Liquidity Owls: A High-Frequency Estimation of Integrated Variance

Gabriel Sánchez Hernández

18. July 2026

Abstract

In this paper, I revisit the problem of measuring the integrated variance of a full trading day when high-frequency data exist only for the active trading session. Following Hansen and Lunde (2005), I combine the squared overnight return with a Newey-West-corrected realized variance of the open-to-close period, with weights chosen to minimize the mean-squared error among conditionally unbiased estimators. I apply the estimator to one-second best bid and offer quotations for the 30 components of the Dow Jones Industrial Average and three index ETFs over 2,034 trading days from May 2018 to June 2026. The qualitative structure of the original study survives: intraday measures dominate, yet the overnight return always retains informative weight. Two decades on, however, the overnight period accounts for roughly 30% of daily variance, up from 20%, and for the diversified ETFs the squared overnight return earns weights several times larger than for any individual equity. Formal specification tests, making use of the longer sample, reject a fixed overnight share.

Contents

1. Introduction

How should the volatility of a full trading day be measured when high-frequency prices exist only for the active trading session? Equivalently, how much weight should the squared overnight return receive when it is combined with a realized variance built from intraday returns? This is the research question of Hansen & Lunde (2005), whose framework I revisit and extend on a modern, high-resolution sample.

The question is relevant because volatility is a latent object at the center of asset pricing, portfolio choice, and risk management, and because the realized variance, the sum of squared intraday returns, estimates it well only over the interval it actually covers. The regular session of the NASDAQ and NYSE spans six and a half hours, so a measure built exclusively from intraday returns is silent about the substantial part of daily volatility that unfolds between the close and the subsequent open. The mismatch is not immaterial, and, as this article documents, the inactive period nowadays accounts for roughly a third of daily variance. Simply adding the squared overnight return is no solution either, since a single squared return is an extremely noisy variance measure; the weight it receives must trade off its information against its noise.

The main finding of Hansen & Lunde (2005) is that this trade-off has an operational optimal solution despite the latent nature of the target. Restricting attention to conditionally unbiased estimators, the weights that minimize the mean-squared error depend only on moments that simple sample averages estimate consistently. Empirically, the optimal combination assigns the intraday measure a weight near unity and shrinks the squared overnight return heavily, yet the overnight component always retains some weight. Their contribution to the previous literature is to replace the practices that preceded them of ignoring the overnight period, such as scaling up the intraday measure, or adding the raw squared overnight return, with a single framework that combines these choices, states the assumptions under which each is justified, and provides tests of those assumptions.

I recover the central results of the original study on a sample that differs deliberately from it. As to methods, I follow the original construction throughout: intraday returns on a previous-tick grid, a Newey-West correction with a Bartlett kernel for the microstructure-induced autocorrelation, bid and ask series averaged to offset the bounce, the same moment-based optimal weights with a similar outlier rule, and the same specification checks, the kitchen sink regressions and conditional-unbiasedness ratios. As to data, I use one-second best bid and offer quotations from the Nasdaq TotalView-ITCH database for the 30 components of the Dow Jones Industrial Average (DJIA) over 2,034 trading days from May 2018 to June 2026, twice the amount of the original panel of 2001 to 2004, and I add three index ETFs (DIA, SPY, and QQQ), an instrument class absent from the original study. As to findings, the structure of the original results survives two decades later: volatility levels vary widely across equities while their intraday-to-overnight composition is comparatively stable, the correlation between the two variance components remains near 0.3, and the intraday measure dominates the optimal combination. Two results are new. The overnight share of daily variance has risen from 20% to 30%, and it is largest at the index level: the diversified overnight returns of the ETFs are precise enough to earn optimal weights several times larger than those of any single stock. And the identifying assumptions, which the original study could not reject on four years of data, are rejected on this longer sample for most instruments examined: the overnight share is not a fixed proportion.

The remainder of the article is organized as follows. Section 2 places the paper in the literature. Section 3 presents the theory and the econometric model: the realized variance and its noise correction, the sampling scheme, and the optimal combination of overnight and intraday variance. Section 4 describes the data and discusses the estimates, first for the full cross-section and then, together with the tests of the identifying assumptions, for the technology equities and the ETFs, closing with limitations and practical relevance. Section 5 concludes. All figures and tables are collected in the Appendix.

