

# ODTE Index Options and Market Volatility: How Large is Their Impact?★

Diego Amaya<sup>a</sup>, Pedro A. Garcia-Ares<sup>b,c</sup>, Neil D. Pearson<sup>d,e</sup>, and Aurelio Vasquez<sup>f</sup>

January 25, 2025

## Abstract

The large volume of trading in ODTE S&P 500 index options suggests that options market makers (OMMs) have large positions in these options. Because short-dated options have large gammas, OMMs' hedge rebalancing trades might be large enough to impact the index. We estimate the maximum impact of OMM gamma on index volatility using proprietary trade data to determine the aggregate position of OMMs, and their gamma. We then estimate models that relate market volatility to the OMM gamma, and simulate the models under the counterfactual assumption that OMM gamma does not impact volatility to identify the impact of OMM gamma.

*JEL Classification:* G12, G13, G23, G24

*Keywords:* ODTE options, SPX options, S&P 500 index options, hedge rebalancing, gamma squeeze, options trading

---

★ We thank Cboe Global Markets for providing the data on the SPX and SPXW options trades, and Algoseek LLC for providing the quote data we use.

<sup>a</sup> Lazaridis School of Business & Economics Business, Wilfrid Laurier University, 75 University Avenue West, Waterloo, Ontario, Canada, N2L 3C5. Tel.: +1 (519) 886-9351 ext. 4428. E-mail: damaya@wlu.ca.

<sup>b</sup> Mendoza School of Business, University of Notre Dame, E-mail: pgarciaa@nd.edu.

<sup>c</sup> Instituto Tecnológico Autónomo de México (ITAM), Río Hondo No. 1, Col. Progreso Tizapán, México, D.F., México, 01080. Tel.: +52 (55) 5628 4000 ext. 6518. E-mail: pedro.garcia@itam.mx.

<sup>d</sup> Department of Finance, University of Illinois at Urbana-Champaign, 1206 South Sixth Street, Champaign, Illinois, 61820. Tel.: +1 (217) 244-0490. E-mail: pearson2@illinois.edu.

<sup>e</sup> Canadian Derivatives Institute, 3000, Chemin de la Côte-Sainte-Catherine, Montréal, Québec H3T 2A7, Canada.

<sup>f</sup> Instituto Tecnológico Autónomo de México (ITAM), Río Hondo No. 1, Col. Progreso Tizapán, México, D.F., México, 01080. Tel.: +52 (55) 5628 4000 ext. 6518. E-mail: aurelio.vasquez@itam.mx

# **0DTE Index Options and Market Volatility: How Large is Their Impact?**

January 25, 2025

## **Abstract**

The large volume of trading in 0DTE S&P 500 index options suggests that options market makers (OMMs) have large positions in these options. Because short-dated options have large gammas, OMMs' hedge rebalancing trades might be large enough to impact the index. We estimate the maximum impact of OMM gamma on index volatility using proprietary trade data to determine the aggregate position of OMMs, and their gamma. We then estimate models that relate market volatility to the OMM gamma, and simulate the models under the counterfactual assumption that OMM gamma does not impact volatility to identify the impact of OMM gamma.

*JEL Classification:* G12, G13, G23, G24

*Keywords:* 0DTE options, SPX options, S&P 500 index options, hedge rebalancing, gamma squeeze, options trading

## 1. Introduction

Whenever stocks drop violently, dealers' gamma positioning is often at the top of any investigation of potential culprits.

The equity market selloff in 2018, dubbed 'volamageddon', and the Covid-related rout of 2020 offer prime examples of the havoc banks' options hedging activities can wreak in jittery markets. Last week's meltdown may be different, however, with market strategists offering competing theories on the role gamma position played when the S&P 500 shed 6% in three days.

Faye Kilburn, "After the selloff, competing theories on dealer gamma,"  
*Risk.net*, August 12, 2024.

If options market makers (OMMs) hold negative gamma options positions and delta hedge those positions, the stock or stock index futures trades they execute to rebalance their delta hedges can exacerbate movements in the underlying market, and thus increase market volatility. Specifically, if the gamma of the aggregate OMM position is negative, OMMs in aggregate will buy as the market moves up and sell as the market moves down. On the other hand, OMM hedge rebalancing can attenuate market movements when OMMs have positive gamma positions. In both cases, the strength of the effect depends on the magnitude of OMM gamma, and on whether OMMs rebalance their delta hedges. If the gamma of the aggregate OMM position is large and OMMs rebalance their delta hedges, the hedge rebalancing trades can result in a relation between market volatility and the gamma of the aggregate OMM options position. Ni et al. (2021) explain this mechanism and provide evidence consistent with the hypothesis that OMM rebalancing of delta hedges of single-name options impacts the volatilities of individual stocks.

The recent dramatic increase in the trading of 0DTE options on the S&P 500 index suggests that OMMs will often have large positions in 0DTE S&P 500 options. Because the gammas of short-dated (for example, 0DTE) at-the-money (ATM) options are very large, market movements in which the S&P 500 index crosses the options' strike prices can result in large changes in delta and OMMs' required delta-hedge positions. This raises the possibility that OMMs might trade large quantities of the underlying, likely through the index futures market, as the S&P 500 index moves. This potentially impacts the entire stock market. If this mechanism is at play, there will be a relation between the gamma of the aggregate position held by OMMs and market volatility.

We estimate this relation, and use it to estimate the maximum impact of OMM gamma on S&P 500 index volatility during the period running from July 2020 through June 2023. Specifically, we estimate the maximum impact of OMM gamma on market volatility over one-day and 30-minute periods, compared to a counterfactual in which OMM gamma does not impact market volatility. We do this by first using data on all trades in Cboe-listed SPX and SPXW options to construct the aggregate OMM net position in every SPX and SPXW option series, and the gamma of the OMM aggregate net position across all series, at one-minute frequency throughout regular trading hours. We then use the estimates of OMM gamma to estimate statistical models that relate market volatility to OMM gamma. Finally, we simulate the models under the counterfactual assumption that market volatility is not related to gamma to estimate what market volatility would have been but for the impact of OMM gamma.

The results show that the gamma of the aggregate OMM position, while typically positive, is often negative. The typical impact of gamma is to reduce volatility, which is to be expected, given the negative relation between volatility and gamma and the typically positive values of gamma. But volatility is elevated when OMM gamma is negative. We find that the maximum impact of OMM gamma on annualized volatility is to increase annualized realized daily volatility by 3.3 percentage points. An increase of 3.3 percentage points is well within the range of daily changes in annualized realized volatility. The maximum impact of OMM gamma on the annualized volatility during 30-minute periods is to increase it by 6.4 percentage points.

Others also examine the issue of the impact of 0DTE option trading on market volatility. Brogaard, Han, and Won (2024) show that stock market return volatility, as measured by the return volatility of an S&P 500 exchange-traded fund, is correlated with the ratio of 0DTE trading volume to total trading volume in S&P 500 index options. Adams, Fontaine and Ornathanalai (2024) carry out a well-identified analysis that shows that market volatility is lower on days when 0DTE options are available for trading.<sup>1</sup> They also show that volatility is related to the sign and magnitude of the gammas of OMM positions, consistent with the analysis in Ni et al. (2021). Dim, Eraker, and Vilkov (2024) also present evidence consistent with the hypothesis that volatility is related to the sign and magnitude of the gammas of OMM positions.

---

<sup>1</sup> To do this, they exploit the fact that SPXW options expiring on Tuesdays and Thursdays were not available for trading until April 26 and May 19, 2022, respectively, other than the times when the quarterly and month-end SPXW expirations fell on a Tuesday or a Thursday. Thus, prior to May 2022 there were many days on which there was no 0DTE trading.

These papers focus on the mean and conditional mean impact of 0DTE trading on volatility, and find that the average impact is to reduce volatility because OMMs typically have positive gamma positions. The finding that the average and conditional average effect is to reduce market volatility is consistent with 0DTE trading positions sometimes causing dramatic increases in volatility. We focus on this possibility. Specifically, we estimate the largest increases in risk, measured by volatility, due to 0DTE positions during our sample period.

Other papers that study aspects of 0DTE options trading include Bandi, Fusari and Reno (2024), who present closed-form pricing formulae based on a local-in-time approach that captures 0DTE prices and the implied-volatility surface. Almeida, Freire, and Hizmeri (2024) explore the risk premia investors can expect to earn holding 0DTE options. Lipson, Tomio, and Zhang (2023) show that 0DTE trading of single-name options is associated with increases in the volatility of the options' underlying stocks. In other work related to the impacts of hedge rebalancing on underlying asset markets, Ni, Pearson, and Poteshman (2005) show that hedge rebalancing trading on expiration dates of single-name options causes stocks to “pin” on option strike prices, and Golez and Jackwerth (2012) show that index option hedge rebalancing and unwinding impacts index futures prices. Baltussen et al. (2021) present evidence consistent with the hypothesis that hedge rebalancing contributes to intraday market momentum. Barbon et al. (2023) present related results.

We use trade records provided by Cboe Global Markets (Cboe) to construct the aggregate OMM net position in each SPX and SPXW options series at one-minute frequency during the period running from July 2020 through June 2023. The trade records provide, for each trade, the option series, the trade date and time, the trade price and quantity, and the trading capacities of the traders who bought and sold. For example, the trade records might indicate that a Customer sold and a Market Maker bought. For each option series, we aggregate the trades of each capacity from the inception of trading in that series to compute the aggregate position of each capacity, including the Market Maker capacity, at one-minute frequency.

We next compute the gammas of each option series at the same one-minute frequency, multiply by the aggregate OMM net position in that series, to compute the OMM gamma in that series, and sum across the series to compute the OMM gamma across all series. To compute the gamma of each series, we obtain best bid and offer (BBO) quotes at one-minute frequency for each series. Using the BBO midpoint as an estimate of the option value, we use the Black-

Scholes-Merton formula (the version of the Black-Scholes formula that includes a continuous dividend yield), together with data on interest rates and option-implied dividend yields, to compute the implied volatility of each option series at one-minute frequency. We then use the implied volatility, together with data on interest rates and dividend yields, to compute the gamma of each series using the same Black-Scholes-Merton formula. Finally, we aggregate the gammas of the series to compute the aggregate OMM gamma across all series, as indicated above. Gamma is typically positive, implying that hedge rebalancing trades will typically reduce volatility, but sometimes negative and large.

We then use the OMM gammas to estimate GARCH models that combine features of the models proposed by Engle and Sokalska (2012) and Engle, Ghysels, and Sohn (2013). The models specify the conditional variance of each one-minute return to be a function of the previous squared return, the previous conditional variance, and lagged OMM gammas. They also allow the return variance to differ across days, and for a diurnal pattern within each day. The results show that the conditional variance is negatively related to the lagged OMM gammas, as expected. We explore whether the negative relation is robust to the model specification by estimating linear models that relate the one-minute squared return to lagged squared returns and lagged OMM gammas. The results of the linear models are consistent with those of the GARCH models.

We use the models to generate simulated returns under the counterfactual assumption that the return variance is unrelated to the lagged gammas. We then use the simulated returns to compute realized variance for each trading day, and compare these counterfactual realized variances to the realized variances computed from the actual one-minute returns. The differences between the realized variances based on actual and simulated returns provide measures of the impact of OMM gamma on returns. We use the same approach to compute simulated and actual realized variances for non-overlapping 30-minute periods starting 9:30 am each day.

The results show that the typical impact of gamma is to reduce volatility, with the median impact being to reduce daily volatility by 0.08 percentage points. Negative impacts on volatility are to be expected, given the negative relation between variance and gamma and the typically positive values of gamma. But OMM gamma is sometimes negative, and at these times volatility is elevated. We find that the maximum impact of OMM gamma on annualized volatility is to increase annualized daily volatility by 3.3 percentage points. The maximum impact of OMM

gamma on annualized 30-minute volatility is to increase it by 6.4 percentage points. We carry out the same exercises using the linear models, and obtain similar results.

Should one think of a gamma-induced change in annualized realized volatility of 3.3 percentage points as large or small? To address this question, we compute the annualized realized volatility for each day, and examine the distribution of daily changes in that quantity. The standard deviation of daily changes in annualized realized volatility is 4.5 percentage points, and the change in annualized realized volatility from the previous day exceeds 3.0 percentage points on about 20% of the trading days. Thus, changes in annualized realized volatility of about 3.0 percentage points or greater occur reasonably frequently, on about 20% of the days, or on average about once per week. Given this, our interpretation of the results is that a gamma-induced change of about 3.3 percentage points is not large.

