Order Book Implied Moments

JinGi Ha, Jianfeng Hu, and Xingtan Zhang*

March 2026

Preliminary draft

Comments welcome.

Abstract

We develop a theoretical framework that links the limit order book to the moments of the asset payoff distribution. The n-th derivative of the demand schedule at the expected payoff encodes the $(n+1)$-th central moment, combined with investor preference parameters. We show how to invert this relationship to extract order-book-implied volatility, skewness, and, in principle, all higher moments from level-2 market data. We focus empirically on order book implied volatility. Using Nasdaq order book data, we find that order-book-implied volatility at the market open predicts same-day realized volatility, even after controlling for lagged realized volatility and option-implied volatility. This predictive power strengthens on days with pre-market macroeconomic announcements, when forward-looking information is most valuable, while backward-looking measures become less informative. Order-book-implied volatility also predicts cross-sectional variation in stock returns. Because every electronically traded security has an order book, our approach provides a source of forward-looking distributional information that does not require options markets.

*Ha is with Soongsil University and can be reached by email at jgha@ssu.ac.kr. Hu is with Singapore Management University and can be reached by email at jianfenghu@smu.edu.sg. Zhang is with Cheung Kong Graduate School of Business and can be reached by email at xingtanzhang@ckgsb.edu.cn.

1 Introduction

The moments of the return distribution—volatility, skewness, and higher moments—are among the most basic building blocks of financial economics. They shape portfolio choice, determine risk premia, and drive how markets respond to news. Yet measuring these moments accurately remains challenging. Existing approaches are either backward-looking, relying on the history of past returns (e.g., Andersen et al., 2003; Corsi, 2009), or forward-looking but limited to the subset of stocks with actively traded options (e.g., Britten-Jones and Neuberger, 2000; Bakshi et al., 2003).

In this paper, we propose a new way to estimate forward-looking moments directly from the limit order book. The order book records the prices and quantities investors are willing to trade at each moment in time. These bids and offers encode beliefs, risk assessments, and preferences even before transactions occur. By inverting the order book through the lens of portfolio-choice theory, we recover estimates of volatility and skewness that reflect investors' current willingness to bear risk. Unlike realized measures, these moments are forward-looking. Unlike option-implied measures, they apply to virtually all electronically traded stocks, regardless of whether the firm has listed options. Because order books update continuously, our approach can capture changes in beliefs in real time.

The theoretical contribution of this paper is to show that the n-th derivative of the demand schedule at the expected payoff encodes the $(n+1)$-th central moment of the payoff distribution, combined with preference parameters up to the $(n+1)$-th derivative of the utility function. The first derivative (the slope) reveals variance; the second derivative (the curvature) reveals skewness; and, in principle, higher-order derivatives reveal kurtosis, hyperskewness, and all subsequent moments. While we focus our empirical implementation on variance, which can be estimated reliably from standard level-2 data with four or five price levels on each side, the framework provides a theoretical recipe for extracting the

entire payoff distribution from the order book.

Empirically, we construct a forward-looking volatility measure—LOBVol—for a large cross-section of Nasdaq stocks from 1999 to 2008. Our main findings are threefold. First, LOBVol measured at the market open significantly predicts same-day realized volatility, even after controlling for lagged realized volatility, option-implied volatility, and other standard predictors. Second, this predictive power is amplified on days with pre-market macroeconomic announcements, precisely when forward-looking information is most valuable, while lagged realized volatility becomes less informative. This asymmetric pattern provides strong evidence that our measure captures genuinely forward-looking content, not merely a noisy proxy for past volatility. Third, LOBVol is negatively associated with same-day stock returns in the cross-section, consistent with a risk-return tradeoff at the daily frequency.

Related Literature

Our paper contributes to several strands of the literature. First, we contribute to the large body of work on volatility estimation and forecasting. The realized volatility literature, pioneered by Andersen and Bollerslev (1998) and Barndorff-Nielsen and Shephard (2002), uses high-frequency returns to construct nonparametric estimators of price variation. Andersen et al. (2003) develop forecasting models based on realized volatility, and Corsi (2009) proposes the HAR model that captures the multi-horizon persistence of volatility. These methods have been refined to handle microstructure noise (Zhang et al., 2005; Barndorff-Nielsen et al., 2008), jumps (Barndorff-Nielsen and Shephard, 2004; Tauchen and Zhou, 2011), and leverage effects (Patton and Sheppard, 2015). The GARCH family of models (Engle, 1982; Bollerslev, 1986) provides a complementary parametric approach. All of these methods are fundamentally backward-looking: they estimate or fore-

cast volatility from the history of past returns. Our contribution is to extract volatility from the current state of the order book, producing a measure that is forward-looking by construction. We show empirically that this forward-looking content is not redundant with lagged realized volatility, and that it becomes particularly valuable on days when new information renders backward-looking measures less reliable.

Second, we contribute to the literature on extracting forward-looking information from option prices. Britten-Jones and Neuberger (2000) and Jiang and Tian (2005) develop model-free methods for extracting implied variance from option prices. Bakshi et al. (2003) show how to extract skewness and kurtosis from option portfolios. Bollerslev et al. (2009) use the gap between option-implied and realized variance to measure the variance risk premium, which predicts aggregate stock returns. Our approach shares the forward-looking philosophy of this literature but removes its most binding constraint: the requirement that active, liquid option markets exist. This is practically important because options are largely illiquid at the market open—precisely when forward-looking information about the trading day ahead is most valuable—and many stocks, including small-caps, international equities, and fixed-income instruments, have no listed options at all. Our framework therefore extends the reach of forward-looking moment estimation to the entire universe of electronically traded securities.

Third, we contribute to the literature on the informational content of limit order books. Glosten (1994) provides the theoretical foundations for limit order markets. Biais et al. (1995) document empirical regularities in order flow and book dynamics. Cao et al. (2009) show that the order book contains information about future price movements beyond the best quotes. Cont et al. (2014) model the price impact of order book events. More recently, several papers have explored extracting volatility-related quantities from the order book using approaches grounded in financial mathematics or statistical modeling (Horst and Xu, 2019; Cont and Mueller, 2023). Relative to this literature, our contribution is twofold.

We provide an economic theory—rooted in portfolio choice under expected utility—that explicitly links the shape of the demand schedule to the moments of the payoff distribution and the curvature of the utility function. This structural foundation gives each derivative of the demand schedule a clear economic interpretation, and it extends naturally beyond variance to skewness, kurtosis, and all higher moments—a generality that purely statistical approaches do not offer.

Fourth, our work on order-book-implied skewness connects to the literature on skewness preferences and asset pricing. Kraus and Litzenberger (1976) introduce skewness into the CAPM framework. Harvey and Siddique (2000) show that conditional skewness helps explain the cross-section of returns. Conrad et al. (2013) extract ex ante skewness from options and find a negative relationship with future returns. Amaya et al. (2015) study realized skewness as a return predictor. Despite the theoretical importance of skewness, empirical evidence on the skewness premium remains mixed, with different studies reaching conflicting conclusions depending on the skewness measure used. Our order-book-implied skewness offers a new, universally available measure that does not rely on option markets and may therefore help resolve these discrepancies by providing a consistent forward-looking skewness signal across the entire cross-section of stocks.

Fifth, our empirical finding that LOBVol's predictive power strengthens on macroeconomic announcement days connects to the literature on the asset pricing effects of macroeconomic news. Savor and Wilson (2014) show that stocks earn higher average returns on announcement days, consistent with the resolution of macroeconomic risk. Lucca and Moench (2015) document a pre-FOMC announcement drift in equity prices. Our results complement these findings by showing that the order book is a particularly rich source of forward-looking information on announcement days: when investors have had time to process new macroeconomic releases before the opening auction, the order book reflects their updated beliefs, and LOBVol becomes a stronger predictor of realized volatility while

backward-looking measures become weaker. This asymmetric pattern provides the most direct evidence that our measure captures genuinely forward-looking content.

2 Theory

This section develops the theoretical link between the limit order book and the moments of the payoff distribution. We proceed in three steps. First, we show how level-2 market data map into an aggregate demand schedule. Second, we derive a local approximation of the demand schedule that expresses its slope and curvature in terms of payoff variance, skewness, and preference parameters. Third, we show how to invert these relationships to extract implied volatility, implied skewness, and implied prudence from observable order book data.

