential kernel (2) in the specification of the Hawkes model (1). We need to mention that the results of the analysis at 10 minutes intervals are robust to the choice of the kernel. In particular, our tests have shown that using long memory (power

law) kernel results in small bias in the estimations of the reflexivity index n , but does not change the overall secular dynamics presented in the following sections.

In order to characterize the possible long-term evolution of the parameters over the whole investigated period, we have taken averages of the estimates of the parameters (μ, n, τ) over all 10 minutes windows within a one month period. In the following subsections, we report these average estimates together with quantile ranges. Interestingly, when considering either all estimates within one month period, or only the estimates that could not be rejected with the Kolmogorov-Smirnov test, the averages and the quantile intervals remain similar. However, to be consistent, we have excluded from the averages those estimates for which the null hypothesis of the Hawkes model as the generating process for the data could be rejected (corresponding to estimates with p -value below 0.05).

3.2. Financial markets: E-mini S&P 500 futures

Before analyzing commodity market, we revisit the analysis initially performed until August 29, 2010 in (Filimonov and Sornetto, 2012) of the E-mini S&P 500 futures contracts, which are traded on the Chicago Mercantile Exchange (CME). Being introduced in 1997 as a supplement to the regular S&P 500 futures contracts with a reduced size of 50 times the value of the index, the E-mini has attracted many small investors and has become one of the most actively traded derivatives in the world.

Fig. 4(a) and (b) present, respectively, two-month volume and trading activity (measured in number of mid-quote price changes) as well as daily volatility and price dynamics for the E-mini S&P 500 futures contracts between 1998–2012. Together with the dynamics of these traditional measures of activity, Fig. 4(c) and (d) show, respectively, the dynamics of the estimated background intensity $(\hat{\mu})$ and branching ratio (\hat{n}) over the same time period. The estimates for each different Δ (100, 200, 300 milliseconds) are practically undistinguishable. This observation together with the narrowness of the 25%-75% quantile range confirm both the relevance of the Hawkes model as an excellent data descriptor and the robustness of our estimation procedure. Both observations will be later verified with the data analysis performed on commodities futures. Finally, let us note that considering $\Delta = 1$ sec, which corresponds to the uncertainty of exchange timestamps, results in slightly higher branching ratios, but does not change its overall dynamics.

Fig. 4 shows that the number of mid-price changes in panel (a); the daily volatility in panel (b); and the background intensity $\hat{\mu}(t)$ in panel (c) exhibit synchronized peaks that coincide with major episodes of market instabilities. Indeed, one of the first peak coincides with the burst of the ICT dot-com bubble (Johansen and Sornette, 2000). Note the synchronized behavior during the following bearish regime as well as during the financial crisis that started in 2007, including its culmination with Lehman Brothers bankruptcy (Sornette and Woodard, 2010). Note that the increase of trading activity from 1998 to 2012, as proxied by volume in Fig. 4(a), is not accompanied by an increase of the background intensity $\hat{\mu}$ of exogenous events in the market. This makes intuitive sense since μ should reflect the genuine news impacting the market.

In contrast with Fig. 4(a), (b) and (c), the time evolution of the branching ratio \hat{n} presented in Fig. 4(d) exhibits a very different behavior⁵. Importantly, one should note that the branching ratio is not simply another measure of trading activity or the frequency of price changes. Indeed, Fig. 4 illustrates the existence of completely different dynamics of the branching ratio compared with measures of activity such as volume or mid-quote price changes. We address this point in greater details when we discuss several robustness tests in section 3.6. For now, we only highlight the following findings.

- (i) Between 1998 and 2004, the monthly trading volume⁶ increased almost 36 times (from 316'401 contracts in February 1998 to 11'428'371 contracts in February 2004). However the branching ratio increased only slightly from 0.35 to 0.4 during this period.
- (ii) Similarly, despite an almost doubling of the volume from 25'890'923 in June 2007 to 55'251'608 in August 2007, the branching ratio decreased from 0.6 to 0.45.
- (iii) The same period a year later, June–August 2008, could be considered as another example illustrating the decoupling between trading activity and branching ratio. The number of transactions increased 3 times (from 1'346'928 to 4'191'227) and the number of mid-quote price changes doubled (from 230'022 to 580'220) over the same period

⁵Note that the present analysis slightly differs from the one presented in (Filimonov and Sornette, 2012).

⁶Note that here and after we discuss monthly volume and number of transactions, and average branching ratios over one month interval, however for the sake of clarity in fig. 4 we plot dynamics of two-months values.

of time. Nevertheless, this did not lead to an increase of the branching ratio, which, by definition, is estimated based on mid-quote price changes.

- (iv) On the opposite, the dramatic surge of the branching ratio from 0.5 in September 2009 to 0.75 in March 2010 coincided with a moderate increase of volume (from 31'574'403 to 43'320'946) and of mid-quote price changes (from 219'918 to 266'014).
- (v) Finally, one could observe that the spike in volume and background activity in June 2010 did not affect the branching ratio at all, and even a similar spike in September 2011 is associated with a sharp decline of the ratio.

The branching ratio, measuring the level of endogeneity, increased regularly from 2002 onwards and peaked first with the beginning of the bear market on the E-mini in 2007/2008. The successive policy interventions led to a decrease in the level of endogeneity until August 2009. Afterwards, the branching ratio rapidly rose again to reach a high plateau from the beginning of 2010 onwards with a small transient decline at the end of 2011, which coincides with the discussion on the debt ceiling in the US and a deepening of the eurozone crisis. The decline of endogeneity level between 2008 and mid-2009 coincides with a series of financial and economic interventions, when fundamentals, like liquidity provision to avert a credit freeze on financial capital markets and stimulus packages to revive aggregate demand, were prevalent features. Although it is beyond this paper to explain the effect of quantitative easing policies on endogeneity, one notes that the sharp rise of the endogeneity level coincides with the first hints of a possible second round of easy liquidity in August 2009⁷. In parallel, animated discussions about the shape of the economic recovery, either V, U, L or W, from summer 2009, added to economic uncertainties as characterized by the risk-on/risk-off behaviors. These economic uncertainties as well as the quantitative easing policies have remained prevalent in the subsequent years until today, while risks of credit freeze in financial capital markets have receded substantially. The combination of economic uncertainties and unlimited liquidity could rationalize the high plateau of the endogeneity level measured on the E-mini S&P 500 futures.

⁷c.f. "QE2: Will the Fed Surprise the Markets?" <http://www.thestreet.com/story/10909094/qe2-will-the-fed-surprise-the-markets.html>

3.3. The evolution of endogeneity in commodity futures markets

In Fig. 5–9, we report our estimates of the branching ratio for several commodities. Similar to the endogeneity level of the E-mini S&P 500 futures, we usually find the average levels of endogeneity above 50 per cent for all considered commodity markets since the mid-2000s. Moreover, we observe that endogeneity levels are greater in 2012 than when our estimates start. Nevertheless, these increases have not necessarily been monotonous. In the case of oil, both series, Brent and WTI, show a gradual increase before reaching a peak in late-2008 to early-2009. Afterwards, they have partly receded. By contrast, the level of endogeneity of the soft commodities exhibit a marked oscillating pattern around a upward-sloped or constant long-term trend.

In addition, when our data can go back to 2005 or 2006, we usually observe a period of about half a year in the course of 2006 or 2007 when the branching ratio escalates sharply to higher levels from which it does not recede anymore. In particular, we observe this phenomenon for the Brent crude oil in Fig 5(a) during 2006, when the monthly averages of its branching ratio move from roughly 0.4 to 0.6, and for the White Sugar futures market in Europe in Fig. 7(a) in the course of 2007, when this figure rose sharply from about 0.3 to 0.55. These phenomena are similar to the pattern on the E-mini S&P 500 futures markets over 2005 and 2006. Nevertheless, for the commodities, these episodes seem to have taken place over a shorter time span.

As mentioned above, we cannot compute the branching ratios prior to the introduction of full electronic trading and of sufficient liquidity in the mid-2000s. Nevertheless, in all likelihood, the levels of endogeneity in commodity markets in late 1990s and in the first-half of 2000s were not greater than the one of the E-mini S&P 500 futures market at that time. This conjecture makes us believe that the already-high endogeneity levels that we observe for all commodity markets in the second-half of the 2000s have not been a permanent feature in the period prior to the introduction of full electronic trading and the availability of reliable tick data on commodity derivatives.

