V-shapes*

Maria Flora[‡]

Roberto Renò§

November 7, 2022

Abstract

We propose a new methodology to define short-term price reversals (V-shapes), devised to detect flash crashes from price records. Our empirical application shows that the detection methodology we propose is consistent with the forensic definition used by the SEC in legal charges for market access rule violation. Moreover, we show that recent years have seen an increase in the frequency and severity of flash crashes, and that increased algorithmic trading can be associated with these events. We also show that transient inefficiencies are not necessarily short-lived and imply significant wealth redistribution when coupled with frictions such as a supply shock.

JEL classification: G14, G12, C58;

Keywords: Price reversals, overreaction, flash crashes, market violations, algorithmic trading, financial fragility, auction cycle.

^{*}We thank Giacomo Bormetti, Antoine Bouveret, Kim Christensen, Fulvio Corsi, Carsten Gerner-Beuerle, Gerardo Ferrara, Aleksey Kolokolov, Offer Lieberman, Michele Manna, Roberto Marfé, Antonio Mannolini, Gaetano Marseglia, Claudio Pacati, Alfonso Puorro, Lorenzo Schoenleber, Giuseppe Trapani, Olivier David Zerbib, the participants at: the Midwest Finance Association (MFA) 70th Annual Meeting (March 18-20, 2021), and in particular the discussant, Konstantin Sokolov; the 13th Society of Financial Econometrics (SoFiE) Annual Conference (June 14-17, 2021) and in particular the discussant, Bas Werker; the Financial Management Association 2021 Annual Meeting (October 20-23, 2021) and in particular the discussant, Huidi Lin; the QFFE 2022 in Marseille (June 16-17, 2022); the PRIN Kick-off Workshop (Venice, March 21-22, 2022); the 14th International Risk Management Conference (IRMC) (October 1-2, 2021); the XXI Workshop on Quantitative Finance (Napoli, January 29-31 2020); the XLIV AMASES Conference (Padova, September 18 2020); the Ninth Italian Congress of Econometrics and Empirical Economics (Cagliari, January 21-23, 2021); and the participants at the seminar at: Bank of Italy (2019); CREST, ENSAE (2021); Collegio Carlo Alberto (2021); Warwick Business School (2021); ETH, Zurich (2021); European Securities and Markets Authority (2021); University of Cergy (2022); IESEG (2022); ESSEC Business School (2022); GFRI, Geneve (2022); and University of Lugano (2022); for comments and discussions. We thank Borsa Italiana S.p.A. for providing the transaction data on the secondary Italian bond market. R. Renò acknowledges PRIN funding from MIUR, Italy, HiDEA Project 2017RSMPZZ. 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. All errors and omissions are sole responsibility of the Authors.

[‡]Corresponding Author. CREST, CNRS, Institut Polytechnique de Paris, e-mail: MARIA.FLORA@ENSAE.FR

[§]ESSEC Business School and University of Verona, Department of Economics, e-mail: RENO@ESSEC.EDU

1 Introduction

Inefficient markets and market fragility are a threat to the financial system. A now well-recognized form of market inefficiency first hit the headlines with the Flash Crash of May 6, 2010 in the U.S. stock market. This type of events is characterized by a price that overshoots past fundamental values, and then rebounds back, tracing a V-shape. In their recent ESMA report, Bouveret et al. (2022) write that "flash events do raise questions about how developed markets react under stressed circumstances, especially if prolonged in time." However, it is hard to determine when, and for how long, market prices are away from fundamentals. Existing methods to detect transient inefficiencies of this kind mainly rely on detecting extreme returns or significant serial correlation, e.g. using variance ratios. Nevertheless, neither extreme price volatility nor variance ratios can easily disentangle overreaction from underreaction. Indeed, they both generate high volatility and positive serial correlation; and return-based measures of price reversals are plagued by the noise inherent in volatile prices.

This paper adds to the growing literature on flash crashes, and similar price reversals, by providing a novel methodology for their detection. We develop a strategy to decrypt additional information hidden in the traded price dynamics based on the distinction between a "drift move" (due to price trends) and a "variance move" (due to price uncertainty). Among the "drift move" cases, we are able to detect sudden drift changes, which leads to our formal definition of a V-shape. The main advantage of this strategy with respect to alternative techniques for the detection of price reversals is that our measure can effectively discern between overreaction and underreaction.

We present several applications of our methodology. First, we show that our method can be used to locate reliably instances of market abuse. We thus provide market authorities with a tool that can be useful for monitoring financial markets and enhance their orderly functioning. Second, we deliver a rigorous description on the extent of mini-flash crashes in the US equity index, showing that these events increased in both frequency and severity during the period considered (2013-2020) and that their increase was associated with increased algorithmic trading activity. Finally, we show that V-shapes are not limited to flash crashes, but that they also comprise transient inefficiencies due to well known market frictions. We provide a specific example in the sovereign bond market in which the friction is the auction cycle, and we highlight its implications in terms of wealth redistribution.

Figure 1 frames our intuition. The rows of Figure 1, from top to bottom, exhibit the expected price, price volatility, and price drift respectively. The first two columns (Model 1 and Model 2) represent well-known forms of tail risk, namely price jumps and volatility spikes. Model 2

¹This study by the European Securities and Markets Authority (ESMA) uses the V-statistic proposed in this paper to study flash crashes in the cash and futures markets of European sovereign bonds.

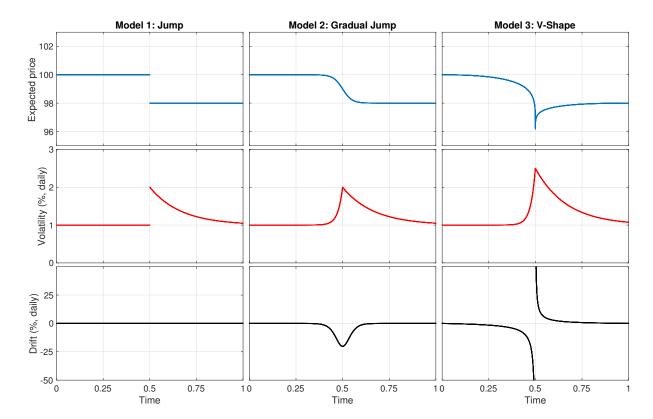


FIGURE 1: Stylized representation of three idealized markets hit by news. Left column: jump. Center column: gradual jump (underreaction). Right column: V-shape (overreaction). Each case is characterized by the expected price (top row), price volatility (center row) and price drift (bottom row). Precise formulations of the models are provided in Section 5.

displays underreaction. We add on the existing literature on tail risk (see, e.g., Bollerslev and Todorov, 2011; Kelly and Jiang, 2014; Weller, 2019) by focusing on Model 3 (the V-shape), where the expected price overreacts. As it is clear from Figure 1, a V-shape is characterized by a discontinuity in the sign of the price drift, which is the distinctive feature that differentiates it from the efficient ups-and-downs of erratic prices. Using this novel insight, and building on the technical results on drift bursts of Christensen, Oomen, and Renò (2022) (henceforth, COR), we develop the *V-statistic*, a new tool that can reveal the overshooting in prices. The focus on the drift proves to be crucial to set our V-statistic apart from existing measures of distress, typically based on serial correlation, volatility or price jumps. None of these measures, including the variance ratio (Lo and MacKinlay, 1988), can distinguish Model 2 from Model 3, as illustrated in Figure 1 and shown in our simulation experiments.

In the data, V-shapes come in different durations, and this paper provides a set of empirical applications that encompasses events that last only fractions of seconds, like mini-flash crashes (using the terminology of Biais and Foucault, 2014), the Flash Crash of May 6, 2010, with a duration of about 30 minutes, and the auction cycle in bond markets (Lou et al., 2013), which can last several hours.

In our first application, we analyze regulatory charges against financial intermediaries for violation of the Market Access Rule. Our measure reinforces and complements the Securities and Exchange Commission (SEC) evaluation, and proves to be a reliable tool for detecting such events. In particular, we prove that our econometric definition of a mini-flash crash matches the forensic definition used by the SEC in (successful) legal charges against financial intermediaries which, in the SEC words, "harmed the integrity of trading on the securities market" (U.S. Securities and Exchange Commission, 2016). We exploit the intraday resolution of our measure and extend the analysis to the U.S. stock market in the years 2013-2020. We show that mini-flash crashes increased over time, and we relate them to increased algorithmic activity, a persistent increase in non-fundamental volatility and to deteriorated liquidity conditions. The increase is statistically significant, and our simulation experiments, showing that the test is correctly sized, imply that the likelihood of this result being due to false positives is negligible.

Next, we analyze two iconic events. The first one is the eponymous Flash Crash of May 6, 2010. There is strong consensus on the fact that the whole U.S. stock market was inefficient on that day (see, e.g., Madhavan, 2012; Kirilenko et al., 2017, and Menkveld and Yueshen, 2019). This event was indeed triggered by a huge, non-fundamental selling trade in the E-mini futures market (CFTC and SEC, 2010). We show that the V-statistic provides a highly significant signal on the Flash Crash day. Then, as a counterfactual, we analyze the U.S. stock market during the first wave of the COVID-19 pandemic. After the news about rising COVID-19 cases spread out, the market rapidly collapsed and then recovered. The apparent V-shape in the traded price data is however not confirmed as a significant V-shape by our test, suggesting that prices during the first pandemic wave incorporated available information efficiently, at least at the hourly frequency (while we document many flash crashes at higher frequencies).

Finally, we demonstrate that our methodology can be used to detect longer-lived V-shapes caused by classical price anomalies. The event we study is the May 2018 crash in the Italian bond market in a period of political turmoil. We show that sovereign bond prices deviated abnormally while they were auctioned in the primary market. The price reversal lasted several hours. Using a regime-switching model, we estimate that, during the auctions of May 30, about 0.45 billion euros were transferred from the Italian Treasury to primary dealers. The existence of an auction premium (or auction cycle) is well established and rationalized (Lou et al., 2013), and it represents an example of market anomaly that can give rise to a V-shape. In our specific case, what is striking 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 existing theories of the auction cycle.

The paper is organized as follows. Section 2 reviews the literature related to this paper. In

Section 3 we introduce the V-statistic. Section 4 describes its implementation with real data. Section 5 presents results on simulated data meant to illustrate the size and power of the V-statistic under the null and the alternative, and to compare the V-statistic to the variance ratio. Section 6 contains the empirical applications. Section 7 concludes.

2 Literature Review

We relate to the vast economic theory of price formation with frictions leading to market anomalies such as overreaction or price cycles. Several explanations are indeed offered in the literature for the presence of price reversals, typically generated by the strategic interaction of traders with different objectives in a frictional environment. Frictions considered by theoretical papers that may lead to a V-shape include immediacy costs, asymmetric information, anticipated supply shocks, slow moving capital and market fragmentation.

One celebrated example of a mechanism producing a V-shape is presented in Grossman and Miller (1988), who 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 leading to a V-shape 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; 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; 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; 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; and Lou et al. (2013), as discussed, show the presence of an auction cycle of bond prices around Treasury auctions, due to risk-aversion of primary dealers and limited capacity of arbitrageurs. We complement this literature by providing a new reliable measure of market distress that can be used in event studies to shed empirical light on economic theory.

