OFFICE OF GRADUATE STUDIES

NAME OF STUDENT: Yi Wang

DEGREE AWARDED: Master of Science in Management

ACADEMIC UNIT: Faculty of Business Administration

TITLE OF THESIS: Order Book Curvature as a Liquidity Measure: Evidence on High-Frequency Volatility in Crude Oil Futures

This thesis has been prepared under my supervision and the candidate has complied with the Master's regulations.

Signature of Supervisor

Oct. 19, 2025

Date

ORDER BOOK CURVATURE AS A LIQUIDITY MEASURE: EVIDENCE ON HIGH-FREQUENCY VOLATILITY IN CRUDE OIL FUTURES

BY

Yi Wang
Master of Engineering, Peking University, 2019
Bachelor of Economics, Zhejiang Normal University, 2017

THESIS
Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Management
in the Faculty of Business Administration
Lakehead University

Thunder Bay, Ontario, 2025

TITLE: Order Book Curvature as a Liquidity Measure: Evidence on High-Frequency Volatility in Crude Oil Futures

APPROVED BY

SUPERVISOR

Name	Mehdi Arzandeh	Signature		Date	Oct. 19, 2025

READER

Name	Xiaoping Song	Signature		Date	Oct 19, 2025

Declaration

I certify that I am the author of this project and that any assistance I received in its preparation is fully acknowledged and disclosed in this project. I have also cited any source from which I used data, ideas, or words, either quoted or paraphrased. I also certify that this paper was prepared by me specifically for this course.

No portion of the work referred to in this study has been submitted in support of an application for another degree or qualification to this or any other university or institution of learning.

Student Name Yi Wang Signature Yi Wang Date October 19, 2025

Abstract

This thesis proposes the limit order book (LOB) curvature as a new measure of liquidity. Using high-frequency data on crude oil futures traded on the CME, curvature is estimated from a power-law relation between normalized cumulative price distance and normalized cumulative depth, capturing how liquidity is distributed across the order book. This thesis uses a vector autoregression (VAR) model with impulse-response functions, examining the dynamic links among curvature, depth, spreads, and returns, while two-scale realized volatility (TSRV) is used to assess its predictive power for short-term volatility. Results show that curvature shocks reduce depth, widen spreads, and that curvature significantly improves volatility forecasts beyond traditional liquidity measures. The effect of scheduled news such as U.S. Energy Information Administration (EIA) Weekly Petroleum Status Report is insignificant after controlling the intraday seasonality, implying that liquidity dynamics are largely endogenous and resilient, with structural adjustments restoring equilibrium after transient imbalances rather than reacting to new information. Overall, curvature provides a unified non-linear perspective for understanding liquidity dynamics and volatility in high-frequency futures markets, offering new insight for market microstructure research and practice.

Acknowledgments

My deepest thanks to my supervisor, Professor Dr. Mehdi Arzandeh, for his professional guidance and for introducing me to the field. His insight, thoughtful comments, and fruitful discussions strengthened my understanding and sharpened my thinking, and his mentorship helped me navigate the complexities of research along the way.

I am also grateful to all the teachers, colleagues, and friends for their help throughout this journey, and to my family for their steadfast emotional and financial support. I could not have completed this degree without them.

Table of Contents

1. Introduction

Market liquidity determines how efficiently assets can be traded with minimal price impact, directly influencing market price efficiency, volatility, and overall market stability, related to well-functioning financial markets. Traditional liquidity measures can be broadly classified into trade-based and order-based approaches. Trade-based measures rely on transaction data such as trading volume, trade frequency, and turnover ratio to infer the historical cost of trading, while order-based measures use limit order book information, such as bid-ask spreads and quoted depths, to evaluate the availability of liquidity and potential transaction costs. Order-based measures have become the primary tools for evaluating liquidity conditions and information in the limit order book. Among these measures, the bid-ask spread (Roll, 1984), and the total quoted depth (Glosten, 1994), serve as standard tools for assessing market liquidity. However, these measures primarily focus on individual price levels or volumes of the book rather than the overall distribution of liquidity in the order book. In current studies of market liquidity, it is necessary not only to determine "how much" liquidity is available but also to understand "how" it is distributed in the order book. Building on these benchmarks, further measures such as order book slope (Næs & Skjeltorp, 2006), and relative liquidity (Valenzuela et al., 2015) aim to capture distributional aspects of liquidity and highlight the shape of liquidity across price levels. Yet these indicators remain linear approximations of the order book geometry, representing liquidity as a straight-line relation between price and depth and fall short of fully characterizing its whole geometric structure.

To address this need, this thesis introduces curvature, as a new liquidity measure that quantifies the non-linear geometric shape of the limit order book. Curvature characterizes the relationship between cumulative price distance from mid-quote and cumulative quantities of limit order book, providing new insights into volatility clustering and trading aggressiveness in high-frequency markets. Compared to traditional liquidity indicators, curvature offers several advantages. First, it captures the distribution of liquidity across the entire order book rather than only at selected price levels as in linear patterns, offering a complete liquidity profile. Second, it uses concave and convex order book shapes to introduce a new perspective in market microstructure research, intuitively representing how liquidity evolves under different market conditions. Third, it reflects how liquidity adjusts over time, especially during market announcements, and reveals how information shocks drive liquidity reallocation. Therefore, this study advances the analysis of market liquidity by extending from linear geometric measures such as spread, depth, and slope to a non-linear curvature-based framework**,** which captures a more complete and realistic representation of the order book's shape.

In market microstructure literature, order-based measures have been widely applied to explain intraday liquidity dynamics. For instance, order flow composition (Ahn et al., 2001) and depth at and away from the best quotes (Pascual & Veredas, 2010) are used as indicator of liquidity provision. In addition, the forecasting power of slope (Duong & Kalev, 2008), and relative liquidity (Valenzuela et al., 2015) are documented in predicting short-term volatility. However, they provide limited explanatory power when incorporated into models of volatility forecasting and dynamic interactions. To address these limitations, curvature is incorporated into a vector autoregression (VAR) model with impulse response functions (IRF) to analyze the short-run dynamic relationships among curvature, depth, spread, and returns. This design identifies how liquidity shocks propagate through the market and demonstrates the structural and resilient nature of liquidity adjustments. Second, curvature is integrated into a two-scale realized volatility (TSRV) forecasting model, providing empirical evidence that curvature significantly improves volatility predictions beyond traditional liquidity measures. Finally, by testing scheduled announcement effects of EIA Weekly Petroleum Status Reports, the thesis finds that information variables have limited significance once intraday seasonality is controlled. The insignificance of the announcement variable does not imply that information releases have no market impact. Rather, it suggests that short-term variations in curvature and depth are primarily driven by endogenous liquidity reallocation within the order book, rather than by new information. In other words, traders' adjustments to liquidity supply such as posting or cancelling limit orders in response to spread changes or depth imbalances, dominate over direct responses to scheduled announcements. This finding aligns with the idea that in high-frequency futures markets, liquidity dynamics are largely structural and self-correcting, with predictable information already incorporated before the report release. Collectively, these contributions advance the literature on market microstructure by connecting non-linear liquidity geometry, dynamic interactions, and volatility forecasting, offering a unified geometric perspective for understanding market resiliency in high-frequency futures markets.

2. Literature Review

2.1 Market Liquidity

Liquidity is an important characteristic of efficient financial market. It determines how easily securities can be traded without causing significant price changes. In theory, a perfectly liquid market allows securities to be exchanged for cash instantaneously and without cost. However, markets deviate from this ideal condition. Transaction costs, microstructure frictions, and information asymmetries introduce constraints that shape trading outcomes (Kyle, 1985). As a result, market liquidity is not determined solely by immediate transactions, but also by how orders are distributed across the book, and therefore, its measurement requires careful attention to market design, trading protocols, and regulatory rules. Key factors that influence liquidity include transaction costs (e.g., bid-ask spreads, price impact), market microstructure (e.g., order book depth, trading protocols) and regulatory factors (e.g., tick size rules, short-selling restrictions). Traditional liquidity measures capture some of these dimensions but remain limited, as they often reduce liquidity to a single snapshot, ignoring how it is distributed across multiple price levels.

2.2 Measures of Liquidity

Liquidity measurement is inherently complex, as no single metric fully captures its multiple dimensions. The literature generally categorizes liquidity measures into trade-based and order-based measures, each providing different insights into market conditions (Aitken & Comerton-Forde, 2003). Trade-based liquidity measures include trading volume, trading value, trade frequency, and turnover ratio. Higher trading volume is generally associated with greater liquidity (Chordia et al., 2001). Frequent trading suggests market efficiency but does not reflect order book conditions (Amihud & Mendelson, 1986). Trade-based liquidity measures rely on historical transaction data and are commonly used due to their simplicity and availability of data. However, they often reflect ex-post liquidity conditions rather than the ability to execute trades in the future (Pástor & Stambaugh, 2003). Trade-based measures are useful for analyzing liquidity trends but fail to capture the cost of trading and the availability of liquidity at different price levels, which is critical for institutional traders, as their trade size often exceed the available depth at the best quotes (Hasbrouck, 2009). Consequently, researchers often turn to order-based measures for a more detailed understanding of liquidity. Order-based liquidity measures utilize limit order book data, capturing the cost and depth of available liquidity in the market. Among these, the bid-ask spread, and quoted depth are the most common measures of the cost of liquidity, while the order book slope is often used to capture the distributional aspects of liquidity. In addition to these standard measures, this thesis introduces curvature, a novel indicator designed to fully characterize the geometric shape of the limit order book and provide additional insights into market dynamics.

2.2.1 Bid-Ask Spread

The bid-ask spread (BAS) is the difference between the best bid and ask prices, and it is a primary liquidity measure (Roll, 1984). If a trader wants to buy an asset, he must pay the market's ask price at which others are willing to sell, and if a trader wants to sell an asset, he can only sell at the market's bid price at which others are willing to buy. For instance, if a stock's market quote is $100 in the ask price (sell price), and $99 in the bid price (buy price), the bid-ask spread is $1. It represents the cost of trading and reflects market liquidity conditions. Studies consistently show that the BAS is a key indicator of market liquidity: a narrower spread indicates lower transaction costs and higher liquidity, whereas a wider spread suggests illiquidity in the trading markets (Stoll, 1989; Brockman & Chung, 1999). The greater the market volatility, the wider the spread and the higher the transaction costs, indicating that liquidity decreases when market uncertainty intensifies (Plerou et al., 2005).

Researchers have identified several key factors influencing spread levels. Price and trading volume are among the most significant determinants, and studies consistently find that the spreads are positively related to price and negatively related to trading volume (Demsetz, 1968; Tinic, 1972). The number of competing dealers and market structure also play a role, as spreads tend to be narrower in markets with more competition among dealers. The spreads are higher in market-maker dominated markets (Nasdaq) and lower in specialty markets (NYSE) (Stoll, 1989). Additionally, institutional ownership and investor composition influence spreads, with higher institutional ownership generally associated with lower bid-ask spreads due to increased market efficiency and liquidity (Chiang & Venkatesh, 1988). Risk factors, such as price variance and unsystematic risk, contribute to wider spreads, reflecting higher uncertainty in security valuation (Barnea & Logue, 1975). Lastly, the bid-ask spread is also sensitive to information asymmetry. When information asymmetry is higher, the bid-ask spread in the market is larger. Higher quality accounting information reduces spreads, and earnings announcements increase bid-ask spreads temporarily, consistent with increased uncertainty around new information (Callahan et al., 1997). Firms with greater analyst coverage tend to exhibit lower adverse selection costs and narrower spreads, even after controlling for other liquidity determinants (Brennan & Subrahmanyam, 1995).

Together, these findings underscore the bid-ask spread's role as a core measure of market liquidity, shaped by price, trading volume, market structure, risk factors, and information asymmetry. However, despite its widespread use, bid-ask spread alone may not fully capture liquidity fluctuations in modern electronic markets (Madhavan et al., 2005). For instance, spreads remain tight with small tick sizes, limiting their usefulness in capturing liquidity dynamics (Chordia et al., 2005). The spread only measure top-level liquidity and ignores the order book depth (Stoll, 2000). The limitation highlights the need for more comprehensive liquidity measures.

2.2.2 Order Book Depth

Limit order book (LOB) depth refers to the volume of outstanding buy and sell orders available beyond the best bid and ask prices at different price levels. It is a key order-based measure of market liquidity, as it reflects the market ability to absorb trades without significant price impact (Biais et al., 1995). Depth is a core dimension of liquidity as defined by Kyle (1985), who decomposed liquidity into resiliency, tightness, and depth. Among these, depth specifically refers to the quantity of orders available at different price levels and how well the market can handle large transactions without major price shifts (Kyle, 1985).

In early studies, LOB depth was often restricted to the best bid and ask levels due to data limitations (Frino et al., 2008; Chordia et al., 2011). However, with advancements in electronic trading and data availability, depth is now analyzed as the core dimension of LOB. It can provide essential information on the distribution of liquidity and the order flow dynamics across price levels (Biais et al., 1995). It can be regarded as the "buffer layer" of the market, which determines the capacity and price elasticity of the market in the face of large orders (Biais et al., 1995). Knowledgeable traders are more likely to submit limit orders than market orders in electronic markets (Bloomfield et al., 2005). Patient investors typically favor limit orders, whereas impatient investors are more likely to use market orders (Foucault et al., 2005). In addition, the LOB depth offers useful information that may be used to forecast future price swings, and professionals frequently profit from it (Harris & Panchapagesan, 2005; Madhavan et al., 2005). Compared to trade-based liquidity metrics, LOB depth can capture unexecuted liquidity (Harris & Panchapagesan, 2005). It can also reflect the market impact of large trades with minimal price movement, indicating higher market stability (Madhavan et al., 2005). It can help us understand how order book and order types, bids and asks, liquidity, and trader's behaviour shape the market outcomes, which is essential for Market Microstructure Theory (Easley et al., 2012; Kyle, 1985).

Empirically, LOB depth provides a more detailed view of market liquidity than simpler indicators like bid-ask spread, which only captures price-based liquidity without considering order volume. The distribution of LOB depth often follows a power-law pattern, where the number of orders decreases as prices deviate further from the best bid or ask (Bloomfield et al., 2005; Glosten, 1994). This pattern reflects market participants' expectations and risk preferences.

Despite its advantages, LOB depth has notable limitations. Many studies measure LOB depth at specific time points, providing only a snapshot rather than capturing liquidity's continuous evolution (Goettler et al., 2005). Traders frequently place and cancel large limit orders to manipulate perceived liquidity, especially in high-frequency trading environments (Harris & Panchapagesan, 2005). Prior research also suggests that LOB depth is not always evenly distributed between bid and ask sides. Some studies find symmetry (Aidov & Daigler, 2015), whereas others suggest that bid-side liquidity is often larger due to asymmetric order submission behavior (Yong, 2009). Therefore, measures of distribution of liquidity along the LOB prices are introduced in the literature.

2.2.3 Order Book Slope

Beyond bid-ask spread and limit order book depth, several other order-based liquidity measures have been developed to capture distributional aspects of market liquidity. One of them is order book slope. The slope represents the inclination of the order book across multiple price levels. A steeper slope indicates more concentrated liquidity while a flatter slope suggests more dispersed liquidity. The order book slope is systematically related to price volatility and the average daily slope of the order book is gentler with the greater disagreement among analysts (Næs & Skjeltorp, 2006). The slope of the order book is also an important microstructure parameter with predictive power, capable of significantly forecasting future price volatility, trade prices, and trading speed (Jain et al., 2011). Slope can be incorporated alongside relative liquidity and other liquidity measures to capture the shape of the order book. A higher slope indicates that liquidity is concentrated near the best quotes, allowing the market to absorb larger order flows with stronger liquidity; this, in turn, is associated with lower volatility (Jain & Jiang, 2014; Valenzuela et al., 2015).While slope improves upon traditional liquidity measures by considering liquidity dispersion, it still does not provide a geometric representation of the whole order book shape. This motivates the introduction of curvature as an alternative measure, which explicitly quantifies liquidity concentration vs. dispersion across multiple price levels.

2.2.4 Resiliency

Resiliency is another important dimension of market liquidity. It describes how rapidly liquidity conditions recover after being disturbed by temporary shocks. From a price recovery perspective, Kyle (1985) defines resiliency as the rate at which temporary pricing deviations caused by order-flow imbalances are corrected in the market. From a liquidity replenishment perspective, Harris (1990) describes a resilient market as one where new order arrive promptly to offset temporary liquidity shortages, restoring the order book to equilibrium. Building on these ideas, this thesis defines resiliency as the time required for one liquidity variable to recover following a specific liquidity shock. Early studies provided preliminary evidence that liquidity shocks tend to be short-lived, with order book spreads and depths typically reverting within seconds in equity markets and within minutes in agricultural commodity markets, indicating a strong short-term recovery capacity in modern financial systems. (Kempf et al., 2009; Large, 2007; Marshall et al., 2012) . However, these approaches cannot fully reflect the dynamics observed under algorithmic and high-frequency trading, where liquidity adjustments have intensified to occur at millisecond intervals. Lo and Hall (2015) advance this literature by employing millisecond data and estimating a high-frequency VAR model to examine how order submissions and cancellations affect the dynamics of the limit order book. By quantifying resiliency through impulse response functions expressed in clock time, their study reveals how quickly depth and spreads revert aftershocks, offering an intuitive picture of the market's capacity to absorb liquidity demand (Lo & Hall, 2015). Building on this framework, He et al., (2021) extend the analysis to agricultural futures during flash events. They show that even under extreme intraday price swings, market depth and spreads rebound rapidly, underscoring the strong resiliency of these markets and highlighting the role of order book structure in maintaining stability (He et al., 2021).

Together, these studies suggest that resiliency reflects an internal adjustment mechanism of the LOB, emphasizing how quickly liquidity recovers after being disturbed. Yet resiliency alone does not explain why such recovery unfolds in different ways across market states. This points to the importance of curvature, which links the static geometry of the order book to its dynamic resiliency, providing deeper insight into the internal drivers of market stability.

2.3 Curvature as a New Liquidity Measure

Early studies on market liquidity primarily treated depth as a static measure, often focusing on cumulative order quantities at specific price levels (Biais et al., 1995; Goettler et al., 2005; Hillman et al., 2001). However, more recent research suggests that order book depth is not fixed, and it dynamically evolves based on market conditions, trader sentiment, and external events (Bouchaud et al., 2008; Cont et al., 2010). Importantly, evidence shows that even limit order book steps beyond the best bid and ask spread to contribute substantially over 27% to price discovery in futures markets, highlighting that deeper layers of the book also embed valuable information (Arzandeh & Frank, 2019). A key insight from these studies is that the shape of the order book reflects traders' willingness to provide liquidity under different conditions. Unlike traditional measures such as spread, depth or slope, which only quantify liquidity at specific points, order book geometry reveals deeper insights into trader behavior and market resiliency (Tóth et al., 2011).

Current literature has observed that order book depth follows a power-law relationship with price distance from the mid-quote (Bouchaud et al., 2002; Farmer et al., 2008; Tóth et al., 2011). However, most studies have described the shape of the liquidity curve in qualitative terms (e.g., hump-shaped or V-shaped), rather than proposing a quantitative measure to systematically describe the geometry of the LOB curves (Bouchaud et al., 2002; Roşu, 2009; Tóth et al., 2011). For instance, empirical studies document that the average order book often exhibits a hump-shaped structure, with more orders submitted at prices away from the best bid and ask (Bouchaud et al., 2002). This hump shape arises because different types of traders contribute differently to liquidity provision, as patient traders prefer to post limit orders further from the mid-quote to reduce execution risk, while impatient traders concentrate their orders near the best quotes (Roşu, 2009). However, these studies fall short of introducing a consistent metric to quantify geometric properties of the order book, focusing instead on empirical patterns within specific markets or asset classes. Beyond hump-shaped observations, scholars demonstrate that the order book exhibits a V-shaped structure around the current price. Liquidity is scarce at the best bid and ask, but increases linearly as prices move away from the mid-quote (Tóth et al., 2011). This finding highlights that order book geometry is not merely a static snapshot, but instead reflects how traders submit, cancel, and fragment orders under uncertainty to sustain diffusive price behavior.

Several empirical studies have identified two contrasting structural forms of order books. This thesis categorizes them into concave shape and convex shape. A concave depth profile emerges when a larger proportion of limit orders are placed farther from the mid-quote, producing cumulative depth functions with diminishing marginal increases as prices move away from the mid-quote (Farmer et al., 2008). This concavity indicates greater trader disagreement and heightened uncertainty regarding the asset's fundamental value, and is often associated with informed traders who strategically place orders based on superior information or asymmetric market insights (Bouchaud et al., 2002; Tóth et al., 2011; Cont & Kukanov, 2017).