We also use the models to compute the impact on realized volatility during 30-minute periods running from 9:30 a.m. through 4:30 p.m. on each day. After annualizing the changes, the standard deviation of the gamma-induced change in 30-minute realized volatility is 0.9 percentage points. This is smaller than the standard deviation of changes in annualized 30-minute realized volatility due to all causes, which is 4.6 percentage points. The maximum gamma-induced change in annualized 30-minute realized volatility is 6.4 percentage points, which is much less than the maximum change in annualized 30-minute realized volatility of 63.4 percentage points.

The maximum gamma-induced change in annualized 30-minute realized volatility of 6.4 percentage points is slightly less than the 95<sup>th</sup> percentile of the distribution of changes in realized volatility due to all causes. Each non-holiday each week there are 65 30-minute periods between 9:30 a.m. and 4:00 p.m. Using the fact that more than 5% of the changes in annualized 30-minute realized volatility exceed the maximum gamma-induced change of 6.4 percentage points, during a typical week there are more than three 30-minute periods in which the change in annualized realized volatility exceeds 6.4 percentage points. Thus, changes in annualized 30-minute realized volatility that exceed our estimate of the maximum gamma-induced change occur frequently. Similar to our conclusion about the maximum gamma-induced change in daily volatility, we conclude that the maximum gamma-induced change in 30-minute volatility is not unusual.

The remainder of the paper is as follows. Section 2 describes the data we use, and how we use the data to determine OMM positions and compute the gamma of the aggregate OMM

position. Section 3 presents the results of estimating the models. Section 4 describes how we simulate the models to estimate the maximum impact of gamma on realized variance, while Section 5 examines how the estimated gamma-induced increases in volatility compare to typical changes in volatility. Section 6 briefly concludes.

## **2. Data and computation of the gamma of the aggregate market maker net position**

The main data we use are records of all trades in S&P 500 index options with symbols SPX and SPXW during the period running from January 2020 through June 2023. We first describe these data, and how we use them to compute the aggregate OMM net position in every SPX and SPXW options series at one-minute frequency throughout the trading day. Then we describe how we compute the gamma of the OMM position. This involves using quote data to compute the implied volatilities at one-minute frequency of all SPX and SPXW option series in the aggregate OMM net position, and then using the implied volatilities and other data to compute the position gammas.

### *2.1 SPX and SPXW options trade data*

The records of trades in S&P 500 index options are provided by Cboe Global Markets. Each single-leg trade, and each leg of each complex trade, appears as at least two trade records, one for the buyer and one for the seller. A trade or trade leg appears as more than two trade records when the passive side of the trade is shared among two or more traders. Each trade record includes the option series traded, the trade date and time, trade price, number of contracts, trading capacity of the trader, and whether the option is purchased or sold. We group the trading capacities provided in the Cboe data into the broader categories of customer, market maker, professional customer, firm, and broker-dealer. For the non-market maker capacities, each trade record also indicates whether the trade opened a new options position or closed an existing position. The dataset contains a total of 442,589,991 trade records.

Table 1 shows the average percentages of trading volume due to the buys and sells of each trading capacity, for options with different numbers of days to expiration (DTE). There are four panels, for trading in options with 0DTE, 1-5DTE, 6-21 DTE, and > 21 DTE. Other than the sums in the last row and column of each panel, each entry is the average percentage of total SPX and SPXW volume due to buys by the trading capacity indicated in the row label and sells by the capacity indicated in the column heading. We compute the average percentages by first computing the percentage for each trading date, where the denominator is the total trading

volume on that date, across all SPX and SPXW option series. We then average the daily percentages across the trading days during our sample period. The last row and last column of each panel display the sums of the entries across rows or columns, respectively.

For example, the entry in the row labelled “Customer” and column headed “Market maker” in Panel A shows that on average 14.952% of total SPX and SPXW option trading volume is due to trades in 0DTE options which a customer bought and a market maker sold. The entry in the row “Customer” and column “Customer” indicates that on average 1.659% of volume comes from trades in which a customer bought and another customer sold. The other entries in this row show that 0.201%, 0.272%, and 0.066% of volume comes from trades in which a customer bought and a firm, professional customer, or broker dealer sold, respectively. The last entry in the row shows that 17.150% of total SPX and SPXW trading volume is due to trades in which a customer bought a 0DTE option. The entry in the column headed “Customer” and row labelled “Sum” indicates that a similar, but slightly smaller, percentage of trading volume, 15.988% consisted of trades in which a customer sold a 0DTE option. The lower right-hand entry in Panel A is the sum of the 25 entries for the different trading capacities. It shows that on the average day trading in 0DTE options accounted for 34.771% of total SPX and SPXW trading volume. The results displayed in Panels B-D are interpreted similarly.<sup>2</sup>

The results in Table 1 show that trading in SPX and SPXW options is dominated by customers and market makers. On the average day, trades in which a customer buys and a market maker sells account for 34.835% ( $= 14.952\% + 9.516\% + 5.115\% + 5.225\%$ ) of trading volume, and trades in which a customer sells and a market maker buys account for 33.882% ( $= 14.062\% + 9.436\% + 5.094\% + 5.290\%$ ) of volume. Thus, on the average day, 68.717% of volume involves a customer trading with a market maker. An additional 8.959% ( $= 1.477\% + 2.625\% + 1.860\% + 2.997\%$ ) involves a market maker trading with another market maker, and 2.184% ( $= 1.659\% + 0.389\% + 0.075\% + 0.061\%$ ) involves a customer trading with another customer. Trading volume due to customers trading with other customers is a much larger fraction of 0 DTE volume as compared to trading volume in options with one or more days to expiration. Customers and market makers of course also traded with other capacities. Summing the entries

---

<sup>2</sup> The sum of the entries in the lower right-hand corners of the four panels,  $34.771\% + 29.215\% + 18.277\% + 27.095\% = 109.358\%$ , does not sum to 100% because the results are averages of ratios and total trading volume in SPX and SPXW options differs considerably across dates.

in the rows and columns labelled “Firm,” “Professional Customer,” and “Broker Dealer” across the four panels, on the average day only 2.866% of trades involved neither a customer nor a market maker.

## *2.2 Computing the aggregate OMM net positions*

Because the trade records include every trade in which an OMM participated, and indicate whether the OMM bought or sold, we can compute the aggregate OMM net position at any time for most options series. The idea is that, at the inception of trading in each option series, we know that the aggregate OMM net position in that series must be zero. Then each OMM buy increases the aggregate OMM net position in that option series by the trade size, and each market maker sell decrements the aggregate net position by the trade size. Thus, we can determine the aggregate net position at any time simply by cumulating the market maker trades from the inception of trading up to that time, where a buy is a positive increment and a sell is a negative increment.

We cannot compute the aggregate OMM net position for every option series because we cannot compute it for options series that began trading before January 2, 2020, which is the first date for which we have the Cboe trade data. For options that began trading before that date, computing market maker positions by cumulating the trades starting from January 2020 results in an error equal to the aggregate OMM net position at the beginning of January 2020. We limit the magnitude of this error by carrying out the statistical analysis using our estimates of market maker positions and gammas only starting from July 2020. Of the option series that are available for trading at the beginning of July 2020, relatively few are available for trading on January 2, 2020. These are the SPX options with expirations in September, October, November, and December 2020, March, June, September, and December 2021, and December 2022. For the series expiring at these times, the error consists of the market maker positions that existed at the beginning of January 2020, more than six months before the expiration of the first of these option series to expire. Because open interest in options with six months or more to expiration tends to be small relative to open interest in shorter expirations, and the aggregate OMM net position must be less than or equal to open interest, our approach of carrying out the statistical analysis using only the estimates starting from July 2020 limits the magnitude of the error in our estimates of aggregate OMM net positions. During the latter part of the sample period the errors

are smaller or non-existent, because there are fewer or no options that began trading before January 2020.

Having decided to carry out the statistical analysis using estimates starting from July 2020, we perform the computation described above for every minute during regular trading hours starting from July 2020. Specifically, for each minute starting from the beginning of trading on July 1, 2020, and for each option series, we cumulate the signed market maker trading volumes starting from January 2020 to estimate the aggregate market maker net position in that series. The result is, for each option series, a time-series of aggregate market maker net positions at one-minute frequency that consists of the positions at each minute from 9:30 am to 4:15 pm on each date, starting from the first day the option series traded.

### *2.3 Computing the gammas of the aggregate OMM net position*

Then, for each series and minute, we compute the gamma by multiplying the position (in contracts) by the contract multiplier of 100 and the gamma of the series as of that minute. The gamma is the second derivative of the Black-Scholes-Merton formula. Computing the gamma requires a dividend yield, an interest rate, and an implied volatility.

For each trade date, OptionMetrics provides the option-implied S&P 500 index dividend yield for each SPX or SPXW expiration date. For each option series  $j$  and minute on date  $d$ , we use the OptionMetrics dividend yields from the preceding date  $d - 1$ . OptionMetrics also provides, for each date, a term structure of zero-coupon interest rates, for a set of maturity dates that are not necessarily option expiration dates. For each option series  $j$  and minute on date  $d$ , we interpolate the date  $d$  OptionMetrics interest rates to obtain a zero-coupon rate for the expiration date of series  $j$ .

In contrast to the dividend yields and interest rates, which are treated as constant within each date, we compute implied volatilities for each series for each minute during the trading day. We do this using “minute bar” data for the SPX and SPXW options obtained from Algoseek LLC. Each minute bar includes, among other data items, the prevailing BBO quotes at the end of the minute. For each minute at each time  $t$ , and each SPX and SPXW series in which the aggregate market maker net position is non-zero, we use the midpoints of these quotes to compute the implied volatility of the option series as of the end of the minute by (numerically) inverting the Black-Scholes-Merton formula. We obtain the spot price needed for this calculation from the price of the E-mini futures contract (symbol ES) that has the greatest trading volume on

that date, adjusted from the futures maturity to the current date. The calculation of implied volatility also requires a dividend yield and an interest rate that are obtained from OptionMetrics, as described above.

Finally, for each minute on each date, we sum the gammas of the option series that are available for trading on that date to obtain the gamma of the aggregate market maker net position as of that date and minute. The result is a time-series of market maker gammas at one-minute frequency that consists of the gammas at each minute from 9:30 am to 4:00 pm on each date, during the period running from July 2020 through June 2023.<sup>3</sup>

#### *2.4 Summary statistics*

We report summary statistics that describe some features of the estimates of OMM gamma in Table 2. We first scale all the OMM gamma estimates by dividing them by  $10^9$ . For each date, we compute the mean, median, standard deviation, maximum, and minimum, and autocorrelation of the (scaled) gamma estimates at one-minute frequency. The result is a sample of daily values of each statistic, for example, a sample of daily means. We then compute the mean, standard deviation, and several percentiles of the samples of daily values, and report the results in Table 2.

Panel A displays the results for all days during the period running from July 2020 through June 2023. The results show that gamma is typically positive, with the mean of the daily means being 341.2 (that is,  $341.2 \times 10^9$ ). It also displays considerable variation, with the standard deviation of the daily means being 319.6, almost as large as the mean of the daily means. The distribution of daily minima shows that gamma is negative at some point on at least 25% of the days (the 25<sup>th</sup> percentile of the daily minima is  $-94.9$ ). The gamma is negative for the entire day on at least 1% of the days, as the first percentile of the daily maxima is  $-55.1$ . The minimum of the daily minima is  $-5,426.6$ , and the maximum of the daily maxima is even larger, 16,743.9.

Panels B and C display the results for subsamples of days during the periods running from July 2020 through April 2022 and from May 2022 through July 2023. The dividing point, the beginning of May 2022, is approximately the time when the Tuesday and Thursday

---

<sup>3</sup> The 9:30 am gamma reflects all trades up through 9:30 am, that is, it includes the trades during the overnight period. Even though our gamma estimates include all SPX and SPXW options, we stop computing gamma at 4:00 pm on each date because our analysis is motivated by the possible impact of large ODTE positions, and the ODTE options expire at 4:00 pm.

expirations were introduced.<sup>4</sup> The results in these panels show that OMM gamma is generally greater before the Tuesday and Thursday expirations were introduced. For example, the means of the daily means are 408.8 before, and 234.2 after, the Tuesday and Thursday expirations were introduced. Negative gammas are also more frequent. After the Tuesday and Thursday expirations were introduced, the median of the daily minima is negative (−45.0), implying that gamma is negative at some point during at least half of the trading days.

