Retail Option Trading and Market Quality: Evidence from High-Frequency Data

James O'Donovan, Gloria Yang Yu, Jinyuan Zhang *

January 13, 2023

Abstract

Retail option trading has become an important feature of modern financial markets. Since option contracts are in zero net supply, net imbalances are delta hedged by financial intermediaries. We leverage a unique dataset that allows us to categorize option trading by trader type to show that option delta hedge re-balancing trades driven by uninformed retail traders affect market quality. This effect is seen in multiple measures of liquidity measured using high-frequency data.

*O'Donovan is at the Economics and Finance Department, City University of Hong Kong, james.odonovan@cityu.edu.hk, Yu is at Singapore Management University gloriayu@smu.edu.sg Zhang is at the UCLA Anderson School of Management jinyuan.zhang@anderson.ucla.edu.

1 Introduction

The last two decades have seen market access become easier for many groups of traders. While the effects of high-frequency and algorithmic trading are now well-studied, the effects of retail trader activity on markets is less well understood. Option trading in particular has become easier for retail investors, and U.S. regulators have begun to worry about how easily investors can trade options using apps like Robinhood. [1] [2] While cheaper and easier option trading seems like positive developments for investors this trading can affect the market quality of the underlying asset because equity options trading can create latent liquidity supply and demand for the underlying security. [3]

Retail option trades create risk exposures for intermediaries which must be hedged by trading in the underlying market - depending on the risk exposures of the intermediary this trading can supply or demand liquidity from the stock market. This liquidity effect is latent and predictable: The hedge rebalancing trade of the intermediary is a known function of the realized movement in the underlying stock. In this paper we show that retail option trading from 2011-2019 has been affecting the quality of the market for the underlying stocks as measured by market liquidity.

The equity option market is imbalanced, for example, single stock call options are on average oversupplied by retail investors implying that equity option market makers must take the other side of the trade and manage the risk of the option until maturity. When

[1]Robinhood offers commission-free option trading, its average customer is 31 years old and has a median account balance of $240 (reuters.com).

[2]https://www.wsj.com/articles/investors-are-using-robinhood-other-platforms-to-jump-into-options-trades-worrying-u-s-regulators , The S.E.C. has solicited feedback on the current situation - see Securities Exchange Act Release No. 92766 (August 27, 2021) and "Staff Report on Equity and Options Market Structure Conditions in Early 2021".

[3]Retail option traders are certainly inexperienced Bryzgalova et al. (2022) estimated that retail investors lost $1.14 billion trading options from November 2019 to June 2021 (as reported at bloomberg.com).

the market maker is long the stock exposure through the net long option position they need to buy (sell) more stock when the underlying price decreases (increases) to keep delta neutral. In this case, delta hedge rebalance trades of market makers act as a stabilizing force in the stock market. In situations where the equity option market maker is net short the call option the mechanism works in reverse – the market maker demands liquidity when the price increases or decreases which acts as a destabilizing force. This mechanism suggests that retail option trading can directly affect market quality as measured by liquidity.

In this paper, we empirically quantify the impact of retail option trading on liquidity through the hedging activities of intermediaries. We first construct data on market quality by calculating quoted, effective, and realized spreads and price impact using millisecond resolution data for the years from 2011 to 2019. We combine this data with Open/Close data from the CBOE C1 option exchange as well as the NASDAQ ISE exchange, which allows us to estimate the amount of delta-hedging in the equity market due to retail option trading.[4] Our main variable *net* Γ captures the intensity of delta hedging of option market makers induced by a change in the stock price.[5] We then analyze the relationship between the stock trading required to delta hedge options and stock market liquidity at a daily level.

Our main finding is that when retail traders on net sell (buy) more options which increases (decreases) net gamma that market liquidity is improved (deteriorates). For example, a one standard deviation decrease in net-gamma causes a 11% increase in realized spreads

[4]The granularity of trade classification needed in this study has only been available since 2009 and 2011 from NASDAQ and CBOE, respectively.

[5]In option pricing terminology the delta is the first derivative of the option price with respect to the underlying price and is used as a hedge ratio to remove first order risk from a portfolio of options. Gamma is the term used for the second derivative of the option price with respect to the underlying stock price, or the first derivative of the delta with respect to the stock price. This Gamma (denoted Γ) captures the change in the hedge ratio induced by changes in prices and will be the main focus of our study.

and a 10% increase in price impact relative to the mean level of these measures. These magnitudes are particularly large in light of the fact that the mechanism that is only activated in periods when there is movement in the underlying security that induces re-hedging of option portfolios by market makers.

To enhance our understanding of the impact of retail trading on market quality, we also consider market conditions in which we expect the effect of latent liquidity to be stronger. We use the setting of earnings announcements as prior research such as Johnson and So (2018) has shown market makers adjust their spreads before earnings announcements as they are exposed to more adverse selection risk. We expect that liquidity supply and demand in the stock market due to the hedging of retail option trades will have more impact when equity market makers supply less liquidity because these intermediary re-hedging trades are not conditioning on information. We find the relationship between our measure of net gamma and liquidity to be much larger in the period leading up to earnings announcements, for example the main results on realized spreads are approximately twice as large in these periods. These results suggest that the effects of retail option trading are more pronounced when markets are already fragile.

We further consider the effects of information asymmetry using a split sample analysis. We consider analyst dispersion of earnings forecasts, and analyst coverage as measures of information asymmetry and measure whether the relationship between gamma and liquidity is higher in stocks with higher information asymmetry. We find that the impact of uninformed liquidity supply and demand driven by retail option trading is much larger for stocks where the market making for the stock is subject to more adverse selection risks. These results complement the findings for earnings periods - at times when adverse selection risk is higher the impacts of hedge-rebalancing trades (as measured by net gamma) on liquidity are larger.

We consider several alternative explanations for our results and provide robustness tests. First, we use the method of Ni et al. (2021) to show that our results are robust to removing from our measure variation that could be related to informed trade in options by retail investors, in fact, our results are even stronger when using this method. Another concern is that retail option trading predicts volatility because retail traders are informed about option volatility and trade on that information. Lakonishok et al. (2007) show that very little trading is volatility-based – i.e. opening straddles or strangles - this is consistent with anecdotal evidence that trading profitably on volatility information is extremely difficult for non-sophisticated option traders. Nonetheless, we rule out this explanation by showing that volatility trading strategies built using this information are never profitable even when we assume that investors can trade at the midpoint of the bid-ask spread and thus don't have to pay large option transaction costs. We also show that the fact that we use an estimate of net gamma has minimal impact on our results and that the result is not driven by option expiration periods or the particular measure of net gamma that we use.

Our results have policy implications. While earlier literature has studied whether the listing of options can affect the underlying asset and found mixed results, our findings make clear that this effect depends on the time-varying level of *net* Γ. Thus, an option introduction could cause an increase in non-informational liquidity supply and demand which improves market quality if retail investors are net short the options but this effect could work in the opposite direction if retail investors start to take a net long position in the options. This means regulators should pay attention not only to the existence of an option market, but the current state of the market. Currently, option trading is still dominated by the CBOE C1 exchange and the NASDAQ ISE exchange, however, the number of trading venues are increasing over time so that market fragmentation could make this hedge-rebalancing effect more difficult to track. This creates incentives for regulators to aggregate information relating to option trading on multiple exchanges. [6]

2 Related Literature

Our results contribute to a large literature on the relationship between option trading and market quality. Portfolio insurance strategies based on option replication were blamed for the infamous October 1987 crash. Fedenia and Grammatikos (1992) find that options listing decreases spreads on small stocks while increasing spreads on large stocks. Results in Kumar et al. (1998) suggest that option listing increases various aspects of market quality. Both Conrad (1989) and Detemple and Jorion (1990) find a volatility decrease after option introduction, while Mayhew et al. (2000) find that option introduction increases volatility when one controls for the endogeneity of option listing. Sorescu (2000) looks at the effect of option introduction on stock prices and find that prices sometimes increase and sometimes decrease. In a more recent study Hu (2018) finds that option listing increases uninformed trading in the US between 2001 and 2010. Our results add to this literature by explaining why options can affect market quality, while explaining that the result can vary over time and depends on the general trading patterns of retail option investors.

Financial technology can obscure this activity of market makers from many option end-users. In particular, retail investors often don't appreciate that as long as a market maker is net-long or short options in a particular stock, their option trade has the same effect as submitting a particular liquidity demand or supply schedule to the market with their order. The delta hedging program will be a mechanical function of the underlying asset

[6]Evidence in Bryzgalova et al. (2022) and Ernst and Spatt (2022) suggests that relationships between retail brokers, exchanges, and wholesalers and affiliated market makers are becoming increasingly complex and that certain retail traders orders may end up concentrated on particular exchanges in the future.

price change.[7] This liquidity is latent in the sense that it is only realized if the underlying price changes and in this way can pose a hidden threat to the quality of the underlying market. [8]

We also speak to the research on market access and market quality. Technology has changed markets dramatically in recent decades, Hendershott et al. (2011) finds algorithmic trading improves liquidity and enhances the informativeness of quotes and Brogaard et al. (2010) find that high-frequency-traders add substantially to the price discovery process and often provide the best bid and ask quotes. Brogaard et al. (2018) find that high-frequency-traders are reliable in times of market stress and do not seem to cause extreme price movements. Boehmer et al. (2021) use an international sample and find that algorithmic trading improves informational efficiency but increases short-term volatility Chakrabarty et al. (2021) show that completely unfiltered market access may not be beneficial for traders who demand liquidity. Our paper suggests retail option trading that is completely unconstrained can cause changes in market quality. To the extent that technology continues to remove frictions to retail option trading, the empirical effects we document should become even stronger.