Conversely, order book depth exhibits a convex shape in certain asset classes, where greater proportion of orders are clustered near the mid-quote, particularly around the best bid and ask prices. This pattern suggests a higher level of agreement among market participants, implying lower uncertainty about price discovery. Under such conditions, incoming orders tend to concentrate further at top-of-book levels, increasing convexity. This behavior is typically associated with liquidity traders, who are more concerned with execution efficiency than informational advantages (Bouchaud et al., 2006; Menkveld & Yueshen, 2019).

Building on the empirical findings of concave and convex order books, this study introduces curvature as a novel liquidity measure that captures the geometric properties of order book depth. Traditional measures such as spread and depth describe liquidity at fixed points but fail to reflect how liquidity is distributed across different price levels or why it may vanish at the best quotes. Curvature provides a continuous, non-linear geometry-based representation of liquidity distribution, while resiliency characterizes its temporal recovery after shocks. Together, they form a unified framework linking the static and dynamic dimensions of market liquidity. This integration advances microstructure theory by illustrating how the structural allocation of liquidity conditions its subsequent adjustment process.

2.4 Liquidity in the Crude Oil Futures Market

Market participants trade in the futures markets with the intention of managing risk and hedging (Kyle, 1985). Futures markets frequently exhibit significant depth above and beyond the best bid and ask prices and this more extensive distribution over a range of price points provides insightful information about trading activity (Aidov & Daigler, 2015). Crude oil futures rank among the most actively traded contracts in global derivatives markets (Coppola, 2008). Beyond their role as a benchmark for physical oil inventories, these contracts are central to short-term price discovery, which explains why many international institutions incorporate futures data into oil price forecasts (Alquist et al., 2013). For market participants, crude oil futures serve not only as a tool for producers and consumers to hedge against price fluctuations, but also as an investment vehicle. Given oil's status as a globally traded commodity and the depth of its futures market, crude oil futures have become attractive to investors. Their historical co-movements with equities, bonds, and the U.S. dollar further strengthen their role as a portfolio diversification instrument (Hedi Arouri & Khuong Nguyen, 2010). As a result, the liquidity structure of the crude oil market is highly sensitive to geopolitical risks, macroeconomic uncertainty, and OPEC production decisions, often resulting in significant intraday volatility and frequent shifts in the order book (Lang & Auer, 2020).

A distinctive feature of the crude oil futures market is that its total order book depth is shaped not only by outright orders directly placed by traders but also by implied orders. The aggregate depth in the futures market includes two components: outright depth, which refers to observable limit orders submitted directly to the order book, and implied depth that are algorithmically generated by the exchange from spread and inter-month contracts. This distinction allows for a more detailed analysis of how market participants adjust their liquidity provision behavior between explicit and implied channels, particularly in response to external shocks such as scheduled announcements. The implied orders often dominate top-of-book liquidity when outright orders retreat under uncertainty, making the depth profile of crude oil futures fundamentally different from that of equities or other asset classes.

The literature on the response of futures markets to news announcements is far from settled. Some studies report little or no significant market reaction to U.S. macroeconomic news, arguing that prices largely incorporate expectations in advance (Kilian & Vega, 2011). However, these results are often based on monthly or low-frequency data, which may obscure short-lived but economically meaningful intraday adjustments. In contrast, exogenous information shocks such as macroeconomic announcements can significantly influence informed traders to be more aggressive in their trading strategies when reports are released during the trading sessions (Arzandeh et al., 2025). Energy prices and trading volumes adjust within minutes of monetary policy announcements or central bank surprises (Basistha & Kurov, 2015; Rosa, 2014). Liquidity studies similarly produce conflicting results. Most existing liquidity studies emphasize bid-ask spreads, particularly in equity and bond markets. Evidence shows that spread tends to widen before announcements and normalize shortly afterward (Balduzzi et al., 2001; Ederington & Lee, 1993; Rühl & Stein, 2015). However, spread alone does not capture all aspects of liquidity. Market participants also manage risk by adjusting quote sizes, which affects the depth of the order book. Research has shown that wider spreads often coincide with lower depth, indicating reduced willingness to provide liquidity (Lee, Mucklow, & Ready, 1993). Some researchers argue that liquidity is influenced by market-wide information factors such as macroeconomic news, which may explain correlated liquidity shifts across different asset classes (Chordia et al., 2011). This controversy provides a natural motivation for re-examining liquidity responses with new measures that move beyond spreads, such as order book curvature, which may reveal dimensions of liquidity not fully captured by traditional metrics.

Crude oil futures volatility exhibits distinct statistical properties that are widely documented in the literature. Volatility in crude oil markets is highly sensitive to external shocks. It tends to fluctuate between distinct regimes, often shifting in response to geopolitical tensions, weather-related demand changes, or production disruptions (Fong & See, 2002). Short-term shocks to volatility tend to be transitory. Empirical evidence suggests that oil market volatility reverts to its mean over a horizon of five to ten weeks, indicating a relatively stable structure over time (Pindyck, 2004). Another salient feature of crude oil returns is the frequent occurrence of price jumps. These large, sudden movements are often triggered by short-term supply-demand imbalances and are consistent with jump diffusion dynamics (Askari & Krichene, 2008). Subsequent studies confirm that such jump behavior coexists with continuous volatility processes in the crude oil market (Larsson & Nossman, 2011). Connecting volatility persistence and jump intensity with the evolving geometry of the order book provides new insights into how short-term shocks propagate and how quickly markets revert to equilibrium.

The U.S. Energy Information Administration (EIA) Weekly Petroleum Status Report, released every Wednesday at 10:30 a.m. EST, represents the most closely watched scheduled announcement in the crude oil market. Empirical literature shows that inventory surprises often lead to statistically significant responses in crude oil futures markets. For instance, futures and options prices react strongly to inventory announcements from EIA announcements (Miao et al., 2018). In addition, inventory information shocks rather than actual inventory changes negatively affect crude oil returns on the day the EIA releases the inventory information (Bu, 2014). Regarding limit order book dynamics, when volatility is expected to increase around weekly inventory announcements, the time-weighted order book slope decreases significantly, which also has predictive power for volatility one day ahead, highlighting that inventory news does affect the shape of the book (Tian et al., 2019). Overall, these findings highlight that the EIA weekly inventory report is a key source of exogenous information shocks in crude oil futures markets, with potential implications for both price dynamics and order book structure.

3. Rationale and Hypotheses

An increase in curvature is associated with a change in the composition of the aggregate depth, particularly, in the face of wider spreads. A highly convex order book reflects a depth that consists more of implied orders. Implied orders add depth exclusively to the first three steps[1] of the LOB in the crude oil futures market. A highly convex order book, thus, reflects the condition where outright positions along the 10 steps of the book[2] have shrunk and the depth composition is in favour of the implied orders. This occurs when market conditions are such that outright traders move away from the top of the book and/or exit their positions at the time of heightened uncertainty and information asymmetry to protect themselves from being "picked up" by informed traders. The spread in such conditions widens, luring more activity in the spread market. Together, a reduction in outright orders and an increase in implied orders rooted in increased activity in the spread market changes the composition of orders and increases the LOB convexity. The added depth from implied orders, however, is not sufficient to offset the reduction in outright orders. As a results, a positive shock to curvature is expected to reduce the overall depth. On the other hand, an increase in overall depth changes the composition of the orders such that more outright orders are present in the LOB compared to implied orders. Outright orders are more uniformly distributed along the 10 steps of the LOB, and they arrive at the time of more certainty and stability in the market. Therefore, the increase in depth reduces the convexity of the LOB as well as the bid-ask spread.

[1] In the CME limit order book, each step (or level) represents a unique price level at which orders are resting. The first three steps refer to the top three bid and ask levels closest to the mid-quote, which typically concentrate the majority of executable liquidity.

[2] The 10 steps of the book correspond to the ten best bid and ask price levels publicly disseminated by the CME. Depth and price information beyond the tenth level are, although visible to the public feed in the MBO data, are commonly eliminated in microstructure studies as a relatively small portion of contracts exists beyond the tenth step.

Hypothesis 1: Increase in curvature is associated with lower overall depths and wider spreads and vice versa.

Curvature reflects a nonlinear geometric distribution of liquidity. When curvature is high, indicating a more convex structure, the market exhibits stronger consensus among traders, and its ability to absorb order imbalances is enhanced, leading to lower short-term volatility. Conversely, when curvature decreases, indicating greater concavity, liquidity becomes concentrated at more distant price levels, reflecting higher disagreement and greater uncertainty, which results in stronger price fluctuations. Therefore, curvature is expected to show a negative relationship with realized volatility. Beyond this directional relationship, curvature captures structural variations across different price levels within the order book, so it contains incremental information relative to conventional liquidity metrics and may help explain short-term volatility fluctuations that are not captured by spread or depth alone.

Hypothesis 2: Curvature is negatively associated with market volatility, providing incremental explanatory power beyond traditional liquidity measures.

Market microstructure theory implies that if the report conveys new information, market participants will reprice and rebalance liquidity (Kyle, 1985; Glosten & Milgrom, 1985). In the crude oil futures market, the U.S. Energy Information Administration (EIA) Weekly Petroleum Status Report has been shown to trigger statistically significant reactions, affecting returns, volatility, and order book dynamics (Bu, 2014; Miao et al., 2018; Tian et al., 2019). While EIA reports are potential sources of information shocks, the thesis argues that short-term volatility variations in crude oil futures are primarily driven by endogenous liquidity reallocation rather than new information. In the presence of microstructure frictions, liquidity suppliers continuously adjust their limit orders in response to spread widening, depth thinning, or curvature shifts. These adjustments reflect structural rebalancing of liquidity, where traders manage execution risk and queue position rather than repricing based on new public information. Accordingly, if EIA reports introduce genuinely new information, volatility should increase on announcement days even after controlling for liquidity and intraday seasonality. Conversely, if market adjustments are largely endogenous, the announcement dummy should be insignificant once liquidity variables are included.

Hypothesis 3: After controlling for liquidity measures and intraday seasonality, the scheduled EIA announcements do not exert a significant incremental effect on short-term volatility in crude oil market.

4. Methodology

4.1 Curvature Calculation and Visualization

This study introduces curvature as a distributional measure of market liquidity, which is derived from the relationship between price distance (H, height) and cumulative quantity (Q, depth) for any snapshot of the book, using the power-law equation:

$$H = A \cdot Q^{\beta} \tag{4.1}$$

The parameter β is the curvature. It gauges how quickly price distance grows as cumulative depth increases. Let the mid-quote be mid=1/2 $(P_1^{bid}+P_1^{ask})$, which is the average of the best bid and best ask. At level j=1,…,10, the bid-side height is H_j^{bid} = mid - P_j^{bid} , and the ask-side height is H_j^{ask} = P_j^{ask} - mid. Cumulative depth is $Q_j=\sum_{k=1}^{j} q_k$, showing the total displayed quantity from the best quote up to level j. To remove units and allow comparisons across time, each snapshot is normalized by its tenth level: $H_j^{\%}=H_j/H_{10}$ and $Q_j^{\%}=Q_j/Q_{10}$. This normalization results in $H_{10}^{\%}=Q_{10}^{\%}=1$ within a snapshot, which places all curves on a common 0-1 scale. In addition, heights and depths are multiplied by 100 to express them in percentages.

Curvature is estimated by snapshot, that is, for each snapshot, β is obtained from the log-log regression in the H on Q direction:

$$\ln H_j^{\%} = \alpha + \beta \ln Q_j^{\%} + \varepsilon_j \tag{4.2}$$

where j=1,…,10. Bid and ask side are estimated separately using robust standard errors. Log-transformation helps stabilize variance and handle extreme values, thereby improving the reliability of regression models.

The economic interpretation follows directly from the normalized curves. Because the plots place $Q^{\%}$ on the horizontal axis and $H^{\%}$ on the vertical axis, the power law $H^{\%} \propto (Q^{\%})$ implies: If

$\beta>1$, it represents a convex limit order book indicating that a relatively large portion of orders is concentrated near the mid-quote. This pattern suggests that traders have a greater degree of consensus about the true value of the asset. Under such conditions, new incoming orders also tend to accumulate close to the mid-quote, which further increases the convexity of the order book. These environments are typically dominated by liquidity traders who are more willing to provide liquidity close to the best quotes, reflecting lower uncertainty and stronger execution willingness. If $\beta<1$, it represents a concave limit order book, indicating that a greater number of limit orders are placed farther away from the mid-quote. This reflects higher levels of disagreement among market participants regarding the asset's fair value. In such cases, incoming orders are also submitted at prices further from the mid-quote, deepening the concave shape. These conditions are commonly associated with informed trading and increased levels of information asymmetry. If $\beta=1$，a linear relationship exits that indicates uniform distribution. This behavioral distinction forms the theoretical basis for interpreting curvature as a measure of market consensus and information structure.

To display representative shapes while preserving intraday structure, curves are aggregated and visualized by trading day. For each snapshot, the ten points ($Q_j^{\%}$, $H_j^{\%}$) are first computed. Within each trading day, pointwise medians across snapshots are then taken at each level j, yielding a daily median curve with ten points. All 272 daily curves are plotted as thin gray polylines, and the cross-day pointwise median serves as a red benchmark curve overlaid on top. The axes are labeled "Cumulative Quantity (%)" and "Price Distance (%)." This pipeline ensures that the normalized curvature used for estimation is consistent with the percent-axis curves used in visualization, and that convexity/concavity corresponds one-for-one to the magnitude of β.

4.2 Vector Autoregression and Impulse Response Functions

This study employs a vector autoregression (VAR) framework to analyze the interrelationships among price variation, order-book geometry, depth and spread. The VAR model can flexibly capture the feedback effects among multiple time series variables, while the impulse response function (IRF) explains the transmission path of shocks over time. Following He et al. (2021), we include returns in the VAR framework, which is also a conventional choice in market microstructure VARs (Lütkepohl, 2005; Hasbrouck, 1991). The VAR and IRF constructed in this paper are used to interpret the dynamic interactions among returns, order book curvature, market depth, and bid-ask spread. Using the four-dimensional endogenous vector:

$$\mathbf{y}_t = \begin{bmatrix} \mathrm{RETURN}_t \\ \beta_t \\ \mathrm{DEPTH}_t \\ \mathrm{SPREAD}_t \end{bmatrix} \tag{4.3}$$

where RETURN_t is the first difference of the log mid-quote, calculated by $\ln(\mathrm{mid}_t) - \ln(\mathrm{mid}_{t-1})$, β_t is limit order book curvature calculated by the average of β_t^{bid} and β_t^{ask}, DEPTH_t is order book depth and it is the average depth of $\ln Q_t^{bid}$ and $\ln Q_t^{ask}$, and SPREAD_t is the relative spread, standardized by the mid-quote, that is, relative spread = $(P_t^{ask}-P_t^{bid})$ / mid-quote *100.

The VAR(p) model is then expressed as:

$$\mathrm{y}_t = \mathrm{c} + \sum_{\ell=1}^{p} \mathrm{A}_\ell \, \mathrm{y}_{t-\ell} + \mathrm{u}_t \tag{4.4}$$

where c is a constant vector, $A_\ell(\ell=1,\ldots,p)$ are coefficient matrices of dimension 4×4, and u_t is a residual vector with $E(u_t)=0$ and covariance matrix Σu.

The lag length p is determined using standard information criteria, including the Akaike Information Criterion (AIC), the Bayesian Information Criterion (BIC), and the Hannan-Quinn

Criterion (HQ) (Akaike, 1974; Hannan & Quinn, 1979; Schwarz, 1978). The selected p balances model fit against parsimony and ensures stable estimation.

To interpret the dynamic responses of variables to structural innovations, following (Lütkepohl, 2005), the residual covariance matrix Σu is orthogonalized using Cholesky decomposition, $\Sigma u = PP'$. The IRF of a unit shock in variable k on variable m at horizon h is then defined as:

$$\mathrm{IRF}_{m,k}(h) = \mathbf{e}_m' \, \Psi_h \, P \, \mathbf{e}_k \tag{4.5}$$

where Ψ_h denotes the moving-average coefficient matrices, and e_m and e_k are selection vectors identifying the responding and shocked variables, respectively.

4.3 Predicting Volatility

Building on the VAR in the previous section, we have identified the dynamic transmission among return, curvature, depth and spread. We now turn to volatility forecasting, where the goal is to quantify and predict latent "true" volatility. When market microstructure noise is present, standard realized volatility (RV) estimators become inconsistent at higher sampling frequencies, since they capture both the underlying price variation and the variance induced by noise (Bandi & Russell, 2008). Realized volatility that is not explicitly noise-robust is prone to upward bias and strong sampling-frequency sensitivity due to high-frequency microstructure noise (Andersen et al., 2001; Zhang et al., 2005), which can treat noise as signal in a multivariate system and weaken structural identification and the interpretability of propagation paths. To address this bias, this thesis adopts the two-scale realized volatility (TSRV) estimator, which combines realized variances from two sampling schemes: a 1-second sample "fast" RV for noise variance estimation, and a 15-minute sample "slow" RV that reduces noise contamination (Zhang et al., 2005; Aït-Sahalia et al., 2011).

For each 15-minute interval, we first compute the all-sample realized variance based on the mid-quote:

$$\mathrm{RV}^{\text{fast}} = \sum_{i=1}^{n_{\text{fast}}} \left(\ln \mathrm{mid}_i - \ln \mathrm{mid}_{i-1}\right)^2 \tag{4.6}$$

where n_{fast} is the number of returns in the interval. Although $\mathrm{RV}_{\text{fast}}$ is upward biased due to noise, it provides an estimator of the noise variance. Specifically, the variance of the noise term is estimated as:

$$\hat{\eta}^2 \;=\; \frac{1}{2\, n_{\text{fast}}} \mathrm{RV}^{\text{fast}} \tag{4.7}$$

Next, by applying a moving-window sparse sampling scheme within each 15-minute interval, one observation every 30 seconds is captured following Valenzuela et al. (2015). For each subsample k=1,…,K, the slow realized variance is computed as:

$$\mathrm{RV}_{\text{slow}}^{(k)} = \sum_{i=1}^{n_{\text{slow}}^{(k)}} \left(\ln \mathrm{mid}_i^{(k)} - \ln \mathrm{mid}_{i-1}^{(k)}\right)^2 \tag{4.8}$$

where $n_{\text{slow}}^{(k)}$ denotes the number of returns in the k-th slow subsample, formed within each 15-minute window by taking every k-th observation starting at offset k.

The interval-level slow realized variance is then obtained as the average across all valid subsamples:

$$\mathrm{RV}_{\text{slow}} = \frac{1}{K} \sum_{k=1}^{K} \mathrm{RV}_{\text{slow}}^{(k)} \tag{4.9}$$

The TSRV estimator of the integrated variance of the latent efficient price process is obtained by removing the estimated noise bias from the slow-frequency RV:

$$\mathrm{TSRV} = \mathrm{RV}_{\mathrm{slow}} - \frac{n_{\mathrm{slow}}}{n_{\mathrm{fast}}}\,\mathrm{RV}_{\mathrm{fast}} \tag{4.10}$$

where n_{slow} is the average number of returns in each slow-frequency subsample, $n_{slow} = 1/K\sum_{k=1}^{K} n_{slow}^{(k)}$. The ratio n_{slow}/n_{fast} scales the fast-sample noise estimate so that it matches the effective sample size of the slow scheme.

Following Zhang et al. (2005), we assume the noise terms η_t are i.i.d. with zero mean and constant variance. In this study, we construct subsamples by selecting one observation every 30 seconds within each 15-minute interval, which yields approximately K = 30 valid subsamples per interval. We further require at least 30 observations in each subsample to ensure a balance between reducing noise contamination and retaining sufficient price information. By applying the TSRV estimator to high-frequency LOB data for crude oil futures, this study obtains a robust measure of integrated volatility even in the presence of significant market microstructure noise. This approach mitigates the divergence problem of RV at ultra-high frequencies and provides a stable and unbiased volatility input for subsequent event-window regressions.

Intraday volatility patterns are known to follow strong seasonal cycles (Nikkinen & Rothovius, 2019). To control for these effects, we transform the raw trading timestamps into a continuous “session time” that starts at 17:00 CT and extends to the next day’s close, ensuring that no trading day wraps around within a session. This thesis trims the trading session to 17:10 to 15:50 CT, subtracting the first and last 10 minutes from the opening and closing periods to reduce the influence of opening sequence and settlement process of the exchange. Next, we create 30-minute slot dummy variables from the session open in the regression models to control for intraday seasonality.

