School of Finance

University of St.Gallen

«Gamma Fragility»

Andrea Barbon
Andrea Buraschi

Working Papers on Finance No. 2020/05

Swiss Institute of Banking and Finance (s/bf - HSG)

November 2020

Gamma Fragility

Andrea Barbon Andrea Buraschi

November 5, 2020

Abstract

We build on a growing literature that studies the impact of market frictions on the dynamics of stock markets, such as momentum, price spirals, excess volatility, and investigate the potential feedback effects of delta-hedging in derivative markets on the underlying market. We document a link between large aggregate dealers' gamma imbalances in illiquid markets and intraday momentum/reversal and market fragility. This link is distinct from information frictions (adverse selection and private information) and funding liquidity frictions (margin requirement shocks). We test our joint hypothesis using a large panel of index and equity options that we use to compute a proxy of aggregate gamma imbalance. We find supporting evidence that intra-day momentum (reversal) is explained by the interaction of negative (positive) aggregate ex-ante gamma imbalance and market illiquidity. The effect is stronger for the least liquid underlying securities. The result helps to explain both intra-day volatility and autocorrelation of returns. Moreover, we find that gamma imbalance correlates to the frequency and the magnitude of flash crash events.

Keywords: Frictions, Momentum, Option Markets, Risk Management, Gamma Imbalance, Flash Crashes, Liquidity

First version: January 2019

This version: 5 November 2020

Andrea Buraschi is Chair of Finance at Imperial College Business School; Andrea Barbon is Assistant Professor in fiancne at the University of St.Gallen. Emails: andrea.buraschi@imperial.ac.uk, andrea.barbon@unisg.ch.

A growing literature has documented the existence of temporary and permanent components in equity returns (Fama and French (1988), Daniel, Hirshleifer, and Subrahmanyam (1998)).[1] Some of these effects are found at daily frequencies (price spirals) while others manifest over longer horizons (momentum). In this paper, we focus on the first set of properties. Indeed, at intraday frequencies, several studies document that stock prices deviate from a random walk.

Figure 1 provides a simple summary of the properties of the autocorrelation coefficient ρ of intraday returns sampled at 5-minutes frequency using 12 non-overlapping observations in a 60 minutes window based on a large sample of liquid stocks in the TAQ dataset in the period 1996 to 2014. [2] On the y-axis the figure reports the probability that the absolute value $|\rho|$ is larger than $\bar{\rho}$ with $\bar{\rho} = 0.10, 0.20, \ldots, 0.90$. If stock prices were random walks, the autocorrelation coefficient ρ would be zero. The results strongly reject this hypothesis and show evidence of non-zero intraday autocorrelation. For instance, 34% of the time the autocorrelation coefficient is either larger than 0.30 or lower than -0.30. Moreover, if we run a regression of ρ on its lag we find that the slope coefficient is positive with a t-statistics equal to 54.17, suggesting a significant persistence of ρ at the stock level.

These properties are not only statistically strong but also interesting from an economic point of view as they deviate from what one would expect in a simple frictionless market. Recently, some of the biggest investment banks have started to offer investable financial products based on strategies that use intraday movements (momentum and reversion) of underlying assets.[3] Therefore, a natural question that emerges is whether these patterns are the outcome of some deeper structural properties of financial markets.

[1] Examples include under- and over-reaction to news, momentum, mean reversion, excess volatility, and price spirals.

[2] This provides us with 4,777 trading days and 6,039,248 hourly observations of ρ_i across all stocks i.

[3] See, for instance, the Morgan Stanley Intraday Indexes (quoted on Bloomberg) MSCBL1ES, MSCBL1NQ, and MSCBL1RT which are constructed on the SPX, Nasdaq, and Russell Indexes, respectively.

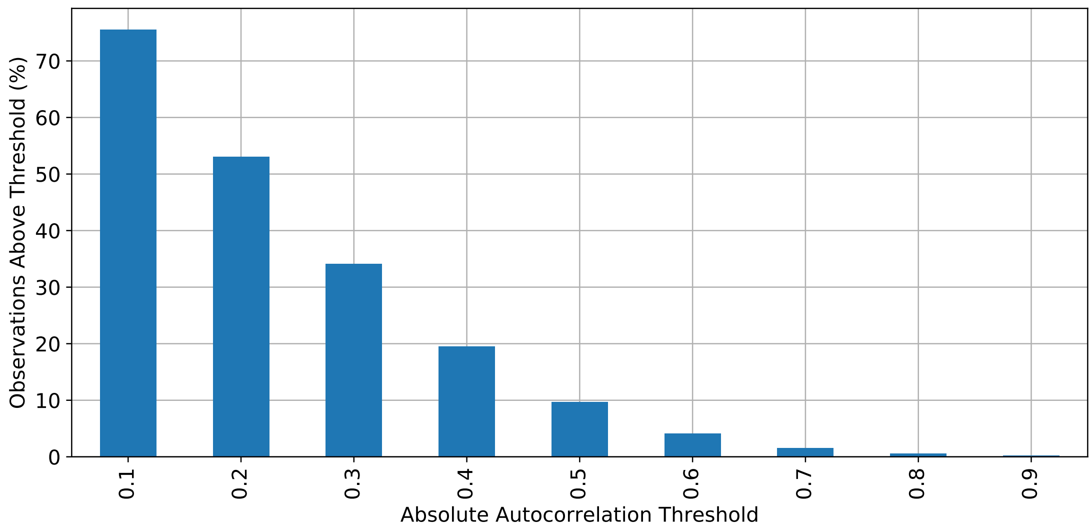

Figure 1. Magnitude of Intra-day Autocorrelation
The figure describes the magnitude of autocorrelation of intra-day stock returns. We define $\rho = \rho(j, d, h)$ as the auto-correlation coefficients of 5-minutes returns of stock j during the trading hour h of day d, thus based on 12 non-overlapping observations. The figure displays the probability that $|\rho| > \bar{\rho}$ for $\bar{\rho} \in [0.1, 0.2, \ldots, 0.9]$, based on the empirical distribution of ρ sampled for a set of 276 individual U.S. equity stocks over the period from 1996 to 2017 (6,039,248 observations).

The literature has proposed a few explanations to explain some of these regularities. A first explanation emphasizes the role played by economic frictions, such as margin requirements and capital constraints, in the formation of price spirals (Brunnermeir and Pedersen (2009)) and excess volatility. However, these explanations find it difficult to explain the existence, at the same time, of both positive and negative autocorrelations. A second explanation builds on how information gets incorporated into asset prices. Easley, O'Hara, and Srinivas (1998) formally study a model that allows for endogeneous participation of informed traders in the options market. They show the existence of a pooling equilibrium in which informed investors use both option and stock markets if the leverage implicit in options is large and the liquidity in the stock market is low. As a result, equity returns may display predictability. Pan and Poteshman (2006) find supporting evidence for the existence of an information-based channel linking the trading activity of privately informed traders in the options market and predictability in stock returns. Indeed, several studies have argued that informed investors might choose to trade derivatives because of the higher leverage offered by such instruments, e.g. Black (1975). Easley, Lopez de Prado, and O'Hara (2012) propose a measure of *flow toxicity* which captures states of the world in which there is a high probability that market demand is affected by the presence of informed investors. In these states, order flow adversely selects market makers, who may be unaware they are providing liquidity at a loss. As a consequence, markets can become fragile, volatility and autocorrelation may increase, and flash crashes are more likely to occur. A third growing literature emphasizes the importance played by behavioral biases.[4] This paper focuses on a different type of friction that is linked to the role played by the derivative market during periods of market illiquidity. It differs from information-based explanations and it may occur even in the absence of margin and capital constraints.

In the last 20 years, the use of derivatives for hedging has seen a massive increase due both to growing emphasis on asset-liability management and to explicit regulatory constraints. Solvency II, Dodd-Frank, and the EMIR Risk Mitigation Regulation increased the cost of capital in favor of risk mitigation techniques including hedging and reduction of counterparty risk. In particular, insurance companies have increased their use of derivatives to manage the embedded optionalities which are normally provided as part of their contracts and reduced the use of reinsurance. A life insurer, for instance, with a large portfolio of guaranteed minimum death benefits (GMDB) and variable annuities (VA) can hedge against a drop in equity markets using put options. As hedging reduces required economic capital, derivatives have become a widely used ex-ante mechanism for insurance companies to relieve excessive con-

[4]Initial underreaction to new fundamental information has been linked to the disposition effect (Frazzini (2006), Barberis and Xiong (2009), Shefrin and Statman (1985), Ben David and Hirshleifer (2012)). Daniel, Hirshleifer, and Subrahmanyam (1998) and Luo, Subrahmanyam, and Titman (2018) argue that the overestimation of one's own precision of private information signals (own ability) can help explain post-corporate post-earnings announcement stock price 'drift', negative long-lag autocorrelations (long-run 'overreaction'), and excess volatility of asset prices. Moreover, if in addition agents are affected also by a self-attribution bias, one may justify momentum (positive short-lag autocorrelations).

centration of risk and generate business value.[5] As a consequence, in 2017 the notional value of insurance industry derivatives reached $2.3 trillion, with roughly 95% held by life/annuity insurers and with options constituting 43% of this exposure.[6]

The aggregate net positions of options is by definition equal to zero. However, there is significant heterogeneity in the positions of different types of options (over maturity and strikes) across different types of institutions. This reflects institutional differences in preferences, technologies and possibly different strategic objectives across institutions and retail investors in the supply chain of options. Indeed, several empirical studies find that individual agents are on average net buyers of index put options, both directly and indirectly via insurance and structured products. Insurance companies provide these optionalities and buy protection using liquid options in centrally cleared markets. In equilibrium, the risk is transferred over the option supply chain to market makers who are left, depending on their risk aversion, with the ultimate task to manage their risk exposure using dynamic hedging techniques. See Figure 2.[7]

In absence of any risk management, a dealer profit profile is potentially very volatile and non-linear. If he sold put options, he would be long delta and short gamma: if the value of the underlying asset drops, his put options become more in-the-money and their deltas increase. To limit their market exposure, most dealers delta-hedge by selling shares of the underlying asset. The optimal amount of shares to short, however, changes depending on the intraday fluctuations of the underlying price. When the aggregate market gamma imbalance of dealers is large, the overall intraday activity of financial intermediaries may be substantial

[5]An example of the preference of using derivatives to hedge instead of reserve-capital to back guarantees is made explicit by Generali INA-Assitalia (a large insurance company), which states: "we view our life insurance policies as nothing less than a portfolio of embedded options... We use liquid derivatives to replicate the liability portfolio." (*Risk*, November 1999)

[6]See, Best's Special Report (2018), "Growing Use of Derivatives for Liability Risk Management".

[7]No market maker would simply sell puts (or call) options to hold the naked position on the underlying asset overnight. Such directional bets would simply be too large and not justifiable.

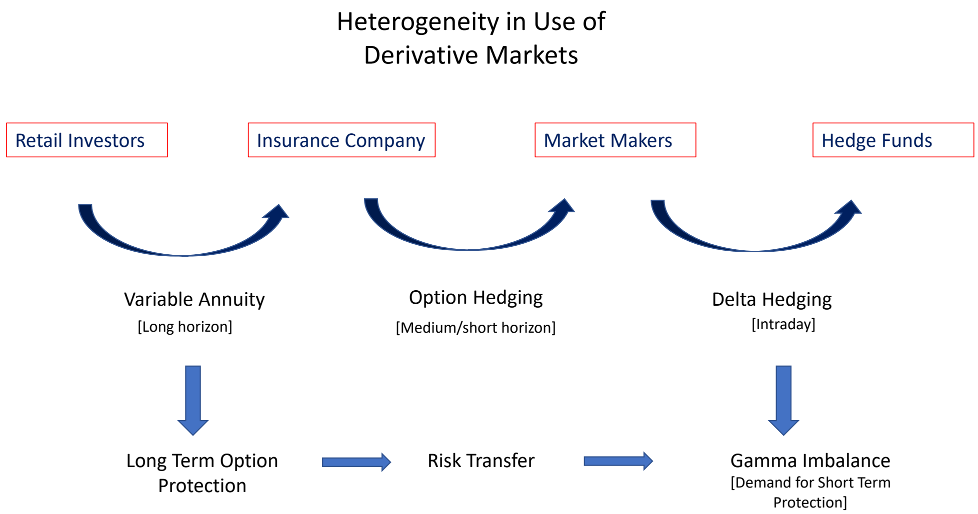

Figure 2. Risk Transfer in the Supply Chain of Options Markets

and may add additional pressure to an initial move of the underlying asset, giving rise to intraday *momentum*. If the dealer's book is made, on the other hand, of long positions in out-the-money call options, the dealer is long delta and long gamma. In this case, an initial price drop induces the dealer to buy the underlying asset, reducing pressure to the underlying asset and giving rise to intraday *mean-reversion*. The potential role of this channel to help to explain these observed effects depends on the extent of heterogeneity in both the supply chain of options and on the incentives by market makers to hedge intra-daily. Differently than individual investors, most market makers have an institutional mandate to hedge their derivative positions by the end of the trading day. However, the specific dynamics during the day typically vary among dealers.

The sign of the intraday serial correlation is opposite depending on the aggregate composition of the book. This provides an interesting set of joint testable implications. Let us define the aggregate dollar value of all market makers outstanding gamma with their clients as the *gamma imbalance*, which accounts for all options positions of different strike prices and maturity. When the aggregate gamma imbalance is large and negative, one should observe larger market volatility and short-term momentum (positive serial autocorrelation). On the other hand, when the gamma imbalance is positive, one should observe lower than average volatility and short-term mean reversion (negative serial autocorrelation). Notice that this specific effect is independent of the presence of *information* frictions, such as order flow toxicity, and of the additional role played by economic frictions such as margin constraints. It does, however, require that the market is not infinitely liquid to absorb intraday market makers demand shocks. This suggests an additional testable implication: the effect should be stronger (in the cross-section) for less liquid stocks and (in the time-series) during less liquid times since in a perfectly liquid and frictionless market options are redundant assets and no feedback effect should exist between options and their underlying.