2. Literature Review

This article draws on three main aspects: the literature on market microstructure noise, which explains why a realized variance cannot simply be computed at the highest frequency the data allow, the literature on modelling irregularly spaced financial data, and the work on extending intraday variance measures to the whole trading day, in which Hansen & Lunde (2005) is the central contribution and the paper on which this article builds. The first two have grown into a field of their own, with dedicated books by Hautsch (2011) and Ait-Sahalia & Jacod (2014).

The economic idea that observed prices deviate from efficient ones goes back to Black (1986): noise trading is what makes continuous markets possible in the first place, yet the same noise keeps prices from fully revealing fundamentals, so every observed price carries a transitory component. The theory of how prices emerge from the trading process is laid out by O'Hara (1998), and its empirical and econometric counterpart by Hasbrouck (2007); for the present purpose, three concrete mechanisms matter most. Roll (1984) showed that the mere existence of a bid-ask spread makes observed price changes negatively autocorrelated, as transactions bounce between the two sides of the quote, and turned this observation into his implicit spread estimator; Harris (1990) characterized the statistical properties of that estimator, showing how noisy the serial covariance of observed returns is in finite samples. Gottlieb & Kalay (1985) added a second mechanism, the rounding of prices onto a discrete grid, and quantified the biases that discreteness induces in the moments of returns. A third mechanism is adverse selection: part of the spread compensates liquidity providers for trading against better-informed agents, so the noise also carries an information-driven component. For volatility measurement these frictions are not a side issue but a constraint: the finer the sampling, the larger the share of each squared return that is noise rather than efficient-price variation, so the realized variance becomes increasingly biased exactly where it should become increasingly precise.

The remedy adopted here corrects the realized variance with its own empirical autocovariances, weighted by a Bartlett kernel so that the resulting measure cannot turn negative, in the spirit of the covariance matrix estimator of Newey & West (1986). Getting the measure right matters beyond the measure itself: Hansen & Lunde (2006) show that when volatility models are evaluated against a noisy proxy, the empirical ranking of the models can contradict the true one unless the evaluation criterion satisfies restrictive conditions, and Ghysels & Sinko (2011) document that the treatment of the noise likewise affects the accuracy of the volatility forecasts built on realized measures. An accurate measure of daily variance is therefore a prerequisite for model comparison and forecasting, not a refinement of them. Ait-Sahalia & Jacod (2014) synthesize this literature, from noise-robust estimation to the treatment of jumps, and document the pronounced non-normality of high-frequency returns, a regularity known since Mandelbrot (1963) proposed stable distributions for speculative prices, one that resurfaces in the descriptive statistics of Section 4, and one taken to its extreme by modelling high-frequency prices as pure jump processes (Jing et al., 2012).

A separate methodological choice concerns the irregular timing of the raw data. Quotes do not arrive on a grid, and one branch of the literature treats the arrival times themselves as the object of interest: the autoregressive conditional duration model of Engle & Russell (1998) models the times between transactions much as GARCH models returns, and Bauwens & Hautsch (2009) test the resulting point-process framework. The realized variance tradition takes the simpler route followed in this paper, interpolating quotes onto a calendar-time grid; for a day-level aggregate built from thousands of intervals this is adequate, though the duration literature is a reminder that the interpolation discards the information carried by the timing of quotes. Li et al. (2014) show that when sampling times are endogenous, correlated with the price process they represent, the standard asymptotics of realized volatility no longer apply.

The third body of work concerns the part of the day without prices. Volatility does not stop at the close: earnings are released after the bell and macroeconomic news before the open, and whatever accumulates overnight must be included in the opening quote. Applied work initially handled this by ignoring the overnight period or by rescaling the intraday measure up to the close-to-close level, as in Koopman et al. (2005), who scale the realized variance of the S&P 100 in a forecast comparison against historical and implied volatility measures. Hansen & Lunde (2005) replaced these conventions with a single framework: the squared overnight return and the noise-corrected intraday measure are two noisy readings of the same latent daily variance, the class of conditionally unbiased linear combinations of the two is characterized by a small set of identifying conditions, and the mean-squared-error-optimal weights within that class are estimable from simple sample moments. Scaling and adding then emerge as special cases, and the identifying conditions become testable, which is the part of the framework this article pushes hardest, on a sample twice as long as the original and extended to an instrument class, the index ETFs, that the original study did not cover.