Figure 1 illustrates one feature of the variation in gamma by plotting the average OMM gamma on each trading date. Specifically, for each trading date we average the OMM gammas at one-minute frequency, resulting in a single average value for each day. We then scale these by dividing by  $10^{12}$ , and plot these scaled daily average values in the figure. The figure shows considerable day-to-day variation in OMM gamma. One feature apparent in the plot is that daily average gamma is rarely negative before 2022, but is often negative starting in early 2022.

Figure 2 illustrates the intraday variation in OMM gamma. This figure consists of seven panels, each of which displays the time-series of gamma at one-minute frequency on a selected day. As in Figure 1, the gammas are scaled by dividing by  $10^{12}$ . The seven days are chosen to be representative of the possible variation in OMM gamma. Specifically, they are the days with the standard deviation of the one-minute gammas closest to the 1<sup>st</sup>, 5<sup>th</sup>, 25<sup>th</sup>, 50<sup>th</sup>, 75<sup>th</sup>, 95<sup>th</sup>, and 99<sup>th</sup> percentiles of the distributions of daily standard deviations.

### 3. Model estimates

The delta ( $\Delta$ ) of an options position is the derivative of the position value with respect to the underlying price  $S$ , that is,  $\Delta = \partial V / \partial S$ , where  $V$  is the theoretical value of the position. Options market makers often buy or sell the underlier to make the net delta of the combined position of options and underlier be within an acceptable range, that is, to be within some neighborhood of zero. This entails holding a position in the underlier such that the number of units of the underlying asset  $n$  is approximately equal to the negative of the delta of the options positions, i.e., such that  $n \approx -\Delta$ . Under the assumption that market makers hold such (approximately) delta-neutral positions, and rebalance their delta hedges to maintain (approximate) delta neutrality, the hedge rebalancing trades are determined by the negative of the

---

<sup>4</sup> The first Tuesday expiration was on April 26, 2022 and the first Thursday expiration was on May 19, 2022. These options began trading on April 18 and May 11, 2022, respectively. See <https://ir.cboe.com/news/news-details/2022/Cboe-to-Add-Tuesday-and-Thursday-Expirations-for-SPX-Weekly-Options-04-13-2022/default.aspx>.

changes in the delta of the options position. It is well understood that this is given by the gamma of the options position.<sup>5</sup> This is true both for an individual option and an options portfolio, including the aggregate portfolio held by OMMs collectively.

### 3.1 GARCH model

Our goal is to estimate the impact of the aggregate OMM net gamma on volatility, or equivalently, variance. We use a flexible version of a generalized autoregressive conditional heteroskedasticity (GARCH) model that incorporates lagged OMM gammas as additional explanatory variables, to allow for a relation between the conditional variance of returns and lagged OMM gammas. The model also incorporates ideas from the GARCH model proposed by Engle and Sokalska (2012) to allow the level of variance to differ across days, and for a diurnal variance pattern within each day.

Specifically, the one-minute index return is modelled as

$$r_{t,d} = h_d s_t q_{t,d}, \quad (1)$$

where  $(h_d s_t q_{t,d})^2$  is the return variance. In this equation  $d$  indexes the day,  $t$  indexes the minute within the day,  $r_{t,d} = \ln(F_{t,d}/F_{t-1,d})$  is the log change in the midpoint price of the S&P 500 futures contract with highest trading volume on day  $d$ ,  $\varepsilon_{t,d}$  is an innovation with mean zero and variance one,  $h_d^2$  is the daily variance component,  $s_t^2$  is the diurnal variance pattern, and  $q_{t,d}^2$  is the intraday variance component.

Following the two-step estimation approach of Engle and Sokalska (2012), we first estimate the daily variance component  $h_d^2$  and the diurnal variance pattern  $s_t^2$  using non-parametric methods. For the daily component, we compute the daily realized variance using one-minute returns. For the diurnal pattern, we use the daily realized standard deviation to standardize the one-minute returns on each day, and then compute the sample variance for each minute across the full sample.

In the second step, we normalize returns by the daily and diurnal volatility components, and model the residual intraday variance component  $q_{t,d}^2$  as a multi-component GARCH model:

$$\sigma_{t,d}^2 = (1 - \alpha - \beta) + \alpha \frac{z_{t-1,d}^2}{\tau_{t-1,d}^2} + \beta \sigma_{t-1,d}^2, \quad (2)$$

---

<sup>5</sup> The appendix carries out a careful analysis of this.

$$\ln(\tau_{t,d}^2) = m + \gamma \sum_{l=1}^K \varphi_l(\omega_1, \omega_2) \Gamma_{t-l,d}, \quad (3)$$

$$q_{t,d}^2 = \tau_{t,d}^2 \sigma_{t,d}^2, \quad (4)$$

$$z_{t,d} = \frac{\hat{r}_{t,d}}{\hat{h}_d \hat{\sigma}_t}. \quad (5)$$

The specification for the  $\tau$  component builds on the MIDAS specification proposed in Engle, Ghysels, Sohn (2013), which allows the model to incorporate information from an economic variable in the spirit of a MIDAS regression. In our model, the economic variable of interest is  $\Gamma_{t,i}$ , representing the gamma exposure the aggregate OMM net position at minute  $t$  on day  $d$ . Lags of this variable are included in the  $\tau$  component according to the weighting scheme  $\varphi_l$ , governed by two parameters that can represent various shapes, including monotonically increasing (or decreasing) patterns and unimodal, hump-shaped schemes. We use the ‘‘Beta’’ weighting function presented in Engle, Ghysels, Sohn (2013).

We also introduce parameters representing the initial variance for each day. Including initial variance as a parameter to be estimated, together with the daily variance pattern  $h_d^2$ , implies that only information for a given day is used in the variance forecast, as no information about past variables is carried over across days. Our primary prediction is that the  $\gamma$  coefficients are negative and that the weighting scheme exhibits a decreasing magnitude with increasing lag.

We estimate the model in equations (1)-(5) with five lags of gamma for each month from July 2020 to June 2023, using the one-minute returns and quasi-maximum likelihood estimation based on the assumption that  $\varepsilon_{t,d} \sim N(0, 1)$ . We obtain model parameters for each month, and compute average parameter values across the 36 months. The OMM gammas are scaled by dividing them by  $10^{12}$ , and the returns are measured in basis points. With this scaling the magnitudes of the parameter estimates are convenient for reporting.

### 3.2 GARCH model results

In Table 3 we report the results of estimating the GARCH model in equations (1)-(5). The table presents averages across the 36 months of the estimates of the coefficients on  $r_{t-1,d}^2$  and  $\sigma_{t-1,d}^2$ , the parameters  $m$  and  $\gamma$ , and the implied weights on the lagged gammas  $\Gamma_{t-l,d}$ . Standard errors of the average coefficient estimates are in parentheses below the coefficient estimates. These standard errors are computed as the standard deviation of the 36 monthly estimates, divided by the square root of 36. Column (1) displays the results for a version of the

model that does not include the lagged gammas. Columns (2) and (3) present the results of specifications that include either one or five lags of gamma, respectively. To assess the statistical performance of the models with lagged gammas in these two columns, for each month, and each model, we compute a likelihood ratio test statistic for the test of the hypothesis that the coefficients  $m$  and  $\gamma$  are jointly zero. The last row of the table displays the averages across the 36 months of these likelihood ratio test statistics.

Column (4) displays the results of a specification in which the initial variance is not allowed to differ across the days of the month. This specification is more restrictive than that in column (3). This column does not include the likelihood ratio test statistic, because the specification in this column does not nest the model in column (1).

The averages of the coefficient estimates on  $r_{t-1,d}^2$  and  $\sigma_{t-1,d}^2$  are approximately 0.044 and greater than 0.9, which are unsurprising values for GARCH(1,1) models estimated using index returns. The standard errors of these averages indicate that there is relatively little variation in the estimates of these coefficients across months, which is also unsurprising.

Turning to the estimates of the coefficient  $\gamma$ , the average across the 36 months is negative for all models, indicating a negative relation between OMM net gamma and the variance of index returns. The standard errors indicate that there is not much variability in these estimates. These results are consistent with the prediction that delta-hedging activity by market makers in the S&P 500 index options market impacts the variance of the index at the minute level. This impact persists, as evidenced by the non-negative weights  $\phi_1, \dots, \phi_5$ , and the effect diminishes with longer lags, as expected.

The likelihood ratio test statistics in the last row are tests of the hypotheses that the coefficients  $m$  and  $\gamma$  are zero. Under the null hypotheses that these coefficients are zero, these test statistics are distributed  $\chi^2$ , with two degrees of freedom. The 5%, 2.5%, and 1% critical values of the  $\chi_2^2$  test statistic are 5.99, 7.38, and 9.21, respectively. The average values of the test statistic in columns (2) and (3) are 8.94 and 9.09, providing on average some statistical evidence against the hypothesis that  $m = \gamma = 0$ . Thus, the likelihood ratio tests provide evidence against the hypothesis of no relation between index return variance and OMM gamma.

These likelihood ratio test statistics also reveal that, based on this statistical criterion, there is little benefit to adding lags of gamma beyond the first lag. To see this, note that the difference between the likelihood ratio test statistics for the models in columns (2) and (3)

provides a test of the hypothesis that adding additional lags of gamma improves the performance of the model. Because the “Beta” lag structure we borrow from Engles, Ghysels, and Sohn (2013) involves two free parameters, moving from the model in column (2) to that in column (3) relaxes two restrictions, and, for each month, the difference between the likelihood test statistics is distributed  $\chi_2^2$  under the null hypothesis. The difference in the average values of the test statistics in columns (3) and (2) is  $9.09 - 8.94 = 0.15$ , far from to the 5% critical value, suggesting that few of the 36 test statistics included in the average exceed the 5% critical value.

### 3.3 Linear models explaining squared returns

We explore whether the results are robust to the choice of statistical model by estimating linear models that allow variance to depend on the aggregate OMM net gamma. These models relate the squared return to lagged squared returns and lagged gammas, as follows:

$$r_{t,d}^2 = \omega + \sum_{l=1}^5 \alpha_l r_{t-l,d}^2 + \sum_{l=1}^L \gamma_l \Gamma_{t-l,d} + \delta_d + \sum_{h=1}^7 \theta_h I(t \in A_h) + e_{t,d}, \quad (6)$$

where  $r_{t,d}$  is the log return on the e-mini futures contract with highest trading volume during minute  $t$  of day  $d$ ,  $\Gamma_{t-l,d}$  denotes the gamma of the aggregate OMM net position at minute  $t-l$  of date  $d$ ,  $\delta_d$  is a date fixed effect, and  $\theta_h I(t \in A_h)$ , for  $h = 1, \dots, 7$ , are fixed effects for  $t$  in each of the intervals  $A_1 = 9:30 - 10:30$  a.m.,  $A_2 = 10:30-11:30$  a.m., ..., and  $A_7 = 3:30 - 4:00$  p.m. The estimation follows that of the GARCH model. We estimate the linear model in equation (6) with one-minute data for each month, using the same scaling we use in estimating the GARCH model. We obtain model parameters for each month from July 2020 to June 2023 and compute average parameter values across 36 months. Our primary prediction is that the  $\gamma$  coefficients are negative, in line with the GARCH model results.

We report the results of the linear regression models in Table 4. The table presents averages across months of point estimates and standard errors for these averages, each of which is computed from the 36 monthly parameter estimates. Column (1) displays estimates without net gamma position variables. Additionally, for each additional model in columns (2)-(6) we include lagged values of the gamma of the aggregate OMM net position. To assess the statistical performance of the model that incorporates lagged net gammas, we also report the average F-statistic relative to the model in the first column. These F-statistics have one degree of freedom for model (2) to five degrees of freedom for model (6).

The results for these linear models are consistent with those of the GARCH models. For all models, the average values of the gamma coefficients are negative for all lags. For the models in columns (2)-(6) that successively include an additional lagged gamma, the average values of the F-statistics in columns (2) and (3) exceed the 1% critical value of the  $\chi^2_2$  distribution, while those in columns (4)-(6) exceed the 5% critical value.