We estimate a series of linear regressions to examine the incremental explanatory power of order book variables in predicting short-term volatility. The dependent variable in all specifications is the one-step-ahead two-scale realized volatility ($TSRV_{t+1}$). Explanatory variables are sequentially added across models (I-VI) to identify their marginal contribution to model fit. The benchmark specification (Eq. 4.11) regresses the one-step-ahead TSRV only on its own lag, capturing the persistence of realized volatility. In the next specification (Eq. 4.12), we incorporate the curvature of the order book by adding both bid-side and ask-side curvature measures. Building on this, Eq. (4.13) further introduces the average order book depth, thereby accounting for the quantity dimension of liquidity. Eq. (4.14) extends the model by including the bid-ask spread and Eq. (4.15) augments the regression with the slope variable. Eq. (4.16) adds Wed_t dummy measure. This sequential modeling approach allows us to identify the marginal contribution of each liquidity dimension to the predictive power for short-term volatility. In addition, it reveals if curvature adds explanatory power to the volatility prediction after controlling for other determinants of volatility.

$$\mathrm{TSRV}_{t+1} = \alpha + \gamma_1\,\mathrm{TSRV}_t + \varepsilon_t \tag{4.11}$$

$$\mathrm{TSRV}_{t+1} = \alpha + \gamma_1\,\mathrm{TSRV}_t + \gamma_2\,\beta_t^{\mathrm{bid}} + \gamma_3\,\beta_t^{\mathrm{ask}} + \varepsilon_t \tag{4.12}$$

$$\mathrm{TSRV}_{t+1} = \alpha + \gamma_1\,\mathrm{TSRV}_t + \gamma_2\,\beta_t^{\mathrm{bid}} + \gamma_3\,\beta_t^{\mathrm{ask}} + \gamma_4\,\mathrm{Depth}_t + \varepsilon_t \tag{4.13}$$

$$\mathrm{TSRV}_{t+1} = \alpha + \gamma_1\,\mathrm{TSRV}_t + \gamma_2\,\beta_t^{\mathrm{bid}} + \gamma_3\,\beta_t^{\mathrm{ask}} + \gamma_4\,\mathrm{Depth}_t + \gamma_5\,\mathrm{Spread}_t + \varepsilon_t \tag{4.14}$$

$$\mathrm{TSRV}_{t+1} = \alpha + \gamma_1\,\mathrm{TSRV}_t + \gamma_2\,\beta_t^{\mathrm{bid}} + \gamma_3\,\beta_t^{\mathrm{ask}} + \gamma_4\,\mathrm{Depth}_t + \gamma_5\,\mathrm{Spread}_t + \gamma_6\,\mathrm{Slope}_t + \varepsilon_t \tag{4.15}$$

$$\mathrm{TSRV}_{t+1} = \alpha + \gamma_1\,\mathrm{TSRV}_t + \gamma_2\,\beta_t^{\mathrm{bid}} + \gamma_3\,\beta_t^{\mathrm{ask}} + \gamma_4\,\mathrm{Depth}_t + \gamma_5\,\mathrm{Spread}_t + \gamma_6\,\mathrm{Slope}_t + \gamma_7\,\mathrm{WED}_t + \varepsilon_t \tag{4.16}$$

In addition to spread, depth, and public announcements, slope is found to explain volatility in the microstructure literature (Næs & Skjeltorp, 2006; Jain & Jiang, 2014; Valenzuela et al., 2015). In each snapshot, the slope on each book side is defined as the average discrete elasticity of log-quantity with respect to relative price distance across the top ten levels (the bid side uses absolute price gaps) (Næs & Skjeltorp, 2006), with the mid-quote denoted by mid as given in Eq. (4.17) and Eq. (4.18):

$$\mathrm{slope}_{\mathrm{bid}} = \frac{1}{10}\left[\frac{\ln Q_1^{\mathrm{bid}}}{\frac{P_1^{\mathrm{bid}}}{mid} - 1} + \sum_{j=2}^{10}\frac{\ln Q_j^{\mathrm{bid}} - \ln Q_{j-1}^{\mathrm{bid}}}{\frac{P_j^{\mathrm{bid}}}{P_{j-1}^{\mathrm{bid}}} - 1}\right] \tag{4.17}$$

$$\mathrm{slope}_{\mathrm{ask}} = \frac{1}{10}\left[\frac{\ln Q_1^{\mathrm{ask}}}{\frac{P_1^{\mathrm{ask}}}{mid} - 1} + \sum_{j=2}^{10}\frac{\ln Q_j^{\mathrm{ask}} - \ln Q_{j-1}^{\mathrm{ask}}}{\frac{P_j^{\mathrm{ask}}}{P_{j-1}^{\mathrm{ask}}} - 1}\right] \tag{4.18}$$

where P_j and Q_j are the level-j price and visible depth on bid- and ask-side.

By comparing adjusted R^2 across Models I to VI, we assess the incremental explanatory contributions of curvature, market depth, bid-ask spread, order book slope, EIA announcements, and seasonality effects.

5. Data Collection

5.1 Date and Time

This study exclusively focuses on Light Sweet Crude Oil futures (ticker symbol: CL) traded on the Chicago Mercantile Exchange (CME). The study uses the CME's Market by Order (MBO) data that provides every order's detail required to reconstruct the LOB of CL with nanosecond granularity. The CL contract is one of the most actively traded energy futures contracts globally, serving as a benchmark for U.S. crude oil prices. Each contract represents 1,000 U.S. barrels of West Texas Intermediate (WTI) crude oil, quoted in U.S. dollars per barrel, with a minimum tick size of $0.01, equivalent to $10 per contract. Trading is conducted electronically via the CME Globex platform, offering 23-hour trading, from 5:00 p.m. to 4:00 p.m. CT, from Sunday through Friday. The sample period in this study spans from December 3, 2018, to December 31, 2019, covering a total of 272 trading days.

5.2 Roll Dates

Futures traders at the Chicago Mercantile Exchange are allowed to shift their open positions from one contract month to another at any time. In practice, they typically complete this roll forward process as the current contract nears expiration and market liquidity begins to decline. According to CME specifications, WTI crude oil futures are listed for trading in every calendar month, including all twelve months of the current year and each month over the next ten years. Trading in each WTI crude oil futures contract ends on the third business day prior to the 25th calendar day of the month preceding the delivery month. If the 25th day falls on a non-business day, the last trading day is adjusted to the fourth business day prior. In this study, we define a product-specific roll date rule based on trading activity. When the daily open interest of the front-month contract falls below that of the next-nearest-month contract, we identify this point as the "roll date". From that day onward, the next-nearest-month contract is treated as the new front-month contract for analysis purposes.

It is important to note that calendar spread trading is widely practiced in the crude oil futures market. Traders often engage in strategies that involve simultaneously buying and selling contracts with different expiration dates, typically for the purposes of arbitrage or hedging. These transactions give rise to a distinct mechanism known as implied orders. Implied orders refer to limit orders that are not directly submitted to the order book of a specific contract. Instead, they are algorithmically generated by the CME's matching engine based on calendar spread quotes placed in related contracts. Implied orders represent real and executable liquidity, and they participate in the matching process alongside traditional orders. Therefore, they play a critical role in the market's microstructure. In contrast, outright orders are explicit limit orders that are directly submitted to a contract's order book, reflecting traders' actual intended price and quantity. Currently CME disseminates a two-step deep implied order book for the crude oil futures. We reconstruct the combined limit order book by merging the outright and implied order books, consistent with the approach of Arzandeh and Frank (2019).

5.3 Data Cleaning

To ensure data quality and computational efficiency, the raw LOB data is pre-processed using Stata with the following step: (1) Identify and remove duplicate observations that share the same timestamp and price level, in order to prevent double counting. (2) Filter out erroneous records that result in temporarily negative bid-ask spreads due to nanosecond-level data capture.

For econometric analysis, the thesis employs two sampling frequencies. First, the 1-second snapshots are used to retain the rich information embedded in high-frequency order book dynamics, which allows us to capture immediate liquidity responses. Under the first sampling frequency and after the cleaning process, the dataset consists of 1-second snapshots over the 272 trading days. Therefore, the first sample contains 19,114,926 observations, capturing prices and quantities of levels 1 through 10 on both bid and ask sides. Second, the data is collapsed into 15-minute snapshots, where all the order book variables are recorded as last-tick levels within each 15-minute interval. This procedure yields 24,947 15-minute intervals, which serve as the primary unit of analysis for regression models. The dual-frequency sampling strategy follows the market microstructure literature (Zhang et al., 2005; Bandi & Russell, 2008), balancing the informational advantage of high-frequency data against the need to mitigate microstructure noise.

6. Empirical Findings

6.1 Curvature Estimates and Visualization

Across the ten best bid and ask levels, the average price centers around $5,645 to $5,655 per barrel in Table 6.1. The standard deviations are nearly identical across all levels (419.98 to 419.99), indicating that price volatility is uniformly distributed across the depth of the book, reflecting the consistency of order placement within the top 10 layers of the limit order book.

From quantity side, Table 6.2 shows that quantities are modest at the best quotes (Bid 1 = 22.59; Ask 1 = 22.84), with standard deviations around 15.7, implying high variability at the top of the book. From Level 2 to Level 5, order sizes increase significantly (around 31 to 46 contracts per side), suggesting clustering of liquidity just behind the best quotes. Beyond Level 6, depth stabilizes around 30 contracts per level on both sides, with standard deviations (23 to 25), reflecting a relatively flat depth profile at deeper levels. Based on the mean and standard deviation of the full-sample crude oil LOB depth, the total bid and ask quantities are broadly balanced. At the near levels, bid quantities are slightly lower than the corresponding ask quantities, indicating a slight order imbalance.

Table 6.1 Mean and Standard Deviation of the Crude Oil LOB Price

Bid 1	Ask 1	Bid 2	Ask 2	Bid 3	Ask 3	Bid 4	Ask 4	Bid 5	Ask 5
5645.47	5646.55	5644.47	5647.55	5643.46	5648.56	5642.46	5649.56	5641.45	5650.57
(419.99)	(419.98)	(419.99)	(419.98)	(419.99)	(419.98)	(419.99)	(419.98)	(419.99)	(419.98)
Bid 6	Ask 6	Bid 7	Ask 7	Bid 8	Ask 8	Bid 9	Ask 9	Bid 10	Ask 10
5640.45	5651.58	5639.44	5652.58	5638.43	5653.59	5637.43	5654.59	5636.42	5655.60
(419.99)	(419.98)	(419.99)	(419.98)	(419.99)	(419.98)	(419.99)	(419.98)	(419.99)	(419.98)

Price unit: U.S. cents per barrel
Obs. = 19,114,926

Table 6.2 Mean and Standard Deviation of the Crude Oil LOB Quantity

Bid 1	Ask 1	Bid 2	Ask 2	Bid 3	Ask 3	Bid 4	Ask 4	Bid 5	Ask 5
22.59	22.84	46.21	46.31	42.14	42.08	31.87	32.00	31.19	31.35
(15.70)	(15.75)	(22.49)	(23.57)	(25.90)	(26.89)	(21.44)	(22.17)	(22.08)	(22.66)
Bid 6	Ask 6	Bid 7	Ask 7	Bid 8	Ask 8	Bid 9	Ask 9	Bid 10	Ask 10
30.62	30.88	30.54	30.96	29.99	30.35	30.16	30.50	30.20	30.66
(22.83)	(23.49)	(23.27)	(24.34)	(23.52)	(24.38)	(24.23)	(25.14)	(24.60)	(25.42)

Quantity (Contracts) unit: 1,000 barrels
Obs. = 19,114,926

Table 6.3 reports the descriptive statistics for the estimated beta coefficients (β^{bid} and β^{ask}), corresponding p-values, and R^2 values obtained from individual regressions. Overall, the bid-side regression results provide evidence of statistical reliability and explanatory power. β^{bid} is 1.106 on average with a standard deviation of 0.340. The mean of p-value is 0.00002, and the 99th percentile p-value observed is 0.00021, showing the statistical significance of the bid-side regressions. In terms of model fit, the mean and median R^2 are 0.960 and 0.988, respectively. The 99th percentile approaches 0.999 and the minimum value is 0.737.

The estimation results on the ask side also indicates stable model performance in terms of coefficient magnitude and model fit. The estimated coefficients β^{ask} have a mean of 1.120 and a median of 1.062. The interquartile range spans from 0.920 (25%) to 1.228 (75%). The average p-value is 0.00002, and the 99th percentile remains 0.00021. The mean and median R^2 values are 0.961 and 0.988, respectively. The R^2 values are concentrated near one, with a minimum of 0.734. Notably, according to the distribution of beta estimates from 15-minute snapshots, more than 50% of cases exhibit beta > 1. However, the distribution is relatively balanced, and no single concave or convex pattern dominates across the entire sample.

Table 6.3 Descriptive Statistics of Curvature Estimates

Variable	Mean	Std. Dev.	1%	25%	50%	75%	99%
β^{bid}	1.106	0.340	0.547	0.917	1.061	1.217	2.258
P-value	0.00002	0.00017	9e-15	1e-10	1e-8	3e-06	0.00021
R^2	0.960	0.061	0.737	0.958	0.988	0.995	0.999
β^{ask}	1.120	0.371	0.552	0.920	1.062	1.228	2.372
P-value	0.00002	0.00030	8e-15	1e-10	1e-8	3e-06	0.00021
R^2	0.961	0.061	0.734	0.959	0.988	0.995	0.999

Obs. = 24,947 (after 15-minute interval)
Note: 1%, 25%, 50%, 75% and 99% refer to the 1^{st} 25^{th} 50^{th} 75^{th} and 99^{th} percentiles, respectively. (See Appendix Table A.1 Distribution of Curvature Estimates for full summary statistics.)

Figure 6.1 plots the daily median bid-side order book curves over 272 trading days. The horizontal axis reports the cumulative quantity percentage obtained by summing bid sizes from the best bid outward through the 10-step limit order book. The vertical axis reports the cumulative height percentage, defined as the cumulative price distance from the best bid outward across those steps. Thin gray lines show the median curve for each day, and the thick red line shows the overall median across the 272 days.

Generally, when it turns to convex order book, in regions near the best bid, each additional unit of order volume corresponds to a relatively small change in price, reflecting a high degree of liquidity concentrated near the top of the table. Beyond the top levels, limit orders are also distributed in an organized manner across deeper price levels. As prices move away from the mid-quote, the additional quote required for each unit of price change tends to decline, indicating increased sensitivity of depth to price movements. The characteristics of curves in Figure 6.1 show the same pattern: when cumulative quantity increases from 20% to 40%, cumulative height rises by only about 10%, whereas cumulative quantity increases from 80% to 100%, cumulative height must rise by more than 40%. These numbers indicate strong clustering of bids near the mid-quote and substantial depth at the top of the book. Farther from the top, traders need much larger price concessions to obtain the same amount of liquidity, so orders become sparse at deeper levels. This pattern fits the crude oil futures market: heavy trading and queue competition near the top, as many traders posting near the mid-quote to raise execution probability and queue priority, and a smaller group placing orders deeper to reduce adverse-selection risk. Together these forces create a flat front section and a steep tail, producing the observed convex profile.

We also observe certain concave shapes, which indicate that as the price moves further away from the best bid, the order volume associated with each tick increases. This implies that, relative to the best bid, a substantial number of orders are accumulated at deeper price levels. In such cases, traders exhibit greater disagreement about the market, and the curvature measure falls below 1.

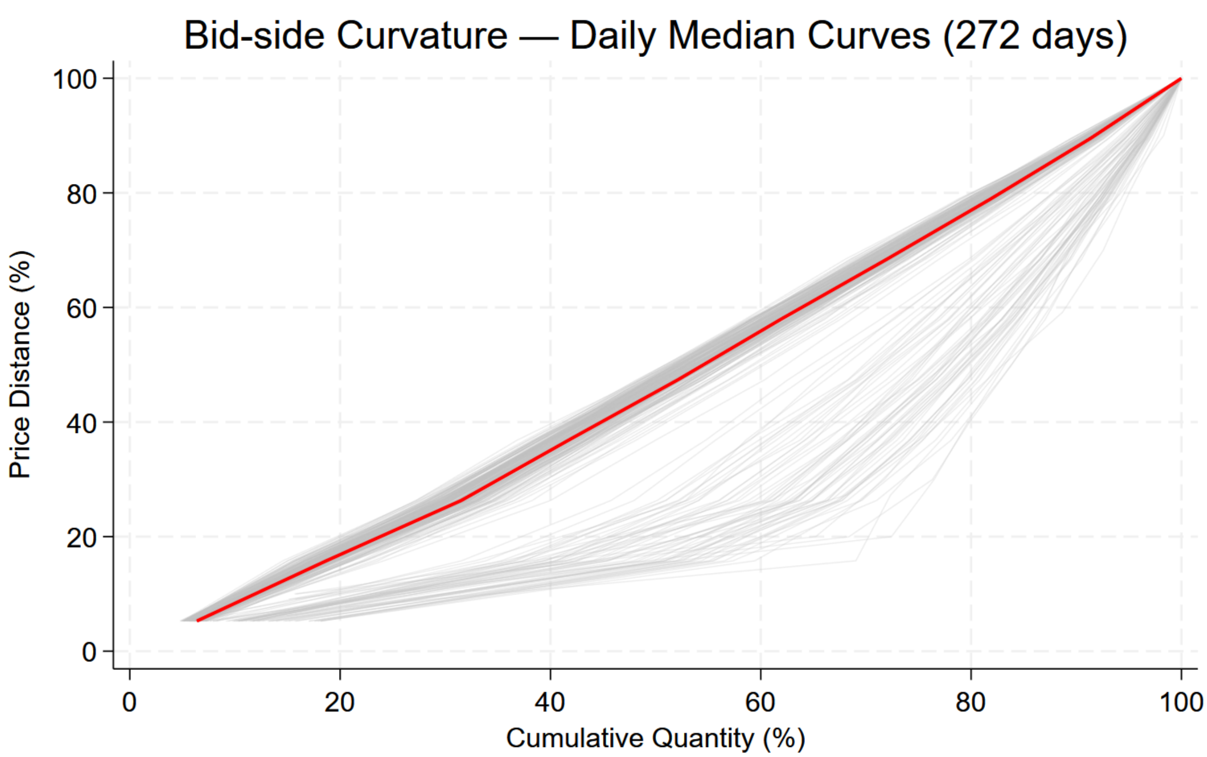

Figure 6.1 Bid-side Curvature of the Crude Oil LOB (272 days)
Note: Gray lines are daily median LOB curves, obtained by collapsing 15-minute snapshots within each day; the red line is the cross-sectional median across all 272 trading days. The horizontal axis denotes cumulative quantity share, and the vertical axis denotes price distance share.

The ask-side limit order book also exhibits a convex structure in most cases, consistent with the pattern observed on the bid side, suggesting that a large concentration of orders is clustered around the best ask and its neighboring quotes, thereby reflecting strong liquidity. According to the descriptive statistics of the curvature coefficient (β), the average β on the ask side is 1.120, slightly higher than the bid-side average of 1.106, with nearly identical median values across both sides. However, the distribution of β on the ask side demonstrates greater dispersion, indicating higher variability compared to the bid side. Given that crude oil is a highly liquid and volatile commodity, ask-side liquidity may be more responsive to market sentiment and risk expectations.