It should also be recalled that, at first sight, our reflexivity indices are not particularly designed to capture longer-term herding mechanisms, which are responsible for bubble formation on time scales of months to years. Since we need to calibrate the model in running windows of 10 minutes, our branching ratio does not have by definition a long-term memory. However, we surprisingly find that the mechanisms working at longer time scales sometimes seem to cascade down to the shorter intervals on which we compute our indices.

In specific cases discussed below, we show in fact that growing average levels of endogeneity during several months sometimes coincide with bubble-bust cycles.

In addition, some common shocks seem to impact commodity markets simultaneously. For example:

- Among the commodities examined in this paper, many US commodities exhibit a decline of the branching ratio for the period around June/July 2011 (WTI, Wheat, Corn and Soybean). This period coincides with discussion on the US debt ceiling and fears that no agreement could trigger a worldwide new economic downturn. It relates also to oil reserve releases by the IEA Members and better than expected weather conditions in the US.
- Interestingly, the second half of 2012 exhibits a sharp and synchronized increase of the branching ratio for all US traded commodities here examined, possibly on expectation of QE3.
- Oil, including Brent and WTI, seem to follow the same pattern as the E-mini from their endogeneity peak of 2008 to the low of 2009, which coincides with the August 2009 hint of further quantitative easing (see discussion above in section 3.2)
- For the WTI, Corn and Wheat, one can observe a slight decline in the branching ratio early 2008 before the bubble bursts. In the USA, ethanol is mostly produced from corn and some substitution effects (oil/corn and corn/wheat) could explain this common feature in the branching ratio.

After these general observations, we concentrate on each different commodity market specifically.

3.4. Crude oil: Brent and WTI

We start our analysis of commodity futures with oil prices, which exhibited a record rise followed by a spectacular crash in 2008. The peak of Brent daily close prices at USD 146.08 (daily high of USD 146.69) per barrel was observed on July 3, 2008. Six months later, Brent prices reached a trough of USD 44 (daily low of USD 43.03) on December 19, 2008, a level not seen since 2004.

Fig. 5 presents the dynamics of the price and of the daily volatility (estimated with the Garman&Klass open-high-low-close estimator (Garman and Klass, 1980) for two major futures contracts on light sweet crude oil: Brent Crude (fig. 5(a)) and WTI (fig. 5(b)). Along with the price dynamics and the price volatility, the evolution of the estimated branching ratio (effective degree of reflexivity) is presented. Different symbols on the plots correspond to different values of Δ : 100, 200, 300 milliseconds and 1 second.

Fig. 5 documents the following regimes for both Brent Crude (Europe) and WTI (US):

- The branching ratio has shown an upward trend over the whole bubble period until early 2009, having an intermediate peak in July 2008 coinciding with the end of the oil bubble.
- The branching ratio exhibited three large periods of stabilization, which were preceded by a small drop: in Q1-2007, Q1-2008 and from mid-2008 to the end of 2008.
- In the last period of the run-up (December 2007—April 2008), the branching ratio showed a pronounced drawdown for the WTI.
- The branching ratio started accelerating again until the price peak in July 2008.
- The branching ratio stayed high (at values of 0.7 for WTI and 0.75–0.78 for Brent) during the whole period of the price fall until the bottom in December 2008, even exhibiting a maximum at the price bottom higher than its previous peak reached on July 2008. This illustrates that the branching ratio is also independent from price trajectory, in addition to being unrelated to volume or mid-quotes changes (see the discussion on the E-mini in section 3.2).
- Thereafter, the branching ratio starts decreasing until mid-2009. Afterwards, the dynamics of the branching ratio for Brent and WTI slightly diverged:
 - (i) The branching ratio for Brent was falling until December 2010 (see the note below on the sharp fall) and then changed to a sideways trend.

- (ii) The branching ratio for WTI was falling until mid-2009, and started rising again, peaked on April–May 2010, before a sharp (but small) fall in May 2010 occurred. This is similar to the pattern observed on the E-mini between mid 2009 to mid 2010 (see the discussion on the E-mini in section 3.2).
- The branching ratio sharply decreased for the Brent at the end of 2010 (also visible on the WTI) although the price increased. This might be attributed to the unusual cold weather in Europe that lift up oil demand unexpectedly and reduced herding mechanism temporarily.
- Beginning 2011, the branching ratio increased sharply for the WTI and also for the Brent following the start of the Arab Spring and speculation on oil output of some producing countries like Libya. While the cross-market correlations between these commodity futures and the US equity markets collapsed during that period and suggested a growing role of the fundamentals (Bicchetti and Maystre, 2012), the endogeneity levels on each specific market grew during that period.

The most remarkable result obtained from the calibration of the branching ratio is its very large increase during the period when oil prices started to accelerate. The fact that our methodology identifies a growing reflexivity during the ascent of the price and, even more so, during its collapse, is particularly interesting in view of other analyses that documented strong evidence for the existence of a bubble during that period. Since the beginning of 2008, a growing number of specialists⁸, bankers⁹ and academics¹⁰ were considering the possibility that oil may have entered a bubble regime. The tormenting question was: how to justify the quadrupling of oil prices since 2003? Some attributed it mainly to the growing demand from the emergent China and India markets, a claim that former Chinese President Jiang Ze-Min himself debunked at least for China (see Fig. 3 with caption in English in Jiang (2008)). Comparing the values on World liquid fuel supply and

⁸See e.g. Zumbun, J., Soros tells congress to pop an oil bubble, *Forbes*, 3 June 2008.

⁹Credit Suisse, The Investment Committee Meeting of May 27, 2008.

¹⁰See e.g. Siegel, J. and W. Henisz, What's Behind the Flare-ups in Oil Prices? Jeremy Siegel and Witold Henisz Weigh In, Knowledge@Wharton, May 28, 2008; also see Krugman, P., More on oil and speculation (The Conscience of a Liberal), *The New York Times*, May 13, 2008.

demand reported by the International Energy Agency (IEA) and by the US Energy Information Administration (EIA), Sornette et al. (2009) noted that, until the end of 2005, both agencies reported consistent numbers showing that supply was systematically exceeding demand. Since 2006, there was a significant discrepancy between the numbers presented by the two agencies, ushering a period of uncertainty or opaque reporting, with no clear conclusion on whether an excess demand versus supply was the cause of the appreciation of oil prices. One can argue that the lack of clarity of the oil supply versus demand during that period ushered a period of growing speculation both in the literary sense of “forming conjectures” and in the financial sense, based on the general fact that the more imprecise is the estimation of the fundamental value of an asset, the more room there is for “stories” and “new economy” thinking that can justify speculative bubble prices (Kindleberger and Aliber, 2005; Sornette, 2003).

Sornette et al. (2009) further support the hypothesis that the 2007-2008 oil price run-up was amplified by speculative behaviors of the type found during a bubble-like expansion. They analyzed oil prices in USD and in other major currencies and found clear diagnostics of speculation. Based on the mechanism of positive feedbacks and the concept of emergent phase transitions (or bifurcation) to another regime using analogies with statistical physics and complexity theory, Sornette et al. (2009) used an approach that diagnoses bubbles as transient super-exponential regimes Sornette (2003). In a nutshell, the methodology aims at detecting the transient phases where positive feedbacks operating on some markets or asset classes create local unsustainable price run-ups. The mathematical signature of these bubbles is a log-periodic power law (LPPL, see e.g. Sornette and Johansen (1998); Johansen and Sornette (1999); Johansen et al. (2000); Sornette and Johansen (2001)). The power law finite-time singular process models the faster-than-exponential growth culminating in finite time at some critical time t_c . The log-periodic oscillations reflect hierarchical structures (Johansen and Sornette, 1999; Johansen et al., 2000) as well as competition between the trading dynamics of fundamental value and momentum investors (Ide and Sornette, 2002). Reproduced from Fig. 5 in (Sornette et al., 2009), Fig. 6 shows the calibration of the LPPL model to the oil price (NYMEX Light Sweet Crude, Contract 1, from the Energy Information Administration of the U.S. Government). The shaded box shows the 80 per cent confidence interval of the critical time t_c indicating the end of the bubble. Note that this analysis was performed ex-ante before the oil price did peak and was presented as a

genuine real-time prediction, which turned out to be successful, as recounted by Sornette et al. (2009). The overall conclusion of this analysis is that the geopolitical events unfolded in 2007 and 2008 have participated in raising the level uncertainty, which worked as a fertilizer for speculation, leading to oil prices increasingly decoupled from fundamental valuation (the hallmark of a bubble).