#### 4. Estimating the impact of OMM gamma on market return variance

We now turn to using the GARCH and linear models to estimate the impact of OMM gamma on the variance of S&P 500 index returns. One cannot do this by computing the partial derivatives of variance with respect to the lagged gammas in the GARCH model or by examining the coefficients on lagged gammas in the linear model, because the lagged gamma time  $t$  impacts not only the time  $t$  conditional variance but also all future variances through the dynamic variance models. For example, in the GARCH model, at time  $t$  lagged gamma  $\Gamma_{t-1,d}$  impacts  $\sigma_t^2$  and thus  $r_t$  directly through the coefficient  $\gamma$ . Then  $\sigma_{t,d}^2$  and  $r_{t,d}$  impact  $\sigma_{t+1,d}^2$  and  $r_{t+1,d}$  through the coefficients  $\beta$  and  $\alpha$ ; these in turn impact  $\sigma_{t+2,d}^2$  and  $r_{t+2,d}$ , which in turn impact  $\sigma_{t+3,d}^2$  and  $r_{t+3,d}$ , and so forth. The process continues until the impact of  $\Gamma_{t-1,d}$  eventually decays. Through this process, the variance at any time  $t$  is impacted by many lags of gamma, with the magnitude of the impact determined by the coefficients on the lagged gammas and the other model parameters in a complicated way.

Similarly, in the linear model, lagged gamma  $\Gamma_{t-1,d}$  impacts  $r_t^2$  directly through the coefficient  $\gamma_1$ . Then  $r_t^2$  impacts  $r_{t+1}^2$  through the coefficient on  $r_t^2$ , and  $r_{t+1}^2$  in turn impacts  $r_{t+2}^2$ , which in turn impacts  $r_{t+3}^2$ , and so forth. Again, the process continues until the impact of  $\Gamma_{t-1,d}$  on squared returns eventually decays.

We estimate the impact of gamma on return variance by simulating the GARCH and linear models to generate simulated returns or squared returns under the counterfactual assumption that the lagged gammas are all zero, or equivalently, under the assumptions that the coefficients  $m$  and  $\gamma$  are all zero. These simulations yield time series of simulated squared returns at one-minute frequency based on the counterfactual assumption that OMM gammas do not impact variance. We then use the simulated squared returns to compute realized variance over horizons of one day and 30 minutes, and compare them to the realized variances computed from

the actual observed one-minute returns. The differences between the realized variances based on actual and simulated returns provide measures of the impact of market maker gamma on returns.

#### 4.1 Simulating the GARCH models

We begin by constructing a comprehensive measure of daily total variability, which includes the price impact induced by OMM gammas. This involves generating one-minute returns derived from bid-ask midpoint futures prices, followed by the computation of the realized variance for each trading day and for each half hour during the day. The realized variance captures the total variability observed in the market, reflecting both the volatility due to fundamental factors (for example, news) and any additional volatility arising from market dynamics, including the impact of hedge rebalancing by OMMs.

Letting  $t$  index the minute during regular trading hours on a day  $d$ , the realized variance (during regular trading hours) for day  $d$  is

$$RV_d = \sum_{t=1}^n r_{t,d}^2, \quad (7)$$

where  $r_{t,d}$  represents the one-minute return at minute  $t$  of date  $d$ , and  $n$  is the total number of one-minute intervals during regular trading hours.

We identify the effect of gamma by simulating a gamma-free realized variance. For each minute, we compute an estimate  $\hat{\varepsilon}_{t,d}$  of the GARCH model innovation  $\varepsilon_{t,d}$ . This is done by recovering the estimated innovation  $\hat{\varepsilon}_{t,d}$  from equation (1) using the actual return and the model estimate of the conditional standard deviation, that is,  $\hat{\varepsilon}_{t,d} = z_{t,d} / (\hat{\sigma}_{t,d})$ . We then use these innovations to simulate one-minute returns,  $\hat{r}_{t,d}$ , under the counterfactual assumption that returns are not influenced by the gammas. The simulation involves multiplying the recovered one-minute innovations by the estimated conditional variance from the GARCH model in Table 3 column (3), while setting the coefficients associated with lagged values of net gamma,  $m$  and  $\gamma$ , to zero. This removes the influence of gamma from the returns, thereby yielding estimated gamma-free returns.

Then, we use the simulated returns to compute a gamma-free realized variance,

$$RV_d^{GF} = \sum_{t=1}^n \hat{r}_{t,d}^2, \quad (8)$$

for each day  $d$ . The impact of gamma on index volatility is then estimated by comparing the annualized gamma-free realized variance and the observed realized variance. The day  $d$  impact is calculated as the difference between the annualized gamma-free realized volatility and the realized volatility:

$$Impact_d = \sqrt{252 \times RV_d} - \sqrt{252 \times RV_d^{GF}}. \quad (9)$$

With this definition of the impact, positive impacts mean that the actual realized volatility, which reflects the impact of OMM delta hedging, exceeds the counterfactual volatility computed under the assumption that OMM delta hedging does not impact returns.

We use the same approach to compute the gamma impacts over shorter, 30-minute intraday periods, allowing for a more granular analysis of gamma's influence on market volatility throughout the trading day.

In Table 5 Panel A we report summary statistics describing the distribution of the impacts of OMM gamma annualized daily volatility and annualized 30-minute volatility. The average daily and 30-minute impacts are negative for both horizons, with a means of  $-0.19$  and  $-0.17$  percentage points for the daily and 30-minute horizons, respectively. That is, on average, observed realized volatility, which reflects the impact of OMM delta hedging, is less than what it would have been in the absence of delta hedging. This occurs because return variance is negative related to OMM gamma, and OMM gamma is most often positive (Table 2). However, the 90<sup>th</sup> percentiles of the distributions of impact are positive for both the daily and 30-minute horizons. The interpretation of this result is that OMM delta hedging increases volatility during both 10% of the days and 10% of the 30-minute periods.

The first row in the right-hand column of Table 5 Panel A shows that the largest impact on daily volatility is to increase annualized volatility by 3.3 percentage points. The second row in the right-hand column shows that the greatest impact on 30-minute volatility is to increase annualized volatility by 6.4 percentage points. This impact of 6.4 percentage points is an outlier—the 99<sup>th</sup> percentile is an increase of only 2.0 percentage points.

Figure 3 displays the daily time series of impacts on the annualized daily volatilities. These impacts are generally negative, consistent with the results that the OMM gammas are typically positive. Comparing the daily gamma-induced impacts in Figure 3 to the average daily

gammas in Figure 1, one can see that the gamma-induced impacts tend to be positive when the gammas are negative, which is expected.

Panels A and B of Figure 4 displays histograms of the distributions of the daily and 30-minute impacts.

### 5.2 Simulating the linear models

We follow a similar approach to use the linear model in Table 4 column (6) to develop alternative estimates of the impact of gamma on the realized volatility of the S&P 500 index. We first recover estimates of the residual  $e_t$  using the relation

$$\hat{e}_{t,d} = r_{t,d}^2 - \hat{\omega} - \sum_{l=1}^5 \hat{\alpha}_l r_{t-l,d}^2 - \sum_{l=1}^L \hat{\gamma}_l \Gamma_{t-l,d} - \delta_d - \sum_{h=1}^7 \theta_h I(t \in A_h), \quad (10)$$

where parameters with “hats,” for example,  $\hat{\omega}$ , are estimates. We then set the coefficients on the lagged market maker gammas,  $\gamma_1$  through  $\gamma_5$ , to zero, and use equation (10) and the recovered innovations  $\hat{e}_{t,d}$  to simulate the squared return  $\hat{r}_{t,d}^2$ . This process allows us to remove the influence of gamma on the squared returns, thereby yielding estimated gamma-free squared returns. We then use the simulated return and equation (8) to compute a gamma-free realized variance. The impact of gamma on S&P 500 index volatility is calculated using equation (9) at daily and 30-minute frequencies.

Table 5 Panel B reports the impact of OMM delta hedging on daily volatility and 30-minute volatility. Overall daily impact is on average negative for both horizons, with means of  $-2.1$  and  $-2.3$  percentage points for the daily and 30-minute impacts. On average, gamma-free realized volatility is larger than the observed realized volatility. These means, and most of the percentiles displayed, are of similar magnitude to the GARCH model estimates in Panel A. The maximum impact for the 30-minute horizon in Panel B of 7.0 of percentage points is similar to the corresponding maximum in Panel A of 6.4 percentage points.

Panels A and B of Figure 6 displays histograms of the distributions of the daily and 30-minute impacts based on the linear model. As in Figure 4, the right-hand tail of the distribution in Panel B is truncated, in that the maximum increase of 7.0 percentage points is not displayed.

## 6. Discussion

Should one think of a gamma-induced increase in annualized daily realized volatility of 3.3 percentage points as large or small? Should one think of a gamma-induced increase in annualized 30-minute realized volatility of 6.4 percentage points as large or small?

The distributions of changes in annualized daily and 30-minute realized volatilities provide benchmarks for thinking about these questions. For example, suppose daily changes in annualized realized volatility of about 3.3 percentage points or greater occur frequently, for whatever reason. If this is the case, then market participants are likely used to dealing with such changes, and the fact that OMM delta hedging can increase annualized realized volatility by 3.3 percentage points might not be a source of concern. But if increases of about 3.3 percentage points are rare, then the fact that OMM delta hedging can cause such an increase might merit more attention.

Figure 7 Panel A displays the histogram of daily changes in annualized daily realized volatility. To construct this histogram, we use one-minute log futures returns to calculate the realized variance during regular trading hours, annualize the realized variance by multiplying by 252, and take the square root to convert the result to a volatility. Then we subtract each day's annualized realized volatility from the same quantity on the previous day, and construct the histogram displayed in Figure 7 Panel A. Table 6 presents various percentiles of the distribution of daily changes in annualized daily realized volatility, along with the mean, standard deviation, minimum, and maximum of the daily changes.

The standard deviation of daily changes in annualized realized volatility is 4.5 percentage points, and the change in annualized realized volatility from the previous day exceeds three percentage points on about 20% of the trading days. Thus, changes in annualized realized volatility of about three percentage points or greater occur reasonably frequently, on about 20% of the days, or on average about once per week. Given this, our interpretation of the results is that a gamma-induced change of 3.0 percentage points is not large.

Figure 7 Panel B displays the histogram of changes in annualized 30-minute realized volatilities. For this figure, each daily change is the difference between the annualized realized volatility during a 30-minute period, say 10:00-10:30 am, and the previous 30-minute period, 9:30-10:00 am. The distribution does not display the tails of the distribution. Specifically, it displays changes between  $-20$  and  $20$  percentage points, and there are a few changes outside that

range. Table 6 includes several percentiles of the distribution of these daily changes, along with the mean, standard deviation, minimum, and maximum of the daily changes. It shows that the minimum and maximum changes are  $-52.2$  and  $63.4$  percentage points.

Table 6 shows that the standard deviation of the daily changes in annualized 30-minute realized volatility,  $4.6\%$ , and the 1<sup>st</sup> and 99<sup>th</sup> percentiles of the daily changes in annualized 30-minute realized volatilities are  $-11.6$  and  $14.2$  percentage points. Using these as benchmarks, the maximum gamma-induced change of  $6.4$  percentage points does not seem large.

## **6. Conclusion**

We estimate the maximum impact of OMM gamma on S&P 500 index volatility during the period running from July 2020 through June 2023. Specifically, we estimate the maximum impact of OMM gamma on market volatility over one-day and 30-minute periods, as compared to a counterfactual situation in which OMM gamma does not impact market volatility.

We use the GARCH models to generate simulated returns under the counterfactual assumption that the return variance is unrelated to the lagged gammas. We then use the simulated returns to compute realized variance for each trading day, and compare these counterfactual realized variances to the realized variances computed from the actual one-minute returns. The differences between the realized variances based on actual and simulated returns provide measures of the impact of market maker gamma on returns. We use the same approach to compute simulated and actual realized variances for non-overlapping 30-minute periods starting 9:30 am each day.

The results show that the typical impact of gamma is to reduce volatility, with the mean impact being to reduce volatility by  $0.2$  percentage points. But OMM gamma is sometimes negative, and at these times volatility is elevated. We find that the maximum impact of OMM gamma on annualized daily volatility is to increase annualized daily volatility by  $3.3$  percentage points. The maximum impact of OMM gamma on annualized 30-minute volatility is to increase it by  $6.4$  percentage points.