To investigate the extent to which gamma imbalance can help to explain the observed intraday patterns in equity markets, we use the IvyDB dataset from OptionMetrics to construct a comprehensive dataset of equity options that include both index and a large cross-section of individual stock options. Information on both open interest and option specific greeks are used to obtain a panel on day- and security-specific gamma imbalance. The dataset covers the period between 1996-2017. Intra-day returns and the trading volume for each of the underlying assets are from TAQ. The cross-sectional dimension of the dataset is important since at each given point of time, a large number of assets have negative and positive gamma imbalance. This helps to identify the effect and increases the power of our tests.

We focus on three main questions. First, what is the extent to which gamma imbalance explains asset volatility? Limited market liquidity and the existence of an institutional friction suggests that volatility is larger when gamma imbalance is more negative. We run a panel regression of the absolute value of the return of underlying asset j on day t on our proxy for the gamma imbalance of options dealers measured at time $t-1$. As predicted, the data shows a negative relationship between daily return volatility and gamma imbalance.

The results are strongly statistically significant and robust to a series of controls, such as alternative definitions of intraday volatility and after conditioning for time and other common effects. We also find that most of the impact of this friction is transitory and is not present after 2 trading days. Moreover, using the cross-sectional dimension of the dataset, we find that the increase in volatility during negative gamma days is significantly stronger for the least liquid assets. This is consistent with the hypothesis of a dealer order flow friction that becomes more binding in illiquid markets.

Second, what is the extent to which gamma imbalance explains intra-day momentum and reversal? We run a panel regression of the autocorrelation coefficients of h-minute returns of stock j onto date and security-specific gamma imbalance. We find that the slope coefficients are negative and statistically significant, supporting the null hypothesis. We also find that the magnitude and significance of the coefficients is increasing with the frequency and peaks at a horizon of $h = 30$ minutes, consistent with an economy in which dealers adjust their delta-hedged portfolios intraday at a similar frequency.

Third, we investigate the link between gamma imbalance and flash crashes, that is, large price drops materializing over a short time period. Using the "drift burst" detection methodology proposed by Christensen, Oomen, and Renò (2018), we identify a panel of flash crash events and study their relationship to gamma imbalance. We find that conditional on negative ex-ante gamma imbalance flash crashes are more likely to occur and to be larger in magnitude. This effect is economically significant in the first part of our sample, while it becomes less relevant after the introduction of circuit breakers in 2010.

Finally, we compare the economic importance of institutional frictions such as gamma imbalance versus alternative explanations, such as flow toxicity. Easley, Lopez de Prado, and O'Hara (2012) argue that when order flow adversely selects market makers, who may be unaware they are providing liquidity at a loss, it can become toxic. We use a large dataset on volume imbalance and trade intensity to construct Easley, Lopez de Prado, and O'Hara (2012) proxy of order flow toxicity (VPIN) sampled at *volume* time, instead of traditional clock time. We study the relative economic and statistical importance of these two channels and find that gamma imbalance retains its significance even after controlling for VPIN, suggesting that it captures an additional independent driver of price volatility.

Related Literature. This paper is related to several streams of the asset pricing literature. A first stream studies the relation between economic frictions, liquidity, and market fragility – see, among others, Holmstrom and Tirole (1997) and Brunnermeir and Pedersen (2009). The last investigate the link between an asset's market liquidity (i.e., the ease with which it is traded) and traders' funding liquidity (i.e., the ease with which they can obtain funding). They show that, under certain conditions, margins are destabilizing and market liquidity and funding liquidity are mutually reinforcing, leading to liquidity spirals. They argue that speculators' capital can become a driver of market liquidity and risk premiums. Their work is part of a growing literature. See among others, Adrian, Moench, and Shin (2014), Gromb and Vayanos (2002), Garleanu and Pedersen (2001), Adrian, Etula, and Muir (2014), Adrian and Shin (2010), Adrian and Boyarchenko (2015), Adrian, Colla, and Shin (2012).

Our work is also related to an important literature that studies the impact of asymmetric information and adverse selection in the context of informed trading on asset prices. Information-based models (e.g., Glosten and Milgrom (1985), Easley, O'Hara, and Srinivas (1998)) suggest that while stock prices will fully adjust when all public information is revealed, they may gradually adjust to the private component of information. As a result, stock returns may display predictability and serial correlation. Pan and Poteshman (2006) test this hypothesis using information from options markets and show that stocks with low put-call ratios outperform stocks with high put-call ratios by more than 40 basis points on the next day and more than 1% over the next week. They interpret the economic source of this predictability as private information possessed by options traders rather than market inefficiency. Easley, Lopez de Prado, and O'Hara (2012) contribute to this literature proposing a new proxy of order toxicity that is sensitive to the probability that market dealer face privately informed traders. Their measure is based on volume increments, as opposed to time increments, and find that it helps to explain market volatility. Engle and Lange (2001) propose an intraday measure of market liquidity (VNET) that captures the depth of the market corresponding to a particular price deterioration. They use the excess volume of buys or sells (order imbalance) measured over price increments to obtain VNET and find that it helps to explain volatility.

Our paper provides further evidence suggesting that trading on derivatives may affect the price of the underlying securities. Ben-David, Franzoni, and Moussawi (2018) shows that the arbitrage activity of market makers in the ETF market generates excess volatility on the underlying stocks. Shum, Hejazi, Haryanto, and Rodier (2015) find that leveraged ETFs drive up volatility near the market's close because of hedging activities of ETF providers. More recently, growing interest has focused on the role played by high-frequency traders (HF) on flash crashes and other intraday price dynamics. Hendershott and Riordan (2013) present evidence on HF trading on the Deutsche Borse. Brogaard (2010) and Hasbrouck and Saar (2013) analyze the role and strategies of HF traders in U.S. equity markets. Kirilenko, Kyle, Samadi, and Tuzun (2017) characterize the behavior of HF traders and other market participants in the S&P 500 futures market. Hendershott, Jones, and Menkveld (2011) study the empirical relationship between algorithmic trading and liquidity, finding that algorithmic trading improves liquidity for large stocks.

Our paper relates also to a third stream of the literature that investigates the effect of institutions on asset prices. In this area, important works include Vayanos and Woolley (2013) and Gorton, Hayashi, and Rouwenhorst (2013). The first suggests a theory of momentum and reversal based on flows between investment funds. They argue that flows are triggered by changes in fund managers' efficiency, which investors either observe directly or infer from past performance. Momentum arises if flows exhibit inertia, and because rational prices under-react to expected future flows. Reversal arises because flows push prices away from fundamental values. Hendershott and Seasholes (2007) study explicitly the link between non-informational order imbalances (buy minus sell volume) to predict daily stock returns at the market level. Additional evidence on the role of institutions in the correlations of asset returns is also discussed in Anton and Polk (2014), Chang, Hong, and Liskovich (2013), Greenwood and Thesmar (2011), Lou and Polk (2013), and Jotikasthira, Lundblad, and Ramadorai (2012). In the context of commodity markets, Gorton, Hayashi, and Rouwenhorst (2013) argue that momentum can be explained by specific market frictions in the context of the *Theory of Storage* and *Normal Backwardation*. Indeed, they show that momentum portfolios take positions in similar commodities as those obtained from basis-sorted portfolios. Thus, high momentum portfolios select commodities with below normal inventories and relatively high bases, while the Low Momentum portfolio does the opposite. They argue that, in the event of a stock-out the convenience yield is positive to reflect a spot price increase due to a shortage of goods. They interpret the result as evidence of a friction-based explanation of momentum.

I. An Example: General Motors

To illustrate the potential effect of gamma imbalance on intraday stock price dynamics, consider the time series properties of General Motors stock price on October 13^{th} 2014 and on January 9^{th} 2009. Figure 3 summarizes the two time series. The first panel displays the evolution of the price on October 13^{th} 2014, when the gamma imbalance of option dealers on GM at the beginning of the trading day was deeply negative, amounting to more than 11 million dollars. This is a significant amount, equivalent to more than 21% of the average hourly volume for the underlying stock. If the imbalance is due to short OTM put options positions, intermediaries can limit their market exposure by delta-hedging with short positions in the stock. At around 13:30 the price starts dropping sharply, driving a large amount of put options closer to being ATM. Thus, to maintain a constant delta exposure,

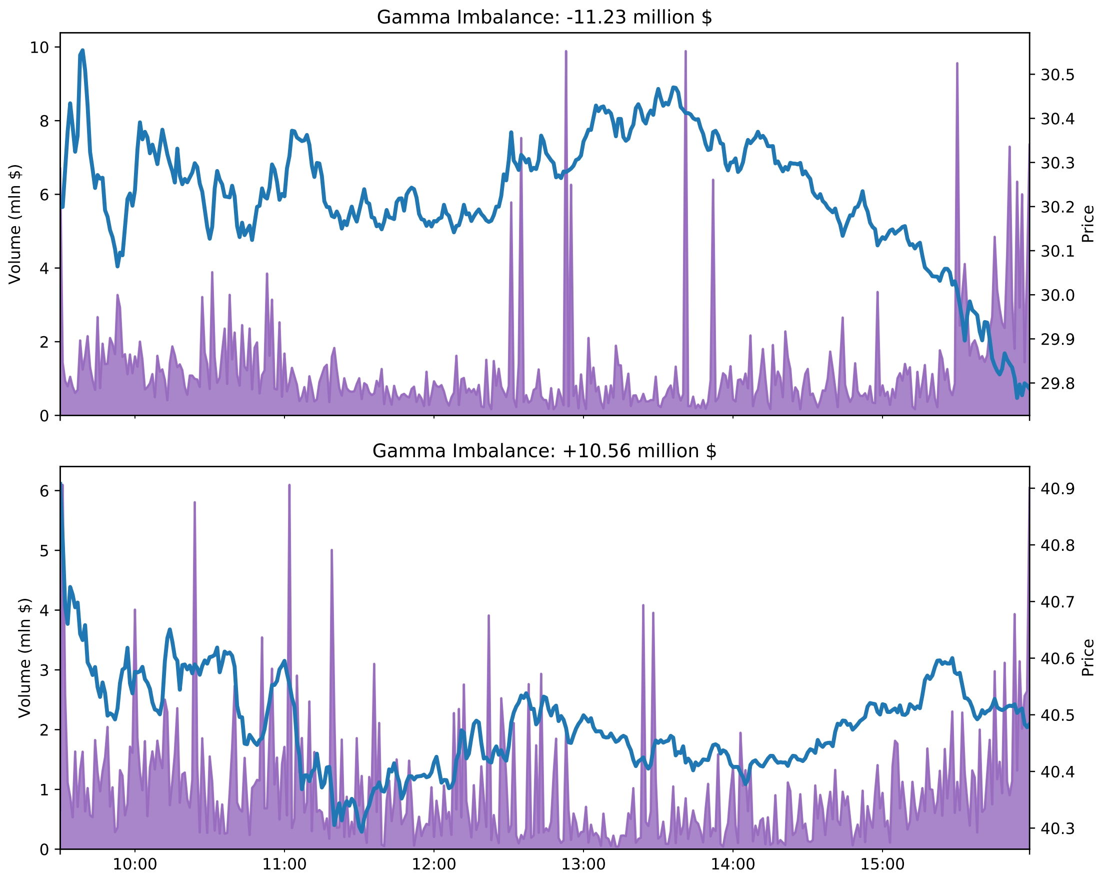

Figure 3. General Motors Stock Price

intermediaries should increase their short positions by selling more stocks. As the stock price continues to fall throughout the rest of the trading day, delta-hedging would induce dealers to increase even further their short positions. The effect is particularly strong during the last 30 minutes of the trading day when we observe a significant rise in trading volume. On this day, we find that the autocorrelation of hourly returns is +43%. Our conjecture is that the automatic risk-management order flow has contributed to the intraday momentum.

The bottom panel shows another, albeit opposite, example. At the beginning of January 9^{th} 2009, the aggregate gamma imbalance was significantly positive, due to long OTM call options positions. If the intermediary wanted to hedge market exposure, delta hedging would require the dealer to short the underlying asset. In this case a price drop reduces the call option's delta, thus inducing intermediaries to buy back some equity. If the flow is significant relative to the available liquidity, it could dampen the initial negative price shock. Consistently, the plot shows evidence of strong intra-day mean reversion with an autocorrelation coefficient equal to -63%. In these two examples, the sign of the autocorrelation coefficient is negatively related to the sign of the ex-ante aggregate gamma imbalance. Moreover, while in the first example the intraday (annualized) volatility is 24%, which is above the unconditional average, in the second case the volatility is at 18%, below the average. Are these two examples simply a coincidence or are they revealing the existence of an important friction?

II. Data and Summary Statistics

This section discusses the data sources, the definition of the main variables we use and documents the main summary statistics.

A. Data Sources

We obtain data on daily options prices, open interest, and greeks from the IvyDB dataset provided by OptionMetrics. We focus on stocks and indices that have liquid option trading by selecting the 300 underlying asset with largest average dollar open interest reported in OptionMetrics over the sample period from 1996 to 2017.

On each day, we compute the total open interest for each asset in our universe for both call and puts across all strike price levels and maturity available. Then, we merge this dataset with the CRSP and TAQ databases by ticker to obtain information on daily and intra-daily returns and trading volume for the underlying assets, respectively.

We also obtain data on market fundamentals from the News Analytics dataset by RavenPack, which provides sentiment scores that is useful to identify asset-days with significant news.

The final sample includes four market indices (S&P 500, Dow Jones, Nasdaq 100 and Russell 2000) and 276 individual U.S. equity stocks. Our panel is, by construction, unbalanced because some of the single-name options enter the sample only when they start being traded.

Figure 7 shows the distribution of intra-day autocorrelation coefficients for equity returns sampled at 5-minutes intervals. The median is slightly positive and equal to 10 percent. We cannot reject the null hypothesis that the unconditional average of the autocorrelation coefficient is zero. However, we find strong evidence both of conditional intra-day momentum (positive autocorrelation) and reversal (negative autocorrelation). Indeed, about 24% (10%) of the days had an intra-day autocorrelation greater than $+0.30$ (lower than -0.30).