The common practice to detect price reversals from price records is related to detecting market inefficiency (for a discussion on market efficiency, see Campbell et al., 1997, Chapter 1, Section 5, and Chapter 2 for a review of inefficiency measures) and mostly relies on serial correlation

(Lo and MacKinlay, 1988; Poterba and Summers, 1988; Campbell et al., 1993; Chaudhuri and Wu, 2003; Griffin et al., 2010). Financial distress has also often been identified with price dispersion measures, such as high volatility (Shiller, 1981; Christensen et al., 2014; Brogaard et al., 2018; Bates, 2019) and price jumps (Bates, 2019; Calcagnile et al., 2018). We deviate substantially from this literature by focusing on the drift part of the price dynamics. We argue that sustained volatility or jumps are actually compatible with efficient markets and do not imply the presence of price reversals. Moreover, measures based on serial correlation, like the variance ratio, can fail to tell apart overreaction from underreaction. Drift is instead a much more reliable signal of distress, and the way we use it is able to isolate V-shapes from concurrent dynamics which still generate significant serial correlation.

The growing empirical literature on flash crashes shed light on a market vulnerability that now appears to be endemic to all financial markets (Gao and Mizrach, 2016; Golub, Keane, and Poon, 2017; Laly and Petitjean, 2020). We contribute to this literature by providing a precise quantification of the phenomenon in recent times. Theoretical literature also tries to understand the source of frequent flash crashes. For example, Cespa and Vives (2022) investigate the role of lack of transparency in generating illiquidity spirals; Easley et al. (2012) investigate the role of order flow toxicity. The precise mechanism leading to a flash crash is still unknown, which makes hard for regulators to design remedies. Our paper provides a tool which helps studying flash crash events in a systematic fashion.

Finally, a large literature studies the impact of algorithmic trading on market liquidity, see for example Menkveld (2016) for a review of the vast literature on electronic financial intermediation, and Brogaard et al. (2018) and the quoted literature for the analysis of the behavior of high frequency traders during distress. Anecdotically, algorithmic traders are are typically indicated as the main culprit for the presence of flash crashes. However, our paper is the first, to the best of our knowledge, to study the association of algorithmic trading to flash crashes directly and on a long time span. Bellia et al. (2022) are more granular in studying designated liquidity provision during drift bursts as defined by Christensen et al. (2022), but on a much shorter time span.

On a policy angle, our findings are important for 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). As early as in U.S. Securities and Exchange Commission (2010), the SEC noted that short-term price swings can benefit short-term market traders against long term investors. The V-statistic constitutes a critical step towards the rigorous detection of genuine flash episodes. More broadly, 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. Confirming the attention from market authorities, the V-statistics has been used in the mentioned study of Bouveret et al. (2022) to study flash crashes in the Italian bond and futures market.

3 The V-statistic: theory

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 assume that the log-price process p_t , observed in the interval [0,1], is driven by the equation:

$$dp_{t} = \underbrace{\mu_{t} dt + \sigma_{t} dW_{t} + dJ_{t}}_{\text{traditional semimartingale}} + \underbrace{\frac{c_{1,t}^{-}}{(\tau_{\mathcal{V}} - t)^{\alpha}} \mathbb{1}_{\{t < \tau_{\mathcal{V}}\}} dt + \frac{c_{1,t}^{+}}{(t - \tau_{\mathcal{V}})^{\alpha}} \mathbb{1}_{\{t > \tau_{\mathcal{V}}\}} dt}_{\text{drift bursts}} + \underbrace{\frac{c_{2,t}}{|\tau_{\mathcal{V}} - t|^{\beta}} dW_{t}}_{\text{volatility burst}}, \quad (3.1)$$

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 $\tau_{\mathcal{V}} \in]0,1[$ 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 $\tau_{\mathcal{V}}$. For ease of exposition, we use the same rate of explosion $\alpha \in [0,1[$ before and after, a harmless assumption 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 $\tau_{\mathcal{V}}$ is still consistent with p_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 .

The null (H0) and alternative (H1, that is the V-shape) hypothesis are formulated as follows:

H0:
$$c^{\pm} = c_{1,\tau_2}^- \cdot c_{1,\tau_2}^+ = 0$$
 or $c^{\pm} \neq 0$ and $\alpha - \beta \leq 3/4$.

H1 :
$$c^{\pm} \neq 0$$
 and $\alpha - \beta > 3/4$.

Thus, under the null we can have drift explosion, but *either* at the left or at the right of the explosion point, not at both $(c^{\pm} = 0)$; or we can have explosion at both sides $(c^{\pm} \neq 0)$, provided that the drift explosion is not "too strong" with respect to the explosion in volatility

 $(\alpha - \beta \le 3/4)$. Under the alternative, which we label a V-shape, we do have drift explosion both at the left and at the right of the explosion time, and the explosion is strong enough with respect to the explosion in volatility. The model is flexible in allowing for explosion in volatility or jumps at the same time of the explosion in the drift, both under the null and the alternative.

To test **H1** versus **H0** at a given time point $\tau \in]0,1[$, we propose to use the *V-statistic*, defined as:

$$\mathcal{V}_{\tau,n} = \sqrt{h_n} \cdot T_{\tau,n}^+ \cdot T_{\tau,n}^-, \tag{3.2}$$

where we use a sample of n+1 log-price observations p_0, \ldots, p_n observed at times t_0, \ldots, t_n satisfying Assumption 3 in Appendix A, and

$$T_{\tau,n}^{-} = \sqrt{\frac{h_n}{K_2^{-}}} \frac{\hat{\mu}_{\tau,n}^{-}}{\hat{\sigma}_{\tau,n}^{-}}, \tag{3.3}$$

where

$$\hat{\mu}_{\tau,n}^{-} = \frac{1}{h_n} \sum_{i=1}^{n} K^{-} \left(\frac{t_{i-1} - \tau}{h_n} \right) (p_i - p_{i-1}), \quad \text{for } \tau \in (0, 1],$$
(3.4)

is a localized estimator of the drift, and

$$\hat{\sigma}_{\tau,n}^{-} = \left(\frac{1}{h_n} \sum_{i=1}^{n} K^{-} \left(\frac{t_{i-1} - \tau}{h_n}\right) (p_i - p_{i-1})^2\right)^{\frac{1}{2}}, \quad \text{for } \tau \in (0, 1],$$
(3.5)

is a localized estimator of the volatility. The $T_{\tau,n}^+$ statistics is defined accordingly as

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

with $\hat{\mu}_{\tau,n}^+$ and $\hat{\sigma}_{\tau,n}^+$ computed as in Eqs. (3.4) and (3.5) with $K^-(\cdot)$ replaced by the rightsided kernel $K^+(\cdot)$, that is $K^+(x) = 0$ when $x \leq 0$. In equation (3.2), 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 $K_2^- = \int_{\mathbb{R}} (K^-(x))^2 dx$ 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. It can also be interpreted as the inverse of the left-sided local coefficient of variation (CV), appropriately scaled and where the mean and the standard deviation are estimated nonparametrically. The bandwidth h_n has thus the important interpretation of being proportional to the duration of the V-shape, that is of the transient market inefficiency. 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 τ , that is a trending price. When negative (the case we are interested in), it identifies a swing, that is a time-point τ 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-tests 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 in Appendix A, 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.
$$c^{\pm} = c_{1,\tau_{1}}^{-} \cdot c_{1,\tau_{2}}^{+} \neq 0$$
,

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

Proof. See Appendix A.
$$\Box$$

Theorem 1 shows that the V-statistic can formally capture two-sided drift explosions (when $c^{\pm} \neq 0$), diverging to infinity, when three sufficient 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 (for example, if the explosion is only one-sided), 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 price trend.

Importantly, the first condition in Theorem 1 ($\alpha-\beta>3/4$) implies that the absence of arbitrage principle is violated (see the discussion in COR, where it is shown that absence of arbitrage is preserved by $\alpha-\beta<1/2$). This implies that, under the alternative, the market is "broken" in this specific sense. Violation of the absence of arbitrage principle can be taken as our technical definition of market fragility.

To tackle the multiple testing issue, we compute the minimum value of the V-statistics across several time points (τ_1, \ldots, τ_m) as:

$$\mathcal{M}\operatorname{in}\mathcal{V}_{\tau_1,\dots,\tau_m,n} = \min_{i=1,\dots,m} \mathcal{V}_{\tau_i,n}^{\pm}.$$
(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:

$$\tilde{p}_t = p_t + \varepsilon_t, \tag{3.8}$$

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 to for details.

4 The V-statistic: implementation

How can we separate the two cases of diverging and annihilating $\mathcal{V}_{\tau,n}$ in small samples? And how can we obtain confidence bands when computing the minimum of multiple tests, as in Eq. (3.7)? In case of neither drift nor volatility explosions $(c_{1,\tau}^- = c_{1,\tau}^+ = c_{2,\tau} = 0)$, 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).² 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 unreliable in practice for realistic data generating processes, and useless against gradual jumps (when $c^{\pm} = 0$ but either $c_{1,\tau}^+$ or $c_{1,\tau}^-$ is nonzero) and multiple testing biases. Moreover, using the asymptotic values cannot disentangle one-sided explosions from V-shapes.

For these reasons, instead of relying on the asymptotic theory, we advocate the adoption of a simulated bootstrap to evaluate confidence bands for the $\mathcal{V}_{\tau,n}$ and $\mathcal{M}\text{in}\mathcal{V}_{\tau_1,\dots,\tau_m,n}$. The idea of the bootstrap is to compute the distribution of $\mathcal{V}_{\tau,n}$ (or $\mathcal{M}\text{in}\mathcal{V}_{\tau_1,\dots,\tau_m,n}$) on simulations based on the data at hand. Simulations should have zero drift, but still reproduce the volatility observed in the data. In this paper, we simulate using an EGARCH(1,1) model with t-student residuals to filter the observed variance. The EGARCH(1,1) has two appealing features: it is simple to be estimated, and it can reproduce the most important features of the price dynamics, namely stochastic volatility (through the classical autoregressive dynamics for the logarithmic variance), the leverage effect (since it allows for correlated price and volatility shocks) and jumps (through the fat-tailed distributions of the residuals). Thus, it represent a good compromise to obtain reliable confidence bands in manageable computing time. Appendix C describes the

²If $c_{1,\tau_{\mathcal{V}}}^- = c_{1,\tau_{\mathcal{V}}}^+ = 0$, but $c_{2,\tau_{\mathcal{V}}} > 0$, that is if there is volatility explosion without drift explosion, the product $T_{\tau,n}^- \cdot T_{\tau,n}^+$ is still a modified Bessel function but with a larger variance which grows with the volatility explosion rate β , see Theorem 2 in COR.

implementation of the bootstrap in detail. One limitation of our approach is that we do not prove formal validity of the bootstrap, which is a challenging theoretical question better left for future research. Our bootstrap methodology is instead tested on simulated data in Section 5.