This paper also contributes to a burgeoning literature on the effects of derivative hedging on stock market return characteristics. Baltussen et al. (2021) find evidence that short gamma positions of option hedgers is related to intraday momentum in futures markets, while Barbon et al. (2022) show that both momentum and mean reversion are driven by leverage effects and hedge-rebalancing in ETF markets. Ni et al. (2021) show that market maker option gamma is negatively related to daily absolute returns suggesting

[7] In some situations we may not observe other positions held by the market maker "an-axe" which allows them to avoid delta-hedging, but we don't expect this to have a large practical impact on our results.

[8] Another example of a latent factor affecting liquidity is the existence of margins - Foley et al. (2022) find that sudden increase in margin requirements during the covid-19 crisis resulted in withdrawn liquidity.

a feedback effect to from options to return volatility. There is also a large literature on the effects of more exotic derivative products on financial market quality Auh and Cho (2022) show that when payoffs of structured equity derivatives change, it can cause significant price pressure of the underlying stock upon an event of dramatic payoff change and that other market participants try to front-run these changes. Boehmer et al. (2015) that the introduction of single-name credit default swap (CDS) contracts reduces market quality of the underlying asset. The results in this paper complement these studies by providing direct microstructure evidence that option trading affects market quality. A distinct advantage of our setting is that by calculating market quality measures using trades and quotes we can link the option gamma to quantities of interest to regulators and prospective traders, allowing us to make statements about likely ex-ante market quality rather than ex-post market characteristics.

We also connect to a theoretical literature on liquidity. Theories of liquidity typically map liquidity to state variables like information asymmetry (Glosten and Milgrom (1985), and Easley and O'hara (1987)), intermediary inventory costs (Amihud and Mendelson (1980) and Ho and Stoll (1981)) and raw trade costs (Demsetz (1968)). However, some modern theories include derivatives within models of liquidity. Ronnie Sircar and Papanicolaou (1998) show theoretically how portfolio insurance strategies can increase the volatility of the underlying when program traders must delta hedge their inventory. Wilmott and Schönbucher (2000) also study the feedback effects of option replication trading. In a recent study Huang et al. (2021) embed derivative trading into a model with informed traders and show that delta hedging can reduce price impact in the underlying. Our findings provide empirical evidence in support of these theoretical mechanisms.

3 Retail option trading and latent liquidity

While equity securities have net-positive supply, equity options have net-zero supply - for every option buyer, there must be an option writer. The classic theory of option pricing of Scholes and Black (1973) and Merton (1973) shows that option pricing is intimately linked with stock trading through a delta hedging argument in which a dynamically varying position in the stock is used to offset the risk exposure of the option position. Thus, options are created by replication and this replication activity will require liquidity to be demanded or supplied depending on the case situation. Early theories ruled out market impact of the trading needed for option replication by the assumption of frictionless markets without transaction costs. [9] Removing the assumption of frictionless markets, allows the stock-trading from dynamic delta hedging to impact the market for the underlying.

To fix ideas we focus on an individual stock and detail the relationships between the liquidity supply and demand from option trading and market quality. First consider the delta a given option contract [10]

$$\Delta_{i,t} = \frac{\partial C_{i,t}}{\partial S_{i,t}}$$

This represents the rate of change of the option price with respect to the stock price and is the hedge ratio of first-order importance for an option market maker. The hedging of options would be trivial if not for changes in this hedge ratio, but since the option price is a nonlinear function of the stock price, this hedge ratio will change over time.

[9] In reality, equity option market-makers act as producers in the option market rather than mere matching mechanisms, market makers for equities aim to make the spread and end the day without holding any stock inventory. In an ideal world, equity option market makers would operate in a similar fashion, matching equity option buyers and sellers, and taking the spread as compensation for risks such as inventory, operational costs, and information asymmetry risk.

[10] We will use the Black-Scholes-Merton option price function to calculate closed-form values for the option sensitivities

The first derivative of the option delta with respect to the stock price tells us how much the hedge ratio will change when the stock price changes, and thus, allow us to predict the mechanical trading of the option market maker given what we know about their net option position and total net gamma.

$$\Gamma_{i,t} = \frac{\partial \Delta_{i,t}}{\partial S_{i,t}} = \frac{\partial^2 C_{i,t}}{\partial S_{i,t}^2}$$

Later we will detail the calculation of the net option position to be hedged by market makers, for now we take this as given and use an example to fix ideas. To operationalize this idea that changes in price will induce trading we start with a measure of net share gamma which captures the number of shares that need to be purchased by the market makers given a \$1 increase in the stock price. Here it is crucial that we use the net open interest of the option market makers (or more generally delta-hedgers).

$$\text{Net Share Gamma}_t = \sum_{j=1}^{J_t} \text{Net Open Interest}_{j,t}^{\text{delta hedgers}} \Gamma_j(t, S_t)$$

We then rescale this measure so that it is expressed in terms of dollar volume traded per 1% move in the stock price, a measure commonly used by practitioners

$$\$ \text{ Gamma } 1\%_t = \text{Net Share Gamma}_t \times S_t \times \frac{S_t}{100}$$

This tells us the dollar volume of stock that would be traded given a hypothetical move of 1% in the stock price. We can compare this to the total dollar volume traded in the stock on a given day

$$\$ \text{ Volume} = Volume_t \times S_t$$

Finally, we can also calculate the realized dollar volume of stock traded on day t due to delta hedging, by combining the Net Share Gamma_t and the realized return [11]

$$\text{\$ Gamma Realized}_t = \text{Net Share Gamma}_t \times S_t \times R_t$$

Figure 1 contains two visual examples of the relationship between Dollar Net Gamma and liquidity (Percent Price Impact) for one of our sample firms (Meta Platforms Inc.). Figure 1 Panel A contains a kernel density estimate for the distribution of price impact in two different states - when Dollar Net Gamma is greater than zero, and less than zero respectively. Given the discussion above, we expect liquidity conditions and market quality to be better when market maker Net Gamma is higher because of increased liquidity supply. It is clear from the graph that the distribution of Price Impact is much tighter and concentrated closer to zero when Dollar Net Gamma is positive, when Dollar Net Gamma is negative we see a much longer tail in the distribution of price impacts, suggesting that the mechanical liquidity demand of option hedgers in the stock market affects market quality. Panel B contains a time series plot of the Net (dollar) Gamma and Price Impact. We calculate one-month moving averages of the variables to make the data easier to visualize (we also lag the Dollar Net Gamma by one day relative to the realization of price impact). We can see that on average periods with high Net Gamma have relatively lower price impact which suggests that the results in Panel A are not coming from some time effects (such as both dollar gamma and price impact being subject to a time trend).

[11] This is an estimate which is accurate for small returns, when returns are extreme there may also be changes in implied volatility which will induce more trading through the sensitivity of delta to implied volatility (we ignore this second-order sensitivity which is named vanna in the option trading community).

4 Data and Variables

We use three main databases. We use the Millisecond Trade and Quote data, the "Daily Product" from TAQ, Open/Close data from the Chicago Board Options Exchange (CBOE) and the NASDAQ International Securities Exchange (ISE), and stock price and volume data from CRSP. We include common stocks with share codes 10 or 11 and stocks with exchange codes 1, 2 or 3 corresponding to the NYSE, NYSE MKT (formerly AMEX), NASDAQ. After calculating liquidity measures, and net gamma for each stock we keep only stocks for which we have at least 500 trading days.

4.1 Liquidity Measures

We use the Daily Trade And Quote (DTAQ) database to calculate the liquidity measures. Holden and Jacobsen (2014) show that DTAQ is the first best solution for calculating liquidity measures and we refer interested readers to their paper for detailed sample cleaning and institutional details. [12]

The percent quoted spread for time interval s is defined as

$$\text{Percent Quoted Spread}_s = \frac{A_s - B_s}{M_s}$$

where A_s is the National Best Ask and B_s is the National Best Bid assigned to time interval s by a particular trade classification technique and M_s is the midpoint, which is the average of B_s and A_s. We aggregate to the daily level and calculate for each stock the time-weighted average of the Percent Quoted Spread over all time intervals and denote this variable as the *Quoted Spread*.

[12] We thank the authors for making their code available online.

The percent effective spread for a given stock on trade k is defined as

$$\text{Percent Effective Spread}_k = \frac{2D_k(P_k - M_k)}{M_k}$$

where D_k is an indicator variable that equals +1 if the kth trade is a purchase, −1 if the kth trade is a sale, and P_k is the trade price. We aggregate to the daily level by calculating for each stock the dollar-volume-weighted average of the Percent Effective Spread over all trades and denote this variable as the *Effective Spread*.