Our present analysis summarized in Fig. 5 of the effective degree of reflexivity estimated with high frequency data has shown that, during the bubble period, the herding between investors existed not only at scales of years but was also accompanied with short-term herding of the algorithmic trading strategies. Combining the evidence of Fig. 5 and Fig. 6, we conclude that the positive feedback mechanisms working at large time scales, which are at the origin of the oil bubble, cascaded down to the minute time scales and were reflected in the abnormal increase of the branching ratio that occurred concomitantly with the development of the bubble and its burst. Such cascade processes (Arneodo et al., 1998) are captured by the concept of multifractality that has been found to provide a remarkably powerful description and predictor of asset return dynamics (Muzy et al., 2001; Calvet and Fisher, 2008; Sornette et al., 2003; Lux, 2008).

3.5. Soft commodities: Soybean, Sugar, Corn and Wheat

3.5.1. Sugar (Europe and US)

One can distinguish four main regimes in the dynamics of the branching ratio for Sugar (Europe) shown of Fig. 7(a):

- Before 2007, the branching ratio is hovering around 0.3, with a rather large standard deviation due to the limited size of the data set resulting from relatively low trading activity.
- Starting at the beginning of 2007, the branching ratio increases rapidly and doubles in less than three quarters, stabilizing around the value 0.5-0.6 in the third quarter of 2007.
- From the fourth quarter of 2007 till mid 2011, the branching ratio is practically stable and remains in the range 0.5-0.6 with some excursions higher up.
- From mid-2011 to the end of 2011, the branching ratio increases and passes over 0.75. Thereafter, it decreases but remains in the range 0.6-0.7.

The rough pattern is similar for Sugar #11 shown in Fig. 7(b), except that (i) the story is shorter, (ii) there is a divergence in endogeneity dynamics with respect to Sugar (Europe) in spring 2009: sugar (US) experiences a spike in endogeneity levels from around 0.55 to 0.7. The peak of the branching ratio observed for Sugar (Europe) at the very end of 2011 is absent for Sugar #11 and its dynamics in 2012 is characterized by a monotonous increase of the branching ratio in the second part of 2012 to the extreme high level of 0.8 (possibly on expectations of QE3 discussed above).

The volatility of the branching ratio in 2009 possibly owes a lot to speculation on government intervention and on sugar supply deficit in usually net exporter countries like India and Brazil. The fluctuation of the branching ratio in 2010–2012, which is not synchronized between European and American sugar, may be rationalized by divergent internal discussions on both sides of the Atlantic regarding quotas affecting sugar imports and production.

3.5.2. Corn (US), Soybean (US) and Wheat (US)

The dynamics of the prices and volatilities of Corn (US), Soybean (US) and Wheat (US) shown of Fig. 8(b), Fig. 8(a) and Fig. 9 share many similar features. They exhibit a very large peak in mid-2008, only surpassed by very recent price surges for Corn and Soybean, followed by a deflating price until mid-2010. These peaks in mid-2008 are coincident with the peak of oil price previously discussed and are symptomatic of the commodity bubble that developed in 2007 and 2008. One can also notice a precursory peak in the first quarter of 2008, which is especially pronounced for Soybean (US) and is actually a dominant price feature for Wheat.

Interestingly, in contrast with the behavior of the branching ratio for oil, which accompanied by its increase the growth of the oil price bubble, the branching ratios for these commodities remained in the range 0.5–0.6, with some spikes before the price peak and spikes associated with the price correction following the peak in mid-2008. However, since the branching ratio could not be computed prior to the end of 2006 for these commodities, one can reasonably assume based on the branching ratio measured on the E-mini S&P 500 and Sugar (Europe) that the reflexivity level is likely to be around 0.3–0.4 for the earlier years. Therefore, the level measured just before the commodity bubble burst is already relatively high and does not exhibit the jump seen on the Brent or the E-mini in early 2006.

The branching ratios of Corn, Soybean and Wheat remained approxi-

mately constant after this bubble episode, in the range 0.40–0.6 from mid-2009 to mid-2010 for Soybean and Wheat and in the range 0.5–0.65 for Corn from mid-2008 to mid-2010. The branching ratios of both Soybean and Corn exhibited a sharp increase from 0.5 to 0.6 for Soybean and even to 0.7 for Corn and Wheat, in the third quarter of 2010, which can be associated with a change to a phase of rising prices. In fact, in the summer of 2010, Russia announced an export ban for wheat and Ukraine followed by announcing export restriction¹¹. As wheat, soybean and corn are substitutes for feed grain, the export restrictions created a shock that impacted simultaneously these markets. The substitutability between these commodities through their use as feedstock creates correlation between them.

Thereafter, the branching ratio for Soybean showed a steady decline from 0.6 to slightly above 0.5 in October 2012. In contrast, the branching ratio for Corn exhibited much larger volatility with drops down to 0.4 and peaks up to 0.6 from 2011 to September 2012.

3.6. Robustness tests

The branching ratio represents a standalone measure of reflexivity, which is not affected by a simple increase or decrease of trading activity (measured in the number of transactions or volume) or by price changes. As discussed in sections 2.2 and 3.1, the input for the calibration of the Hawkes model is the series of timestamps of mid-quote price changes, independently of their directions. Thus, the branching ratio is insensitive to the presence and direction of trends, whether the price is rising, falling or moving sideways. An increase of the branching ratio qualifies an increase of self-excitation in the price formation mechanisms and, as explained in section 2.2.2, could signal the development of short-term instabilities and of incoming crises.

Similarly to the effect of the direction of price moves, neither transactions nor volume enter directly into the formulation of the Hawkes model, since individual transactions do not necessary result in a change of the mid-price. As an example, doubling the number of transactions by splitting each of them into two independent transactions (to keep the daily volume constant)

¹¹See Reuters article “Snap analysis — Race for Russia’s grain business after export ban” (<http://www.reuters.com/article/2010/08/05/uk-russia-grain-export-ban-idUKTRE6744E720100805>) and BBC article “Ukraine sets grain export quotas following drought” (<http://www.bbc.co.uk/news/business-11495369>).

does not affect the dynamics of the mid-price at all. Similarly, keeping the number of transactions constant and doubling the volume of each of them (doubling the volume of each incoming market order) while simultaneously doubling the volume of all incoming limit orders again would not change the dynamics of the mid-price. The number of events (mid-quote price changes) also does not directly affect the parameters of the model except for the background intensity parameter μ . For instance, doubling the number of events by superimposing two identical clusters from fig. 3 (or two clusters generated with identical parameter set (n, τ)) will result in doubling the background intensity μ that quantifies the rate of exogenous (zero-order) events in the system, but the branching ratio n will not be changed.

The above is theory, but does it hold in practice, in particular in the statistical estimation of the branching ratio with limited data and for different parameters? In order to reject the possibility that the observed dynamics of the branching ratio could reflect an increase of trading activity, we performed the following test. Fixing the number of mid-quote price changes per day, we redistributed these events in time such that, within one day, their dynamics was described by a Poisson process. This “redistribution” of the time series amounts to keeping the price trajectories, the daily volume, the number of price and mid-quote price changes per day unchanged, i.e., keeping the same trajectories as shown in Fig. 5–9 while distorting time such that the intervals between consecutive mid-quote price changes within one day become uncorrelated and exponentially distributed. Then, we performed exactly the same procedure as described in section 2.2 and 3.1. Namely, we divided each day in 10 minutes intervals, rounded timestamps to nearest end of sub-interval of size Δ of 100, 200, 300 milliseconds and 1 second (which would correspond to introducing uncertainty in timestamps), implemented the procedure described in section 3.1 and estimated the parameters (μ, n, τ) of the Hawkes process (1) with an exponential kernel (2) within each of these intervals. Fig. 10 presents the results of this robustness test. Despite the increase of activity (measured, for instance, in the number of transactions) and the increase of trading volume (see table 2), as well as the existence of a highly nontrivial seasonal volume dynamics (see fig. 1), the random shuffling of the time stamps have completely erased the self-excited structure of the time series. Indeed, the estimated branching ratio in the randomized time series is consistently found very small, as it should. Its average and median values are always $\hat{n} \lesssim 0.08$ and the 75%-quantile is below 0.1. One can thus clearly reject the hypothesis that the branching ratio is sensitive

to, or equivalently provides another measure of, trading activity and trading volume. This quantitative result supports the key property of the Hawkes model, which is that the branching ratio is not determined by the average rate of events but by the degree of self-excitation of the system.