5 **Simulations**

We test the V-statistic and its implementation, as described in Section 4, on synthetic data obtained by simulating four models: the first one is a simple geometric Brownian motion, and the last three are those illustrated in Figure 1. 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 (Brownian motion): $\begin{cases} \overline{p}_t = \log\left(P_1\right) \\ \sigma_t = \sigma_0 \end{cases}$
- Model 1 (price jump) $\begin{cases} \overline{p}_t = \log \left(P_1 \cdot I_{\{t \le 0.5\}} + P_2 \cdot I_{\{t > 0.5\}} \right) \\ \sigma_t = \sigma_0 + (\sigma_1 - \sigma_0) e^{-(2t-1)/\tau_1} \cdot I_{\{t > 0.5\}} \end{cases}$
- Model 2 (gradual jump with volatility spike):

$$\begin{cases} \overline{p}_t = \log\left(P_1 - \frac{P_1 - P_2}{1 + e^{-(2t-1)/\tau_2}}\right) \\ \sigma_t = \sigma_0 + (\sigma_1 - \sigma_0)\left(e^{-(2t-1)/\tau_1} \cdot I_{\{t > 0.5\}} + e^{(2t-1)/\tau_2} \cdot I_{\{t \le 0.5\}}\right) \end{cases}$$

$$\begin{cases} \overline{p}_t = & \log(P_1) + \frac{1-\alpha}{\alpha} \int_0^t \left((\log(P_1) - \log(P_3)) \left(1 - (-(2s-1))^{-\alpha} \right) I_{\{s \le 0.5\}} \right. \\ & + \left(\log(P_2) - \log(P_3) \right) \left((2s-1)^{-\alpha} - 1 \right) I_{\{s > 0.5\}} \right) ds \\ \sigma_t = & \sigma_0 + (\sigma_2 - \sigma_0) \left(e^{-(2t-1)/\tau_1} \cdot I_{\{t > 0.5\}} + e^{(2t-1)/\tau_2} \cdot I_{\{t \le 0.5\}} \right) \end{cases}$$

Parameters are set as follows: $\sigma_0 = 1\%$ (the baseline daily volatility), $\sigma_1 = 2\%$ (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 volatility), $\tau_2 = 0.05$ (the decay time of the price, and the exponential rise time of volatility in Model 2 and Model 3), $\alpha = 0.8$ (the explosion rate for drift in Model 3). For Model 3, use either a heightened volatility at the peak ($\sigma_2 = 2.5\%$) or the same volatility of Model 2 $(\sigma_2 = 2\%, \text{ in this case we call it Model 3'})$. 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. We simulate two settings. The first one has $P_1 = P_3 = 100$, and $P_2 = 98$, that is no fundamental news and 2%overshooting in Models 3 and 3'. In this setting without a fundamental price jump, Model 1 and Model 2 differ only in the volatility shape before t = 0.5. The second one has $P_1 = 100$,

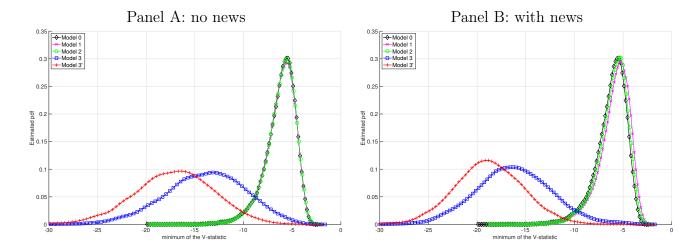


FIGURE 2: We simulate Model 0 (Brownian motion), 1 (price jump), 2 (gradual jump with volatility spike), 3 and 3' (V-shape). Model 3' has the drift of model 3 and the (smaller) volatility of model 2. Panel A: $P_1 = P_3 = 100$, $P_2 = 98$ (no change in the efficient price, 2% overshooting). Panel B: $P_1 = 100$, $P_3 = 98$, $P_2 = 96$ (2% change in the efficient price, 2% overshooting). The panels show the distributions of the minimum of the V-statistics (3.2) on 10,000 replications of the four models.

 $P_3 = 98$ (2% change in the fundamental price) and $P_2 = 96$ (2% overshooting in Models 3 and 3').

We generate 10,000 simulations of n = 23,400 observations (corresponding to 1-second observations in a day with 6.5 hours of trading). We then compute the minimum of the V-statistic and its confidence bands on each simulation as described in Appendix C. We use exponential kernels and a bandwidth $h_n = 300$ seconds for the numerator, and $h'_n = 1500$ seconds for the denominator, in line with the empirical application below.

Figure 2 shows the distribution of $\mathcal{M}in\mathcal{V}_{\tau_1,\dots,\tau_m,n}$ on the simulations of the five models. Panel A reports the distribution in the setting with no news; and Panel B in the setting with news. Panel A shows that, when there are no news, the distribution of the minimum of the V-statistic in the simple Brownian motion (Model 0), in Model 1 and in Model 2 (a volatility spike) is almost identical, while in Model 3 (the V-shape) the distribution is clearly shifted to the left. The distribution is even more shifted to the left in Model 3', that is under the alternative with smaller volatility. The same results are confirmed in Panel B, that is when there is a jump in the fundamental price. In this case, the distance between the null and the alternative is even more pronounced, and the distinction between Model 1 (a price jump) and Model 2 (a gradual price jump) from the simple Brownian motion is virtually immaterial. Figure 2 thus shows two crucial properties of the V-statistic. The first is that it is able to clearly disentangle the null (Model 0, 1 and 2) from the alternative (Model 3 and 3'). The second is that it is extremely robust to a vast set of null hypothesis, and in particular to jumps or volatility spikes.

TABLE 1: Percentage of rejections at 1% and 5% of the $\mathcal{M}in\mathcal{V}_{\tau_1,...,\tau_m,n}$ statistic (3.7), based on 1,000 bootstrapped simulations with an EGARCH(1,1) model. Panel A: $P_1 = P_3 = 100$, $P_2 = 98$ (no change in the efficient price, 2% overshooting). Panel B: $P_1 = 100$, $P_3 = 98$, $P_3 = 96$ (2% change in the efficient price, 2% overshooting). Model 3' has the same drift of Model 3 but the (lower) volatility of Model 2.

Panel A: no fundamental price change, 2\% overshooting.

	Model 0		Model 2		Model 3'	
Quantile	Percentage of rejections					
0.1%	0.22	0.21	0.20	42.04	72.05	
1%	0.92	0.92	0.91	67.42	89.98	
5%	4.50	4.66	4.58	84.47	96.87	

Panel B: 2% fundamental price change, 2% overshooting.

	Model 0	Model 1	Model 2	Model 3	Model 3'	
Quantile	Percentage of rejections					
0.1%	0.22	0.11	0.15	46.46	68.05	
1%	0.91	0.48	0.83	79.49	93.75	
5%	4.43	2.48	3.98	93.03	98.85	

It really takes a discontinuity in the sign of the drift to drive the V-statistic away from the null hypothesis.

Table 1 reports the number of rejections at the 0.1%, 1% and 5% confidence level, estimated using a simulated bootstrap based on the EGARCH(1,1) model, in the two settings. The table shows that the test is very well sized in both settings, with a slight oversizing in the extreme tail, and slight undersizing for model 1 in the case with news. Power is extremely good in both settings, and it increases when the volatility spike is smaller, or when there is fundamental news, or if the price overshooting is larger (not shown).

Finally, we show that our test has more power in detecting V-shapes than the Variance Ratio (Lo and MacKinlay, 1988). We focus on this specific test because it can be considered the most popular test for the violation of the random walk hypothesis, which is typically associated with the weak form of market efficiency. Under the null hypothesis that prices follow a random walk with uncorrelated increments, the Variance Ratio should equal one at all lags. We implement

the Variance Ratio statistics on a modification of Model 0 (the benchmark), a modification of Model 2 (the gradual jump, capturing underreaction) and Model 3 (the V-shape, capturing overreaction). The modification of Models 0 and 2 consists in using the volatility of Model 3 (σ_2) in place of their respective volatility, σ_0 and σ_1 . In such a way, the three models share the same volatility, σ_2 . For the modified Model 2, we also simulate the same drift of Model 3 up to 0.5, and no drift thereafter. These modifications allow for a fair comparison of the results of the three models.

To compute a Variance Ratio (VR), one needs to specify a frequency at which returns are sampled and a horizon at which returns are aggregated. The VR is then computed as the ratio between the variance of aggregated returns and the variance of the sampled returns, times the ratio between the sampling frequency and the aggregation horizon. We subtract 1 to center the VR at zero. Figure 3 shows the distributions for a 60-second frequency and 300-second aggregation horizon (Panel A) and for a 10-second frequency and 300-second aggregation horizon (Panel B) for the three models. It is clear that the Variance Ratio cannot disentangle Model 2 from Model 3. Indeed, both models generate a significantly positive serial correlation (generated by a single trend in the case of Model 2, and by two opposite trends in the case of Model 3), which materializes in Variance Ratio statistics shifted to the right, without capturing the change in the trend sign of the V-shape in Model 3. In this relevant example, the Variance Ratio statistics are even less powerful than the V-statistic. Indeed, the Variance Ratio test would reject, at 1% confidence level, in 54.06% of the replications of Model 3 in the 10/300 case, and just in 15.57% of the replications in the 60/300 case, thus significantly less than the V-statistic (79.49% as displayed in Table 1).

For robustness, we compute variance ratio statistics from a variety of aggregation horizons for the two sampling frequencies considered in Figure 3, that is 10 and 60 seconds. We report the average and 95% confidence bands of (centered) VRs in Figure 4. This robustness check shows that the finding of Figure 3 does not depend on the sampling frequency nor on the aggregation horizon we use to implement the Variance Ratio.

Summarizing, our simulation experiments show that the V-statistics, and in particular their minimum, are able to disentangle a V-shape from a random fluctuation accompanied by a jump (discontinuous or gradual) and a burst in volatility on synthetic data mimicking realistic situations in financial markets. The ability of the V-statistic to detect V-shapes increases when price volatility is lower and when price overshooting is higher. By contrast with our test, the classical Variance Ratio statistic (based on serial correlation) is instead not able to tell apart underreaction from overreaction, and is less powerful than the V-statistics. Corroborated by these results, we now turn to real financial market data.

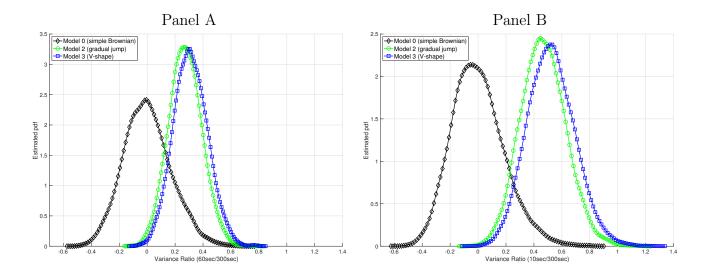


FIGURE 3: We simulate a modification of Model 0 (Brownian motion), a modification of Model 2 (gradual jump with volatility spike), and Model 3 (V-shape) with $P_1 = P_3 = 100$, $P_2 = 98$. Model 0 and Model 2 have now the same volatility of Model 3. Model 2 has the same drift of Model 3 up to the center time point, and then zero drift. Panel A: The Variance Ratio is implemented with sampling frequency 60 seconds and aggregation horizon 300 seconds. Panel B: The Variance Ratio is implemented with sampling frequency 10 seconds and aggregation horizon 300 seconds. The panels show the distributions of the Variance Ratio (minus 1) on 10,000 replications of the three models.