2.1 From the Order Book to the Demand Schedule

A level-2 order book snapshot lists the prices and quantities at each level on the bid and ask sides:

$$\begin{array}{cc} \vdots & \\ bp_2 & bv_2 \\ bp_1 & bv_1 \\ \hline ap_1 & av_1 \\ ap_2 & av_2 \\ \vdots & \end{array}$$

where bp_i denotes the i-th best bid price, bv_i the corresponding bid volume, ap_i the i-th best ask price, and av_i the corresponding ask volume. Prices are ordered as $\cdots < bp_2 < bp_1 < ap_1 < ap_2 < \cdots$.

We interpret the order book as a discretized demand schedule. Each price level reveals

how many shares investors are willing to buy or sell *at or better than* that price. Specifically, at bid price bp_k, the cumulative quantity demanded is $\sum_{i=1}^{k} bv_i$; at ask price ap_k, the cumulative quantity supplied is $-\sum_{i=1}^{k} av_i$. The resulting demand schedule is:

$$\begin{array}{cc} \vdots & \\ bp_2 & bv_1 + bv_2 \\ bp_1 & bv_1 \\ ap_1 & -av_1 \\ ap_2 & -(av_1 + av_2) \\ \vdots & \end{array}$$

We adopt the convention that positive quantities denote demand (bid side) and negative quantities denote supply (ask side). This mapping transforms the order book into a function $x(p)$ that records the net number of shares investors wish to hold at each price p—precisely the object that standard portfolio-choice theory predicts.

2.2 A Model of the Demand Schedule

Consider a two-period financial market, $t \in \{0, 1\}$. A risk-free asset has both its price and payoff normalized to one. A single risky asset has time-0 price p and time-1 payoff $\tilde{\theta}$. We denote the mean payoff by $\mu = E[\tilde{\theta}]$, the payoff variance by $\sigma^2 = E[(\tilde{\theta} - \mu)^2]$, and the (standardized) skewness by $s = E[(\tilde{\theta} - \mu)^3]/\sigma^3$.

A representative investor with von Neumann–Morgenstern utility $u(w)$ and initial wealth w_0 submits a demand schedule $x(p)$—the number of shares she would buy or sell at each price p. We impose the following standard assumptions on u.

Assumption 1. The utility function $u : \mathbb{R} \to \mathbb{R}$ is strictly increasing, strictly concave, and thrice continuously differentiable. It exhibits nonincreasing absolute risk aversion

(NARA): $-u''(w)/u'(w)$ is nonincreasing in w.

The NARA condition ensures that risky assets with positive expected excess returns are normal goods. The demand schedule solves

$$x(p) = \arg\max_x \; E\Big[u\Big(w_0 + x(\tilde{\theta} - p)\Big)\Big]. \tag{1}$$

Because u is strictly concave, the first-order condition (FOC) is both necessary and sufficient:

$$E\Big[u'\Big(w_0 + x(p)(\tilde{\theta} - p)\Big)\,(\tilde{\theta} - p)\Big] = 0. \tag{2}$$

Demand at the expected payoff. When $p = \mu$, setting $x = 0$ satisfies the FOC (2), since $E[u'(w_0)(\tilde{\theta} - \mu)] = u'(w_0) \cdot 0 = 0$. Strict concavity guarantees uniqueness, so $x(\mu) = 0$: the investor neither buys nor sells when the price equals the expected payoff.

We are interested in the local behavior of $x(p)$ near $p = \mu$. The following lemma, proved in Appendix A, characterizes the first- and second-order terms.

Lemma 1. *Under Assumption 1, the demand schedule satisfies*

$$x(p) = \underbrace{\frac{u'(w_0)}{u''(w_0)}\frac{1}{\sigma^2}}_{\equiv\,\alpha}(p - \mu) \;-\; \underbrace{\frac{1}{2}\frac{u'''(w_0)}{u''(w_0)}\left(\frac{u'(w_0)}{u''(w_0)}\right)^2\frac{s}{\sigma^3}}_{\equiv\,\beta/2}(p - \mu)^2 \;+\; o\big((p - \mu)^2\big), \tag{3}$$

where $o((p-\mu)^2)$ denotes terms that vanish faster than $(p-\mu)^2$ as $p \to \mu$.

The demand schedule is locally characterized by two quantities: its *slope* α and its *curvature* β. We discuss each in turn.

Slope and volatility. The slope of the demand schedule at $p = \mu$ is

$$x'(\mu) = \alpha = \frac{u'(w_0)}{u''(w_0)}\frac{1}{\sigma^2}. \tag{4}$$

Because $u' > 0$ and $u'' < 0$, the slope is negative: demand falls as the price rises, as expected. The slope is the product of two components: the investor's *risk tolerance* $-u'(w_0)/u''(w_0)$ and the *inverse payoff variance* $1/\sigma^2$. A more risk-tolerant investor or a less volatile asset produces a steeper (more negative) demand schedule—intuitively, the investor is willing to absorb larger positions for a given price concession.

As a benchmark, consider the textbook case of CARA utility $u(w) = -\exp(-\gamma w)$ and normally distributed payoffs $\tilde{\theta} \sim N(\mu, \sigma^2)$. In this case, skewness is zero, and the demand schedule is exactly linear:

$$x(p) = \frac{\mu - p}{\gamma \sigma^2}, \tag{5}$$

with slope $x'(\mu) = -1/(\gamma\sigma^2)$.

Curvature and skewness. The curvature of the demand schedule at $p = \mu$ is

$$x''(\mu) = \beta = \frac{u'''(w_0)}{-u''(w_0)} \left(\frac{u'(w_0)}{-u''(w_0)} \right)^2 \frac{s}{\sigma^3}. \tag{6}$$

Under NARA, $u'' < 0$ and $u''' > 0$ (the latter following from nonincreasing absolute risk aversion). Therefore, the sign of $x''(\mu)$ is determined entirely by the sign of skewness s:

- If $s > 0$ (positive skewness), then $x''(\mu) > 0$: the demand schedule is locally *convex*.
- If $s < 0$ (negative skewness), then $x''(\mu) < 0$: the demand schedule is locally *concave*.
- If $s = 0$ (zero skewness), then $x''(\mu) = 0$: the demand schedule is locally *linear*.

The economic intuition is illustrated in Figure 1. A NARA investor prefers positive skewness. When the underlying asset is positively skewed, the investor demands weakly more shares at every price relative to the zero-skewness benchmark, holding other moments

fixed.[1] Since $x(\mu) = 0$ regardless of skewness (by the FOC), the demand curve for a positively skewed asset must lie weakly above the linear benchmark on both sides of μ—which requires positive curvature.

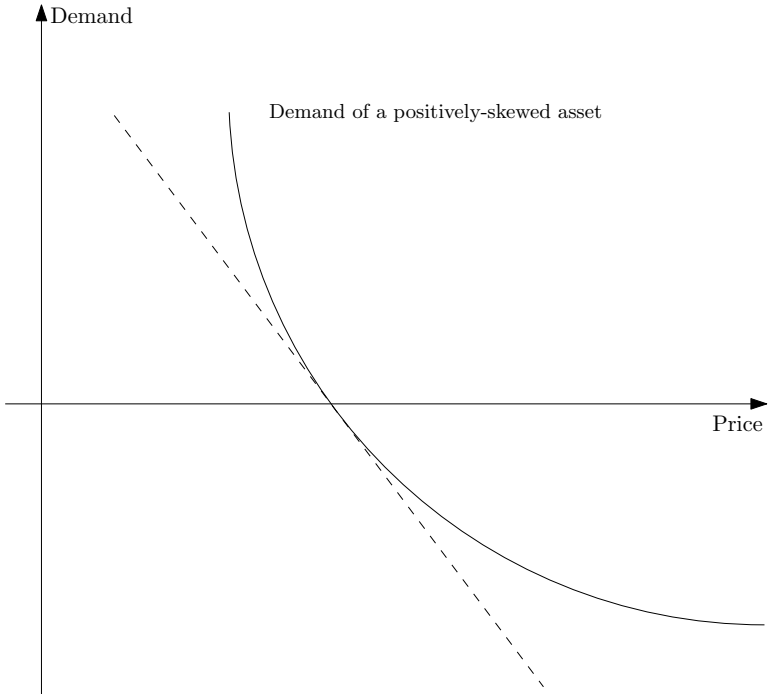


Figure 1: The Demand Schedule of a Positively Skewed Asset

The dashed line represents the linear demand schedule that obtains under CARA utility and normally distributed payoffs. The solid curve represents the demand schedule when the risky asset is positively skewed. Both curves pass through $(p, x) = (\mu, 0)$ and are tangent at that point, but the solid curve is convex, lying weakly above the dashed line. This is consistent with Lemma 1: positive skewness implies positive curvature of the demand schedule at $p = \mu$.