4. Conclusion

Using the Hawkes self-excited conditional Poisson process, we have quantified the degree of endogeneity in the price dynamical generating process of a number of highly-traded commodity futures markets. For all analyzed markets, we have found high levels of endogeneity. On average, our conservative estimates show that more than one out of two price changes is due to another preceding price change since the second-half of the 2000s, and not due to an exogenous piece of news. In other words, price dynamics on these commodity markets are partly driven by self-reinforcing mechanisms. In our view, this evolution partly reflects the development of algorithmic trading and of high frequency trading in particular. Using the insights obtained from the properties of the Hawkes self-excited conditional Poisson process calibrated to the commodity futures markets, we infer that these high levels of endogeneity are likely to make the price formation process less efficient, because higher endogeneity implies a longer convergence process. Moreover, it also points to a growing instability of the system, as we explained in section 2.2.

Our robustness tests show that our measure of endogeneity is independent of other factors that have also experienced significant changes over the last decade. More importantly, it is also independent of the background intensity of exogenous events in these markets. Interestingly, we do not observe a long-term increase of this variable in parallel of the developments that we observe for the other variables, like transactions volumes. This suggests that the rate of genuine news impacting the market, reflected in our measure of background intensity, has remained relatively constant over the analyzed period. Our results also do not support the view that the financialization of commodity markets has allowed to process a greater set of relevant information than what previous market participants considered before the rise of electronic trading.

While our index does not have a long-term memory, interestingly, we find that it can still provide some interesting insights when the mechanisms working at longer time scales cascade down to shorter terms, as occurred

for during the Oil bubble that culminated in July 2008 and crashed until December 2008.

Acknowledgments

We would like to thank Dr. Mika Kastenholtz for fruitful discussions. We also would like to thank Dr. Heiner Flassbeck, former Director of the Division on Globalization and Development Strategies, UNCTAD, without whom this research would have not been possible.

References

- Aït-Sahalia, Y., Mykland, P. A., Zhang, L., 2005. How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise. *Review of Financial Studies* 18 (2), 351–416.
- Aite Group, 2009. *New World Order: The High Frequency Trading Community and Its Impact on Market Structure*. Tech. rep., Aite group.
- Arneodo, A., Muzy, J.-F., Sornette, D., 1998. “Direct” causal cascade in the stock market. *The European Physical Journal B* 2 (2), 277–282.
- Azizpour, S., Giesecke, K., Schwenkler, G., 2011. Exploring the sources of default clustering.
- Bauwens, L., Hautsch, N., 2009. Modelling Financial High Frequency Data Using Point Processes. In: Mikosch, T., Kreiß, J.-P., Davis, R. A., Andersen, T. G. (Eds.), *Handbook of Financial Time Series*. Springer, pp. 953–979.
- Bicchetti, D., Maystre, N., 2012. The synchronized and long-lasting structural change on commodity markets: evidence from high frequency data. UNCTAD Discussion Paper, No. 208, 1–31.
- Black, F., 1986. Noise. *The Journal of Finance* 41 (3), 529–543.
- Bouchaud, J.-P., Farmer, J. D., Lillo, F., 2009. How markets slowly digest changes in supply and demand. In: *Handbook of Financial Markets: Dynamics and Evolution*. North Holland, Amsterdam, pp. 57–160.

- Bowsher, C. G., 2002. Modelling Security Market Events in Continuous Time: Intensity Based, Multivariate Point Process Models. Nuffield College Economics Discussion Papers,, 1–55.
- Bowsher, C. G., 2007. Modelling security market events in continuous time: Intensity based, multivariate point process models. *Journal of Econometrics* 141 (2), 876–912.
- Calvet, L. E., Fisher, A. J., 2008. *Multifractal Volatility: Theory, Forecasting, and Pricing*. Elsevier - Academic Press.
- CFTC, 2011. Large Trader Net Position Changes. Tech. rep.
- Chavez-Demoulin, V., Davison, A. C., McNeil, A. J., 2005. Estimating value-at-risk: a point process approach. *Quantitative Finance* 5 (2), 227–234.
- Cont, R., 2011. Statistical Modeling of High Frequency Financial Data: Facts, Models and Challenges. *IEEE Signal Processing* 28 (5), 16–25.
- Daley, D. J., Vere-Jones, D., 2008. *An Introduction to the Theory of Point Processes. Volume II: General theory and structure*, 2nd Edition. Vol. 2 of *Probability and Its Applications*. Springer Verlag.
- Duhigg, C., 2009. Traders Profit With Computers Set at High Speed.
URL <http://www.nytimes.com/2009/07/24/business/24trading.html>
- Eisler, Z., Kertész, J., 2006. Size matters: some stylized facts of the stock market revisited. *The European Physical Journal B* 51 (1), 145–154.
- Embrechts, P., Liniger, T., Lu, L., 2011. Multivariate Hawkes Processes: an Application to Financial Data. *J. Appl. Probab.* 48A, 367–378.
- Engle, R. F., 2000. The Econometrics of Ultra-High-Frequency Data. *Econometrica: Journal of the Econometric Society* 68 (1), 1–22.
- Engle, R. F., Russell, J. R., 1997. Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model. *Journal of Empirical Finance* 4 (2-3), 187–212.

- Engle, R. F., Russell, J. R., 1998. Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data. *Econometrica: Journal of the Econometric Society* 66 (5), 1127–1162.
- Errais, E., Giesecke, K., Goldberg, L. R., 2010. Affine Point Processes and Portfolio Credit Risk. *SIAM Journal on Financial Mathematics* 1 (1), 642.
- Fama, E. F., 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance* 25 (2), 383–417.
- Fama, E. F., 1991. Efficient capital markets: II. *Journal of Finance* 46 (5), 1575–1617.
- Filimonov, V., Sornette, D., 2012. Quantifying reflexivity in financial markets: Toward a prediction of flash crashes. *Physical Review E* 85 (5), 056108.
- Garman, M. B., Klass, M. J., 1980. On the Estimation of Security Price Volatilities from Historical Data. *The Journal of Business* 53 (1), 67–78.
- Hardiman, S. J., Bercot, N., Bouchaud, J.-P., 2013. Critical reflexivity in financial markets: a Hawkes process analysis.
URL <http://arxiv.org/abs/1302.1405>
- Harris, T. E., 2002. *The Theory of Branching Processes*. Dover Phoenix Editions.
- Hasbrouck, J., 1991. Measuring the Information Content of Stock Trades. *The Journal of Finance* 46 (1), 179–207.
- Hawkes, A. G., 1971. Point Spectra of Some Mutually Exciting Point Processes. *Journal of the Royal Statistical Society. Series B (Methodological)* 33 (3), 438–443.
- Helmstetter, A., Sornette, D., 2003. Importance of direct and indirect triggered seismicity in the ETAS model of seismicity. *Geophysical Research Letters* 30 (11), 1576.
- Hewlett, P., 2006. Clustering of order arrivals, price impact and trade path optimisation. In *Workshop on Financial Modeling with Jump processes*, Ecole Polytechnique.

- Iati, R., 2009. The Real Story of Trading Software Espionage.
URL <http://www.advancedtrading.com/algorithms/the-real-story-of-trading-software-espio/218401501>
- Ide, K., Sornette, D., 2002. Oscillatory finite-time singularities in finance, population and rupture. *Physica A: Statistical Mechanics and its Applications* 307 (1-2), 63–106.
- Irwin, S. H., Sanders, D. R., 2012. Financialization and Structural Change in Commodity Futures Markets. *Journal of Agricultural and Applied Economics* 44 (3), 371–396.
- Ivanov, P. C. C., Yuen, A., Podobnik, B., Lee, Y., 2004. Common scaling patterns in intertrade times of U. S. stocks. *Physical Review E* 69 (5), 056107.
- Jiang, Z.-m., 2008. Reflections on energy issues in China. *Journal of Shanghai Jiaotong University (Science)* 13 (3), 257–274.
- Jiang, Z.-Q., Chen, W., Zhou, W.-X., 2009. Detrended fluctuation analysis of intertrade durations. *Physica A: Statistical Mechanics and its Applications* 388 (4), 433–440.
- Johansen, A., Ledoit, O., Sornette, D., 2000. Crashes as Critical Points. *International Journal of Theoretical and Applied Finance* 3 (2), 219–255.
- Johansen, A., Sornette, D., 1999. Critical Crashes. *Risk* 12 (1), 91–94.
- Johansen, A., Sornette, D., 2000. The Nasdaq crash of April 2000: Yet another example of log-periodicity in a speculative bubble ending in a crash. *The European Physical Journal B* 17 (2), 319–328.
- Kindleberger, C. P., Aliber, R., 2005. *Manias, Panics, and Crashes: A History of Financial Crises*. Wiley.
- Kyle, A. S., Obizhaeva, A. A., 2012. Large Bets and Stock Market Crashes.
- Large, J., 2007. Measuring the resiliency of an electronic limit order book. *Journal of Financial Markets* 10 (1), 1–25.