## Appendix. Delta-hedge rebalancing trades and gamma

The delta ( $\Delta$ ) of an options position is the derivative of the position value with respect to the underlying price  $S$ , that is,  $\Delta = \partial V / \partial S$ , where  $V$  is the theoretical value of the position. Options market makers often buy or sell the underlier to make the net delta of the combined position of options and underlier be within an acceptable range, that is, to be within some neighborhood of zero. This entails holding a position in the underlier such that the number of units of the underlying asset  $n$  is approximately equal to the negative of the delta of the options positions, i.e. such that  $n \approx -\Delta$ . Under the assumption that market makers hold such (approximately) delta-neutral positions, and rebalance their delta hedges to maintain (approximate) delta neutrality, the hedge rebalancing trades are determined by the negative of the changes in the delta of the options position. We consider the change in delta.

### A.1 Analysis of delta-hedge rebalancing trades

We carry out the analysis for a single option. Because the gamma of an options portfolio is the sum of the gammas of the positions that comprise the portfolio, and the gamma of the aggregate portfolio held by all OMMs collectively is the sum of the gammas of the portfolios of the several OMMs, the same conclusion holds for both the options portfolio held by each OMM and for the aggregate OMM portfolio.

Assume that the value of an option  $V(S, t)$  is a function of the underlying price  $S$  and time  $t$ ; this implies that the delta  $\Delta(S, t)$  is also a function of the underlying price and time. Suppose the movement of the underlying price is described by the stochastic differential equation

$$dS = \mu S dt + \sigma S dW, \quad (11)$$

where  $\mu$  is the expected return on the underlier and  $W$  is a standard Brownian motion. Itô's lemma implies that the dynamics of  $\Delta$  are

$$d\Delta(S, t) = \frac{\partial \Delta}{\partial S} dS + \frac{1}{2} \frac{\partial^2 \Delta}{\partial S^2} (dS)^2 + \frac{\partial \Delta}{\partial t} dt \quad (12)$$

$$= \frac{\partial \Delta}{\partial S} dS + \frac{1}{2} \frac{\partial^2 \Delta}{\partial S^2} (\sigma^2 S^2 dt) + \frac{\partial \Delta}{\partial t} dt \quad (13)$$

$$= \Gamma dS + \frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt. \quad (14)$$

Examination of equation (9) suggests why industry practitioners are sometimes interested in speed ( $\partial \Gamma / \partial S = \partial^2 \Delta / \partial S^2 = \partial^3 V / \partial S^3$ ) and charm ( $\partial \Delta / \partial t = \partial^2 V / \partial S \partial t$ ), as these terms contribute to the change in delta, and thus to the required hedge rebalancing trade. The sum of the first two

terms  $\frac{\partial \Delta}{\partial S} dS + \frac{1}{2} \frac{\partial^2 \Delta}{\partial S^2} (dS)^2$  provides an estimate of how the required delta hedge changes when the stock price changes, with the speed term capturing the impact of large stock price changes on  $\Delta$ . The third term  $\frac{\partial \Delta}{\partial t} dt$  provides an estimate of how the required delta hedge changes as time passes, if the stock price does not change.

As we illustrate below, for 0DTE options the terms  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt$  and  $\frac{\partial \Delta}{\partial t} dt$  are larger than  $\Gamma dS$  for reasonable values of the volatility  $\sigma$  and other parameters. However, these terms almost offset each other, so that their sum is small. In addition, their sum is proportional to the option gamma, implying that hedge rebalancing trades given by (12)-(14) are proportional to gamma.

To show that the sum  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt$  is both proportional to the gamma and small, we need to exploit the partial differential equation that appears in option valuation. The version of the Black-Scholes formula with a continuous dividend yield, originally due to Merton (1973), is widely used for the valuation and hedging of index options. However, it is not helpful in the current context. The Merton (1973) assumption of a continuous dividend yield is a modelling convenience that is very useful at some time scales, but not useful intraday. At intraday time scales, one must recognize that dividend payments are discrete, and stocks go ex-dividend during the overnight period; there are no intraday dividends.<sup>6</sup> Once one recognizes that dividends are discrete, the option value is the solution of a system of partial differential equations, connected by time boundary conditions across the ex-dividend times.<sup>7</sup>

Intraday, option values satisfy the Black-Scholes (1973) partial differential equation

$$\frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS \frac{\partial V}{\partial S} + \frac{\partial V}{\partial t} - rV = 0. \quad (16)$$

Differentiate this partial differential equation with respect to  $S$ , yielding

$$\frac{1}{2} \sigma^2 S^2 \frac{\partial^3 V}{\partial S^3} + \sigma^2 S \frac{\partial^2 V}{\partial S^2} + rS \frac{\partial^2 V}{\partial S^2} + r \frac{\partial V}{\partial S} + \frac{\partial^2 V}{\partial S \partial t} - r \frac{\partial V}{\partial S} = 0. \quad (17)$$

Rearranging, the sum  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt$  is given by

$$\frac{1}{2} \sigma^2 S^2 \frac{\partial^3 V}{\partial S^3} + \frac{\partial^2 V}{\partial S \partial t} = -\sigma^2 S \frac{\partial^2 V}{\partial S^2} - rS \frac{\partial^2 V}{\partial S^2} - r \frac{\partial V}{\partial S} + r \frac{\partial V}{\partial S} \quad (18)$$

<sup>6</sup> The lack of a flow of dividends during the intraday period explains why equations (11) and (16) do not include a dividend yield.

<sup>7</sup> See, for example, Wilmott (1998), Section 47.10.

$$= -(\sigma^2 + r)S \frac{\partial^2 V}{\partial S^2}. \quad (19)$$

Substituting into equation (11), the hedge rebalancing trades are proportional to  $\Gamma$ :

$$d\Delta(S, t) = \Gamma(dS - (\sigma^2 + r)Sdt). \quad (20)$$

The second term  $\Gamma(\sigma^2 + r)Sdt$  is clearly small relative to  $\Gamma dS$ .

### A.2 Illustrative calculations

On the expiration date, no dividends remain and the Black-Scholes formula solves equation (13). The gamma, speed, and charm are

$$\Gamma = \frac{1}{S\sigma\sqrt{\tau}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{d_1^2}{2}\right), \quad (21)$$

$$\text{speed} = -\frac{1}{S^2\sigma\sqrt{\tau}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{d_1^2}{2}\right) \left(\frac{d_1}{S\sigma\sqrt{\tau}} + 1\right) = -\frac{\Gamma}{S} \left(\frac{d_1}{S\sigma\sqrt{\tau}} + 1\right) \quad (22)$$

$$\text{charm} = -\frac{1}{S^2\sigma\sqrt{\tau}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{d_1^2}{2}\right) \frac{2r\tau - d_2\sigma\sqrt{\tau}}{2\tau\sigma\sqrt{\tau}}, \quad (23)$$

where  $d_1 = \frac{\ln(\frac{S}{K}) + (r + \frac{\sigma^2}{2})\tau}{\sigma\sqrt{\tau}}$ ,  $d_2 = \frac{\ln(\frac{S}{K}) + (r - \frac{\sigma^2}{2})\tau}{\sigma\sqrt{\tau}}$ ,  $t = T - t$  is the remaining time to expiration, and  $N(\bullet)$  is the standard normal distribution function.

Using these formulas for gamma, speed, and charm, we compute the three terms  $\Gamma dS$ ,  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt$ , and  $\frac{\partial \Delta}{\partial t} dt$  for a call option on a single share with  $K = 100$ ,  $97 \leq S \leq 103$ ,  $dt = 1$  minute  $= 1/(60 \times 24 \times 365)$  year,  $dS$  equal to the one-minute standard deviation, that is  $dS = \sigma\sqrt{dt} = 0.4\sqrt{dt}$ , and four values of  $\tau$ , these being  $\tau = 6.75$  hours  $= 6.75/(24 \times 365)$  year,  $\tau = 3$  hours  $= 3/(24 \times 365)$  year,  $\tau = 1$  hour  $= 1/(24 \times 365)$  year, and  $\tau = 0.25$  hour  $= 0.25/(24 \times 365)$  year. The interest rate is  $r = 0.04$ . We also compute the sum of the second and third terms,  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt$ .

The solid black and red lines in Figure A1 Panel A display the values of  $\Gamma dS$  and the sum  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt$ , for  $\tau = 6.75$  hours  $= 6.75/(24 \times 365)$  year. The sum  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt$  is much smaller than the term  $\Gamma dS$ ; at the peaks of both, it is only about 6.2% as large as  $\Gamma dS$ . The two dashed lines in Panel A display the values of  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt$  and  $\frac{\partial \Delta}{\partial t} dt$ , for  $\tau = 6.75$  hours.

These terms are larger than  $\Gamma dS$ , but the two terms offset each other, resulting in a small value of

the sum  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt$ . While it is difficult to see because the red line is close to zero, examination of the figure reveals that the term  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt$  displayed in red appears to be proportional (with a negative coefficient) to the values of the gamma term displayed in black. Panels B-D display the values of the same terms for  $\tau = 3$  hours =  $3/(24 \times 365)$  year,  $\tau = 1$  hour =  $1/(24 \times 365)$  year, and  $\tau = 0.25$  hour =  $0.25/(24 \times 365)$  year. In these figures the term  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt$  displayed in red also are proportional (with a negative coefficient) to the gamma term displayed in black.

### A.3 The gamma of the aggregate market maker position

The analysis above is for a single option series. An OMM's net delta across all series is the linear combination

$$\Delta = \sum_{j=1}^J n_j \Delta_j, \quad (24)$$

where  $n_j$  is the number of contracts of series  $j$ ,  $\Delta_j$  is the delta of option series  $j$ , and  $\Delta$  is the delta of the aggregate position across all series. The net hedge rebalancing trade is then

$$d\Delta = \sum_{j=1}^J n_j d\Delta_j \quad (25)$$

$$= \sum_{j=1}^J n_j \Gamma_j (dS - (\sigma^2 + r)S dt) \quad (26)$$

$$= \Gamma (dS - (\sigma^2 + r)S dt), \quad (27)$$

where  $\Gamma_j$  is the gamma of series  $j$  and  $\Gamma = \sum_{j=1}^J n_j \Gamma_j$  is the gamma of the OMM's options portfolio.

A similar analysis leads to the conclusion that that aggregate hedge rebalancing trade across all OMMs is proportional to the gamma of the aggregate OMM position.

## References

- Adams, Greg, Jean-Sebastien Fontaine, and Chayawat Ornthanalai. 2024. The market for 0-days-to-expiration: The role of liquidity providers in volatility attenuation. Working Paper, Bank of Canada and University of Toronto.
- Almeida, Caio, Gustavo Freire, and Rodrigo Hizmeri. 2024. 0DTE asset pricing. Working Paper, Princeton University.
- Baltussen, Guido, Zhi Da, Sten Lammers, and Martin Martens, 2021, Hedging demand and market intraday momentum, *Journal of Financial Economics* 142, 377–403.
- Bandi, Frederico M., Nicola Fusari, Roberto Reno. 2024. 0dte option pricing. Working paper, Johns Hopkins University.
- Barbon, Andrea, Heiner Beckmeyer, Andrea Buraschi, and Mathis Moerke. 2022. Liquidity Provision to Leveraged ETFs and Equity Options Rebalancing Flows: Evidence from End-of-Day Stock Prices Swiss Finance Institute Research Paper Series No. 22-40
- Bank of America. 2023. The truth on 0dte intraday gamma imbalances, Bank of America Global Research Paper.
- Beckmeyer, Heiner, Nicole Branger, and Leander Gayda, 2023, Retail traders love 0dte options... but should they? Working paper, University of Munster.
- Black, Fischer, and Myron Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, No. 3, 637-654.
- Brogaard, Jonathan, Jaehee Han, and Peter Young Won. 2024. Does 0DTE Options Trading Increase Volatility? Working Paper, University of Utah.
- Dim, Chukwuma, Bjorn Eraker, and Grigory Vilkov, 2024, 0dtes: Trading, gamma risk and volatility propagation, Working paper, George Washington University.
- Engle, R.F., Ghysels, E. and Sohn, B., 2013. Stock market volatility and macroeconomic fundamentals. *Review of Economics and Statistics*, 95, No. 3, 776–797.
- Engle, R.F. and Sokalska, M.E., 2012. Forecasting intraday volatility in the US equity market. Multiplicative component GARCH. *Journal of Financial Econometrics*, 10, No. 1, 54–83.
- Hu, J., A. Kirilova, and D. Muravyev. 2023. Options Market Makers. Working Paper, Singapore Management University
- Kilburn, Faye. 2024. After the selloff, competing theories on dealer gamma. *Risk.net* (August 12). Available at [https://www.risk.net/markets/7959811/after-the-selloff-competing-theories-on-dealer-gamma?check\\_logged\\_in=1](https://www.risk.net/markets/7959811/after-the-selloff-competing-theories-on-dealer-gamma?check_logged_in=1)