2.3 Extracting Implied Volatility from the Order Book

Lemma 1 establishes that the slope of the demand schedule encodes both the investor's risk tolerance and the payoff variance. To extract an implied volatility measure, we must

[1] Strictly speaking, it is impossible to increase skewness without affecting higher moments. However, as the proof of Lemma 1 shows, the curvature of the demand schedule at $p = \mu$ depends only on the first three moments, so the graphical argument is exact in the neighborhood of $p = \mu$.

separate the two.

2.3.1 Measuring the Slope

The slope $x'(\mu)$ can be estimated directly from the order book. With only the best bid and ask (NBBO), a simple difference quotient gives

$$\hat{x}'(\mu) \approx -\frac{bv_1 + av_1}{ap_1 - bp_1}. \tag{7}$$

Using deeper levels of the book provides a more robust estimate:

$$\hat{x}'(\mu) \approx -\frac{\sum_{i=1}^{K}(bv_i + av_i)}{ap_K - bp_K}, \tag{8}$$

where K is the number of levels used (e.g., $K = 4$ in our empirical implementation).

2.3.2 From Dollar Returns to Rate-of-Return Volatility

The payoff variance σ^2 in equation (4) refers to dollar payoffs $\tilde{\theta} - p$. Empirically, it is more natural to work with rate-of-return volatility. Let σ_r denote the standard deviation of the rate of return $\tilde{\theta}/p - 1$, so that $\sigma \approx \mu\,\sigma_r$. Substituting into equation (4) and rearranging yields

$$x'(\mu)\,\mu^2 = \frac{u'(w_0)}{u''(w_0)}\frac{1}{\sigma_r^2}. \tag{9}$$

The left-hand side—the demand slope scaled by the squared mid-quote—is observable. The right-hand side is the ratio of risk tolerance to return variance.

2.3.3 Identifying Volatility

Equation (9) contains two unknowns: the risk tolerance $u'(w_0)/u''(w_0)$ and the return variance σ_r^2. We identify the risk tolerance parameter by assuming that it is approximately

constant over a rolling estimation window of T trading days. Under this assumption, equation (9) implies a linear relationship between the observable $x'(\mu)\,\mu^2$ and the inverse realized return variance $1/\sigma_r^2$. This motivates a two-step procedure:

1. **Calibrate risk tolerance.** Over the past T trading days, regress the daily values of $\hat{x}'(\mu)\,\mu^2$ on $1/\hat{\sigma}_r^2$, where $\hat{\sigma}_r^2$ is realized return variance measured from intraday data. The estimated coefficient $\hat{b}$ is a proxy for $u'(w_0)/u''(w_0)$.

2. **Compute implied volatility.** On the current day, the implied return variance is

$$\hat{\sigma}_r^2 = \frac{\hat{b}}{\hat{x}'(\mu)\,\mu^2}, \tag{10}$$

   and the implied return volatility is $\hat{\sigma}_r = \sqrt{\hat{b}/(\hat{x}'(\mu)\,\mu^2)}$.

The key identifying assumption—that risk tolerance is stable over the estimation window—is analogous to the assumption underlying rolling-window realized volatility estimators (that the data-generating process is locally stationary). The advantage is that the resulting volatility measure is forward-looking: it reflects the beliefs and risk assessments embedded in today's order book, not yesterday's price path.

2.3.4 Comparison with Alternative Approaches

Several papers in the market microstructure and financial mathematics literatures have proposed methods for extracting volatility-related quantities from the limit order book. These approaches typically model the order book as a stochastic process and derive volatility from the dynamics of order flow, the shape of the book, or the rate of order arrivals. While valuable, these methods are primarily statistical: they estimate volatility from observable patterns in the data without imposing structure from economic theory.

Our approach differs in that it derives the relationship between the order book and volatility from a first-principles model of investor behavior. The slope of the demand schedule is not merely correlated with volatility; it is a direct consequence of the investor's optimization problem, with the functional form dictated by the utility function and the payoff distribution. This theoretical grounding has two practical advantages. First, it provides a clear decomposition of the demand slope into preference and distributional components, enabling us to extract volatility after controlling for risk tolerance. Second, it extends naturally to higher moments (skewness) and preference parameters (prudence), which purely statistical approaches do not accommodate.

2.4 Extracting Implied Skewness from the Order Book

The curvature of the demand schedule encodes payoff skewness (equation 6). To extract implied skewness, we proceed analogously to the volatility case.

2.4.1 Measuring the Curvature

The second derivative $x''(\mu)$ can be estimated from the order book by fitting a quadratic to the demand schedule. Using K levels on each side, we regress the cumulative demand x_i on $(p_i - \hat{\mu})$ and $(p_i - \hat{\mu})^2$, where $\hat{\mu} = (bp_1 + ap_1)/2$ is the mid-quote and p_i runs over all $2K$ price levels. The coefficient on the quadratic term provides an estimate $\hat{x}''(\mu)$.

Alternatively, with four points on the demand schedule (two bid levels and two ask levels), one can construct a finite-difference approximation to the second derivative. In practice, the regression-based approach is more robust because it averages over all available levels.

2.4.2 Identifying Skewness

From equation (6), the curvature is

$$x''(\mu) = \underbrace{\frac{u'''(w_0)}{-u''(w_0)}\left(\frac{u'(w_0)}{-u''(w_0)}\right)^2}_{\equiv \eta}\frac{s}{\sigma^3},$$

where η is a preference parameter (the product of absolute prudence and squared risk tolerance). The implied volatility step already provides an estimate of σ, and the same rolling-window logic identifies η: over the past T days, regress $\hat{x}''(\mu)$ on $\hat{s}/\hat{\sigma}^3$ (using realized skewness and volatility), and the estimated coefficient $\hat{\eta}$ proxies for the preference component. The implied skewness on the current day is then

$$\hat{s} = \frac{\hat{x}''(\mu)\,\hat{\sigma}^3}{\hat{\eta}}. \tag{11}$$

Order-book-implied skewness complements option-implied skewness measures in two important respects. First, it is available for all stocks with electronic limit order books, not just those with actively traded options. Second, because it is measured at the market open, it provides a high-frequency, forward-looking skewness signal that can be compared to subsequent realized skewness within the same trading day.

2.5 Implied Prudence

An additional output of the estimation procedure is a measure of investor preferences. Recall that the slope of the demand schedule identifies risk tolerance $-u'(w_0)/u''(w_0)$ (from equation 4), and the curvature identifies the composite parameter $\eta = [u'''(w_0)/(-u''(w_0))]\cdot[-u'(w_0)/u''(w_0)]^2$ (from equation 6). Dividing $\hat{\eta}$ by the square of estimated risk tolerance

yields an estimate of *absolute prudence*:

$$\widehat{\text{Prudence}} = \frac{\hat{\eta}}{\left(-\hat{b}\right)^2} \approx \frac{u'''(w_0)}{-u''(w_0)}, \tag{12}$$

as defined by Kimball (1990). Prudence governs precautionary saving behavior and plays a central role in models of consumption under uncertainty. To our knowledge, our approach provides the first method for estimating prudence at the individual stock level and at daily frequency, opening a new window into investor preferences that has been largely inaccessible to empirical researchers.

2.6 Extension to All Higher Moments

While our empirical implementation focuses on the first two nontrivial moments—variance (from the slope) and skewness (from the curvature)—the theoretical framework extends naturally to the entire sequence of payoff moments. The demand schedule $x(p)$ is a smooth function of p near μ (given sufficient smoothness of u), and each successive derivative of x at $p = \mu$ encodes one additional central moment of the payoff distribution. The following lemma makes this precise.