The percentage effective spread can be decomposed into a permanent and a transitory component, the temporary component which is the realized spread, and the permanent component which is the price impact.

The percent realized spread on the kth trade on a given stock is defined as

$$\text{Percent Realized Spread}_k = \frac{2D_k(P_k - M_{k+5})}{M_k}$$

where M_{k+5} is the midpoint five minutes after the midpoint M_k. We aggregate to the daily level by calculating for each stock the dollar-volume- weighted average of the Percent Realized Spread over all trades and denote this variable as the *Realized Spread*.

The percent price impact on the kth trade on a given stock is defined as

$$\text{Percent Price Impact}_k = \frac{2D_k(M_{k+5} - M_k)}{M_k}$$

We aggregate to the daily level by calculating for each stock the dollar-volume-weighted average of the Percent Price Impact over all trades and denote this variable as the *Price Impact*.

The Percent Realized Spread and Percent Price Impact are a function of D_k which requires the use of a trade classification algorithm. There are three widely used trade classification classification algorithms: First is that of Lee and Ready (1991) in which a trade is classified buy if $P_k > M_k$, a sell if $P_k < M_k$, or tick test is used if $P_k = M_k$. Using the tick test, a trade is classified as a buy (sell) if the last trade at a different price was at a lower (higher) price than P_k; second is that of Ellis et al. (2000) in which a trade is classified buy if $P_k = A_k$, a sell if $P_k = B_k$, and the tick test is used otherwise; the third method is Chakrabarty et al. (2007) in which a trade is a buy if $P_k \in [0.3B_k + 0.7A_k, A_k]$, a sell if $P_k \in [B_k, 0.7B_k + 0.3A_k]$, and the tick test is used otherwise. We do not take a stance on the correct algorithm, rather we show the main results for all three methods, and show other results using the Lee and Ready (1991) algorithm while verifying that research conclusions do not change with alternative classification methods.

These various measures tell us about different dimensions of liquidity. Specifically, the quoted spread is a good indicator of trading costs for a small investor who does not expect their trade size to move prices. The effective spread takes into account latent liquidity or movements in the price due to a trade absorbing all of the depth at the given quote and executing on higher or lower limit orders. Realized spreads pick up price changes from the trade-price to the post-trade value and can be thought about as a proxy for market maker revenues. Price impact reflects the cost market makers face when trading with informed traders, after such a trade the price will not reverse as the information is incorporated into the price. We can see the realized spread is equal to the effective spread less price impact. Some studies refer to price impact and realized spreads as the 'permanent" and "transitory" price impacts of a trade or as the "informational" and "non-informational" impact of trading.

4.2 Calculation of market maker gamma

CBOE and ISE classify trades by trader group and type. Trader groups are public customers (retail traders and professional customers), firms (firm proprietary traders and broker dealers) and market makers. Trade types are either open buy, open sell (an option is written), close buy (a written option is closed out) or close sell (a purchased option is closed out). By tracking the daily movements of each category and type of trade we can follow the total open interest of each trader group. We are interested in the total open interest of likely delta hedgers which represents the net option demand of retail traders that must be delta hedged by intermediaries.

$$\text{Net Open Interest}_{j,t}^{\Delta\text{hedgers}} = \text{ Open Interest}_{j,t}^{Sell,Retail} - \text{ Open Interest}_{j,t}^{Buy,Retail}$$

where for each option series j

$$\text{Open Interest}_{j,t}^{Sell,Retail} = \text{ Open Interest}_{j,t-1}^{Sell,Retail} + Volume_{j,t}^{OpenSell} - Volume_{j,t}^{CloseBuy}$$

and

$$\text{Open Interest}_{j,t}^{Buy,Retail} = \text{ Open Interest}_{j,t-1}^{Buy,Retail} + Volume_{j,t}^{OpenBuy} - Volume_{j,t}^{CloseSell}$$

Where Open Buy and Open Sell represent new purchased and written options by customers and Close Buy and Close Sell represent the closure of existing positions.

To calculate the net gamma for an individual stock, we account for the fact that each option contract is written on 100 shares of the underlying stock and sum the gamma across all open series on day t

$$\text{Net Share Gamma}_t = \sum_{j=1}^{J_t} 100 \times \text{Net Open Interest}_{j,t}^{\text{delta hedgers}} \Gamma_j(t, S_t)$$

In this form, larger firms with more option trading will mechanically have higher *Net Share Gamma*, so we rescale the net gamma to make the variable comparable across firms, we choose the rescaling in Ni et al. (2021) to make $Net\Gamma_t$ dimensionless.

$$Net\Gamma_t = \sum_{j=1}^{J_t} 100 \left(\frac{S_t}{M_t}\right) \times \text{Net Open Interest}_{j,t}^{\text{delta hedgers}} \Gamma_j(t, S_t) \quad (1)$$

Where J_t is the number of open option series on the stock, S_t is the stock price and M_t is the number of shares outstanding all measured at t. Since each individual stock will have somewhat unique liquidity conditions, we don't expect the relationship between $Net\Gamma$ and liquidity to be identical across stocks. For this reason we will later estimate the relationship firm-by-firm allowing the coefficients for each stock to differ.

We can only estimate the positions of likely delta hedgers because the trading volume data are from the CBOE and ISE together compose only about 50% of the options trading volume. Using data from these two exchanges will induce classical measurement error to our gamma variable which will bias us against finding any relationship between $Net\Gamma$ and liquidity.

Our measure of $Net\Gamma$ depends on whether we include firm proprietary traders in the set of likely delta hedgers. Market makers and broker dealers are almost certain to hedge their inventory while proprietary traders are assumed to mostly delta hedge their inventory.

5 Empirical Analysis

5.1 Likely Delta Hedgers and Retail Trading

The option trades in our dataset can be classified into three broad categories which include public customers, firms, and market-makers. Trades initiated by firms can be broken down further into trades by broker dealers and proprietary traders. Crucially for our study, transactions initiated by public customers can be further classified into retail investors and those by institutional investors such as hedge funds (called "professional customers"). [13] The estimate of $Net\Gamma_{t-1}$ depends on our assumption about the identities of likely delta hedgers. The literature either uses market-makers and broker-dealers or market-makers, broker-dealers and firm proprietary traders as the likely delta hedgers. Since our focus is on trades of retail investors with intermediaries we follow Chen et al. (2019) to merge firm proprietary traders and market-makers as one group and observe how their trades with public investors affect the market.

Our focus is on the effect of retail investors on market liquidity through delta hedging of net option positions. However, our measure for $Net\Gamma$ includes both retail traders and "professional customers" who may be considered more sophisticated investors. To alleviate the concern that our measure of retail trading is distorted by these investors, we show that they make up a negligible portion of overall public customer trading volume. [14]

Table 1 contains data on the percentage of total trading volume for each firm which is

[13] This further granularity which allows us to distinguish between professional customers and retail traders is only possible after 2009 for the NASDAQ data and after 2011 for the CBOE data.

[14] It is also worth highlighting that many of these professional customer trades are likely coming from products such as CallWrite and PutWrite funds which are ultimately sold back to retail investors.

accounted for by each category of trader. We calculate the percentage of non-delta-hedger volume for each category of trader at the firm level, and then calculate averages of each statistic across firms. If we assume likely delta hedgers are firm investors and market makers we see that almost all of the trading volume comes from retail investors and not professional customers. Table 1 Panel A shows that the average firm has median retail trading of 97% and professional customer trading of 3%, even at the first percentile 83% of trading is accounted for by retail trades. Table 1 Panel B shows that even if we remove firm investors from the set of likely delta hedgers, the vast majority of non-delta-hedger trading is coming from retail investors with the average firm having a median of 78% retail trading. Later we show that our results are robust to assuming that the likely delta-hedgers do not contain firm proprietary traders.

5.2 Descriptive Statistics

Our final sample contains 4,279,421 daily observations for 2,639 unique firms from 2011 to 2019.

Table 2 contains summary statistics for the main variable and control variables used in the paper. We calculate the statistics for each firm, and then average across firms. We can see that the mean and median of $Net\Gamma$ is positive, meaning that on average for each firm, retail investors are net short options which have to be delta hedged by market makers. We also present multiple spread measures. The average quoted spread is 0.30% while the average effective spread is 0.24% which can be decomposed into 0.08% realized spread and 0.16% price impact using the method of Lee et al. (1993), with alternative trade classification measures giving qualitatively similar values.