- Lux, T., 2008. The Markov-Switching Multifractal Model of Asset Returns: GMM Estimation and Linear Forecasting of Volatility. *Journal of Business and Economic Statistics* 26 (2), 194–210.
- Marsan, D., Lengline, O., 2008. Extending Earthquakes’ Reach Through Cascading. *Science* 319 (5866), 1076–1079.
- Meyer, G., July 5, 2011. CFTC data reveal day traders’ role in volatile oil markets.
 URL <http://www.ft.com/intl/cms/s/0/b29b2b1e-a743-11e0-b6d4-00144feabdc0.html#axzz288MmVAIz>
- Muzy, J.-F., Sornette, D., Delour, J., Arneodo, A., 2001. Multifractal returns and Hierarchical Portfolio Theory. *Quantitative Finance* 1 (1), 131–148.
- Ogata, Y., 1978. The asymptotic behaviour of maximum likelihood estimators for stationary point processes. *Annals of the Institute of Statistical Mathematics* 30 (1), 243–261.
- Ogata, Y., 1988. Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association* 83 (401), 9–27.
- Oswiecimka, P., Kwapien, J., Drozd, S., 2005. Multifractality in the stock market: price increments versus waiting times. *Physica A: Statistical Mechanics and its Applications* 347, 626–638.
- Ozaki, T., 1979. Maximum likelihood estimation of Hawkes’ self-exciting point processes. *Annals of the Institute of Statistical Mathematics* 31 (1), 145–155.
- Papangelou, F., 1972. Integrability of Expected Increments of Point Processes and a Related Random Change of Scale. *Transactions of the American Mathematical Society* 165, 483–506.
- Perelló, J., Masoliver, J., Kasprzak, A., Kutner, R., 2008. Model for interevent times with long tails and multifractality in human communications: An application to financial trading. *Physical Review E* 78 (3), 036108.

- Politi, M., Scalas, E., 2008. Fitting the empirical distribution of intertrade durations. *Physica A: Statistical Mechanics and its Applications* 387 (8-9), 2025–2034.
- Samuelson, P. A., 1965. Proof That Properly Anticipated Prices Fluctuate Randomly. *Industrial Management Review* 6, 41–49.
- Scheffer, M., 2009. *Critical Transitions in Nature and Society*. Princeton Studies in Complexity.
- Sheppard, D., March 3, 2011. NYMEX oil trade 45 percent computer-driven. URL <http://www.reuters.com/article/2011/03/03/us-finance-summit-nymex-volume-idUSTRE7225RV20110303>
- Sornette, D., 2003. *Why Stock Markets Crash: Critical Events in Complex Financial Systems*. Princeton University Press.
- Sornette, D., 2006. *Critical Phenomena in Natural Sciences. Chaos, Fractals, Selforganization and Disorder: Concepts and Tools*. Springer Series in Synergetics. Springer Series in Synergetics.
- Sornette, D., Johansen, A., 1998. A hierarchical model of financial crashes. *Physica A: Statistical Mechanics and its Applications* 261 (3-4), 581–598.
- Sornette, D., Johansen, A., 2001. Significance of log-periodic precursors to financial crashes. *Quantitative Finance* 1 (4), 452–471.
- Sornette, D., Malevergne, Y., Muzy, J.-F., 2003. What causes crashes? *Risk* 16 (2), 67–71.
- Sornette, D., Utkin, S., 2009. Limits of declustering methods for disentangling exogenous from endogenous events in time series with foreshocks, main shocks, and aftershocks. *Physical Review E* 79 (6), 061110.
- Sornette, D., Woodard, R., 2010. Financial bubbles, real estate bubbles, derivative bubbles, and the financial and economic crisis. Proceedings of APFA7 (Applications of Physics in Financial Analysis), Conference series entitled Applications of Physics in Financial Analysis focuses on the analysis of large-scale economic data.

- Sornette, D., Woodard, R., Zhou, W.-X., 2009. The 2006–2008 oil bubble: Evidence of speculation, and prediction. *Physica A: Statistical Mechanics and its Applications* 388 (8), 1571–1576.
- Soros, G., 1987. *The Alchemy of Finance: Reading the Mind of the Market*. John Wiley & Sons, NY.
- Stoll, H. R., Whaley, R. E., 2010. Commodity Index Investing and Commodity Futures Prices. *Journal of Applied Finance* 20 (1), 7–46.
- Stoll, H. R., Whaley, R. E., 2011. Commodity Index Investing: Speculation or Diversification? *The Journal of Alternative Investments* 14 (1), 50–60.
- Sussman, A., Tabb, L., Iati, R., 2009. *US Equity High Frequency Trading: Strategies, Sizing and Market Structure*. Tech. rep., TABB group.
- Tang, K., Xiong, W., 2010. Index Investment and Financialization of Commodities. NBER Working Paper No. 16385.
- Toke, I. M., 2011. “Market making” in an order book model and its impact on the spread. In: *Econophysics of Order-Driven Markets*. Springer Verlag, pp. 49–64.
- UNCTAD, 2009. *Trade and Development Report 2009, Chapter II: The Financialization of Commodity Markets*. United Nations publications.
- UNCTAD, 2011. *Price Formation In Financialized Commodity Markets: The Role of Information*. United Nations publications, New York and Geneva.
- Vere-Jones, D., 1970. Stochastic Models for Earthquake Occurrence. *Journal of the Royal Statistical Society. Series B (Methodological)* 32 (1), 1–62.
- Vere-Jones, D., Ozaki, T., 1982. Some examples of statistical estimation applied to earthquake data I. Cyclic Poisson and self-exciting models. *Annals of the Institute of Statistical Mathematics* 34 (1), 189–207.
- Zhuang, J., Ogata, Y., Vere-Jones, D., 2002. Stochastic declustering of space-time earthquake occurrences. *Journal of the American Statistical Association* 97 (458), 369–380.

Table 1: Description of the selected instruments.

Abbreviations and TRTH RIC	Specification	Exchange and trading platform	Contract month	Introduction of electronic trading	Hours of active trading
Brent Crude (LCOc1)	1,000 barrels of light sweet crude oil	ICE Europe / ICE electronic platform	Every month	April 7, 2005	BST 15:15–19:45 ^a ; 14:00–19:45 ^b
WTI (CLc1)	1,000 barrels of light sweet crude oil	NYMEX / CME Globex	Every month	September 4, 2006	EST 10:00–14:45 ^a ; 9:00–14:45 ^b
Soybean (Sc1)	5,000 bushels (~136 metric tons)	CBOT / CME Globex	January, March, May, July, August, September, November	August 1, 2006	CDT 9:45–13:30
Sugar #11 (SBc1)	112,000 pounds	ICE US / ICE electronic platform	March, May, July, October	January 12, 2007 ^c	EST 8:15–13:45
Corn (Cc1)	5,000 bushels (~127 metric tons)	CBOT / CME Globex	March, May, July, September, December	August 1, 2006	CDT 9:45–13:30
Wheat (Wc1)	5,000 bushels (~136 metric tons)	CBOT / CME Globex	March, May, July, September, December	August 1, 2006	CDT 9:30–13:30
Sugar (LSUc1)	50 metric tons	LIFFE / NYSE Euronext	March, May, August, October, December	November 27, 2000	BST 9:30–17:30 ^d ; 8:30–17:30 ^e ;
E-mini S&P 500 (ESc1)	50 x E-mini S&P 500 futures price	CME / CME Globex	March, June, September, December	September 9, 1997	EST 9:30–16:15;

^abefore January 22, 2007

^bafter January 22, 2007

^cHowever before the March 2, 2008 data was disaggregated into RICs “SBc1” and “1SBc1” for pit and electronic trading and real time bid, ask, volume, and settlement values are not provided due to feed limitations.