Lemma 2 (General Moment-Derivative Correspondence)**.** *Suppose u is $(n+1)$-times continuously differentiable for some integer $n \geq 1$. Denote the central moments by $m_k = E[(\tilde{\theta} - \mu)^k]$. Then $x^{(n)}(\mu)$ is determined by the utility derivatives $\{u^{(j)}(w_0)\}_{j=1}^{n+1}$ and the central moments $\{m_k\}_{k=2}^{n+1}$. Moreover, the $(n+1)$-th central moment enters linearly:*

$$x^{(n)}(\mu) = -\frac{u^{(n+1)}(w_0)\left(x'(\mu)\right)^n}{u''(w_0)\,\sigma^2}\, m_{n+1} \;+\; \Psi_n, \tag{13}$$

where Ψ_n depends only on the lower-order moments $\{m_k\}_{k=2}^{n}$ and the lower-order utility derivatives $\{u^{(j)}(w_0)\}_{j=1}^{n}$.

The proof is in Appendix B. The first three cases of Lemma 2 are:

n	Derivative	New moment	New utility derivative
1	$x'(\mu)$	$m_2 = \sigma^2$ (variance)	u''
2	$x''(\mu)$	m_3 (skewness $\times\ \sigma^3$)	u'''
3	$x'''(\mu)$	m_4 (kurtosis $\times\ \sigma^4$)	$u^{(4)}$

Remark 1 (Recursive identification). Lemma 2 implies that the moments of the payoff distribution can be recovered one at a time from the derivatives of the demand schedule. At each stage n, the n-th derivative $x^{(n)}(\mu)$ introduces exactly one new unknown—the $(n+1)$-th central moment m_{n+1}—provided that all lower moments and preference parameters have already been identified. Because $x^{(n)}(\mu)$ depends *linearly* on m_{n+1}, the inversion is straightforward: one can solve equation (13) for m_{n+1} given estimates of $x^{(n)}(\mu)$, $x'(\mu)$, and the preference parameter $u^{(n+1)}(w_0)/(u''(w_0)\sigma^2)$.

Remark 2 (CARA–Normal benchmark). Under CARA utility $u(w) = -\exp(-\gamma w)$ and normally distributed payoffs, the demand schedule is exactly linear (equation 5), so $x^{(n)}(\mu) = 0$ for all $n \geq 2$. This is consistent with Lemma 2: normal distributions have $m_k = 0$ for all odd $k \geq 3$, and the even moments m_k for $k \geq 4$ are determined by σ^2, so the CARA utility derivatives and the vanishing odd moments together force all curvature terms to zero.

In practice, the feasibility of estimating higher-order derivatives from discrete order book data diminishes rapidly. Estimating $x'(\mu)$ requires only two points (e.g., the best bid and ask); estimating $x''(\mu)$ requires at least three; and estimating $x^{(n)}(\mu)$ requires at least $n+1$ points. More importantly, numerical differentiation of noisy data amplifies estimation error at each successive order, so that the signal-to-noise ratio deteriorates quickly. With standard level-2 data providing four or five levels on each side, reliable estimation of derivatives beyond the second is challenging. Nevertheless, the theoretical result is important: it establishes that the order book is, in principle, a *complete* source of information about the

payoff distribution—every moment is encoded in the shape of the demand schedule, not just variance and skewness.

This completeness result distinguishes our framework from option-based methods, which also provide access to higher moments but require a continuum of strikes (Bakshi et al., 2003). In practice, the number of available strikes constrains option-based moment estimation just as the number of price levels constrains order-book-based estimation. The conceptual advantage of the order book approach is that it is grounded in a structural model of demand, which provides a clear economic interpretation for each derivative and a principled way to separate moments from preferences.

3 Empirical Analysis

3.1 Data

We use the Nasdaq ITCH (Inter-market Trading System) message-level data, which records every order submission, cancellation, and execution on the Nasdaq exchange. From this data, we reconstruct the full limit order book at any point in time. Our sample covers stocks traded on Nasdaq from 1999 to 2008. We exclude penny stocks (those with prices below $1) and stocks with low data coverage.

We focus on the limit order book at the market open, immediately following the opening auction, rather than during continuous trading. The opening auction aggregates overnight information and produces a single clearing price, after which the order book reflects the residual demand and supply schedules of investors who did not trade in the auction. This setting offers a clean snapshot of investor beliefs at the start of the trading day, before intraday trading activity has the opportunity to contaminate the book with transient liquidity or high-frequency strategies.

3.1.1 Variable Construction

Using the first four levels of the order book after the opening auction, we construct the slope of the demand schedule as follows:

$$\text{DemandSlopeN4} = -\frac{\frac{bv_1+av_1+bv_2+av_2+bv_3+av_3+bv_4+av_4}{\text{SharesOutstanding}}}{\frac{ap_4-bp_4}{\text{MidQuote}}}, \tag{14}$$

where bv_i and av_i denote the bid and ask volumes at level i, bp_i and ap_i denote the bid and ask prices at level i, and the mid-quote is defined as $(bp_1 + ap_1)/2$. This construction normalizes the demand slope by shares outstanding (to make volumes comparable across firms) and by the mid-quote (to express the price range in percentage terms). By construction, DemandSlopeN4 is negative, reflecting the downward-sloping nature of demand.

Our primary measure of limit-order-book implied volatility is

$$\text{LOBVol} = \frac{1}{1000}\frac{1}{\sqrt{-\text{DemandSlopeN4}}}. \tag{15}$$

The simple theory in Section 2 suggests that LOBVol $\propto \sqrt{\gamma}\sigma$, where γ is the representative investor's risk aversion and σ is the standard deviation of the asset's payoff. Thus, LOBVol captures a combination of payoff uncertainty and investor risk preferences. The scaling factor $\frac{1}{1000}$ is chosen to make the units comparable to conventional volatility measures.

We measure realized volatility using the standard deviation of 5-minute returns within the trading day (from open to close), denoted StdevRQ. Lagged realized volatility (StdevRQ$_{t-1}$) serves as the primary benchmark predictor. Additional control variables include: the bid-ask spread (SPRD), defined as $(ap_1 - bp_1)/\text{MidQuote}$; share turnover (TURN); the overnight return from the previous close to the current open; and lagged daily returns (RET or RETQ) over the preceding five trading days. When available, we also control for option-implied volatility (IVol), obtained from OptionMetrics.

3.2 Results

3.2.1 LOBVol Predicts Realized Volatility

Table 1 reports the results of panel regressions of daily realized volatility (StdevRQ_t) on LOBVol_t and control variables. All specifications include firm-year fixed effects and cluster standard errors at the firm-year level. Because volatility is persistent, OLS standard errors may understate uncertainty. Our firm-year clustering accounts for within-group serial correlation. As a robustness check, we also compute Newey–West standard errors with automatic bandwidth selection; the results are qualitatively unchanged.[2]

Column (1) shows that LOBVol measured at the market open is a significant predictor of same-day realized volatility, with a coefficient of 0.102 ($t = 4.80$), even after controlling for lagged realized volatility, the bid-ask spread, turnover, overnight returns, and lagged returns. Importantly, lagged realized volatility (StdevRQ_{t-1}) remains significant with a coefficient of 0.277 ($t = 11.03$), indicating that LOBVol captures incremental information beyond what is already embedded in the recent history of realized volatility.

Column (2) adds option-implied volatility (IVol_{t-1}) as an additional control. The sample shrinks to 341,661 observations because IVol is available only for the subset of stocks with actively traded options. Even in this restricted sample, LOBVol retains a highly significant coefficient of 0.032 ($t = 13.39$). This result is important because it demonstrates that the order book contains forward-looking information about volatility that is not subsumed by option-implied volatility. As discussed in Section 2.3.4, our theory-motivated LOBVol differs from purely statistical or financial-mathematics-based approaches to extracting volatility from the order book. The theoretical grounding ensures that the demand slope is not merely correlated with volatility but is a structural consequence of the investor's

[2]The persistence of realized volatility raises the concern that standard inference may overstate significance. Firm-year clustering allows for arbitrary within-cluster correlation and is our baseline approach. Driscoll–Kraay standard errors, which are robust to both cross-sectional and temporal dependence, yield similar conclusions.

optimization problem, providing a clean decomposition into preference and distributional components.

