V-shapes*

Maria Flora‡ Roberto Renò§

September 17, 2020

Abstract

An insidious form of market inefficiency, by which prices lose their informativeness and wealth is distributed arbitrarily, translates into V-shapes, that is sudden changes of the sign of the price drift. We use this insight to develop a new tool for the detection of reverting drift, the V-statistic. We apply this tool to (i) quantify the extent of this kind of market inefficiency in the U.S. stock market during the Covid-19 pandemic; and (ii) show the harmful consequences of V-shapes on financial stability by estimating the huge loss suffered by Italian taxpayers (0.45B euros) in May 2018, when a transient crash hit the secondary bond market during a Treasury auction.

JEL classification: G14, G12, C58;

Keywords: Market inefficiency, financial fragility, price drift, financial stability, flash crash.

*We thank Giacomo Bormetti, Antoine Bouveret, Kim Christensen, Fulvio Corsi, Gerardo Ferrara, Aleksey Kolokolov, Michele Manna, Antonio Mannolini, Gaetano Marseglia, Claudio Pacati, Alfonso Puorro, Giuseppe Trapani, the participants at the XXI Workshop on Quantitative Finance (Napoli, January 29-31 2020), and the participants at the seminar at the Bank of Italy (2019) for comments and discussions. We thank Borsa Italiana S.p.A. for providing the transaction data. The software used to compute the V-statistic is available upon request. Any use, including commercial, of this copyrighted software is prohibited without the explicit consent of all Authors. This is a working paper, so feedback is welcome. All errors and omissions are sole responsibility of the Authors only.

‡University of Verona, Department of Economics, e-mail: MARIA.FLORA@UNIVR.IT

§University of Verona, Department of Economics, e-mail: ROBERTO.RENO@UNIVR.IT

A downward-pointing "V", which is similar to the geometric configuration of the face in angry expressions, is perceived as threatening.
Larson, Aronoff and Stearns – Emotion (2007), Vol. 7, No. 3

1 Introduction

Inefficient markets and market fragility are a threat to investors. However, it is hard to determine when, and for how long, market prices are away from fundamentals. This paper fills this gap by introducing the concept of *V-shapes* in traded prices, and by designing a novel econometric test for their detection, the *V-statistic*, which can reliably detect such instances. V-shapes, defined as a sudden discontinuity in the sign of the price drift, are observed when the traded price deviates from fundamentals and retraces back. We clarify their relation with market efficiency and market liquidity by the light of standard economic models of price formation. We apply our new test to study the extent of market fragility in the U.S. during the Covid-19 pandemic. We further analyze a case study about the Italian sovereign bond market in May, 2018. This market was hit by a shock due to macro-political news (new Italian cabinet formation) just before a large Treasury auction. We show that the secondary market was inefficient during the shock, and we quantify the harsh consequences for Italian taxpayers.

Our findings are connected to several literatures. They directly relate to *flash crashes*, a peculiar form of market disruption that recently attracted the attention of investors, academics and regulators. In the July 2019 Financial Stability Report, the Bank of England defines a flash crash as a "*large and rapid change in the price of an asset that does not coincide with – or in some cases substantially overshoots – changes in economic fundamentals, before typically retracing those moves shortly afterwards"* (Bank of England, 2019). Thus, flash crashes are a specific source of V-shapes. The most notorious example of a flash crash is that of May 6, 2010 in the U.S. stock market (see, e.g., Madhavan, 2012; Kirilenko et al., 2017; Menkveld and Yueshen, 2019). This event was triggered by a huge, non-fundamental selling trade in the E-mini futures market (CFTC and SEC, 2010). Shocking occurrences like this shed light on a market vulnerability that now appears to be endemic to most financial markets, and to occur increasingly often over time (Christensen, Oomen, and Renò, 2017; Golub, Keane, and Poon, 2017; Laly and Petitjean, 2020). The V-statistic constitutes a further step toward rigorous detection of genuine flash episodes.

V-shapes also represent a new form of tail risk (Bollerslev and Todorov, 2011; Kelly and Jiang, 2014; Weller, 2019) in which the change of the traded price does not completely reflect the change in the fundamental price and in some instances, like May 6, 2010, it does not reflect it at all. We

differ from existing measures of tail risk because our V-statistic can reveal the inefficient part of the price change, that is the overshooting unrelated to fundamentals. In our empirical analysis, we provide evidence of several significant V-shapes per year, with durations up to 30 minutes for the U.S. stock index, and up to two days for Italian bonds. Our analysis thus makes two important points: first, market inefficiency can last longer than the typical duration of a flash crash (Biais and Foucault, 2014 use the terminology "mini flash crashes"). Second, this kind of inefficiency has detrimental consequences for financial stability. Financial stability can indeed be defined as the "*ability to facilitate and enhance economic processes, manage risks, and absorb shocks*" (Schinasi, 2004). As a matter of fact, a V-shape is a prominent signature of disrupted market, a non-fundamental risk and a shock which is not absorbed immediately. A market exposed to this kind of fragility (in the Allen and Gale, 2004 sense) can be inefficient for a worryingly long time. This paper contributes to this debate by describing a prolonged V-shape with huge financial stability implications in a central market (Italian sovereign bonds), and its aftermaths in terms of inefficiency, volatility, illiquidity and wealth redistribution.

Our paper also relates, more broadly, to the vast economic theory of price formation with frictions leading to overshooting. Several explanations are indeed offered in the literature for the presence of V-shapes, typically generated by the strategic interaction of traders with different objectives in a frictional environment. Frictions considered by theoretical papers include transaction costs, asymmetric information, anticipated shocks, and market fragmentation. For example, Grossman and Miller (1988) predict a large and localized price decline (and subsequent reversal to the initial price level) in the presence of selling pressure looking for immediacy. The price decline of their model is proportional to the trade size and inversely proportional to the liquidity of the market. This simple mechanism can be exacerbated by co-existing frictions. Bernardo and Welch (2004) show that the fear of future liquidity shocks can induce traders to liquidate their positions during a run; Brunnermeier and Pedersen (2005) show that predatory buyers could follow the initial sell orders to push the price downward inefficiently in an illiquid market; Huang and Wang (2009) show that market monitoring costs can prevent agents to synchronize trades, giving rise to abnormal selling and large and inefficient price declines; Cespa and Foucault (2014) point at liquidity spillovers, showing that feedback liquidity loops can generate transient crashes; Colliard (2017) shows that flash crashes can be exacerbated by the presence of traders with superior information on liquidity; and Menkveld and Yueshen (2019) show that cross-arbitrage may break during a severe liquidity shock, and point at fragmented markets as a potential source of flash crashes. We complement this literature by providing a new reliable measure of market distress which can be used to shed empirical light on theoretical assumptions.

Our econometric strategy, which builds on Christensen, Oomen, and Renò (2017) (henceforth,

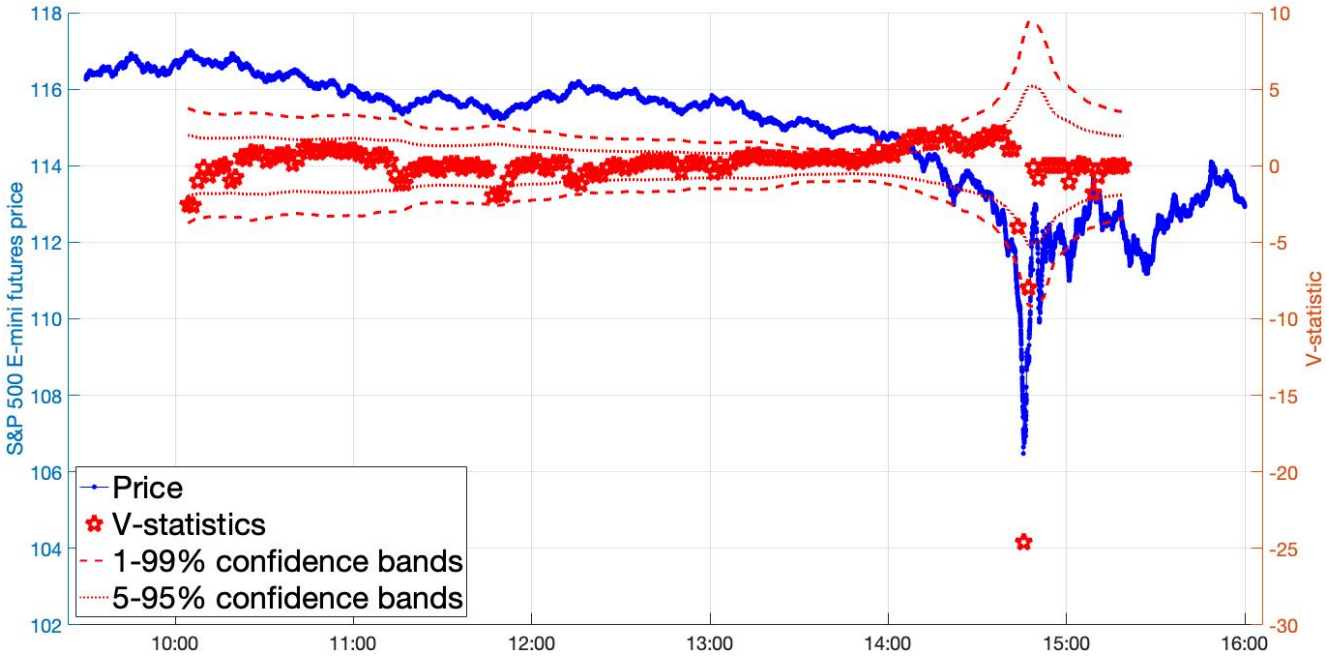

Figure 1: Blue dots: transaction prices for E-mini S&P500 futures on May 6, 2010. Red pentagrams: V-statistics computed using an exponential kernel and a bandwidth $h_n = 1000$ seconds. Confidence bands for the individual tests are computed bootstrapping the V-statistics on 50,000 simulations of an E-GARCH(1,1) process without drift fitted on observed returns. The minimum of the V-statistics is -24.6, signaling a strong V-shape. The p-value of the minimum, which we interpret as the probability that the flash crash was due to a pure volatility move, and thus that the market was not inefficient in that day, is estimated to be 0.006%.

COR), is purposely designed to capture V-shapes. The test introduced by COR is indeed meant to detect "drift bursts", that is abnormally large trends. We follow their logic that associates financial distress with drift explosions. However, the V-statistic proposed in this paper allows to disentangle the case of a rapid change to a new price level ("gradual jump", in the terminology of Barndorff-Nielsen et al., 2008) with that of a V-shaped price pattern. When significantly positive, our test still captures trends, as the COR test. When significantly negative, our test instead unambiguously captures V-shapes (for transient crashes) or Λ-shapes (for transient surges).

Figure 1 illustrates the ability of the V-statistic to detect inefficient states. We apply it to E-mini S&P500 futures prices recorded on May 6, 2010, that is the eponymous flash crash. At the bottom of the peak, right before 3 p.m., the V-statistic reaches the value -24.6, which is strongly significant when compared to the displayed confidence bands, obtained with a simulated bootstrap technique which is robust to time-varying volatility. When we account for multiple testing (as explained below), we estimate that the probability that such value is reached because of a statistical fluctuation due to volatility only is around 0.006% (the figure caption provides technical details about the implementation). There is strong consensus on the fact that the market was inefficient

in that day. The V-statistic does its job in providing a strong signal.

We apply the V-statistic to two further data sets. The first one is the set of all trades of SPY in 2019 and 2020 (up to August). The point of the analysis is to illustrate the significant presence of V-shapes in the U.S. stock index, and also to see how this was impacted by the advent of the Covid-19 pandemic. We find significant price inefficiencies which, however, last less than 30 minutes and are scattered uniformly in the sample period. This suggests that the U.S. financial market appeared quite resilient to this shock, in the sense that it incorporated all the negative news about future expected cash flows in a substantially efficient way, at least with respect to this form of market fragility.

Our second application is an event study of Italian bonds in the last week of May, 2018. We show a clear V-shape in bond prices on May 30 which, according to our econometric procedure, again deviated abnormally from an efficient statistical fluctuation. The time scale of the inefficiency ranged several hours, which can be explained by the joint effect of political turmoil (due to the change of cabinet) and the presence of a large Treasury auction held on May 30 issuing more than 6 billion euros of medium and long-term bonds. Using a regime-switching model, we estimate that, during the auctions of May 30, about 0.45 billion euros were transferred from Italian taxpayers to primary dealers. While the existence of an auction premium is well established and rationalized (Lou et al., 2013), what is shocking is the size of the wealth transfer, which is equivalent, in a single day, to the whole premium paid by the U.S. Treasury to primary dealers in one year. We also document that the V-shape was associated with increased volatility and deteriorated liquidity conditions, that persisted for several months afterwards, as predicted by the quoted theories of price formation.

The paper is organized as follows. In Section 2, we discuss how periods of financial fragility leading to potential flash crashes are better identified by looking at fast-reverting drifts instead of the usual measures typically used in the literature, based on price dispersion. In Section 3 we introduce our new test for fast-reverting trends. Section 4 contains result on simulated data meant to illustrate the size and power of the V-statistics. Section 5 contains the empirical application to SPY during the Covid-19 pandemic. In Section 6, we describe the case study on Italian bonds in 2018. Section 7 concludes.

2 Market inefficiency and V-shapes

The common practice to detect financial distress is to look at price dispersion measures, such as volatility (Christensen et al., 2014; Bates, 2019; Brogaard et al., 2018) and, in some instances,

price jumps (Calcagnile et al., 2018). In this section, we argue that sustained volatility and/or jumps are compatible with efficient markets, and that instead price V-shapes, whose definition is based on the drift component of the price process, are a much more reliable signal of market inefficiency.

To formalize our intuition, we work with the following model for the traded logarithmic price p_t:

$$\underbrace{p_t}_{\text{traded price}} = \underbrace{p_t^e}_{\text{efficient price}} + \underbrace{f_t}_{\text{frictional component}} + \underbrace{z_t}_{\text{random component}}, \tag{2.1}$$

where p_t^e is the efficient price, or fundamental price, that is the expected value of future cash flows properly adjusted for risk; f_t is a market-dependent adjustment that pollutes the efficient price because of frictions; and z_t is the random shock which deviates the traded price from the expected price, whose dynamics is given by:

$$dz_t = \sigma_t dW_t, \tag{2.2}$$

with W_t being a standard Brownian motion. The coefficient σ_t is the standard deviation of shocks to returns. In model (2.1), p_t^e, f_t and σ_t are thought to be deterministic functions, an assumption which is made only for ease of exposition here, and relaxed in the model (3.1) used for formal econometric analysis. The expected price $\overline{p}_t$ is given by

$$\overline{p}_t = p_t^e + f_t \,.$$

We assume that $\overline{p}_t$ is differentiable. The *drift* in this model is defined by:

$$\mu_t = \frac{\partial \overline{p}_t}{\partial t} \,, \tag{2.3}$$

so that Eq. (2.1) can be re-expressed as

$$dp_t = \mu_t dt + \sigma_t dW_t.$$

In Section 3 we will add a complementary form of friction, in the form of a non-differentiable shock ε_t to the traded price, which is typically called market microstructure noise (Diebold and Strasser, 2013), without, again, changing neither the logic below nor the econometric analysis.

We can now define formally a V-shape in this model.