5.3 Regression Specification

Our primary focus is the relationship between delta-hedger gamma from retail-trading and measures of market quality. To this end we specify, for each stock the following regression model:

$$LIQ_t = a + b \times Net\Gamma_{t-1} + \sum_{l=1}^{L} c_l \times X_{l,t-1} + \varepsilon_t \tag{2}$$

The regression model is estimated for each stock, and we report the cross-sectional average coefficient estimates. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks (see Thompson (2011)). The key prediction is that $Net\Gamma_{t-1}$ is negatively related to measures of liquidity at time t. [15]

As we run the regression stock-by-stock, this coefficient captures the average relationship across all stocks. Although we normalize the net-gamma measure to make it comparable across stocks, we cannot guarantee that the relationship has the same coefficient for each stock, thus we use the same regression specification as Ni et al. (2021). Another advantage of this specification is that we do not need to worry about time-invariant stock characteristics that determine liquidity. Some firms may have relatively higher or lower levels of $Net\Gamma$ relative to others which could have a time-invariant level effect on market quality and we cannot identify these effects as they can not be distinguished from other

[15] Many authors have noted historical time trends in liquidity as well as sharp drops in spreads related to tick-size changes (from eights to sixteenths in 1997 and decimalization in 2001) which induced them to use changes or deviations from trends in their specifications (see e.g. Comerton-Forde et al. (2010). The sample from 2011 to 2019 does not display any time trend.

time-invariant firm characteristics. We control for volatility by including up to 10 lags of absolute returns and absolute returns interacted with a positive return dummy to account for the asymmetric response of volatility to negative returns, we also control for 10 lags of the dependent variable to account for any persistence in liquidity.

5.4 Main results

5.4.1 Overall Sample

The results are contained in Table 3. We can see that for each measure of liquidity there is a statistically negative relationship between $Net\Gamma_{t-1}$ and that measures of liquidity on day t. The results are intuitive, first the economic significance of the relationship between $Net\Gamma$ and the *Quoted Spread* is quite small, amounting to about 1.5% of the daily average value. This is sensible as the liquidity demand and supply mechanism related to option trading needs some movement in the price of the underlying before the option delta is changed and the mechanism is activated. Nonetheless, market makers of the underlying asset who detect greater uninformed trading may be willing to reduce spreads a little. For the regression with *Effective Spread* as the dependent variable the coefficient on $Net\Gamma$ is -0.033, since a 1 standard deviation move in gamma is 0.006 this translates to a -2 basis point change in the *Effective Spread* for every standard deviation move in gamma which is approximately 8% of the overall average level of the *Effective Spread*. The effect on *Price Impact* is slightly larger at 10% while the largest impact of option trading is on *Realized Spread* with a one-standard deviation move explaining over 11% of the mean percent realized spread. This result is intuitive - when $Net\Gamma$ is positive the equity option hedge trading represents mechanical liquidity supply that competes with existing market makers.

The overall message from these results is that, on average across stocks, retail option trading that results in changes in the $Net\Gamma$ position of equity option market makers can make a significant difference to the market quality of the underlying stock. Increases (decreases) in $Net\Gamma$ lead to improvement or deterioration in all measured aspects of market quality. Market marker profits as measured by the transitory component of the spread as well as the price impact of potential informed traders are most affected. The economic size of these results is large, considering these effects are latent and contingent on sufficient movements in the underlying asset to elicit changes in the optimal option hedge. In the next section we move on to empirical tests which highlight time periods and stocks where the effect is expected to be particularly large.

5.4.2 Evidence From Earnings Announcement Periods

Our main results show that delta hedgers in the option market affect underlying market quality through their mechanical liquidity demand and supply. In this section, we show the impact of this liquidity plays a more important role when existing market makers are reluctant to provide liquidity.

The period leading up to an earnings announcement is a perfect setting to test the effect of latent liquidity on measures of liquidity, because these effects can only be seen in cases when equity option market makers delta hedge trades are the marginal trades in the market. This is more likely when traditional liquidity providers have a reduced incentive to make the market. Krinsky and Lee (1996) show that the adverse selection component of the bid-ask spread increases around earnings announcements and Johnson and So (2018) show that market makers supply liquidity asymmetrically leading up to earnings announcements. Lee et al. (1993) show that spreads widen in the lead-up to earnings announcements suggesting liquidity providers use spreads to actively manage information asymmetry risk.

To see why stock-trading coming from the re-hedging of a market maker's options portfolio should have more impact during periods of relatively higher information asymmetry consider typical models of trading with information asymmetry such as Kyle (1985) or Glosten and Milgrom (1985) in which uninformed investors and informed investors trade with an optimizing market maker (or auctioneer). These models rely on a market maker (or auctioneer) who behaves rationally and conditions the price (or the bid ask spread) on available information about the ratio of informed and uninformed traders. The introduction of an options market maker who is hedging a positive net gamma position would effectively introduce a competitor for the market maker who supplies liquidity at given price points (those points at which the option portfolio is re-hedged) without conditioning on any information. This non-optimizing trader is much more likely to be the marginal liquidity supplier in a market where information asymmetry risk is higher and optimizing market makers hedge this risk by reducing the quantity of shares in the limit order book. Likewise, the introduction of a market maker who is hedging a positive net gamma position would effectively introduce a non-strategic trader who systematically takes liquidity at certain price points (re-hedging the delta) which will impact prices/spreads as long as equity market makers cannot distinguish their non-informational trades from other trades in the market.

To test the hypothesis that latent liquidity effects driven by retail option trading are stronger in the lead-up to earnings announcements we estimate the regression model of the previous section augmented with an interaction term between $Net\Gamma$ and $EARN$ where the variable $EARN$ is equal to 1 in the three trading days leading up to the earnings announcement day, and zero otherwise. The specification is:

$$LIQ_t = a + b \times Net\Gamma_{t-1} + c \times Net\Gamma_{t-1} \times EARN_t + \sum_{l=1}^{L} c_l \times X_{l,t-1} + \varepsilon_t \quad (3)$$

To the extent that the impact of latent liquidity is stronger in the lead-up to earnings we expect the estimated average coefficient $\hat{c}$ to be negative and economically significant.

The results are contained in Table 4. We can see that all four liquidity measures are impacted more heavily by option-hedge trading during earnings periods. This suggests that retail option trading can have a particularly large impact on stocks with already fragile trading environments due to perceived information asymmetry. Adding the baseline effect to the interaction effect we see that the magnitude of the result is twice as large for *Quoted Spread*, *Realized Spread*, and *Price Impact*, while the result is three times as large for the *Effective Spread* measure.

5.4.3 Evidence from high information asymmetry stocks

The evidence above suggests that the effect of latent liquidity from retail option trading is stronger at times when adverse selection risk is higher. A related question is the importance of stock level information asymmetry, we would expect this effect to also be stronger in firms where the overall information environment makes information asymmetry risk a larger concern for market makers in the equity. We test this idea with a commonly used measure of information asymmetry, namely analyst disagreement. Sadka and Scherbina (2007) use this measure to investigate the link between mispricing and liquidity. We calculate analyst disagreement in each year as the standard deviation of all outstanding fiscal year earnings forecasts scaled by the absolute value of the mean forecast. In stocks where analysts disagree more about potential earnings there is greater scope for informed traders to take advantage of their information, and thus market makers are change the price of liquidity. However, latent liquidity demand and supply that arises from retail option trader imbalances and delta hedging of market maker inventory must be demanded or supplied as a function of the delta hedging needs, without taking into account adverse selection.

The results are contained in Table 5. We can see that for every measure of liquidity, stocks with high analyst disagreement - which we expect to have more information asymmetry risk for market makers - are more affected by delta hedging of net retail option positions. The differences in the effect size between low and high dispersion stocks 20% for *Quoted Spread*, for *Effective Spread* the effect size is approximately 4 times as large. For the *Price Impact* variable the effect size is approximately twice as large, giving an effect size of 13% for high dispersion stocks relative to 6% for low dispersion stocks. This suggests that retail option trading can be particularly impactful in stocks with already fragile trading environments due to perceived information asymmetry.

As an alternative measure of information asymmetry we also consider analyst forecast coverage, we expect stocks with low coverage to have a poorer information environment which presents better opportunities for informed traders and thus more risk averse behavior of market makers when providing liquidity. Results in Table 6 show the relationship between $Net\Gamma_{t-1}$ and various measures of liquidity. Except for *Quoted Spread* which has a similar effect size for high and low analyst coverage stocks, the effect sizes for the other variables are three to ten times as large for low analyst coverage stocks relative to high analyst coverage stocks.

5.5 Decomposition of $Net\Gamma$

While the effects of $Net\Gamma$ on liquidity will exist regardless of the source of the net option imbalance that must be delta hedged, the interpretation of this statistical coefficient estimate depends on whether the variation comes from informed or uninformed trading. If informed traders anticipate changes in volatility which thus anticipates changes in liquidity, it could bias our estimate. To deal with this we decompose our $Net\Gamma$ measure to remove variation potentially related to informed trading.