6 Empirical applications

The empirical applications we showcase in this section are ordered according to the duration of the transient inefficiency we capture. We first show the ability of the V-statistic in detecting mini-flash crashes, a phenomenon that has recently drawn the attention of researchers and regulators (see, e.g. Biais and Foucault, 2014, Laly and Petitjean, 2020 and Bouveret et al., 2022). Here, we use a forensic benchmark to test the correctness of our specification, that is we investigate the events involved in the SEC legal actions against financial intermediaries for market access rule violations. We then study mini-flash crashes in the Google stock in 2013 and in the SPDR S&P 500 trust exchange-traded fund (SPY) in the period 2013-2020. Next, we study the Flash Crash of May 6, 2010, and the first wave COVID-19 market crash, the former case providing an example of a strong signal, and the latter providing an example of a price fluctuation that is not recognized by the V-statistic as a distressed state. We conclude with an event study focusing on the auction cycle in the Italian sovereign bond market and on its redistribution effects.

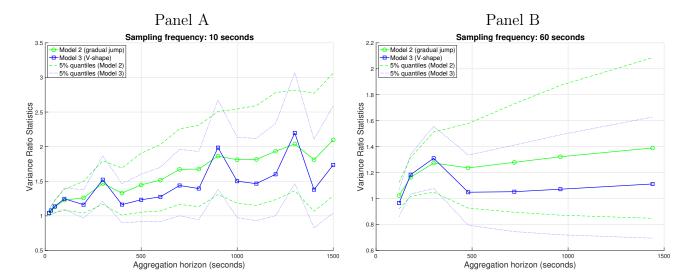


FIGURE 4: We simulate a modification of Model 0 (Brownian motion), a modification of Model 2 (gradual jump with volatility spike), and Model 3 (V-shape) with $P_1 = P_3 = 100$, $P_2 = 98$. Model 0 and Model 2 have now the same volatility of Model 3. Model 2 has the same drift of Model 3 up to the center time point, and then zero drift. Panel A: The Variance Ratio is computed with sampling frequency of 10 seconds and the indicated aggregation horizon. Panel B: The Variance Ratio is computed with sampling frequency 60 seconds and the indicated aggregation horizon. The panels show the averages of Variance Ratio (minus 1) and their 5% quantiles on 10,000 replications of the three models.

6.1 Mini-flash crashes and market violations

While there is clear indication that mini-flash crashes are a menace to financial markets,³ there is no widespread consensus on how to define a mini-flash crash yet. We argue that the SEC legal actions against financial intermediaries for market access rule violations provide an excellent benchmark to assess the correctness of our specification, and we apply the V-statistic to mini-flash crashes as determined by regulatory charges against traders or financial intermediaries (the flash crash of May 6, 2010 in Figure 9 is a prominent example in this category, since the trader that caused the crash was eventually prosecuted).

On September 26, 2016 the SEC announced that "Merrill Lynch has agreed to pay a 12.5 million penalty for maintaining ineffective trading controls that failed to prevent erroneous orders from being sent to the markets and causing mini-flash crashes". The occasions investigated by the SEC are listed in Table 2.⁴

Table 2 also shows that the V-statistic is able to detect mini-flash crashes in 12 out of 14

³See, for example, the speech given by Andrew Hauser on January 7, 2021: "Why central banks need new tools for dealing with market dysfunction", available at https://www.bankofengland.co.uk/speech/2021/january/andrew-hauser-speech-at-thomson-reuters-newsmaker.

⁴We exclude one event, namely EBND on 4-5 December 2012, since the first trade of the day was involved, and the price preceding the event is the prior day's closing price.

Table 2: Mini-flash crashes reported by the SEC as evidence of violations by Merrill Lynch of Section 15(c)(3) of the Exchange Act and Exchange Act Rule 15c3-5 ("Market Access Rule" or "Rule"). The first three columns report the stock ticker, the date of the event and the start of the event through their Height/Depth, as reported by the SEC (U.S. Securities and Exchange Commission, 2016). The fourth column reports the daily minimum of the V-statistic, and the fifth column its corresponding p-value, computed by simulated bootstrap using 5,000 replication of an EGARCH(1,1) model calibrated each day. The last column is the location of the minimum of the V-statistics.

Ticker	Date	Time	V-statistics	p-value	Estimated time
		(from SEC)	(daily minimum)	(%)	
ACRE	18-jun-2013	15:58:53-15:58:54	-7.53	0.10	15:58:54.083
APC	17-may- 2013	15:59:59-15:59:59	-5.73	0.72	15:59:59.489
AXE	8-mar- 2013	11:48:51-11:48:52	-10.44	0.02	12:54:27.323
BCEI	14-may- 2013	15:40:05-15:40:06	-13.62	0.02	15:40:06.233
DEO	7-nov-2012	15:45:47-15:45:49	-29.67	0.00	15:45:49.282
GOOG	22-apr-2013	09:37:11-09:37:12	-22.76	0.00	09:37:12.132
GVA	2-jan- 2013	11:45:56-11:45:57	-11.02	0.02	11:45:57.826
PCRX	1-feb- 2013	15:38:17-15:38:18	-12.61	0.00	15:38:18.182
PLD	11-jun-2013	09:11:48-09:11:48	-4.55	4.68	10:46:47.105
PPG	18-jan- 2013	15:00:50-15:00:51	-15.83	0.00	15:00:51.868
QLYS	25-apr- 2013	09:35:38-09:35:39	-13.09	0.04	10:15:23.172
ROI	27-nov-2012	09:30:00-09:30:29	-5.85	8.74	10:49:30.909
TYC	19-sep- 2012	10:18:38-10:18:39	-24.21	0.00	10:18:39.650
ZBRA	17-jun- 2014	15:59:03-15:59:51	-14.66	0.00	15:59:14.029

instances. The two events that are not detected are characterized by extremely thin trading, with the mini-flash crashes consisting in one or two trades only, see the display of the events in Figures 20 and 21 in Appendix E. Also, the detected timing corresponds exactly with that indicated by the SEC for all events. Moreover, three stocks (AXE, QLYS and ROI) also display a second flash crash later on during the same day. The V-statistic detects both events (see again Figures 20 and 21 in Appendix E) and is more negative for the second crash in all the three cases. Thus, our econometric definition of mini-flash crashes virtually coincides with the forensic definition adopted by the SEC to charge Merrill Lynch.

Were there additional mini-flash crashes for these stocks? The SEC itself seems to suggest that this is the case, since the events in Table 2 are defined as "Examples of market incidents that resulted from erroneous orders that Merrill Lynch sent to the market during the relevant period" (U.S. Securities and Exchange Commission, 2016).

To detect mini-flash crashes, we adopt an algorithm which complements the V-statistic with a companion measure which has an immediate economic interpretation. Denote by $\tau_{\mathcal{V}}$ the time when a significant V-shape is detected using the bandwidth h_n , and for this point define the

(standardized) price convexity $\mathcal{C}_{\tau_{\mathcal{V}},h_n}$ as

$$C_{\tau_{\mathcal{V}},h_n} = \frac{p_{\tau_{\mathcal{V}}-h_n} + p_{\tau_{\mathcal{V}}+h_n} - 2p_{\tau_{\mathcal{V}}}}{\sqrt{2h_n\sigma_{\tau_{\mathcal{V}}-h_n}^2}},$$
(6.1)

where p_t and σ_t^2 are the logarithmic price and the spot variance at time t respectively. Price convexity is just the difference of logarithmic returns before and after the V-shape peak, standardized with the variance at the beginning of the time window. This dimensionless quantity has the same structure of a t-statistic and is clearly related to the V-statistic (3.2), but with a major difference: in equation (6.1), we standardize with volatility before the peak, not at the peak. It makes economic sense to define mini-flash crashes as V-shapes with high price convexity. We estimate price convexity using the filtered volatility from the same EGARCH model we use for the bootstrap, and detect a mini-flash crash if the two following conditions are met:

- 1. The minimum of the V-statistics with bandwidth h_n is significant at the confidence level α ;
- 2. The absolute value of price convexity at the time of the minimum $\tau_{\mathcal{V}}$ is greater than a given threshold $\theta_{\mathcal{C}}$.

We apply the V-statistic to data of the Google stock (ticker: GOOG) in 2013 with a bandwidth of $h_n = 10$ seconds and, again, we use simulated bootstrap of an EGARCH(1,1) model to determine confidence bands. With $\theta_C = 10$ and $\alpha = 1\%$, we reveal 3 mini-flash crashes, displayed in Figure 5. Out of these three, one (displayed in the bottom panel of Figure 5, on 22 April 2013) has been used by the SEC for its legal actions (see Table 2).

We do not need to restrict our analysis to flash crashes that span a few seconds only. We apply our methodology also using $h_n = 30,60,300$ seconds, and the number of events detected by our algorithm with $\theta_C = 5$ on the Google stock in 2013 grows to 59 and includes notorious examples like the "Twitter flash crash" of April 23, 2013, displayed in Figure 18. Figures 22-24 in Appendix E shows the 19 detected events at $\alpha = 1\%$ confidence level with a price convexity larger than $\theta_C = 10$ in absolute value. All of these are clear examples of the pervasive presence of flash crashes in stock data, which reinforces the findings of COR that flash crashes are not isolated events, but rather a stylized fact of the price dynamics.

Next, we apply the detection technology to SPY from 2013 to 2020. We again use $h_n = 10, 30, 60, 300$ seconds and $\alpha = 1\%$. When using the threshold $\theta_C = 5$, we detect 240 events in

⁵Using price convexity as defined in Eq. (6.1) to detect a mini-flash crash only works if we do so after conditioning to the significance of the V-statistic. The sole statistic (6.1) would not indeed be robust to jumps and volatility spikes, see e.g. Lee and Mykland (2008) who use a test statistic similar to (6.1) purposely to detect jumps.

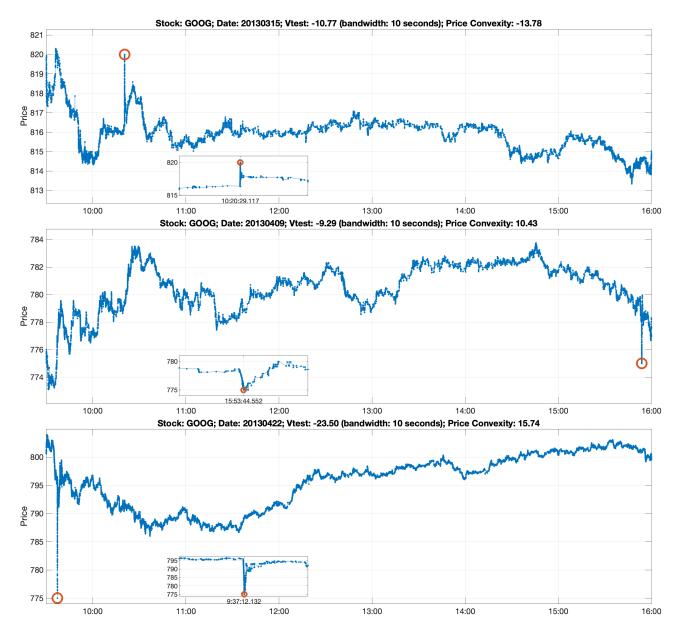


FIGURE 5: The three detected mini-flash crashes on the Google stock in 2013 using a bandwidth $h_n = 10$ seconds. In the inset, a zoom around the few seconds around the 40 seconds around the crash. The event displayed in the bottom panel has been used by the SEC in legal charges against Merrill Lynch for causing mini-flash crashes (U.S. Securities and Exchange Commission, 2016).