Definition 1. *The traded price p_t has a V-shape at the time point τ if*

$$\lim_{t \to \tau^-} \mu_t < 0, \quad \textit{and} \quad \lim_{t \to \tau^+} \mu_t > 0. \tag{2.4}$$

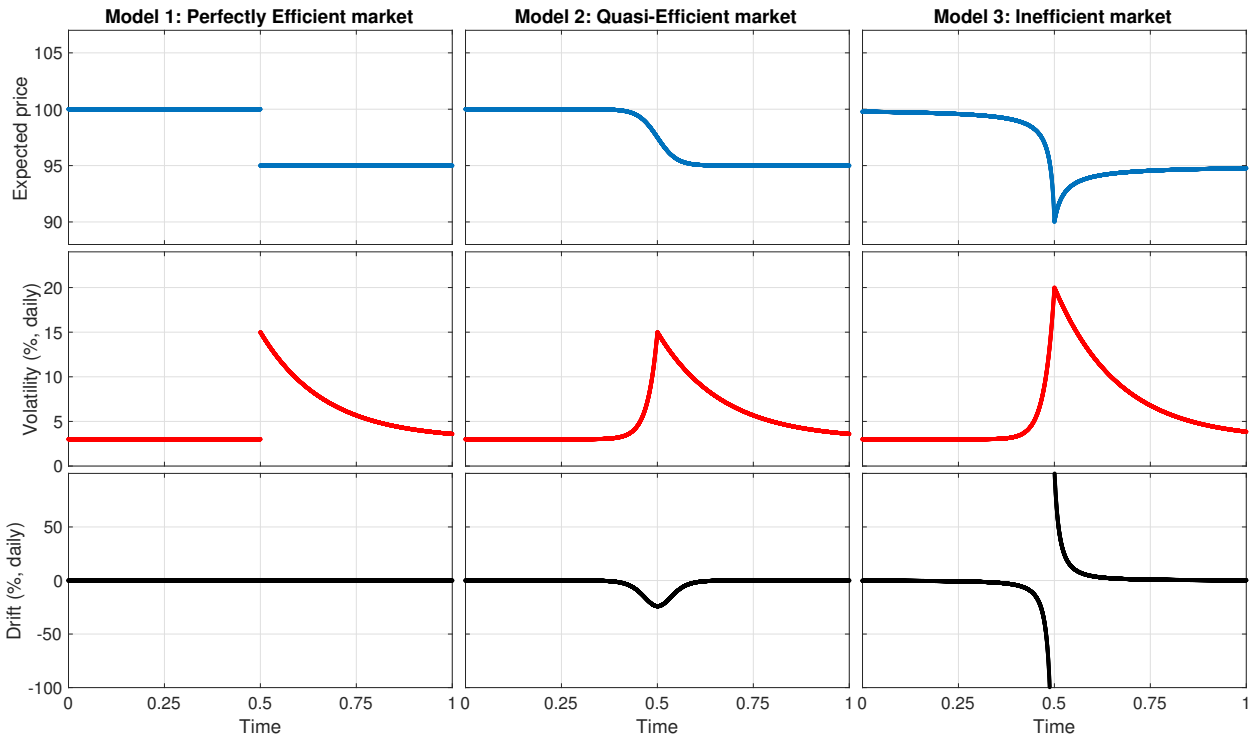

FIGURE 2: Pictorial representation of three idealized markets: a perfectly efficient one (left column), a frictional quasi-efficient one (center column) and a frictional inefficient one (right column). Each market is characterized by the expected price (top row), price volatility (center row) and price drift (bottom row). Precise formulations of the models are provided in Section 4. The markets are described here in a setting in which a fundamental piece of news arrives. Model 3 displays a V-shape.

The signs are switched for a Λ-shape. When μ_t is a stochastic process, as we will assume in Section 3, the left and right limits defining a V-shape should be interpreted as almost sure.

Figure 2 describes a situation in which unexpected fundamental information hits three different markets, one for each column. The rows of Figure 2, from top to bottom, exhibit the expected price $\overline{p}_t$, the volatility σ_t, and the drift μ_t, as defined above. Precise formulations of the three models are provided in Section 4.

The left column represents the perfectly efficient case in which the market has no frictions, that is $f_t = 0$. As the news hits, the expected price incorporates it immediately and jumps to the new fundamental level (top-left panel). Typically, the price change comes with higher uncertainty, that is price volatility instantaneously soars (center-left panel), and then slowly declines as uncertainty about the new fundamental level dissipates away. Soaring volatility and jumps are thus totally consistent with an efficient market in which the traded price and the efficient price are the same. The drift is zero (bottom left-panel), since in the presence of a price jump both the left-limit and the right-limit of the first derivative of the expected price are zero. Thus, in this idealized setting, drift is unaffected by the fundamental news.

The center column is the quasi-efficient case, in which the expected price smoothly declines to the new fundamental level instead of jumping to it instantaneously. In this case, when the new information arrives, it still gets embedded in the expected price, which frictionally ($f_t \neq 0$) decreases to the new fundamental level through a "gradual jump" (Barndorff-Nielsen et al., 2008). Accordingly, volatility smoothly spikes, and then declines. A negative drift appears whose magnitude depends on the speed of reaction of the price to the news. This drift is always negative during the expected price path to the new fundamental level.

The right column represents an inefficient market, in which the expected price overshoots the efficient price and displays a V-shape. This is the kind of "rogue frictional market" we want to detect. The shape that the price volatility displays is essentially similar to that of the quasi-efficient case, albeit higher to capture more sustained volatility which could be associated with the crash. In Section 3, we even allow volatility to explode to infinity, without compromising our reasoning. The cusp in the expected price, the V-shape, is associated with a drift coefficient which gets first largely negative, and then largely positive immediately after, hence Definition 1. The following proposition shows that, in our setting, a V-shape implies market inefficiency.

Proposition 1. *Assume that* $p_0^e = p_1$ *and* $p_1^e = p_2 \leq p_1$*, with* $p_1, p_2 \in \mathbb{R}$*. If the following conditions hold:*

1. *The price* p_t *exhibits a V-shape in* $\tau \in]0, 1[$*;*
2. $\overline{p}_0 = p_1$ *and* $\overline{p}_1 = p_2$ *(that is, the market is efficient in the long run);*

then $\overline{p}_\tau < p_2$ *(that is, the market is inefficient).*

Proof. If we had $\overline{p}_\tau > p_2$, by Lagrange theorem, it should exist a t^* with $\tau < t^* < 1$ where

$$\mu_{t=t^*} = \frac{p_2 - \overline{p}_\tau}{1 - \tau} < 0\,,$$

but this is in contrast with Assumption 1. □

The flash crash of May 6, 2010, displayed in Figure 1, is an example of a V-shape in which there were no specific news hitting the market. The market in the third column has been widely described by both the empirical and theoretical literature. Examples of typical distressed events in which the asset price follows a V-shaped trajectory are reported, for example, in Bellia et al. (2019), Figure 2, in Christensen et al. (2017), Figure 1, in Kirilenko et al. (2017), or described in Brunnermeier and Pedersen (2005), Figure 2. A realistic market with frictions is either represented by the quasi-efficient case, when the market can still be considered to be well-functioning, or by

the inefficient case, when market efficiency breaks. We argue that volatility cannot really tell apart these two settings. Volatility can spike (and even explode to infinity) in the inefficient setting, in the quasi-efficient setting, and even in the perfectly efficient setting. High volatility cannot be uniquely associated to a distressed state. Thus, it does not appear to be a reliable indicator of inefficiency. Neither jumps can help. Indeed, price jumps – either sudden or gradual – are natural in efficient markets when news come in.

Importantly, Figure 2 is agnostic about the duration of the inefficient state, which depends on the liquidity of the markets and on its ability to recover quickly enough. The May 6, 2010 U.S. equity market flash crash lasted approximately 30 minutes only, and most flash crashes are typically short-lived. Short-lived inefficiencies are customary for stocks, futures, and foreign exchange rates, which are characterized by large liquidity and limited fragmentation. However, we can observe flash-crash type behavior spanning several hours, or even days, when several frictions – leading to price overshooting – concur jointly. For example, Duffie (2010) describes V-shapes ranging several days, for deleted stocks, secondary equity issuances, U.S. Treasury and corporate bond issuances, and even ranging few months for mutual funds experiencing large redemptions. Thus, frictions play a role in determining the duration of V-shapes, as our empirical application below confirms.

3 The V-statistic

We work with a traditional continuous-time model encompassing virtually all popular models in financial economics and broadening them to allow for an explosive drift and volatility. We enrich the model (2.1) by assuming that the log-price process X, observed in the interval $[0,1]$, is driven by the equation:

$$\mathrm{d}X_t = \underbrace{\mu_t\,\mathrm{d}t + \sigma_t\,\mathrm{d}W_t + \mathrm{d}J_t}_{\text{traditional semimartingale}} + \underbrace{\frac{c^-_{1,t}}{(\tau_{\mathrm{db}}-t)^\alpha}\mathbb{1}_{\{t<\tau_{\mathrm{db}}\}}\,\mathrm{d}t + \frac{c^+_{1,t}}{(t-\tau_{\mathrm{db}})^\alpha}\mathbb{1}_{\{t>\tau_{\mathrm{db}}\}}\,\mathrm{d}t}_{\text{drift burst}} + \underbrace{\frac{c_{2,t}}{|\tau_{\mathrm{db}}-t|^\beta}\,\mathrm{d}W_t}_{\text{volatility burst}}, \quad (3.1)$$

where μ_t and σ_t are bounded drift and volatility respectively, W_t is a Brownian motion, J_t is a jump process, and $\tau_{\mathrm{db}} \in]0,1[$, is the explosion time point. We allow for different drift explosion coefficients $c^-_{1,t}$ and $c^+_{1,t}$ before and after the explosion point τ_{db}. For ease of exposition, we use the same rate of explosion $\alpha \in [0,1[$ before and after, an assumption that is harmless and that can be relaxed without altering our results. We also allow for explosion in volatility with rate $\beta \in [0,1/2[$. The technical conditions that the coefficients in Eq. (3.1) should meet are extremely mild and spelled out in Assumption 1 in Appendix A. We note that an infinite drift and volatility at point τ_{db} is still consistent with X_t being a semi-martingale (Jacod and Protter, 2012). Indeed,

what matters is that these quantities can be integrated over finite intervals including the explosion point, which is warranted by our assumptions on the coefficients α and β.

We are interested in detecting V-shapes as defined in Definition 1. Here we show that we are able to do so when we have explosion both before and after τ_{db} (that is, $c^-_{1,\tau_{\text{db}}} \neq 0$ and $c^+_{1,\tau_{\text{db}}} \neq 0$ and they have different signs), against the null in which the explosion is only at the left or at the right (that is, we have a gradual jump), or there is no explosion at all. Please note that in Definition 1 we are not requiring the drift to explode to define a V-shape, but just to change sign discontinuously. This disconnect is due to the fact that, for the continous-time semi-martingale (3.1), it is impossible to estimate the sign of the drift consistently (Bandi, 2002; Kristensen, 2010) unless it explodes, as was shown in COR. Indeed, estimation of μ_t in Eq. (3.1) is jeopardized by the Brownian component. This is however empirically not restrictive, since in small samples (where prices are discrete) an exploding drift in continuous time should be interpreted just as a "large drift", as usually done with asymptotic theory. Alternatively, one may interpret our econometric procedure as a way to detect a subset of V-shapes as defined in Definition 1, that in which the change in the drift sign is accompanied with explosion.

We borrow from the approach of COR in adopting a fully non-parametric technique. The advantage of non-parametric approaches is that the detected evidence is robust to model specification. COR propose a test statistic to detect the alternative of drift bursts versus the null of bounded drift. Their test uses $n+1$ log-price observations $X_0, \ldots, X_n$ observed at times $t_0, \ldots, t_n$ satisfying Assumption 3 in Appendix A. The test can be formally expressed, at time-point τ, as:

$$T^-_{\tau,n} = \sqrt{\frac{h_n}{K_2^-}} \frac{\hat{\mu}^-_{\tau,n}}{\hat{\sigma}^-_{\tau,n}}\,, \tag{3.2}$$

where

$$\hat{\mu}^-_{\tau,n} = \frac{1}{h_n} \sum_{i=1}^{n} K^- \left(\frac{t_{i-1} - \tau}{h_n} \right) (X_i - X_{i-1})\,, \quad \text{for } \tau \in (0,1]\,, \tag{3.3}$$

is a localized estimator of the drift, and

$$\hat{\sigma}^-_{\tau,n} = \left(\frac{1}{h_n} \sum_{i=1}^{n} K^- \left(\frac{t_{i-1} - \tau}{h_n} \right) (X_i - X_{i-1})^2 \right)^{\frac{1}{2}}\,, \quad \text{for } \tau \in (0,1]\,, \tag{3.4}$$

is a localized estimator of the volatility. In equations (3.3) and (3.4), h_n is a bandwidth parameter measuring the extent of the localization, and $K^-(\cdot)$ is a suitable left-sided kernel, that is $K^-(x) = 0$ when $x \geq 0$, and $K_2^- = \int_{\mathbb{R}} (K^-(x))^2 \, \mathrm{d}x$ is a kernel-specific constant. The properties of the bandwidth and the kernel are specified in Assumption 2 in Appendix A.

From a statistical point of view, the statistic $T^-_{\tau,n}$ can be interpreted as the ratio between the part of the log-return, between $\tau - h_n$ and τ, due to the drift and the part due to volatility. The bandwidth h_n has thus an important interpretation as the duration of the V-shape, that is of the potential market inefficiency. Under the null of non-exploding drift, that is, $c^-_{1,\tau} = c^+_{1,\tau} = 0$ in model (3.1), the statistic converges to a standard normal distribution, even in the presence of jumps or volatility bursts. A large value of the test statistic thus signals an abnormal trend, or a "drift burst" in the COR language.

The test we propose is similar in spirit. We propose indeed to use the product of the test-statistic (3.2) with a similar one, now computed using a right-sided kernel, suitably scaled by the square root of the bandwidth. We label it *V-statistic*. Formally, it takes the form:

$$\mathcal{V}_{\tau,n} = \sqrt{h_n} \cdot T^+_{\tau,n} \cdot T^-_{\tau,n} \tag{3.5}$$

where

$$T^+_{\tau,n} = \sqrt{\frac{h_n}{K_2^+}} \frac{\hat{\mu}^+_{\tau,n}}{\hat{\sigma}^+_{\tau,n}}, \tag{3.6}$$

with $\hat{\mu}^+_{\tau,n}$ and $\hat{\sigma}^+_{\tau,n}$ computed as in Eqs. (3.3) and (3.4) with $K^-(\cdot)$ replaced by the right-sided kernel $K^+(\cdot)$, that is $K^+(x) = 0$ when $x \leq 0$. The interpretation of $T^+_{\tau,n}$ is the same as that of $T^-_{\tau,n}$, now for the log-return between τ and $\tau + h_n$.

The V-statistic $\mathcal{V}_{\tau,n}$ has thus a simple interpretation. When positive, it identifies a drift having the same sign before and after t, that is a trending price, either downward (a crash) or upward (a surge). When negative (the interesting case for us), it identifies a switch, that is a time-point t in which the drift changes sign and the price experiences a V-shape or a Λ-shape. The scaling by $\sqrt{h_n}$ is needed to disentangle the case in which both $T^-_{\tau,n}$ and $T^+_{\tau,n}$ explode, so that the statistic diverges to infinity, from the gradual jump case in which only one of the two t-test explodes, and the scaling annihilates the statistic asymptotically. We can indeed prove the following:

Theorem 1. *Assume X is driven by model* (3.1)*. Under Assumption 1, 2 and 3, if the following conditions are met:*

1. $\alpha - \beta > 3/4$;
2. $\frac{1}{n^{2-2\alpha} h_n} \to 0$ *as* $n \to \infty$;
3. $c^{\pm} = c^-_{1,\tau_{\text{db}}} \cdot c^+_{1,\tau_{\text{db}}} \neq 0$,

then, as $n \to \infty$, $\mathcal{V}_{\tau_{\text{db}},n} \to sign(c^{\pm})\infty$. *If instead* $c^{\pm} = 0$, *then as* $n \to \infty$, $\mathcal{V}_{\tau_{\text{db}},n} \to 0$.