Insert Figure 7 here

B. Gamma Imbalance

Let the value of the underlying asset at time t be S_t. The delta Δ_t of an option $C_t(S_t; K, T)$ is defined as the first derivative of the option price with respect to the underlying price $\Delta_t = \frac{\partial C}{\partial S_t}$. At time t, delta-hedging of an option portfolio requires buying or selling an amount of the underlying asset equal to $-\Delta_t$. Since changes in S_t changes the value of Δ_t, delta hedging strategies require a dynamic adjustment of the position on the underlying asset. The greek Γ_t measure the rate of change of Δ_t given a change in the underlying asset, i.e. $\Gamma_t = \frac{\partial \Delta_t}{\partial S_t}$, and is proportional to the convexity of the value of the derivative security with respect to S_t. Both call and put options are convex in S_t. Thus, a long portfolio of options implies a positive Γ_t, which implies that the size of the delta-hedging position is positively related to the underlying price. On the other hand, a short portfolio of options implies a negative Γ_t, i.e. the size of the delta-hedging position is negatively related to the underlying price.

Financial institutions heavily rely on option markets to transfer some of the risk embedded in the contracts offered to their clients. The traded options market has become a solution of choice, due to their liquidity and limited counterparty risk. A large literature documents that in certain periods the book of financial intermediaries can be imbalanced. This would occur, for instance, when the aggregate demand of their customers contributes to a build-up of an excess demand for either puts or calls in a particular index or individual stock. We are interested in a measure of this imbalance (the gamma imbalance) and compute it at the day and asset-specific level.

First, we compute the dollar open interest for each option contract, taking the product of the open interest $OI(t)$ and the contract price $P(t)$[8]

$$OI^{\$}_{Type,K,T}(t) = OI_{Type,K,T}(t) \cdot P_{Type,K,T}(t), \qquad Type \in Call, Put. \tag{1}$$

Second, since each option has a gamma which depends on its moneyness and time to expiration, we compute the dollar gamma of each specific call and put option on each trading day t. Then we compute the value-weighted sum of the gammas of all option contracts to obtain $\Gamma^{\$}$ as follows:

$$\Gamma^{\$}_{Type}(t) = \sum_{K,T} OI^{\$}_{Type,K,T}(t) \cdot \Gamma_{Type,K,T}(t), \qquad Type \in Call, Put. \tag{2}$$

Finally, since regulatory restrictions limit the amount of information available about the specific identity of derivative holders, we build two alternative proxies of dealers gamma imbalance depending on alternative assumptions about their positions. If dealers are long call options and short put options, the aggregate imbalance is proportional to the difference between the aggregate gamma of the call minus the aggregate gamma of the put options,

[8]the contract price is computed as the arithmetic average of the best bid and ask offers for the contract.

$\left[\Gamma^{\$}_{Call}(t) - \Gamma^{\$}_{Put}(t)\right]$. Accordingly, we define our first proxy $\Gamma^{\$}_{1}(t)$ as a fraction of the average daily volume for the underlying security:

$$\Gamma^{IB}_{1}(t) = \left[\Gamma^{\$}_{Call}(t) - \Gamma^{\$}_{Put}(t)\right] / \text{ADV}(t) \times 100. \tag{3}$$

The average daily volume ADV(t) is computed as the rolling average of the dollar trading volume over the quarter previous to t.

If dealers are long both call and put options, the aggregate imbalance is proportional to the sum of the aggregate gamma of calls and puts options, $\left[\Gamma^{\$}_{Call}(t) + \Gamma^{\$}_{Put}(t)\right]$. Thus, we define

$$\Gamma^{IB}_{2}(t) = \left[\Gamma^{\$}_{Call}(t) + \Gamma^{\$}_{Put}(t)\right] / \text{ADV}(t) \times 100. \tag{4}$$

We design the empirical test to be agnostic with respect to the specific structural imbalance of dealers. However, Pan and Poteshman (2006) compare the total open-buy and open-sell option volume and document that put (call) options on the S&P500 and Nasdaq are on average bought (sold) by investors that are not market makers.[9] According to their results, Γ^{IB}_{1} would be the preferred proxy of gamma imbalance. On each trading day, we use Option Metrics to obtain data on the open interest and gamma of all strike and maturity-specific put and call options, as of the end of trading day. In all empirical applications, we lag Γ^{IB}_{1} by one trading day to avoid any forward-looking bias in the analysis.

Table I reports summary statistics on Γ^{IB}_{1} and the open interest for the option contracts in our sample. We also report statistics on the implied volatility, computed as the average across all outstanding contracts for each day-asset pair. Equity and index options exhibit similar statistical distributions for our gamma imbalance proxy while, as expected, index options have a significantly larger open interest and a substantially lower implied volatility relative to single name options. Our identification strategy is based on both the time-series

[9]For individual stock options, the difference varies over time and in the cross section.

and the cross-sectional variation of Γ_1^{IB}. In Figure 4, the darker line highlights the cross-sectional average on each date t; the light (dark) shaded area represents the 10th and 90th (25th and 75th) percentiles interval of the daily distribution of Γ_1^{IB}.

Two main properties emerge. First, on almost every day in the sample Γ_1^{IB} ranges from significant negative to positive values across different individual stocks. Moreover, the cross-sectional average displays rich time-variation. These two properties are important since our identification strategy uses both time-series and cross-sectional variation.

Insert Figure 4 and Table I here

Two channels drive the time-variation in Γ_1^{IB}. The first one is due to a quantity effect: the portfolio decision of institutional and retail investors to buy/sell options and the willingness of option dealers to accept the trade. The second one is due to a price effect: the endogenous impact on Γ_1^{IB} induced by a variation in the price of the underlying asset. This second effect depends on the entire distribution of the moneyness and maturities of the call and put options held in dealers' book; intraday shocks to the underlying price have immediate implications on the aggregate Γ_1^{IB}. Figure 5 shows the relative importance of the two channels. On November 15th, 2018, the S&P500 index was 2,730.27 and the aggregate Γ_1^{IB} outstanding was negative. As the blue curve shows, an increase in the S&P500 would also increase Γ_1^{IB} and, in absence of other portfolio changes, the aggregate Γ_1^{IB} would change sign and become positive if the S&P500 were to move above 2762.5 (a 2.3 percentage move). The shape of the curve depends on the specific distribution across moneyness and maturity of the dealers' options book. Keeping constant the S&P500 level, one can compare the relative importance of these two channels, on this particular day, by comparing the vertical distance between the two curves (quantity effect). Our data allows us to track this information at the daily frequency.

Insert Figure 5 here

III. Empirical Results

In this section, we present the main empirical results of the paper. We articulate our analysis around three main null hypotheses.

A. *Volatility and Gamma Imbalance*

The first empirical null hypothesis relates to the assumption that dealers' order flow in the underlying assets has a non-trivial intra-day price impact.

$H0_1$: *Does Gamma Imbalance help to explain intraday market volatility?*

In particular, we investigate the existence of a negative link between daily volatility and dealers' gamma imbalance. When dealers' gamma is positive (negative), their delta increases (drops) when the underlying asset increases. Thus, their delta-hedging strategy requires selling (buying) more of the underlying asset following an increase in the underlying price. Therefore, dealers order-flow acts as a contrarian (reinforcing) force, thus limiting (strengthening) the magnitude of initial price movements.

It is important to notice that there are many reasons why one may expect the null hypothesis to be rejected. First, dealers may be less risk-averse than commonly perceived and less interested to aggressively hedge their delta imbalances. Second, the underlying asset market could be sufficiently liquid and frictionless so that dealers' delta hedging strategies have no price or volatility impact. Third, dealers could be extremely rational and technologically sophisticated to be able to implement hedging strategies with no aggregate price impact. Fourth, while this channel might be realistic, it might not be sufficiently strong to dominate other equally important channels. Finally, the construction of a proxy of gamma imbalance is notoriously difficult due to regulatory restrictions that, in some cases, limit the amount of information necessary to identify all trading counterparties. The noisier the proxy, the less powerful the test, and the more likely one may reject the null hypothesis $H0_1$ even if there were a relationship, thus biasing the test against finding a result.

To test this hypothesis we run a panel regression of absolute returns on stock j conditional on the aggregate gamma imbalance of option dealers $\Gamma^{IB}_{j,t-1}$ for security j at time $(t-1)$:

$$|R_{j,t}| = \alpha + \beta_0 \Gamma^{IB}_{j,t-1} + \beta_1 IVOL_{j,t-1} + FE_{j,t} + e_{j,t}, \tag{5}$$

where the term $FE_{j,t}$ refers to different time and asset fixed-effects. Standard errors are clustered by asset and month to allow for arbitrary correlation structures both in the time series and cross-section. We control for the previous day implied volatility $IVOL_{j,t-1}$ computed as the average across all the outstanding options.

The results are summarized in Table II and show that the coefficient on $\Gamma^{IB}_{j,t-1}$ is negative and highly significant with a t-statistics ranging between -4.54, after controlling for lagged implied volatility, and -5.32, after controlling also for asset and month fixed effects. Indeed, consistent with our null hypothesis, the data shows a negative relationship between daily volatility of returns and gamma imbalance. In column two, three, and four we control for asset, time, and both time and asset fixed effects, respectively. The statistical significance of the results are even stronger. These results are also significant from an economic standpoint. The estimates imply that one standard deviation in the gamma imbalance is associated to a decrease of roughly 12 basis points in the absolute value of the daily return for the underlying stock. This constitutes a non-negligible impact, given that absolute daily returns average at around 109 bps in our sample.

Insert Table II

Figure 6 (top panel) shows the distribution of t-statistics of the estimator of β_0 when we run time-series regressions for each of the 280 stocks in our sample. T-statistics are computed using the heteroskedasticity and autocorrelation consistent (HAC) standard errors from Newey and West (1986). We find that the distribution is strongly shifted to the left (of zero): for 206 out of 280 stocks (73% of the population) the t-statistic is more negative than -2.5. This is indeed a large proportion which highlights a strong negative link between aggregate Gamma Imbalance and stock volatility. Figure 6 (bottom panel) shows a scatter plot of the relation between the t-statistics and the logarithm of the open interest of options written on the corresponding security. Consistent with the main economic argument, the plot shows that a large open interest correlates with a large negative β_0.

Insert Figure 6

Next, we investigate whether there are differences in the impact of the gamma imbalance for options written on indices versus individual stocks. We expect the effect to be statistically and economically stronger for index options for both statistical and economic reasons. Indeed, the open interest on index options is significantly larger than that of single-name options and the aggregate gamma imbalance measure is likely to be a more precise proxy of the real gamma imbalance of dealers. Results are summarized in Table III and show that the coefficient on Γ^{IB} is negative and significant in all specifications. As expected, the slope coefficient is significantly larger for index options than for single-name options, ranging from -14.8 (after controlling for both asset and month fixed effects) to -5.5 (after controlling only for asset fixed effects). Moreover, the statistical significance is still high, and the t-statistics ranges between -4.4 (after controlling for lagged implied volatility and abnormal volume but no fixed effects) to -3.3 (after controlling for lagged implied volatility, abnormal volume, and both asset and month fixed effects). This is rather impressive since the sample is about 50 times smaller than the panel of individual stocks. The economic magnitude of the effect is even more striking, as one standard deviation increase in the gamma imbalance is associated to a decrease in absolute return of the underlying index of more than 20 basis points (about 20% of a standard deviation)

Insert Table III

Finally, since return volatility is known to be a highly persistent process, we re-run the panel regressions adding alternative variables to control for persistent components in the volatility process. Table IV summarizes the results. We find that in all these alternative specifications the coefficient on Γ^{IB} continues to be highly significant and roughly of the same (negative) magnitude.

B. Heterogeneity in Illiquidity

Empirical findings supporting $H0_1$ might suggest the existence of frictions in the derivative and underlying asset market. A potential channel for such a friction relies on option dealers generating price impact when rebalancing their delta-hedging positions. A testable implication of such a theory is that the effect of gamma imbalance on stock returns is more pronounced for less-liquid stocks, for which the price impact arising from option dealers trades should be larger in magnitude. Hence, to study in further detail the nature of these potential frictions, we investigate whether the effect of gamma imbalance is stronger in less liquid assets and in less liquid time periods:

$H0_2$: *Is the impact of Gamma Imbalance stronger during time periods and for stocks that are more illiquid?*

To test the second hypothesis we compute the Amihud illiquidity ratio for each underlying asset at the daily frequency, using a rolling window of two years ending one month before the day of interest. We then compute the median value of each asset across the entire sample period and use the resulting asset-level illiquidity measure to split our universe into two equally-sized groups.

We run a difference-in-differences regression of absolute returns on $I_\Gamma \times I_\mathcal{L}$, defined as the interaction of the dummy indicating the group of the most illiquid asset $I_\mathcal{L}$ and the dummy indicating days in which the gamma imbalance Γ^{IB} on the asset is negative I_Γ.

Table V summarizes the results and unveils a positive and significant coefficient on the interaction term across all specifications. These results imply that the increase in absolute returns during negative gamma days is significantly larger for the most illiquid assets. This finding supports our theory on the channel supporting the empirical relationship between gamma imbalance and stock return dynamics, that is, intraday momentum and reversal are induced by the price impact of option dealers rebalancing their delta-hedging positions.

Insert Table V

More generally, this result can be interpreted as evidence of the importance of economic frictions for short-term asset price dynamics. It relates to a growing literature that argues the existence of a link between asset's market liquidity and market fragility. Brunnermeir and Pedersen (2009) argue that market liquidity and funding liquidity are mutually reinforcing, leading to excess volatility and liquidity spirals. Our results suggest that this may occur even in the absence of (in addition to) margin constraints through the interaction of market illiquidity and the risk-limiting behavior of risk-averse market makers.[10]

C. *Intraday Autocorrelation: The Role of Gamma Imbalance for Momentum and Mean-Reversion.*

An additional implication of the channel underlying hypothesis $H0_1$ directly relates to the existence of intraday market momentum or mean-reversion depending on the gamma imbalance. Indeed, dealers delta-hedging strategies require selling more of the underlying asset following an increase in the underlying price if their Gamma Imbalance is positive. The opposite effect occurs when the Gamma Imbalance is negative. Thus, dealers order-flow should act as a contrarian force when the Gamma Imbalance is positive giving rise to intraday mean-reversion, namely negative intraday autocorrelation. The opposite effect should

[10] See also Adrian, Moench, and Shin (2014), Gromb and Vayanos (2002), Garleanu and Pedersen (2001), Adrian, Etula, and Muir (2014), Adrian and Shin (2010), Adrian and Boyarchenko (2015), Adrian, Colla, and Shin (2012).

emerge when the Gamma Imbalance is negative when one should expect intra-day momentum, i.e. positive intraday autocorrelation.