Following Ni et al. (2021) we decompose $Net\Gamma$ into 3 components. The $Net\Gamma$ at any time t can be decomposed into a component that comes from net gamma τ periods ago, a component that comes from new option trades between $t-\tau$ and t, and a component that comes from changes in gamma that arise from changing stock prices τ days ago until today. Specifically, we can first decompose $Net\Gamma_t$ into a component that comes from trades in the last τ days and the residual:

$$Net\Gamma_t = \underbrace{Net\Gamma_t(t-\tau, S_t)}_{\text{Residual Gamma}} + \underbrace{[\text{Net}\Gamma_t - \text{Net}\Gamma_t(t-\tau, S_t)]}_{\text{Information Gamma}} \tag{4}$$

The first term gives the net gamma coming from positions opened $t-\tau$ periods ago using the day t price. A potential concern is that the residual component still may contain long-lived information about liquidity at time t, under the assumption that the investor does not anticipate the change in the net gamma coming exclusively from the change in price between time $t-\tau$ and time t we can further decompose into the gamma at period $t-\tau$ using the stock price at $t-\tau$ and the gamma at period $t-\tau$ using the stock price at t and the gamma at period $t-\tau$ using the stock price at $t-\tau$:

$$NetΓ_t = \underbrace{NetΓ_t(t-\tau, S_{t-\tau})}_{\text{Period } t-\tau \text{ Gamma}} + \underbrace{[\text{Net}Γ_t - \text{Net}Γ_t(t-\tau, S_t)]}_{\text{Information Gamma}} + \underbrace{[\text{Net}Γ_t(t-\tau, S_t) - \text{Net}Γ_t(t-\tau, S_{t-\tau})]}_{\text{Hedge Gamma}} \tag{5}$$

Given $\tau = 5$ the hedge gamma component captures hedge rebalancing trades due to movement in the stock price from t-5 to t which could not have been predicted by the option trader and this makes this component relatively immune to concerns about informed trading.

Table 7 contains the results of using the first decomposition. We can see that the non-information gamma is negatively related to all four measures of liquidity. The economic and statistical significance is quantitatively similar to the results reported in Table 3. Table 8 contains the results of using the second decomposition. Using this decomposition we see that the hedge gamma component is not significant for *Quoted Spread*. However, for the *Effective Spread*, *Realized Spread*, and *Price Impact*, the effect sizes are even larger than in our main specification. The effect size for *Effective Spread* is now 15% which is almost twice as large as the effect size of 8% in Table 3. The effect sizes for the *Realized Spread* and *Price Impact* are larger by 15% and 55% respectively. These results help alleviate concerns that information driven changes in net-gamma are generating the statistical relationship between net-gamma and liquidity.

5.6 Robustness

In this section we consider the robustness of our main results to a variety of potential explanations.

5.6.1 Alternative Trade Classification Algorithms

In Table 9 we assess the impact of alternative trade classification algorithms on our results. We re-run regressions model in equation (2) using the methods of Ellis et al. (2000) (EMO) and Chakrabarty et al. (2021) (CLNV) to calculate the Percent Realized Spread and Percent Price Impact which rely on a trade classification algorithm. The results on Percent Price Impact are statistically and economically slightly stronger with the CLNV classification while being weaker with the EMO classification, while the results on Percent Realized Spread are qualitatively similar, the economic effect sizes are lower with the alternative trade classification algorithms.

5.6.2 Excluding Proprietary Traders from Likely delta Hedgers

Equation (1) makes clear that the measure of net-gamma depends on the group of likely delta hedgers. While it is natural to assume firm proprietary traders are hedging the options they hold, one can also consider a measure of gamma that considers the firm proprietary traders to be more similar to retail investors and hold naked option positions. We replicate Table 2 using this alternative definition of $Net\Gamma$ in Table 10 and find that the results are largely similar to those using the alternative classification. There is a negative and statistically significant relationship between $Net\Gamma_{t-1}$ and all four liquidity measures as in Table 3. The effect sizes for *Quoted Spread* and *Effective Spread* are almost identical, the effect size for *Realized Spread* is lower (at 4% relative to 11% in Table 3), while the effect size for *Price Impact* is higher (at 13% compared to 11% in Table 3).

5.6.3 Do retail investors predict volatility?

One alternative explanation for these results is that retail investors have private information about volatility, which they express by taking positions in the options market. While the results of the decomposition show our results are robust to such a concern, in this section we show a trading-strategy profitability based robustness test.

Our measure of $Net\Gamma_{t-1}$ connects the behavior of option delta-hedgers today to the decisions of option traders multiple periods ago. To the extent that the aggregated decisions of option traders reflect information about future outcomes in the market, there is a concern that this variable picks up information as well as a delta-hedging effect.[16] For example, if retail investors expect low volatility, they may sell options on net, leaving market makers net long gamma and if this low volatility environment is associated with low liquidity it could explain the relationship between gamma and liquidity. [17]The mechanism described in this paper will exist regardless of any relationship between private information and volatility, but the policy implications are different if non-informational trading can cause changes in liquidity as we claim.

There is much evidence to suggest retail traders are uninformed and unsophisticated, Bryzgalova et al. (2022) suggests that retail investors lose quite a lot on their option

[16]It should be noted that the delta hedging effect is mechanical and the question is one of quantifying the effect. The concern is whether the entire coefficient can be attributed to this delta-hedging effect or some of the magnitude should be attributed to another channel such as volatility information trade.

[17]Ni et al. (2008) provide empirical evidence to suggest investors trade on volatility information, however, their measure focuses on demand for vega which is concentrated in long-maturity options while gamma is concentrated in short-maturity options.

trades, and Lakonishok et al. (2007) estimate an upper bound of approximately 1% of open volume on options devoted to purchased or written straddles which suggest volatility trading. Hu (2018) also find that option introduction disproportionately increases uninformed trading relative to informed trading which is consistent with the findings in this paper. Nonetheless, we test this hypothesis in our paper by considering the profitability of a variety of volatility trading strategies using retail investor data.

If the private-volatility-information hypothesis is true, then retail investors should profit from trading on volatility information. For an investor with private information about volatility to profit in the options market, she needs to take a position that can profit from information that realized volatility will be above or below market expectations (as captured in option implied volatility). If retail investors expect realized volatility to be high (low) they should go short (long) volatility by selling (buying) a call option and delta-hedging the directional risk. To capture this idea we construct a time-series of delta hedged returns at the stock level and consider a variety of trading strategies that would profit if $Net\Gamma$ is informative about volatility. In particular we consider strategies which take a long or short position depending on the level of $Net\Gamma$. We consider strategies which are long (short) volatility when $Net\Gamma$ is negative (positive). As levels of $Net\Gamma$ near zero may reflect a lack of conviction by retail investors about the direction of volatility, we also consider strategies when investors only take long-short positions when gamma is at extreme levels such as above the 75th, 90th, and 95th percentile values or below the 25th, 10th, or 5th percentile values. We consider both equal-and value-weighted returns and we assume that investors can trade at the midpoint of the bid-ask spread to give the best possible chance of achieving economically significant returns.

The results are presented in Table 11. If investors were trading options based on private information on volatility, we would expect strategies based on this information to display economically and statistically significant positive returns. The results in Panel A and B reveal that an investor who tried to exploit information in gamma would typically make significantly negative returns (even in the absence of transaction costs which are large for such option trading strategies).

5.7 Option Expiration Effects

It is well known that for an at-the-money option, gamma increases strongly as time to maturity decreases. As such, one might think that the effects described in this paper are concentrated only in periods where options expire. While this would still be a notable effect, it would be easier to predict and monitor for regulators and market participants because options have a consistent expiration cycle. All options with an expiration greater than one week, will expire on the third Friday of the expiration month. If all of the liquidity effects in this paper are concentrated in this period, then without knowing net-gamma, an investor or regulator could predict the existence of liquidity effects (but not the direction which would require knowing whether retail investors are net-long or net-short gamma).

We explore the hypothesis that the liquidity effects are concentrated in option expiration weeks by estimating the following specification

$$LIQ_t = a + b \times Net\Gamma_{t-1} + c \times Net\Gamma_{t-1} \times OPEX_t + \sum_{l=1}^{L} c_l \times X_{l,t-1} + \varepsilon_t \tag{6}$$

where $OPEX_t$ is equal to 1 if it is an option expiration week and zero otherwise. If option expiration effects explained the earlier results, we would expect the average coefficient on $Net\Gamma$ to be statistically and economically insignificant after controlling for the interaction of $Net\Gamma_{t-1}$ and the $OPEX_t$ indicator.

Results are contained in Table 12. Looking across the columns for the coefficient on the interaction term, we see that it is positive for three of our four liquidity measures, indicating that the relationship between gamma and liquidity is typically weaker in these option expiration weeks. While the effect of $Net\Gamma_{t-1}$ on the Realized Spread is stronger in Option Expiration weeks, it does not control for the baseline effect. After controlling for these periods, the main coefficients are still statistically and economically significant. Furthermore, the effect sizes are either economically close to, or larger than the effect sizes in the main table.

5.8 Sample Splits

In this section we investigate if particular stock characteristics are associated with our results.

5.8.1 Coverage of CBOE ISE data

Our variable $Net\Gamma_{t-1}$ is an estimate of the actual net-gamma at any point in time as we do not observe the totality of option trading on our sample stocks on U.S. exchanges. As we do not have any reason to suspect option contracts traded by retail investors are systematically opened or closed on particular exchanges (which would affect our estimate of net gamma), we expect our estimates to contain only classical measurement error that would bias our coefficient estimates toward zero. To verify that our results are robust to using stocks with various levels of data coverage, we calculate coverage by calculating the total option volume for each stock-day in our sample traded on the CBOE C1 exchange and the NASDAQ ISE exchange, and divide by the total option volume on that stock-day as recorded by OptionMetrics. We then take an average across all days in the sample for each stock to get a stock-level measure of data-coverage. We then estimate equation (2) separately for high-and low-coverage stocks.