8 years, 174 with positive convexity (crashes) and 66 with negative convexity (surges). With $\theta_C = 10$, we find 63 instances, (45 crashes and 18 surges). The temporal distribution of these instances is shown in Panel A of Figure 6. Panel B shows the average cumulative return across these events, which displays exactly the V-shape the test is designed to capture. There is a clear asymmetry between crashes and surges, with crashes being relatively deeper. Panel A shows that the number of flash events rose over the past few years, especially for events with

 $^{^6}$ All the instances with absolute value of price convexity larger than 10 are shown at http://dse.univr.it/hidea/researchoutput.html.

larger price convexity, which is a rather worrying phenomenon from the perspective of financial stability. The number of detected events in 2013 is somewhat smaller than that observed for Google (31 vs 59 respectively), indicating that mini-flash crashes are more likely for individual stocks than for the index. Our results complement the analysis of Gao and Mizrach (2016), who show an increase of market breakdowns in equities from 1993 to 2008 and a decline in 2009 and 2010, which they associate with a regulatory change, namely the adoption of the Regulation National Market System.

In mini-flash crashes detected by the SEC and used as evidence in legal proceedings, the events were always associated with abnormal volume. This is the case also for mini-flash crashes in SPY. Panel C of Figure 6 shows the median traded volume (in millions of U.S. dollars) per second in the thirty minutes preceding and following the event. We can see a huge spike exactly at the time when the mini-flash crashes occur, with the volume rising by more than 25 times with respect to the average level of trading.⁸ This corroborates the view expressed by the SEC that mini-flash crashes are generated by trading abnormal sizes. Figure 7 shows how the volume at the second of the peak of a detected flash crash compares with the unconditional distribution of traded volumes on SPY. Clearly, flash events are in the very right tail. This supports the policy of the SEC of preventing the entry of orders that would exceed appropriate credit or capital thresholds. Indeed, the reason why Merrill Lynch was charged was "failing to establish to establish pre-trade risk management controls reasonably designed to prevent the entry of erroneous orders" (U.S. Securities and Exchange Commission, 2016).

Panel D of Figure 6 shows the average volatility around the events, measured with the same EGARCH(1,1) model used to filter volatility for the bootstrap and for the estimation of price convexity in Eq. (6.1). The observed behavior of SPY volatility around flash crashes is similar to that depicted in model 3 in Figure 1. Volatility rises few minutes before the event, and then declines much more slowly. The effect is stronger with stronger price convexity, and the impact on price volatility is much more persistent than the impact of the price itself, with volatility remaining higher than the previous level for longer time.⁹

Finally, we investigate algorithmic trading (AT) activity during flash crashes. We follow Weller (2018) and compute two measures of algorithmic trading which can be inferred by trades,

⁷The t-statistics of the regression lines in Panel A of Figure 6 are 1.52 for events with price convexity larger than 5 and 4.03 for events with price convexity larger than 10.

⁸The difference in trading intensity is strongly statistically significant. The average trading size in SPY in the period 2013-2020 was \$0.90 million dollar per second (we use the same units below). In the second in which we detect the peak of the crash, average trading size was \$7.96 for events with price convexity larger than 10, with a t-test for comparison between means of 5.55, and \$25.44 for events with price convexity larger than 5, with a t-test of 1.44. If we include the previous and the next second in the calculation, the average trading size for events with price convexity greater than 5 becomes \$13.02, and the t-test becomes 2.13.

⁹This effect is, again, statistically significant. Comparing the average volatility 15 minutes after and 15 minutes before the peak of the event, the two-sample mean equality t-test is 3.51 for V-shapes with price convexity larger than 5, and 4.15 for V-shapes with price convexity larger than 10.

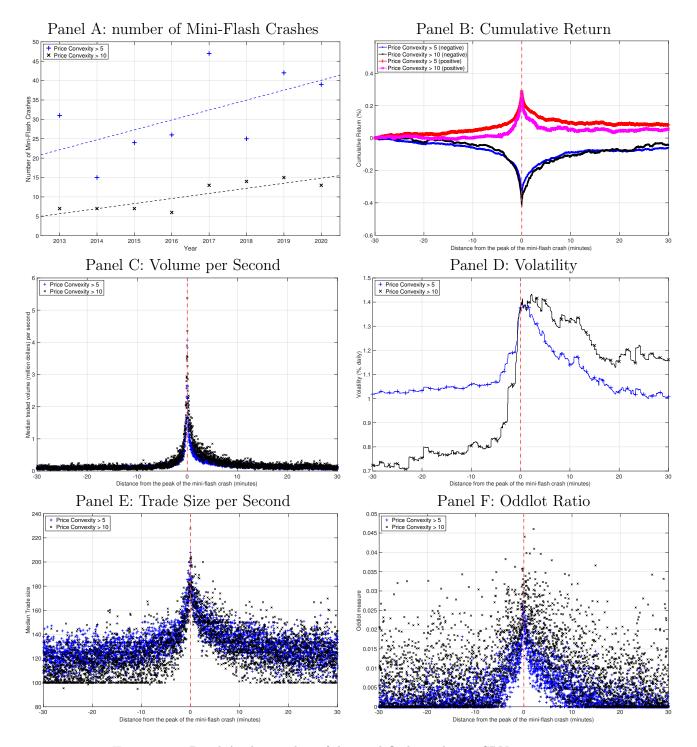


FIGURE 6: Panel A: the number of detected flash crashes in SPY per year, from 2013 to 2020. Panel B: the average cumulative return around detected mini-flash crashes in the 30 minutes around the peak. Panel C: the median traded volume (in U.S. dollars) per second. Panel D: the average volatility (daily units, percent, estimated from the EGARCH(1,1) model). Panel D: the median trade size (in U.S. dollars) per second. Panel E: the average percentage of traded volume in trades with less than 100 contracts (oddlot ratio). In all figures we show averages with $\theta_C = 5$ and $\theta_C = 10$, where θ_C is the price convexity threshold used for detection. Namely, we select V-shapes with bandwidths $h_n = 10, 30, 60, 300$, significant at the 1% level, with

$$|\mathcal{C}_{\tau_{\mathcal{V}},h_n}| = \left| \frac{p_{\tau_{\mathcal{V}}-h_n} + p_{\tau_{\mathcal{V}}+h_n} - 2p_{\tau_{\mathcal{V}}}}{\sqrt{2h_n\sigma_{\tau_{\mathcal{V}}-h_n}^2}} \right| > \theta_C.$$

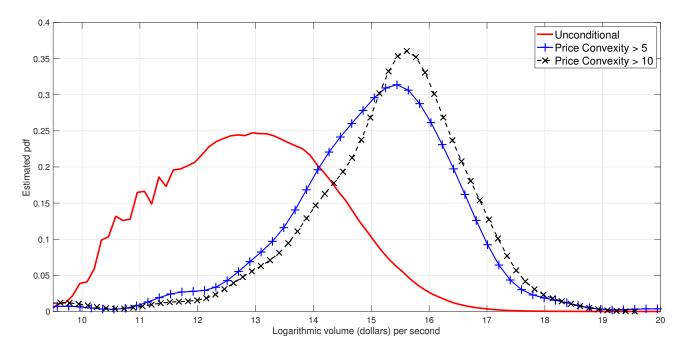


FIGURE 7: Estimated probability distributions of logarithmic traded dollar volume (in U.S. dollars) per second between 9.30 and 16:00 on SPY in the period 2013-2018: unconditional (red-solid line), and conditional to the detection of a V-shape in that second. For conditional distributions, we select V-shapes with bandwidths $h_n=10,30,60,300$, significant at the 1% level, with $|\mathcal{C}_{\tau_{\mathcal{V}},h_n}|=\left|\frac{p_{\tau_{\mathcal{V}}-h_n}+p_{\tau_{\mathcal{V}}+h_n}-2p_{\tau_{\mathcal{V}}}}{\sqrt{2h_n}\sigma_{\tau_{\mathcal{V}}-h_n}^2}\right|>\theta_C$. The two distributions are for $\theta_C=5$ and $\theta_C=10$.

namely average trade size (the number of shares traded divided by the number of trades) and oddlot ratio (the total volume executed in quantities smaller than 100 shares divided by total volume traded). The behavior of these two measures is displayed in Panels E and F respectively. In presence of increasing algorithmic trading, the trade size is expected to decrease, and this is not the case, on average, during our detected events. This might be explained by the fact that, during flash events, non-algorithmic traders could increase their trade size. For the oddlot ratio instead the evidence is unambiguous: the measure peaks during mini-flash crashes, indicating an increasing presence of algorithmic trading. Figure 8 shows the time series of yearly averages of the two algorithmic trading activity measures. We see that, over the considered years, the oddlot ratio increased and the trade size decreased. In this case, both measures indicate that algorithmic trading increased on SPY. While the association with the increase of mini-flash crashes (Panel A of Figure 6) does not imply causality, it represents an interesting empirical fact revealed by our detection methodology in line with the widespread consensus on the fact that mini-flash crashes and algorithmic trading are two related phenomena (see, e.g., the discussion in U.S. Securities and Exchange Commission, 2020).

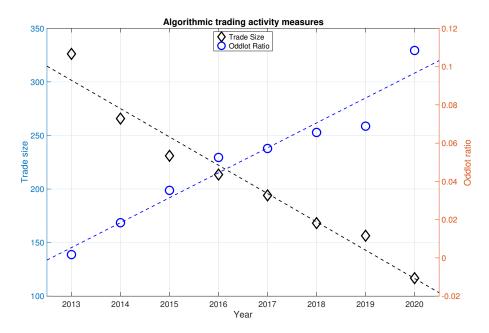


FIGURE 8: Yearly averages of algorithmic trading measures (oddlot ratio and trade size) from 2013 to 2020. The dashed line are the output of linear regression.

6.2 The flash crash of May 6, 2010

Here we apply the V-statistic to E-mini S&P500 futures prices recorded on May 6, 2010, that is the day of 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 its displayed confidence bands. When we account for multiple testing (as explained in Appendix C), we estimate that the probability that such a value is reached because of a statistical fluctuation due to volatility is only 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.

6.3 V-shapes during the COVID-19 pandemic?

When the COVID-19 pandemic started in February 2019, the U.S. stock index dropped dramatically and then recovered through 2020, as shown in Figure 10. Can the observed price dynamics be considered a V-shape?