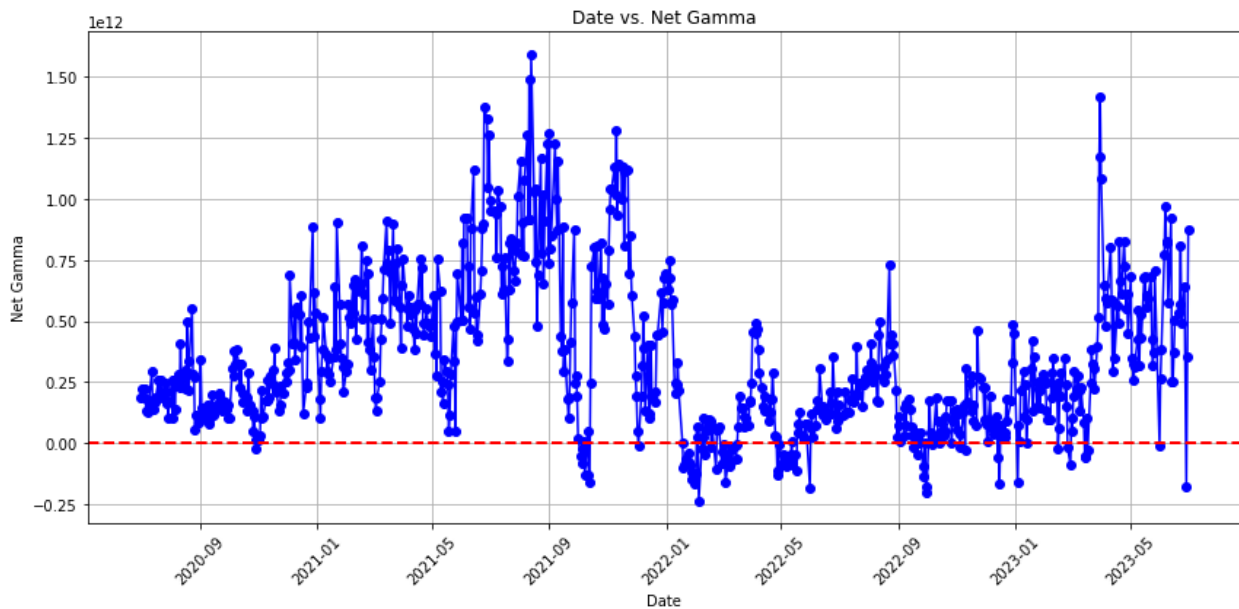
Lipson, Marc Lars, Davide Tomio, and Jiang Zhang, 2023, A real cost of free trades: Retail option trading increases the volatility of underlying securities, Working paper, Darden School of Business.

Merton, Robert C. 1973. Theory of Rational Option Pricing. *Bell Journal of Economics and Management Science* 4, No. 1,141–183.

Ni, Sophie, Neil D. Pearson, and Allen M. Poteshman, 2005, Stock price clustering on option expiration dates, *Journal of Financial Economics* 78, 49–87.

Ni, Sophie Xiaoyan, Neil D. Pearson, Allen M. Poteshman, and Joshua White. 2021. Does Option Trading Have a Pervasive Impact on Underlying Stock Prices??" *Review of Financial Studies* 34, No. 4, pp.1952–1986.

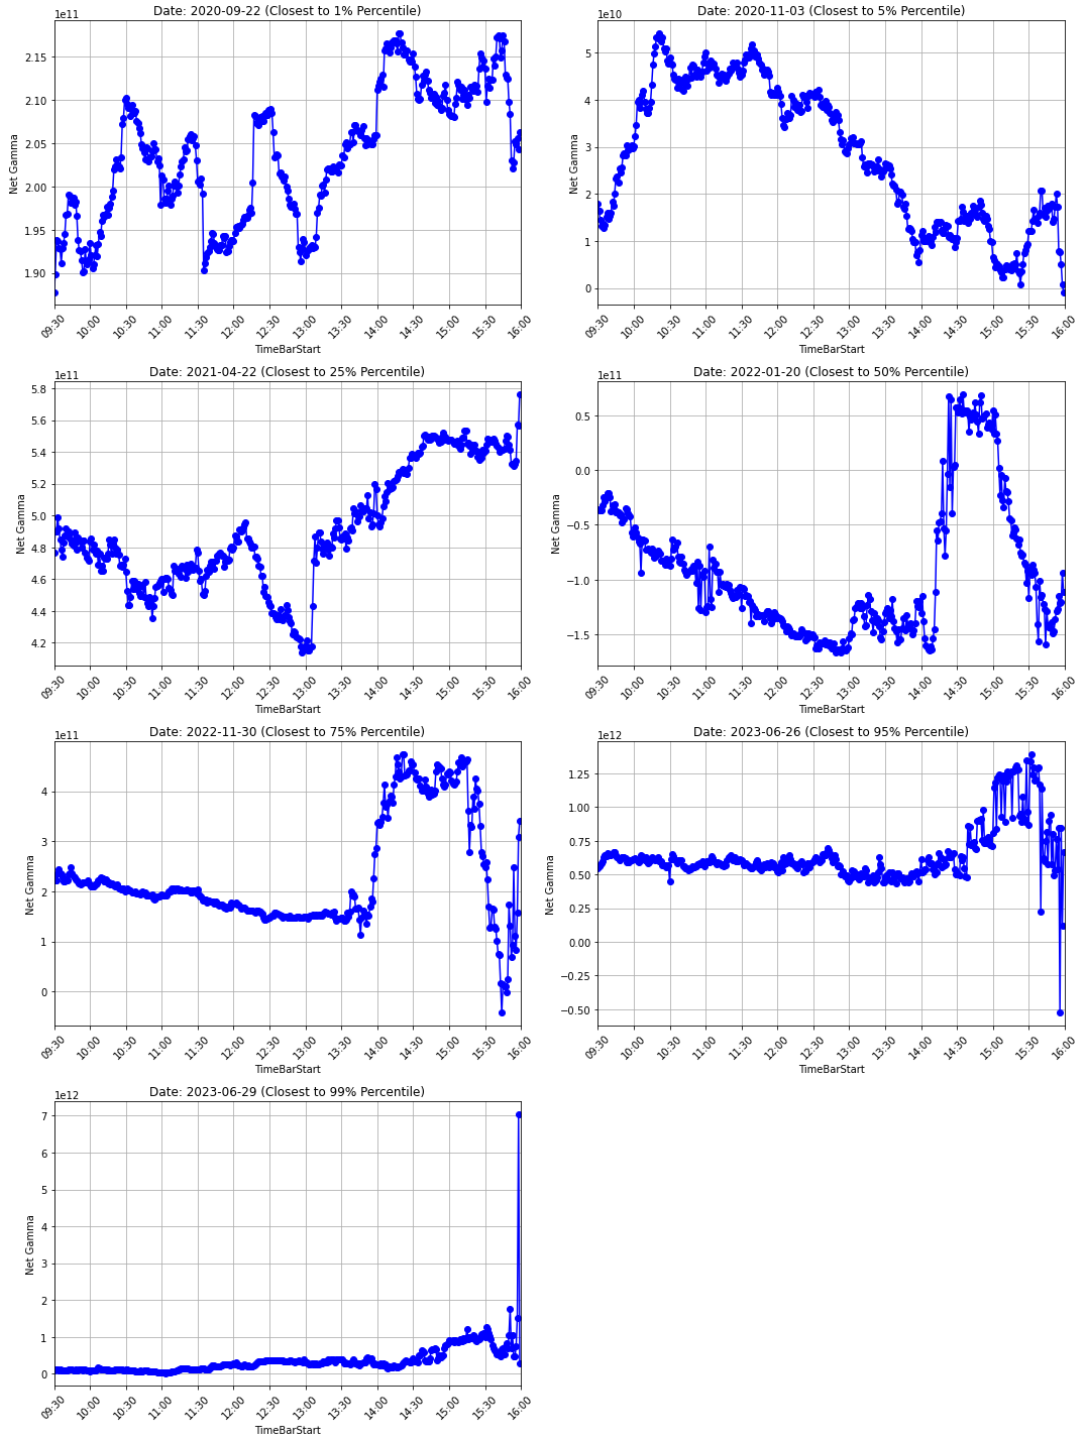
Wilmot, Paul. 1998. *Derivatives: The Theory and Practice of Financial Engineering*. New York: John Wiley & Sons, Inc.



**Figure 1**

**Time series of daily averages of gamma at one-minute frequency**


For each date, we compute the average of the gammas at one-minute frequency, resulting in a single average gamma for each date. The figure plots these daily averages. The gammas are scaled by dividing by  $10^{12}$ .



**Figure 2**

**Time series of gammas at one-minute frequency for representative days**

This figure displays the time-series of gammas at one-minute frequency for seven representative days. The seven days are those with the standard deviation of the one-minute gammas closest to the 1<sup>st</sup>, 5<sup>th</sup>, 25<sup>th</sup>, 50<sup>th</sup>, 75<sup>th</sup>, 95<sup>th</sup>, and 99<sup>th</sup> percentiles of the distributions of daily standard deviations. The gammas are scaled by dividing by  $10^{12}$ .




**Figure 3**


**GARCH model estimates of gamma-induced impacts on annualized daily realized volatilities**

This figure plots the daily time series of impacts on the annualized daily realized volatilities. The gamma-induced impact is the difference between the annualized daily realized volatility and an estimate of what the annualized daily realized volatility would be under the counterfactual assumption that the GARCH model conditional variance is unrelated to OMM gamma. The GARCH model is the version with five lags of gamma, for which average coefficient estimates are reported in column (3) of Table 3.

Panel A




Panel B



**Figure 4**

**Histograms of GARCH model estimates of gamma-induced impacts on realized variances**

Panels A and B display histograms representing the distributions of the gamma-induced impacts on annualized daily (Panel A) and 30-minute (Panel B) realized variances. The gamma-induced impact is the difference between the annualized daily or 30-minute realized volatility and an estimate of what the realized volatility would be under the counterfactual assumption that variance is unrelated to gamma. The GARCH model is the version with five lags of gamma, for which average coefficient estimates are reported in column (3) of Table 3.




**Figure 5**

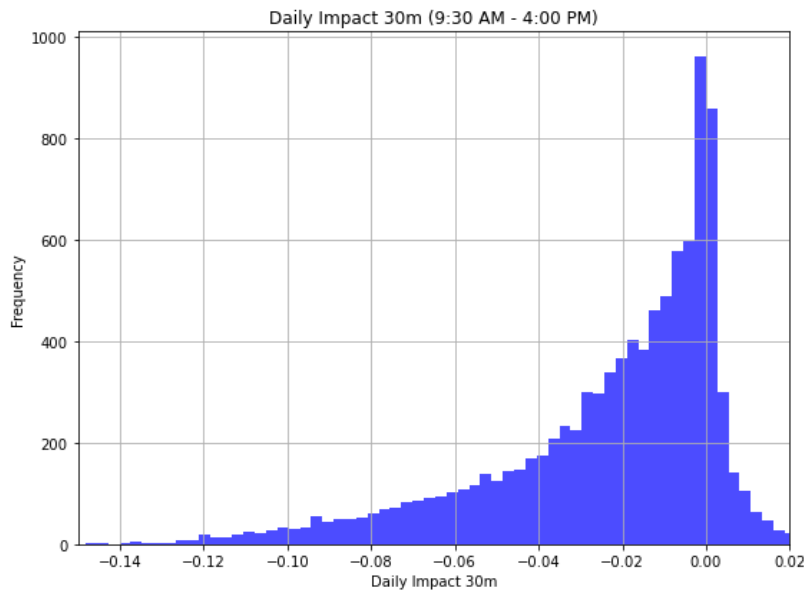
**Linear model estimates of gamma-induced impacts on annualized daily realized volatilities**

This figure plots the daily time series of impacts on the annualized daily realized volatilities. The gamma-induced impact is the difference between the annualized daily realized volatility and an estimate of what the annualized daily realized volatility would be under the counterfactual assumption that the linear model conditional variance is unrelated to OMM gamma. The linear model is the version with five lags of gamma, for which average coefficient estimates are reported in column (6) of Table 4.