The results are contained in Table 13. We can see that we lose the power to identify the effect using *Effective Spread* and *Price Impact* for stocks with low data coverage, however can see that the results for stocks with high data coverage are statistically and economically significant. The effect sizes for the *Effective Spread*, *Realized Spread*, and *Price Impact* are economically significant, a one standard deviation change in $Net\Gamma_{t-1}$ translates into a 11.85%, 10.75% and 16.35% change in the mean of the respective dependent variables. These results are either similar or stronger than our baseline results suggesting that measurement error is not driving our results.

5.8.2 Size and volume effects

We also consider whether the effects are concentrated in small stocks with relatively low trading volume. Our sample stocks are already quite large and well traded because of the requirements that the stocks have traded options. Nonetheless, it seems likely that the effects of delta-hedging on liquidity should be stronger in smaller stocks which have lower trading volume. We investigate these ideas in Table 14 and Table 15 respectively. Comparing the results for Table 14 Panels A and B we can see that the results exist for *Effective Spread*, *Realized Spread* and *Price Impact* for both large and small stocks. The effect sizes are approximately three times larger for small stocks at 9%, 12% and 13% of the mean value of the dependent variable. Comparing the results for Table 15 Panels A and B we see that volume seems more important than size for explaining heterogeneity in the main effect size. While the results still hold in high volume stocks for *Quoted Spread*, *Realized Spread*, and *Price Impact*, the effect sizes are much lower.

6 Conclusion

This paper shows that delta-hedging of net option positions driven by retail trading has a pervasive effect on market quality as measured by daily liquidity. These effects are stronger in stocks with low volume or higher information asymmetry and are not explained by informed trading, option expiration effects, or other channels. The impact depends on the net option position of market makers as well as the movement of the underlying stock which means that data from multiple option exchanges needs to be centralized and the measure of potential destabilization calculated dynamically for effective oversight by regulators.

References

Amihud, Y. and H. Mendelson (1980). Dealership market: Market-making with inventory. *Journal of financial economics 8*(1), 31–53.

Auh, J. K. and W. Cho (2022). Liquidation cascade and anticipatory trading: Evidence from the structured equity product market. *The Review of Asset Pricing Studies*.

Baltussen, G., Z. Da, S. Lammers, and M. Martens (2021). Hedging demand and market intraday momentum. *Journal of Financial Economics 142*(1), 377–403.

Barbon, A., H. Beckmeyer, A. Buraschi, and M. Moerke (2022). Liquidity provision to leveraged etfs and equity options rebalancing flows. *Available at SSRN*.

Boehmer, E., S. Chava, and H. E. Tookes (2015). Related securities and equity market quality: The case of cds. *Journal of Financial and Quantitative Analysis 50*(3), 509–541.

Boehmer, E., K. Fong, and J. J. Wu (2021). Algorithmic trading and market quality: International evidence. *Journal of Financial and Quantitative Analysis 56*(8), 2659–2688.

Brogaard, J. et al. (2010). High frequency trading and its impact on market quality. *Northwestern University Kellogg School of Management Working Paper 66*.

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

Bryzgalova, S., A. Pavlova, and T. Sikorskaya (2022). Retail trading in options and the rise of the big three wholesalers. *Available at SSRN*.

Chakrabarty, B., P. K. Jain, A. Shkilko, and K. Sokolov (2021). Unfiltered market access and liquidity: evidence from the sec rule 15c3-5. *Management Science 67*(2), 1183–1198.

Chakrabarty, B., B. Li, V. Nguyen, and R. A. Van Ness (2007). Trade classification algorithms for electronic communications network trades. *Journal of Banking & Finance 31*(12), 3806–3821.

Chen, H., S. Joslin, and S. X. Ni (2019). Demand for crash insurance, intermediary constraints, and risk premia in financial markets. *The Review of Financial Studies 32*(1), 228–265.

Comerton-Forde, C., T. Hendershott, C. M. Jones, P. C. Moulton, and M. S. Seasholes (2010). Time variation in liquidity: The role of market-maker inventories and revenues. *The journal of finance 65*(1), 295–331.

Conrad, J. (1989). The price effect of option introduction. *The Journal of Finance 44*(2), 487–498.

Demsetz, H. (1968). The cost of transacting. *The quarterly journal of economics 82*(1), 33–53.

Detemple, J. and P. Jorion (1990). Option listing and stock returns: An empirical analysis. *Journal of Banking & Finance 14*(4), 781–801.

Easley, D. and M. O'hara (1987). Price, trade size, and information in securities markets. *Journal of Financial economics 19*(1), 69–90.

Ellis, K., R. Michaely, and M. O'Hara (2000). The accuracy of trade classification rules: Evidence from nasdaq. *Journal of financial and Quantitative Analysis 35*(4), 529–551.

Ernst, T. and C. S. Spatt (2022). Payment for order flow and asset choice. Technical report, National Bureau of Economic Research.

Fedenia, M. and T. Grammatikos (1992). Options trading and the bid-ask spread of the underlying stocks. *Journal of Business*, 335–351.

Foley, S., A. Kwan, R. Philip, and B. A. Ødegaard (2022). Contagious margin calls: How covid-19 threatened global stock market liquidity. *Journal of Financial Markets 59*, 100689.

Glosten, L. R. and P. R. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics 14*(1), 71–100.

Hendershott, T., C. M. Jones, and A. J. Menkveld (2011). Does algorithmic trading improve liquidity? *The Journal of finance 66*(1), 1–33.

Ho, T. and H. R. Stoll (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial economics 9*(1), 47–73.

Holden, C. W. and S. Jacobsen (2014). Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *The Journal of Finance 69*(4), 1747–1785.

Hu, J. (2018). Option listing and information asymmetry. *Review of Finance 22*(3), 1153–1194.

Huang, S., B. Z. Yueshen, and C. Zhang (2021). Derivatives and market (il) liquidity. *Available at SSRN 3296699*.

Johnson, T. L. and E. C. So (2018). Asymmetric trading costs prior to earnings announcements: Implications for price discovery and returns. *Journal of Accounting Research 56*(1), 217–263.

Krinsky, I. and J. Lee (1996). Earnings announcements and the components of the bid-ask spread. *The Journal of Finance 51*(4), 1523–1535.

Kumar, R., A. Sarin, and K. Shastri (1998). The impact of options trading on the market quality of the underlying security: An empirical analysis. *The Journal of Finance 53*(2), 717–732.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 1315–1335.

Lakonishok, J., I. Lee, N. D. Pearson, and A. M. Poteshman (2007). Option market activity. *The Review of Financial Studies 20*(3), 813–857.

Lee, C. M., B. Mucklow, and M. J. Ready (1993). Spreads, depths, and the impact of earnings information: An intraday analysis. *The Review of Financial Studies 6*(2), 345–374.

Lee, C. M. and M. J. Ready (1991). Inferring trade direction from intraday data. *The Journal of Finance 46*(2), 733–746.

Mayhew, S., V. Mihov, et al. (2000). Another look at option listing effects. *Institute for Quantitative Research in Finance*.

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of economics and management science*, 141–183.

Ni, S. X., J. Pan, and A. M. Poteshman (2008). Volatility information trading in the option market. *The Journal of Finance 63*(3), 1059–1091.

Ni, S. X., N. D. Pearson, A. M. Poteshman, and J. White (2021). Does option trading have a pervasive impact on underlying stock prices? *The Review of Financial Studies 34*(4), 1952–1986.

Ronnie Sircar, K. and G. Papanicolaou (1998). General black-scholes models accounting for increased market volatility from hedging strategies. *Applied Mathematical Finance 5*(1), 45–82.

Sadka, R. and A. Scherbina (2007). Analyst disagreement, mispricing, and liquidity. *The Journal of Finance 62*(5), 2367–2403.

Scholes, M. and F. Black (1973). The pricing of options and corporate liabilities. *Journal of Political Economy 81*(3), 637–654.

Sorescu, S. M. (2000). The effect of options on stock prices: 1973 to 1995. *The Journal of Finance 55*(1), 487–514.

Thompson, S. B. (2011). Simple formulas for standard errors that cluster by both firm and time. *Journal of financial Economics 99*(1), 1–10.

Wilmott, P. and P. J. Schönbucher (2000). The feedback effect of hedging in illiquid markets. *SIAM Journal on Applied Mathematics 61*(1), 232–272.

Figure 1: Price Impact and Dollar Gamma for Meta Platforms Inc.

This figure contains in Panel A a Kernel Density Estimate of the distribution of price impact for the case when Dollar Net Gamma is positive and negative respectively. Panel B contains a time series graph of average Price Impact (in basis points) and dollar Net Gamma (in millions of USD) the time series are constructed as 20 period moving averages. Both Panel A and Panel B contain graphs constructed using data for Meta Platforms Inc. (formerly Facebook Inc.).