To answer this question, we apply the V-statistic to all transacted prices of SPY from January 1st, 2019 to August 17, 2020 (see Figure 10). The price nadir was reached on March 23, 2020, at 11:46:11 am (the recorded SPY price was 218.29). We compute the V-statistic precisely at that time. We use different bandwidths h_n ranging from 5 minutes to more than 40 hours. The bandwidths can be roughly interpreted, as discussed above, as the duration of the state of

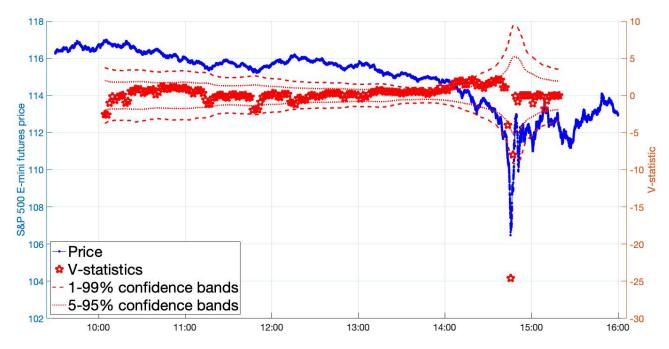


FIGURE 9: Blue dots: transaction prices for E-mini S&P500 futures on May 6, 2010. Red pentagrams: V-statistics computed using an exponential kernel and a bandwidth $h_n = 1000$ seconds. Confidence bands for the individual tests are computed bootstrapping the V-statistics on 50,000 simulations of an E-GARCH(1,1) 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%.

transient market inefficiency. For consistency, we apply the V-statistic to the minimum reached by simulations in the bootstrap procedure. The result is shown in the inset of Figure 10. The V-statistics are negative, as expected, but not significant. Thus, we cannot conclude that the U.S. stock market was driven by drift during the turmoil due to the pandemic. This result is consistent with the market being hit by a huge wave of uncertainty, which translated in enormous volatility in 2020.

This section thus shows an example in which an "apparent V-shape" in the data is not actually a V-shape (that is a discontinuity in the sign of the price drift) when judged with the formal statistical procedure we introduce in this paper. On the other hand, the COVID-19 pandemic did have an adverse effect on the fragility of the stock market. On SPY, for example, in 2020 we detect 12 V-shapes with price convexity greater then 10, and the one with the largest convexity (28.83) was observed on March 12, 2020, during the first wave. However, the flash crashes we detect are all at high-frequency in nature and not qualitatively different from those in previous years, and our technology does not allow to conclude that there was market inefficiency over a longer time scale.

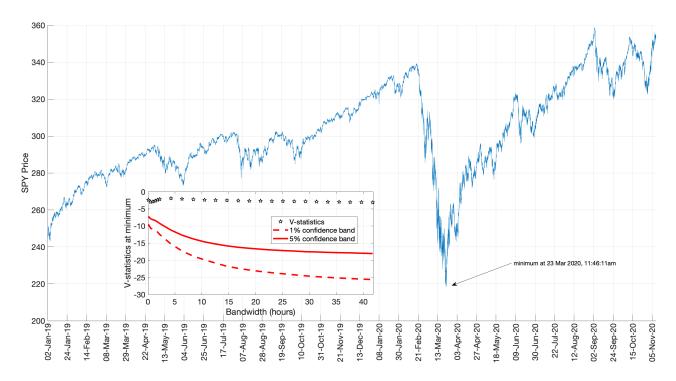


FIGURE 10: Traded prices of SPY interpolated to the nearest second for the sample period considered. In the inset, we show the V-statistics, together with 1% and 5% confidence bands computed via simulated bootstrap of an EGARCH model calibrated on the same data, as computed at the instant in which SPY 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 .

6.4 Event study: Italian bonds in May, 2018

While the previous section makes it clear that the phenomenon of flash crashes in financial markets is endemic and increasing over the last few years, it still remains to discuss whether such dysfunctions threaten financial stability on a time scale that lasts more than a few minutes. The event study in this section serves to this purpose in illustrating an impactful V-shape in the Italian sovereign bond market in May, 2018. The choice of this specific event is not accidental. First, the Italian bond market, one of the largest in the world (the nominal size of the Italian public debt being roughly 10% of the U.S. one), is central to the global economy. Second, our analysis focuses on the auction cycle, that is on the price decline and subsequent rise typically observed when an anticipated supply shock (the auction) hits the bond market (Lou et al., 2013; Sigaux, 2020). The auction cycle is an example of price reversal originated by economic forces which are different from those originating the flash crashes analyzed in the previous section.

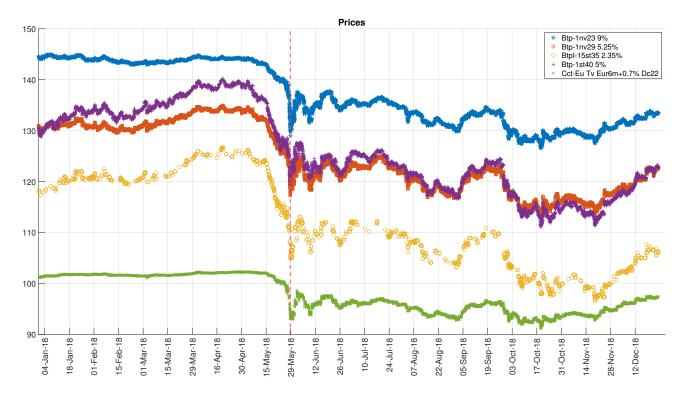


FIGURE 11: Tick-by-tick transaction prices (clean price) observed in the MOT (secondary market) for five representative Italian bonds. The dashed-red line is May 29, 2018.

6.4.1 Prices, volatility, volume and V-statistic

Figure 11 shows the tick-by-tick transaction prices of selected Italian government bonds in 2018.¹⁰ The figure highlights the extreme price swing that occurred on May 29. The observed pattern is consistent with that shown in the top-right panel of Figure 1.

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 12 shows the square root of daily RV for the instruments considered in Figure 11 (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 1, and thus it does not allow to disentangle a V-shape from the normal dynamics of the bond prices.

Figure 13 reports the daily minimum of the V-statistics on selected bonds, which is largely

¹⁰These bonds represent the spectrum of all available Italian bonds: BTPs are fixed rate bonds, CCTs are floating rate bonds and BTPis are inflation linkers. Appendix D provides details about the data used for the analysis in this section, as well as the macroeconomic background during the crash.

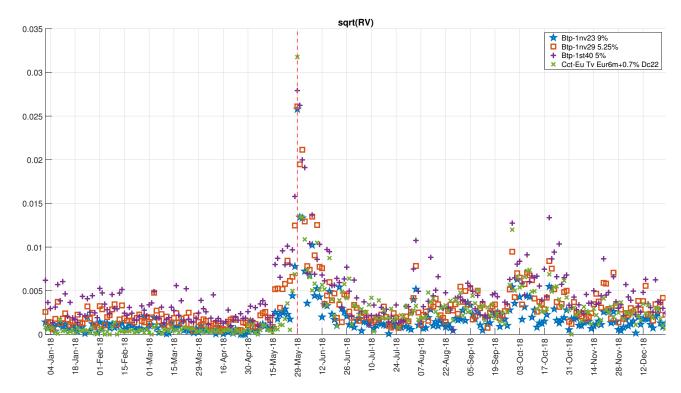


FIGURE 12: Square root of daily five-minutes realized volatility (RV) for four representative Italian bonds sampled in the MOT. The dashed-red line is May 29, 2018.

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-statistic is close to the 1% confidence limit. The p-values for the minimum of the V-statistics, estimated by simulated bootstrap (see Section 4) are 1.5%, 2.4%, 1.8% for the fixed-rate bonds and 9.8% for the floating rate bond. It is also interesting to note the positive peak of the V-statistics in mid-May, which, according to the theory, is a signal of strong trend.

Commenting on the crash, Financial Times¹¹ 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 drift. The distinction between a volatility move and a drift move, as argued across this paper, is not immaterial. Indeed, large volatility is possible even in an efficient and perfectly liquid market. A V-shape has instead to be associated with distress, in this case with prices pushed away from fundamentals because of the auction cycle.

Market liquidity is affected too. The daily traded volumes of the five assets, showcased in Figure 14, show an increase around the week of May 29, 2018 which is similar to that observed around flash crashes (Figure 6). The number of transactions, shown in the inset of the same figure, also peaks, and even more than volume itself. This dynamics is similar to that of mini-flash

¹¹ Italian bonds' extreme volatility exposes liquidity strains", Financial Times, June 1, 2018

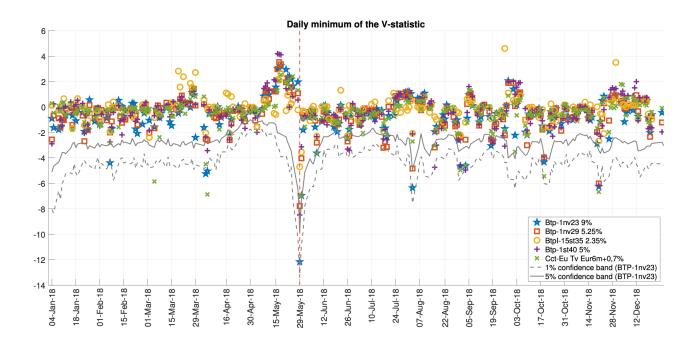


FIGURE 13: Daily minimum of V-statistic computed as in Eq. (3.7) at all intraday transaction times. proposed in this paper, using a bandwidth $h_n = 2$ days. The reported confidence bands are relative to the Btp-1nv23. The vertical dashed-red line is May 29, 2018.

crashes but affects the market on a much longer time scale.

Finally, we compute an explicit economic metric associated with this V-shape, namely the redistribution effect in terms of wealth suffered by the Italian taxpayers in favor of primary dealers. The Italian Treasury held auctions on May 28, 29 and 30. Details on the auction mechanism are provided in Appendix D. This section quantifies the wealth transfer in the auctions of May 30, the impact on the other two auctions having been minor.

For each auctioned Treasury bond, we define the loss Italian taxpayers had to endure as the difference 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 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 allow for a third regime $(s_t = 3)$, which we expect to occur after the distressed state, where log-returns have again zero mean, but a possibly different (and

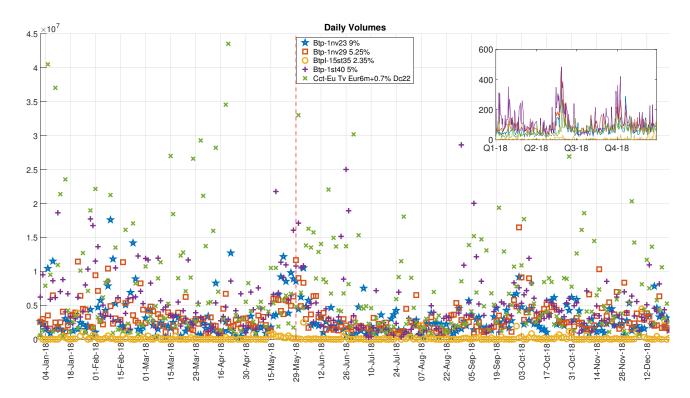


FIGURE 14: Daily volumes observed in the MOT (secondary market) for five representative Italian bonds. In the inset the number of transactions per trading day for the same five bonds. The dashed-red line is May 29, 2018.

supposedly higher) volatility, σ_3 . Allowing for a different volatility level in the third state is meant to capture the typical behavior of volatility when news arrive 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 16 shows the historical behavior of the Btp-1Feb28 2% log-returns (upper panel), and the 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 bonds, 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

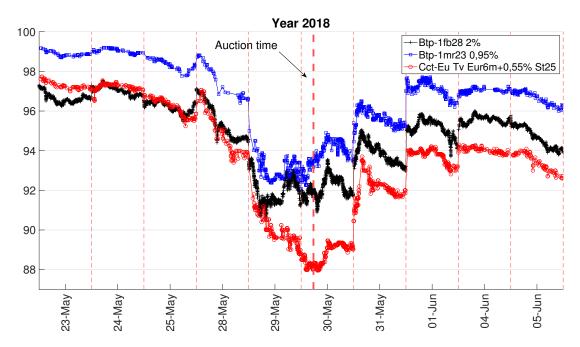


FIGURE 15: Secondary market tick-by-tick transaction prices of the three securities auctioned on May 30, 2018, during the debt crash week. The dashed-red line is the time at which the auction was concluded.