^dbefore June 29, 2009

^eafter June 29, 2009

Table 2: Number of transactions, annual volume (in contracts) and volume per transaction (VPT): average (A), median (M) and 90%-quantile (Q₉₀) of analyzed contracts.

Year	Transactions	Volume	VPT			Transactions	Volume	VPT			
			A	M	Q ₉₀			A	M	Q ₉₀	
Brent Crude (Europe)						WTI (US)					
2005	2'266'953	12'324'431	5.4	1	10	919'941	24'431'479	26.6	2	13	
2006	5'723'522	17'543'910	3.1	1	5	2'468'946	29'541'698	12.0	2	8	
2007	8'619'436	22'091'574	2.6	1	3	11'960'866	58'268'584	4.9	1	6	
2008	13'413'832	26'408'342	2.0	1	2	21'429'745	66'766'312	3.1	1	4	
2009	12'789'309	28'241'439	2.2	1	3	21'104'592	66'833'089	3.2	1	4	
2010	17'690'209	38'581'454	2.2	1	3	31'570'311	79'334'457	2.5	1	3	
2011	25'033'310	46'720'379	1.9	1	3	41'855'040	78'088'015	1.9	1	2	
2012*	18'875'419	36'397'876	1.9	1	3	27'420'055	47'640'155	1.7	1	2	
Soybean (US)						Sugar #11 (US)					
2006	437'313	7'389'376	16.9	2	10						
2007	1'512'818	11'886'079	7.9	2	10	853'963	11'082'111	13.0	3	14	
2008	3'218'183	13'443'592	4.2	1	5	2'884'089	13'010'845	4.5	2	10	
2009	2'870'535	13'365'613	4.7	1	6	2'167'801	12'424'883	5.7	1	7	
2010	5'522'405	13'385'860	2.4	1	2	4'572'232	12'767'545	2.8	1	4	
2011	7'023'025	16'435'216	2.3	1	2	4'513'119	10'867'352	2.4	1	4	
2012*	5'043'826	11'191'303	2.2	1	2	3'244'271	8'864'245	2.7	1	5	
Corn (US)						Wheat (US)					
2005	919'941	24'431'479	26.6	2	13	116'059	4'540'024	39.1	2	10	
2006	2'468'946	29'541'698	12.0	2	8	306'472	6'296'176	20.5	2	10	
2007	11'960'866	58'268'584	4.9	1	6	1'126'338	7'897'908	7.0	2	10	
2008	21'429'745	66'766'312	3.1	1	4	2'060'348	8'120'508	3.9	1	5	
2009	21'104'592	66'833'089	3.2	1	4	1'765'353	8'123'123	4.6	1	6	
2010	31'570'311	79'334'457	2.5	1	3	3'887'485	9'447'008	2.4	1	3	
2011	41'855'040	78'088'015	1.9	1	2	5'099'530	10'128'749	2.0	1	2	
2012*	27'420'055	47'640'155	1.7	1	2	3'677'335	8'582'026	2.3	1	2	
Sugar (Europe)						E-mini S&P 500					
2005	58'524	849'928	14.5	8	30	11'439'420	183'667'226	16.1	2	35	
2006	82'688	891'134	10.8	5	23	11'095'507	223'402'685	20.1	2	48	
2007	148'815	1'217'541	8.2	4	20	22'183'920	362'881'400	16.4	2	31	
2008	158'151	925'481	5.9	3	12	49'488'715	551'544'452	11.1	2	23	
2009	294'445	919'343	3.1	1	6	41'655'339	492'581'685	11.8	2	21	
2010	400'850	977'312	2.4	1	5	107'143'664	497'545'699	4.6	1	9	
2011	485'522	870'938	1.8	1	3	120'700'428	540'010'834	4.5	1	9	
2012*	350'957	686'958	2.0	1	3	72'728'681	316'597'629	4.4	1	9	

* Year-To-Date: Datasets of Brent Crude, WTI, Soybean, Sugar #11 (US), Corn and E-mini S&P 500 contained data until September 30, 2012. Datasets of Wheat (US) and Sugar (Europe) contained data until May 30, 2012.

Table 3: Average and median uncertainty of the timestamps of events resulting from the nature of the FAST/FIX feed. Dash lines (—) correspond to the time periods before the introducing of electronic trading for the given contract (see table 1).

(A) Average uncertainty (in milliseconds)

Contract	2005	2006	2007	2008	2009	2010	2011	2012
Brent (EU)	332	222	105	98	107	115	165	167
WTI (US)	—	326	208	133	144	137	141	110
Soybean (US)	—	267	240	174	192	146	125	141
Sugar #11 (US)	—	—	—	235	199	183	243	242
Corn (US)	—	268	267	186	207	164	142	144
Wheat (US)	—	287	281	211	213	146	147	141
Sugar (EU)	309	272	303	344	230	212	200	185
E-mini S&P 500	173	195	168	112	129	87	92	103

(B) Median uncertainty (in milliseconds)

Contract	2005	2006	2007	2008	2009	2010	2011	2012
Brent (EU)	227	118	35	26	24	30	65	68
WTI (US)	—	199	80	62	61	62	59	22
Soybean (US)	—	149	130	71	77	32	22	23
Sugar #11 (US)	—	—	—	112	58	43	127	135
Corn (US)	—	151	174	75	106	45	32	26
Wheat (US)	—	174	179	91	86	29	30	22
Sugar (EU)	223	197	190	245	119	85	84	69
E-mini S&P 500	127	121	79	51	60	31	32	41

Table 4: Fraction of total calibrations (per year) of the Hawkes model that could be rejected with 95% confidence on the basis of the Kolmogorov-Smirnov test (see text). Dash lines (—) correspond to the years before the introduction of electronic trading for the given contract (see table 1). Stars (***) denote the years when the corresponding reliability of timestamps (Δ) is not applicable (see text).

(A) $\Delta = 50$ milliseconds

Contract	Fraction of rejected estimates per year								Total
	2005	2006	2007	2008	2009	2010	2011	2012	
Brent (EU)	***	***	32.60%	41.28%	31.35%	31.40%	***	***	34.16%
WTI (US)	—	***	***	***	***	***	***	34.26%	34.26%
Soybean (US)	—	***	***	***	***	33.20%	28.17%	31.03%	30.75%
Sugar #11 (US)	—	—	—	***	***	34.53%	***	***	34.53%
Corn (US)	—	***	***	***	***	33.39%	48.87%	29.92%	38.45%
Wheat (US)	—	***	***	***	***	33.50%	32.78%	28.42%	31.93%
Sugar (EU)	***	***	***	***	***	***	***	***	***
E-mini S&P 500	***	***	***	***	***	45.52%	60.78%	36.97%	48.97%
Total	***	***	32.60%	41.28%	31.35%	35.96%	45.64%	32.96%	37.29%

(B) $\Delta = 100$ milliseconds

Contract	Fraction of rejected estimates per year								Total
	2005	2006	2007	2008	2009	2010	2011	2012	
Brent (EU)	***	***	15.11%	18.94%	20.52%	20.06%	14.71%	13.17%	17.24%
WTI (US)	—	***	10.09%	6.81%	7.57%	17.63%	9.91%	22.41%	11.96%
Soybean (US)	—	***	***	13.99%	25.73%	23.77%	19.65%	21.97%	21.01%
Sugar #11 (US)	—	—	—	***	16.50%	25.69%	***	***	21.48%
Corn (US)	—	***	***	28.29%	***	22.85%	35.41%	23.93%	28.16%
Wheat (US)	—	***	***	26.46%	30.28%	25.72%	24.86%	20.67%	26.24%
Sugar (EU)	***	***	***	***	***	36.10%	38.10%	42.70%	38.08%
E-mini S&P 500	***	***	9.81%	9.11%	9.30%	39.56%	51.66%	27.99%	24.75%
Total	***	***	11.62%	15.62%	16.85%	27.37%	29.01%	24.50%	22.13%

Table 4: (continued).