Panel A

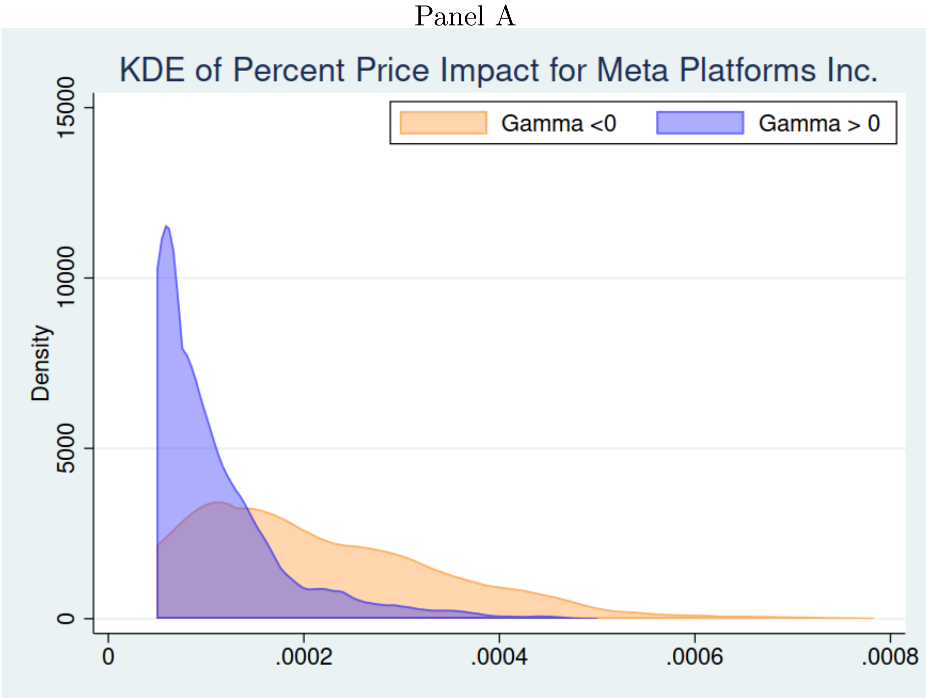

Panel B

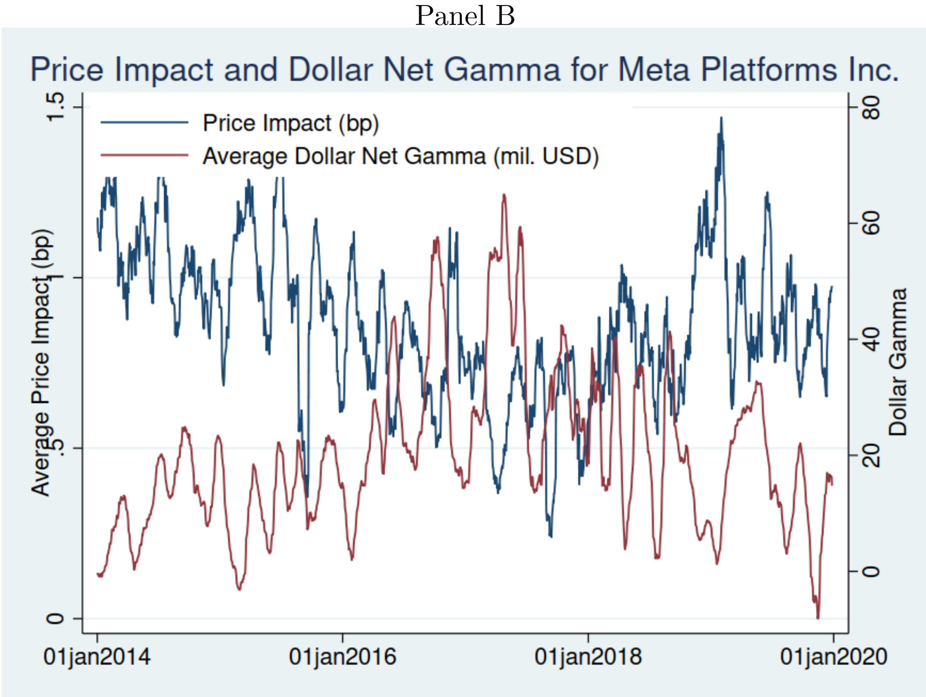

Panel A: Market-makers and proprietary traders as likely delta-hedgers

	mean	p1	p5	p25	p50	p75	p95	p99
Retail	0.96	0.83	0.91	0.95	0.97	0.98	1.00	1.00
Pro. Customer	0.04	0.00	0.00	0.02	0.03	0.05	0.09	0.17

Panel B: Market makers as likely delta hedgers

	mean	p1	p5	p25	p50	p75	p95	p99
Prop. Trader	0.19	0.02	0.05	0.12	0.19	0.24	0.34	0.42
Retail	0.78	0.54	0.63	0.72	0.78	0.85	0.93	0.97
Pro. Customer	0.03	0.00	0.00	0.01	0.03	0.04	0.07	0.14

Table 1: This table contains summary statistics on the percentage of non-delta-hedger volume accounted for by different groups of traders. Pro. Customer means Professional Customer while Retail trade is non professional customer trade. In Panel A the likely delta hedgers are market-makers and firm proprietary traders and in Panel B the likely delta hedgers are market-makers. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

	mean	sd	p1	p10	p25	p50	p75	p90	p99
$Net\Gamma$	0.001	0.006	-0.013	-0.003	-0.001	0.001	0.003	0.006	0.019
absolute return	1.813	2.087	0.007	0.210	0.567	1.280	2.396	3.882	9.334
quoted spread (%)	0.302	0.461	0.097	0.138	0.175	0.238	0.339	0.475	0.905
effective spread (%)	0.243	0.237	0.067	0.100	0.131	0.187	0.280	0.421	1.061
realized spread LR (%)	0.078	0.378	-0.373	-0.049	0.004	0.053	0.123	0.227	0.796
realized spread EMO (%)	0.051	0.213	-0.266	-0.050	-0.005	0.034	0.087	0.165	0.525
realized spread CLNV (%)	0.073	0.308	-0.303	-0.046	0.003	0.049	0.114	0.211	0.698
price impact LR(%)	0.160	0.412	-0.294	0.025	0.068	0.121	0.202	0.330	1.052
price impact EMO(%)	0.102	0.191	-0.175	0.012	0.045	0.084	0.138	0.218	0.540
price impact CLNV (%)	0.147	0.339	-0.228	0.023	0.061	0.109	0.183	0.302	0.924

Table 2: This table contains summary statistics of the main variables used in the paper. The statistics are calculated at the stock level, and the table contains averages of the firm level statistics. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.001	0.000	0.001
	(10.206)	(10.171)	(8.399)	(14.379)
$Net\Gamma_{t-1}$	-0.008	-0.033	-0.015	-0.029
	(-4.982)	(-7.009)	(-6.914)	(-7.202)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2243	2243	2243	2243
Avg. R2	0.835	0.850	0.201	0.394

Table 3: This table contains the results of estimating model (2). The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.001	0.000	0.001
	(12.934)	(13.590)	(10.870)	(14.179)
$Net\Gamma_{t-1}$	-0.013	-0.007	-0.006	-0.008
	(-7.008)	(-7.477)	(-7.230)	(-7.353)
$Net\Gamma_{t-1}$ x $EARN$	-0.017	-0.016	-0.006	-0.006
	(-3.823)	(-4.798)	(-2.702)	(-3.166)
EARN	0.000	0.000	-0.000	0.000
	(2.540)	(5.526)	(-1.681)	(7.257)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2243	2243	2243	2243
Avg. R2	0.835	0.850	0.204	0.409

Table 4: This table contains the results of estimating model (3). EARN is an indicator variable that takes a value of 1 for the period (t-3,t) where t is the firm's quarterly earnings date and is 0 otherwise. The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

Panel A: Low Dispersion of Analyst Forecasts

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.001	0.000	0.001
	(8.090)	(8.632)	(8.778)	(11.383)
$Net\Gamma_{t-1}$	-0.005	-0.005	0.000	-0.010
	(-3.481)	(-4.487)	(0.484)	(-5.291)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	1920	1920	1920	1920
Avg. R2	0.194	0.553	0.045	0.479

Panel B: High Dispersion of Analyst Forecasts

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.002	0.001	0.001	0.001
	(10.593)	(10.913)	(11.023)	(16.230)
$Net\Gamma_{t-1}$	-0.013	-0.063	-0.022	-0.050
	(-4.431)	(-7.179)	(-6.978)	(-7.500)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2243	2243	2243	2243
Avg. R2	0.835	0.850	0.201	0.394

Table 5: This table contains the results of estimating model (2) separately for groups of stocks with low (Panel A) and high (Panel B) dispersion of analyst forecasts. Low (High) analyst dispersion firms are defined as those with below (above) median levels of analyst dispersion. The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

Panel A: Low Analyst Coverage Stocks

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.002	0.001	0.001	0.002
	(10.076)	(10.551)	(10.000)	(18.187)
$Net\Gamma_{t-1}$	-0.008	-0.075	-0.037	-0.061
	(-2.747)	(-5.392)	(-5.338)	(-5.473)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	1467.000	1467.000	1467.000	1467.000
Avg. R2	0.048	0.235	0.098	0.048