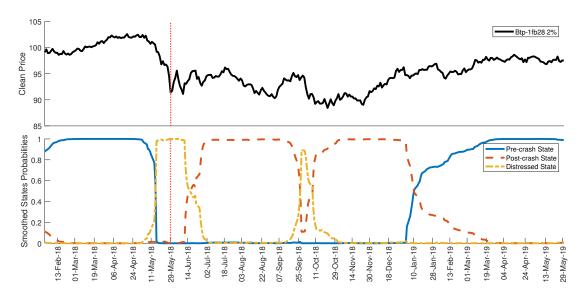


FIGURE 16: Top Panel: Daily prices of Btp-1Feb28 2%. Bottom panel: smoothed state probabilities. The dotted red line indicates 29 May 2018.

for 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 6.4.1.

To estimate the distribution of \tilde{P}_t , we assume that a functioning market would have switched from state $s_t = 1$ to state $s_t = 3$, without going through the distressed state $s_t = 2$. That is, we assume that the average price moves to a new fundamental level, that the volatility spikes,

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.

	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 · 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% $(2.53 \cdot 10^{-8})$

but the inefficient market (that of the right column of Figure 1), 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 a market dominated by drift.

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, show-cased in Figure 17. 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 primary dealers. To gauge the enormous size of this wealth transfer, it is enough to mention that this single auction costed 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.

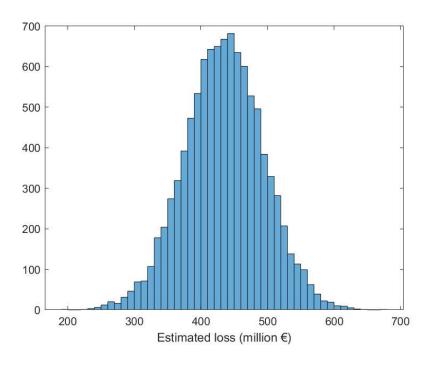


FIGURE 17: Distribution of the estimated loss for the Italian Treasury due to the auction cycle of May 30, 2018.

7 Conclusions

When the market is transitorily moving away from fundamentals, we observe a V-shape in prices, defined as a discontinuity in the sign of the price drift. The V-statistic that we introduce in this paper can reliably detect such V-shapes. We present several applications of this new tool. We provide new empirical evidence relative to the presence of frequent (and increasing over time) mini-flash crashes in U.S. stocks. As highlighted by SEC legal proceedings against financial intermediaries for violation of the Market Access Rule, V-shaped events represent not only a potential violation of market rules, but also a threat to investor confidence in the market. Our results prove that our tool is of help for the regulators' mission to maintain orderly and efficient markets. In a specific event study (the crash of Italian bonds on May, 2018), we show the harmful implications of V-shapes for financial stability. Specifically, we quantify the extent of the wealth transfer that the V-shape-related inefficiency (the auction cycle magnified by political turmoil) caused.

Understanding what gives rise to flash crashes is of clear relevance, especially for regulators of markets on which trading is performed almost exclusively by algorithms. We position our paper as an intermediate step to achieve this goal. We indeed advocate that, to study these events, one first has to identify that they indeed occurred. Our paper provides a solid methodology to do so, which allows to answer important questions to be addressed in future research. Why

do mini flash crashes occur? Why did their number increase over time? Can we identify price changes that are likely to be reversed using information available before the reversal occurs? Our paper provides ground for this line of research by providing investors, regulators, market authorities, academics, and practitioners with a new, promising 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.
- Andrews, D. W. K. (1991). Heteroscedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica* 59(3), 817–858.
- Bank of England (2019, July). Financial Stability Report.
- 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 (2022). Do designated liquidity providers provide liquidity during a flash crash? 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.
- Bouveret, A., M. Haferkorn, G. Marseglia, and O. Panzarino (2022). Flash crashes on sovereign bond markets EU evidence. ESMA Working Paper No. 1.
- 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.
- Campbell, J., S. Grossman, and J. Wang (1993). Trading volume and serial correlation in stock returns. Quarterly journal of economics 108(4), 905–939.
- Campbell, J., A. Lo, and C. McKinlay (1997). Econometrics of Financial Markets. Princeton University Press. Cespa, G. and T. Foucault (2014). Illiquidity contagion and liquidity crashes. Review of Financial Studies 27(6), 1615–1660.
- Cespa, G. and X. Vives (2022). Market opacity and fragility. Working paper.
- CFTC and SEC (2010). Findings regarding the market events of May 6, 2010.
- Chaudhuri, K. and Y. Wu (2003). Random walk versus breaking trend in stock prices: Evidence from emerging markets. *Journal of Banking & Finance* 27(4), 575–592.
- 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 (2022). The drift burst hypothesis. *Journal of Econometrics* 227(2), 461–497.
- 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.
- Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. The Journal of finance 65(4), 1237–1267.
- Easley, D., M. M. L. de Prado, and M. O'Hara (2012). Flow toxicity and liquidity in a high-frequency world. Review of Financial Studies 25(3), 1457–1493.
- Gao, C. and B. Mizrach (2016). Market quality breakdowns in equities. *Journal of Financial Markets* 28, 1–23. Golub, A., J. Keane, and S.-H. Poon (2017). High Frequency Trading and Mini Flash Crashes. Working paper.
- Griffin, J. M., P. J. Kelly, and F. Nardari (2010). Do market efficiency measures yield correct inferences? A comparison of developed and emerging markets. *The Review of Financial Studies* 23(8), 3225–3277.
- 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. Protter (2012). Discretization of Processes. Springer-Verlag.
- 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.
- 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.
- Lee, S. and P. Mykland (2008). Jumps in financial markets: A new nonparametric test and jump dynamics. Review of Financial Studies 21(6), 2535–2563.
- Lo, A. W. and A. C. MacKinlay (1988). Stock market prices do not follow random walks: evidence from a simple specification test. *The Review of Financial Studies* 1, 41–66.
- 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. (2016). The economics of high-frequency trading. *Annual Review of Financial Economics* 8, 1–24.
- Menkveld, A. J. and B. Z. Yueshen (2019). The flash crash: A cautionary tale about highly fragmented markets. $Management\ Science\ 10(10),\ 4470-4488.$
- Poterba, J. and L. Summers (1988). Mean reversion in stock returns: evidence and implications. *Journal of Financial Economics* 22, 27–60.
- 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.
- Shiller, R. J. (1981). The use of volatility measures in assessing market efficiency. The Journal of Finance 36(2), 291-304.
- Sigaux, J.-D. (2020). Trading ahead of Treasury auctions. Working Paper.
- Tsai, I.-C. (2018). Flash crash and policy uncertainty. *Journal of International Financial Markets, Institutions and Money* 57, 248–260.
- U.S. Securities and Exchange Commission (2010). SEC Concept Release on Equity Market Structure. Release No. 34-61358 (available at https://www.sec.gov/rules/concept/2010/34-61358.pdf).
- U.S. Securities and Exchange Commission (2016). SEC Docket. Vol. 115, No. 2 (Sep. 26-30, 2016), Release No. 34-78929 (available at https://www.sec.gov/litigation/admin/2016/34-78929.pdf).
- U.S. Securities and Exchange Commission (2020). Algorithmic Trading in U.S. Capital Markets. Staff Report, (available at https://www.sec.gov/files/marketstructure/research/algo_trading_report_2020.pdf).
- Weller, B. M. (2018). Does algorithmic trading reduce information acquisition? The Review of Financial Studies 31(6), 2184–2226.
- 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) The log-price p_t is a stochastic process defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathcal{P})$ and assumed to be an $It\bar{o}$ semimartingale described by the dynamics in Eq. (3.1), where p_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:

$$J_{t} = \int_{0}^{t} \int_{\mathbb{R}} \delta(s, x) I_{\{|\delta(s, x)| \le 1\}} \left(\nu(ds, dx) - \tilde{\nu}(ds, dx) \right) + \int_{0}^{t} \int_{\mathbb{R}} \delta(s, x) I_{\{|\delta(s, x)| > 1\}} \nu(ds, dx), \tag{A.1}$$

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 $u \in B_{\epsilon}(t)$,

$$E_{u \wedge s} \left[|\mu_u - \mu_s|^2 + |\sigma_u - \sigma_s|^2 \right] \le C_t^{(m)} |u - s|^{\Gamma},$$
 (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) $\int_{-\infty}^{+\infty} K(x) dx = 1$ and $K_2 = \int_{-\infty}^{+\infty} K^2(x) dx < \infty$;
- (K3) It holds that for every positive sequence $G_n \to \infty$, $\int_{|x|>G_n} K(x) dx \le CG_n^{-B}$ for some B>0 and C>0 (i.e., K has a fast vanishing tail);
- (K4) $m_K(\alpha) = \int_{-\infty}^{+\infty} K(x)|x|^{\alpha} dx < \infty$, for all $\alpha > -1$; $m_K'(\alpha) = \int_{-\infty}^{+\infty} K^2(x)|x|^{\alpha} dx < \infty$, 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 $H_n(t) = \frac{1}{\Delta_n} \sum_{t_i \le t} (\Delta_{i,n})^2$, we assume H(t) exists and is Lebesgue-almost surely differentiable in (0,T) with derivative H' such that: $|H'(t_i) - \Delta_{i,n}/\Delta_n| \le C\Delta_{i,n}$, for any t_i in which H is differentiable, for a suitable constant $C \ge 0$ which does not depend on i and n.

Proof of Theorem 1. Consider the $T_{\tau,n}^-$ statistic in Eq. (3.3). 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

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

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.

When dealing with SPY, given the huge amount of recorded transactions (roughly 480K transactions per day, more than 195M in total) we interpolate prices to a 1-second grid to ease data management.

C Implementation of the V-statistic

This section describes how the V-statistic is implemented on a single set of data to identify V-shapes. We describe our implementation using a specific day of trades, that is the Google stock on April 23, 2013, displayed in Figure 18, when a spectacular mini-flash crash, labelled the "Twitter flash crash", occurred (this followed the mini-flash crash of April 22, 2013, which the SEC used as evidence to charge Merrill Lynch, see Section 6.1). The procedure inputs are the trade times t_1, \ldots, t_n , the prices p_1, \ldots, p_n and the bandwidth h_n .

Step 1: fit the EGARCH(1,1)

We fit the EGARCH(1,1) model with a starting sample of 1-minute log-returns by maximum likelihood and we filter the volatility using estimated parameters. We use a t-student distribution for the residuals, where the degrees of freedom is a parameter to be estimated. Figure 18 shows the estimate of daily volatility obtained on the selected date. We see that the mini-flash crash is associated to a spike in volatility, since in the EGARCH(1,1) model there is no drift, so a misspecified driftless model can only attribute the price swing to volatility. The EGARCH(1,1) filtered variances will be used in the step 3 below, but also to compute price convexity using formula (6.1).