Proof. See Appendix A. □

Theorem 1, shows that the V-statistic can formally capture two-sided drift explosions (when $c^{\pm} \neq 0$), diverging to infinity, when three necessary conditions are met: existence of an exploding drift both at the left and at the right of the exploding point ($c^{\pm} \neq 0$), sufficiently fast drift explosion ($\alpha - \beta > 3/4$) and sufficiently large bandwidth h_n ($n^{2-2\alpha} h_n \to \infty$). If one of these three conditions is not met, the statistic will go to zero (or be bounded at the boundaries of the conditions, e.g. when $\alpha - \beta = 3/4$). When $c^{\pm}$ is negative, we have a V-shape or a Λ-shape (when $T^{-}_{\tau,n} < 0$ and $T^{+}_{\tau,n} > 0$, or when $T^{-}_{\tau,n} > 0$ and $T^{+}_{\tau,n} < 0$, respectively). When $c^{\pm}$ is positive, we have a sustained trend.

Since in practice it is natural to test on multiple points, it makes sense also to define aggregate statistics. When looking for V-shapes, we will look at the minimum value of the V-statistics as computed on several points $(\tau_1, \ldots, \tau_m)$ as:

$$\mathcal{M}\text{in}\mathcal{V}_{\tau_1,\ldots,\tau_m,n} = \min_{i=1,\ldots,m} \mathcal{V}^{\pm}_{\tau_i,n}. \tag{3.7}$$

Finally, with the intention to apply this technique to tick-by-tick data, we further enrich the model with a classical market microstructure noise component, by assuming that the observed log-price $\tilde{X}_t$ is given by:

$$\tilde{X}_t = X_t + \varepsilon_t, \tag{3.8}$$

where ε_t satisfies Assumption 4 in Appendix A. The presence of market microstructure noise in the data is tackled with a combination of pre-averaging (Jacod et al., 2009) and HAC correction (Andrews, 1991) as in COR, to which the reader is referred for details.

3.1 Small-sample confidence bands: simulated bootstrap with stochastic volatility

In small samples, how can we separate the two cases of diverging and annihilating $\mathcal{V}_{\tau,n}$? And how can we obtain confidence bands when aggregating multiple tests, as in Eq. (3.7)? Under the null of neither drift nor volatility explosions ($c^{-}_{1,\tau_{\text{db}}} = c^{+}_{1,\tau_{\text{db}}} = c_{2,\tau_{\text{db}}} = 0$), the product $T^{-}_{\tau,n} \cdot T^{+}_{\tau,n}$ is asymptotically distributed as the product of two independent standard normal variates, that is it has the distribution of a modified Bessel function of the second kind of order zero (Craig, 1936). The asymptotic mean and variance of $\mathcal{V}_{\tau,n}$ are still 0 and 1 respectively, and their (two-sided) 95% and 99% confidence limits are 2.18 and 3.60 respectively. Choosing time units such that $\sqrt{h_n} = 1$, as we do in the empirical application, these values can be loosely used as a replacement

for traditional confidence limits. However, as our simulations below will show, these values would be too small to be used in practice for realistic data generating processes, and useless against gradual jumps (when $c^{\pm} = 0$ but either $c^{+}_{1,\tau_{\text{db}}}$ or $c^{-}_{1,\tau_{\text{db}}}$ is nonzero) and multiple testing biases.

We advocate instead the adoption of a simulated bootstrap to evaluate confidence bands for the $\mathcal{V}_{\tau,n}$ and aggregated statistics as $\mathcal{M}\text{in}\mathcal{V}_{\tau_1,\ldots,\tau_m,n}$ in Eq. (3.7). The idea of the bootstrap is to compute the distribution of $\mathcal{V}_{\tau,n}$ (or $\mathcal{M}\text{in}\mathcal{V}_{\tau_1,\ldots,\tau_m,n}$) on simulations based on the data at hand. Simulations should have *zero drift*, and reproduce the stochastic volatility observed in the data. Attributing all the potential drift observed in the sample to volatility will generate confidence bands which are automatically robust to stochastic volatility. In this paper, we use an EGARCH(1,1) model to filter the observed variance, to be used for simulated bootstrap. The EGARCH(1,1) has two appealing features: it is simple, and it can reproduce stochastic volatility and the leverage effect. Confidence bands in Figure 1 are obtained this way.

4 Simulations

We test the V-statistics on synthetic data obtained by simulating four models: the first one is a simple geometric Brownian motion, and the last three are those described in Section 2 and illustrated in Figure 2 (perfectly efficient market, quasi-efficient market, inefficient market). More precisely, the models we simulate in the time span $[0, 1]$ (all units are meant to be daily, so that the time span represents one day of trading) are:

- Model 0: $\begin{cases} \overline{p}_t = \log(P_1) \\ \sigma_t = \sigma_0 \end{cases}$
- Model 1: $\begin{cases} \overline{p}_t = \log\left(P_1 \cdot I_{\{t\leq 0.5\}} + P_2 \cdot I_{\{t>0.5\}}\right) \\ \sigma_t = \sigma_0 + (\sigma_1 - \sigma_0)e^{-(2t-1)/\tau_1} \cdot I_{\{t>0.5\}} \end{cases}$
- Model 2: $\begin{cases} \overline{p}_t = \log\left(P_1 - \frac{P_1-P_2}{1+e^{-(2t-1)/\tau_2}}\right) \\ \sigma_t = \sigma_0 + (\sigma_1 - \sigma_0)\left(e^{-(2t-1)/\tau_1} \cdot I_{\{t>0.5\}} + e^{(2t-1)/\tau_2} \cdot I_{\{t\leq 0.5\}}\right) \end{cases}$
- Model 3: $\begin{cases} \overline{p}_t = \log(P_1) + \frac{1-\alpha}{\alpha}\int_0^t \Big((\log(P_1) - \log(P_3))\left(1 - (-(2s-1))^{-\alpha}\right)I_{\{s\leq 0.5\}} \\ \qquad\qquad + (\log(P_2) - \log(P_3))\left((2s-1)^{-\alpha} - 1\right)I_{\{s>0.5\}}\Big)\, ds \\ \sigma_t = \sigma_0 + (\sigma_2 - \sigma_0)\left(e^{-(2t-1)/\tau_1} \cdot I_{\{t>0.5\}} + e^{(2t-1)/\tau_2} \cdot I_{\{t\leq 0.5\}}\right) \end{cases}$

Parameters are set as follows: $\sigma_0 = 2\%$ (the baseline daily volatility), $\sigma_1 = 10\%$ (the peak of volatility after the change in the efficient price in Model 1 and 2), $\tau_1 = 1/3$ (the decay time of the

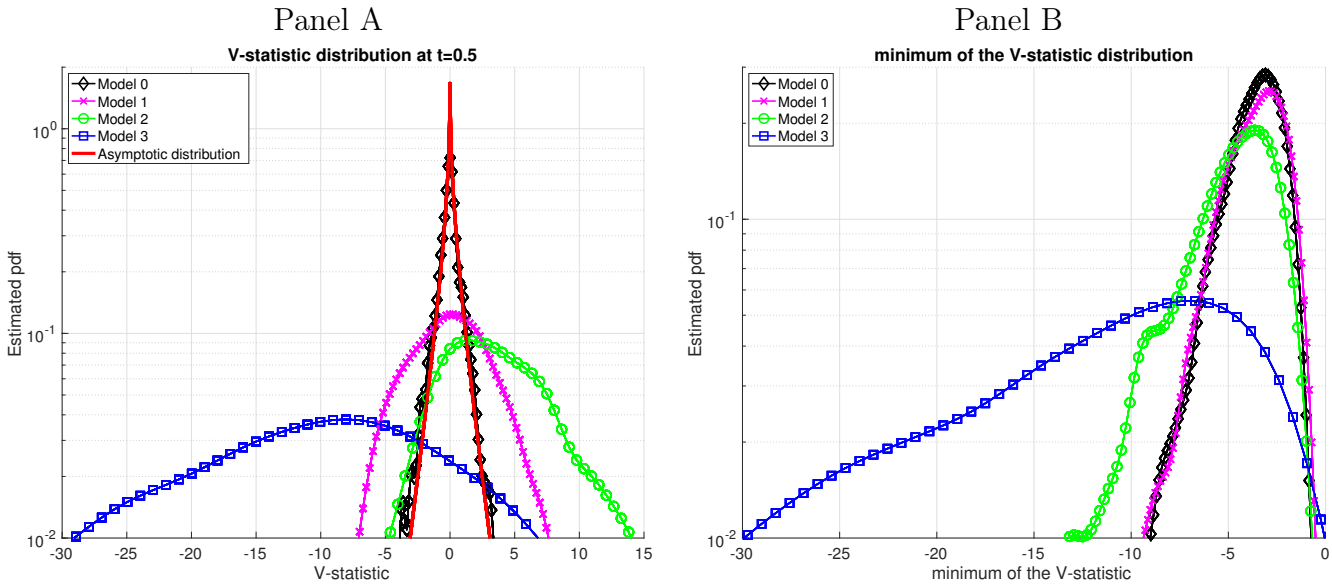

Figure 3: We simulate Model 0, 1, 2, 3 with $P_1 = 100$, $P_2 = 95$, $P_3 = 90$, that is with a 5% drop in the efficient price and a 5% overshooting of the traded price. Panel A: Distributions of the V-statistics at $t = 0.5$ on $1,000$ replications of Model 0, 1, 2, 3. The asymptotic distribution is the modified Bessel function of second kind of order zero. Panel B: Distributions of the minimum of the V-statistics over 101 points in the interval $[0.1, 0.9]$ on $1,000$ replications of Model 0, 1, 2, 3.

volatility), $\tau_2 = 0.05$ (the decay time of the price, and the exponential rise time of volatility in Model 2 and Model 3), $\alpha = 0.7$ (the explosion rate for drift in Model 3). The parameter P_1 is the starting price, P_2 is the final price, P_3 is the lowest reached price and $P_2 - P_3$ is the overshooting. For Model 3, we also use a heightened volatility at the peak, $\sigma_2 = 20\%$.

We generate 1,000 simulations of the three models on paths of $n = 23,400$ observations (corresponding to 1−second observations in a day with 6.5 hours of trading). We compute the $V-$statistics on 101 points equally spaced in the interval $[0.1, 0.9]$. We use an exponential kernel and a bandwidth $h_n = 0.1$ for the numerator, and $h'_n = 0.5$ for the denominator, which is consistent with the recommendations of COR.

Panel A of Figure 3 shows the distribution of $\mathcal{V}_{\tau,n}$ at the middle point $t = 0.5$ in the case $P_1 = 100$, $P_2 = 95$, $P_3 = 90$, that is when we assume a 5% drop in the efficient price and 5% overshooting (see the top-left panel of Figure 2). The simple geometric Brownian motion (Model 0) has a small-sample distribution of the V-statistics perfectly in line with that of the modified Bessel function of second kind of order zero, which may be considered as a sanity check. Models 1, 2 and 3 instead deviate from the Bessel distribution. Model 1 features a pre-announced jump, that is a jump in all the generated trajectories. Thus, Model 1 is not under the null of Theorem 1. Panel A of Figure 3 shows that the distribution of the V-statistics is inflated with respect to the Bessel distribution

Table 1: Percentage of rejections at 1% and 5% of the $\mathcal{M}\text{in}\mathcal{V}_{\tau_1,\ldots,\tau_m,n}$ statistic, based on 1,000 bootstrapped simulations with an EGARCH(1,1) model. In Panel A we have $P_1 = P_2 = 100$, $P_3 = 95$. In Panel B we have $P_1 = 100$, $P_2 = 95$, $P_3 = 90$. Model 3' has the same drift of Model 3 but the (lower) volatility of Model 2.

Panel A: no fundamental price change, 5% overshooting.

Quantile	Model 0	Model 1	Model 2	Model 3	Model 3'
	Percentage of rejections				
1%	0.70	0.60	1.20	11.60	52.40
5%	5.30	5.90	6.10	25.60	70.30

Panel B: 5% fundamental price change, 5% overshooting.

Quantile	Model 0	Model 1	Model 2	Model 3	Model 3'
	Percentage of rejections				
1%	0.70	2.20	2.50	38.90	87.40
5%	5.30	8.30	11.10	56.50	92.10

but still centered around zero. In the case of Model 2, which is under the null of Theorem 1, the downward trend in the price results in a *positive* value of the V-statistics, as expected. In the case of Model 3, which is the alternative in Theorem 1, the V-shape results in a large *negative* value of the V-statistics. Thus, Figure 3 shows that the V-statistics behave as expected from the theory. Importantly, the V-statistics have considerable power in disentangling a V-shape (as in Model 3) from a gradual jump (as in Model 2) and a pure jump (as in Model 1).

Panel B of Figure 3 shows the distribution of $\mathcal{M}\text{in}\mathcal{V}_{\tau_1,\ldots,\tau_m,n}$ on the same simulations of Panel A. As said, the minimum is computed over $m = 101$ testing points. Interestingly, in this case Model 0 and Model 1 deliver similar result: the price jump does not change the distribution of the minimum V-statistic much with respect to a simple Brownian motion. Model 2 can produce some false positives in detecting V-shapes (i.e., negative values of the statistics), even if of limited amount, its distribution being slightly skewed to the left, due to the large volatility after the gradual jump. Model 3 is clearly skewed to the left, indicating its statistical power in detecting V-shapes using the minimum statistic. Table 1 reports the number of rejections at 1% and 5%, with the confidence level computed using a simulated bootstrap based on the EGARCH(1,1) model. We now have five models: we add Model 3', which has the same drift of Model 3 but the (lower)

volatility of Model 2. We also simulate two settings: Panel B of Table 1 refers to the same setting as above ($P_1 = 100$, $P_2 = 95$, $P_3 = 90$) with 5% fundamental drop and 5% overshooting; Panel A refers to a setting with no news ($P_1 = P_2 = 100$, $P_3 = 95$) with 5% overshooting only. Model 0 is nicely sized in both cases. Model 1 and 2 are correctly sized when there are no fundamental news, but slightly oversized when there are fundamental news, since the fundamental drop biases the V-statistics downward. Table 1 also shows the power of the test in the case of Model 3, and illustrates its determinants. We can see that power increases from Panel A to Panel B (that is, the downward fundamental news "helps" the statistics a little bit to identify the V-shape) and from Model 3 to Model 3' (that is, with a lower volatility).

Summarizing, our simulation experiments show quite clearly that the V-statistics, and in particular their minimum, are able to disentangle a V-shape by a random fluctuation accompanied by a jump and/or a burst in volatility on synthetic data. The ability of the V-statistics to detect V-shapes increases when price volatility is lower and, not surprisingly, when the V-shape is accompanied by an actual change in the price fundamentals in the same direction. Corroborated by these results, we now turn to real financial market data.

5 V-shapes in the U.S. stock index

We apply the V-statistics to all transacted prices of SPY (the SPDR S&P 500 trust exchange-traded fund) which can be considered a reliable and liquid proxy of the U.S. stock index. The sample spans 2019 and 2020 up to August 17, for a total of 408 trading days. We clean the data using standard procedures (see Appendix B) and, given the huge amount of recorded transactions (roughly 480K transactions per day, more than 195M in total) we interpolate the prices to a 1-second grid to ease data management. Importantly, the data set includes the turmoil due to the spread of the Covid-19 pandemic, which reached its climax on March 23, 2020. Figure 4 shows the time series of 1-second prices for our sample. A sensible question is whether the U.S. stock market was efficient during those dramatic times.