Panel A



Panel B




**Figure 6**


**Histograms of linear model estimates of gamma-induced impacts on realized variances**

Panels A and B display histograms representing the distributions of the gamma-induced impacts on annualized daily (Panel A) and 30-minute (Panel B) realized volatilities. The gamma-induced impact is the difference between the annualized daily or 30-minute realized volatility and an estimate of what the realized volatility would be under the counterfactual assumption that variance is unrelated to gamma. The model is the version with five lags of gamma, for which average coefficient estimates are reported in column (6) of Table 4. The right-hand tail of the distribution in Panel B is truncated, as the increase of 0.07 is not displayed.

Panel A



Panel B

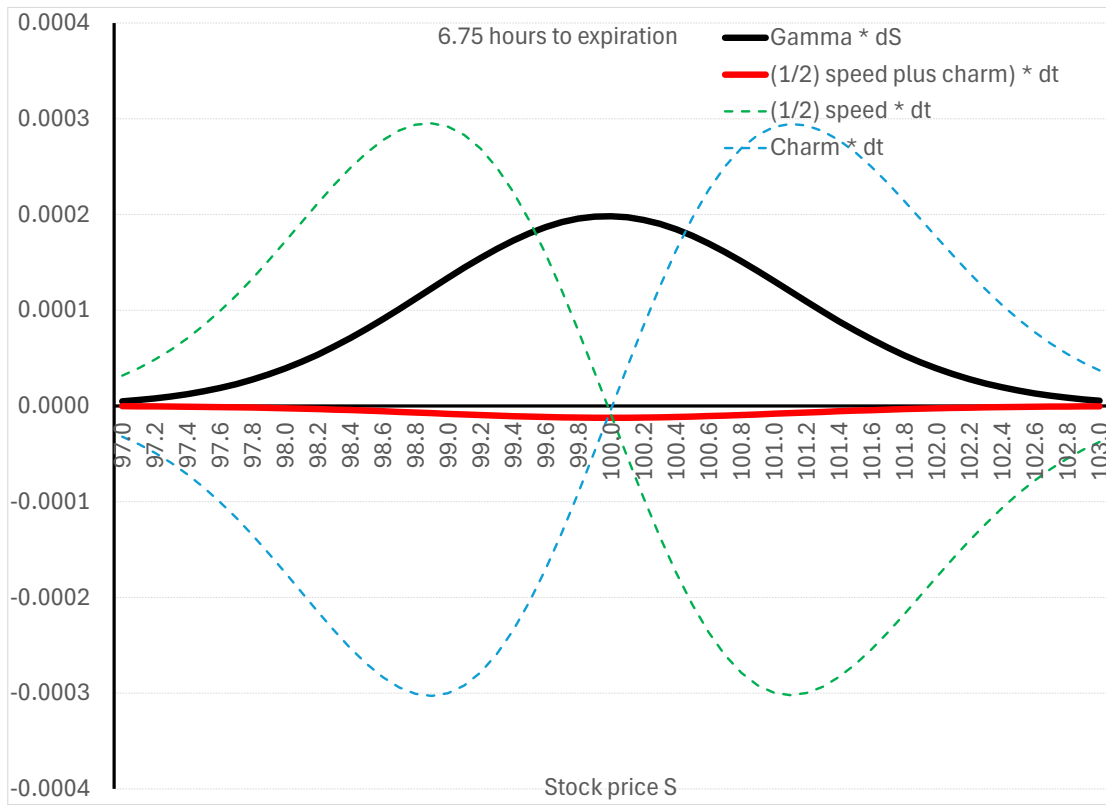


**Figure 7**


**Histograms of changes in annualized daily and annualized 30-minute realized volatility**

Panels A and B display histograms representing the distributions changes in annualized daily and annualized 30-minute realized volatility, respectively. The realized volatilities are computed using log returns on the futures contract with highest trading volume on each date. The sample period runs from July 2020 through June 2023.

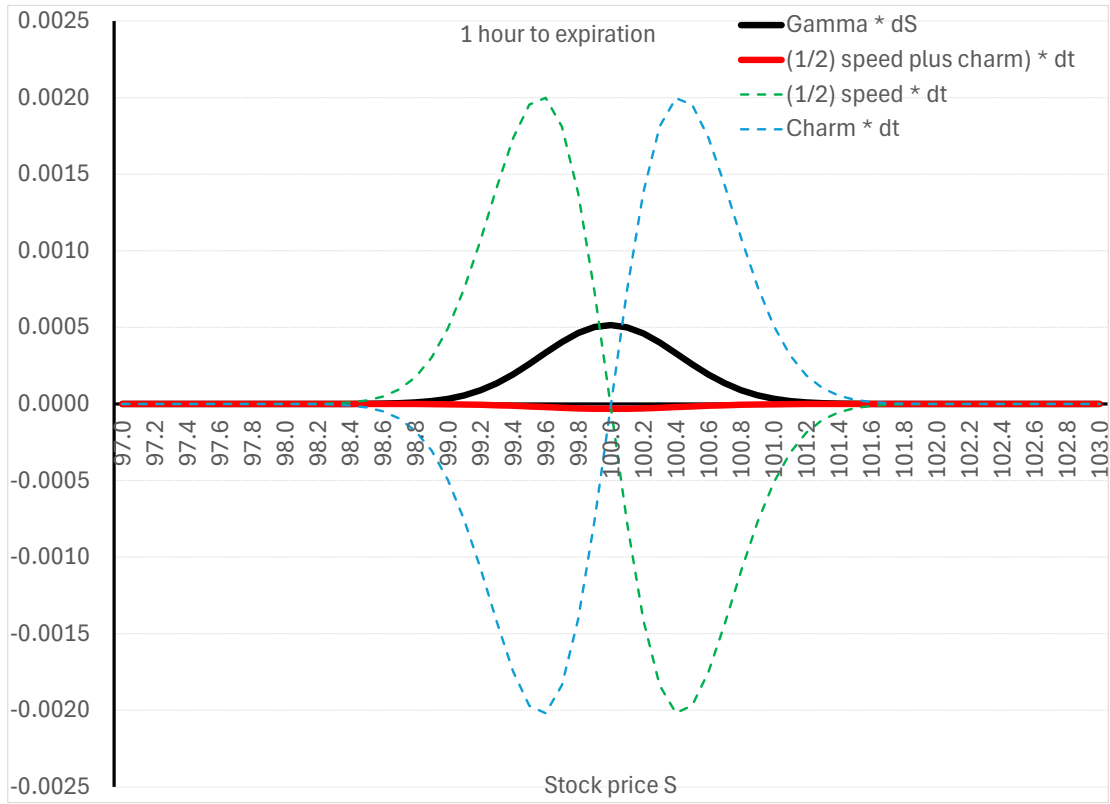
Panel A



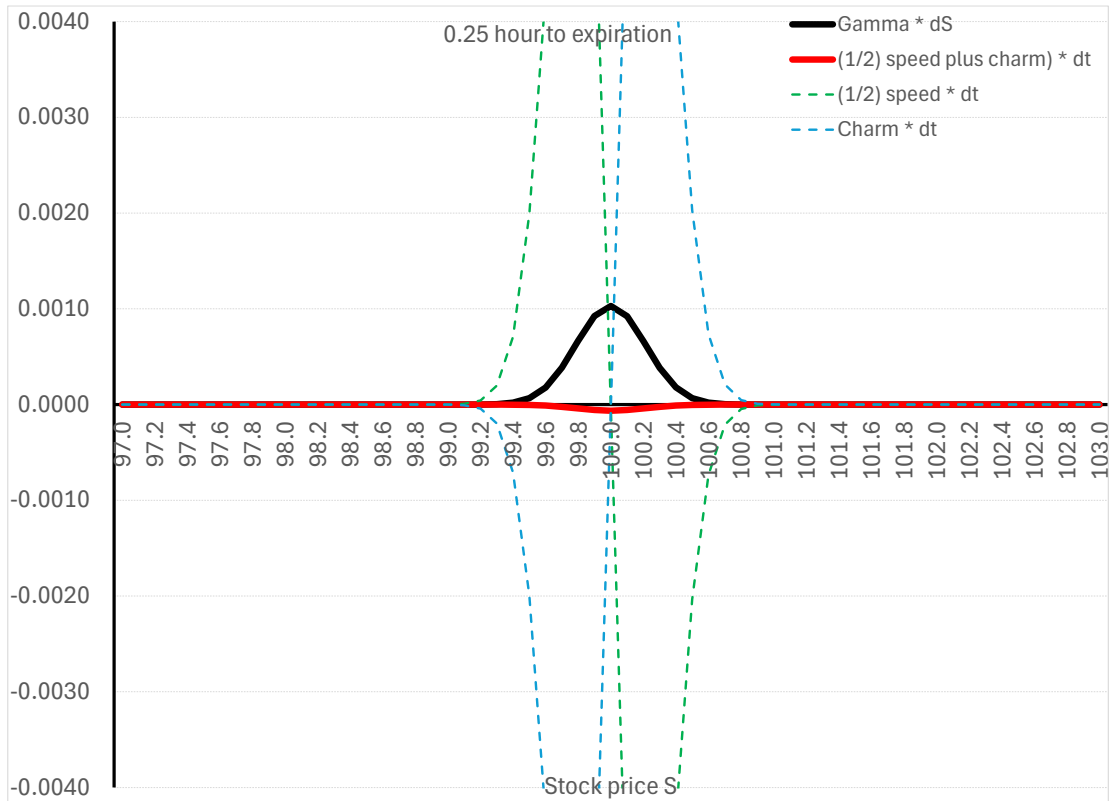
Panel B



Panel C



**Panel D**



**Figure A1**

**Terms that comprise  $d\Delta$**

In each panel, the solid black and red lines display the values of  $\Gamma dS$  and the sum  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt + \frac{\partial \Delta}{\partial t} dt$ , and The two dashed lines display the values of  $\frac{1}{2} \frac{\partial \Gamma}{\partial S} \sigma^2 S^2 dt$  and  $\frac{\partial \Delta}{\partial t} dt$ . In Panel A the time to expiration is  $\tau = 6.75$  hours =  $6.75 / (24 \times 365)$  years, in Panel B it is  $\tau = 3$  hours =  $3 / (24 \times 365)$  years, in Panel C it is  $\tau = 1$  hour =  $1 / (24 \times 365)$  years, and in Panel D it is  $\tau = 0.25$  hours =  $0.25 / (24 \times 365)$  years. The term  $dt = 1$  minute =  $1 / (60 \times 24 \times 365)$  years. The term  $dS$  is set equal to the standard deviation per minute,  $\sigma \sqrt{dt}$ . The other parameters are  $K = 100$ ,  $\sigma = 0.4$ , and  $r = 0.04$ .

**Table 1****Percentages of daily trading volume due to the buys and sells of each trading capacity**

This table shows the percentages of trading volume due to the buys and sells of each trading capacity. There are four panels, for trading in options with 0DTE, 1-5DTE, 6-21DTE, and > 21DTE. Each entry displays the percentage of total SPX and SPXW volume due to buys by the trading capacity indicated in the row label and sells by the capacity indicated in the column heading. For example, the entry in the row labelled “Customer” and column headed “Market maker” in Panel A shows that 17.5% of total SPX and SPXW option trading volume consisted of trades in which a customer bought and a market maker sold a 0 DTE option. The row labelled "Customer" and column headed "Customer" shows that 5.00% of total SPX and SPXW option trading volume consisted of trades in which a customer bought a 0 DTE option fro another customer. The row labelled "Customer" and column headed “Sum” shows that 24.65% of total SPX and SPXW option trading volume consisted of trades in which a customer bought a 0 DTE option. The sample period is January 2020 through June 2023.