Table 1: LOBVol Predicts Realized Volatility

Dependent Variable	**StdevRQ$_t$**	
	(1)	(2)
LOBVol$_t$	0.102***	0.032***
	[4.80]	[13.39]
StdevRQ$_{t-1}$	0.277***	0.431***
	[11.03]	[72.56]
IVol$_{t-1}$		0.011***
		[9.57]
SPRD$_{t-1}$	0.137***	0.060
	[18.50]	[1.24]
TURN$_{t-1}$	0.032***	−0.008
	[3.78]	[−1.48]
Overnight Return$_t$	0.007	−0.061***
	[0.73]	[−3.05]
Overnight Return Squared$_t$	−0.000	0.207***
	[−0.69]	[2.98]
RETQ$_{t-1}$	−0.008*	−0.012***
	[−1.69]	[−14.92]
RETQ$_{t-2}$	−0.001***	−0.001
	[−2.60]	[−0.37]
RETQ$_{t-3}$	−0.001*	−0.004***
	[−1.80]	[−8.33]
RETQ$_{t-4}$	−0.002	0.001
	[−0.69]	[0.97]
RETQ$_{t-5}$	−0.003	−0.001
	[−0.94]	[−0.63]
Firm-Year FE	Yes	Yes
Firm-Year Cluster	Yes	Yes
Adj. R^2	0.173	0.178
Number of Observations	716,744	341,661

This table reports panel regressions of daily realized volatility (StdevRQ$_t$, the standard deviation of intraday 5-minute returns) on LOBVol$_t$ measured at the market open and control variables. LOBVol is defined in equation (15). Column (2) adds option-implied volatility (IVol$_{t-1}$). All specifications include firm-year fixed effects with standard errors clustered at the firm-year level. t-statistics are reported in brackets. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

3.2.2 Macro Announcement Days

The baseline results establish that LOBVol predicts realized volatility in the full sample. But the more subtle question is *why*. If LOBVol merely proxies for slow-moving stock characteristics—such as firm size, industry, or average liquidity—then its predictive power need not reflect forward-looking information. Macroeconomic announcement days provide a powerful setting to discriminate between this alternative and our preferred interpretation.

The logic is as follows. Pre-market macroeconomic announcements—such as non-farm payrolls, CPI releases, GDP reports, or FOMC decisions—inject new information into the market before the opening bell. Investors who observe these announcements update their beliefs about fundamentals and risk before placing orders. The resulting order book at the open therefore reflects updated expectations. If LOBVol captures this forward-looking content, its predictive power should *increase* on announcement days relative to non-announcement days. Conversely, backward-looking measures such as lagged realized volatility should become *less* informative on these days, because the new information renders yesterday's price path a less reliable guide to today's volatility.

This asymmetric prediction—LOBVol gains informativeness when backward-looking measures lose it—is the key empirical test. It directly echoes the logic of Savor and Wilson (2014), who show that the risk-return tradeoff is concentrated on macroeconomic announcement days, and of Lucca and Moench (2015), who document that pre-FOMC announcement returns reflect the resolution of uncertainty. Our setting differs in that we examine volatility prediction rather than return premia, but the underlying mechanism is the same: announcement days are natural experiments that separate forward-looking from backward-looking information.

Table 2 reports the results. We interact LOBVol_t and StdevRQ_{t-1} with an indicator for pre-market macro news days (Macro News_t). The interaction of LOBVol with Macro

News is positive and highly significant (coefficient $= 0.040$, $t = 4.73$), indicating that the predictive power of LOBVol for realized volatility is amplified on macro announcement days. This is exactly the pattern our theory predicts: when investors have had time to process new information before the open, the order book is a richer source of forward-looking signals.

In contrast, the interaction of lagged realized volatility with Macro News is negative and significant (coefficient $= -0.028$, $t = -4.43$), suggesting that backward-looking volatility measures become less informative precisely when new information arrives. This result is intuitive: yesterday's volatility reflects yesterday's information environment, which is partially superseded by the macro announcement. The magnitudes of these interaction effects are economically meaningful—they are comparable to a one-standard-deviation move in DemandSlopeN4 and StdevRQ, respectively.

Taken together, the two interaction effects paint a clear picture. On non-announcement days, lagged realized volatility is the dominant predictor, and LOBVol provides incremental information. On announcement days, the balance shifts: LOBVol becomes more important and lagged realized volatility becomes less important. This crossover pattern is difficult to explain if LOBVol merely captures stock-level fixed characteristics or slow-moving liquidity conditions. It is, however, a natural consequence of a measure that reflects the real-time belief updates of market participants.

These findings also have practical implications for risk management. Standard volatility forecasting models—such as GARCH (Engle, 1982; Bollerslev, 1986) and HAR (Corsi, 2009)—are entirely backward-looking. Our results suggest that on the days when accurate volatility forecasts matter most—days with macroeconomic news, heightened uncertainty, or potential regime changes—these models are least reliable, precisely because the information environment has shifted. Supplementing backward-looking models with order-book-implied volatility can improve forecasts on these critical days.

Table 2: LOBVol Predictability on Pre-Market Macro Announcement Days

Dependent Variable	**StdevRQ$_t$**
LOBVol$_t$ $\times$ Macro News$_t$	0.040***
	[4.73]
LOBVol$_t$	0.101***
	[4.77]
StdevRQ$_{t-1}$ $\times$ Macro News$_t$	$-$0.028***
	[$-$4.43]
StdevRQ$_{t-1}$	0.281***
	[11.76]
Macro News$_t$	0.000
	[0.77]
SPRD$_{t-1}$	0.137***
	[18.41]
TURN$_{t-1}$	0.032***
	[3.78]
Overnight Return$_t$	0.007
	[0.71]
Overnight Return Squared$_t$	$-$0.000
	[$-$0.67]
RETQ$_{t-1}$	$-$0.008*
	[$-$1.68]
RETQ$_{t-2}$	$-$0.001***
	[$-$2.62]
RETQ$_{t-3}$	$-$0.001*
	[$-$1.79]
RETQ$_{t-4}$	$-$0.002
	[$-$0.68]
RETQ$_{t-5}$	$-$0.003
	[$-$0.95]
Firm-Year FE	Yes
Firm-Year Cluster	Yes
Adj. R^2	0.174
Number of Observations	716,744

This table reports panel regressions of daily realized volatility on LOBVol$_t$ interacted with an indicator for pre-market macroeconomic announcement days. All specifications include firm-year fixed effects with standard errors clustered at the firm-year level. t-statistics are reported in brackets. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

3.2.3 LOBVol and Stock Returns

Our theory implies that LOBVol captures a combination of perceived payoff risk (σ) and investor risk aversion (γ), both of which affect expected returns through standard asset pricing channels. A long literature in financial economics—from the CAPM to the intertemporal models of Merton (1973) and French et al. (1987)—predicts that expected returns compensate investors for bearing risk. At the daily frequency, stocks that appear riskier at the open should command a risk premium, but they may also experience price declines as risk-averse investors reduce their exposure, generating a negative contemporaneous relationship between LOBVol and returns.

Table 3 tests this prediction by regressing open-to-close price returns and close-to-close mid-quote returns on LOBVol_t and the standard set of control variables. LOBVol_t enters with a negative and highly significant coefficient in both specifications: -0.003 ($t = -24.59$) for open-to-close returns and -0.004 ($t = -10.96$) for close-to-close returns. These results indicate that stocks with higher order-book-implied volatility at the open tend to experience lower returns during the trading day.

The negative sign is consistent with several interpretations. First, higher LOBVol may signal elevated uncertainty that triggers a contemporaneous drop in the stock price as risk-averse investors demand a larger discount. This is analogous to the well-documented negative relationship between changes in the VIX and stock market returns (Whaley, 2009). Second, to the extent that LOBVol captures an increase in risk aversion rather than (or in addition to) an increase in fundamental risk, the price decline may reflect a shift in the stochastic discount factor, consistent with models in which time-varying risk aversion drives return dynamics (Campbell and Cochrane, 1999). The fact that the coefficient is significant even after controlling for lagged realized volatility and the bid-ask spread suggests that LOBVol contains return-relevant information beyond what standard microstruc-

ture variables capture.