To answer this question, we start by computing the V-statistic precisely at the time in which the price reached its minimum, that is on Mar, 23 at 11:46:11 am (the recorded price was 218.29). We use different bandwidths h_n ranging from 5 minutes to more than 40 hours. The bandwidths can be interpreted, as discussed above, as the duration of the state of potential market inefficiency. The result is shown in the inset of Figure 4. The V-statistics are negative, as expected, but not too large, at least when compared to the values measured during the flash crash of May 6, 2010 (see Figure 1). The V-statistics are just below the 5% confidence limit, and never exceed the 1%

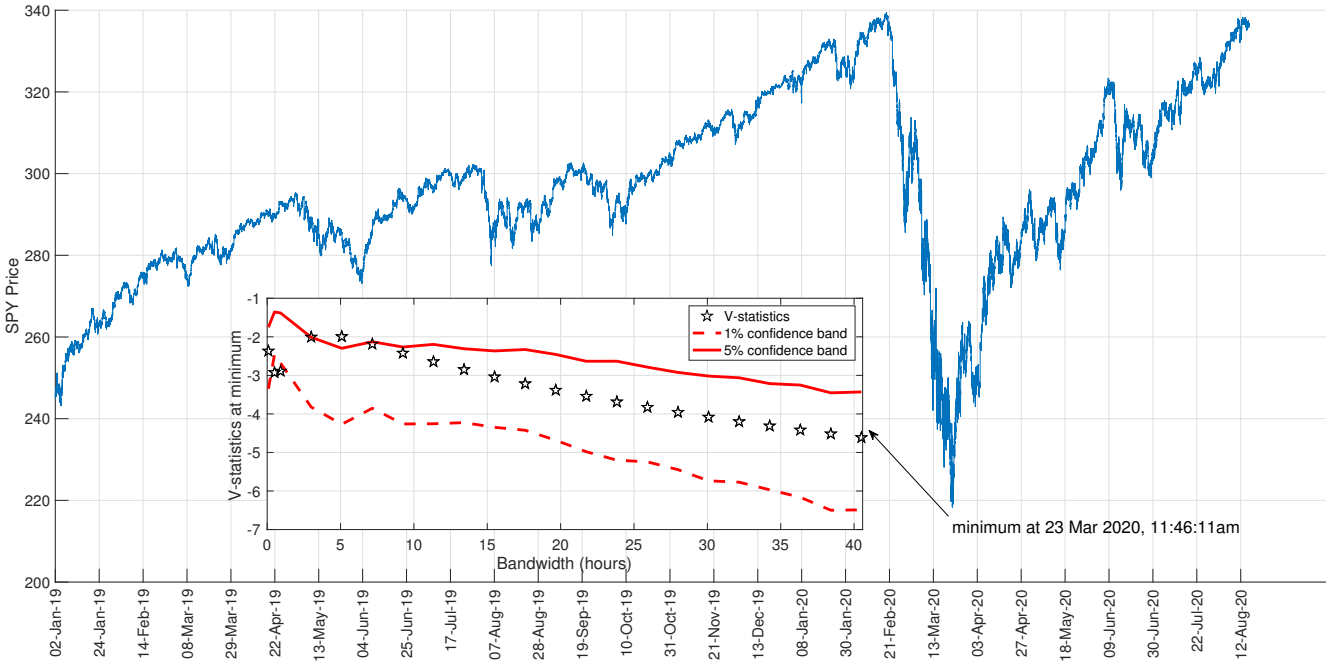

FIGURE 4: Traded prices of SPY interpolated to the nearest second for the sample period considered. In the inset, we show the V-statistics, together with 1% and 5% confidence bands computed via simulated bootstrap, as computed at the instant in which SPY reaches the minimum price, that is on Mar 23, 2020, at 11:46:11 a.m., for different choices of the bandwidth parameter h_n.

confidence limit, if not (slightly) for the 5 and 30 minutes bandwidths. Moreover, the confidence bands reported in the inset of Figure 4 are clearly conservative since we are not testing at a fixed point, but at a data-dependent point (the minimum), which of course biases the measure downward. Thus, we cannot conclude that the U.S. stock market was inefficient on periods longer than 30 minutes, even during the turmoil due to the pandemic.

We next move to a more systematic study and look for V-shapes in the whole sample using the V-statistics. Motivated by the evidence in Figure 4, we test with three bandwidths: 30, 5 and 1 minutes. Figure 5 shows the minimum of the intraday V-statistics for every day in our sample with the 30-minute bandwidth, together with the 5% and 1% confidence bands obtained by simulated bootstrap, using an EGARCH(1,1) model fitted on 1-second returns. We point at two observations: first, we see significant violations of the null, indicating the presence of V-shapes. At the 5% confidence level, we find 35 violations (there should be 20 under the null), 28 being V-shapes and 7 being Λ-shapes.[1] At the 1% confidence level, we find 11 violations (there should be 4 under the null), 10 being V-shapes and only 1 a Λ-shape. The 11 price paths violating the null

[1]The fact that we have more V-shapes than Λ-shapes is another confirmation, both from the statistical and the economic perspective, of the frequent presence of inefficiencies in the market. Indeed, if the violations of the V-statistics were a mere by-product of statistical uncertainty, these would happen with the same probability for crashes and surges. The theoretical literature provides explanations for the asymmetry, see e.g. Huang and Wang (2009).

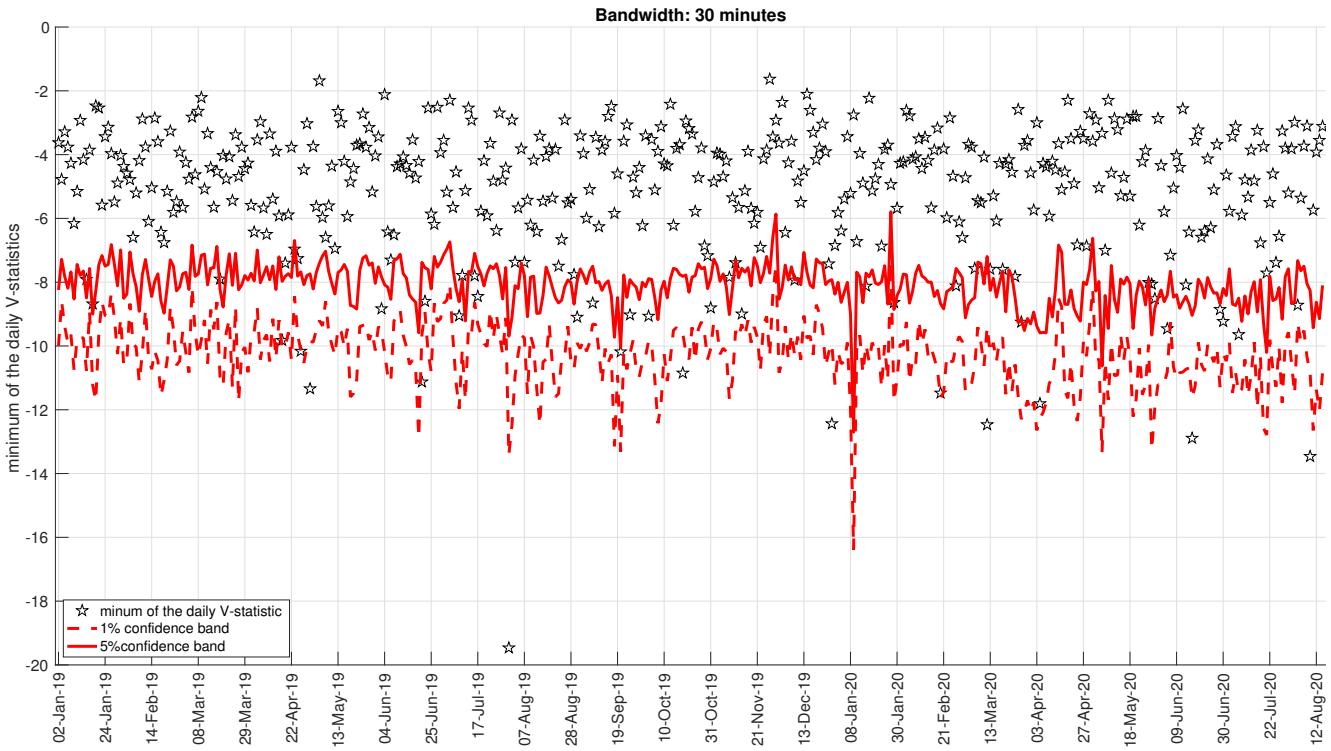

Figure 5: Minimum of the daily V-statistics computed on SPY transaction data with the bandwidth $h_n = 30$ minutes, together with confidence bands at 1% and 5% obtained via simulated bootstrap based on the EGARCH(1,1) model.

at the 30-minute bandwidth are all displayed in Figure 6. The largest V-statistic in our sample occurs on July 31, 2019, when, in a news conference after the Federal Open Market Committee, the FED chairman hinted that further rate cuts were less likely in the near future.

Table 2 reports the numbers of significant violations of the null, also when using the 1-minute and 5-minute bandwidth. We see that, consistently with the literature on mini flash crashes, the number of significant V-shapes grows when the tested duration of the inefficiency decreases. Table 2 also tests for inefficiency using a different metric, which is the average observed return from the peak of the V-shape onward, multiplied by -1 for surges. The statistical significance of positive observed average returns over an horizon up to 4 hours confirms the state of inefficiency of the market during a V-shape.

6 Case study: Italian bonds in May, 2018

The purpose of this section is to apply our methodology to a specific case study: the crash of May, 2018 in the Italian sovereign bond market. The choice of this specific event is not accidental. First, the Italian bond market, one of the largest in the world (nominal size of Italian public debt

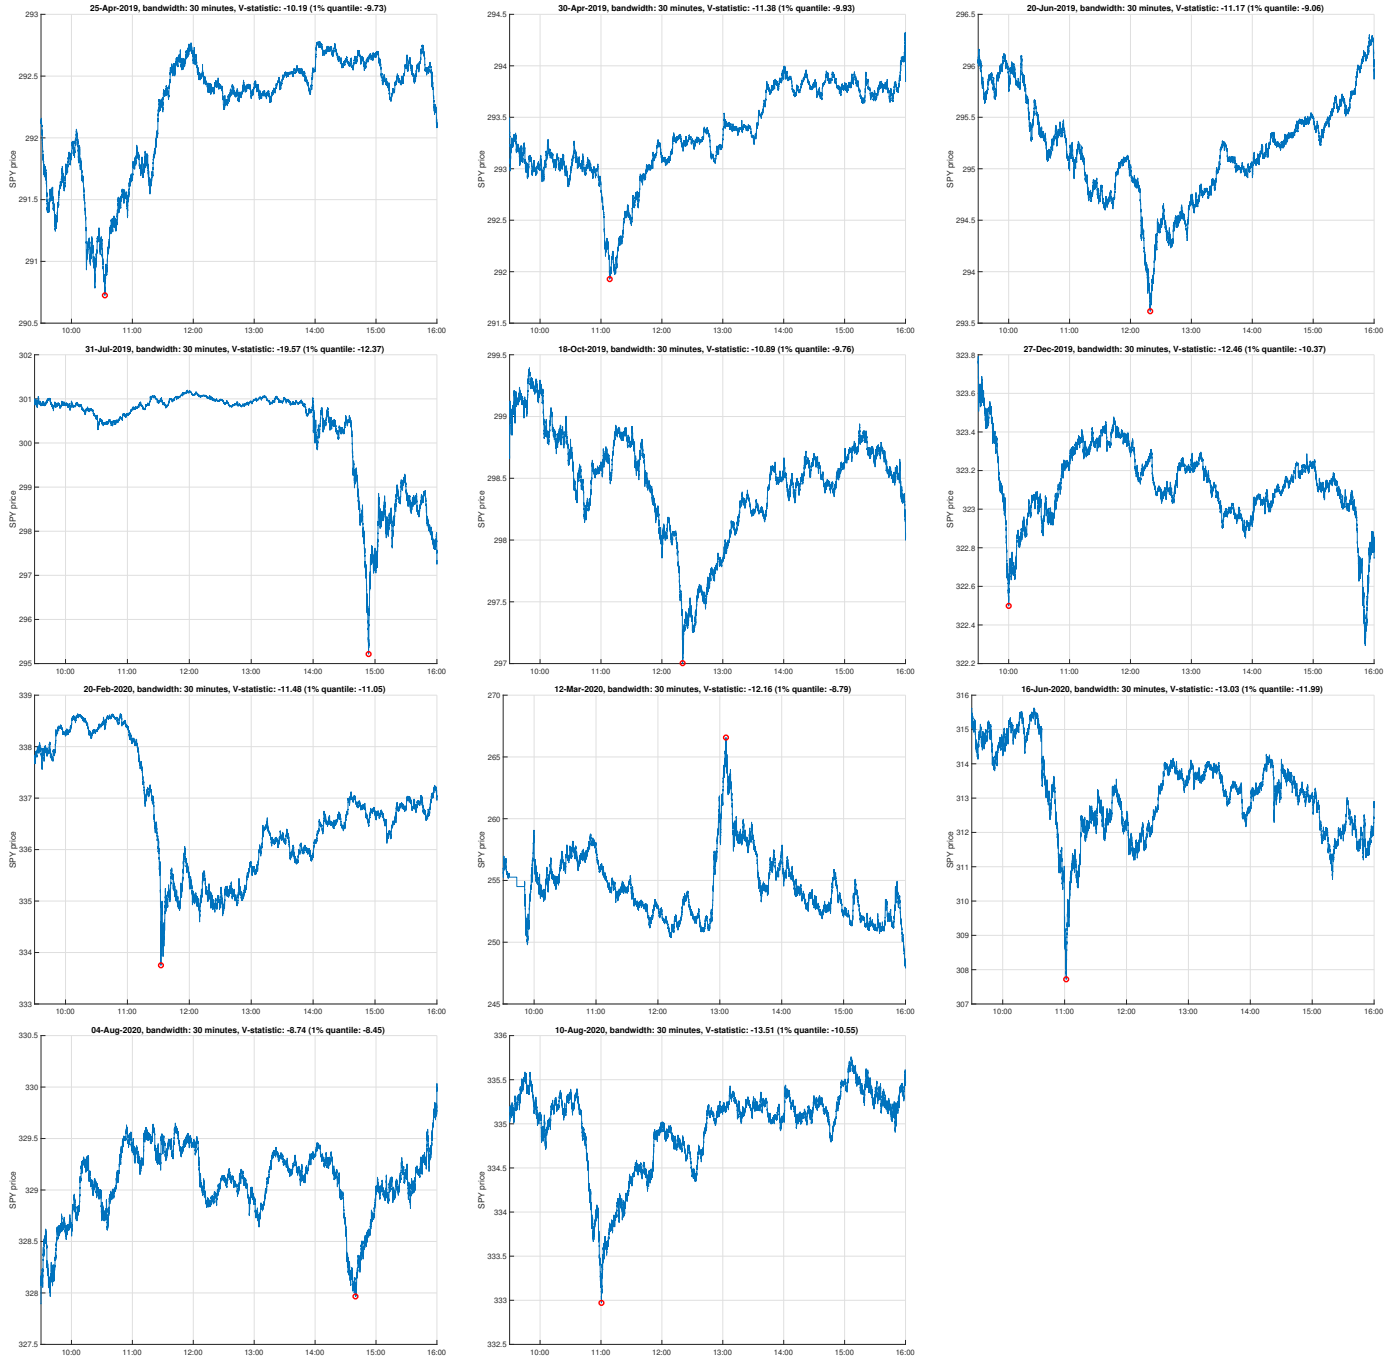

Figure 6: Each panel reports the daily time series of SPY prices for each day with a significant (1% confidence level for the 30-minute bandwidth) V-shape, according to the daily minimum of the V-statistics. The point at which the V-statistics reach their minima is denoted by a red circle.

Table 2: Number of significant test rejections, based on the minima of the intraday V-statistics, and average returns after the peak of the V-shape, multiplied by -1 for surges, for different time horizons. Panel A, B and C report results for the 30, 5 and 1 minute bandwidth, respectively. t-statistics for average returns are in parenthesis. ***, ** and * denote 99%, 95% and 90% significance, respectively.