$H0_3$: *Equity serial autocorrelation is positive (negative) when dealers are short (long) gamma.*

To test this hypothesis, we use data from the TAQ database to compute the autocorrelation of returns for the equity stocks in our universe, at different frequencies. For each day-asset pair (t, j) we estimate the sample autocorrelation coefficient $\rho^h_{j,t}$ of h-minute non-overlapping returns with h = 5, 10, 20, 30, and 60. Then, we run the following panel regressions, for each frequency h

$$\rho^h_{j,t} = \alpha^h + \beta^h_0 \, \Gamma^{IB}_{j,t-1} + \beta^h_1 E_{t-1}[\rho^h_{j,t}] + FE_{j,t} + \varepsilon_{j,t}, \tag{6}$$

where $E_{t-1}[\rho^h_{j,t}]$ is the rolling-average autocorrelation coefficient of h-minute returns of stock j, estimated over a 2-year window ending one month before day t. We test the null hypothesis that $\beta^h_0 < 0$ at the intraday level. Results are reported in Table VI and show that the slope coefficients on $\Gamma^{IB}_{j,t-1}$ are indeed negative and significant for almost all the frequencies, supporting the null hypothesis $H0_3$. It is interesting to observe that the magnitude and significance of the coefficients are increasing with the frequency and they peak at $h = 30$. This is consistent with dealers adjusting their delta-hedged portfolios at an average frequency of about 30 minutes. The economic magnitude of the result, however, is limited. A unitary standard deviation increase in the gamma imbalance proxy is associated to a decrease in autocorrelation of about 1% of a standard deviation.

[Insert Table VI]

IV. Market Fragility and Flash Crashes

A growing literature studies market events that are characterized by a very rapid, deep, and transitory fall in security prices. These events are often referred as “flash crashes” and have been conjectured to originate from the interaction of automated high-frequency trading, spoofing and price manipulation, the fragility of derivative markets, and other automatic hedging strategies.

The literature on flash crashes gained particular vigour in the aftermath of the famous May 6, 2010 U.S. market crash that saw about a trillion dollar of market capitalization being wiped out in 36 minutes. The specific dynamics of how the crash occurred strike a familiar note. When U.S. stock markets opened, the Dow started trending down due to worries about the debt crisis in Greece. By 2:42 p.m. the Dow was down more than 300 points. After 2:42 p.m., the drop accelerated and the index fell an additional 600 points in 5 minutes accumulating a total loss of nearly 1,000 points by 2:47 p.m. Although the sharp drop was later reversed, the economic magnitude of the intraday volatility was remarkable given the absence of significant economic news.

The May 6 2010 Flash Crash was not an isolated event. Gao and Mizrach (2016) argue that flash crashes have occurred in almost every year in the period 1993-2011 that they study. The most notorious of these events have directly involved the index. However, individual stocks have also been exposed to similar events. Our hypothesis is that flash crashes are more likely to occur when markets are fragile. In our context, an important factor contributing to market fragility is the aggregate gamma imbalance.

To investigate whether gamma imbalance is potentially related to these events, first we construct a temporary proxy of intraday volatility that is particularly sensitive to intraday jumps. We use as dependent variable the daily spread, i.e. the difference between the intraday High and Low relative to the mid-point price (e.g., $Spread = \frac{H-L}{H+L}$). Then, for each day-asset pair (t, j) we run a regression of the $Spread(t)$ onto the lagged level of gamma imbalance $\Gamma^{IB}_{j,t-1}$ after controlling for fixed effects and lagged implied volatility:

$$Spread(j,t) = \alpha + \beta_0 \Gamma^{IB}_{j,t-1} + \beta_1 IVOL_{j,t-1} + FE_{j,t} + e_{j,t}. \tag{7}$$

We run four specifications. Specification (1) excludes fixed effects, while specifications (2) and (3) include controls for asset and month fixed effects. Specification (4) include both controls. The estimates are summarized in Table VII. Two main results emerge. First, the slope coefficient of a regression of *Spread* on lagged Γ^{IB} is negative, suggesting that large negative values of the gamma imbalance correlates with increases in intraday price jumps. In all specification, the t-statistics are based on robust standard errors double-clustered at the asset and month level and indicate that lagged Γ^{IB} is statistically significant at the 1% confidence level.

Second, the economic significance of the effect on the intraday *Spread* is important. Indeed, in specification (4), which includes both asset and month fixed effects, the slope coefficient for $\Gamma^{IB}_{j,t-1}$ is -12.3. This estimate implies that a one standard deviation decrease in gamma imbalance is associated to an increase of roughly 15 basis points in the daily spread. This is consistent with the economic hypothesis that periods of large negative gamma imbalances are more fragile. Jumps and flash crashes are more frequent, and volatility is higher.

Insert Table VII

A. *Flash Crash Identification and Tests*

The previous results use the whole data and the dependent variable include both large and small events. In the following, we directly identify the subset of flash crashes in the data set to study the link with the gamma imbalance. Then, we test the hypothesis that the probability of a stock experiencing a flash crash increases when the stock is subject to negative gamma imbalance and decreases when its gamma imbalance is positive. Moreover, we investigate whether, conditional on a flash crash happening, the price drop is more pronounced for stocks exposed to negative gamma imbalances. These tests use both time-series and cross-sectional information.

We use the "drift burst" detection methodology proposed by Christensen, Oomen, and Renò (2018). This approach aims to identify "sudden and extreme movement in price which occurs in relatively short time and then reverts to the initial level" as opposed to simply a period of extreme volatility or a price jump. To illustrate their methodology, let us assume that p_t be the log-price price of the asset follows

$$dp_t = \mu_t dt + \sigma_t dW_t. \tag{8}$$

Moreover, assuming that the price process is observable over $[0, T]$ at some time points $0 = t_0 < t_1 < ... < t_n = T$, returns are defined as

$$r_{t_i} = p_{t_i} - p_{t_{i-1}}, \qquad i = 1, ..., n.$$

Christensen, Oomen, and Renò (2018) propose to compute of the current velocity of the market as defined by the test statistics

$$T_t^n = \sqrt{\frac{h_n}{K_2}} \frac{\hat{\mu}_t^n}{\hat{\sigma}_t^n},$$

where $\hat{\mu}_t^n$ and $\hat{\sigma}_t^n$ are non-parametric estimators of the drift and diffusion of dp_t. These can be obtained using a kernel-weighted average $K(x)$ of observations in the vicinity of t. The bandwidth h_n determines how fast observations are downweighted when they occur farther away from t:

$$\hat{\mu}_t^n = \frac{1}{h_n}\sum_{i=1}^{n} K\left(\frac{t_{i-1}-t}{h_n}\right) r_{t_i}, \qquad and \qquad \hat{\sigma}_t^n = \sqrt{\frac{1}{h_n}\sum_{i=1}^{n} K\left(\frac{t_{i-1}-t}{h_n}\right) r_{t_i}^2}$$

with $K_2 \equiv \int K(x)^2 dx$.

If the price is moving fast relative to the volatility T_t^n is large. Christensen, Oomen, and Renò (2018) show that in absence of a "drift burst" and under the assumption that dp_t follows the process (8), the test statistics T_t^n converges asymptotically to:

$$T_t^n \to \begin{cases} N(0,1) & \text{If no drift burst} \\ \infty & \text{If drift burst} \end{cases},$$

As $n \to \infty, h_n \to 0$, and $nh_n \to \infty$.

We employ a 30-minute bandwidth for the mean and a 60-minute bandwidth for the volatility. This means that, by construction, we are interested in flash crashes which develop on a time span of roughly 30 minutes. We also set $K(x) = exp(-|x|)1_{x\leq 0}$.

We compute the above test statistic every second during the course of a trading session. A significant negative value of T_t^n reveals a large and fast drop in price. We define the "beginning" of the crash t_0 the first time $T_{t_0}^n$ crosses -3 – corresponding to the 1000th quantile of the distribution – from above. The "peak" of the crash, on the other hand, is defined as time t_1 such that $T_{t_1}^n$ reaches its minimum level. Accordingly, $\tau = t_1 - t_0$ is defined as the duration of the crash. However, an additional condition needs to be satisfied in order for an event to be defined as a "Flash Crash", namely that the stock experiences a cumulative price drop of at least 1% during the 30-minutes interval.

We apply the described methodology to a subset of 77 stocks from our original sample, using minute-by-minute price changes constructed from the Trades and Quotes database covering the period from 1997 to 2015. The procedure finds 672 distinct flash-crash events, which are mostly evenly distributed across time and affect 73 stocks (about 1 crash every two years for each stock, on average). These events represent significant price drops materializing in a relatively short time. The average drawdown during the 30-minute window is -3.71% (median -1.17%).

To test our first hypothesis, we construct a dummy variable indicating the presence of a flash-crash in a given stock-day and we run regressions on the panel of daily gamma imbalance estimates for the 73 stocks that experienced a flash crash:

$$I_{Crash}(j,t) = \alpha + \beta_0 \Gamma^{IB}_{j,t-1} + FE_{j,t} + \varepsilon_{j,t},$$

where $I_{Crash}(j,t)$ is the dummy indicating a flash crash for stock j on day t and $\Gamma^{IB}_{j,t-1}$ is the level of gamma imbalance for that stock measured at the end of the previous day. The results are summarized in Table VIII and show that the estimate for β_0 is significantly negative with a t-stat equal to -5.99. The relationship between gamma imbalance and flash crashes is highly significant not only from a statistical perspective, but also economically. Indeed, these estimates imply that a negative shift of one standard deviation in the stock gamma imbalance doubles the probability of flash crash for that stock. Furthermore, the probability of a flash crash increases by 25% when the gamma imbalance is negative. It is important to stress that correlation does not imply causation and we certainly do not interpret these results as evidence that delta-hedging of option dealers can cause flash crashes; rather, gamma imbalance can work as a friction that may exacerbate price movements due to fundamental news or liquidity shocks.

Next, to directly assess the extent to which gamma imbalance is related to the severity of flash crashes. Conditional on a "flash crash" identified by the previous methodology, we create two subsamples depending on the sign of the gamma imbalance before the event. Then, we compute the two event-time average price pattern in the two sub-samples. Figure 8 summarizes the results. We find the flash crashes are more severe when the stock is exposed to negative gamma imbalance. More precisely, the magnitude of the average drawdown is -5% when the gamma imbalance is negative, compared to -3% when the gamma imbalance is positive. The difference is statistically significant at the 1% level. From an economic standpoint, given that the stocks we are considering are large caps, this implies an average loss of value of about one billion US dollars, for each flash crash.

Insert Table VII and Figure 8

B. The May 6th 2010 Flash-Crash

The May 6, 2010 Flash Crash generated significant interest both among regulators and academics. The reasons are simple. In a matter of only 36 minutes the U.S. stock first erased more than a trillion dollar of market capitalization. From a theoretical perspective this is hard to be reconciled with a frictionless market given that, by the time the market closed, the Dow index had already recovered from the initial losses almost completely and that no significant economic news marked the day. Indeed, both the press and regulators immediately announced their intention to investigate the event and conjectured that the cause of the event was the existence of market manipulation.

As we revisit in greater details the unfolding of these events, one fact appears particularly striking to us as it rhymes with the narrative played by the role of gamma imbalance as discussed earlier. Because of worries about the debt crisis in Greece, at the opening of the stock market and well before the occurrence of the "flash crash", the Dow had already started falling. Indeed, by 2:32pm the Dow had already dropped about 2%. After 2:32pm, the equity market entered the "flash crash" climax and accelerated its fall by dropping an additional 600 points in just 5 minutes. By 2:47pm the index had accumulated a loss of nearly 1,000 points. Interestingly, this price fluctuation occurred in the absence of any significant economic news and by 3:00pm the Dow had recovered most of its losses.

The Commodity Futures Trading Commission (CFTC) conducted an investigation and eventually the U.S. Department of Justice pressed charges on "22 criminal counts, including fraud and market manipulation" against Navinder Singh Sarao, a British financial trader. Later, he was found guilty of using spoofing algorithms: just prior to the flash crash, he placed orders for about $200 million worth of bets" on the E-mini S&P 500 stock index futures contracts that the market would fall, which were eventually replaced or modified 19,000 times before they were canceled. The CFTC argued that Sarao actions were "significantly responsible for the order imbalances in the derivatives market which affected stock markets and exacerbated the flash crash."[11] The CFTC argument is directly related to the core question of our paper. Indeed, while price manipulation is difficult in normal markets, spoofing can become effective in the presence of large order imbalances in the derivative market.

According to our theory and the results from the previous section, gamma imbalance could have played a role during the flash crash. To shed light on this issue, we investigate the level of market-wide Gamma Imbalance during the days preceding the flash crash, defined as the daily cross-sectional average of stock-level gamma imbalance. The first Panel of Figure 9 shows that the market gamma imbalance drops and moves into negative territory in the days leading to the crash, and it is recorded at -1.7 standard deviation units during May 6. This corresponds to a negative market-wide dollar imbalance of roughly -431 million U.S. dollars. The second Panel of Figure 9 shows that this value is statistically significantly different from the average and sits in the left tail of the distribution. The drop in gamma imbalance occurred reasonably quickly during the week leading to the flash crash. Indeed, one week before the flash crash the total gamma imbalance was positive at about 383 millions.

Even though we cannot claim a causal relationship between the negative level of gamma imbalance and the flash crash, this evidence suggests that the delta-hedging activity of option dealers can contribute to market fragility and increasing the magnitude of price drops. Given

[11] See, Brush, Silla, Tom Schoenberg, and Suzi Ring (April 22, 2015), "*How a Mystery Trader with an Algorithm May Have Caused the Flash Crash*", Bloomberg News.

the fast price drop and the negative gamma exposure, dealers likely sold a significant amount of index shares to keep their positions hedged. This could have consumed additional liquidity and increased the downward pressure on the price, thus enhancing the drop.