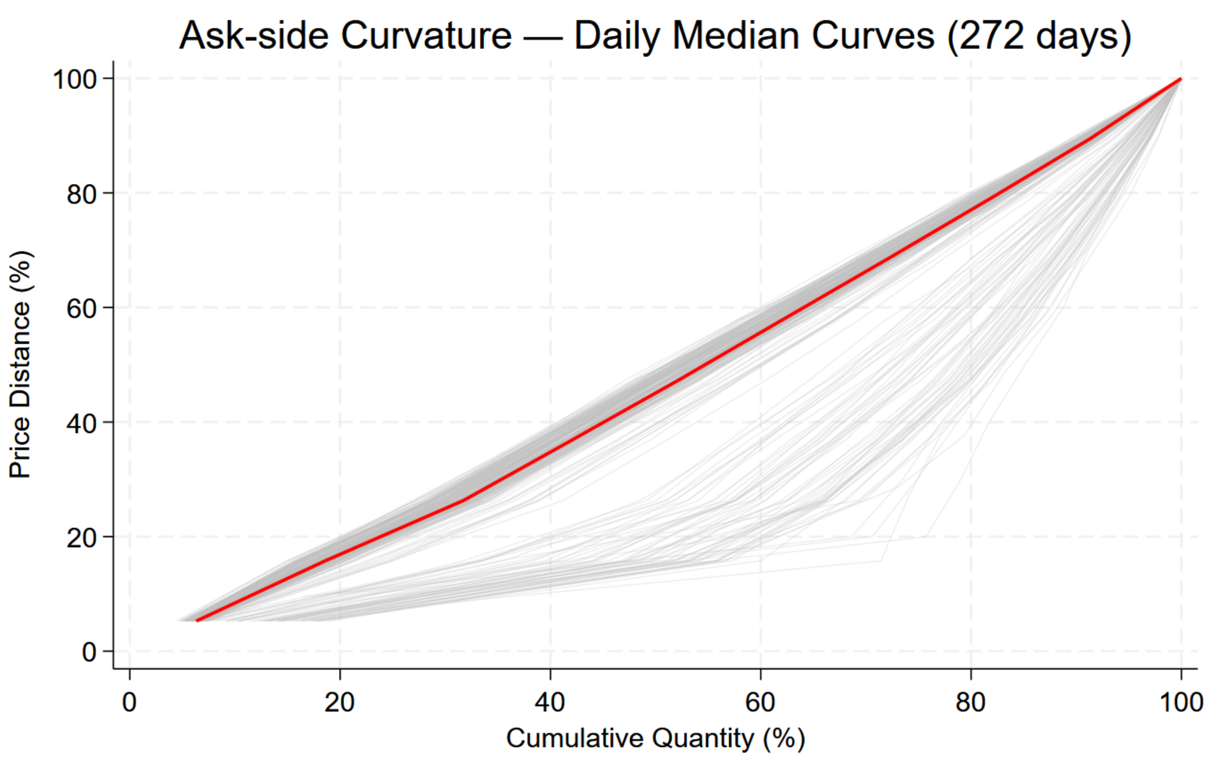

Figure 6.2 Ask-side Curvature of the Crude Oil LOB (272 days)
Note: Gray lines are daily median LOB curves, obtained by collapsing 15-minute snapshots within each day; the red line is the cross-sectional median across all 272 trading days. The horizontal axis denotes cumulative quantity share, and the vertical axis denotes price distance share.

When aggregating daily median curves across 272 trading days, the majority of curves exhibit a convex form. Specifically, on the bid side, 243 daily medians are convex and 29 are concave, while on the ask side, 239 daily medians are convex and 33 are concave. This indicates that compared with 15-minute snapshots, the daily median aggregation smooths intraday short-term fluctuations of the order book and more clearly reveals its dominant convex shape in most periods.

6.2 VAR and IRF Analysis

We estimate a VAR (4) for $RETURN_t$, β_t, $DEPTH_t$, $SPREAD_t$ with a one-second sampling interval and report orthogonalized IRFs using a Cholesky identification. We set p=4 in the model because the Final Prediction Error (FPE), Akaike Information Criterion (AIC), and Hannan-Quinn Information Criterion (HQIC) all attain their minima at four lags, and the Schwarz-Bayes Information Criterion (SBIC) does not overturn this choice (see Appendix Table A.2). Given the sampling frequency, one lag corresponds to one snapshot step, so p=4 captures feedback over four seconds. The endogenous variables include return ($RETURN_t$), average order book curvature (β_t), average market depth ($DEPTH_t$), and relative bid-ask spread ($SPREAD_t$). All endogenous variables exhibit strong persistence across lags (see Appendix Table A.5).

Figures 6.3-6.5 illustrate the dynamic responses of market depth, spread and return to shocks in order book curvature (β). DEPTH drops on impact by about -0.02, with the trough around 5-10 seconds; it then rises monotonically and recovers by roughly 180 seconds. The 95% confidence band excludes zero for the first 80-120 seconds. SPREAD widens immediately following a shock to curvature, peaks around 0.001-0.0015, has a half-life (50% recovery) under 10 seconds, and fully recovers within about 60-90 seconds. This effect is statistically significant primarily during the first few dozen seconds. RETURN reacts only minimally: a small oscillation is observed in the first seconds ($\approx\pm2\times10^{-6}$) followed by a quick reversion to zero. Mechanically, the implied matching functionality of the exchange derives implied orders at the top or near the top of the book. A wider spread is accompanied by a larger share of implied orders out of total depth in the LOB as participants re-post or pull outright orders from deeper tiers. This means that a very convex book if accompanied by wider spreads contains lower depths. The depths near the top of the book relatively thickens while the outer layers thin even more, and thus, the 10-step depth declines on impact (Figure 6.3). As this is a reallocation of supply rather than an information shock, returns show only a very short-lived blip around zero (Figure 6.5). In sum, higher curvature concentrates liquidity near the mid-quote, thins the outer layers of the book (lower average total depth), and triggers a brief widening of the spread, while the effect on returns vanishes within a few seconds.

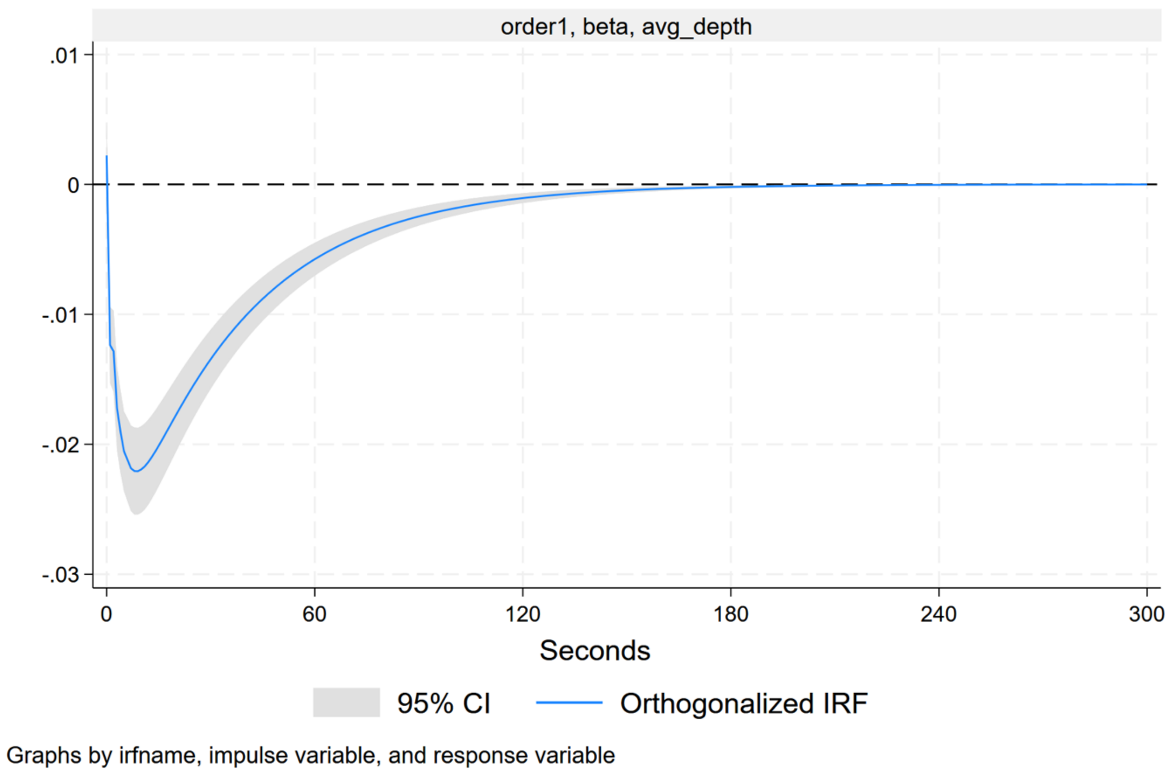

Figure 6.3 IRF: Response of Market Depth to an Order Book Curvature Shock

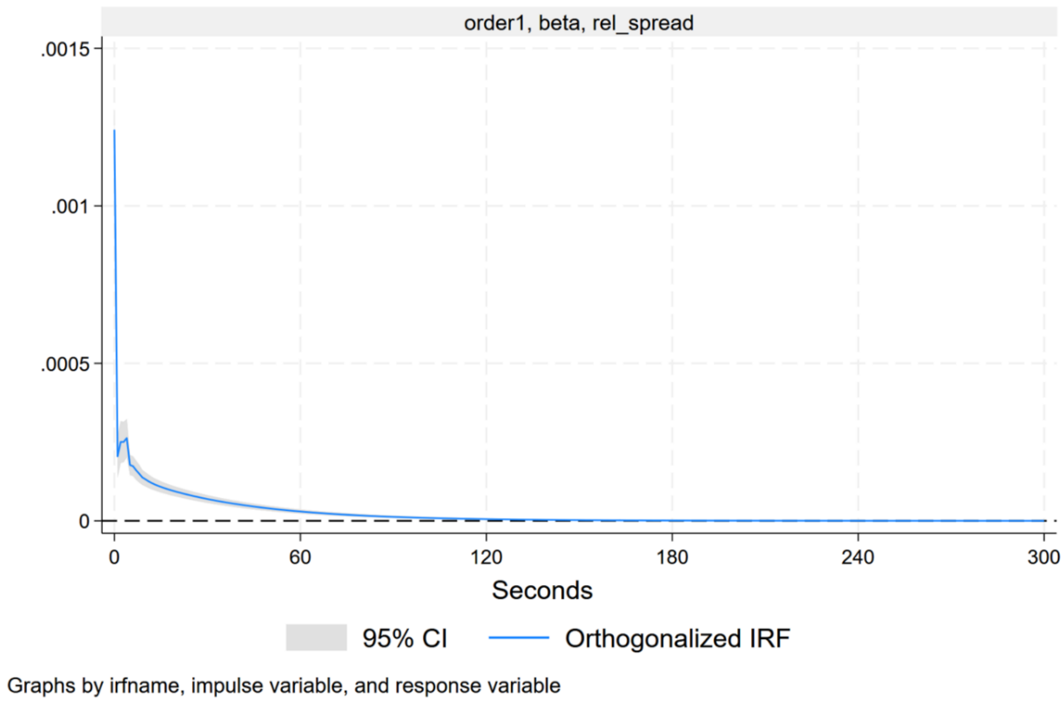

Figure 6.4 IRF: Response of Relative Bid-Ask Spread to an Order Book Curvature Shock

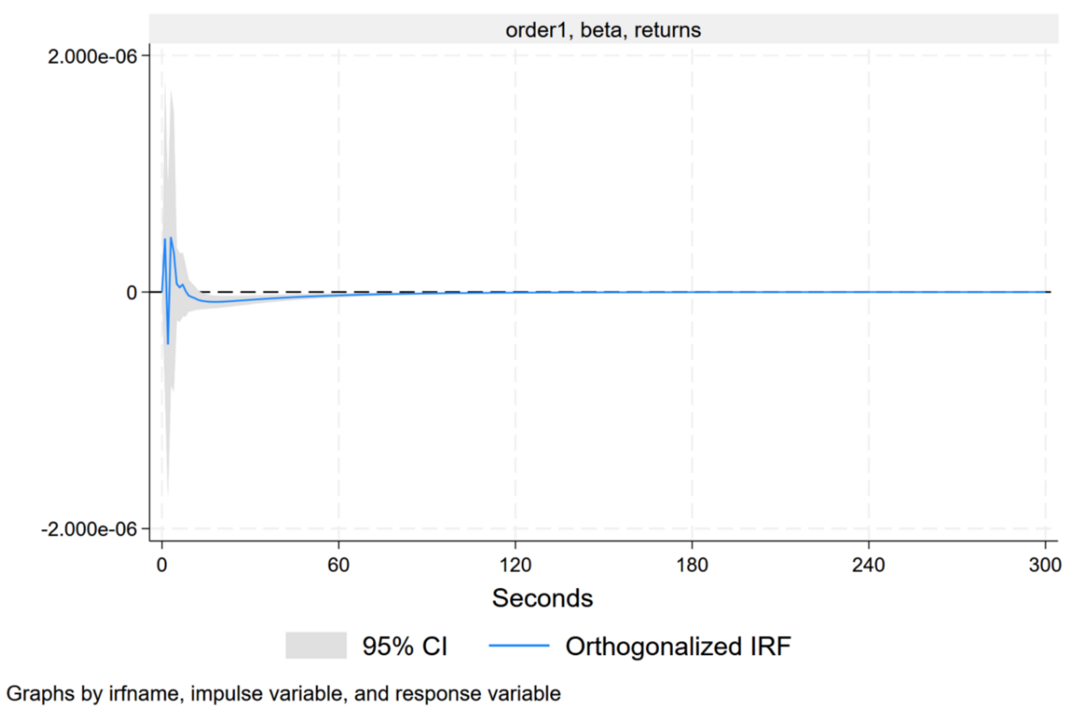

Figure 6.5 IRF: Response of Return to an Order Book Curvature Shock

Figures 6.6-6.8 present the impulse responses of order book curvature (β), spread and return to a positive shock in average market depth. β drops on impact by roughly -0.014 to -0.015 with the trough at around 5-10 seconds; it then rises monotonically and recovers by 120-150 seconds. The 95% band excludes zero for about the first 60-90 seconds, showing a sizable, short-run dilution of convexity as liquidity spreads across more price levels. SPREAD tightens immediately, about -0.0008 to -0.0005, and mean-reverts within 60-120 seconds, consistent with thicker books compressing trading costs near the top of the book. RETURN reacts only weakly, a small blip in the first seconds followed by a quick fade to zero, statistically indistinguishable from zero beyond the very short horizon. Mechanically, greater depth reflects more outright orders supplied across multiple price levels, so the top of book is no longer over-stacked while the outer tiers become better supplied, result in curvature declines. Overall, a depth shock reduces convexity and briefly narrows spreads, while its impact on returns is minimal.

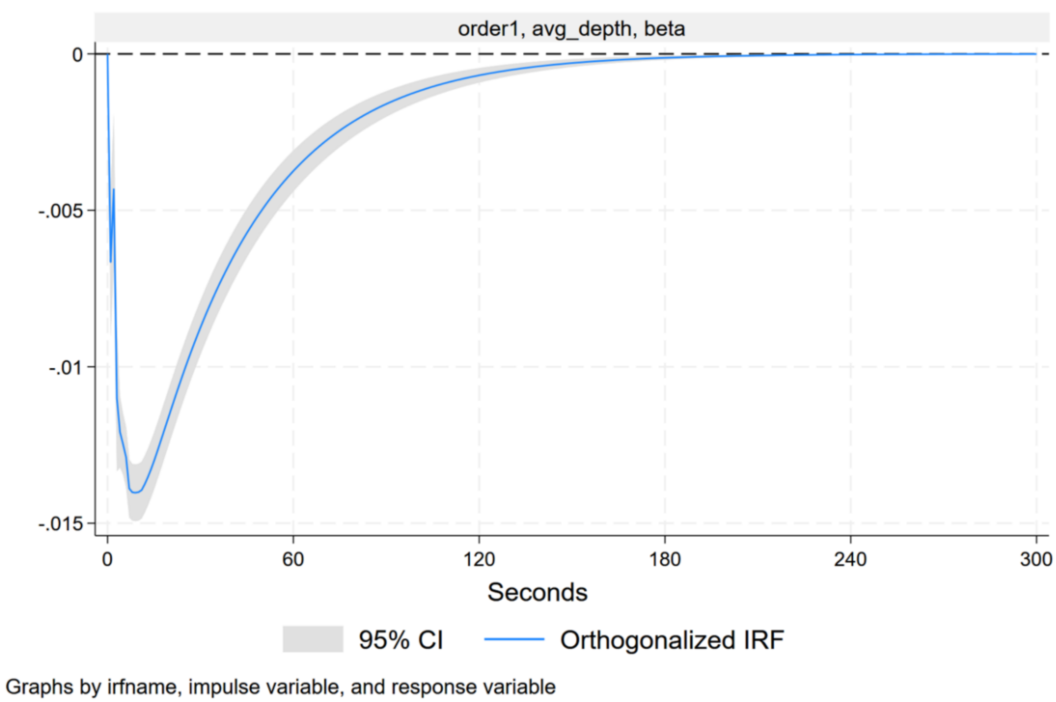

Figure 6.6 IRF: Response of Curvature to Market Depth Shock

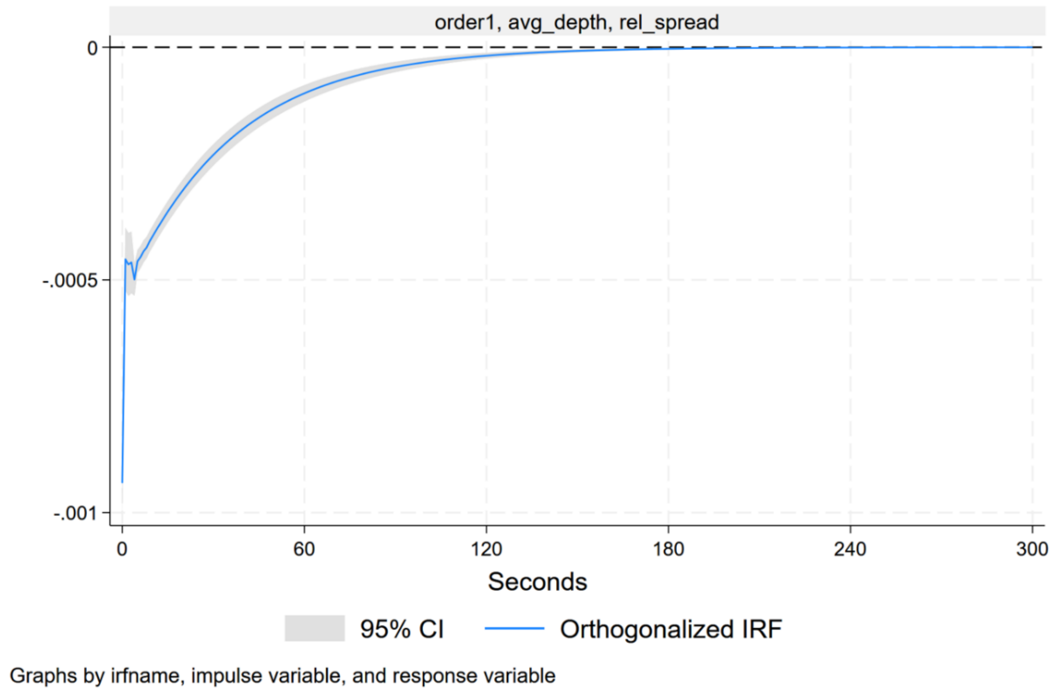

Figure 6.7 IRF: Response of Relative Bid-Ask Spread to Market Depth Shock

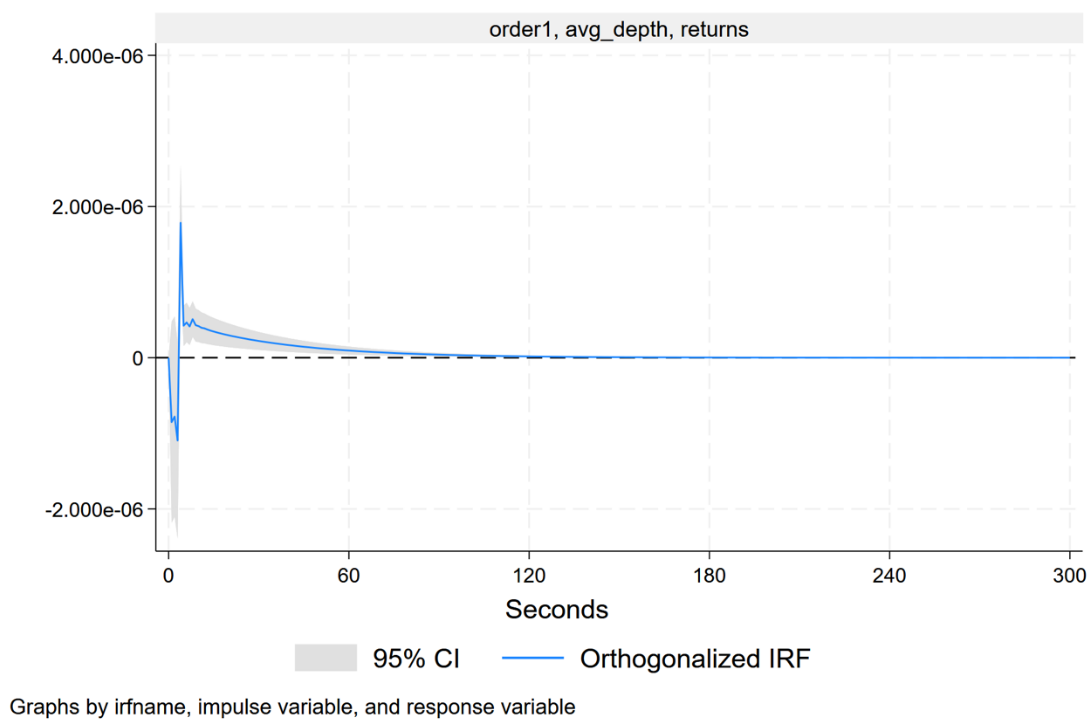

Figure 6.8 IRF: Response of Return to Market Depth Shock

Figure 6.9-6.11 illustrate the dynamic responses of order book curvature (β), market depth, and return to a shock in relative spread. β (curvature) turns up after a brief blip, peaking around 0.002 at 5-10 seconds, then decays smoothly to zero by 120 seconds; the 95% band excludes zero for roughly the first 30-60 seconds. DEPTH falls sharply on impact to about -0.010 to -0.012 with the trough at 5-10 seconds and then mean-reverts toward zero within 120-150 seconds; it remains significantly negative over about 60-90 seconds. RETURN shows only a small oscillation around $\pm 2\times10^{-6}$ in the first seconds and quickly reverts to zero. A wider SPREAD raises the payoff to spread trading and increases activity in spread markets. The matching engine generates more implied orders from spreads, which add resting volume at the best 1-3 price levels. This inflow thickens the book near the mid-quote and makes the book more convex. At the same time, outright liquidity recedes as traders step away from the top of the book or close positions to avoid trading against informed flow, so depth thins across 10 steps. The combination of more implied volume at the top and fewer outright orders through the stack explains the short-run rise in curvature together with the contraction in aggregate depth.