Panel A: Bandwidth 30 minutes

	Number of:		Average Return (daily units) after minutes:				
Confidence level	V-shapes	Λ-shapes	5	30	60	120	240
1%	10	1	30.31^{***}	10.70^{***}	5.61^{***}	3.73^{**}	2.23^{**}
			(2.88)	(2.63)	(2.76)	(2.40)	(2.07)
5%	28	7	22.93^{***}	7.51^{***}	4.24^{***}	2.49^{***}	1.34^{***}
			(5.22)	(5.09)	(5.62)	(4.35)	(3.66)

Panel B: Bandwidth 5 minutes

	Number of:		Average Return (daily units) after minutes:				
Confidence level	V-shapes	Λ-shapes	5	30	60	120	240
1%	12	7	33.44^{***}	7.52^{***}	4.80^{**}	2.67^{***}	1.25^{***}
			(3.71)	(3.24)	(2.31)	(2.60)	(3.09)
5%	40	21	33.90^{***}	7.61^{***}	4.01^{***}	2.53^{***}	1.45^{***}
			(6.20)	(4.34)	(4.21)	(4.55)	(4.40)

Panel C: Bandwidth 1 minute

	Number of:		Average Return (daily units) after minutes:				
Confidence level	V-shapes	Λ-shapes	5	30	60	120	240
1%	18	7	12.76^{***}	1.77^{***}	0.37	0.02	0.06
			(11.05)	(5.44)	(0.81)	(0.07)	(0.32)
5%	47	33	16.81^{***}	2.71^{***}	1.00^{***}	0.32	0.28^{**}
			(6.25)	(3.82)	(3.45)	(1.52)	(2.02)

being roughly 15% of U.S.), is central to the global economy. Second, our analysis focuses on auction days, when an anticipated supply shock (the auction) hits the market and price drift is expected (Lou et al., 2013; Sigaux, 2018). This case study can be considered as an example of how V-shapes can be a serious threat to financial stability. Appendix C provides details about the data used for the analysis in this section, as well as the macroeconomic background during the crash.

6.1 Prices and volatility

Figure 7 shows the tick-by-tick transaction prices of selected Italian government bonds in 2018. The figure highlights the extreme price movement that occurred on May 29. The observed pattern is consistent with that shown in the top-right panel of Figure 2.

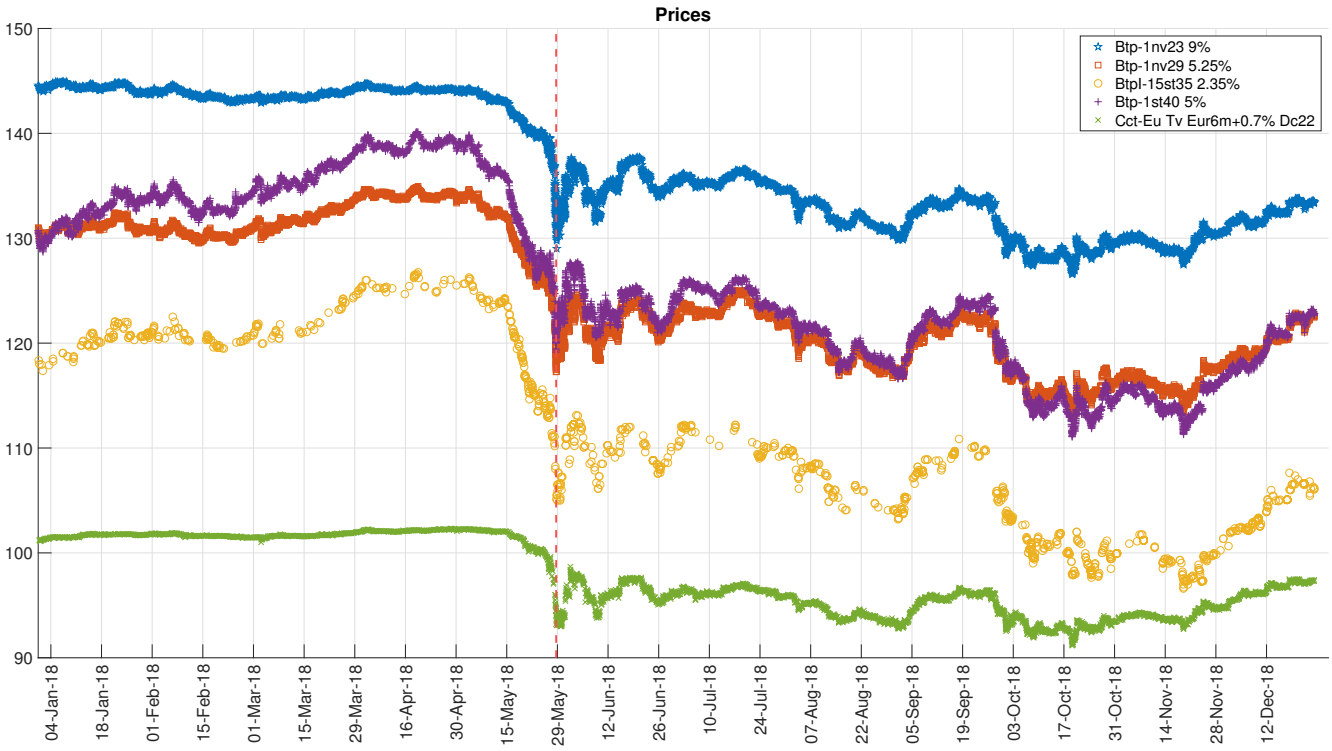

Figure 7: Tick-by-tick transaction prices (clean price) observed in the MOT (secondary market) for five representative Italian bonds. The dashed-red line is May 29, 2018.

For each bond, we estimate daily realized volatility (RV), that is the sum of squared 5-minute intraday returns. For this computation, prices are sampled on the 5-minute grid using the last observed transaction. Figure 8 shows the square root of daily RV for the instruments considered in Figure 7 (given the limited number of transactions, we exclude the inflation-linked BTPi from this plot). The RV paths have a marked spike on May 29, with a transient component that lasts a couple of months after the price crash. A similar long-term impact on volatility was observed after the flash crash of May 6, 2010 (Boulton et al., 2014). The price volatility dynamics of the five sovereign bonds is visually compatible with that represented in both the middle and right columns of Figure 2, and thus it does not unequivocally signal the presence of market inefficiency.

6.2 Liquidity measures

A key prediction of Grossman and Miller (1988) is that the transient mispricing should be more severe in a market with poor liquidity. We thus look at traditional liquidity measures in the Italian bond market.

The daily traded volumes of the five assets, showcased in Figure 9, show an increase around the week of May 29, 2018. The number of transactions, shown in the inset of the same figure,

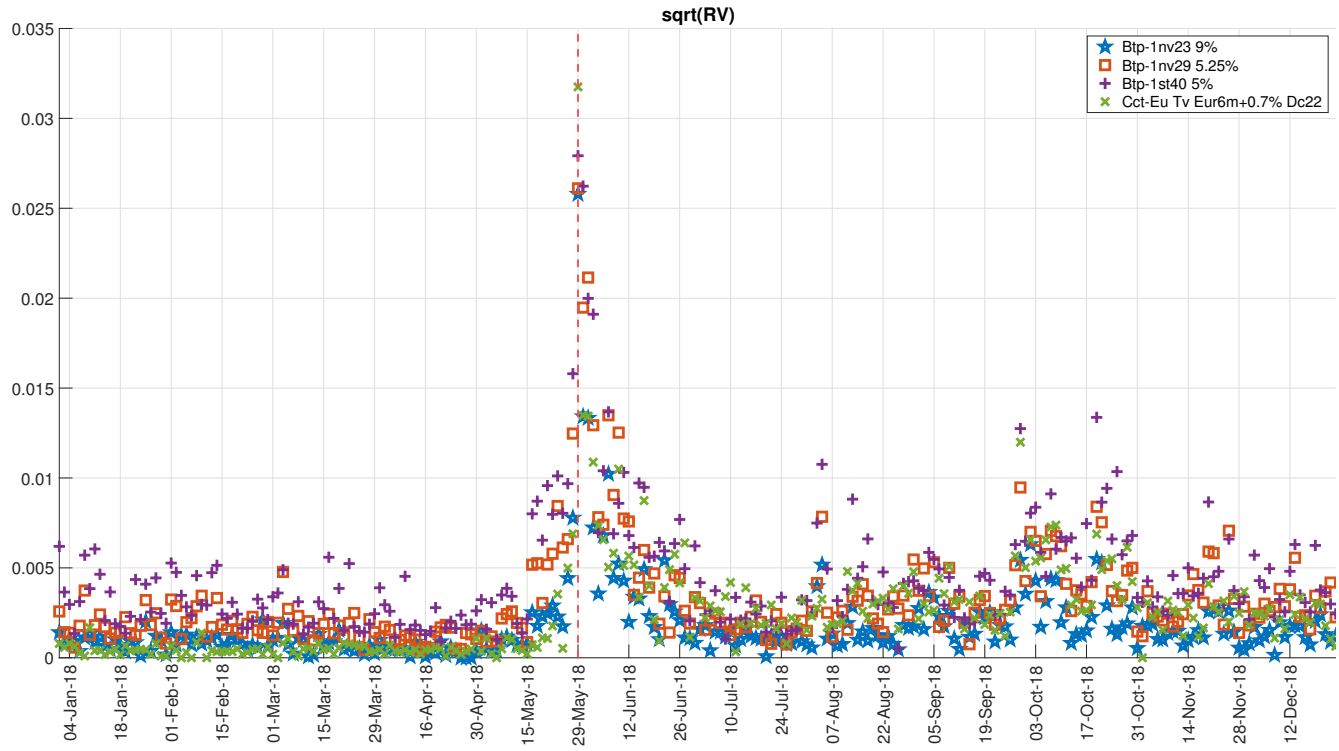

FIGURE 8: Square root of daily five-minutes realized volatility (RV) for four representative Italian bonds sampled in the MOT. The dashed-red line is May 29, 2018.

also peaks, and even more than volume itself, signaling that the average volume per transaction declined.

Figure 10 displays the daily bid-ask spread reported from Bloomberg for a representative 2-years and 10-years BTP. The Figure supports the intuition about a severe liquidity shock hitting the market. We notice indeed a huge spike of bid-ask spreads around and after the crash. We also note that the market was already quite illiquid *before* the crash, which is consistent with the intuition in Weller (2019) that the order book captures relevant forward-looking information for tail risk.

We also compute two realized liquidity measures that can be inferred directly from transaction prices, and which nevertheless confirm the evidence in Figure 10. The first is a measure of price impact which is close to the Amihud (2002) measure (we modify it to take into account the irregular sampling of trades). The measure is implemented as follows. For each trade t and bond i in the sample, we compute

$$\text{Amihud}^*_{i,t} = \frac{|\Delta \log(p_{i,t})|}{V_{i,t}\sqrt{T_{i,t}}}\,, \tag{6.1}$$

where $\Delta \log(p_{i,t}) = \log(p_{i,t}) - \log(p_{i,t-1})$ is the log-return between trade $t-1$ and trade t, $V_{i,t}$ is the volume of the $t-th$ trade, and $T_{i,t}$ is the time, in milliseconds, between trade $t-1$ and trade t. Panel A of Figure 11 shows the daily median of the measure in (6.1) for the four bonds in Figure

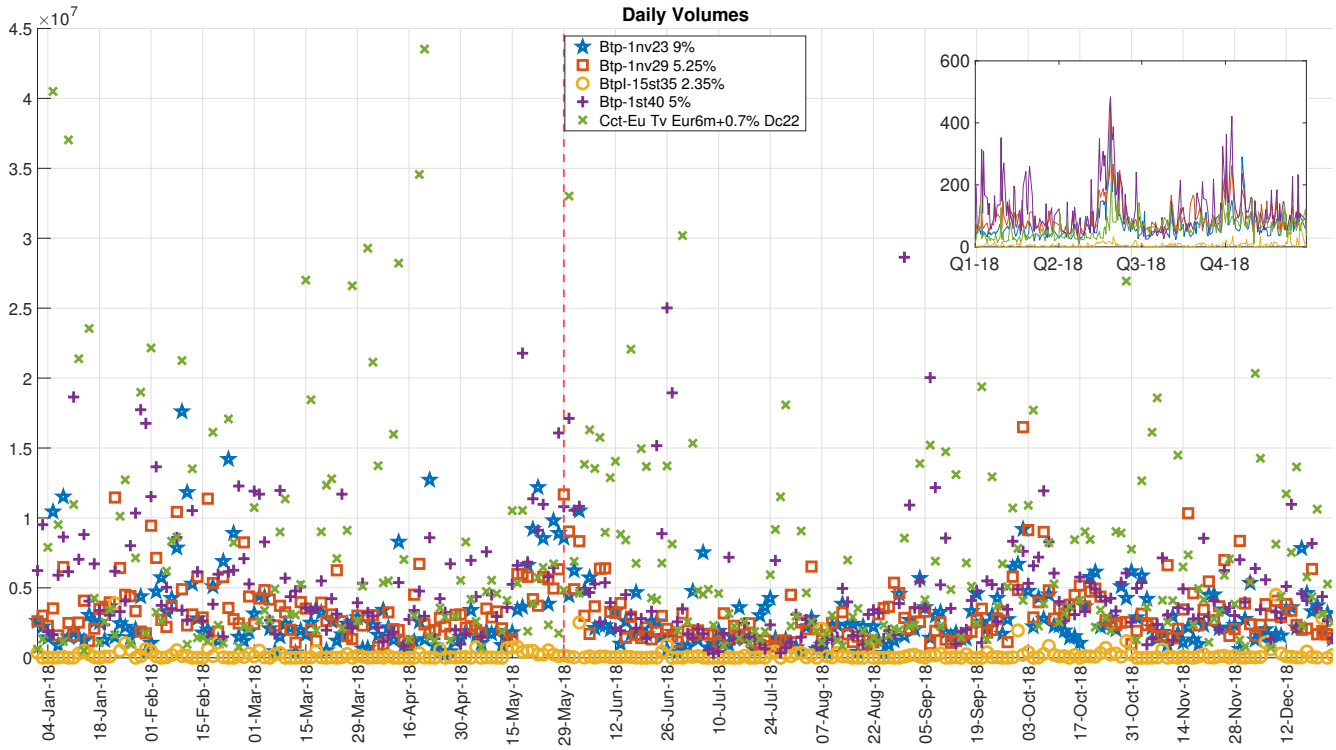

FIGURE 9: Daily volumes observed in the MOT (secondary market) for five representative Italian bonds. In the inset the number of transactions per trading day for the same five bonds. The dashed-red line is May 29, 2018.

8. The pattern is very similar to that of volatility. As expected, the measure peaks in the crash week. Most importantly, the impact of the crash is to increase persistently the illiquidity measure in the market, with a transient effect that lasts several weeks after the crash.

The second statistic we employ is the one proposed in Roll (1984), which we compute day-by-day:

$$\text{Roll} = 2\sqrt{-\text{Cov}\left(\Delta p_{i,t}\Delta p_{i,t-1}\right)}\,, \tag{6.2}$$

where $\Delta p_{i,t}$ is the price difference between two consecutive transactions. This measure can be regarded as a proxy for the effective bid-ask spread: the higher its value, the higher the costs in terms of immediacy for the investors. Panel B of Figure 11 shows the daily value of the Roll measure, which has a very similar dynamics to that of the Amihud* measure, and to that of volatility as well. Again, there is a marked spike during the crash, and a persistent impact on the market illiquidity, with a transient impact that lasts several weeks. A persistent impact on liquidity deterioration was also documented in Boulton et al. (2014) after the flash crash of May 6, 2010.

We thus conclude that the crash of Italian bonds was associated with evaporating market liquidity.