Insert Figure 9

C. Introduction of Circuit Breakers

The events that followed the 2010 flash crash provide an interesting natural counterfactual experiment. Indeed, following the Flash Crash on May 6 2010, the SEC announced that new trading curbs – also known as circuit breakers – would be tested starting from June of that year. This new rule dictates a complete trading suspension for five minutes on any stock whose price moves more than 10% within a 5-minute window. After a short test period, in which the rule was imposed on a subset of pilot stocks, by the end of June circuit breakers were installed for all S&P 500 stocks.

These new measures have been controversial as many claimed that they interfere with the price discovery process. Subrahmanyam (1994, 1995) argues that circuit breakers [12] may have the perverse effect of increasing price variability by forcing agents to advance their trades before trading is halted, a so-called "magnet effect." Moreover, he shows that in a multiple markets setting if the circuit breaker is triggered in the more liquid market, there is the risk of a negative externality in the less liquid market which would experiencre higher volatility and lower liquidity. Subrahmanyam (1997)[13] also suggests that informed traders may reduce their trading in anticipation of a trading halt resulting in higher trading costs for small investors. On the other hand, Greenwald and Stein (1991), argue that circuit breakers can be useful in reducing the uncertainty about execution price. When current prices fail

[12] Subrahmanyam, Avanidhar (1995) On rules versus discretion in procedures to halt trade. Journal of Economics and Business 47(1), 1–16.

Subrahmanyam, Avanidhar (1994) Circuit breakers and market volatility: A theoretical perspective. Journal of Finance 49(1), 237–254.

[13] Subrahmanyam, Avanidhar (1997) The ex ante effects of trade halting rules on informed trading strategies and market liquidity. Review of Financial Economics 6(1), 1–14.

to accurately represent information, a circuit breaker may encourage buyers and sellers to submit orders. Indeed, when a very large volume shock hits a market, transactional risk rises sharply.

Circuit breaker rules were already in effect at the time of the May 6 2010 Flash Crash but the rules applied market-wide and they did not get triggered.[14]

Given the debate on the role of circuit breakers to help avoiding flash crashes, we investigate whether the occurrence and severity of flash crashes change after the introduction of the new circuit breaker rules in June 2010. We partition our sample in two non-overlapping periods, before and after June 2010. Then, we identify all flash crashes that occurred before and after June 2010. Consistent with Greenwald and Stein (1991) we find that the frequency of fire-sales events decreases significantly in the most recent period, from an average of more that 45 to roughly 4.5 per year. The difference is both statistically and economically significant. Moreover, the average drawdown displays a four-fold decline from -4% to -1% after the introduction of trading curbs. Interestingly, as shown in Figure 10, the empirical relationship between the ex-ante gamma imbalances and the magnitude of the crashes is significantly reduced after June 2010.

V. Information-Based Explanations

An important literature studies the role played by asymmetric information and adverse selection in the short-term dynamics of asset prices. Several information-based models (e.g., Glosten and Milgrom (1985), Easley, O'Hara, and Srinivas (1998)) show that when market makers face the risk of privately informed traders, stock returns may display predictability and serial autocorrelation.

In this section, we compare the relative importance of the gamma-imbalance factor versus adverse selection in driving return volatility and autocorrelations.

[14]The first circuit breaker was implemented in October 1988. The original circuit breakers were triggered when the DJIA fell by 250 points and trading would be halt for one hour.

A. Order-Flow Toxicity

Easley, Lopez de Prado, and O'Hara (2012) propose a new proxy of adverse selection that is designed to capture the probability that market dealers face privately informed traders. The idea is that risk-averse market makers withdraw from the market in the presence of informed flow and, as a consequence, liquidity dries out when it is needed the most. They call this measure *order-flow toxicity* (VPIN) and is based on volume, as opposed to time, increments. They find that it helps to explain market volatility and can be used as a predictor of flash crashes and sharp price movements.

Under this view, the baseline source of the price movement is fundamental information, but short-term volatility is amplified by a temporary lack of liquidity. This resonates with our theory, in which initial price movements, motivated by fundamental news, may induce (conditional on negative gamma imbalance) a delta-hedging portfolio rebalancing, thus leading to short-momentum and excess volatility.

To test the relationship and the relative importance of these two channels, we construct the VPIN measure for all assets in our universe and re-run the panel regressions presented in Section A and Section C after controlling for the order-flow toxicity.

The VPIN is constructed following the definition of O'Hara et al (2012) and using intraday data from the TAQ dataset. For asset j on day t we start from the time-series of minute-by-minute prices p_t and construct the standardized first-difference $\Delta p_t = (p_t - p_{t-1})/\sigma$ where σ is the sample standard deviation of price changes estimated over the previous 12 months. We then apply the standard normal CDF to obtain a proxy for the fraction of buy-initiated orders $b_t = Z(\Delta p_t)$. Next, we average b_t at the trade-time frequency τ (which increases by one unit every 10 transactions) to obtain B_τ and its complementary $S_\tau = 1 - B_\tau$. Finally, we compute the average order flow toxicity (VPIN) for each trading day as volume-weighted average of the trade-time (day) imbalances $\mid B_\tau - S_\tau \mid$.

Results are summarized in Table IX and X. Three key results emerge. First, even after controlling for VPIN the slope coefficients on the gamma imbalance factor Γ^{IB} are significantly negative in all specifications. This suggests that even in the presence of a potential information-based channel, gamma imbalance is an important separate driver of short term volatility.

Table IX and X

Second, more negative levels of Γ^{IB} correlate with greater stock market volatility. The slope coefficient of Γ^{IB} is robust to controlling for asset and month fixed effects, as separately reported in columns (2) and (4). We also find that the slope coefficient is more negative when using as dependent variable the intraday difference between the maximum and minimum price (e.g. "*Spread*") than when using the absolute open-to-close return (e.g. "$|Return|$"). In the first case the slope coefficient is -14.5 with a t-statistics of -4.5; in the second case the slope coefficient is -8.9 with a t-statistics of -4.9. Since flash crashes are very high frequency events and often mean-reverts intradaily, the "*Spread*" variable is a better proxy to capture these events. Indeed, our results suggests that large and negative gamma imbalances correlated with these extreme intraday volatility events.

Finally, when we analyze the effect on the serial autocorrelation of asset returns for different time horizons, ranging from 5 to 60 minutes, we find that the slope coefficient on Γ^{IB} is negative, consistent with the null hypothesis. Table X summarizes the results. After controlling for $VPIN$, the slope coefficient of Γ^{IB} is no longer significant at the 5 minutes horizon; however, it continues to be strongly significant at horizons above 10 minutes. The stronger economic and statistical significance occur when the autocorrelation coefficient is obtained from data sampled at 30 minutes intervals. In this case, the slope coefficient is -0.33 with a t-statistics equal to -3.75. This result is consistent with those documented in Table VI.

It should be noticed that while the coefficient on the VPIN measure is highly significant its sign is negative in all specifications. At first, this may be surprising since it indicates that, after controlling for Γ^{IB}, larger order flow toxicity VPIN is associated to a lower absolute daily returns and daily spread. However, VPIN may be a more effective measure to explain excess volatility and flash crash events occurring at much higher frequencies.

B. Fundamental News

Another potential alternative information-based explanation of short-term dynamics may rely on the idea that participants in the options market may have superior information on the fundamentals of the underlying assets. Black (1975) argues that informed investors might choose to trade derivatives because of the higher leverage offered by such instruments. Easley, O'Hara, and Srinivas (1998) formally study a model that allows for endogenous participation of informed traders in the options market in equilibrium. Pan and Poteshman (2006) find supporting evidence for this hypothesis and show a link between trading in option markets and predictability in stock returns. One may argue that our time-series and cross-sectional pattern could be explained by assuming that then investors take large positions on put options when they expect large negative movements in the stock price.

We already partially address this concern by controlling for ex-ante implied volatility. However, to directly address this alternative explanation, we collect fundamental news from RavenPack News Analytics. In this dataset a sentiment score is assigned to each news article, using both RavenPack proprietary techniques and open source natural language processing (NLP) algorithms, as described in Manning et al (2014). We use the sentiment score to classify strong positive news and strong negative news. Finally, we re-run all our regressions by excluding from the sample those day-asset observation (t, j) in which at least one strongly negative or strongly positive news was released on asset j in a window of two weeks around day t. This filter reduces our sample by about 20% of the data.

Table XI summarizes the results and shows that, even after filtering out the most significant fundamental news, the slope coefficient on the gamma imbalance factor Γ^{IB} continues to have the same magnitude and level of statistical significance.

Table XI

C. Leveraged ETF and ETN

A recent literature suggests that a significant increase in adoption of leveraged ETFs has contributed to an increase in market volatility. Since leveraged ETFs provide a daily constant multiple of the underlying index, following any market move the ETF needs to daily rebalance its positions to maintain a constant leverage ratio. For example, a $4\times$ fund is designed to offer four times the daily index return, while a $-4\times$ fund is designed to generate four times the opposite return of the index. These funds achieve their target exposure to the index with a combination of swaps and futures. While leveraged ETFs are synthetic instruments, their swap counterparties have to rebalance to maintain the required exposure, regardless of market conditions. If L is the leverage multiple, the exposure E_t^* needs to be $L \times NAV_t$. Thus, if the return on the index at time $t+1$ is r_{t+1}, the E_{t+1}^* is equal to $L[NAV_t \times (1 + r_{t+1} \times L)]$. This implies that the required rebalancing at time $t+1$ is

$$E_{t+1}^* - E_t^* = L^2(r_{t+1} \times NAV_t)$$

When the market increases (e.g. $r_{t+1} > 0$), both levered long and short funds need to increase their notional principal, thus adding to any existing buying pressure that contributed to the initial price increase. When the market decreases (e.g. $r_{t+1} < 0$), both long and short funds are net sellers in the market. This implies that by the end of the trading day, the demand and supply of leveraged ETFs create additional demand or selling pressure in the same direction as the market move. This argument predicts that a greater adoption of leveraged ETFs and ETNs should correlate with greater price momentum in the hours before market close (Cheng and Madhavan (2009)).

Ben-David, Franzoni, and Moussawi (2018) find evidence that stocks owned by ETFs exhibit significantly higher overall daily volatility. They estimate that an increase of one standard deviation in ETF ownership is associated with an increase of 16% in daily stock volatility. They argue, similarly to us, about the existence of feedback effects due to the hedging activity of arbitrageurs active in the ETF market. Indeed, they find that the effects are stronger for stocks with lower bid-ask spread and lending fees. However, they do not study intraday patterns. Shum, Hejazi, Haryanto, and Rodier (2015) document the existence of a link between the rebalancing activity of leveraged ETFs and market volatility before market closure. Although related to our findings, these papers do not provide predictions about the coexistence of both intraday momentum and reversal. While the aggregate value of leveraged ETFs have increased almost monotonically in the last ten years, the frequency of positive price momentum has not monotonically increased over time. Instead, we observe both intraday price momentum and reversals. We test whether unusual intraday high (low) autocorrelation or/and volatility correlates with unusually high (low) aggregate notional of leveraged ETF and comfortably reject this hypothesis.

Finally, using cross-sectional data on individual stocks we document significant cross-sectional dispersion in price momentum and reversal which cannot be explained by the rebalancing activity of leveraged ETFs. At the same time, this cross-sectional dispersion correlates with cross-sectional differences in gamma imbalance.

VI. Conclusion

This paper provides evidence supporting the view that trading in derivatives may affect the price process of the underlying assets, contributing to intraday stock volatility and autocorrelation by increasing (dampening) the initial effect of news on market fundamentals. Indeed, we unveil a strong negative relationship between the gamma imbalance of option dealers and measures of intraday volatility. This phenomenon is relevant both for market participants and regulatory concerned about market fragility. This effect might become even more important in the future given the increasing incentived for institutions to use derivative products with non-neutral gamma to reduce their required regulatory risk capital.

References

Adrian, Tobias, and Nina Boyarchenko, 2015, Intermediary leverage cycles and financial stability, *Staff Reports, Federal Reserve Bank of New York* 567.

Adrian, Tobias, Paolo Colla, and Hyun Song Shin, 2012, Which financial frictions? parsing the evidence from the financial crisis of 2007-2009, *NBER Working Papers, National Bureau of Economic Research, Inc* 18335.

Adrian, Tobias, Erkko Etula, and Tyler Muir, 2014, Financial intermediaries and the cross-section of asset returns, *Journal of Finance* 69, 2557–2596.

Adrian, Tobias, Emanuel Moench, and Hyun Song Shin, 2014, Dynamic leverage asset pricing, *Staff Reports, Federal Reserve Bank of New York* 625.

Adrian, Tobias, and Hyun Song Shin Shin, 2010, Liquidity and leverage, *Journal of Financial Intermediationx* 19, 418–437.

Anton, Miguel, and Christopher Polk, 2014, Connected stocks, *Journal of Finance* 69, 1099–1127.

Barberis, and Xiong, 2009, What drives the disposition effect? an analysis of a long-standing preference-based explanation, *The Journal of Finance* LXIV.

Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi, 2018, Do etfs increase volatility?, *Journal of Finance* 73, 2.

Ben David, Itzhak, and David Hirshleifer, 2012, Are investors really reluctant to realize their losses? trading responses to past returns and the disposition effect, *The Review of Financial Studies* 25, 2485–2532.

Black, Fisher, 1975, Fact and fantasy in the use of options, *Financial Analysts Journal* 31, 36–41, 61–72.

Brogaard, J, 2010, High frequency trading and its impact on market quality, Working Paper, Northwestern University.

Brunnermeir, Markus, and Lasse Pedersen, 2009, Market liquidity and funding liquidity, *The Review of Financial Studies* 22, 2201–2238.

Chang, Yen-Cheng, Harrison Hong, and Inessa Liskovich, 2013, Regression discontinuity and the price effects of stock market indexing, *NBER working paper 19290*.

Christensen, Kim, Roel CA Oomen, and Roberto Renò, 2018, The drift burst hypothesis, *Available at SSRN 2842535*.

Daniel, Kent, David Hirshleifer, and A Subrahmanyam, 1998, A theory of overconfidence, self-attribution, and security market under and over-reactions, *Journal of Finance* 53, 1839–85.