Table 3: LOBVol and Stock Returns

Dependent Variable	**Open-to-Close Price Return**	**Close-to-Close Mid-Q Return**
LOBVol_t	-0.003^{***}	-0.004^{***}
	$[-24.59]$	$[-10.96]$
StdevRQ_{t-1}	2.474^{***}	1.799
	[2.71]	[1.61]
SPRD_{t-1}	3.611^{***}	4.452^{***}
	[13.87]	[7.72]
TURN_{t-1}	0.773	0.706
	[0.76]	[0.48]
$\text{Overnight Return}_t$	-33.855^{***}	51.104^{***}
	$[-8.36]$	[5.81]
$\text{Overnight Return Squared}_t$	0.867^{***}	1.208^{***}
	[8.33]	[5.37]
RET_{t-1}	-4.633^{***}	-2.419
	$[-5.85]$	$[-1.15]$
RET_{t-2}	-1.903^{***}	-0.276
	$[-24.39]$	$[-1.31]$
RET_{t-3}	-1.235^{***}	-0.195
	$[-6.96]$	$[-1.13]$
RET_{t-4}	-1.490^{***}	-0.721
	$[-9.18]$	$[-1.59]$
RET_{t-5}	-1.226^{***}	-0.468
	$[-4.25]$	$[-1.26]$
Firm-Year FE	Yes	Yes
Firm-Year Cluster	Yes	Yes
Adj. R^2	0.049	0.062
Number of Observations	718,309	718,086

This table reports panel regressions of daily stock returns on LOBVol_t measured at the market open. Column (1) uses the open-to-close price return and Column (2) uses the close-to-close mid-quote return as the dependent variable. All specifications include firm-year fixed effects with standard errors clustered at the firm-year level. t-statistics are reported in brackets. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

4 Conclusion

This paper develops a theoretical framework that links the limit order book to the moments of the asset payoff distribution. We show that the n-th derivative of the demand schedule at the expected payoff encodes the $(n+1)$-th central moment of the payoff distribution, combined with preference parameters. In principle, the entire sequence of payoff moments can be recovered from the shape of the demand schedule, making the order book a complete source of forward-looking distributional information.

We focus our empirical implementation on variance and construct a limit-order-book implied volatility measure (LOBVol) using Nasdaq ITCH data from 1999 to 2008. Our main findings are as follows. First, LOBVol measured at the market open significantly predicts same-day realized volatility, even after controlling for lagged realized volatility, option-implied volatility, and other standard predictors. Second, this predictive power is amplified on days with pre-market macroeconomic announcements, while lagged realized volatility becomes less informative. This asymmetric pattern is the strongest evidence that our measure captures genuinely forward-looking content: on precisely the days when new information renders backward-looking measures less reliable, the order book becomes more informative. Third, LOBVol is negatively associated with same-day stock returns in the cross-section, consistent with a risk-return tradeoff at the daily frequency.

Our approach has several advantages relative to existing methods. Unlike realized volatility and GARCH-type models (Engle, 1982; Bollerslev, 1986), LOBVol is forward-looking. Unlike option-implied volatility, it is available for virtually all electronically traded stocks, regardless of whether the firm has listed options or whether those options trade actively. Unlike purely statistical approaches to extracting information from the order book, our method is grounded in a structural model of investor demand, providing clear economic interpretations and a principled way to separate distributional moments from

preference parameters. Because every electronically traded security has an order book, our method provides a universal tool for measuring risk perceptions, applicable to equities, fixed income, foreign exchange, and commodities alike.

The macro announcement results deserve particular emphasis. Standard volatility forecasting models are entirely backward-looking: they extrapolate from past volatility patterns. Our results show that on the days when accurate forecasts matter most—days with macroeconomic news, heightened uncertainty, or potential regime changes—backward-looking models are least reliable. Supplementing them with LOBVol can improve forecasts precisely when improvement is most valuable. This has direct applications for risk management, option pricing, and portfolio allocation around scheduled macroeconomic events.

Several avenues for future work remain. First, extending the empirical analysis to order-book-implied skewness—as developed in Section 2.4—is a natural next step. The theory predicts that the curvature of the demand schedule reveals payoff skewness, and testing whether implied skewness predicts realized skewness would provide direct validation of the higher-order terms in Lemma 1. Moreover, order-book-implied skewness may help reconcile conflicting findings in the literature on the cross-sectional relationship between expected skewness and stock returns (Harvey and Siddique, 2000; Conrad et al., 2013; Amaya et al., 2015). While some studies document a negative skewness premium, others find the opposite; a forward-looking, universally available skewness measure could shed light on whether these discrepancies reflect measurement limitations of existing approaches.

Second, the theoretical result that all higher moments are in principle recoverable from the demand schedule (Lemma 2) opens a path toward order-book-implied kurtosis and tail-risk measures.

Third, the estimation of absolute prudence from the order book—as outlined in Section 2.5—opens the door to measuring investor preferences at high frequency and at the individual stock level. Prudence governs precautionary behavior and plays a central role in

consumption-based asset pricing (Kimball, 1990), but has been difficult to estimate empirically. Our framework provides a new pathway to this fundamental preference parameter.

Fourth, applying the framework to other asset classes—sovereign bonds, currencies, and commodity futures—would test the generality of the approach and potentially provide new risk indicators for markets where option-based measures are unavailable or unreliable.

References

Amaya, D., Christoffersen, P., Jacobs, K., and Vasquez, A. (2015). Does realized skewness predict the cross-section of equity returns? *Journal of Financial Economics*, 118(1):135–167.

Andersen, T. G. and Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4):885–905.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71(2):579–625.

Bakshi, G., Kapadia, N., and Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1):101–143.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2008). Designing realized kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica*, 76(6):1481–1536.

Barndorff-Nielsen, O. E. and Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B*, 64(2):253–280.

Barndorff-Nielsen, O. E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, 2(1):1–37.

Biais, B., Hillion, P., and Spatt, C. (1995). An empirical analysis of the limit order book and the order flow in the Paris Bourse. *Journal of Finance*, 50(5):1655–1689.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.

Bollerslev, T., Tauchen, G., and Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies*, 22(11):4463–4492.

Britten-Jones, M. and Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. *Journal of Finance*, 55(2):839–866.

Campbell, J. Y. and Cochrane, J. H. (1999). Force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251.

Cao, C., Hansch, O., and Wang, X. (2009). The information content of an open limit-order book. *Journal of Futures Markets*, 29(1):16–41.

Conrad, J., Dittmar, R. F., and Ghysels, E. (2013). Ex ante skewness and expected stock returns. *Journal of Finance*, 68(1):85–124.

Cont, R., Kukanov, A., and Stoikov, S. (2014). The price impact of order book events. *Journal of Financial Econometrics*, 12(1):47–88.

Cont, R. and Mueller, M. S. (2023). Limit order book models and market microstructure. *Working Paper*.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2):174–196.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4):987–1007.

French, K. R., Schwert, G. W., and Stambaugh, R. F. (1987). Expected stock returns and volatility. *Journal of Financial Economics*, 19(1):3–29.

Glosten, L. R. (1994). Is the electronic open limit order book inevitable? *Journal of Finance*, 49(4):1127–1161.

Harvey, C. R. and Siddique, A. (2000). Conditional skewness in asset pricing tests. *Journal of Finance*, 55(3):1263–1295.

Horst, U. and Xu, W. (2019). A second-order stock market model. *Working Paper*.

Jiang, G. J. and Tian, Y. S. (2005). The model-free implied volatility and its information content. *Review of Financial Studies*, 18(4):1305–1342.

Kimball, M. S. (1990). Precautionary saving in the small and in the large. *Econometrica*, 58(1):53–73.

Kraus, A. and Litzenberger, R. H. (1976). Skewness preference and the valuation of risk assets. *Journal of Finance*, 31(4):1085–1100.

Lucca, D. O. and Moench, E. (2015). The pre-FOMC announcement drift. *Journal of Finance*, 70(1):329–371.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica*, 41(5):867–887.

Patton, A. J. and Sheppard, K. (2015). Good volatility, bad volatility: Signed jumps and the persistence of volatility. *Review of Economics and Statistics*, 97(3):683–697.

Savor, P. and Wilson, M. (2014). Asset pricing: A tale of two days. *Journal of Financial Economics*, 113(2):171–201.