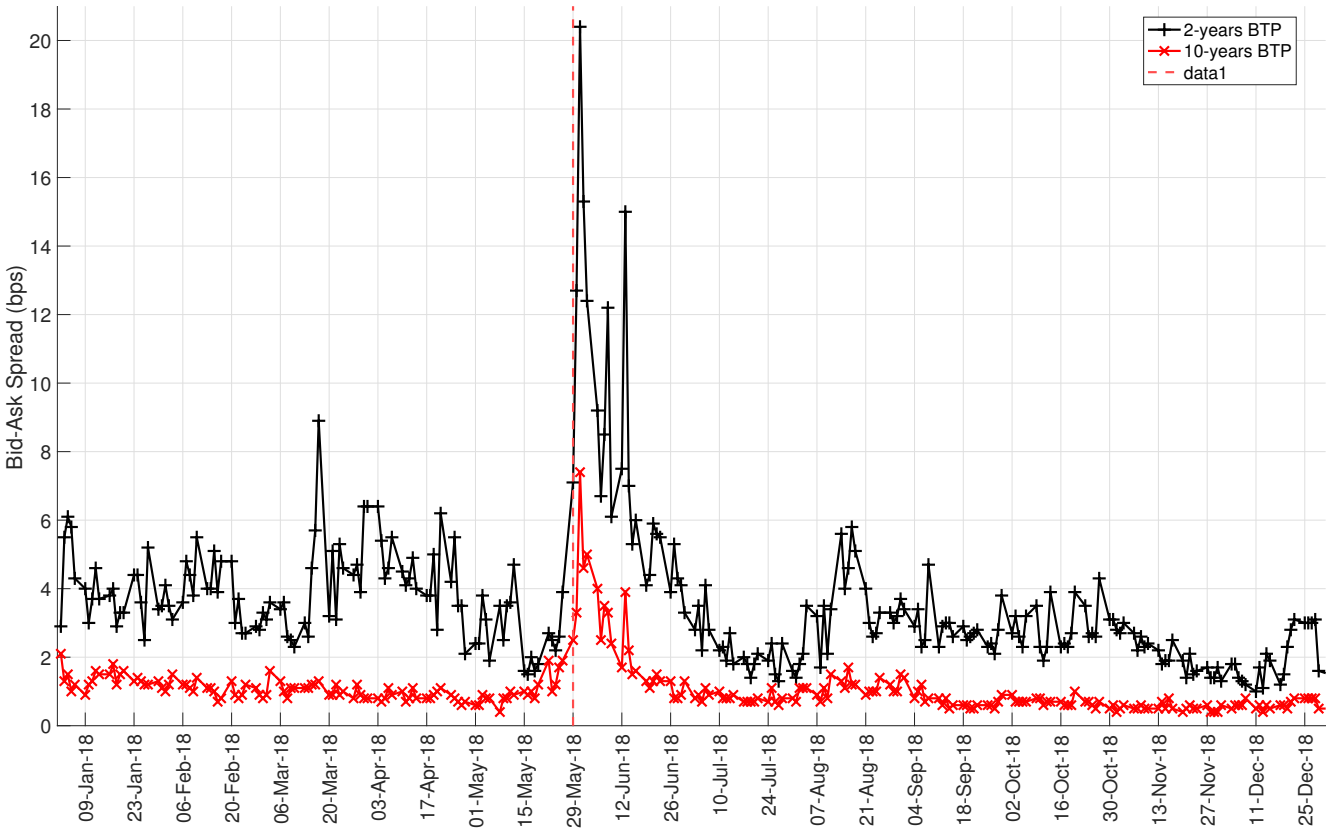

FIGURE 10: Daily bid-ask spreads of two representative Italian bonds. Source: Bloomberg.

Panel A: Amihud

Panel B: Roll

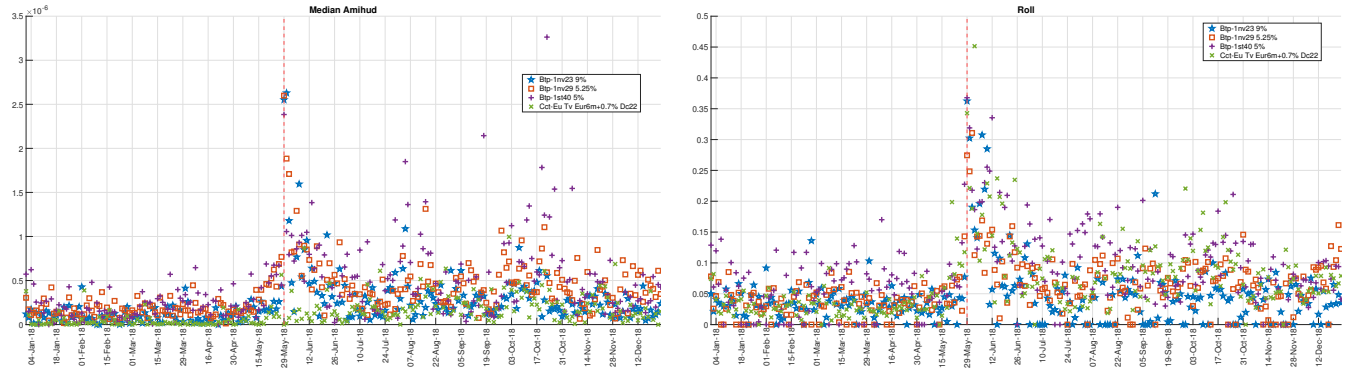

FIGURE 11: Daily median of the Amihud* measure, computed as defined in Eq. (6.1) (panel A), and Roll measure (panel B). The dashed-red line is May 29, 2018.

Reuters[2] reported: "*Several banks have stepped back from primary dealing roles, partly due to regulatory pressures*". Poor market liquidity is particularly relevant during price swings because of what Brunnermeier and Pedersen (2005) call "predatory trading". According to this theory, an informed agent which observes a transient price decline in an illiquid market could exploit the larger

[2]"Dwindling bond liquidity means Italy shock may be just a warning tremor", Reuters, July 3, 2018

price impact (see Figure 11) to sell during decline, even if the fundamental price is higher. This is rational since the high price impact in a deprived market could make the price decline even more, making the later buying more convenient. As mentioned, Menkveld and Yueshen (2019) show that cross-arbitrage may break during a severe liquidity shock, and point at fragmented markets (like the Italian secondary market for Treasury bonds) as a potential source of flash crashes. Their explanation is even more appealing for bond markets, where not only we have fragmentation among trading platforms, but also fragmentation across bonds, since different bonds with close maturity can be regarded as close substitutes.

6.3 V-statistic

The presence of severe liquidity conditions foretells large values of the V-statistic. The neatest economic interpretation of the test comes indeed from the intermediation theory of Grossman and Miller (1988). In their model, a trader looking for immediacy (either informed, or not informed) is willing to sell to M market makers a volume s of a security. Market makers accept to trade immediately but with a price concession, which is given by the formula:

$$\frac{\mu}{\sigma} = \frac{s\gamma}{1+M}\sigma P_0, \tag{6.3}$$

where $\mu = E(P_1/P_0 - 1)$, P_0 is the initial price, P_1 is the price at which the market maker trades, σ is the standard deviation of the price move, and γ is the risk aversion of market makers. Intuitively, the price concession $P_1 - P_0$ will be larger when the traded size is larger, when volatility is larger, when risk-aversion is larger, and when the number of market makers (the unique measure of liquidity in the model) is smaller. Thus, large trading volumes in an illiquid market (precisely the conditions documented in Section 6.2) generate the trends, overshooting and reversals which are captured by the V-statistics. Of course, alternative explanations are possible, as discussed in the introduction. In the case of Italian bonds, the mechanism is reinforced by the auctions (Lou et al., 2013) and by political uncertainty (Tsai, 2018 associates large changes with political turmoils, justified using the model of Kyle, 1985).

Figure 12 reports the daily minimum of the V-statistics on selected bonds. As expected, this statistic is largely negative on May 29 and statistically significant. The figure reports the confidence bands for the BTP-1nv23 only, showing that for this instrument the V-statistics is close to the 1% confidence limit. The p-values for the five statistics, estimated by simulated bootstrap with stochastic volatility as explained in Section 3.1, are 1.5%, 2.4%, 2.2% (for the fixed-rate bonds), 1.8% (for the inflation bond) and 9.8% (for the floating rate bond). It's also interesting to

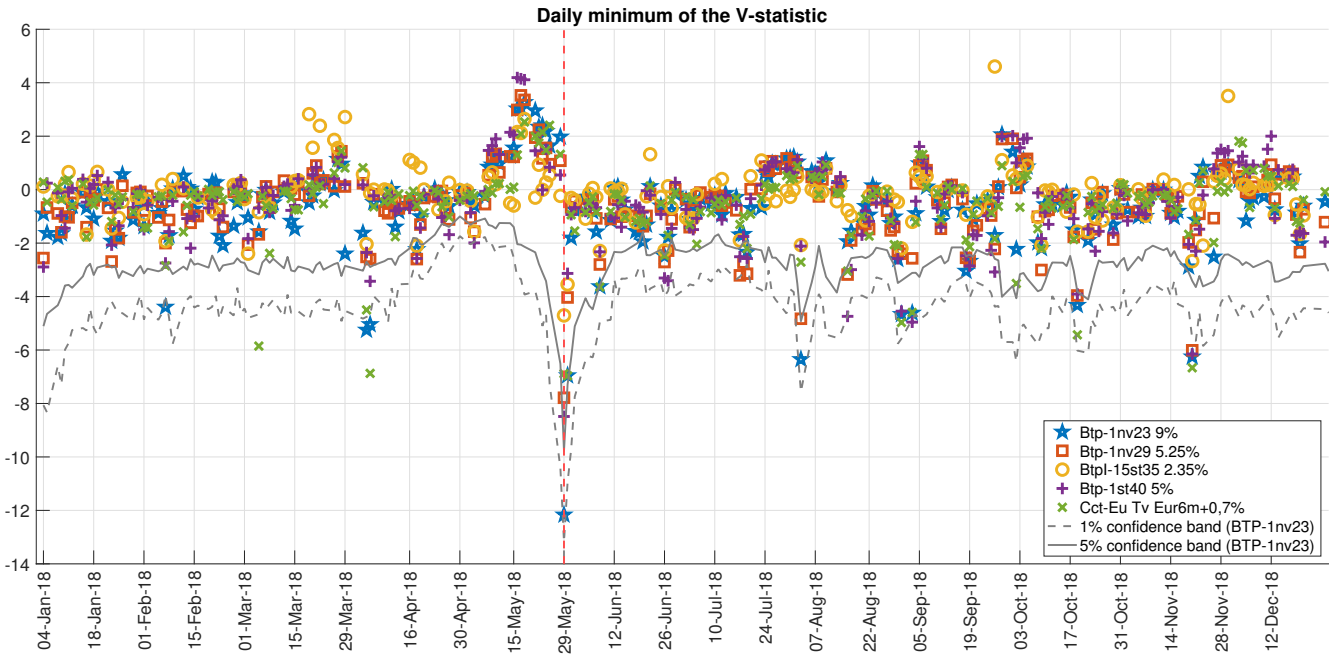

FIGURE 12: Daily minimum of V-statistic computed as in Eq. (3.7) at all intraday transaction times. proposed in this paper, using a bandwidth $h_n = 2$ days. The reported confidence bands are relative to the Btp-1nv23. The vertical dashed-red line is May 29, 2018.

note the positive peak of the V-statistics in mid-May, which, according to the theory, is a signal of strong trend. In light of the discussion in Section 2 and Section 3, this result delivers strong statistical evidence for an inefficient market on May 29.

Commenting on the crash, Financial Times[3] pointed at extreme volatility caused by a deterioration of market liquidity, questioning the proper functioning of the market. The V-statistics show that the crash was not actually due to volatility, but to trend. The distinction between a volatility move and a drift move is not immaterial. Indeed, as argued, large volatility is possible even in an efficient and perfectly liquid market. A V-shape has instead to be associated with distress, that is with prices pushed away from fundamentals because of market frictions.

6.4 Redistribution effects

The Italian Treasury held auctions on May 28, 29 and 30. Details on the auction mechanism are provided in Appendix C. This section quantifies the wealth transfer due to the documented market inefficiency of May 29. The market fragility associated with a V-shape was indeed painfully costly for Italian taxpayers.

For each auctioned title, we define the loss Italian taxpayers had to endure as the difference

[3]"Italian bonds' extreme volatility exposes liquidity strains", Financial Times, June 1, 2018

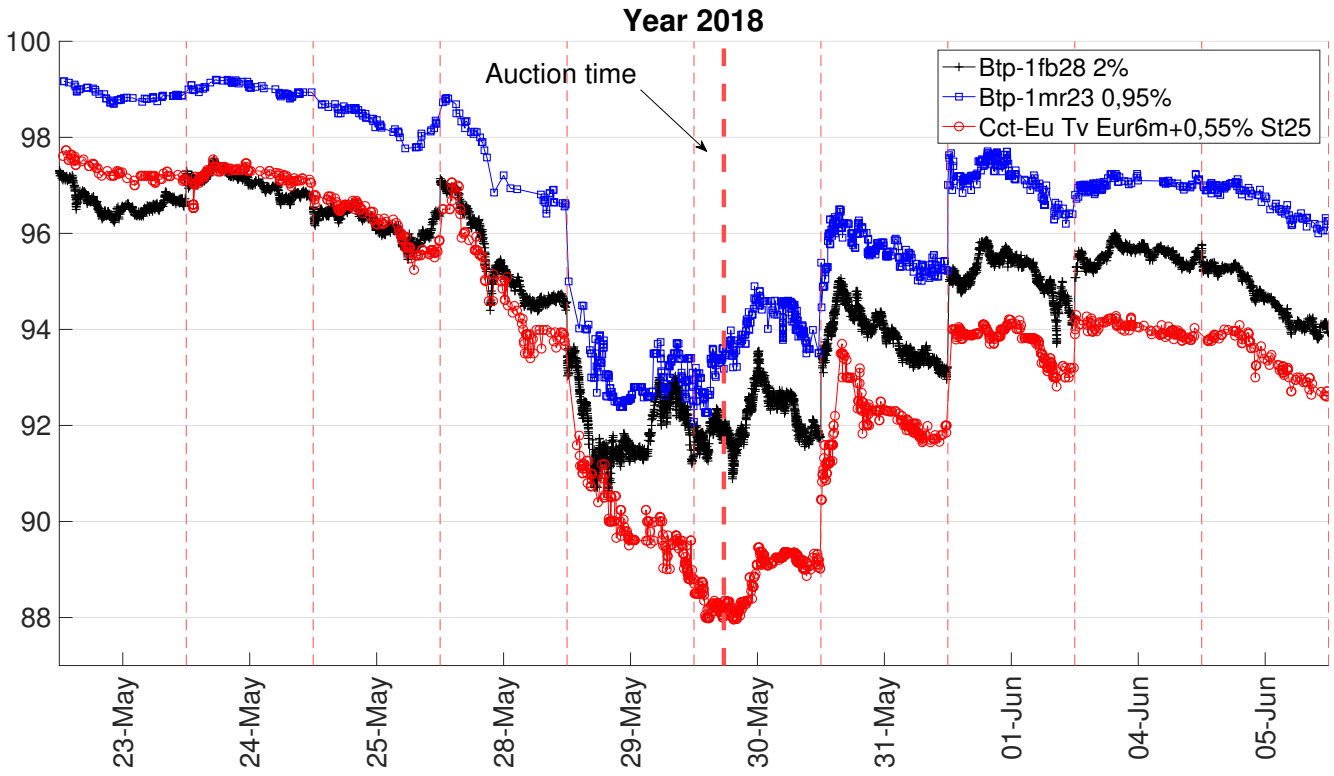

FIGURE 13: Secondary market tick-by-tick transaction prices of the three securities auctioned on May 30, 2018, during the debt crash week. The dashed-red line is the time at which the auction was concluded.

between the realized price P_t and a synthetic price $\tilde{P}_t$, representing the price that would have been traded in a functioning market, times the allocated volume of the issuance, with t equal to the auction date. To determine the distribution of $\tilde{P}_t$ we use a Markov switching model.