		Sell					
		Customer	Market maker	Firm	Professional customer	Broker-dealer	Sum
<i>Panel A 0DTE</i>							
Buy	Customer	1.659	14.952	0.201	0.272	0.066	17.150
	Market maker	14.062	1.477	0.336	0.434	0.100	16.409
	Firm	0.171	0.419	0.050	0.014	0.011	0.665
	Professional customer	0.075	0.298	0.011	0.013	0.004	0.401
	Broker-dealer	0.021	0.095	0.012	0.015	0.003	0.146
	Sum	15.988	17.241	0.610	0.748	0.184	34.771
	<i>Panel B 1-5DTE</i>						
Buy	Customer	0.389	9.516	0.305	0.036	0.047	10.293
	Market maker	9.436	2.625	2.167	0.526	0.181	14.935
	Firm	0.389	1.973	0.275	0.034	0.058	2.729
	Professional customer	0.059	0.543	0.090	0.010	0.012	0.714
	Broker-dealer	0.103	0.302	0.107	0.014	0.018	0.544
	Sum	10.376	14.959	2.944	0.620	0.316	29.215
	<i>Panel C 6-21DTE</i>						
Buy	Customer	0.075	5.115	0.425	0.017	0.060	5.692
	Market maker	5.094	1.860	1.486	0.441	0.171	9.052
	Firm	0.347	1.644	0.251	0.049	0.057	2.348
	Professional customer	0.038	0.480	0.110	0.014	0.021	0.663
	Broker-dealer	0.067	0.270	0.103	0.021	0.061	0.522
	Sum	5.621	9.369	2.375	0.542	0.370	18.277

**Table 1 continued**

		Customer	Market maker	Firm	Sell Professional customer	Broker-dealer	Sum
<i>Panel D &gt; 21DTE</i>							
	Customer	0.061	5.252	0.996	0.042	0.061	6.412
	Market maker	5.290	2.997	3.440	1.095	0.414	13.236
	Firm	0.983	3.382	0.744	0.138	0.105	5.352
Buy	Professional customer	0.051	1.144	0.170	0.028	0.040	1.433
	Broker-dealer	0.063	0.396	0.111	0.026	0.066	0.662
	Sum	6.448	13.171	5.461	1.329	0.686	27.095

**Table 2****Statistics describing the variation in the gamma of the aggregate OMM net options position**

This table displays statistics describing the variation in the gamma of the aggregate OMM options position that we compute at one-minute frequency for each trading date. For each trading date, we compute the mean, median, standard deviation, maximum, and minimum of the gammas on that date. We then compute the statistics describing the distribution (across dates) of the daily statistics. The units are billions. For example, the value 319.6 in the row “Daily std. dev. and column "Mean" indicates that the mean (across days) of the daily standard deviations is 319.6 billion. Panel A presents the statistics for the entire sample running from July 2020 through June 2023. Panels B and C display the statistics for subsamples ending with April 2022 and beginning with May 2022, respectively.

*Panel A July 2020 - June 2023*

Statistic	Mean	Std. dev.	Min.	1%	5%	10%	25%	Median	75%	90%	95%	99%	Max.
Daily mean	341.2	319.6	-238.1	-161.0	-61.8	2.7	111.0	254.8	533.1	806.8	962.7	1,262.0	1,594.1
Daily median	340.2	317.8	-230.4	-152.2	-66.4	5.2	111.2	259.5	521.7	797.6	949.8	1,271.8	1,622.7
Daily std.dev.	86.3	92.0	3.6	7.4	15.1	22.2	37.7	62.2	107.2	158.4	208.4	496.6	956.9
Daily maxima	735.3	1,171.9	-152.7	-55.1	60.4	131.4	280.4	485.6	871.0	1,391.7	1,849.7	4,060.4	16,743.9
Daily minima	60.0	460.2	-5,426.0	-1,639.3	-453.5	-278.5	-94.9	82.2	260.3	482.7	645.0	848.5	1,038.9

*Panel B July 2020 - April 2022*

Variable	Mean	Std. dev.	Min.	1%	5%	10%	25%	Median	75%	90%	5%	1%	Max.
Daily mean	408.8	340.4	-238.1	-139.5	-59.9	26.1	156.8	333.8	620.7	902.3	1,042.2	1,273.7	1,594.1
Daily median	410.4	342.3	-230.4	-146.3	-44.4	24.4	158.1	335.8	629.7	903.9	1,047.5	1,288.6	1,622.7
Daily std.dev.	73.8	72.5	4.2	7.1	11.1	17.8	31.2	53.1	96.8	142.0	183.6	339.6	634.8
Daily maxima	705.8	800.9	-122.7	-49.2	54.1	133.1	295.3	533.4	898.1	1,370.7	1,839.4	3,035.0	11,431.0
Daily minima	183.8	327.9	-1,949.9	-605.9	-267.6	-142.9	20.8	151.2	376.0	581.8	745.7	884.4	942.7

**Table 2 continued***Panel C May 2022 - June 2023*

Variable	Mean	Std. dev.	Min.	1%	5%	10%	25%	Median	75%	90%	5%	1%	Max.
Daily mean	234.2	248.9	-202.0	-180.0	-61.9	-17.8	70.0	174.9	334.0	578.2	714.0	983.1	1,418.0
Daily median	229.1	235.7	-198.0	-164.2	-69.5	-11.0	69.6	181.8	337.3	552.3	688.4	873.8	1,418.7
Daily std.dev.	106.1	113.7	3.6	19.5	28.1	34.2	48.7	80.2	126.2	187.7	237.8	721.2	956.9
Daily maxima	782.1	1,592.3	-152.7	-62.7	73.7	131.3	258.0	430.3	798.6	1,420.2	2,126.0	7,279.6	16,743.9
Daily minima	-136.3	561.2	-5,426.0	-2,733.6	-792.6	-414.8	-207.1	-45.0	75.8	241.1	361.7	638.7	1,038.9

**Table 3**  
**GARCH model results**

This table presents estimation results for intraday GARCH(1,1) models applied to the variance of one-minute returns on the S&P 500 e-mini futures contract (symbol ES). The parameters  $\alpha$ ,  $\beta$ ,  $m$ , and  $\gamma$  are the coefficients in the variance model in equations (2)-(3). The symbols  $\varphi_l$  denote the weights on the lagged OMM gammas in equation (3), where  $l$  is the lag in minutes. The models in columns (1)-(3) estimate the initial daily variance as an additional parameter. The weights  $\varphi_l$  are calculated using the Beta weighting scheme in Engle, Ghysels, and Sohn (2013). Each value reported in the rows for each symbol corresponds to the average across months from July 2020 to June 2023, with standard errors of these estimates shown in parentheses below each value. The final row of the table presents the average values of the likelihood ratio test statistics testing the null hypothesis that  $m = \gamma = 0$ .

Coefficient	(1)	(2)	(3)	(4)
$\alpha$	0.044 (0.002)	0.044 (0.003)	0.044 (0.003)	0.053 (0.003)
$\beta$	0.925 (0.008)	0.924 (0.008)	0.924 (0.008)	0.908 (0.011)
$m$		0.131 (0.023)	0.144 (0.025)	0.157 (0.024)
$\gamma$		-0.482 (0.079)	-0.497 (0.081)	-0.476 (0.087)
$\varphi_1$			0.440 (0.046)	0.452 (0.054)
$\varphi_2$			0.145 (0.008)	0.123 (0.011)
$\varphi_3$			0.108 (0.009)	0.098 (0.011)
$\varphi_4$			0.113 (0.011)	0.105 (0.013)
$\varphi_5$			0.194 (0.028)	0.222 (0.038)
Init. var. estim.	Yes	Yes	Yes	No
LRT ( $m = \gamma = 0$ )		8.94	9.09	

**Table 4****Linear regression model results**

This table displays the results of linear regressions models in which the squared one-minute return ( $r_t^2$ ) on the S&P 500 e-mini futures contract is regressed on lagged gammas and lagged squared returns. The variables  $\Gamma_{t-l,d}$  and  $r_{t-l,d}$  are the gamma of the aggregate OMM net position and the return at time  $t-l$  of day  $d$ , respectively. All models include hourly fixed effects and date fixed effects. The last row of the each displays the values of an F-statistic comparing the model in that column to the model in the first column. The sample period is July 2020 through June 2023.

Variable	(1)	(2)	(3)	(4)	(5)	(6)
$\Gamma_{t-1,d}$		-19.53 (3.76)	-8.82 (8.53)	-7.62 (8.85)	-6.64 (8.49)	-6.81 (8.41)
$\Gamma_{t-2,d}$			-11.55 (9.20)	-3.80 (8.17)	-3.78 (8.50)	-4.15 (8.36)
$\Gamma_{t-3,d}$				-9.34 (7.79)	-3.96 (12.86)	-2.50 (12.71)
$\Gamma_{t-4,d}$					-7.17 (9.96)	-3.96 (11.29)
$\Gamma_{t-5,d}$						-4.45 (8.27)
$r_{t-1}^2$	0.088 (0.01)	0.086 (0.01)	0.085 (0.01)	0.085 (0.01)	0.085 (0.01)	0.085 (0.01)
$r_{t-2}^2$	0.055 (0.01)	0.053 (0.01)	0.053 (0.01)	0.053 (0.01)	0.053 (0.01)	0.053 (0.01)
$r_{t-3}^2$	0.051 (0.01)	0.049 (0.01)	0.050 (0.01)	0.050 (0.01)	0.050 (0.01)	0.050 (0.01)
$r_{t-4}^2$	0.056 (0.01)	0.054 (0.01)	0.054 (0.01)	0.054 (0.01)	0.054 (0.01)	0.054 (0.01)
$r_{t-5}^2$	0.059 (0.01)	0.057 (0.01)	0.057 (0.01)	0.057 (0.01)	0.057 (0.01)	0.057 (0.01)
Hour fxd. effs.	Yes	Yes	Yes	Yes	Yes	Yes
Date fxd. effs.	Yes	Yes	Yes	Yes	Yes	Yes
F-stat. ( $\gamma = 0$ )		16.3	12.0	9.0	7.6	6.7

**Table 5****Summary statistics describing the distributions of the gamma-induced changes in daily and 30-minute realized volatility**

This table displays summary statistics describing the distribution of the impacts of OMM gamma on annualized daily volatility and annualized 30-minute volatility. The impact on annualized daily realized volatility is computed using equation (6), and the impact on annualized 30-minute realized volatility is computed similarly. Panel A and B show the results for the GARCH and linear models, respectively. The sample period is July 2020 through June 2023.

*Panel A GARCH model*

Statistic	Mean	Std. dev.	Min.	1%	5%	10%	25%	Median	75%	90%	95%	99%	Max.
Daily realized vol.	-0.19%	0.86%	-5.61%	-3.12%	-1.67%	-1.16%	-0.52%	-0.08%	0.20%	0.72%	0.99%	1.56%	3.32%
30-min. realized vol.	-0.17%	0.93%	-10.63%	-3.24%	-1.71%	-1.08%	-0.47%	-0.07%	0.22%	0.67%	1.04%	2.01%	6.42%

*Panel B Linear model*

Statistic	Mean	Std. dev.	Min.	1%	5%	10%	25%	Median	75%	90%	95%	99%	Max.
Daily realized vol.	-2.14%	2.48%	-11.33%	-9.32%	-7.55%	-5.78%	-3.27%	-1.38%	-0.21%	0.11%	0.35%	1.39%	3.18%
30-min. realized vol.	-2.32%	2.82%	-17.24%	11.07%	-8.26%	-6.49%	-3.61%	-1.46%	-0.21%	0.16%	0.51%	1.80%	7.00%

**Table 6****Summary statistics describing the distributions of the changes in daily and 30-minute realized volatility**

This table displays summary statistics describing the distributions of annualized realized volatilities and changes in annualized realized volatilities due to all causes. Panels A and B present the results for daily and 30-minute realized volatilities, respectively.

*Panel A Annualized daily realized volatility*

Statistic	Mean	Std. dev.	Min.	1%	5%	10%	25%	Median	75%	90%	95%	99%	Max.
Realized volatility	12.7%	6.1%	3.5%	3.9%	5.0%	6.0%	8.1%	11.4%	16.0%	20.9%	24.1%	33.0%	36.9%
Change in rlzd. vol.	0.0%	4.5%	-16.0%	-13.3%	-7.5%	-5.2%	-2.2%	0.0%	2.2%	5.1%	7.5%	12.2%	18.3%

*Panel B Annualized 30-minute realized volatility*

Statistic	Mean	Std. dev.	Min.	1%	5%	10%	25%	Median	75%	90%	95%	99%	Max.
Realized volatility	12.0%	7.3%	1.7%	2.6%	3.8%	4.8%	6.9%	10.4%	15.2%	21.5%	26.0%	36.3%	94.5%
Change in rlzd. vol.	0.0%	4.6%	-52.2%	-11.6%	-6.2%	-4.3%	-2.0%	-0.1%	1.7%	4.3%	6.5%	14.2%	63.4%