Easley, D., M. M. Lopez de Prado, and M. O'Hara, 2012, Flow toxicity and liquidity in a high-frequency world, *The Review of Financial Studie* 25, 1457–1493.

Easley, D., M. O'Hara, and P. Srinivas, 1998, Option volume and stock prices: Evidence on where informed traders trade, *Journal of Finance* 53, 431–465.

Engle, Robert, and Jow Lange, 2001, Predicting vnet: A model of the dynamics of market depth, *Journal of Financial Markets* 4, 113–142.

Frazzini, Andrea, 2006, The disposition effect and underreaction to news, *Journal of Finance* 61, 2017–2046.

Gao, Cheng, and Bruce Mizrach, 2016, Market quality breakdowns in equities, *Journal of Financial Markets* 28, 1–23.

Garleanu, Nicolae, and Lasse Pedersen, 2001, Margin-based asset pricing and deviations from the law of one price, *Review of Financial Studies* 24, 1980–2022.

Glosten, L., and P. Milgrom, 1985, Bid, ask, and transaction prices in a specialist market with heterogenously informed traders, *Journal of Financial Economics* 14, 71–100.

Gorton, G., F. Hayashi, and G. Rouwenhorst, 2013, The fundamentals of commodity futures returns, *Review of Finance* 17, 35–105.

Greenwood, Robin, and David Thesmar, 2011, Stock price fragility, *Journal of Financial Economics* 102, 471–490.

Gromb, Denis, and Dimitri Vayanos, 2002, Equilibrium and welfare in markets with financially constrained arbitrageurs, *Journal of Financial Economics* 66, 361–407.

Hasbrouck, Joel, and Gideon Saar, 2013, Low-latency trading, *Journal of Financial Markets* 16, 646–679.

Hendershott, Terrence, Charles Jones, and Albert Menkveld, 2011, Does algorithmic trading improve liquidity?, *The Journal of Finance* LXVI, 1–33.

Hendershott, Terrence, and Ryan Riordan, 2013, Algorithmic trading and the market for liquidity, *Journal of Financial and Quantitative Analysis* 48, 1001–1024.

Hendershott, Terrence, and Mark Seasholes, 2007, Market maker inventories and stock prices, *American Economic Review* 97, 210–214.

Holmstrom, Bengt, and Jean Tirole, 1997, Financial intermediation, loanable funds, and the real sector, *The Quarterly Journal of Economics* 112, 663–691.

Jotikasthira, Chotibhak, Christian Lundblad, and Tarun Ramadorai, 2012, Asset fire sales and purchases and the international transmission of financial shocks, *Journal of Finance* p. 20152050.

Kirilenko, Andrei, Albert Kyle, Merhdad Samadi, and Tugkan Tuzun, 2017, The flash crash: Highfrequency trading in an electronic market, *The Journal of Finance* 72, 967–998.

Lou, Dong, and Christopher Polk, 2013, Comomentum: Inferring arbitrage activity from return correlations, *LSE Working paper*.

Luo, Jiang, Avanidhar Subrahmanyam, and Sheridan Titman, 2018, Stock price dynamics with overconfident investors, .

Pan, Jun, and A. M. Poteshman, 2006, The information in option volume for future stock prices, *The Review of Financial Studies* 19, 871–908.

Shefrin, Hersh, and Meir Statman, 1985, The disposition to sell winners too early and ride losers too long, *Journal of Finance* 40, 777–790.

Shum, Pauline, Walid Hejazi, Edgar Haryanto, and Arthur Rodier, 2015, Intraday share price volatility and leveraged etf rebalancing, *Review of Finance* 20, 2379–2409.

Vayanos, Dimitri, and Paul Woolley, 2013, An institutional theory of momentum and reversal, *The Review of Financial Studies* 26, 1087–1145.

VII. Tables

Table I. Option Data - Summary Statistics

This table presents summary statistics of the main variables employed in the analysis. We consider options written on a universe of 280 assets over the period from January 1996 to December 2017. The gamma imbalance proxy Γ^{IB} is defined in Section II.B and expressed in percentage units. The open interest and the trading volume of the underlying assets are expressed in million dollars. Absolute values of returns and daily spreads are in basis points.

All Options	Count	Mean	Std	10%	25%	50%	75%	90%
Dealers Gamma Imbalance	1,036,850	0.157	0.469	-0.525	-0.200	0.214	0.545	0.741
Options Open Interest	1,036,850	3.733	35.720	0.017	0.113	0.484	1.479	3.665
Options Implied Volatility	1,036,850	0.500	0.241	0.274	0.339	0.441	0.591	0.809
Underlying Volume	1,036,850	26.507	260.876	0.003	0.028	0.199	0.906	3.230
Underlying Absolute Return	1,036,850	170.665	212.345	17.584	46.825	109.091	218.567	382.961
Underlying Daily Spread	1,036,850	295.268	253.720	98.855	146.149	227.685	360.144	555.630

Index Options	Count	Mean	Std	10%	25%	50%	75%	90%
Dealers Gamma Imbalance	21,682	0.105	0.421	-0.528	-0.200	0.175	0.447	0.613
Options Open Interest	21,682	101.888	219.035	0.975	2.188	11.058	72.061	486.642
Options Implied Volatility	21,682	0.305	0.097	0.208	0.238	0.283	0.346	0.438
Underlying Volume	21,682	1145.816	1395.593	15.741	57.065	459.680	1861.287	3221.273
Underlying Absolute Return	21,682	95.057	99.731	10.421	27.405	64.040	128.216	216.805
Underlying Daily Spread	21,682	153.221	110.891	54.949	80.660	122.925	189.118	283.163

Equity Options	Count	Mean	Std	10%	25%	50%	75%	90%
Dealers Gamma Imbalance	1,015,168	0.158	0.470	-0.525	-0.200	0.215	0.547	0.743
Options Open Interest	1,015,168	1.637	8.267	0.016	0.108	0.463	1.394	3.302
Options Implied Volatility	1,015,168	0.505	0.241	0.278	0.343	0.445	0.595	0.814
Underlying Volume	1,015,168	2.601	24.152	0.003	0.027	0.187	0.821	2.587
Underlying Absolute Return	1,015,168	172.280	213.813	17.873	47.465	110.345	220.713	386.172
Underlying Daily Spread	1,015,168	298.302	255.041	100.908	148.415	230.403	363.636	560.012

Table II. Gamma Imbalance and Volatility

The dependent variable of this panel regression is the absolute value of day-t return of asset j, expressed in basis points. The main explanatory variable is Γ^{IB}, i.e. the *gamma imbalance* of dealers on put and call options written on j as of the close of day $t-1$, expressed as a fraction of the average daily volume of the underlying. We control for the option-implied volatility as of $t-1$ and we saturate the model with asset and month fixed effects. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and month level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	|Return|	|Return|	|Return|	|Return|
Gamma Imbalance (lag)	-11.902***	-12.212***	-6.276***	-7.777***
	(-4.542)	(-5.505)	(-3.906)	(-5.325)
Implied Volatility (lag)	391.172***	430.140***	360.607***	331.316***
	(18.849)	(12.784)	(24.034)	(19.255)
Constant	-22.493***			
	(-2.870)			
Observations	1,018,199	1,018,199	1,018,199	1,018,199
R-squared	0.171	0.189	0.220	0.237
Asset Fixed Effects	-	Yes	-	Yes
Month Fixed Effects	-	-	Yes	Yes
Ses Clustered By	Asset-Month	Asset-Month	Asset-Month	Asset-Month

Table III. Index Options

These specifications are the same as the ones presented in Table II, but we restrict to to options on market indices. The universe includes S&P 500, Dow Jones, Nasdaq 100 and Russell 2000. We control for the option-implied volatility as of $t-1$ and we saturate the model with asset and month fixed effects. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and month level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	|Return|	|Return|	|Return|	|Return|
Gamma Imbalance (lag)	-14.79***	-14.78**	-6.62**	-5.49**
	(-4.42)	(-3.75)	(-3.64)	(-3.33)
Implied Volatility (lag)	410.28***	402.19***	330.11***	278.81***
	(9.73)	(7.57)	(9.36)	(6.88)
Abnormal Volume	155.77***	155.98**	142.29**	147.11**
	(6.23)	(5.38)	(5.49)	(5.71)
Constant	-138.77***			
	(-5.15)			
Observations	21,434	21,434	21,434	21,434
R-squared	0.25	0.25	0.34	0.35
Asset Fixed Effects	-	Yes	-	Yes
Month Fixed Effects	-	-	Yes	Yes
Ses Clustered By	Asset-Month	Asset-Month	Asset-Month	Asset-Month

Table IV. Control for Past Volatility

In this specification we add controls for past volatility, to account for the fact that volatility is a highly persistence process. We include lagged values of the absolute value of returns by 1, 2 and 3 business days. We also control for the option-implied volatility as of $t-1$ and we saturate the model with asset and month fixed effects. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and month level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	|Return|	|Return|	|Return|	|Return|
Gamma Imbalance (lag)	-8.393***	-9.090***	-5.185***	-6.620***
	(-5.253)	(-6.387)	(-4.359)	(-5.879)
Implied Volatility (lag)	247.954***	284.794***	264.384***	254.037***
	(20.831)	(15.215)	(23.614)	(20.101)
|Ret| (1 day lag)	0.140***	0.129***	0.107***	0.093***
	(16.208)	(15.015)	(11.525)	(10.372)
|Ret| (2 days lag)	0.122***	0.111***	0.090***	0.077***
	(10.401)	(9.439)	(10.746)	(9.682)
|Ret| (3 days lag)	0.096***	0.085***	0.065***	0.051***
	(18.548)	(16.682)	(11.124)	(10.290)
Constant	-13.037***			
	(-2.703)			
Observations	1,018,199	1,018,199	1,018,199	1,018,199
R-squared	0.219	0.227	0.243	0.253
Asset Fixed Effects	-	Yes	-	Yes
Month Fixed Effects	-	-	Yes	Yes
Ses Clustered By	Asset-Month	Asset-Month	Asset-Month	Asset-Month

Table V. Interaction with Illiquidity

In these specifications we further explore the heterogeneity of our baseline results conditional on the level of liquidity of the underlying assets, proxied by the Amihud illiquidity ratio. We interact the *negative gamma imbalance* dummy with a dummy indicating the most illiquid underlying assets, i.e. those below the median Amihud ratio. For each underlying asset, the illiquidity ratio at time t is computed on a two-year rolling window ending one month before day t and then averaged across the sample period. We control for the option-implied volatility as of $t-1$ and we saturate the model with asset and month fixed effects. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and month level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	|Return|	|Return|	|Return|	|Return|
Negative Gamma Imbalance X Low Liquidity	18.663***	15.191***	25.549***	22.547***
	(3.026)	(2.813)	(4.408)	(4.360)
Negative Gamma Imbalance	50.833***	50.758***	20.424***	17.404***
	(6.520)	(6.658)	(6.270)	(6.467)
Low Liquidity	50.866***		52.975***	
	(8.164)		(8.488)	
Constant	126.793***			
	(28.671)			
Observations	1,036,850	1,036,850	1,036,850	1,036,850
R-squared	0.038	0.106	0.128	0.200
Asset Fixed Effects	-	Yes	-	Yes
Month Fixed Effects	-	-	Yes	Yes
Ses Clustered By	Asset-Month	Asset-Month	Asset-Month	Asset-Month

Table VI. Intra-day Auto-Correlation

In these specifications the dependent variable is the auto-correlation coefficient of intra-day returns of the underlying assets at various frequencies (5, 10, 20, 30 and 60 minutes), computed from the TAQ dataset. In every specification we control for the ex-ante level of autocorrelation of each asset, computed as the rolling-average of the autocorrelation coefficient at the relevant frequency. The rolling average for day t is computed over a 2-year window ending one month before day t. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and month level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)	(5)
Dependent Variable	5 Min AC	10 Min AC	20 Min AC	30 Min AC	60 Min AC
Gamma Imbalance (lag)	-0.217	-0.244*	-0.253**	-0.346***	-0.262**
	(-1.583)	(-1.845)	(-2.402)	(-3.735)	(-2.265)
Average AC	0.759***	0.727***	0.604***	0.517***	0.265***
	(36.528)	(33.989)	(25.715)	(22.881)	(6.494)
Observations	767,773	767,773	767,773	767,773	767,773
R-squared	0.089	0.053	0.017	0.010	0.001
Asset Fixed Effects	Yes	Yes	Yes	Yes	Yes
Ses Clustered By	Asset-Month	Asset-Month	Asset-Month	Asset-Month	Asset-Month

Table VII. Gamma Imbalance and Daily Spread

The dependent variable of this panel regression is the *daily spread*, an alternative measure of intra-day volatility defined as $S = (H - L)/(H + L) \times 2 \times 1000$, where H is the highest ask price for asset j in day t and L is the lowest bid price. We control for the option-implied volatility as of $t-1$ and we saturate the model with asset and month fixed effects. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and month level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	Spread	Spread	Spread	Spread
Gamma Imbalance (lag)	-15.977***	-19.463***	-7.068***	-12.318***
	(-3.571)	(-5.083)	(-2.648)	(-4.901)
Implied Volatility (lag)	675.704***	717.048***	632.951***	560.068***
	(20.499)	(12.843)	(27.830)	(19.711)
Constant	-38.528***			
	(-3.161)			
Observations	1,018,199	1,018,199	1,018,199	1,018,199
R-squared	0.349	0.385	0.448	0.479
Asset Fixed Effects	-	Yes	-	Yes
Month Fixed Effects	-	-	Yes	Yes
Ses Clustered By	Asset-Month	Asset-Month	Asset-Month	Asset-Month