Tauchen, G. and Zhou, H. (2011). Realized jumps on financial markets and predicting credit spreads. *Journal of Econometrics*, 160(1):102–118.

Whaley, R. E. (2009). Understanding the VIX. *Journal of Portfolio Management*, 35(3):98–105.

Zhang, L., Mykland, P. A., and 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.

Appendices

A Proof of Lemma 1

Proof. Since $x(\mu) = 0$ (as established in the text), the Taylor expansion of $x(p)$ around $p = \mu$ takes the form

$$x(p) = x'(\mu)(p-\mu) + \tfrac{1}{2}x''(\mu)(p-\mu)^2 + o\big((p-\mu)^2\big)\,.$$

We compute $x'(\mu)$ and $x''(\mu)$ by implicitly differentiating the first-order condition.

Setup. Define the function

$$F(p,x) \;\equiv\; E\Big[u'\Big(w_0 + x(\tilde{\theta}-p)\Big)\,(\tilde{\theta}-p)\Big]\,.$$

The first-order condition (2) requires $F(p, x(p)) = 0$ for all p in a neighborhood of μ. We write $\epsilon \equiv \tilde{\theta} - p$ for the excess payoff. At the point $(p,x) = (\mu, 0)$, we have $\epsilon = \tilde{\theta} - \mu$ and the argument of u' reduces to w_0.

Partial derivatives of F at $(\mu, 0)$. We will need the following partial derivatives, all evaluated at $(p,x) = (\mu, 0)$.

First-order partials.

$$F_x = E\big[u''(w_0 + x\epsilon)\,\epsilon^2\big] \;\xrightarrow{(\mu,0)}\; u''(w_0)\,\sigma^2, \tag{16}$$

$$F_p = E[-x\,u''(w_0 + x\epsilon)\,\epsilon - u'(w_0 + x\epsilon)] \;\xrightarrow{(\mu,0)}\; -u'(w_0). \tag{17}$$

In (16), the derivative with respect to x brings down ϵ from the chain rule, yielding $u'' \cdot \epsilon \cdot \epsilon =$

$u''\epsilon^2$. In (17), the derivative with respect to p acts on both $u'(\cdot)$ (via the $-x\epsilon$ in its argument) and on $\epsilon = \tilde{\theta} - p$ (contributing -1). At $(\mu, 0)$, the first term vanishes because $x = 0$.

Second-order partials.

$$F_{pp} = E\left[x^2 u'''(w_0 + x\epsilon)\,\epsilon + 2x\,u''(w_0 + x\epsilon)\right] \xrightarrow{(\mu,0)} 0, \tag{18}$$

$$F_{px} = E\left[-2\,u''(w_0 + x\epsilon)\,\epsilon - x\,u'''(w_0 + x\epsilon)\,\epsilon^2\right] \xrightarrow{(\mu,0)} -2\,u''(w_0)\,E[\tilde{\theta} - \mu] = 0, \tag{19}$$

$$F_{xx} = E\left[u'''(w_0 + x\epsilon)\,\epsilon^3\right] \xrightarrow{(\mu,0)} u'''(w_0)\,E\left[(\tilde{\theta} - \mu)^3\right] = u'''(w_0)\,s\,\sigma^3. \tag{20}$$

In (18), we differentiate F_p with respect to p; all terms contain a factor of x and vanish at $x = 0$. In (19), we differentiate F_p with respect to x; at $x = 0$, the surviving term is $-2u''(w_0)E[\tilde{\theta} - \mu] = 0$. In (20), we differentiate $F_x = E[u''(w_0 + x\epsilon)\epsilon^2]$ with respect to x, bringing down another factor of ϵ.

Step 1: Computing $x'(\mu)$. Differentiating $F(p, x(p)) = 0$ with respect to p gives

$$F_p + F_x\,x'(p) = 0. \tag{21}$$

Evaluating at $p = \mu$ and using (16)–(17):

$$-u'(w_0) + u''(w_0)\,\sigma^2 \cdot x'(\mu) = 0,$$

so

$$\boxed{x'(\mu) = \frac{u'(w_0)}{u''(w_0)\,\sigma^2}.} \tag{22}$$

Since $u' > 0$ and $u'' < 0$, the slope is negative: $x'(\mu) < 0$. This is the first-order term in the lemma.

Step 2: Computing $x''(\mu)$. Differentiating (21) again with respect to p gives

$$F_{pp} + 2F_{px}\, x'(p) + F_{xx}\,(x'(p))^2 + F_x\, x''(p) = 0. \tag{23}$$

Evaluating at $p = \mu$ and substituting (18)–(20):

$$0 + 0 + u'''(w_0)\, s\, \sigma^3 \cdot \left(\frac{u'(w_0)}{u''(w_0)\,\sigma^2}\right)^2 + u''(w_0)\,\sigma^2 \cdot x''(\mu) = 0.$$

Solving for $x''(\mu)$:

$$\boxed{x''(\mu) = -\frac{u'''(w_0)\,(u'(w_0))^2\, s}{(u''(w_0))^3\,\sigma^3}.} \tag{24}$$

To verify the sign, note that $u'' > 0$ raised to an odd power is negative (since $u'' < 0$), so $(u'')^3 < 0$. Under NARA, $u''' > 0$. Therefore $\operatorname{sgn}(x''(\mu)) = \operatorname{sgn}(s)$.

Step 3: Assembling the Taylor expansion. Substituting (22) and (24) into the Taylor expansion $x(p) = x'(\mu)(p-\mu) + \frac{1}{2}x''(\mu)(p-\mu)^2 + o((p-\mu)^2)$:

$$x(p) = \frac{u'(w_0)}{u''(w_0)}\frac{1}{\sigma^2}(p-\mu) \;-\; \frac{1}{2}\frac{u'''(w_0)\,(u'(w_0))^2}{(u''(w_0))^3\,\sigma^3}\, s\,(p-\mu)^2 + o\big((p-\mu)^2\big)\,.$$

Rewriting the second term by factoring:

$$\frac{u'''(w_0)\,(u'(w_0))^2}{(u''(w_0))^3} = \frac{u'''(w_0)}{u''(w_0)} \cdot \left(\frac{u'(w_0)}{u''(w_0)}\right)^2,$$

we obtain

$$x(p) = \frac{u'(w_0)}{u''(w_0)}\frac{1}{\sigma^2}(p-\mu) - \frac{1}{2}\frac{u'''(w_0)}{u''(w_0)}\left(\frac{u'(w_0)}{u''(w_0)}\right)^2 \frac{s}{\sigma^3}(p-\mu)^2 + o\big((p-\mu)^2\big)\,,$$

which is equation (3).

Validity of implicit differentiation. The implicit function theorem guarantees that $x(p)$ is a well-defined, smooth function of p in a neighborhood of μ, provided $F_x(\mu, 0) \neq 0$. From (16), $F_x(\mu, 0) = u''(w_0)\sigma^2 \neq 0$ because u is strictly concave ($u'' < 0$) and $\sigma^2 > 0$ (the payoff is risky). Moreover, since u is thrice continuously differentiable by Assumption 1, F is twice continuously differentiable in both arguments, so $x(p)$ is twice continuously differentiable and the Taylor expansion with a second-order remainder is justified. □

B Proof of Lemma 2

Proof. We use the same implicit differentiation framework as in Appendix A. Recall that $F(p, x) = E[u'(w_0 + x(\tilde{\theta} - p))(\tilde{\theta} - p)]$ and $F(p, x(p)) = 0$ for all p near μ. We write $\epsilon = \tilde{\theta} - p$ and $m_k = E[(\tilde{\theta} - \mu)^k]$ for the central moments. The proof proceeds in three steps: we first derive a closed-form expression for all mixed partials of F at $(\mu, 0)$, then use repeated implicit differentiation to establish the structure of $x^{(n)}(\mu)$ by induction.

Step 1: Mixed partials of F at $(\mu, 0)$. We claim that for all integers $a \geq 0$ and $b \geq 0$,

$$\left.\frac{\partial^{a+b} F}{\partial p^a \, \partial x^b}\right|_{(\mu,0)} = \begin{cases} (-1)^a \, \dfrac{(b+1)!}{(b+1-a)!} \, u^{(b+1)}(w_0) \, m_{b+1-a} & \text{if } a \leq b+1, \\ 0 & \text{if } a > b+1. \end{cases} \tag{25}$$

Proof of (25). Define $g(p, x, \tilde{\theta}) = u'(w_0 + x\epsilon)\,\epsilon$, so that $F = E[g]$ and $\frac{\partial^{a+b} F}{\partial p^a \partial x^b} = E\left[\frac{\partial^{a+b} g}{\partial p^a \partial x^b}\right]$.