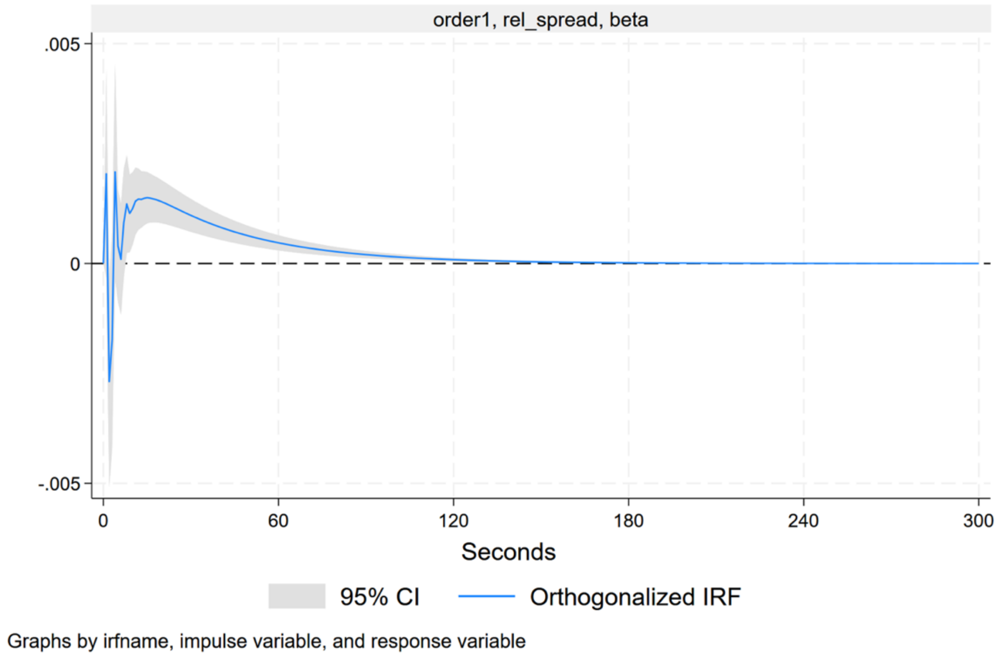

Figure 6.9 IRF: Response of Curvature to Relative Bid-Ask Spread Shock

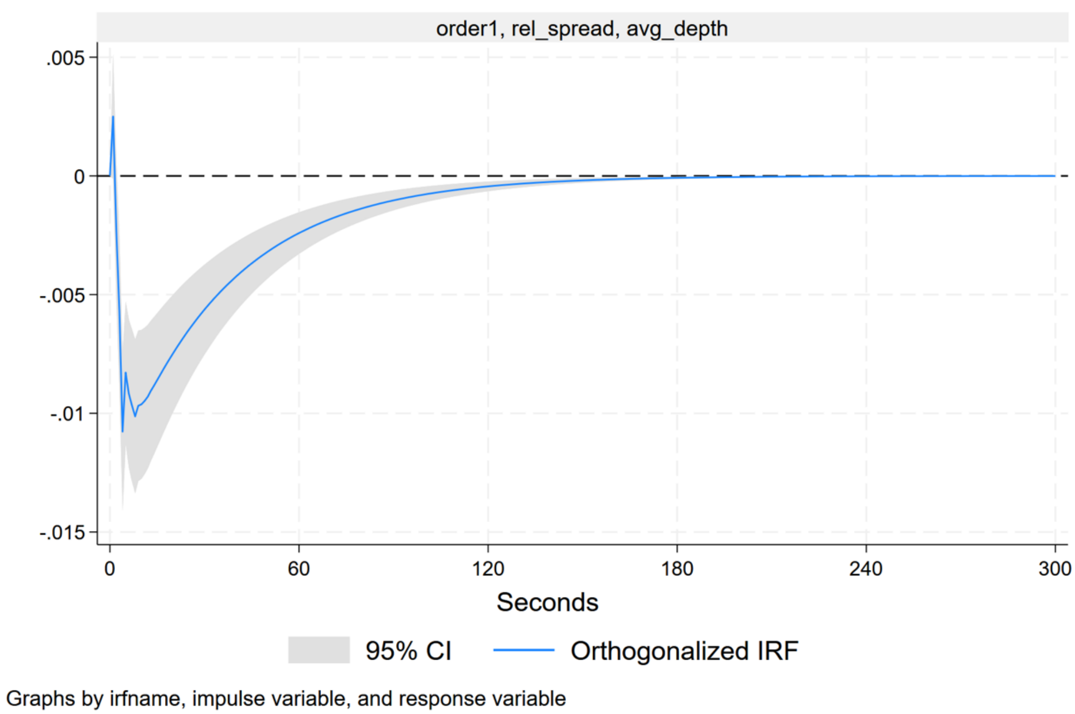

Figure 6.10 IRF: Response of Market Depth to Relative Bid-Ask Spread Shock

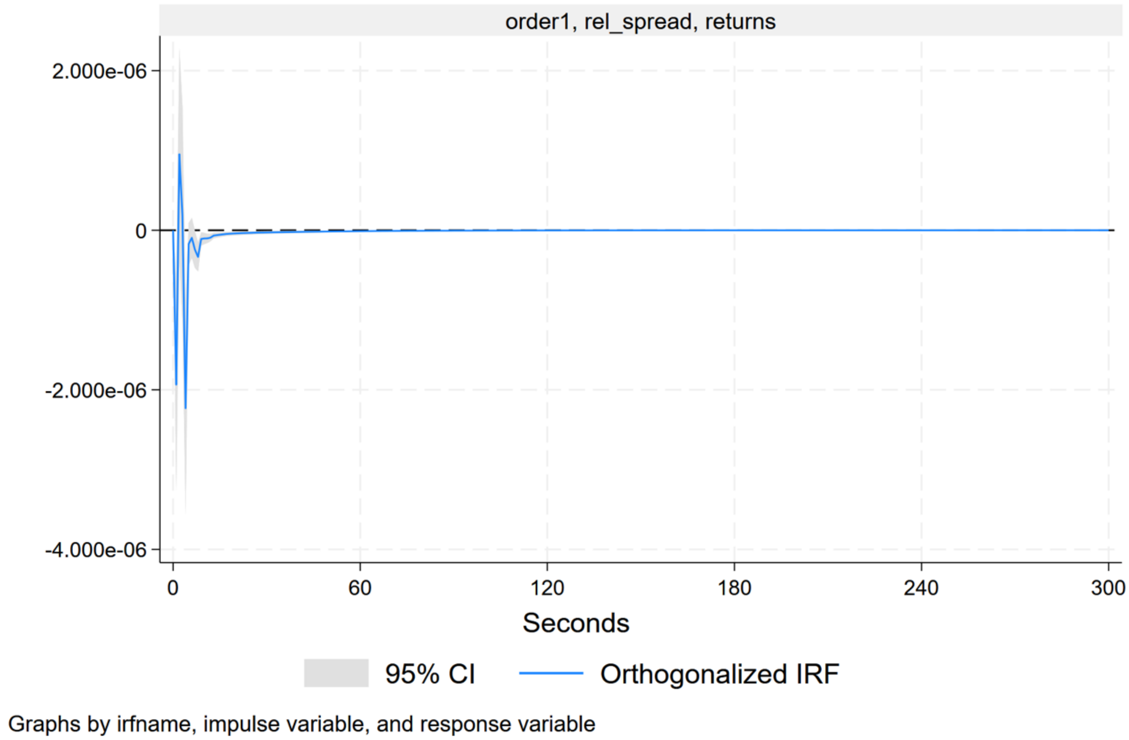

Figure 6.11 IRF: Response of Return to Relative Bid-Ask Spread Shock

Figure 6.12-6.14 illustrate the orthogonalized impulse responses of order book curvature (β), average market depth, and the relative bid-ask spread to a shock in return. β (curvature) dips by about -0.0015 to -0.002 in the first seconds and then mean-reverts to zero by 120 seconds, with significance limited to the very short horizon. DEPTH rises on impact, peaking around 0.005-0.007 at 5-10 seconds, and then decays to zero by 120-150 seconds. The response is statistically different from zero over roughly the first 60-90 seconds. SPREAD shows a brief fluctuation and then a small narrowing (on the order of 10^{-4}), which fades to zero within 60-120 seconds and the confidence band overlaps zero soon after the impact. Therefore, a transient price movement attracts opportunistic liquidity, participants' quotes fan out across multiple price levels, resulting in a short-lived decline in curvature and an increase in aggregate depth.

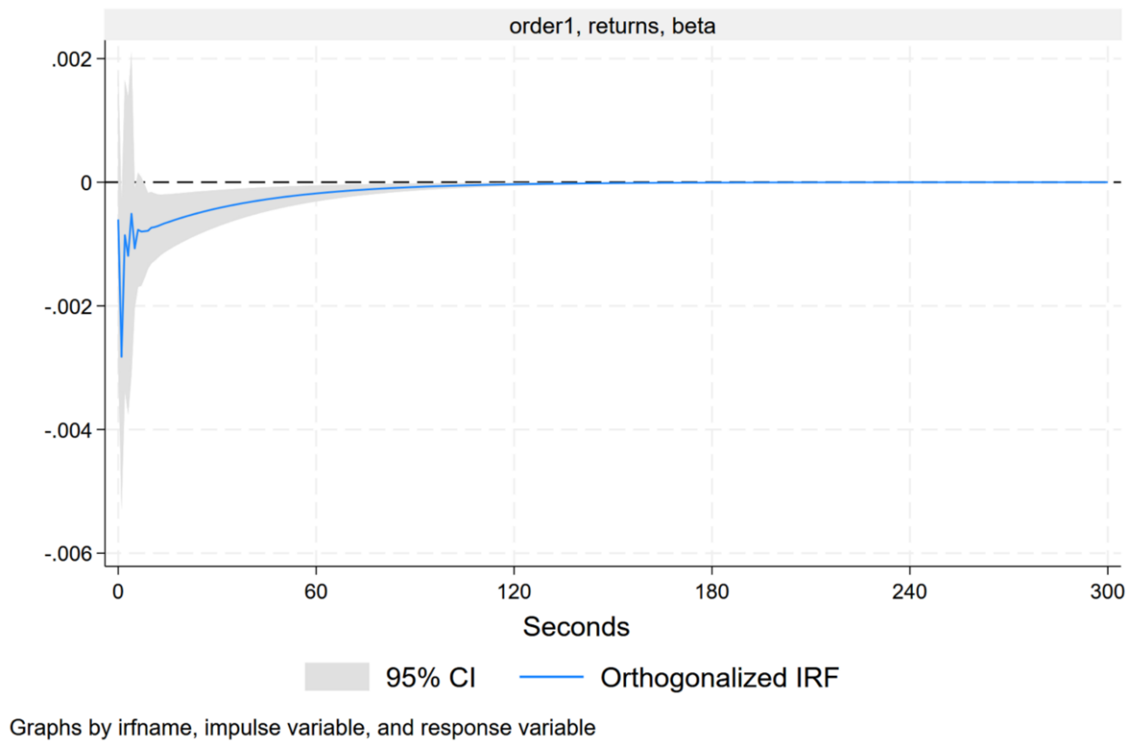

Figure 6.12 IRF: Response of Curvature to Return Shock

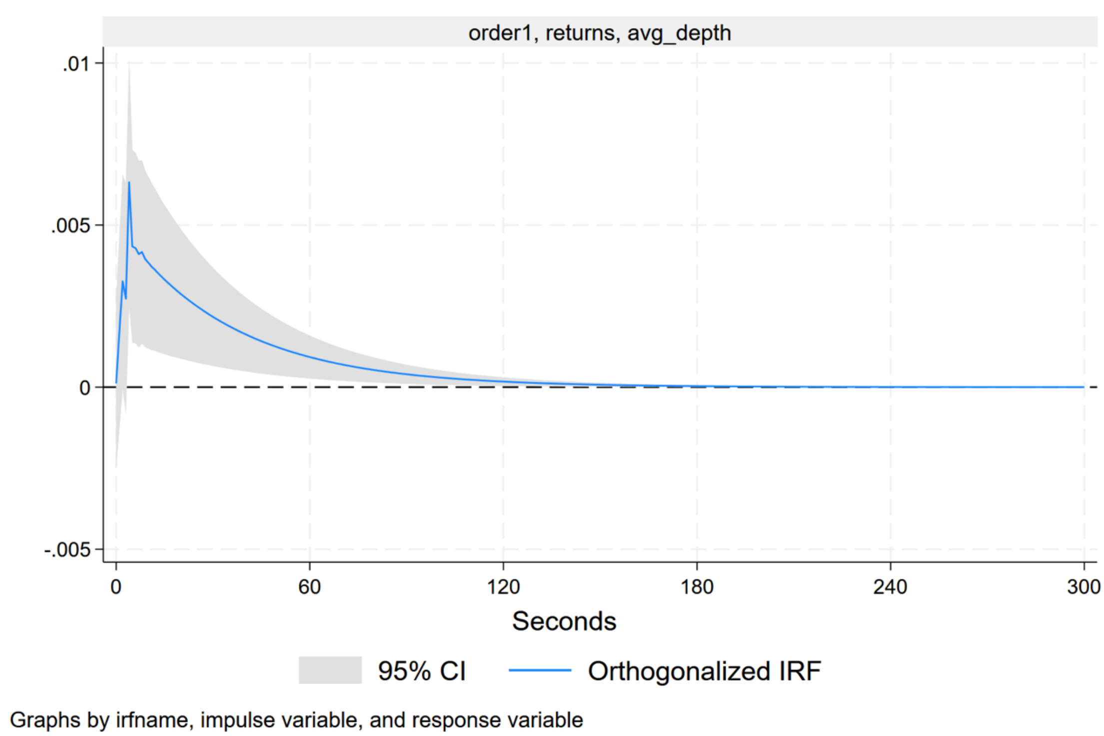

Figure 6.13 IRF: Response of Market Depth to Return Shock

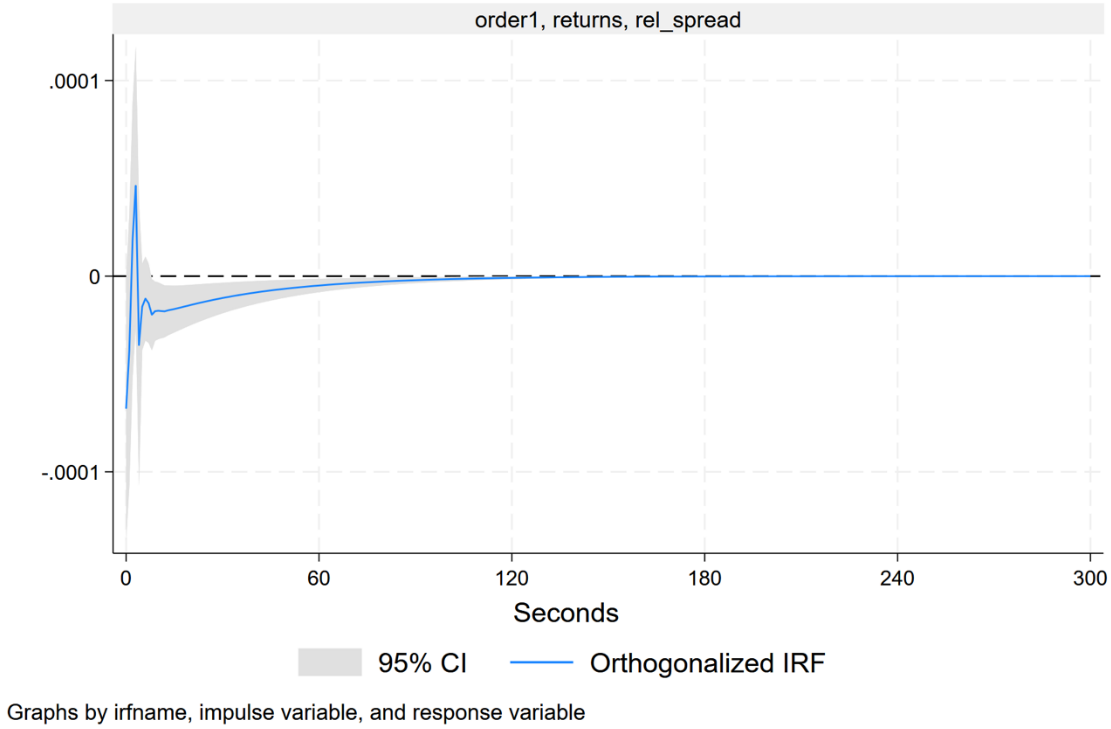

Figure 6.14 IRF: Response of Relative Bid-Ask Spread to Return Shock

The VAR and impulse-response analyses provide strong empirical support for Hypothesis 1, which posits that curvature variations are systematically negative with changes in market depth and spreads. Positive shocks to curvature, reflecting greater convexity in the limit order book are followed by reductions in overall depth and widening spreads. These adjustments indicate temporary liquidity withdrawal and heightened trading asymmetry under increased uncertainty. The effects are short-lived: across all shocks, responses are short-lived and recover from the shocks within 60-180 seconds, indicating that these fluctuations stem from endogenous reallocation of liquidity within the book rather than from exogenous information arrivals. Specifically, the transitory reshuffling between implied and outright orders across the top and outer tiers of the book underscores that curvature shocks represent internal liquidity restructuring. The rapid normalization of depth, spread, curvature, and returns thus provides clear evidence of the structural resiliency of the crude oil futures market.

Compared Lo and Hall (2015), who report that in Australian equities, best-quote variables recover within 1-3 seconds, while deeper depth in mid- and small-cap stocks takes around 20-35 seconds to adjust, our findings indicate that depth in the crude oil futures market requires about 180 seconds to recover, implying a slower adjustment than that observed for mid- and small-cap equities. Relative to He et al. (2021), who show that depth and spreads in corn and lean hog futures generally recover within 1-3 minutes but at accelerated speeds during large price moves (20-80 seconds), our results for crude oil futures are consistent with this range.

6.3 Incremental Explanatory Power

We next measure the two-scale realized volatility for each 15-minute interval. We then forecast 15-minute-ahead market volatility by regressing the next-interval TSRV ($TSRV_{t+1}$) on lagged volatility and liquidity control variables. To examine how external information shocks affect the structure of market liquidity, we include the EIA weekly reports, released each Wednesday at 10:30 a.m. ET. This report is widely regarded as one of the most influential scheduled announcements in the crude oil market and it provides timely and comprehensive information on U.S. petroleum inventories and production (Bu, 2014; Miao et al., 2018; Tian et al., 2019). To capture the potential impact of these scheduled announcements on crude oil volatility, we generate a Wednesday dummy variable. The variable takes the value of 1 if the trading day is Wednesday and 0 otherwise. In the sample, Wed = 1 corresponds to 5,027 observations, while Wed = 0 covers 19,920 observations. This specification enables a clear identification of market differences between announcement days and non-announcement days. To account for intraday seasonality, we also include slot dummy variables for intraday intervals.