Table VIII. Gamma Imbalance and Flash Crashes

The dependent variable of this panel regression is a dummy variable Flash Crash(j, t) which equals to one if stock j experiences a flash crash during day t. Flash crashes are identified applying the algorithm proposed by Christensen, Oomen, and Reno' (2017) to minute-by-minute prices constructed from the TAQ database. The table reports results from a regression of Flash Crash(j, t) onto the Gamma Imbalance of dealers on put and call options written on stock j as of the close of day $t-1$, expressed as a fraction of the average daily volume of the underlying. In specifications (2), (3) and (4) we saturate the model with day and stock fixed effects. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and date level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	Flash Crash	Flash Crash	Flash Crash	Flash Crash
Gamma Imbalance (lag)	-1.89***	-1.97***	-0.56**	-0.59**
	(-5.99)	(-5.90)	(-2.07)	(-1.99)
Constant	0.17***			
	(11.58)			
Observations	375,061	375,061	375,061	375,061
R-squared	0.00	0.00	0.02	0.02
Stock Fixed Effects	-	Yes	-	Yes
Date Fixed Effects	-	-	Yes	Yes
SEs Clustered By	Stock-Date	Stock-Date	Stock-Date	Stock-Date

Table IX. Order-flow Toxicity

In these specifications we re-run our baseline specifications from Table II and Table VII controlling for the average daily order-flow toxicity (VPIN). The VPIN measure is computed from TAQ data following the definition in O'Hara et al (2012) and is then averaged at the daily frequency. We control for the option-implied volatility as of $t-1$ and we saturate the model with asset and month fixed effects. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and month level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	|Return|	|Return|	Spread	Spread
Gamma Imbalance (lag)	-15.552***	-8.921***	-22.385***	-14.515***
	(-4.839)	(-4.920)	(-3.993)	(-4.505)
VPIN	-236.111***	-270.289***	-232.866***	-261.920***
	(-8.631)	(-15.568)	(-5.335)	(-14.141)
Implied Volatility (lag)	413.306***	338.728***	716.880***	576.037***
	(17.585)	(15.411)	(18.923)	(16.153)
Constant	74.766***		54.044***	
	(6.739)		(3.238)	
Observations	764,771	764,770	764,771	764,770
R-squared	0.182	0.244	0.360	0.481
Asset Fixed Effects	-	Yes	-	Yes
Month Fixed Effects	-	Yes	-	Yes
Ses Clustered By	Asset-Month	Asset-Month	Asset-Month	Asset-Month

Table X. Order-flow Toxicity (AC)

In these specifications we re-run the specifications from Table VI controlling for the average daily order-flow toxicity (VPIN). The VPIN measure is computed from TAQ data following the definition in O'Hara et al (2012) and it is then averaged at the daily frequency. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and month level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)	(5)
Dependent Variable	5 Min AC	10 Min AC	20 Min AC	30 Min AC	60 Min AC
Gamma Imbalance (lag)	-0.190	-0.222*	-0.240**	-0.334***	-0.261**
	(-1.467)	(-1.777)	(-2.383)	(-3.748)	(-2.281)
VPIN	-11.850***	-9.524***	-6.354***	-6.358***	-0.229
	(-6.212)	(-5.273)	(-4.365)	(-4.318)	(-0.172)
Average AC	0.759***	0.727***	0.595***	0.510***	0.258***
	(36.546)	(33.992)	(26.195)	(23.426)	(6.755)
Observations	764,770	764,770	764,770	764,770	764,770
R-squared	0.090	0.053	0.016	0.010	0.001
Asset Fixed Effects	Yes	Yes	Yes	Yes	Yes
Ses Clustered By	Asset-Month	Asset-Month	Asset-Month	Asset-Month	Asset-Month

Table XI. Excluding Fundamental News

We run our baseline regression from Table II on a subsample where which excludes days where fundamental news on the underlying assets are released. To proxy for fundamental news we use the RavenPack dataset, covering news about American market listed companies from a number of media outlets and employing a machine learning algorithm is used to assign each news a continuous score based on its direction (0 for very bad news and 1 for very good news). Specification (1) uses the subset of the sample for which the RavenPack data is available, i.e. options on equity stocks from January 2000 to December 2012. In specifications (2) we exclude asset-day observations (i, t) for which RavenPack reports at least one negative (score ≤ 0.25) or positive (score ≥ 0.75) news within a window of 5 trading days centered around day t. This filter excludes roughly 20% of the sample. In specifications (3) and (4) we interact a dummy variable indicating news with the lagged gamma imbalance. In all specifications we control for the option-implied volatility as of $t-1$ and we saturate the model with asset and month fixed effects. T-stats are reported in parentheses, based on robust standard errors double-clustered at the asset and month level. Asterisks denote significance levels (***= 1%, **= 5%, *= 10%).

	(1)	(2)	(3)	(4)
Dependent Variable	|Return|	|Return|	|Return|	|Return|
Gamma Imbalance (lag)	-10.579***	-10.637***	-18.503***	-10.547***
	(-5.293)	(-5.082)	(-4.889)	(-5.292)
Implied Volatility (lag)	373.866***	373.920***	439.669***	373.887***
	(13.406)	(13.160)	(14.856)	(13.405)
News			31.547***	26.849***
			(7.702)	(7.158)
Gamma Imbalance (lag) X News			-3.573	-3.230
			(-0.641)	(-0.593)
Constant			-37.725***	
			(-3.392)	
Observations	596,289	475,404	596,289	596,289
R-squared	0.247	0.252	0.186	0.247
Asset Fixed Effects	Yes	Yes	-	Yes
Month Fixed Effects	Yes	Yes	-	Yes
Ses Clustered By	Asset-Month	Asset-Month	Asset-Month	Asset-Month
Subsample	Full Sample	Excluding News	Full Sample	Full Sample

VIII. Figures

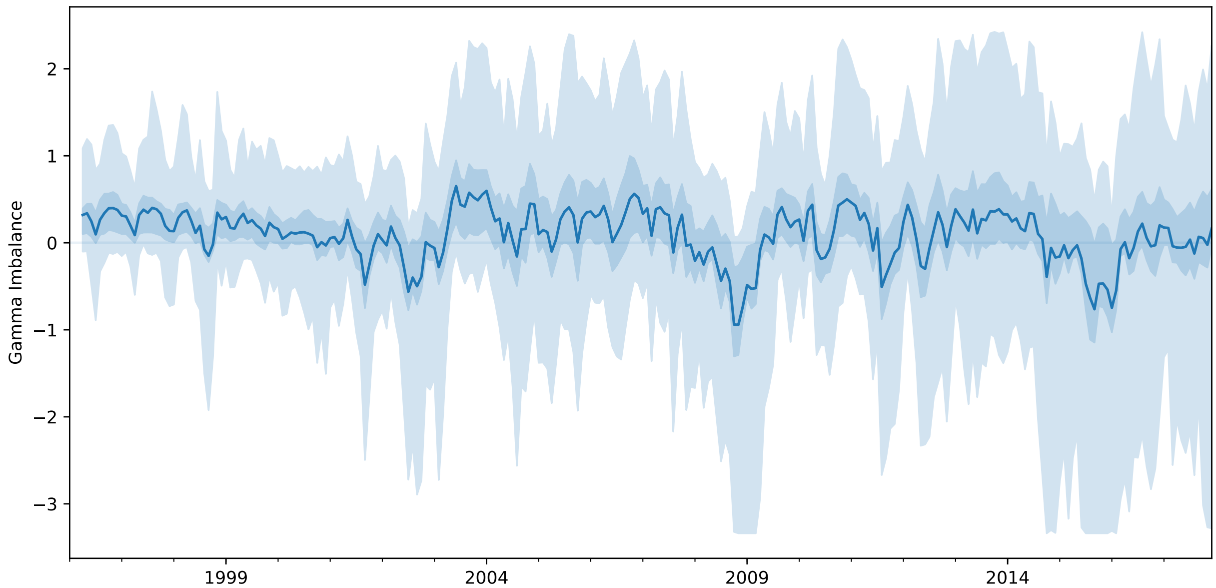

Figure 4. Gamma Imbalance – Panel Variation
The figure displays cross-sectional and time-series variation of gamma imbalance $\Gamma^{IB}(t)$ for the assets in our sample, defined as

$$\Gamma^{IB}(t) = \left(\Gamma^{\$}_{Call}(t) - \Gamma^{\$}_{Put}(t)\right) / \text{ADV}(t) \times 100. \tag{9}$$

The blue line is the cross-sectional average for day t, while the light-blue area represents the interval between the 10th and 90th percentiles in the day-level distributions of the variable, while the darker blue area represents the interval between the 25th and 75th percentiles.

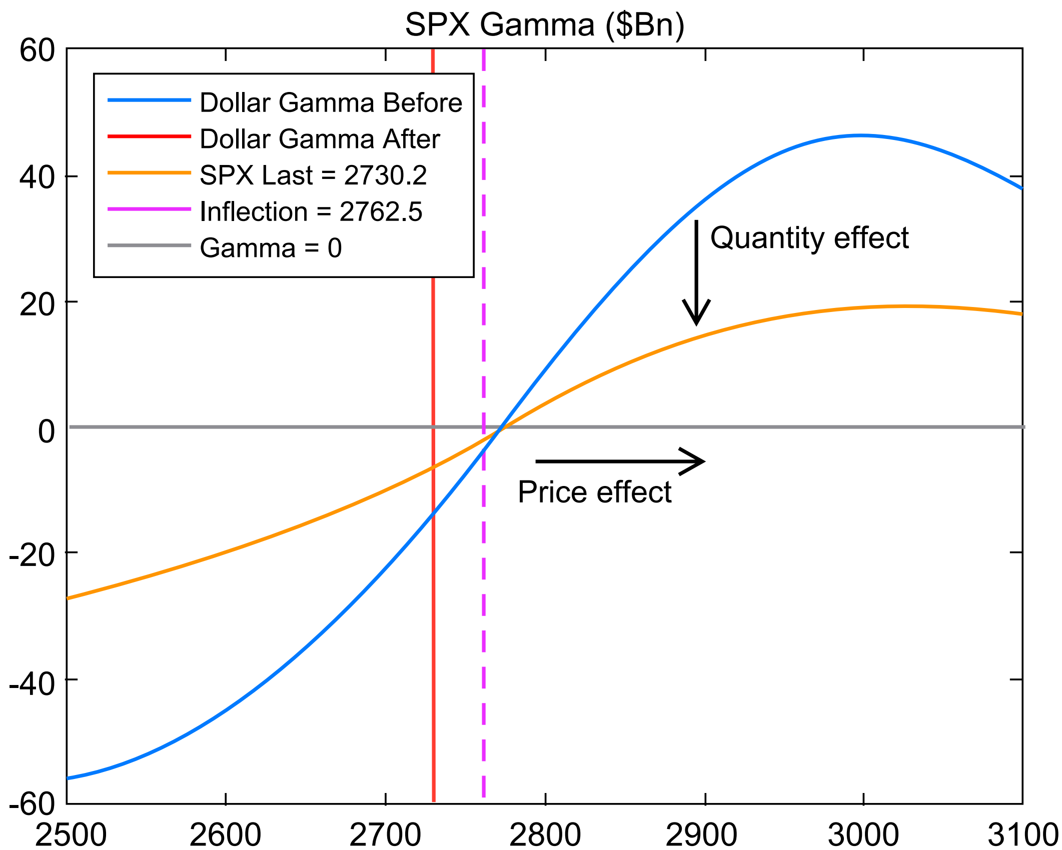

Figure 5. Gamma Inflection Point
This Figure shows how the aggregate Gamma changes as a function of the level of the underlying asset. The figure refers to November 15th ("Before") and 16th 2018 ("After") for options on the SP500. The "Inflection" is defined as the level of the SP500 at which aggregate Gamma is zero.

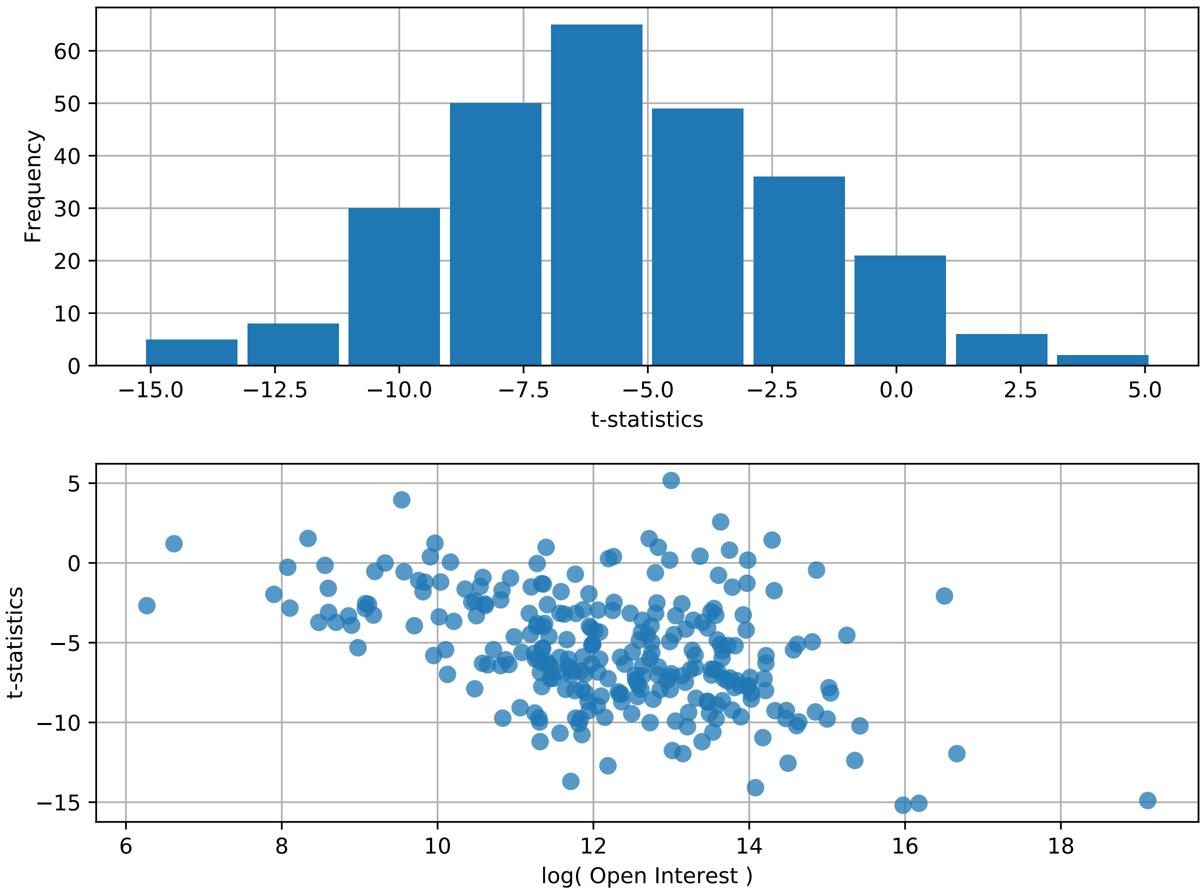

Figure 6. Distribution of t-statistics
The top panel in the figure shows the distribution of the t-statistics on the beta coefficient from the following time-series regressions run on each of the 280 assets in our sample:

$$\mid R_t \mid = \alpha + \beta\,\Gamma_{t-1}^{IB} + \varepsilon_t,$$

where $\mid R_t \mid$ is the absolute value of the asset return on day t and Γ_{t-1}^{IB} is our proxy for the degree of gamma imbalance of dealers as-of the previous day closure, defined in Section B. T-statistics are computed using the heteroskedasticity and autocorrelation consistent (HAC) standard errors from Newey and West (1986). The bottom panel is a scatter plot that summarizes the relation between the implied t-statistics and the log open interest of the options written on the corresponding security.