Since x enters g only through the product $x\epsilon$ in the argument of u', successive x-derivatives give

$$\frac{\partial^b g}{\partial x^b} = u^{(b+1)}(w_0 + x\epsilon)\, \epsilon^{b+1}. \tag{26}$$

At $x = 0$, the argument of $u^{(b+1)}$ reduces to w_0, which does not depend on p. The only remaining p-dependence is through $\epsilon^{b+1} = (\tilde{\theta} - p)^{b+1}$. Therefore, at $x = 0$:

$$\frac{\partial^a}{\partial p^a}\left[u^{(b+1)}(w_0)\ \epsilon^{b+1}\right] = u^{(b+1)}(w_0)\ \frac{\partial^a}{\partial p^a}(\tilde{\theta}-p)^{b+1} = (-1)^a\ \frac{(b+1)!}{(b+1-a)!}\ u^{(b+1)}(w_0)\ \epsilon^{b+1-a}, \tag{27}$$

provided $a \leq b + 1$. If $a > b + 1$, the a-th derivative of the polynomial $(\tilde{\theta} - p)^{b+1}$ in p is zero.

A subtlety arises because we must verify that terms involving x—which vanish pointwise at $x = 0$ in (26)—do not contribute to the mixed partial after p-differentiation. To see this, note that $\frac{\partial}{\partial p}[u^{(b+1)}(w_0 + x\epsilon)\epsilon^{b+1}] = u^{(b+2)}(w_0 + x\epsilon)(-x)\epsilon^{b+1} + u^{(b+1)}(w_0 + x\epsilon)(b + 1)\epsilon^b(-1)$. At $x = 0$, the first term vanishes. Repeating this argument, each further p-derivative generates a term with a factor of x from the chain rule acting on $u^{(k)}(w_0 + x\epsilon)$, plus a term from differentiating ϵ^m. At $x = 0$, only the latter survives. By induction on a, the chain-rule terms contribute zero at $x = 0$ for every a, confirming (27).

Evaluating at $p = \mu$ (so $\epsilon = \tilde{\theta} - \mu$) and taking expectations yields (25). □

As a consistency check, we recover the partials from Appendix A:

$$\begin{aligned}
F_x\ (a = 0, b = 1): &\quad u''(w_0)\, m_2 = u''(w_0)\, \sigma^2, \\
F_p\ (a = 1, b = 0): &\quad (-1)^1 \cdot 1! \cdot u'(w_0) \cdot m_0 = -u'(w_0), \\
F_{xx}\ (a = 0, b = 2): &\quad u'''(w_0)\, m_3, \\
F_{px}\ (a = 1, b = 1): &\quad (-1)^1 \cdot 2! \cdot u''(w_0) \cdot m_1 = -2u''(w_0) \cdot 0 = 0, \\
F_{pp}\ (a = 2, b = 0): &\quad a = 2 > b + 1 = 1\ \Rightarrow\ 0.
\end{aligned}$$

Step 2: Structure of the n-th total derivative. We define $H(p) = F(p, x(p))$, so $H(p) = 0$ for all p near μ, hence $H^{(n)}(\mu) = 0$ for every $n \geq 1$. By the general for-

mula for the n-th total derivative of $F(p, x(p))$ (a consequence of the multivariate chain rule), $H^{(n)}$ can be written as

$$H^{(n)}(\mu) = F_x(\mu, 0)\; x^{(n)}(\mu) + R_n = 0, \tag{28}$$

where R_n collects all terms not involving $x^{(n)}$. Specifically, R_n is a sum of terms, each of which is a product of a mixed partial $\frac{\partial^{a+b} F}{\partial p^a \partial x^b}\big|_{(\mu,0)}$ (with $a + b \leq n$, and $b < n$ when $a = 0$) multiplied by products of lower-order derivatives $x^{(j)}(\mu)$ with $1 \leq j \leq n - 1$. The precise combinatorial coefficients follow from the Faà di Bruno formula, but their exact form is not needed for our argument.

Since $F_x(\mu, 0) = u''(w_0)\sigma^2 \neq 0$ (as shown in Appendix A), we can solve for

$$x^{(n)}(\mu) = -\frac{R_n}{F_x(\mu, 0)} = -\frac{R_n}{u''(w_0)\,\sigma^2}. \tag{29}$$

Step 3: Induction on n. We prove the following claim by induction on $n \geq 1$:

Claim. *$x^{(n)}(\mu)$ depends only on $\{u^{(j)}(w_0)\}_{j=1}^{n+1}$ and $\{m_k\}_{k=2}^{n+1}$, and the $(n+1)$-th moment m_{n+1} enters only through the single term $-u^{(n+1)}(w_0)(x'(\mu))^n\, m_{n+1} \,/\, (u''(w_0)\sigma^2)$.*

Base case ($n = 1$). From Step 1 of Appendix A: $x'(\mu) = u'(w_0)/(u''(w_0)\sigma^2)$. This depends on u', u'' and $m_2 = \sigma^2$. The moment m_2 enters through $F_x = u''(w_0)m_2$, and the leading-order formula reads $x'(\mu) = -F_p/F_x = u'(w_0)/(u''(w_0)m_2)$. Since there is only one term and $\Psi_1 = 0$, the claim holds.

Inductive step. Assume the claim holds for all $k < n$. We analyze R_n in (28). Each term in R_n is a product of:

(i) a mixed partial $\frac{\partial^{a+b} F}{\partial p^a \partial x^b}\big|_{(\mu,0)}$ with $a + b \leq n$, and

(ii) a product of lower-order derivatives $\prod x^{(j_i)}(\mu)$ with each $j_i < n$.

From (25), factor (i) involves $u^{(b+1)}(w_0)$ and m_{b+1-a} (when $a \leq b+1$). The highest moment appearing in any single mixed partial is m_{b+1} (when $a = 0$). Among the terms in R_n with $a = 0$, the largest value of b is $b = n$, arising from the term $\frac{\partial^n F}{\partial x^n}\big|_{(\mu,0)} \cdot (x'(\mu))^n$. This term equals

$$u^{(n+1)}(w_0)\, m_{n+1} \cdot (x'(\mu))^n.$$

It is the *only* term in R_n that involves m_{n+1}, because:

- Any other mixed partial with $a = 0$ has $b \leq n-1$ (since $b = n$ can only pair with the monomial $(x')^n$ to match the total degree n, and this is the term identified above), giving at most m_n.
- Any mixed partial with $a \geq 1$ has $b+1-a \leq b \leq n-1$ (since $a \geq 1$ and $a+b \leq n$), giving at most m_{n-1}.

Similarly, $u^{(n+1)}(w_0)$ appears only in the same term (since all other terms have $b \leq n-1$, giving at most $u^{(n)}(w_0)$).

By the inductive hypothesis, each factor $x^{(j_i)}(\mu)$ in (ii) depends only on $\{u^{(\ell)}(w_0)\}_{\ell=1}^{j_i+1} \subseteq \{u^{(\ell)}(w_0)\}_{\ell=1}^{n}$ and $\{m_\ell\}_{\ell=2}^{j_i+1} \subseteq \{m_\ell\}_{\ell=2}^{n}$.

Combining (i) and (ii) and substituting into (29):

$$x^{(n)}(\mu) = -\frac{u^{(n+1)}(w_0)(x'(\mu))^n}{u''(w_0)\,\sigma^2}\, m_{n+1} \;-\; \frac{1}{u''(w_0)\sigma^2} \sum_{\text{other terms in } R_n},$$

where the sum involves only $\{u^{(j)}(w_0)\}_{j=1}^{n}$ and $\{m_k\}_{k=2}^{n}$. Setting Ψ_n equal to the second group completes the induction.

Smoothness. If u is $(n+1)$-times continuously differentiable, then F is n-times continuously differentiable in (p, x), so the implicit function theorem guarantees that $x(p)$ is n-times continuously differentiable near μ, and all steps above are justified. □