We assume the log-return process of each bond is affected by three unobservable regimes $s_t = \{1, 2, 3\}$. We further assume log-returns are normally distributed in all states. The first regime ($s_t = 1$) corresponds to a Business-as-Usual (BaU) state, with zero mean and a given volatility σ_1. The second state ($s_t = 2$) corresponds to the distressed market state, with expected return μ_2 and volatility σ_2. In this second state we allow for a drift different from zero (supposedly, large and negative). Finally, we assume that, after the distressed state, bond prices reach a third regime ($s_t = 3$), where log-returns have again zero mean, but a different (supposedly higher) volatility, σ_3. The higher volatility in the third state is meant to capture the typical behavior of volatility when news arrive, as discussed in Section 2, as well as the positive trend after the crash. The probability for each regime to occur at time t only depends on the regime at time $t-1$, as e.g. in Kole and Van Dijk (2017). Thus, the unobserved state variable s_t follows a three-state first order ergodic Markov chain. We estimate parameters and filtered state probabilities via maximum likelihood.

Figure 14 shows the historical behavior of the Btp-1Feb28 2% log-returns (upper panel), and the

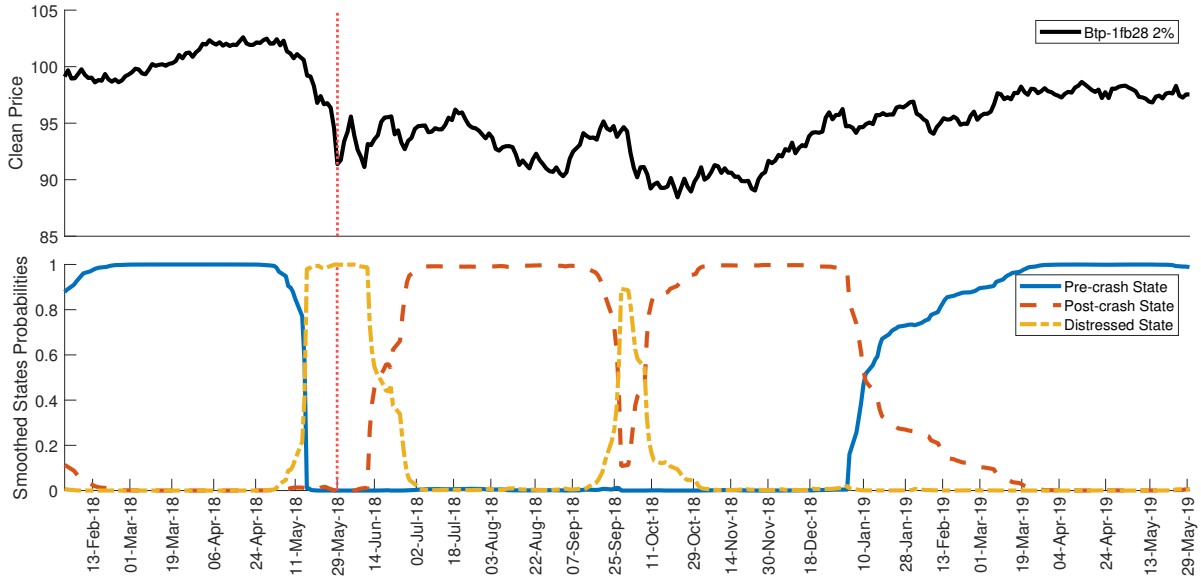

FIGURE 14: Top Panel: Daily prices of Btp-1Feb28 2%. Bottom panel: smoothed state probabilities. The dotted red line indicates 29 May 2018.

smoothed state probabilities for each time t (bottom panel). The dashed red line in the bottom panel indicates the 29 May 2018 date, when the crash occurred. As expected, before the crash occurs, the probability of being in $s_t = 1$ is very high. After that, $s_t = 3$ becomes the most likely state for a long period, until the probability of being in state 1 eventually starts to increase again. Between these two states, we have the distressed period ($s_t = 2$). Similar patterns are observed for the other two bonds.

Table 3 shows the estimates for the parameters in the different regimes. As expected, the volatility estimate for $s_t = 3$ is about twice higher than the one for $s_t = 1$ for all three auctioned instruments, indicating that the distressed event leads to a lingering increased uncertainty in the market. The volatility estimate for $s_t = 2$ is even higher, and the estimate of μ_2 is largely negative (even if hardly significant) for all selected bonds, providing further parametric support to the presence of large local drift (incompatible with an efficient market) during the crash period. Thus, our parametric exercise also confirms the non-parametric empirical evidence shown in Section 6.3.

To estimate the distribution of $\tilde{P}_t$, we assume that a functioning market would have passed from state $s_t = 1$ to state $s_t = 3$, without crossing the distressed state $s_t = 2$. That is, we assume that the average price moves to a new fundamental level, that the volatility spikes, but the inefficient market (that of the right column of Figure 2), in which price overshoots and drift explodes, never takes place. A significant difference of $\tilde{P}_t$ from the actually realized price P_t (observed in a market in which drift is abnormally large, as confirmed by our statistical analysis above) is an additional evidence of broken and inefficient market.

	Btp-1fb28 2%		Btp-1mr23 0.95%		Cct-Eu Tv St25	
	μ_{s_t}	σ_{s_t}	μ_{s_t}	σ_{s_t}	μ_{s_t}	σ_{s_t}
$s_t = 1$	0 -	5.79% (0)	0 -	2.15% $(1.92 \cdot 10^{-5})$	0 -	4.30% $(1.28 \cdot 10^{-4})$
$s_t = 2$	−87.22% (0.0934)	18.69% $(1.61 \cdot 10^{-5})$	−34.55% (0.6492)	22.87% (0.0018)	−126.85% (0.2016)	27.75% (0.0076)
$s_t = 3$	0 -	9.38% $(4.44 \cdot 10^{-16})$	0 -	4.98% $(1.09 \cdot 10^{-11})$	0 -	8.81% $(2.53 \cdot 10^{-8})$

Table 3: Markov Regime Switching annualized parameter estimates for each state. State 1 corresponds to the BaU state, State 2 is the Crash state, and State 3 is the Post-crash unsteady state. P-values are in parentheses.

Following this logic, we simulate 10,000 Brownian bridges from t_1 to t_2, where t_1 and t_2 are, respectively, the end of regime 1 (defined as the first time t when the smoothed state probability of being in $s_t = 1 < 0.95$), and the beginning of regime 3 (defined as the first time t when the smoothed state probability of being in $s_t = 3 > 0.95$). This period includes the crash period, when $\Pr(s_t = 2) > 0.95$. In the simulations, we set the standard deviation of the Brownian Bridge equal to σ_3, that is the volatility in the post-distress state. For each simulated path, $\tilde{P}_t$ is defined as the simulated price at t equal to the auction date. We thus obtain the estimate of the loss for each simulated path. Summing the losses over the three auctioned titles, we obtain the total loss distribution, which is showcased in Figure 15. The estimated loss averages at 434 million euros (145 for Btp-1Feb28 2%, 113 for Btp-1mr23 0.95%, and 176 for Cct-Eu Tv Eur6m+0.55% St25.), and is largely significant, ranging from 250 to 650 millions. This amount of money, according to our estimate, was transferred from the Treasury budget to the primary dealers. To gauge the enormous size of this wealth transfer, it is enough to mention that this single auction costed to the Italian Treasury, *in one day*, roughly the same amount of money ($649 millions) that Lou et al. (2013) estimate the U.S. Treasury loses in *one year* because of the price concessions associated with the supply shocks of repeated auctions.

7 Conclusions

When the market is transitorily moving away from fundamentals, we observe a V-shape in prices, that is a sudden change of the sign of the price drift. The V-statistic that we introduce in this paper can reliably detect such V-shapes, that is it can detect when the market is inefficient. We provide empirical evidence on the presence of frequent such occurrences in the U.S. stock index

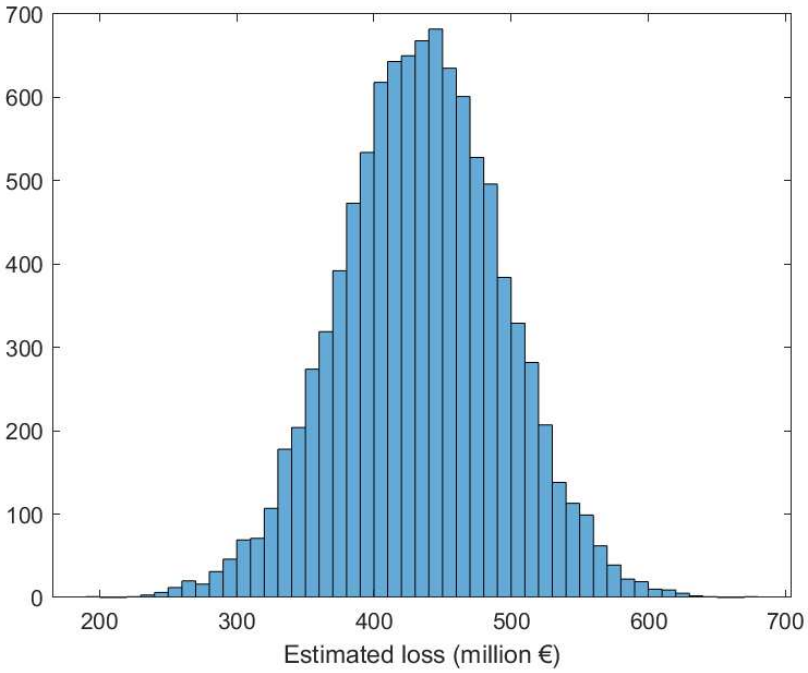

Figure 15: Distribution of the estimated loss for the Italian Treasury due to market inefficiency during the auctions of May 30.

before and during the Covid-19 pandemic. In a specific case-study (the crash of Italian bonds in May, 2018), we show the harmful implications of V-shapes for financial stability. We thus conclude that our analysis enriches regulators, academics and practitioners with a relevant tool for the analysis of financial markets.

References

Allen, F. and D. Gale (2004). Financial fragility, liquidity, and asset prices. *Journal of the European Economic Association 2*(6), 1015–1048.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets 5*(1), 31–56.

Andrews, D. W. K. (1991). Heteroscedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica 59*(3), 817–858.

Bandi, F. (2002). Short-term interest rate dynamics: a spatial approach. *Journal of Financial Economics 65*, 73–110.

Bank of England (2019, July). *Financial Stability Report*.

Barndorff-Nielsen, O., P. Hansen, A. Lunde, and N. Shephard (2008). Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. *Econometrica 76*(6), 1481–1536.

Bates, D. S. (2019). How crashes develop: intradaily volatility and crash evolution. *The Journal of Finance 74*(1), 193–238.

Bellia, M., K. Christensen, A. Kolokolov, L. Pelizzon, and R. Renò (2019). High-frequency trading during flash crashes: Walk of fame or hall of shame? Working paper.

Bernardo, A. E. and I. Welch (2004). Liquidity and financial market runs. *The Quarterly Journal of Economics 119*(1), 135–158.

Biais, B. and T. Foucault (2014). Hft and market quality. *Bankers, Markets & Investors 128*(1), 5–19.

Bollerslev, T. and V. Todorov (2011). Tails, fears, and risk premia. *The Journal of Finance 66*(6), 2165–2211.

Boulton, T. J., M. V. Braga-Alves, and M. Kulchania (2014). The flash crash: An examination of shareholder wealth and market quality. *Journal of Financial Intermediation 23*(1), 140–156.

Brogaard, J., A. Carrion, T. Moyaert, R. Riordan, A. Shkilko, and K. Sokolov (2018). High frequency trading and extreme price movements. *Journal of Financial Economics 128*(2), 253–265.

Brownlees, C. T. and G. M. Gallo (2006). Financial econometric analysis at ultra-high frequency: Data handling concerns. *Computational Statistics & Data Analysis 51*(4), 2232–2245.

Brunnermeier, M. and L. Pedersen (2005). Predatory trading. *The Journal of Finance 60*(4), 1825–1863.

Calcagnile, L. M., G. Bormetti, M. Treccani, S. Marmi, and F. Lillo (2018). Collective synchronization and high frequency systemic instabilities in financial markets. *Quantitative Finance 18*(2), 237–247.

Cespa, G. and T. Foucault (2014). Illiquidity contagion and liquidity crashes. *Review of Financial Studies 27*(6), 1615–1660.

CFTC and SEC (2010). *Findings regarding the market events of May 6, 2010*.

Christensen, K., R. C. Oomen, and M. Podolskij (2014). Fact or friction: Jumps at ultra high frequency. *Journal of Financial Economics 114*(3), 576–599.

Christensen, K., R. C. A. Oomen, and R. Renò (2017). The Drift Burst Hypothesis. Working paper.

Colliard, J.-E. (2017). Catching falling knives: Speculating on liquidity shocks. *Management Science 63*(8), 2573–2591.

Craig, C. C. (1936). On the frequency function of xy. *The Annals of Mathematical Statistics 7*(1), 1–15.

Diebold, F. X. and G. H. Strasser (2013). On the correlation structure of microstructure noise: A financial economic approach. *Review of Economic Studies 80*(4), 1304–1337.

Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. *The Journal of finance 65*(4), 1237–1267.

Golub, A., J. Keane, and S.-H. Poon (2017). High Frequency Trading and Mini Flash Crashes. Working paper.

Grossman, S. and M. Miller (1988). Liquidity and market structure. *Journal of Finance 43*(3), 617–633.

Huang, J. and J. Wang (2009). Liquidity and market crashes. *Review of Financial Studies 22*(7), 2607.

Jacod, J., Y. Li, P. Mykland, M. Podolskij, and M. Vetter (2009). Microstructure noise in the continuous case: the pre-averaging approach. *Stochastic Processes and their Applications 119*(7), 2249–2276.

Jacod, J. and P. E. Protter (2012). *Discretization of Processes* (2nd ed.). Springer-Verlag, Berlin.

Kelly, B. and H. Jiang (2014). Tail risk and asset prices. *The Review of Financial Studies 27*(10), 2841–2871.

Kirilenko, A., A. S. Kyle, M. Samadi, and T. Tuzun (2017). The Flash Crash: High frequency trading in an electronic market. *Journal of Finance 3*, 967–998.

Kole, E. and D. Van Dijk (2017). How to identify and forecast bull and bear markets? *Journal of Applied Econometrics 32*, 120–139.

Kristensen, D. (2010). Nonparametric filtering of the realised spot volatility: a kernel-based approach. *Econometric Theory 26*, 60–93.

Kyle, P. (1985). Continuous auctions and insider trading. *Econometrica 43*, 1315–1335.

Laly, F. and M. Petitjean (2020). Mini flash crashes: Review, taxonomy and policy responses. *Bulletin of Economic Research 72*(3), 251–271.

Lou, D., H. Yan, and J. Zhang (2013). Anticipated and repeated shocks in liquid markets. *The Review of Financial Studies 26*(8), 1891–1912.

Madhavan, A. N. (2012). Exchange-traded funds, market structure and the Flash Crash. *Financial Analysts Journal 68*(4), 20–35.

Menkveld, A. J. and B. Z. Yueshen (2019). The flash crash: A cautionary tale about highly fragmented markets. *Management Science 10*(10), 4470–4488.

Roll, R. (1984). A simple measure of the implicit bid-ask spread in an efficient market. *Journal of Finance 39*, 1127–1139.

Schinasi, G. J. (2004). Defining financial stability. IMF Working Paper.

Schneider, M. and F. Lillo (2019). Cross-impact and no-dynamic-arbitrage. *Quantitative Finance 19*(1), 137–154.

Sigaux, J.-D. (2018). Trading ahead of Treasury auctions. ECB Working Paper.

Tsai, I.-C. (2018). Flash crash and policy uncertainty. *Journal of International Financial Markets, Institutions and Money 57*, 248–260.

Weller, B. M. (2019). Measuring tail risks at high frequency. *The Review of Financial Studies 32*(9), 3571–3616.

A Technical Appendix.

This appendix contains technical details about the model in Eq. (3.1), as well as the proof of Theorem 1.