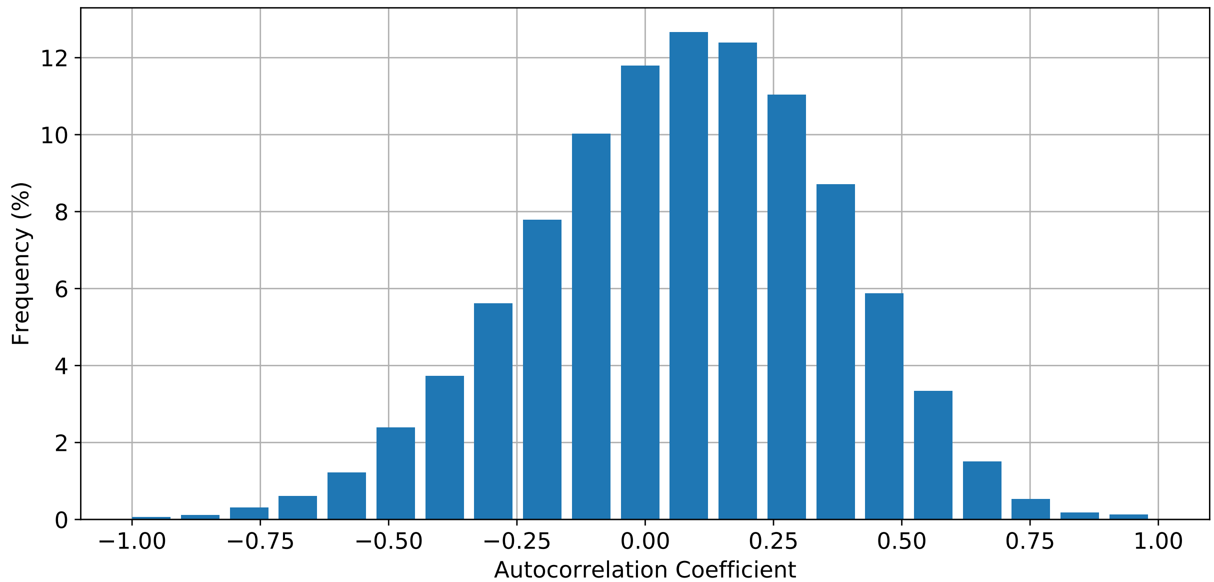

Figure 7. Distribution of Autocorrelation Coefficients
The figure shows the distribution of intraday autocorrelation coefficients for equity returns sampled at the 5-minutes frequency, computed at the stock-hour level.

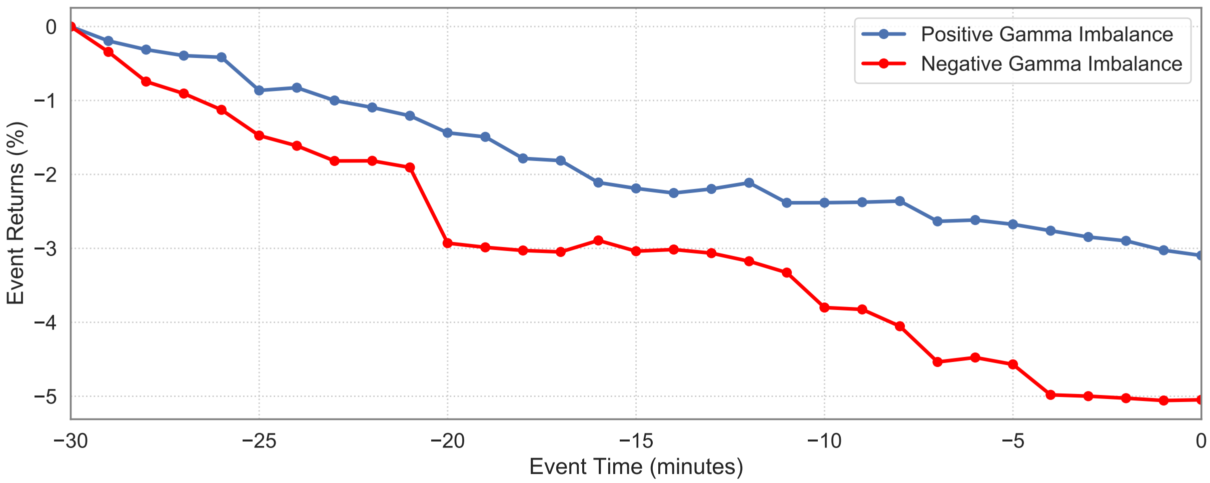

Figure 8. Flash Crashes Drawdowns
The figure presents the event-time average price drop during flash crashes, for two distinct groups of events. The first group (blue solid line) consists of events for which the gamma imbalance for the stock subject to the flash crash was positive, as measured at the previous-day closure. The second group (red dashed line) consists of events for which the gamma imbalance for the stock subject to the flash crash was negative, as measured at the previous-day closure. Flash crashes are identified applying the algorithm proposed by Christensen, Oomen, and Reno' (2017) to minute-by-minute prices constructed from the TAQ database. Time $t = 0$ represents the minute on which the algorithm identifies that a flash crash is starting at $t = -30$.

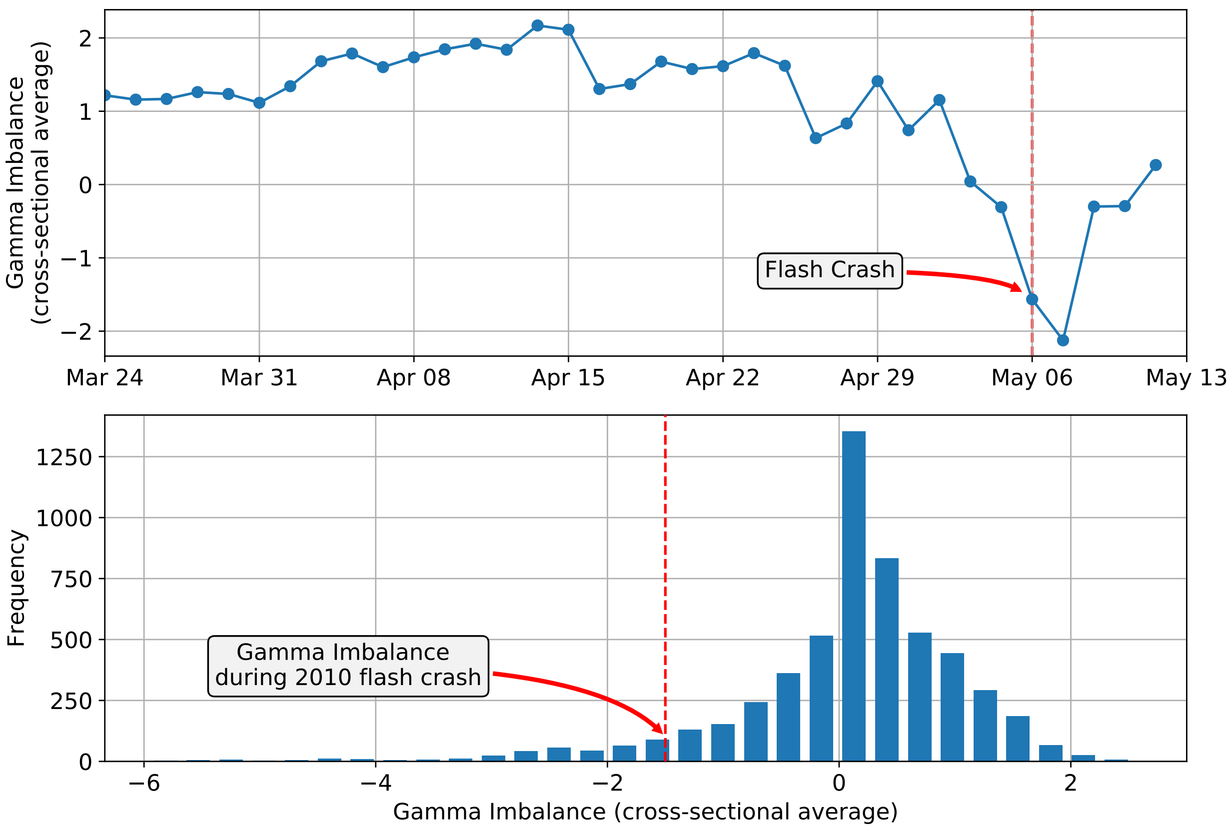

Figure 9. Gamma Imbalance during 2010 Flash Crash
The top panel of the figure shows the time-series evolution of market-wide gamma imbalance, that is the daily cross-section average of stock-level gamma imbalances in standard deviation units, around the flash crash of May 6, 2010. On that date, the total market-wide dollar imbalance was negative at roughly −431 million US Dollars. The bottom panel presents the distribution of market-wide gamma imbalance, with the vertical dashed line representing the value on May 6, 2010.

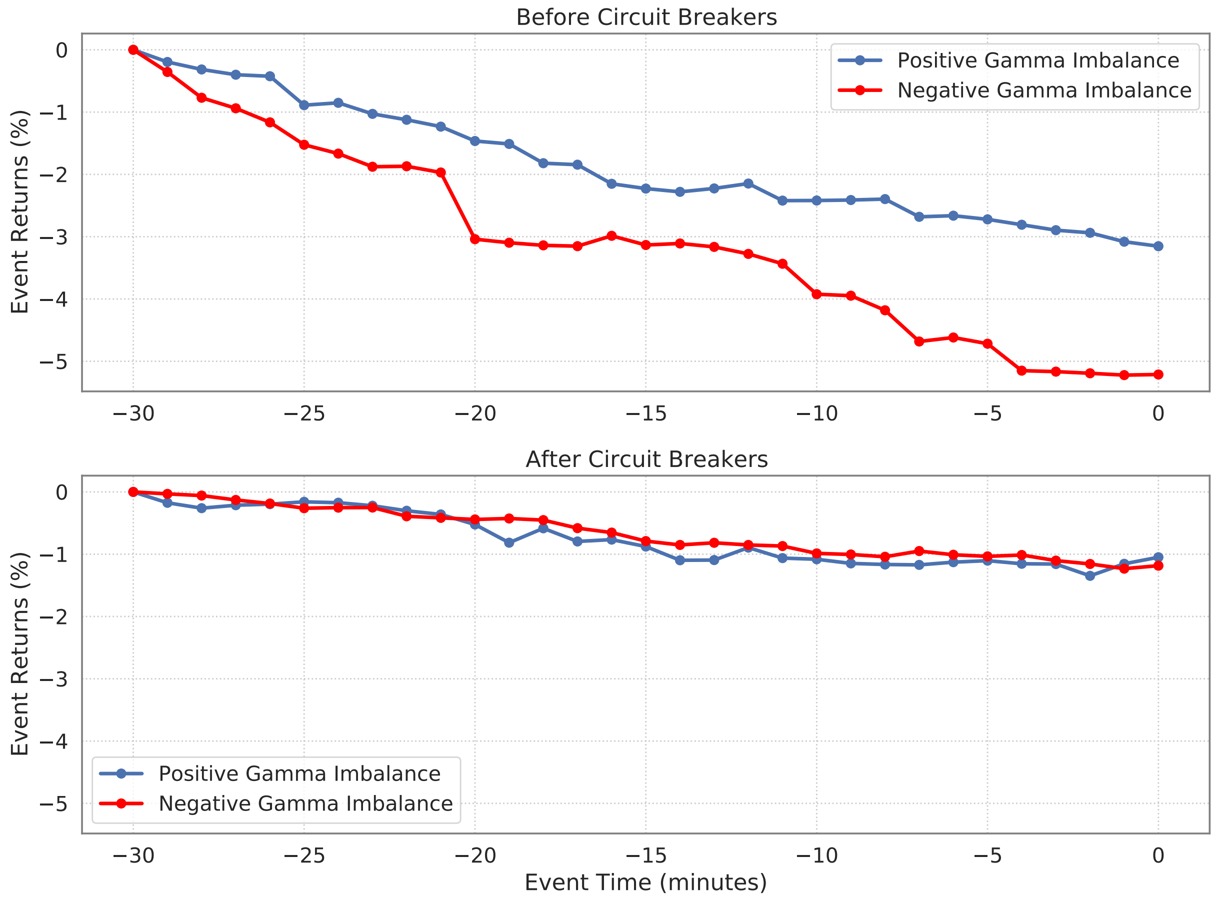

Figure 10. Flash Crashes and Circuit Breakers

The figure presents the event-time average price drop during flash crashes before and after the introduction of circuit reakers (June 2010), for two distinct groups of events. The first group (blue solid line) consists of events for which the gamma imbalance for the stock subject to the flash crash was positive, as measured at the previous-day closure. The second group (red dashed line) consists of events for which the gamma imbalance for the stock subject to the flash crash was negative, as measured at the previous-day closure. Flash crashes are identified applying the algorithm proposed by Christensen, Oomen, and Reno' (2017) to minute-by-minute prices constructed from the TAQ database. Time $t = 0$ represents the minute on which the algorithm identifies that a flash crash is starting at $t = -30$.