3. Theory and Econometric Application

This section presents the econometric framework. I first define the realized variance, the sampling scheme, and the noise correction, and then derive the optimal combination of overnight and intraday variance from the conditions that identify it; the closed-form weights and their sample counterparts are collected in Table 1 of the Appendix.

3.1. Theoretical Framework

Let $\{p^*(t)\}_{t\in[0,\infty)}$ denote the efficient log-price process, where the unit of time is one trading day and integer values of t correspond to market closes. In empirical settings, the efficient price is unobservable due to microstructural frictions, informational or not, such as bid-ask spread bounces, differences in trade sizes, the informational content of price changes, and the discreteness of price changes (Ait-Sahalia & Jacod, 2014). The observed log-price process $p(t)$ is therefore modeled as a noise-contaminated variable: $p(t) = p^*(t) + u(t)$, where $u(t)$ represents an independent and identically distributed (i.i.d.) error with zero mean and time-invariant variance. In the words of Black (1986): *noise causes markets to be somewhat inefficient, but often prevents us from taking advantage of inefficiencies.* Here, those inefficiencies enter additively, although the alternative specifications reviewed in Section 2 tie the noise to specific frictions such as the bid-ask spread or price rounding. The observed log return spanning trading day t is defined as: $r_t \equiv p(t) - p(t-1)$.

I assume that the price satisfies $\mathrm{d}p^*(t) = \mu(t)\,\mathrm{d}t + \sigma(t)\,\mathrm{d}w(t)$, a stochastic differential equation where w is a standard Brownian motion and the drift μ and the spot volatility σ are time-varying random functions independent of w. The target of estimation over a window $[a,b]$ is then the integrated variance (IV) of the efficient price:

$$IV_{[a,b]} \equiv \int_a^b \sigma^2(s)\,\mathrm{d}s \tag{3.1}$$

The realized variance (RV) estimates this target from squared intraday returns. On an equidistant grid with sampling step $\Delta \equiv \frac{b-a}{m}$, the intraday returns are defined as $y_{i,m} = p(t_i) - p(t_i - \Delta)$ and the estimator reads:

$$RV_{[a,b]}^{(\Delta)} = \sum_{i=1}^{m} y_{i,m}^2 = \sum_{i=1}^{m} \left[p(t_i) - p(t_i - \Delta)\right]^2 \tag{3.2}$$

As the grid becomes finer ($\Delta \to 0$), the realized variance converges in probability to the integrated variance of the efficient price; once noise contaminates the observed price, however, the bias grows with the sampling frequency. The raw quotes, moreover, arrive at irregular, random times. To place them on the regular Δ-grid I use previous-tick interpolation (Wasserfallen & Zimmermann, 1985).

A well-documented feature of high-frequency data is that, as the sampling frequency increases ($\Delta \to 0$), the intraday returns $y_{i,m}$ become strongly autocorrelated. The autocorrelation is driven mainly by the bid-ask bounce, which pushes consecutive observed returns in opposite directions. A second feature of the data at this level of aggregation, strongly suggestive of a departure from the Brownian assumption, is non-normality: log-returns sampled at high frequency are far from normally distributed, the regularity first documented by Mandelbrot (1963) that motivated jump models for prices. The departure depends on the sampling step: at the finest grids the distribution of $y_{i,m}$ is sharply peaked with heavy tails, and as Δ grows the aggregation of returns pulls it back toward the Gaussian benchmark. Figure 2 illustrates both facts in the present data.

3.2. A Correction for Microstructure Noise

To obtain an estimator of $IV_{[a,b]}$ that is robust to this serial dependence, I correct the realized variance of Equation (3.2), as in Hansen & Lunde (2005), with a Newey & West (1986) adjustment based on the Bartlett kernel:

$$RV_{NW[a,b]}^{(\Delta)} \equiv RV_{[a,b]}^{(\Delta)} + 2\sum_{h=1}^{q} \left(1 - \frac{h}{q+1}\right) \sum_{i=1}^{m-h} y_{i,m} y_{i+h,m} \tag{3.3}$$

where h is the lag distance between two intraday returns and q is the maximum lag included. The Bartlett weights $\left(1 - \frac{h}{q+1}\right)$ guarantee that the estimator is nonnegative. Following Hansen & Lunde (2005), the lag q is chosen so that the autocovariance window spans approximately 10 minutes of trading time: $q \times \Delta \approx 10$ minutes. This window is long enough for the serial dependence induced by the noise to be reduced, so the correction diminishes the bias without inflating the variance of the estimator.