Step 2: select testing times

To select testing times, we look for local minima and local maxima in the time series with a simple algorithm that localizes the global minimum, excludes observations around the minimum over a given window, and then finds the next minimum. We implement this procedure for 10 minima and maxima with two windows: a short one of length $h_n/5$ (to disentangle between mini-flash crashes that may happen at a distance of a few seconds), and a long one of 20 minutes (to identify all local minima/maxima). In total we could than have up to 40 testing points. In the case in Figure 18, we test in 36 points denoted by a cross.

Step 3: compute the V-statistics at testing times

This is done by applying formula (3.2) with the provided bandwidth h_n . We use $K^-(x) = \exp(-|x|) \mathbb{1}_{\{x < 0\}}$ (left-sided), and $K^+(x) = \exp(-|x|) \mathbb{1}_{\{x \ge 0\}}$ (right-sided). 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.

Step 4: bootstrap

We construct artificial samples of n data with the same trade times $t_1, \ldots t_n$ and logarithmic prices generated by the following formula:

$$p_{t_i} = p_{t_{i-1}} + \varepsilon_t \sqrt{(t_i - t_{i-1})\widetilde{\sigma}_t^2}$$

where the local variances $\tilde{\sigma}_t^2$ are obtained by interpolating the EGARCH(1,1) filtered variances at trading times and ε_t is iid standard normal noise. On each simulated sample, we repeat Step 2 (that is, we test at local minima/maxima of each simulated trajectory) and 3 before, and compute the minimum of the V-statistics across testing times. Figure 19 shows the distribution of the minimum of the V-statistics on 5,000 bootstraps obtained for the day in Figure 18. Estimated quantiles are -11.42 at 1% and -9.23 at 5% (note: much larger than the asymptotic values of the asymptotic Bessel distribution, that would be strongly undersized for a case like this).

Step 5: draw conclusions

On this specific day, the minimum of the V-statistics is -13.47, thus largely significant at 1%, obtained at the testing point immediately after 13:00. This is the only recorded value in this day below the 1% limit.

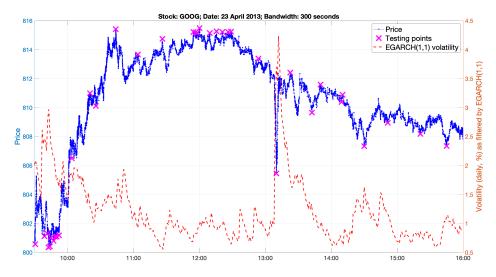


FIGURE 18: Blue dots: transaction prices for Google stock on April 23, 2013 (the day of the "Twitter flash crash"). Magenta crosses: testing points. Red dashed lines: filtered volatility from the EGARCH(1,1) model.

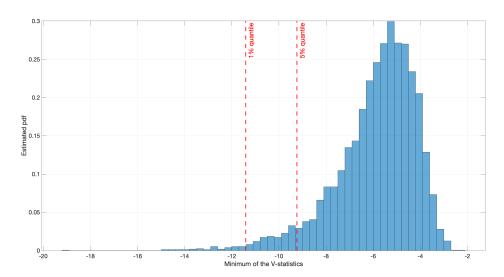


FIGURE 19: Histogram of the minimum of the V-statistics obtained on simulated data with the volatility filtered from the EGARCH(1,1) model on 5,000 replicas.

D The event study: data, macroeconomic background and the auction mechanism

D.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)¹². 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,
- \bullet a 7Y CCT with maturity 2025 and an Euribor-linked coupon rate, CCT-Eu Tv Eur6m+0.55% St25.

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 (3.3) and (3.2) 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;

¹²See CONSOB, Bollettino Statistico n. 14, June 2019, available at http://www.consob.it/web/area-pubblica/bollettino-statistico.

we choose a bandwidth h_n for the drift equal to 2 days¹³, while we base the long-run variance of the drift estimate on a 10-day bandwidth. We use exponential kernels $K^-(x) = \exp(-|x|) \mathbb{1}_{\{x<0\}}$ (left-sided), and $K^+(x) = \exp(-|x|) \mathbb{1}_{\{x\geq0\}}$ (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.

D.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 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. We notice we are not the first to connect V-shapes to political turmoil. Tsai (2018) associates large changes with political turmoils, justified using the model of Kyle (1985).

D.3 The auction mechanism of Italian government bonds

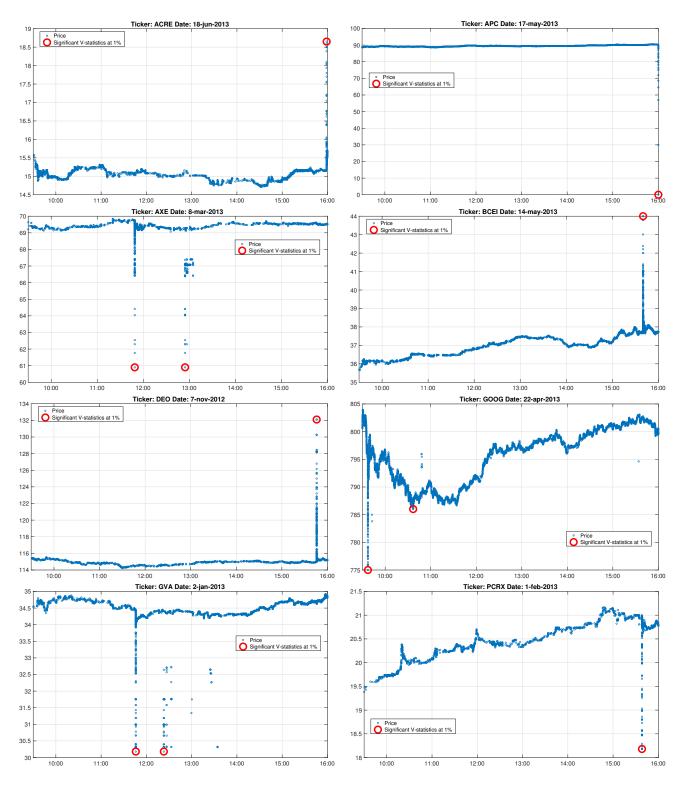
Before trading in the secondary markets, Italian government bonds are first issued on the primary market, where securities are allotted through electronic marginal auctions. The size of the allocations in the primary market in Italy is fairly predictable, since an annual auction calendar with a regular schedule is published at the beginning of the year (Sigaux, 2020). 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)

¹³Results, not shown here, are pretty robust to the choice of h_n .

auctions took place: the $9^{\rm th}$, $7^{\rm th}$, $3^{\rm rd}$ 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 D.1. The total nominal value of the three issuances were 2.159, 2.013 and 2.3 billion euros. Figure 15 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.

E Additional Figures



 ${\bf Figure~20}:$ Mini-flash crashes used by SEC in legal actions against Merrill Lynch.

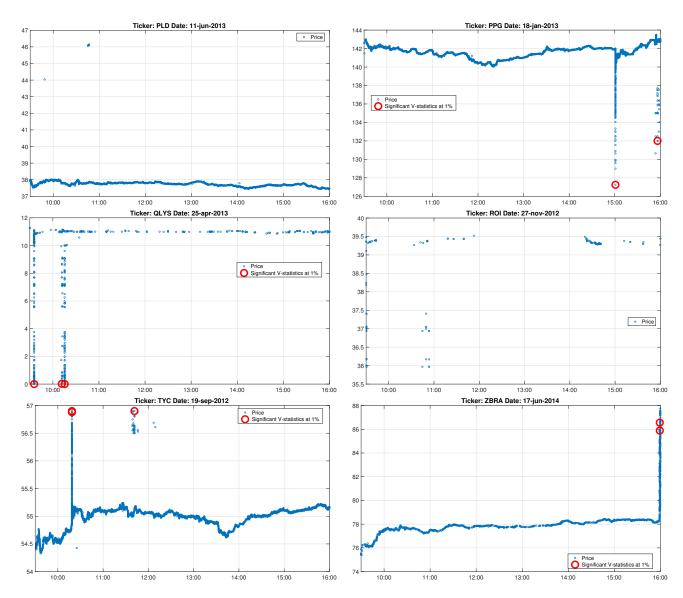
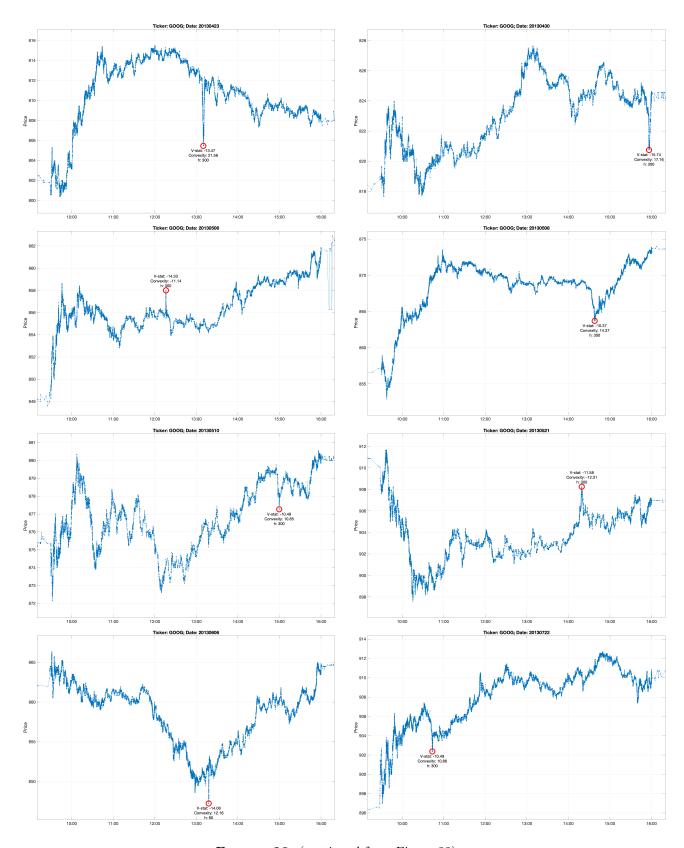


FIGURE 21: (continued from Figure 20)



Figure 22: V-shapes in the Google stock in 2013.



 $\textbf{Figure 23:} \ (\text{continued from Figure 22})$

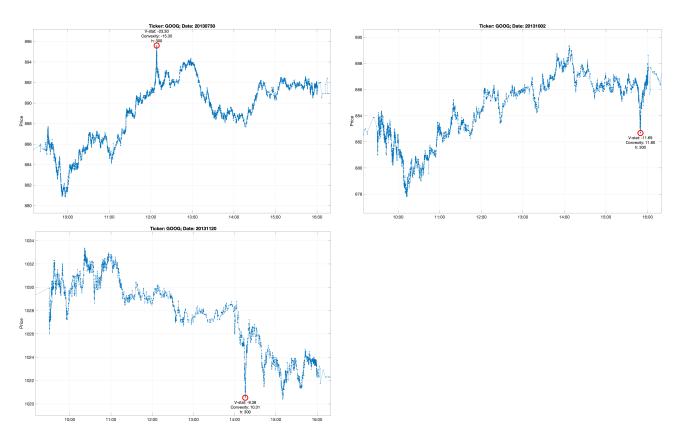


Figure 24 : (continued from Figure 22)