(C) $\Delta = 200$ milliseconds

Contract	Fraction of rejected estimates per year								Total
	2005	2006	2007	2008	2009	2010	2011	2012	
Brent (EU)	***	24.28%	6.27%	12.06%	14.97%	13.00%	5.33%	5.19%	11.38%
WTI (US)	—	6.42%	3.75%	3.11%	3.17%	8.65%	5.87%	14.24%	6.14%
Soybean (US)	—	3.38%	8.30%	8.77%	13.56%	14.09%	11.65%	12.13%	10.98%
Sugar #11 (US)	—	—	—	12.04%	9.86%	16.77%	29.98%	26.50%	19.20%
Corn (US)	—	5.72%	7.36%	16.04%	10.24%	11.82%	18.24%	15.51%	13.14%
Wheat (US)	—	3.27%	8.35%	15.98%	16.59%	15.79%	14.51%	11.80%	13.42%
Sugar (EU)	***	25.72%	13.90%	***	13.38%	22.43%	26.07%	32.98%	23.60%
E-mini S&P 500	1.15%	9.97%	7.05%	4.12%	4.54%	30.66%	34.34%	19.53%	14.36%
Total	1.15%	13.17%	6.70%	9.27%	10.26%	17.55%	18.85%	17.79%	13.58%

(D) $\Delta = 300$ milliseconds

Contract	Fraction of rejected estimates per year								Total
	2005	2006	2007	2008	2009	2010	2011	2012	
Brent (EU)	11.11%	16.42%	3.90%	8.30%	11.14%	9.10%	3.03%	3.41%	8.02%
WTI (US)	—	3.48%	2.45%	1.90%	1.98%	5.85%	3.60%	9.70%	4.01%
Soybean (US)	—	2.61%	5.90%	6.27%	7.58%	9.18%	7.30%	7.64%	7.06%
Sugar #11 (US)	—	—	—	7.10%	6.75%	10.71%	19.13%	18.74%	12.52%
WTI (US)	—	3.48%	2.45%	1.90%	1.98%	5.85%	3.60%	9.70%	4.01%
Wheat (US)	—	2.29%	5.59%	10.01%	9.47%	10.05%	8.65%	7.40%	8.27%
Sugar (EU)	16.92%	24.63%	10.54%	7.84%	9.48%	14.12%	18.73%	28.47%	16.33%
E-mini S&P 500	0.91%	8.76%	6.67%	2.47%	2.67%	23.64%	32.68%	14.10%	12.38%
Total	5.99%	9.37%	4.45%	5.16%	6.03%	11.46%	13.06%	12.90%	8.99%

(E) $\Delta = 1$ second

Contract	Fraction of rejected estimates per year								Total
	2005	2006	2007	2008	2009	2010	2011	2012	
Brent (EU)	2.44%	2.03%	0.85%	1.33%	2.18%	1.21%	0.58%	0.77%	1.36%
WTI (US)	—	0.85%	0.43%	0.33%	0.30%	0.71%	0.45%	1.18%	0.56%
Soybean (US)	—	0.93%	1.41%	1.72%	1.45%	1.94%	1.61%	1.31%	1.55%
Sugar #11 (US)	—	—	—	1.29%	1.34%	1.68%	2.13%	2.95%	1.84%
Corn (US)	—	1.09%	1.42%	2.19%	1.42%	1.97%	2.41%	2.70%	1.99%
Wheat (US)	—	0.61%	1.06%	1.14%	1.54%	1.84%	1.37%	1.18%	1.33%
Sugar (EU)	10.60%	15.14%	2.72%	1.65%	2.12%	1.75%	4.63%	16.84%	4.57%
E-mini S&P 500	0.28%	1.84%	2.19%	0.49%	0.32%	6.03%	7.98%	3.97%	3.10%
Total	1.41%	1.93%	1.25%	1.15%	1.29%	2.23%	3.00%	3.81%	2.08%




Figure 1: Evolution of monthly volume (measured in number of contracts) of the Brent Crude (Europe), WTI (US), Soybean (US), Sugar #11 (US), Corn (US), Wheat (US) and Sugar (Europe) future contracts over the period 2005-2012.

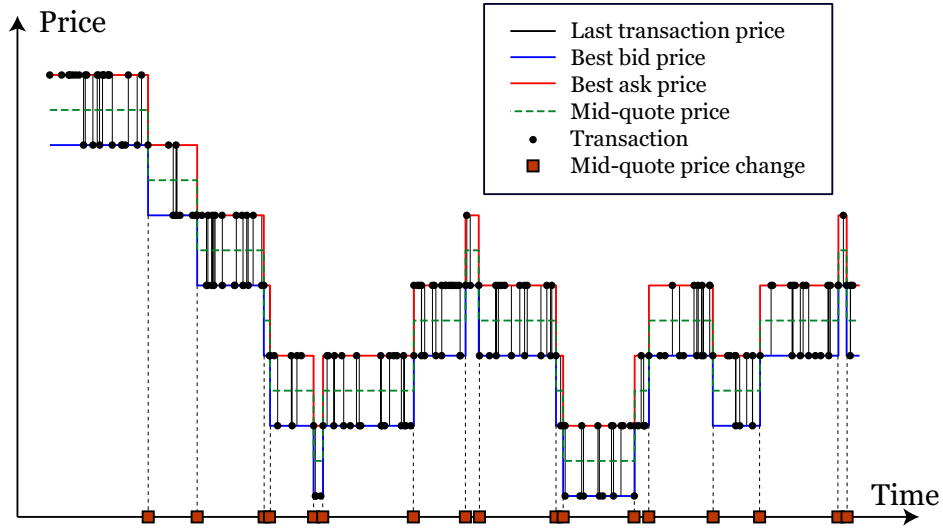


Figure 2: Illustration of the high frequency price dynamics. The black line corresponds to the last transaction price, the red and blue lines correspond to best ask and best bid prices respectively and the dashed green line corresponds to the mid-quote price. Black circles denote transactions and red squares denote timestamps of mid-quote price changes




Figure 3: Illustration of the branching structure of the Hawkes process (top) and events on the time axis (bottom). Different colors of markers correspond to different clusters. The numbers below an event denotes its order within the cluster. This picture corresponds to the branching ration $n = 0.88$.