Assumption 1. *i) X is defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathcal{P})$ and assumed to be an Itō semimartingale described by the dynamics in Eq. (3.1), where X_0 is $\mathcal{F}_0$-measurable, μ_t is a locally bounded and predictable drift, σ_t is locally bounded, adapted, càdlàg and strictly positive (almost surely) volatility, $c^-_{1,t}$, $c^+_{1,t}$, $c_{2,t}$ are continuous and twice differentiable deterministic functions, $W = (W_t)_{t\geq 0}$ is a standard Brownian motion and $J = (J_t)_{t\geq 0}$ is a pure-jump process.*
ii) The jump process J_t is of the form:

$$J_t = \int_0^t \int_{\mathbb{R}} \delta(s,x) I_{\{|\delta(s,x)|\leq 1\}} \left(\nu(ds,dx) - \tilde{\nu}(ds,dx)\right) + \int_0^t \int_{\mathbb{R}} \delta(s,x) I_{\{|\delta(s,x)|>1\}} \nu(ds,dx), \tag{A.1}$$

where ν is a Poisson random measure on $\mathbb{R}_+ \times \mathbb{R}$, $\tilde{\nu}(ds,dx) = \lambda(dx)ds$ a compensator, and λ is a σ-finite measure on $\mathbb{R}$, while $\delta : \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}$ is predictable and such that there exists a sequence $(\tau_n)_{n\geq 1}$ of $\mathcal{F}_t$-stopping times with $\tau_n \to \infty$ and, for each n, a deterministic and nonnegative Γ_n with $min(|\delta(t,x)|, 1) \leq \Gamma_n(x)$ and $\int_{\mathbb{R}} \Gamma_n(x)^2 \lambda(dx) < \infty$ for all (t,x) and $n \geq 1$.
iii) Fix $t \in (0,T]$ and let $B_\epsilon(t) = [t-\epsilon, t]$ with $\epsilon > 0$ fixed. We assume there exists a $\Gamma > 0$ and a sequence of $\mathcal{F}_t$-stopping times $\tau_m \to \infty$ and constants $C_t^{(m)}$ such that for all m, $(\omega, s) \in \Omega \times B_\epsilon(t) \cap [0, \tau_m(\omega)[$, and $u \in B_\epsilon(t)$,

$$E_{u\wedge s}\left[|\mu_u - \mu_s|^2 + |\sigma_u - \sigma_s|^2\right] \leq C_t^{(m)} |u-s|^\Gamma, \tag{A.2}$$

where $E_t[\cdot] = E[\cdot|\mathcal{F}_t]$.

To implement the test, we need a kernel and a bandwidth sequence, satisfying the following assumptions:

Assumption 2. *The bandwidth h_n is a sequence of positive real numbers, such that, as $n \to \infty$, $h_n \to 0$, $nh_n \to \infty$. The kernel $K : \mathbb{R} \to \mathbb{R}_+$ is any function with the properties:*

(K1) K is bounded and differentiable with bounded first derivative; further, $xK(x) \to 0$ and $xK'(x) \to 0$ as $x \to \pm\infty$.

(K2) $\int_{-\infty}^{+\infty} K(x)\mathrm{d}x = 1$ and $K_2 = \int_{-\infty}^{+\infty} K^2(x)\mathrm{d}x < \infty$;

(K3) It holds that for every positive sequence $G_n \to \infty$, $\int_{|x|>G_n} K(x)\mathrm{d}x \leq CG_n^{-B}$ for some $B > 0$ and $C > 0$ (i.e., K has a fast vanishing tail);

(K4) $m_K(\alpha) = \int_{-\infty}^{+\infty} K(x)|x|^\alpha \mathrm{d}x < \infty$, for all $\alpha > -1$; $m'_K(\alpha) = \int_{-\infty}^{+\infty} K^2(x)|x|^\alpha \mathrm{d}x < \infty$, for all $\alpha > -1$.

We consider a left-sided kernel $K^-(x)$ satisfying Assumption 2 with the additional property $K(x) = 0$ when $x > 0$, and a right-sided kernel satisfying Assumption 2 with the additional property $K(x) = 0$ when $x < 0$.

We need an assumption regarding observation times:

Assumption 3. *$(t_i)_{i=0}^n$ is a deterministic sequence. We denote by $\Delta_{i,n} = t_i - t_{i-1}$, $\Delta_n^- = \min_{i=1,\dots,n}\{\Delta_{i,n}\}$, $\Delta_n^+ = \max_{i=1,\dots,n}\{\Delta_{i,n}\}$. We assume that, for a sufficiently large n, and suitable constants C_1, C_2 which do not depend on n,*

$$C_1\Delta_n \leq \Delta_n^- \leq \Delta_n^+ \leq C_2\Delta_n,$$

where $\Delta_n = T/n$. Moreover, denoting the "quadratic variation of time up to t" as $H(t) = \lim_{n\to\infty} H_n(t)$, where $H_n(t) = \frac{1}{\Delta_n}\sum_{t_i\leq t}(\Delta_{i,n})^2$, we assume $H(t)$ exists and is Lebesgue-almost surely differentiable in $(0,T)$ with derivative H' such that: $|H'(t_i) - \Delta_{i,n}/\Delta_n| \leq C\Delta_{i,n}$, for any t_i in which H is differentiable, for a suitable constant $C \geq 0$ which does not depend on i and n.

Proof of Theorem 1. Consider the $T_{\tau,n}^-$ statistic in Eq. (3.2). Using the method of proof of Theorem 2 of COR, we get:

$$T_{\tau,n}^- = \sqrt{h_n}\frac{O_p(h_n^{-\alpha})}{\left(O_p(h_n^{-2\beta}) + O_p\left(\frac{1}{n^{2-2\alpha}h_n}\right)\right)^{\frac{1}{2}}}$$

The term $O_p\left(\frac{1}{n^{2-2\alpha}h_n}\right)$, coming from the bias, is vanishing because of the stated assumption, so that $T_{\tau,n}^-$ diverges with rate $h_n^{1/2-\alpha+\beta}$. The same applies to $T_{\tau,n}^+$, hence when $c^\pm \neq 0$, $\mathcal{V}_{\tau,n}$ is of order $h_n^{3/2-2\alpha+2\beta}$ which diverges if and only if $\alpha - \beta > 3/4$. If instead $c^\pm = 0$, then $\mathcal{V}_{\tau,n}$ is at most of order $h_n^{1-\alpha+\beta} \to 0$. □

We finally provide the assumption on market microstructure noise in Eq. (3.8).

Assumption 4. *$(\varepsilon_{t_i})_{i=0}^n$ is adapted and independent of X. Moreover, $E[\varepsilon_{t_i}] = 0$, $E[(\varepsilon_{t_i})^4] < \infty$, and denoting by $\gamma_k = E[\varepsilon_{t_i}\varepsilon_{t_{i+k}}]$ for any integer $k \geq 0$, we further assume γ_k is finite, independent of i and n, such that $\gamma_k = 0$ for $k > Q$, where $Q \geq 0$ is an integer (i.e., Q-dependent noise).*

B Cleaning of intraday prices

We clean high-frequency transactions according to the following procedure:

a. for each trading day, we discard observations three times larger than the daily price median;

b. at the day-level, we then implement the Brownlees and Gallo (2006) filter to filter out outliers: we keep the j^{th} observation if

$$|p_j - \bar{p}_j(k)| < 3\sigma_j(k) + \gamma\,, \tag{B.1}$$

where $\bar{p}_j(k)$ and $\sigma_j(k)$ denote the δ-trimmed sample mean and standard deviation, respectively, of a neighborhood of k observations around j, while γ is the so-called granularity parameter. We select $k = 50$ observations, $\gamma = 0.02$ (twice the minimum tick), and $\delta = 0.9$.

c. We aggregate transactions with the same time-stamp. We substitute simultaneous tick-by-tick prices with the volume-weighted average price, and simultaneous tick-by-tick volumes with the sum of the simultaneous volumes.

C The case study: data, macroeconomic background and the auction mechanism

C.1 Data

We use tick data for a subsample of Italian government securities traded on the MOT (*Mercato Obbligazionario Telematico*), the electronic Italian-regulated limit order book market for sovereign, bank and corporate bonds. MOT is a retail exchange characterized by many transactions with small volume. It is the main retail trading venue by volume for Italian government bonds (8.25% of all trades on platforms in 2018), even if its volume is much lower than that of the two wholesale platforms, MTS Cash and MTS BondVision (91.24% of all trades on platforms in 2018)[4]. Despite its relatively thin volume, the high number of transactions in the MOT guarantees that absence of cross-market arbitrage and fair security pricing is broadly guaranteed within the bid-ask spreads (Schneider and Lillo, 2019).

The daily trading schedule on the MOT is divided in two segments: an opening auction, from 8:00 a.m. to 9:00 a.m., followed by a continuous trading phase, from 9:00 a.m. to 5:30 p.m.. The opening price is determined during the opening auction phase. We only focus on the continuous trading session, and exclude opening auction activity from the analysis. We select a representative set of Treasury bonds among BTPs (fixed coupon), CCTs (floating + fixed coupon), and BTPi (inflation linked bonds), namely:

- a 10Y BTP with maturity 2023, and 9% coupon rate (BTPs pay semi-annual coupon), BTP-1nv23 9%,
- a 30Y BTP with maturity 2029 and 5.25% coupon rate, BTP-1nv29 5.25%,
- a 30Y BTP with maturity 2040 and 5% coupon rate, BTP-1st40 5%,
- an inflation-linked 30Y BTP with maturity on 2035 and 2.35% coupon rate, BTPi-15st35 2.35%,
- a 7Y CCT with maturity 2022, and with an Euribor-linked coupon rate, CCT-Eu Tv Eur6m+0.7% Dc22,

We also use an additional sample of three government bonds issued in 2018 and auctioned on May 30, 2018:

- a 10Y BTP with maturity 2028 and 2% coupon rate, BTP Tf 2,00% Fb28 Eur
- a 5Y BTP with maturity 2023 and 0.95% coupon rate, BTP Tf 0.95% Mr23 Eur,
- a 7Y CCT with maturity 2025 and an Euribor-linked coupon rate, CCT-Eu Tv Eur6m+0.55% St25.

[4]See CONSOB, Bollettino Statistico n. 14, June 2019, available at `http://www.consob.it/web/area-pubblica/bollettino-statistico`.

We consider the period from January 1, 2018 (except for the last three securities, whose data start on the first issue date, that is January 31, 2018, February 28, 2018 and May 2, 2018, respectively) to May 30, 2019. The data were provided by Borsa Italiana S.p.A., and are recorded with millisecond time-stamps. All transactions are at the clean price and are cleaned using the procedure delined in Appendix B.

To deal with overnight gaps, we construct a new time vector for each time series, associated to the original one, where the time (in milliseconds) elapsed from the closing of day $t-1$ to the next available open price is equal to

$$\tilde{t} = \frac{\sigma_{overnight}}{\sigma_{intraday}}\,\hat{t}\,.$$

Here, $\sigma_{overnight}$ is the standard deviation of overnight returns, defined as the price appreciation or depreciation between market close of day $t-1$ and market open of day t, while $\sigma_{intraday}$ is the standard deviation of intraday returns, defined as the price appreciation or depreciation between market open and close of the same day. Finally, $\hat{t}$ is the time, in milliseconds, elapsed from market open (9 a.m.) to close (5:30 p.m.). We implement the test statistic (3.2) and (3.5) at each point in time of the previously defined time vector with abridged overnight times, which we use in lieu of the original time series one; we choose a bandwidth h_n for the drift equal to 2 days[5], while we base the long-run variance of the drift estimate on a 10-day bandwidth. We adopt exponential kernels $K^-(x) = \exp(-|x|)\,\mathbb{1}_{\{x<0\}}$ (left-sided), and $K^+(x) = \exp(-|x|)\,\mathbb{1}_{\{x>0\}}$ (right-sided). Given that we use all transactions, to account for the presence of market microstructure noise we replace the volatility estimates in the denominator of $T^-_{\tau,n}$ and $T^+_{\tau,n}$ with an HAC estimator of the long-run variance of the local drift estimator, applied to pre-averaged returns, as recommended in COR.

C.2 Macroeconomic background

The Italian Treasury huge bond price movement in the week from May 28 to June 1st reflected the country's political uncertainty in the aftermath of a radical government change which was taking place exactly in that week. On March 4, 2018 Italy held its political elections. The centre-right coalition, led by the League party, got the majority of votes, while the Five Star Movement was the most voted party. However, no political group won an outright majority, leading to a hung parliament. After almost three months of political gridlock, the League and the Five Star Movement reached an agreement and joined a coalition. On May 21, the two parties indicated Mr. Giuseppe Conte, a law professor with no former political experience, as designated Prime Minister. Two days after, Mr. Conte was granted the mandate to form a new cabinet from the Italian President Mr. Mattarella. The leaders of the two coalition parties strongly pushed for the appointment of Mr. Paolo Savona as Minister of Treasury, despite the rumored opposition of the Italian President because of his alleged anti-euro positions. On Sunday, May 27, the leaders of the two coalition parties gave Mr. Mattarella, who must endorse the cabinet, an ultimatum, which triggered him to reject the nomination. As a consequence, during the evening of the

[5]Results, not shown here, are pretty robust to the choice of h_n.

same day, Mr. Conte dropped his bid to form a government. On May 28, Mr. Mattarella appointed a former International Monetary Fund official, Mr. Cottarelli, as designed prime minister, and asked him to form a new cabinet. However, both the Five Star Movement and the League announced their intention not to support a vote of confidence for the new designated Prime Minister, which would have triggered new immediate elections. Finally, on May 31, they agreed upon the composition of the new government in which Mr. Savona was appointed Minister of European Affairs and Mr. Giovanni Tria Minister of Treasury, and on June 1 the new Conte cabinet was formally sworn in.

C.3 The auction mechanism of Italian government bonds

Before trading in the secondary markets, Italian government bonds are first issued on the primary market, where securities are allotted through electronic marginal auctions. The size of the allocations in the primary market in Italy is fairly predictable, since an annual auction calendar with a regular schedule is published at the beginning of the year (Sigaux, 2018). At regular intervals (usually monthly), the Treasury adds to an existing issue (new tranches), so to increase the outstanding amount. In this respect, auctioning of Italian bonds is similar to that of the U.S. Treasury (Lou et al., 2013). After each ordinary offering, usually on the day after, there is a supplementary placement, at the same price of the ordinary offering, reserved to specialists. This is technically considered a subsequent tranche. For each issued bond, there are usually between 12 and 18 tranches. The amount issued is communicated to authorized dealers approximately 3 days before the auction takes place. The main auction concludes at 11:00 of the due date. Most of the bids come to the Treasury in the few minutes before, to exploit the information in the secondary market at best.

During the day following the bonds crash, and precisely at 11:00 a.m. on May 30, three (ordinary) auctions took place: the 9^{th}, 7^{th}, 3^{rd} tranches of BTP Tf 2,00% Fb28 Eur, BTP Tf 0.95% Mr23 Eur, and CCT-Eu Tv Eur6m+0.55% St25 Eur, respectively, were issued. These are the last three debt securities we selected and listed in Section C.1. The total nominal value of the three issuances were 2.159, 2.013 and 2.3 billion euros. Figure 13 shows the price path of the three auctioned instruments from May 23 to June 5. Other auctions took place on May 28 and 29 on a two-years zero coupon bond (CTZ) and two inflation-linked bonds (BTPi), and on May 29 and 30 on a 6-month zero coupon bond (BOT). We assume that the Auctions of May 28-29 were unaffected, the auction price being fixed on the 28th. The 6-months BOT was issued at the sky-rocketing yield of 1.23% (it was -0.421% in the previous issuance of May 26, and 0.092% in the following issuance of June, 27). However the loss due to this specific issuance is negligible, and it can be estimated in roughly 0.5 million euros.