The slot coefficients in Table 6.4 reveal clear intraday volatility patterns in crude oil futures. The signs of the coefficients indicate whether volatility in each slot is higher or lower relative to the baseline period (slot0, 17:10-17:40 CT). Evening to night hours (slots 1-12, 17:40-23:00 CT) show the lowest volatility, with negative and significant coefficients (e.g., 1.slot30 = -0.0177, 11.slot30 = -0.0250). Volatility then gradually increases and peaks during slots 27-33 (06:40-10:10 CT), with slot29 (07:40-08:10 CT) reaching 0.0937 and slot32 (09:10-09:40 CT) at 0.0861, the highest values of the trading day. Midday trading (slots 34-38, 10:10-12:40 CT) exhibits moderate volatility levels, with coefficients between 0.0321 and 0.0731. Before market close (slots 39-41, 12:40-14:10 CT), volatility rises again to a secondary peak around 0.05 and finally drops sharply in the last 10 minutes of trading (slot45, -0.0878). This evidence confirms a U-shaped intraday volatility pattern: the market is relatively calm in the evening, reaches its highest activity in the early U.S. daytime, stabilizes around midday, experiences a secondary rise before settlement, and becomes quiet in the final minutes.

In Table 6.4 Model I, the lagged volatility appears with a coefficient of 0.5145, showing strong persistence in crude oil futures volatility. TSRV alone explains a moderate portion of the variation in future volatility, with an adjusted R^2 of 57.92%. In Model II, order book curvature is added. The coefficients are negative and significant for both sides (β^{bid} = -0.0032*, β^{ask} = -0.0057***). A higher curvature, indicating a steeper concentration of liquidity near the mid-quote, is associated with lower subsequent volatility. This suggests that a thicker book around the best quotes absorbs incoming orders more effectively, reducing price adjustments. Adding curvature (β) in Model II significantly improves the fit to 57.97%, suggesting that the shape of the order book provides additional predictive content beyond volatility persistence. Model III incorporates market average depth. The coefficient is -0.0300, confirming that deeper books stabilize prices by buffering order flow shocks. When average depth is introduced in Model III, explanatory power increases further to 59.77%, indicating that liquidity provision at multiple price levels carries predictive content. Model IV introduces relative spread and its coefficient is positive and highly significant. Model V includes the slope of the book, and its coefficient is small in magnitude but statistically significant. This result is consistent with (Hasbrouck & Seppi, 2001; Jain & Jiang, 2014; Valenzuela et al., 2015), who argue that steeper books lower the price impact of trades, yielding a modest improvement of explanatory power to 60.21%. The forecasting regressions based on two-scale realized volatility (TSRV) lend empirical support to Hypothesis 2. This result underscores the negative relationship between curvature and realized volatility, as well as curvature's incremental explanatory power, highlighting its role as a non-linear geometric measure that captures the market liquidity.

The regression results in Model VI show that the Wednesday dummy, used to capture the scheduled release of the EIA petroleum status report at 09:30 CT, is statistically insignificant. This finding indicates that the occurrence of the report itself does not increase short-term volatility in crude oil futures. In other words, volatility patterns on Wednesdays are not systematically different from other trading days once liquidity controls and time-of-day effects are accounted for. This result is consistent with the evidence from (Kilian & Vega, 2011), who emphasize that asset price reactions are primarily driven by the surprise component of macroeconomic and energy announcements rather than by their scheduled occurrence. Similarly, studies such as (Foucault et al., 2007) highlight that liquidity and order book structure mediate how markets absorb information shocks. Our results therefore reinforce that scheduled announcements only matter when they convey unexpected information, with absent a surprise, the announcement day itself does not lead to higher volatility. Overall, the results concerning the EIA reports provide support for Hypothesis 3. After controlling for intraday seasonality, the volatility dynamics in crude oil futures are predominantly driven by endogenous liquidity reallocation within the order book rather than by exogenous news effects, reinforcing the self-correcting nature of liquidity provision in high-frequency markets.

Table 6.4 Predictability Regression

$TSRV_{t+1}$	**I**	**II**	**III**	**IV**	**V**	**VI**
$TSRV_t$	0.5145***	0.5127***	0.4550***	0.4488***	0.4368***	0.4366***
	(0.0176)	(0.0174)	(0.0161)	(0.0159)	(0.0160)	(0.0159)
Slot30						
1.[17:40-18:10)	-0.0177***	-0.0175***	-0.0122***	-0.0111***	-0.0114***	-0.0114***
	(0.0027)	(0.0027)	(0.0027)	(0.0027)	(0.0027)	(0.0027)
2.[18:10-18:40)	-0.0204***	-0.0201***	-0.0107***	-0.0096***	-0.0104***	-0.0104***
	(0.0030)	(0.0030)	(0.0030)	(0.0031)	(0.0031)	(0.0031)
3.[18:40-19:10)	0.0117***	0.0121***	0.0236***	0.0244***	0.0235***	0.0235***
	(0.0025)	(0.0025)	(0.0025)	(0.0026)	(0.0026)	(0.0026)
4.[19:10-19:40)	-0.0113***	-0.0111***	0.0109***	0.0112***	0.0101***	0.0101***
	(0.0032)	(0.0032)	(0.0034)	(0.0035)	(0.0035)	(0.0035)
5.[19:40-20:10)	0.0028	0.0031	0.0268***	0.0269***	0.0256***	0.0257***
	(0.0030)	(0.0030)	(0.0032)	(0.0032)	(0.0032)	(0.0032)
6.[20:10-20:40)	-0.0063**	-0.0061**	0.0191***	0.0196***	0.0185***	0.0186***
	(0.0028)	(0.0028)	(0.0030)	(0.0030)	(0.0030)	(0.0030)
7.[20:40-21:10)	-0.0101***	-0.0099***	0.0159***	0.0160***	0.0150***	0.0150***
	(0.0028)	(0.0028)	(0.0031)	(0.0032)	(0.0032)	(0.0032)
8.[21:10-21:40)	-0.0174***	-0.0171***	0.0079**	0.0076**	0.0063**	0.0064***
	(0.0028)	(0.0028)	(0.0032)	(0.0032)	(0.0032)	(0.0032)
9.[21:40-22:10)	-0.0174***	-0.0173***	0.0080**	0.0080**	0.0067**	0.0067**
	(0.0027)	(0.0027)	(0.0031)	(0.0031)	(0.0031)	(0.0031)
10.[22:10-22:40)	-0.0239***	-0.0238***	0.0005	0.0004	-0.0010	-0.0010
	(0.0027)	(0.0027)	(0.0031)	(0.0031)	(0.0031)	(0.0031)
11.[22:40-23:10)	-0.0250***	-0.0247***	0.0001	0.0001	-0.0012	-0.0011
	(0.0026)	(0.0026)	(0.0030)	(0.0030)	(0.0030)	(0.0030)
12.[23:10-23:40)	-0.0186***	-0.0184***	0.0065**	0.0067**	0.0053*	0.0054*
	(0.0028)	(0.0028)	(0.0032)	(0.0032)	(0.0032)	(0.0032)
13.[23:40-00:10)	-0.0224***	-0.0223***	0.0030	0.0033	0.0022	0.0023
	(0.0032)	(0.0031)	(0.0036)	(0.0036)	(0.0036)	(0.0036)
14.[00:10-00:40)	-0.0087***	-0.0083***	0.0177***	0.0181***	0.0169***	0.0170***
	(0.0029)	(0.0029)	(0.0033)	(0.0034)	(0.0034)	(0.0034)
15.[00:40-01:10)	-0.0006	-0.0006	0.0241***	0.0241***	0.0227***	0.0228***
	(0.0031)	(0.0031)	(0.0034)	(0.0035)	(0.0035)	(0.0035)
16.[01:10-01:40)	-0.0014	-0.0014	0.0278***	0.0276***	0.0265***	0.0266***
	(0.0034)	(0.0034)	(0.0041)	(0.0041)	(0.0041)	(0.0041)
17.[01:40-02:10)	0.0246***	0.0245***	0.0550***	0.0543***	0.0534***	0.0535***
	(0.0033)	(0.0033)	(0.0039)	(0.0039)	(0.0039)	(0.0039)
18.[02:10-02:40)	0.0138***	0.0136***	0.0458***	0.0455***	0.0451***	0.0452***
	(0.0035)	(0.0034)	(0.0041)	(0.0041)	(0.0041)	(0.0041)
19.[02:40-03:10)	0.0310***	0.0310***	0.0642***	0.0634***	0.0629***	0.0630***
	(0.0040)	(0.0039)	(0.0046)	(0.0045)	(0.0045)	(0.0045)
20.[03:10-03:40)	0.0284***	0.0282***	0.0627***	0.0620***	0.0619***	0.0620***
	(0.0040)	(0.0039)	(0.0044)	(0.0043)	(0.0044)	(0.0044)
21.[03:40-04:10)	0.0295***	0.0293***	0.0631***	0.0626***	0.0626***	0.0627***
	(0.0041)	(0.0040)	(0.0047)	(0.0047)	(0.0048)	(0.0048)
22.[04:10-04:40)	0.0208***	0.0207***	0.0555***	0.0546***	0.0544***	0.0545***
	(0.0041)	(0.0041)	(0.0047)	(0.0046)	(0.0047)	(0.0047)
23.[04:40-05:10)	0.0205***	0.0205***	0.0545***	0.0539***	0.0536***	0.0537***

	(0.0036)	(0.0035)	(0.0041)	(0.0041)	(0.0042)	(0.0041)
24.[05:10-05:40)	0.0204***	0.0204***	0.0550***	0.0542***	0.0539***	0.0540***
	(0.0038)	(0.0038)	(0.0045)	(0.0044)	(0.0045)	(0.0045)
25.[05:40-06:10)	0.0221***	0.0220***	0.0563***	0.0556***	0.0553***	0.0554***
	(0.0039)	(0.0039)	(0.0045)	(0.0045)	(0.0046)	(0.0046)
26.[06:10-06:40)	0.0176***	0.0174***	0.0521***	0.0513***	0.0512***	0.0513***
	(0.0037)	(0.0037)	(0.0044)	(0.0043)	(0.0043)	(0.0043)
27.[06:40-07:10)	0.0382***	0.0381***	0.0728***	0.0723***	0.0721***	0.0722***
	(0.0041)	(0.0040)	(0.0051)	(0.0051)	(0.0051)	(0.0051)
28.[07:10-07:40)	0.0352***	0.0351***	0.0696***	0.0690***	0.0696***	0.0697***
	(0.0041)	(0.0041)	(0.0049)	(0.0048)	(0.0048)	(0.0048)
29.[07:40-08:10)	0.0937***	0.0938***	0.1331***	0.1328***	0.1337***	0.1338***
	(0.0049)	(0.0048)	(0.0054)	(0.0053)	(0.0054)	(0.0054)
30.[08:10-08:40)	0.0703***	0.0706***	0.1167***	0.1162***	0.1180***	0.1182***
	(0.0053)	(0.0053)	(0.0063)	(0.0062)	(0.0063)	(0.0063)
31.[08:40-09:10)	0.0770***	0.0768***	0.1194***	0.1191***	0.1201***	0.1202***
	(0.0064)	(0.0064)	(0.0068)	(0.0068)	(0.0069)	(0.0068)
32.[09:10-09:40)	0.0861***	0.0860***	0.1282***	0.1260***	0.1286***	0.1287***
	(0.0060)	(0.0059)	(0.0060)	(0.0058)	(0.0060)	(0.0060)
33.[09:40-10:10)	0.0764***	0.0763***	0.1223***	0.1208***	0.1227***	0.1228***
	(0.0056)	(0.0056)	(0.0064)	(0.0063)	(0.0064)	(0.0064)
34.[10:10-10:40)	0.0731***	0.0731***	0.1193***	0.1185***	0.1198***	0.1200***
	(0.0062)	(0.0061)	(0.0067)	(0.0066)	(0.0067)	(0.0067)
35.[10:40-11:10)	0.0347***	0.0348***	0.0789***	0.0784***	0.0793***	0.0795***
	(0.0051)	(0.0050)	(0.0058)	(0.0057)	(0.0057)	(0.0057)
36.[11:10-11:40)	0.0321***	0.0322***	0.0735***	0.0726***	0.0733***	0.0734***
	(0.0040)	(0.0040)	(0.0049)	(0.0049)	(0.0049)	(0.0049)
37.[11:40-12:10)	0.0402***	0.0404***	0.0804***	0.0798***	0.0803***	0.0804***
	(0.0043)	(0.0043)	(0.0052)	(0.0051)	(0.0052)	(0.0051)
38.[12:10-12:40)	0.0404***	0.0407***	0.0809***	0.0804***	0.0809***	0.0810***
	(0.0046)	(0.0046)	(0.0057)	(0.0056)	(0.0057)	(0.0056)
39.[12:40-13:10)	0.0500***	0.0501***	0.0883***	0.0879***	0.0883***	0.0884***
	(0.0043)	(0.0043)	(0.0048)	(0.0048)	(0.0049)	(0.0048)
40.[13:10-13:40)	0.0497***	0.0499***	0.0966***	0.0962***	0.0978***	0.0979***
	(0.0056)	(0.0055)	(0.0064)	(0.0063)	(0.0064)	(0.0064)
41.[13:40-14:10)	-0.0057	-0.0055	0.0303***	0.0296***	0.0294***	0.0295***
	(0.0037)	(0.0037)	(0.0044)	(0.0044)	(0.0044)	(0.0044)
42.[14:10-14:40)	-0.0022	-0.0023	0.0293***	0.0286***	0.0279***	0.0280***
	(0.0037)	(0.0037)	(0.0045)	(0.0045)	(0.0045)	(0.0045)
43.[14:40-15:10)	-0.0133***	-0.0132***	0.0128***	0.0128***	0.0121***	0.0122***
	(0.0033)	(0.0032)	(0.0037)	(0.0037)	(0.0037)	(0.0037)
44.[15:10-15:40)	-0.0061*	-0.0065*	0.0040	0.0023	0.0024	0.0024
	(0.0035)	(0.0034)	(0.0032)	(0.0032)	(0.0033)	(0.0033)
45.[15:40-15:50)	-0.0878***	-0.0872***	-0.0694***	-0.0687***	-0.0683***	-0.0683***
	(0.0042)	(0.0041)	(0.0040)	(0.0041)	(0.0041)	(0.0041)
β_t^{bid}		-0.0032*	-0.0185***	-0.0219***	-0.0170***	-0.0170***
		(0.0018)	(0.0021)	(0.0022)	(0.0022)	(0.0022)
β_t^{ask}		-0.0057***	-0.0213***	-0.0244***	-0.0197***	-0.0197***
		(0.0015)	(0.0020)	(0.0020)	(0.0022)	(0.0022)
$DEPTH_t$			-0.0300***	-0.0255***	-0.0193***	-0.0194***

			(0.0026)	(0.0026)	(0.0029)	(0.0029)
$SPREAD_t$				1.0307***	0.5586***	0.5567***
				(0.1315)	(0.1634)	(0.1631)
$SLOPE_t$					-9.33e-06***	-9.27e-06***
					(2.03e-06)	(2.04e-06)
WED_t						-0.0020
						(0.0024)
adj. R^2(%)	57.92	57.97	59.77	60.04	60.21	60.22

Note: The first and last 10 minutes of trading are excluded from the sample. Variable 0. slot30 (17:10-17:40 CT) is used as the baseline with coefficients normalized to zero and therefore is omitted from the table.

6.4 Robustness Tests

We diagnose multicollinearity using variance-inflation factors (VIFs) and pairwise Pearson correlation coefficients for all regressors. The core regressors have VIFs between 1.23 and 2.01 ($TSRV_t$=2.01, β^{ask}=1.26, β^{bid}=1.23, $DEPTH_t$=2.00, $SPREAD_t$=1.44); the 45 intraday Slot dummies have VIFs around 1.50-2.43; the mean VIF is 2.06, and all tolerances (1/VIF) are > 0.4. These levels are well below conventional cutoffs (5 or 10), indicating no material multicollinearity (see Appendix Table A.3). From Pearson correlations results, curvature shows generally weak correlations with volatility, depth and spread (e.g., β^{ask} vs. TSRV ρ=-0.07, β^{bid} vs. TSRV ρ=−0.06, β^{ask} vs. depth ρ=−0.29, β^{ask} vs. spread ρ=0.26). Overall, the correlation structure is moderate and does not pose a collinearity concern (see Appendix Table A.4).

Table 6.5 reports the baseline regression results where the one-step-ahead volatility proxy, $TSRV_{t+1}$, is regressed on order book curvature (β), relative spread, depth, and control variables. The coefficients show that curvature is negatively and significantly associated with future volatility from both bid- and ask-side, suggesting that when liquidity is more concentrated around the mid-quote, prices become less sensitive to order flow imbalances, dampening short-term volatility. At the same time, the coefficient on depth is negative and highly significant, indicating that a thicker order book enhances market capacity and provides a buffer effect that reduces realized volatility. Relative spread is positively related to $TSRV_{t+1}$, reflecting that wider spreads signal higher trading costs and adverse selection risks, which in turn contribute to higher realized volatility.

To evaluate the robustness of these findings, this thesis conducts 3 additional models in this part. First, replacing TSRV with 1-second sample realized volatility (RV_{fast}) yields qualitatively similar results in column 2, confirming that the explanatory power of curvature is not specific to the TSRV measure. Second, using next-day $TSRV_{t+1}$ as the dependent variable shows that the predictive role of curvature weakens over longer horizons, consistent with its short-term nature in column 3. Finally, when a Wednesday news dummy (WED_t) is added to the regression, its effect is statistically insignificant, which aligns with the results in section 6.3, suggesting that public news announcements do not directly alter realized volatility in column 4.