Panel B: High Analyst Coverage Stocks

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.000	0.000	0.000
	(9.425)	(8.256)	(6.577)	(10.799)
$Net\Gamma_{t-1}$	-0.005	-0.003	-0.000	-0.004
	(-6.008)	(-5.975)	(-1.641)	(-6.758)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2243.000	2243.000	2243.000	2243.000
Avg. R2	0.835	0.850	0.201	0.394

Table 6: This table contains the results of estimating model (2) separately for groups of stocks with low (Panel A) and high (Panel B) analyst coverage. Low (High) analyst coverage firms are defined as those with below (above) median levels of analyst coverage. The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.001	0.000	0.001
	(10.213)	(10.155)	(8.615)	(14.442)
Residual Γ_{t-1}	-0.008	-0.037	-0.015	-0.032
	(-5.236)	(-6.941)	(-6.975)	(-7.131)
Information Γ_{t-1}	0.006	-0.019	0.017	-0.045
	(2.328)	(-8.374)	(5.811)	(-9.705)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2243	2243	2243	2243
Avg. R2	0.835	0.850	0.201	0.399

Table 7: This table contains the results of estimating model (2) using the decomposition of $Net\Gamma$ of (4). The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,633 stocks and 4,257,544 stock-date observations.

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.001	0.000	0.001
	(10.221)	(10.161)	(8.629)	(14.489)
Hedge Γ_{t-1}	0.002	-0.085	-0.023	-0.056
	(0.923)	(-4.385)	(-4.552)	(-4.486)
Information Γ_{t-1}	0.007	-0.021	0.019	-0.048
	(2.454)	(-7.355)	(5.911)	(-9.415)
t-τ Γ_{t-1}	-0.013	-0.025	-0.013	-0.024
	(-6.104)	(-5.325)	(-7.141)	(-6.938)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2243	2243	2243	2243
Avg. R2	0.835	0.850	0.201	0.399

Table 8: This table contains the results of estimating model (2) using the decomposition of $Net\Gamma$ of (5). The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,633 stocks and 4,257,544 stock-date observations.

	(EMO) Realized Spread	(CLNV) Realized Spread	(EMO) Price Impact	(CLNV) Price Impact
Constant	0.000	0.000	0.001	0.001
	(9.886)	(8.649)	(13.381)	(14.001)
$Net\Gamma_{t-1}$	-0.003	-0.010	-0.007	-0.032
	(-6.050)	(-6.893)	(-7.623)	(-7.196)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2243	2243	2243	2243
Avg. R2	0.219	0.221	0.357	0.388

Table 9: This table contains the results of estimating model (2) using the trade classification algorithms of Ellis et al. (2000) (EMO) and Chakrabarty et al. (2021) (CLNV) to calculate the Percent Realized Spread and Percent Price Impact. The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.001	0.000	0.001
	(10.599)	(10.663)	(10.373)	(15.786)
$Net\Gamma_{t-1}$	-0.009	-0.044	-0.007	-0.052
	(-3.985)	(-4.343)	(-4.160)	(-4.334)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2494	2494	2494	2494
Avg. R2	0.840	0.851	0.196	0.381

Table 10: This table contains the results of estimating model (2) using an alternative definition of $Net\Gamma_{t-1}$ which includes firm proprietary traders in the set of likely delta hedgers. The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

Panel A: Equal Weighted Portfolio returns

	A	B	C	D
	zero	p50	p75 (p25)	p90 (p10)
Return	0.78	-0.87*	-0.44	-0.56**
	(1.49)	(-1.93)	(-1.30)	(-2.19)

Panel B: Value Weighted Portfolio returns

	A	B	C	D
	zero	p50	p75(p25)	p90(p10)
Return	0.02	-0.46	-0.65**	-0.60***
	(0.25)	(-1.09)	(-2.22)	(-3.48)

Table 11: This table contains equal- and value-weighted annualised returns for portfolios that trade using rules derived from the data on $Net\Gamma_{t-1}$. The basic strategy takes a long (short) position in each stock when it's level of gamma is above or below a given threshold. Column A (B) contains results for a long short strategy with a long (short) position when $Net\Gamma$ is above (below) zero (the 50th percentile, p50). Column C (D) contains results for a long short strategy that takes a long position when $Net\Gamma$ is above the 75th percentile (90th percentile) and a short position when $Net\Gamma$ is below the 10th percentile (5th percentile). Portfolio returns are an averages across stocks in each time period to calculate an average return. Panel A contains results for an equal weight average across stocks while and Panel B contains results for a market-capitalization value weighted average across stocks in each time period. The sample period is from 2011 to 2019. Returns are annualized percentage returns. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.001	0.000	0.001
	(10.265)	(10.145)	(8.546)	(14.497)
$Net\Gamma_{t-1}$	-0.010	-0.043	-0.016	-0.039
	(-5.846)	(-7.856)	(-7.979)	(-8.194)
$Net\Gamma_{t-1}$ x $OPEX_t$	0.011	0.053	-0.004	0.060
	(6.860)	(7.202)	(-1.152)	(6.574)
$OPEX_t$	-0.000	0.000	-0.000	0.000
	(-1.774)	(1.261)	(-0.567)	(0.388)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2243	2243	2243	2243
Avg. R2	0.835	0.850	0.200	0.394

Table 12: This table contains the results of estimating model (6). The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

Panel A: Stocks with Low Data Coverage				
	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.001	0.000	0.001
	(9.252)	(9.296)	(8.304)	(12.541)
$Net\Gamma_{t-1}$	-0.011	-0.001	-0.010	0.001
	(-5.667)	(-1.382)	(-5.601)	(0.997)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg. N. Obs.	2243	2243	2243	2243
Avg. R2	0.835	0.850	0.201	0.394

Panel B: Stocks with High Data Coverage				
	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.002	0.001	0.001	0.001
	(9.461)	(9.449)	(6.924)	(13.933)
$Net\Gamma_{t-1}$	-0.004	-0.068	-0.021	-0.060
	(-1.687)	(-6.562)	(-6.290)	(-6.690)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg N. Obs.	1467	1467	1467	1467
Avg R2	0.048	0.235	0.098	0.048

Table 13: This table contains the results of estimating model (2) separately for groups of stocks with below median (Panel A) and above median (Panel B) coverage of data for estimating $Net\Gamma$. The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

Panel A: Below Median Size Stocks

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.002	0.001	0.001	0.002
	(9.989)	(11.397)	(10.835)	(19.925)
$Net\Gamma_{t-1}$	-0.017	-0.070	-0.033	-0.061
	(-4.960)	(-6.148)	(-6.143)	(-6.300)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg N. Obs.	1467	1467	1467	1467
Avg R2	0.048	0.235	0.098	0.048

Panel B: Above Median Size Stocks

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.000	0.000	0.000
	(11.712)	(9.460)	(9.303)	(11.950)
$Net\Gamma_{t-1}$	-0.000	-0.004	-0.001	-0.003
	(-0.480)	(-5.649)	(-3.687)	(-6.235)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg N. Obs.	2243	2243	2243	2243
Avg R2	0.835	0.850	0.200	0.391

Table 14: This table contains the results of estimating model (2) separately for groups of stocks with below median (Panel A) and above median (Panel B) dispersion of analyst forecasts. The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.

Panel A: Below Median Volume Stocks

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.002	0.001	0.001	0.002
	(10.273)	(10.803)	(10.149)	(19.275)
$Net\Gamma_{t-1}$	-0.014	-0.071	-0.034	-0.061
	(-3.776)	(-5.508)	(-5.411)	(-5.612)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg N.Obs.	1467	1467	1467	1467
Avg R2	0.047	0.233	0.099	0.046

Panel B: Above Median Volume Stocks

	Quoted Spread	Effective Spread	Realized Spread	Price Impact
Constant	0.001	0.000	0.000	0.000
	(9.503)	(7.884)	(6.933)	(10.727)
$Net\Gamma_{t-1}$	-0.003	-0.001	0.000	-0.001
	(-4.341)	(-2.702)	(0.996)	(-5.140)
Lag volatility controls	Yes	Yes	Yes	Yes
Lag liquidity controls	Yes	Yes	Yes	Yes
Avg N. Obs.	2243	2243	2243	2243
Avg R2	0.835	0.850	0.200	0.391

Table 15: This table contains the results of estimating model (2) separately for groups of stocks with below median (Panel A) and above median (Panel B) volume. The regression equation is estimated for each stock, and the table contains the cross-sectional average coefficient estimates. Avg. N. Obs. and Avg. R2 contain the cross-section average of the R-squared and number of observations from the stock-by-stock regressions. Standard errors are constructed from a covariance matrix for the average coefficients that is formed by clustering observations by date and firm with common shocks, as detailed in Thompson (2011). The lag volatility controls and lag liquidity controls use lags from 1 to 10 days before the realization of the dependent variable. The sample includes 2,639 stocks and 4,279,421 stock-date observations.