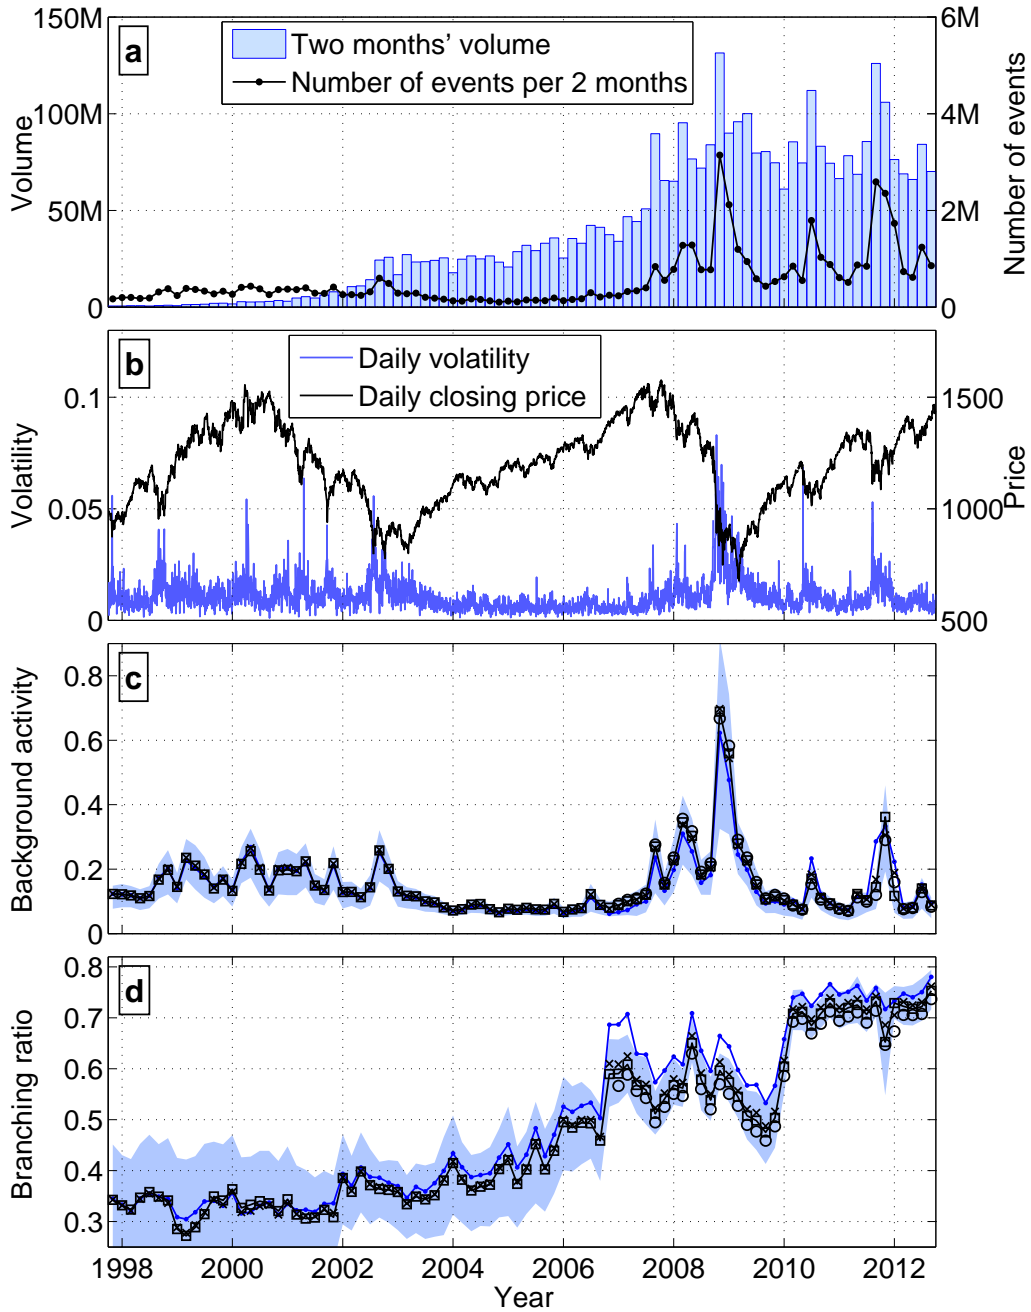




Figure 4: Dynamics of (a) volume and activity measured in number of mid-quote price changes, (b) daily closing price and daily volatility, (c) estimated background intensity ($\hat{\mu}$, see text) and (d) branching ratio ($\hat{\nu}$, see text) for the E-mini S&P 500 futures over the period 1998–2012. Each point in panels (c) and (d) represents averaged estimates over two months interval prior to the point in time windows of 10 minutes for $\Delta = 100$ msec (squares), $\Delta = 200$ msec (crosses with black line), $\Delta = 300$ msec (circles) and $\Delta = 1$ sec (dots with blue line). The shaded area gives the 25%–75% quantile range obtained with the same two months estimates for $\Delta = 200$ msec. In the analysis we have considered only estimates performed within hours of active trading (see table 1).



(a) Brent Crude (Europe)



(b) WTI (US)

Figure 5: (i) Daily closing price and daily volatility estimated with the Garman&Klass estimator and (ii) estimation of the branching ratio (\hat{n}) of the flow of mid-quote price changes of the (a) Brent Crude and (b) WTI futures on light sweet oil. Each point at a given time t in the panels showing the branching ratio represents an averaged over one month prior totime t of windows of 10 minutes for $\Delta = 100$ msec (squares), $\Delta = 200$ msec (crosses with black line), $\Delta = 300$ msec (circles) and $\Delta = 1$ sec (dots with blue line). The shaded area indicates the 25%–75% quantile range obtained with the same one month estimates for $\Delta = 200$ msec. In the analysis we have considered only estimates performed within hours of active trading (see table 1).






Figure 6: Price time series of NYMEX Light Sweet Crude (front month contract) and simple log-periodic power law (LPPL) fits (see Sornette et al. (2009) for details). The shaded box shows the 80% confidence interval of the forecast performed at the time indicated by the vertical dashed line.




(a) Sugar (Europe)

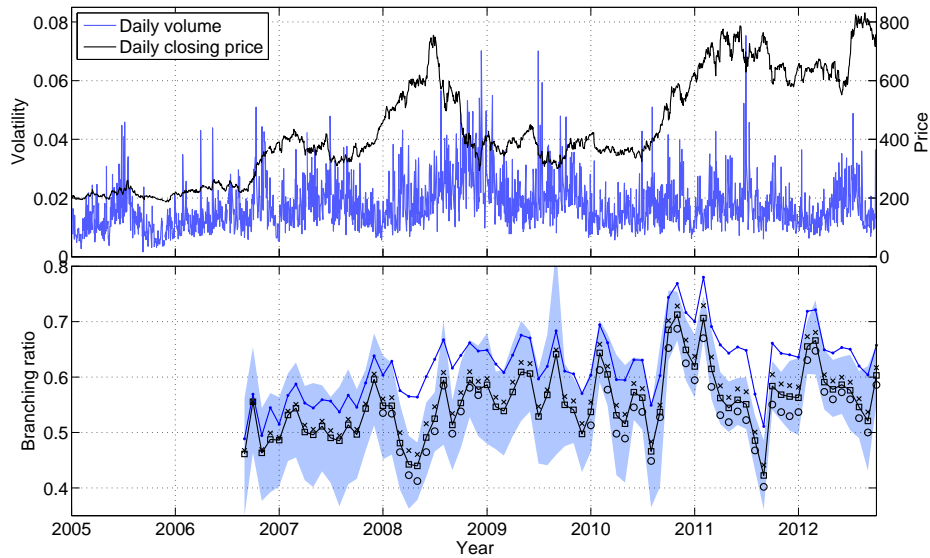


(b) Sugar #11 (US)

Figure 7: (i) Daily closing price and daily volatility estimated with the Garman&Klass estimator and (ii) estimation of the branching ratio (\hat{n}) of the flow of mid-quote price changes of the (a) European Sugar and (b) Sugar #11 (US) futures. Each point at a given time t in the panels showing the branching ratio represents an averaged over one month prior to time t of windows of 10 minutes for $\Delta = 100$ msec (squares), $\Delta = 200$ msec (crosses with black line), $\Delta = 300$ msec (circles) and $\Delta = 1$ sec (dots with blue line). The shaded area delineates the 25%–75% quantile range obtained with the same one month estimates for $\Delta = 200$ msec. In the analysis we have considered only estimates performed within hours of active trading (see table 1).




(a) Soybean (US)



(b) Corn (US)

Figure 8: (i) Daily closing price and daily volatility estimated with the Garman&Klass estimator and (ii) estimation of the branching ratio (\hat{n}) of the flow of mid-quote price changes of the (a) Soybean and (b) Corn futures. Each point at a given time t at the plot of branching ratio represents an averaged over one month interval prior to time t of windows of 10 minutes for $\Delta = 100$ msec (squares), $\Delta = 200$ msec (crosses with black line), $\Delta = 300$ msec (circles) and $\Delta = 1$ sec (dots with blue line). The shaded area corresponds to 25%–75% quantile range obtained with the same 2 months of estimates for $\Delta = 200$ msec. In the analysis we have considered only estimates performed within hours of active trading (see table 1).



(a) Wheat (US)

Figure 9: (i) Daily closing price and daily volatility estimated with the Garman&Klass estimator and (ii) estimation of the branching ratio (\hat{n}) of the flow of mid-quote price changes of the Wheat futures. Each point at a given time t in the panels showing the branching ratio represents an averaged over one month prior to time t of windows of 10 minutes for $\Delta = 100$ msec (squares), $\Delta = 200$ msec (crosses with black line), $\Delta = 300$ msec (circles) and $\Delta = 1$ sec (dots with blue line). The shaded area shows the 25%–75% quantile range obtained with the same one month estimates for $\Delta = 200$ msec. In the analysis we have considered only estimates performed within hours of active trading (see table 1).




Figure 10: Estimation of the branching ratio (\hat{n}) of the flow of randomly redistributed mid-quote price changes (see text) of the Brent Crude (Europe), WTI (US), Soybean (US), Sugar #11 (US), Corn (US), Wheat (US) and Sugar (Europe) future contracts. Each point at a given time t in the panel showing the branching ratio represents an average over one month prior to time t of windows of 10 minutes for $\Delta = 100$ msec (squares), $\Delta = 200$ msec (crosses with black line) and $\Delta = 300$ msec (circles). The shaded area gives the 25%–75% quantile range obtained with the same one month estimates for $\Delta = 200$ msec. In the analysis we have considered only 56 estimates performed within hours of active trading (see table 1).