Table 6.5 $TSRV_{t+1}$ Model, RV Robustness Model, Next-day $TSRV_{t+1}$ Robustness Model and WED_t Model

Dependent Variable:	**(1) $TSRVt_{+1}$**	**(2) RV_{t+1}**	**(3) Next-day $TSRV_{t+1}$**	**(4) WED_t**
$TSRV_t$	0.4488***	0.0001***	0.2952***	0.4486***
	(0.0121)	(2.04e-06)	(0.0121)	(0.0159)
β_t^{ask}	-0.0244***	-2.16e-06***	-0.0157***	-0.0219***
	(0.0017)	(2.09e-07)	(0.0013)	(0.0022)
β_t^{bid}	-0.0219***	-2.01e-06***	-0.0181***	-0.0244***
	(0.0017)	(2.02e-07)	(0.0014)	(0.0020)
$DEPTH_t$	-0.0256***	-1.97e-06***	-0.0222***	-0.0256***
	(0.0013)	(1.51e-07)	(0.0011)	(0.0026)
$SPREAD_t$	1.0307***	0.0001***	0.5029***	1.0369***
	(0.1165)	(0.0001)	(0.0958)	(0.1305)
Slot30				
1.[17:40-18:10)	-0.0111***	-9.29e-07***	0.0192***	-0.011***
	(0.0034)	(3.11e-07)	(0.0031)	(0.0027)
2.[18:10-18:40)	-0.0096***	-8.17e-07***	0.0243***	-0.0096**
	(0.0036)	(3.04e-07)	(0.0028)	(0.0031)
3.[18:40-19:10)	0.0244***	7.84e-07**	0.0272***	0.0244***
	(0.0038)	(3.24e-07)	(0.0030)	(0.0026)
4.[19:10-19:40)	0.0112**	-5.22e-07	0.0622***	0.0113***
	(0.0045)	(4.92e-07)	(0.0031)	(0.0035)
5.[19:40-20:10)	0.0270***	6.64e-07**	0.0515***	0.0270***
	(0.0036)	(3.32e-07)	(0.0045)	(0.0032)
6.[20:10-20:40)	0.0196***	-1.09e-07	0.0599***	0.0197***
	(0.0036)	(3.41e-07)	(0.0030)	(0.0030)
7.[20:40-21:10)	0.0161***	-4.04e-08	0.0551***	0.0161***
	(0.0035)	(3.66e-07)	(0.0030)	(0.0032)
8.[21:10-21:40)	0.0076**	-8.04e-08	0.0444***	0.0076**
	(0.0035)	(3.22e-07)	(0.0027)	(0.0032)
9.[21:40-22:10)	0.0080**	1.37e-07	0.0420***	0.0081***
	(0.0035)	(3.26e-07)	(0.0029)	(0.0031)
10.[22:10-22:40)	0.0004	1.18e-08	0.0358***	0.0005
	(0.0033)	(3.08e-07)	(0.0027)	(0.0030)
11.[22:40-23:10)	0.0001	4.57e-07	0.0275***	0.0002
	(0.0032)	(3.02e-07)	(0.0025)	(0.0030)
12.[23:10-23:40)	0.0067**	7.45e-07**	0.0265***	0.0068**
	(0.0033)	(3.11e-07)	(0.0025)	(0.0032)

13.[23:40-00:10)	0.0034	5.15e-07	0.0352***	0.0034
	(0.0036)	(4.27e-07)	(0.0027)	(0.0036)
14.[00:10-00:40)	0.0182***	5.59e-07*	0.0402***	0.0182***
	(0.0034)	(3.17e-07)	(0.0026)	(0.0034)
15.[00:40-01:10)	0.0241***	2.74e-07	0.0486***	0.0242***
	(0.0035)	(3.24e-07)	(0.0028)	(0.0035)
16.[01:10-01:40)	0.0276***	4.20e-07	0.0611***	0.0277***
	(0.0038)	(3.49e-07)	(0.0030)	(0.0041)
17.[01:40-02:10)	0.0543***	1.76e-06***	0.0663***	0.0544***
	(0.0038)	(3.56e-07)	(0.0034)	(0.0039)
18.[02:10-02:40)	0.0456***	1.12e-06***	0.0917***	0.0457***
	(0.0038)	(4.03e-07)	(0.0034)	(0.0041)
19.[02:40-03:10)	0.0634***	1.85e-06***	0.0871***	0.0635***
	(0.0043)	(4.31e-07)	(0.0032)	(0.0045)
20.[03:10-03:40)	0.0620***	1.36e-06***	0.1074***	0.0621***
	(0.0044)	(4.51e-07)	(0.0037)	(0.0043)
21.[03:40-04:10)	0.0626***	1.24e-06**	0.1051***	0.0627***
	(0.0045)	(4.89e-07)	(0.0038)	(0.0047)
22.[04:10-04:40)	0.0546***	8.41e-07*	0.1076***	0.0547***
	(0.0045)	(4.58e-07)	(0.0038)	(0.0046)
23.[04:40-05:10)	0.0539***	7.13e-07*	0.0960***	0.0540***
	(0.0040)	(4.09e-07)	(0.0037)	(0.0041)
24.[05:10-05:40)	0.0542***	7.00e-07*	0.0952***	0.0543***
	(0.0041)	(4.09e-07)	(0.0034)	(0.0044)
25.[05:40-06:10)	0.0556***	9.29e-07**	0.0952***	0.0557***
	(0.0043)	(4.34e-07)	(0.0036)	(0.0045)
26.[06:10-06:40)	0.0513***	7.24e-07*	0.0956***	0.0514***
	(0.0041)	(4.14e-07)	(0.0036)	(0.0043)
27.[06:40-07:10)	0.0723***	2.43e-06***	0.0943***	0.0724***
	(0.0047)	(6.82e-07)	(0.0034)	(0.0051)
28.[07:10-07:40)	0.0690***	2.01e-06***	0.1141***	0.0691***
	(0.0047)	(5.16e-07)	(0.0043)	(0.0048)
29.[07:40-08:10)	0.1328***	6.13e-06***	0.1264***	0.1330***
	(0.0054)	(5.36e-07)	(0.0046)	(0.0053)
30.[08:10-08:40)	0.1162***	4.55e-06***	0.1731***	0.1164***
	(0.0061)	(6.69e-07)	(0.0047)	(0.0062)
31.[08:40-09:10)	0.1191***	5.50e-06***	0.1814***	0.1193***
	(0.0059)	(7.00e-07)	(0.0054)	(0.0067)
32.[09:10-09:40)	0.1260***	6.60e-06***	0.1792***	0.1262***
	(0.0061)	(7.58e-07)	(0.0054)	(0.0058)
33.[09:40-10:10)	0.1208***	5.84e-06***	0.1893***	0.1210***
	(0.0057)	(7.92e-07)	(0.0055)	(0.0063)
34.[10:10-10:40)	0.1185***	5.91e-06***	0.1893***	0.1186***
	(0.0057)	(6.82e-07)	(0.0053)	(0.0066)
35.[10:40-11:10)	0.0784***	2.41e-06***	0.1636***	0.0785***
	(0.0050)	(6.01e-07)	(0.0051)	(0.0056)
36.[11:10-11:40)	0.0726***	1.67e-06***	0.1292***	0.0727***
	(0.0045)	(4.90e-07)	(0.0043)	(0.0048)
37.[11:40-12:10)	0.0798***	2.23e-06***	0.1179***	0.0800***
	(0.0047)	(4.79e-07)	(0.0040)	(0.0051)
38.[12:10-12:40)	0.0804***	2.41e-06***	0.1268***	0.0805***

	(0.0050)	(5.51e-07)	(0.0045)	(0.0056)
39.[12:40-13:10)	0.0879***	3.86e-06***	0.1212***	0.0881***
	(0.0048)	(4.72e-07)	(0.0052)	(0.0048)
40.[13:10-13:40)	0.0962***	5.22e-06***	0.1789***	0.0963***
	(0.0069)	(7.87e-07)	(0.0055)	(0.0062)
41.[13:40-14:10)	0.0296***	-1.05e-08	0.0944***	0.0297***
	(0.0041)	(4.48e-07)	(0.0039)	(0.0044)
42.[14:10-14:40)	0.0286***	3.51e-07	0.0671***	0.0287***
	(0.0042)	(3.87e-07)	(0.0032)	(0.0045)
43.[14:40-15:10)	0.0128***	-3.31e-07	0.0604***	0.0129***
	(0.0036)	(3.42e-07)	(0.0037)	(0.0037)
44.[15:10-15:40)	0.0023	8.30e-07*	0.0274***	0.0024
	(0.0036)	(4.64e-07)	(0.0028)	(0.0032)
45.[15:40-15:50)	-0.0687***	-2.16e-06***	0.0167***	-0.0686***
	(0.0040)	(3.68e-07)	(0.0028)	(0.0041)
WED_t				-0.0023
				(0.0025)
cons	0.1970***	0.0001***	0.1547***	0.1981***
	(0.0096)	(1.08e-06)	(0.0075)	(0.0163)

Note: The first and last 10 minutes of trading are excluded from the sample. Variable 0. slot30 (17:10-17:40 CT) is used as the baseline with coefficients normalized to zero and therefore is omitted from the table.

The empirical results in this chapter imply that higher curvature is associated with lower depth and wider spreads, while lower curvature corresponds to higher depth and narrower spreads. These effects are not permanent, as all of them typically revert after the shock, indicating the resiliency of the market. The analysis further links curvature to the underlying composition of liquidity between implied depth and outright depth. Convexity is exhibited when orders dominate top-of-book liquidity with market agreement. It increases especially when implied orders are concentrated in the first few steps and outright orders reduced across deeper levels. Conversely, when curvature decreases and the order book becomes more concave, a larger share of orders is placed at deeper levels away from the mid-quote. Since most of implied orders are concentrated near the best quotes, these additional orders mainly come from outright traders. Thus, concavity reflects increased participation of outright orders at distant price levels, which is typically associated with greater market disagreement and heterogeneity of expectations. In this sense, curvature provides a unified framework for interpreting both the static geometry of liquidity and its dynamic recovery process after the shocks, thereby linking the geometric features of the order book with its underlying market structure and adjustment dynamics. At the same time, curvature provides incremental explanatory power in volatility forecasting models. Combined with the conclusions of Hypothesis 1, this incremental contribution indicates that curvature carries information about the market's recovery capacity. In other words, curvature offers forward-looking insights into market stability and serves as an important complement to traditional liquidity measures. Finally, the external information shows that scheduled announcements do not significantly affect the model outcomes in the absence of "surprise" components, the true drivers of volatility remain internal liquidity structures.

7. Discussion

This thesis introduces order book curvature as a novel measure of liquidity in the crude oil futures market and examines its dynamic role and predictive content. The empirical analysis proceeds in three stages. First, curvature is estimated and visualized across 272 trading days of CME crude oil futures, introducing curvature as a new liquidity measure. The results show that both bid and ask sides of the book predominantly exhibit convex structures, with more than 50% of snapshots yielding $\beta > 1$. This indicates that liquidity is typically concentrated near the best quotes, although concavity is also observed, reflecting moments of greater disagreement among market participants. Second, this thesis embed curvature into a high-frequency VAR framework with impulse response functions (IRF), demonstrating the dynamic interactions between curvature, depth, spread, and returns. A positive curvature shock leads to a short-lived decline in overall depth and a temporary widening of spreads, with both measures returning to equilibrium within 60-180 seconds. Depth shocks, conversely, reduce convexity and briefly narrow spreads, while spread shocks raise curvature but lower depth. Across all cases, returns remain largely unaffected, underscoring that these adjustments represent liquidity reallocation rather than informational shocks. Importantly, the rapid recovery of order book variables confirms the resiliency of crude oil futures, though with slower adjustment compared to equities and broadly comparable to other futures markets. Third, the thesis incorporates curvature into volatility forecasting models using the two-scale realized volatility (TSRV) estimator. The results reveal that curvature significantly improves short-horizon volatility forecasts beyond traditional liquidity measures. By contrast, scheduled announcements such as the U.S. EIA inventory reports do not systematically increase short-term volatility once intraday seasonality is accounted for, highlighting that volatility responses are largely conditional on the surprise component of news rather than the announcement itself.

7.1 Implications for Market Microstructure Research

This study contributes to the microstructure and liquidity literature in several important ways. First, it introduces curvature as a novel liquidity metric that complements traditional point-based measures, revisiting liquidity from linear to non-linear representations. Previous studies have typically described the geometry of the limit order book using linear measures such as slope and relative liquidity (Næs & Skjeltorp, 2006; Valenzuela et al., 2015). While slope-based method captures how liquidity changes proportionally across price levels, it cannot differentiate between order books that have similar linear gradients but fundamentally different shapes. By modeling the power-law relationship between cumulative price distance and cumulative depth, the curvature measure introduced in this thesis captures these non-linear characteristics. It thus generalizes the slope framework and offers a higher-order representation of order-book geometry.

Second, the study establishes a log-log curvature estimation framework based on normalized cumulative depth and price deviation. The empirical analysis shows that convex order books dominate in the crude oil futures market, while concave structures appear less frequently. Convex shapes are more common during high-liquidity and high-activity periods, when traders exhibit greater consensus and concentrate their orders near the best quotes. In contrast, concave shapes tend to emerge during less active or uncertain periods. This methodological innovation allows for a standardized evaluation of curvature across markets and improves the empirical robustness of curvature as a liquidity proxy.

Third, the study employs VAR and IRF analyses to evaluate the interrelationship of curvature with other market characteristics of crude oil futures and implements a stepwise volatility prediction framework that sequentially incorporates curvature, depth, spread and slope. It reveals that a positive curvature shock reduces overall depth and temporarily widens spreads, with effects dissipating within two to three minutes. Importantly, these adjustments occur without significant correspond to new information arrival but to internal liquidity restructuring. Economically, a curvature shock can be interpreted as the outcome of traders repositioning their limit orders, either withdrawing outright orders from deeper levels or adding implied orders near the mid-quote. Such behavior is driven by endogenous factors rather than by informational updates. This interpretation aligns with the concept of market resiliency that liquidity disruptions are transient and self-correcting, with structural mechanisms restoring equilibrium after temporary imbalances.

7.2 Limitations and Directions for Further Research

Although this study introduces curvature as a novel liquidity measure and provides empirical evidence of its incremental explanatory power, several limitations should be acknowledged. First, while curvature captures the non-linear distribution of liquidity across the entire order book, it inevitably simplifies the multi-dimensional dynamics of market microstructure into a single numerical value. For instance, high-frequency algorithmic trading strategies, rapid order cancellations, and hidden liquidity may interact with visible depth in complex ways that cannot be fully represented by curvature alone. Second, the regression framework in this study mainly focuses on a limited set of order book variables without incorporating broader market-level factors such as cross-asset liquidity spillovers, market segmentation effects, or volatility clustering. Moreover, the impulse response function results show that the impact of curvature shocks is relatively short-lived, suggesting that other mechanisms may play a more dominant role in longer-term market adjustments. Finally, the scope of this analysis is restricted to the futures market, and it remains uncertain whether curvature can be generalized to other asset classes such as equities, bonds, or cryptocurrencies. While curvature demonstrates both theoretical and empirical value, its practical applications in trading strategy optimization, risk management, and real-time monitoring of market stress have yet to be fully tested in more diverse contexts. Future research may extend the use of curvature to broader market environments to deepen our understanding of liquidity dynamics.

Building on these limitations, several directions for future research can be identified. The incremental value of curvature as a liquidity measure aligns with recent literature that LOBs exhibit consistent shape patterns and scaling laws (Bouchaud et al., 2008), suggesting that shape-based measures may better account for the resiliency and fragility of markets. Empirically, we observe that curvature becomes steeper (more convex) even as average depth declines. This is not contradictory; a stronger convexity indicates that orders cluster more densely near the mid-quote, while liquidity at more distant price levels recedes. This implies a redistribution of liquidity rather than a uniform expansion or contraction. Kyle's (1985) classical framework distinguishes between tightness, depth, and resiliency, and the findings suggest that improvements in tightness (more liquidity close to the mid) may come at the cost of reduced resiliency (weaker buffers against large shocks). This interpretation resonates with the idea that order-book configurations often balance immediacy against longer-horizon absorption capacity (Kyle, 1985).

Finally, while this thesis finds no significant effects of scheduled announcements such as Wednesday EIA reports on predicting volatility, prior research shows that price discovery is primarily driven by the surprise component of announcements rather than their occurrence (Kilian & Vega, 2011). Future work should therefore incorporate forecast-based surprise measures (e.g., survey expectations or futures-implied signals) to test whether volatility responds more strongly to unexpected shocks. More broadly, expanding the empirical scope beyond energy futures to include equities, fixed income, and cryptocurrencies would test the generalizability of curvature as a robust, shape-based proxy for market liquidity.

8. References

Adrangi, B., Chatrath, A., Dhanda, K. K., & Raffiee, K. (2001). Chaos in oil prices? Evidence from futures markets. *Energy Economics*, *23*(4), 405–425. https://doi.org/10.1016/S0140-9883(00)00079-7

Ahn, H., Bae, K., & Chan, K. (2001). Limit orders, depth, and volatility: Evidence from the stock exchange of hong kong. *The Journal of Finance*, *56*(2), 767–788. https://doi.org/10.1111/0022-1082.00345

Aidov, A., & Daigler, R. T. (2015). Depth characteristics for the electronic futures limit order book. *Journal of Futures Markets*, *35*(6), 542–560. https://doi.org/10.1002/fut.21706

Aitken, M., & Comerton-Forde, C. (2003). How should liquidity be measured? *Pacific-Basin Finance Journal*, *11*(1), 45–59. https://doi.org/10.1016/S0927-538X(02)00093-8

Aït-Sahalia, Y., Mykland, P. A., & Zhang, L. (2011). Ultra-high frequency volatility estimation with dependent microstructure noise. *Journal of Econometrics*, *160*(1), 160–175. https://doi.org/10.1016/j.jeconom.2010.03.028

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*(6), 716–723. https://doi.org/10.1109/TAC.1974.1100705

Alquist, R., Kilian, L., Vigfusson, R. J., Elliott, G., & Timmermann, A. (2013). Handbook of economic forecasting. Forecasting the Price of Oil, 2, 427-507.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, *96*(453), 42–55. https://doi.org/10.1198/016214501750332965

Arzandeh, M., & Frank, J. (2019). Price discovery in agricultural futures markets: Should we look beyond the best bid‐ask spread? *American Journal of Agricultural Economics*, *101*(5), 1482–1498. https://doi.org/10.1093/ajae/aaz001

Arzandeh, M., Frank, J., & Daniels, J. (2025). Real‐time tracking of public announcements in the limit order book. *Journal of Futures Markets*, *45*(6), 569–599. https://doi.org/10.1002/fut.22577

Askari, H., & Krichene, N. (2008). Oil price dynamics (2002–2006). *Energy Economics*, *30*(5), 2134–2153. https://doi.org/10.1016/j.eneco.2007.12.004

Balduzzi, P., Elton, E. J., & Green, T. C. (2001). Economic news and bond prices: Evidence from the U.S. treasury market. *The Journal of Financial and Quantitative Analysis*, *36*(4), 523. https://doi.org/10.2307/2676223

Bandi, F. M., & Russell, J. R. (2008). Microstructure noise, realized variance, and optimal sampling. *Review of Economic Studies*, *75*(2), 339–369. https://doi.org/10.1111/j.1467-937X.2008.00474.x

Barnea, A., & Logue, D. E. (1975). The effect of risk on the market maker's spread. *Financial Analysts Journal*, *31*(6), 45–49. https://doi.org/10.2469/faj.v31.n6.45

Basistha, A., & Kurov, A. (2015). The impact of monetary policy surprises on energy prices. *Journal of Futures Markets*, *35*(1), 87–103. https://doi.org/10.1002/fut.21639

Biais, B., Hillion, P., & Spatt, C. (1995). An empirical analysis of the limit order book and the order flow in the Paris Bourse. *the Journal of Finance, 50*(5), 1655-1689. https://doi.org/10.1111/j.1540-6261.1995.tb05192.x

Bloomfield, R., O'Hara, M., & Saar, G. (2005). The "make or take" decision in an electronic market: Evidence on the evolution of liquidity. *Journal of Financial Economics, 75*(1), 165-199. https://doi.org/10.1016/j.jfineco.2004.07.001

Bouchaud, J. P., Farmer, J. D., & Lillo, F. (2009). How markets slowly digest changes in supply and demand. In *Handbook of financial markets: dynamics and evolution* (pp. 57-160). North-Holland.

Bouchaud, J.-P., Kockelkoren, J., & Potters, M. (2006). Random walks, liquidity molasses and critical response in financial markets. *Quantitative Finance*, *6*(2), 115–123. https://doi.org/10.1080/14697680500397623

Bouchaud, J. P., Mézard, M., & Potters, M. (2002). Statistical properties of stock order books: empirical results and models. *Quantitative finance*, *2*(4), 251. https://doi.org/10.1088/1469-7688/2/4/301

Brennan, M. J., & Subrahmanyam, A. (1995). Investment analysis and price formation in securities markets. *Journal of Financial Economics*, *38*(3), 361–381. https://doi.org/10.1016/0304-405X(94)00811-E

Brockman, P., & Chung, D. Y. (1999). Bid‐ask spread components in an order‐driven environment. *Journal of Financial Research*, *22*(2), 227–246. https://doi.org/10.1111/j.1475-6803.1999.tb00724.x

Bu, H. (2014). Effect of inventory announcements on crude oil price volatility. *Energy Economics*, *46*, 485–494. https://doi.org/10.1016/j.eneco.2014.05.015

Chiang, R., & Venkatesh, P. C. (1988). Insider holdings and perceptions of information asymmetry: A note. *The journal of finance*, *43*(4), 1041-1048. https://doi.org/10.2307/2328153

Chordia, T., Roll, R., & Subrahmanyam, A. (2011). Recent trends in trading activity and market quality. *Journal of Financial Economics*, *101*(2), 243–263. https://doi.org/10.1016/j.jfineco.2011.03.008

Cont, R., Stoikov, S., & Talreja, R. (2010). A stochastic model for order book dynamics. *Operations research, 58*(3), 549-563.https://doi.org/10.1287/opre.1090.0780

Cont, R., Kukanov, A., & Stoikov, S. (2014). The price impact of order book events. *Journal of financial econometrics*, *12*(1), 47-88. https://doi.org/10.1093/jjfinec/nbt003

Cont, R., & Kukanov, A. (2017). Optimal order placement in limit order markets. *Quantitative Finance*, *17*(1), 21–39. https://doi.org/10.1080/14697688.2016.1190030

Coppola, A. (2008). Forecasting oil price movements: Exploiting the information in the futures market. *Journal of Futures Markets*, *28*(1), 34–56. https://doi.org/10.1002/fut.20277

Duong, H. N., & Kalev, P. S. (2008). The samuelson hypothesis in futures markets: An analysis using intraday data. *Journal of Banking & Finance*, *32*(4), 489–500. https://doi.org/10.1016/j.jbankfin.2007.06.011

Easley, D., López De Prado, M. M., & O'Hara, M. (2012). Flow Toxicity and Liquidity in a High-frequency World. *Review of Financial Studies*, *25*(5), 1457–1493. https://doi.org/10.1093/rfs/hhs053

Ederington, L. H., & Lee, J. H. (1993). How markets process information: News releases and volatility. *The Journal of Finance, 48*(4), 1161-1191. https://doi.org/10.1111/j.1540-6261.1993.tb04750.x

Farmer, J. D., Gerig, A., Lillo, F., & Waelbroeck, H. (2008). A new approach to understanding the market impact of large trading orders. *Preprint*.

Fong, W. M., & See, K. H. (2002). A markov switching model of the conditional volatility of crude oil futures prices. *Energy Economics*, *24*(1), 71–95. https://doi.org/10.1016/S0140-9883(01)00087-1

Foucault, T., Kadan, O., & Kandel, E. (2005). Limit Order Book as a Market for Liquidity. *Review of Financial Studies*, *18*(4), 1171–1217. https://doi.org/10.1093/rfs/hhi029

Foucault, T., Moinas, S., & Theissen, E. (2007). Does anonymity matter in electronic limit order markets? *Review of Financial Studies*, *20*(5), 1707–1747. https://doi.org/10.1093/rfs/hhm027

Frino, A., Lepone, A., & Wearin, G. (2008). Intraday behavior of market depth in a competitive dealer market: A note. *Journal of Futures Markets*, *28*(3), 294–307. https://doi.org/10.1002/fut.20300

Glosten, L. R. (1994). Is the electronic open limit order book inevitable? *The Journal of Finance*, *49*(4), 1127-1161. https://doi.org/10.1111/j.1540-6261.1994.tb02450.x

Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, *14*(1), 71–100. https://doi.org/10.1016/0304-405X(85)90044-3

Goettler, R. L., Parlour, C. A., & Rajan, U. (2005). Equilibrium in a dynamic limit order market. *The Journal of Finance*, *60*(5), 2149–2192. https://doi.org/10.1111/j.1540-6261.2005.00795.x

Hannan, E. J., & Quinn, B. G. (1979). The determination of the order of an autoregression. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *41*(2), 190–195. https://doi.org/10.1111/j.2517-6161.1979.tb01072.x

Harris, L. (1990). Liquidity, trading rules and electronic trading systems. Working Papers, Southern California School of Business Administration, New York. (No. 91-8). https://econpapers.repec.org/paper/ fthsocabu/91-8.htm

Harris, L. E., & Panchapagesan, V. (2005). The information content of the limit order book: Evidence from NYSE specialist trading decisions. *Journal of Financial Markets*, *8*(1), 25–67. https://doi.org/10.1016/j.finmar.2004.07.001

Hasbrouck, J. (1991). Measuring the information content of stock trades. *The Journal of Finance*, *46*(1), 179-207. https://doi.org/10.1111/j.1540-6261.1991.tb03749.x

Hasbrouck, J. (2009). Trading Costs and Returns for U.S. Equities: Estimating Effective Costs from Daily Data. *The Journal of Finance*, *64*(3), 1445–1477. https://doi.org/10.1111/j.1540-6261.2009.01469.x

Hasbrouck, J., & Seppi, D. J. (2001). Common factors in prices, order flows, and liquidity. *Journal of financial Economics*, *59*(3), 383-411. https://doi.org/10.1016/S0304-405X(00)00091-X

He, X., Serra, T., & Garcia, P. (2021). Resilience in “flash events” in the corn and lean hog futures markets*. American Journal of Agricultural Economics, 103*(2), 743-764. https://doi.org/10.1111/ajae.12146

Hedi Arouri, M. E., & Khuong Nguyen, D. (2010). Oil prices, stock markets and portfolio investment: Evidence from sector analysis in europe over the last decade. *Energy Policy*, *38*(8), 4528–4539. https://doi.org/10.1016/j.enpol.2010.04.007

Hillman, R., Marsh, I. W., & Salmon, M. (2001). Liquidity in a limit order book foreign exchange trading system. *FERC Discussion Paper.*

Huang, R. D., & Stoll, H. R. (1996). Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics*, *41*(3), 313–357. https://doi.org/10.1016/0304-405X(95)00867-E

Jain, P. K., Jain, P., & McInish, T. H. (2011). The predictive power of limit order book for future volatility, trade price, and speed of trading. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.1787625

Kempf, A., Mayston, D., & Yadav, P. K. (2009). Resiliency in limit order book markets: A dynamic view of liquidity. *In AFA 2009 San Francisco Meetings Paper* (pp. 1-39). https://doi.org/10.2139/ssrn.967249

Kilian, L., & Vega, C. (2011). Do energy prices respond to US macroeconomic news? A test of the hypothesis of predetermined energy prices. *Review of Economics and Statistics, 93*(2), 660-671. https://doi.org/10.1162/REST_a_00086

Kyle, A. S. (1985). Continuous Auctions and Insider Trading. *Econometrica*, *53*(6), 1315. https://doi.org/10.2307/1913210

Lang, K., & Auer, B. R. (2020). The economic and financial properties of crude oil: A review. *The North American Journal of Economics and Finance*, *52*, 100914. https://doi.org/10.1016/j.najef.2019.01.011

Large, J. (2007). Measuring the resiliency of an electronic limit order book. *Journal of Financial Markets, 10*(1), 1-25. https://doi.org/10.1016/j.finmar.2006.09.001

Larsson, K., & Nossman, M. (2011). Jumps and stochastic volatility in oil prices: Time series evidence. *Energy Economics*, *33*(3), 504–514. https://doi.org/10.1016/j.eneco.2010.12.016

Lee, C. M., Mucklow, B., & Ready, M. J. (1993). Spreads, depths, and the impact of earnings information: An intraday analysis. *The Review of Financial Studies, 6*(2), 345-374. https://doi.org/10.1093/rfs/6.2.345

Lo, D. K., & Hall, A. D. (2015). Resiliency of the limit order book. *Journal of Economic Dynamics and Control*, *61,* 222-244. https://doi.org/10.1016/j.jedc.2015.09.012

Lütkepohl, H. (2013). Introduction to multiple time series analysis. Springer Science & Business Media.

Madhavan, A., Porter, D., & Weaver, D. (2005). Should securities markets be transparent? *Journal of Financial Markets*, *8*(3), 265–287. https://doi.org/10.1016/j.finmar.2005.05.001

Marshall, B. R., Nguyen, N. H., & Visaltanachoti, N. (2012). Commodity liquidity measurement and transaction costs. *The Review of Financial Studies, 25*(2), 599-638. https://doi.org/10.1093/rfs/hhr075

Menkveld, A. J., & Yueshen, B. Z. (2019). The Flash Crash: A Cautionary Tale About Highly Fragmented Markets. *Management Science*, *65*(10), 4470–4488. https://doi.org/10.1287/mnsc.2018.3040

Miao, H., Ramchander, S., Wang, T., & Yang, J. (2018). The impact of crude oil inventory announcements on prices: Evidence from derivatives markets. *Journal of Futures Markets*, *38*(1), 38–65. https://doi.org/10.1002/fut.21850

Næs, R., & Skjeltorp, J. A. (2006). Order book characteristics and the volume–volatility relation: Empirical evidence from a limit order market. *Journal of Financial Markets*, *9*(4), 408-432 https://doi.org/10.1016/j.finmar.2006.04.001

Nikkinen, J., & Rothovius, T. (2019). Market specific seasonal trading behavior in NASDAQ OMX electricity options. *Journal of Commodity Markets*, *13*, 16–29. https://doi.org/10.1016/j.jcomm.2018.05.002

Pascual, R., & Veredas, D. (2010). Does the open limit order book matter in explaining informational volatility? *Journal of Financial Econometrics*, *8*(1), 57-87. https://doi.org/10.1093/jjfinec/nbp021

Pindyck, R. S. (2004). Volatility and commodity price dynamics. *Journal of Futures Markets*, *24*(11), 1029–1047. https://doi.org/10.1002/fut.20120

Plerou, V., Gopikrishnan, P., & Stanley, H. E. (2005). Quantifying fluctuations in market liquidity: Analysis of the bid-ask spread. *Physical Review E*, *71*(4), 046131. https://doi.org/10.1103/PhysRevE.71.046131

Pástor, Ľ., & Stambaugh, R. F. (2003). Liquidity risk and expected stock returns. J*ournal of Political Economy*, 111(3), 642-685. https://doi.org/10.1086/374184

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of finance*, 39(4), 1127-1139. https://doi.org/10.1111/j.1540-6261.1984.tb03897.x

Rosa, C. (2014). The high-frequency response of energy prices to U.S. monetary policy: Understanding the empirical evidence. *Energy Economics*, *45*, 295–303. https://doi.org/10.1016/j.eneco.2014.06.011

Roşu, I. (2009). A dynamic model of the limit order book. *Review of Financial Studies*, *22*(11), 4601–4641. https://doi.org/10.1093/rfs/hhp011

Rühl, T. R., & Stein, M. (2015). Discovering and disentangling effects of US macro-announcements for european stocks. *SSRN Electronic Journal.* https://doi.org/10.2139/ssrn.2648340

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*(2). https://doi.org/10.1214/aos/1176344136

Stoll, H. R. (1989). Inferring the components of the bid-ask spread: Theory and empirical tests. *the Journal of Finance*, *44*(1), 115-134. https://doi.org/10.1111/j.1540-6261.1989.tb02407.x

Stoll, H. R. (2000). Presidential address: Friction. *The Journal of Finance*, *55*(4), 1479–1514. https://doi.org/10.1111/0022-1082.00259

Tian, X., Duong, H. N., & Kalev, P. S. (2019). Information content of the limit order book for crude oil futures price volatility. *Energy Economics*, *81*, 584–597. https://doi.org/10.1016/j.eneco.2019.04.026

Tóth, B., Lempérière, Y., Deremble, C., De Lataillade, J., Kockelkoren, J., & Bouchaud, J.-P. (2011). Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, *1*(2), 021006. https://doi.org/10.1103/PhysRevX.1.021006

Valenzuela, M., Zer, I., Fryzlewicz, P., & Rheinländer, T. (2015). Relative liquidity and future volatility. *Journal of Financial Markets, 24,* 25-48. https://doi.org/10.1016/j.finmar.2015.03.001

Yeo, W. Y. (2009). Limit order book liquidity and liquidity imbalance. *Available at SSRN 1342228*.

Zhang, L., Mykland, P. A., & Aït-Sahalia, Y. (2005). A tale of two-time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association*, *100*(472), 1394–1411. https://doi.org/10.1198/016214505000000169

9. Appendix

Table A.1 Distribution of Curvature Estimates

		β^{bid}		
	Percentiles	Smallest		
1%	0.547106	0.3055803		
5%	0.6832379	0.3260697		
10%	0.7723898	0.3354709		
25%	0.9170065	0.4051679	Obs.	24,947
50%	1.06137		Mean	1.106773
		Largest	Std. dev.	0.340306
75%	1.216991	5.167073		
90%	1.470036	5.816435	Variance	0.115808
95%	1.707222	6.390694	Skewness	2.758549
99%	2.258166	6.671064	Kurtosis	23.03837
		p-value of β^{bid}		
	Percentiles	Smallest		
1%	8.52e-15	9.04e-18		
5%	3.91e-13	9.65e-18		
10%	3.40e-12	1.01e-17		
25%	1.24e-10	1.08e-17	Obs.	24,947
50%	1.28e-08		Mean	0.0000166
		Largest	Std. dev.	0.0001685
75%	2.56e-06	0.0070398		
90%	0.0000337	0.0113239	Variance	2.84e-08
95%	0.0000853	0.0133735	Skewness	67.32083
99%	0.0002115	0.0152892	Kurtosis	5294.715
		R^2 of β^{bid}		
	Percentiles	Smallest		
1%	0.7367421	0.37127		
5%	0.819912	0.500702		
10%	0.8650419	0.5412996		
25%	0.9583351	0.5499451	Obs.	24,947
50%	0.9878259		Mean	0.959891
		Largest	Std. dev.	0.0611349
75%	0.9950385	0.9998246		
90%	0.9975522	0.9998671	Variance	0.0037375
95%	0.9983655	0.9998863	Skewness	-2.287116
99%	0.9992505	0.9998871	Kurtosis	8.580561
		β^{ask}		
	Percentiles	Smallest		
1%	0.552433	0.3137686		
5%	0.689023	0.3274305		
10%	0.778427	0.3582820		
25%	0.919933	0.3668813	Obs.	24,947
50%	1.062695		Mean	1.120379
		Largest	Std. dev.	0.3706074
75%	1.227624	6.561538		
90%	1.514221	6.648465	Variance	0.1373499

95%	1.758515	7.72354	Skewness	4.087238
99%	2.372239	11.9297	Kurtosis	57.71422
		p-value of β^{ask}		
	Percentiles	Smallest		
1%	7.63e-15	1.44e-18		
5%	3.28e-13	3.44e-18		
10%	2.87e-12	7.18e-18		
25%	1.19e-10	1.98e-17	Obs.	24,947
50%	1.29e-08		Mean	0.0000178
		Largest	Std. dev.	0.0003018
75%	2.46e-06	0.005451		
90%	0.0000333	0.0138912	Variance	9.11e-08
95%	0.0000801	0.0147075	Skewness	111.9502
99%	0.0002124	0.0413445	Kurtosis	14499.26
		R^2 of β^{ask}		
	Percentiles	Smallest		
1%	0.7337643	0.4727701		
5%	0.8197082	0.5361312		
10%	0.8694011	0.554452		
25%	0.9585199	0.5582398	Obs.	24,947
50%	0.9879638		Mean	0.9606
		Largest	Std. dev.	0.0605572
75%	0.9951115	0.9998446		
90%	0.9975872	0.9998616	Variance	0.0036672
95%	0.998401	0.9998818	Skewness	-2.327935
99%	0.9992537	0.9999145	Kurtosis	8.63118

Note: β denotes the bid-side and ask-side curvature (estimated at each snapshot from the regression lnH%~lnQ%). P-value is the two-sided p-value for this curvature; R^2 is the corresponding regression's R^2. The table reports percentiles (1%, 5%, …, 99%), extrema, the mean, and the standard deviation.

Table A.2 Lag Order Selection for VAR Model

Lag	**LogL**	**LR**	**df**	**p-value**	**FPE**	**AIC**	**HQIC**	**SBIC**
0	306915				2.3e-17	-26.978	-26.977	-26.976
1	335772	57713	16	0.000	1.8e-18	-29.513	-29.510	-29.506
2	338925	6306.7	16	0.000	1.4e-18	-29.789	-29.784	-29.776
3	339868	1885.1	16	0.000	1.3e-18	-29.87	-29.864	-29.852
4	340225	714.77*	16	0.000	1.2e-18*	-29.9*	-29.892*	-29.876*

Note: * indicates the optimal lag length selected by each criterion.

Table A.3 Multicollinearity Diagnostics: VIF and Tolerance

Variable	VIF	Tolerance (1/VIF)
$TSRV_t$	2.01	0.496552
β_t^{ask}	1.26	0.792290
β_t^{bid}	1.23	0.816016
$DEPTH_t$	2.00	0.499572
$SPREAD_t$	1.44	0.692271
Slot30		
1	1.96	0.509230
2	1.97	0.507204
3	1.98	0.506173
4	2.03	0.492499
5	2.03	0.492140
6	2.04	0.489010
7	2.05	0.488528
8	2.04	0.491332
9	2.04	0.490831
10	2.03	0.492132
11	2.03	0.491964
12	2.03	0.491874
13	1.53	0.655719
14	2.04	0.489359
15	2.04	0.490353
16	2.07	0.482147
17	2.09	0.479206
18	2.12	0.470812
19	2.13	0.470423
20	2.16	0.462761
21	2.15	0.464560
22	2.16	0.462061
23	2.14	0.466842
24	2.15	0.466092
25	2.14	0.466496
26	2.15	0.465586
27	2.15	0.466173
28	2.17	0.460423
29	2.23	0.448904
30	2.39	0.418760
31	2.38	0.420354
32	2.38	0.420872
33	2.43	0.411383
34	2.43	0.411380
35	2.35	0.426004
36	2.24	0.445466
37	2.21	0.451715
38	2.23	0.447962
39	2.19	0.455864
40	2.39	0.419034
41	2.13	0.468946
42	2.08	0.481733

43	2.03	0.491441
44	1.96	0.510185
45	1.50	0.666164
Mean VIF	2.06	

Note: VIFs for all regressors. Intraday fixed effects (Slot30 dummies) range 1.50-2.43. Tolerance = 1/VIF.

Table A.4 Pairwise Correlation Matrix - Trimmed Window

	$TSRV_t$	**RV_t**	**β_t^{ask}**	**β_t^{bid}**	**$DEPTH_t$**	**$SPREAD_t$**	**Slope**
$TSRV_t$	1.00						
RV_t	0.87	1.00					
β_t^{bid}	-0.07	-0.06	1.00				
β_t^{ask}	-0.06	-0.05	-0.13	1.00			
$DEPTH_t$	0.08	-0.01	-0.29	-0.25	1.00		
$SPREAD_t$	0.08	0.09	0.26	0.25	-0.49	1.00	
Slope	-0.11	-0.14	-0.04	-0.03	-0.63	-0.56	1.00

Note: Pearson correlations computed on the trimmed window (N=24,947).

Table A.5 VAR Estimation of Return, Curvature, Depth and Spread

Predictor	**L1**	**L2**	**L3**	**L4**
Panel A: Dependent Variable = $RETURN_t$				
$RETURN_t$	-0.0091 (0.0068)	0.0010 (0.0068)	0.0128* (0.0068)	0.0033 (0.0068)
β_t	-5.23e-06 (3.82e-06)	-4.73e-06 (3.80e-06)	1.92e-06 (3.76e-06)	3.60e-06 (3.67e-06)
$DEPTH_t$	-6.04e-06* (3.42e-06)	-4.96e-07 (3.82e-06)	-2.33e-06 (3.94e-06)	0.00001*** (3.46e-06)
$SPREAD_t$	-0.0004*** (0.0001)	0.0002* (0.0001)	0.0001 (0.0001)	-0.0004*** (0.0001)
Cons	-9.36e-06 (0.00001)			
Panel B: Dependent Variable = β_t				
$RETURN_t$	-25.9498** (12.1159)	-2.6974 (12.1282)	-4.7717 (12.1620)	4.6867 (12.1411)
β_t	0.1657*** (0.0068)	0.1387*** (0.0068)	0.1364*** (0.0067)	0.1032*** (0.0066)
$DEPTH_t$	-0.0310*** (0.0061)	-0.0026 (0.0068)	-0.0316*** (0.0070)	-0.0105* (0.0062)
$SPREAD_t$	0.4222* (0.2458)	-0.6395*** (0.2457)	-0.2917 (0.2529)	0.5360** (0.2503)
Cons	0.9282*** (0.0263)			

Panel C: Dependent Variable = DEPTH$_t$				
RETURN$_t$	15.9977 (13.3807)	20.6092 (13.3943)	4.0813 (13.4317)	35.5856*** (13.4085)
β_t	-0.0773*** (0.0075)	-0.0219*** (0.0075)	-0.0197*** (0.0074)	-0.0019 (0.0072)
DEPTH$_t$	0.5042*** (0.0067)	0. 2422*** (0.0075)	0.0875*** (0.0077)	0.0796*** (0.0068)
SPREAD$_t$	0.5052* (0.2715)	-0.7216*** (0.2714)	-1.0661*** (0.2793)	-1.4507*** (0.2764)
Cons	0.6801*** (0.0291)			
Panel D: Dependent Variable = SPREAD$_t$				
RETURN$_t$	-0.2780 (0.3427)	0.3366 (0.3431)	0.6367* (0.3440)	-0.1865 (03434)
β_t	0.0005** (0.0002)	0.0005** (0.0002)	0.0002 (0.0002)	-0.0001 (0.0002)
DEPTH$_t$	-0.0018*** (0.0002)	-0.0008*** (0.0002)	-0.0002 (0.0002)	-0.0001 (0.0002)
SPREAD$_t$	0.1012*** (0.0070)	0.0905*** (0.0070)	0.0948*** (0.0072)	0.1040*** (0.0071)
Cons	0.0258*** (0.0007)			

Note: The optimal lag length for the VAR model was determined using standard lag-order selection criteria. As shown in Appendix Table A.2, the information criteria (FPE, AIC, HQIC, and SBIC) all suggest four lags as optimal. Accordingly, we estimated a VAR (4) model including return, order book curvature (β), market depth, and the relative bid-ask spread.