3.3. Optimal Combination of Variances

The first component is the squared overnight return $r_{1,t}^2$, where the overnight return is the log-difference between the first mid-quote of day t and the last mid-quote of day $t-1$; it captures the variance accumulated while the market is closed. The second is a noise-robust intraday realized variance, computed as the average of the Equation (3.3) estimator applied separately to the ask and bid quote series,

$$RV_{2,t} \equiv \frac{1}{2} RV_{NW[a,b],\text{ ask}}^{(\Delta)} + \frac{1}{2} RV_{NW[a,b],\text{ bid}}^{(\Delta)} \tag{3.4}$$

The daily integrated variance is then estimated by the linear combination

$$RV_t(\omega) = \omega_1 r_{1,t}^2 + \omega_2 RV_{2,t} \tag{3.5}$$

Because IV_t is never observed, the weights are only identified once the two observable components are tied to the latent target. Following Hansen & Lunde (2005), I assume that a fixed share of the daily integrated variance happens during the trading session, whatever the level of volatility, and that each component measures its own segment of the day, overnight or intraday, with a proportional bias. The first assumption still allows the price dynamics of the two periods to differ, and calendar patterns such as elevated weekend volatility remain admissible provided they scale both periods alike; it is therefore the restrictive condition, the fixed split of the day, and Section 4 puts it to the test.

Under this condition[1], the combinations that are conditionally unbiased for IV_t are exactly those whose weights satisfy $\omega_1\mu_1 + \omega_2\mu_2 = \mu_0$, where $\mu_0 \equiv \mathbb{E}(IV_t)$, $\mu_1 \equiv \mathbb{E}(r_{1,t}^2)$, and $\mu_2 \equiv \mathbb{E}(RV_{2,t})$; and for a conditionally unbiased estimator, minimizing the mean-squared error is equivalent to minimizing the unconditional variance of the combination subject to that constraint, a problem that involves no latent quantity. The solution takes the form

$$\omega_1^* \equiv (1-\varphi)\frac{\mu_0}{\mu_1}, \qquad \omega_2^* \equiv \varphi\frac{\mu_0}{\mu_2} \tag{3.6}$$

where the importance factor φ approaches one as the intraday measure becomes more precise relative to the overnight one; its closed form, together with the sample counterparts of all the moments involved, is collected in Table 1 of the Appendix.

[1] Two secondary assumptions are that neither component needs to be unbiased, and that a law of large numbers applies for the daily sequences, so that sample averages estimate the moments consistently.

Replacing the moments by their sample averages yields the feasible optimal estimator of daily volatility:

$$RV_t^* \equiv \hat{\omega}_1 r_{1,t}^2 + \hat{\omega}_2 RV_{2,t} \tag{3.7}$$

One practical detail completes the estimation: the moment estimates are sensitive to the most extreme overnight observations, since prices occasionally jump between close and open in response to news. Following the outlier rule of Hansen & Lunde (2005), detailed in the note to Table 1, I exclude about 2% of the trading days from the estimation of the weights.

4. Empirical Results

For the application, I use quotation data come from Databento, specifically the *Best Bid and Offer (BBO) on interval* dataset extracted from the *Nasdaq TotalView-ITCH* feed: for every second of the session in which activity occurs, the record reports the last best bid, best offer, and sale price. I retain the bid and ask quotes with their timestamps in Eastern Time and restrict each series to the regular 9:30 to 16:00 session. The clean-up removes non-positive and crossed quotes, filters out quotes whose relative spread exceeds 5%, and discards the 18 early-close sessions in the sample; quotations are noisier for the equities whose primary listing is outside NASDAQ, since they originate from a non-primary venue. To remove the price discontinuities induced by corporate actions, which matter only for overnight returns, I back-adjust the quotes for stock splits, spin-offs, and cash dividends, such as Walmart's 3-for-1 split of February 2024 and 3M's spin-off of Solventum in April 2024.

The sample spans May 2018 to June 2026, 2,034 trading days, and covers the 30 DJIA equities, following the original application of Hansen & Lunde (2005), plus three ETFs, listed with the full cross-section in Table 2 of the Appendix: DIA and SPY as proxies of the Dow Jones Industrial Average and the S&P 500, and QQQ as a proxy of technology companies. Two descriptive features of the cleaned data, illustrated for AAPL and DIA in Figure 1 and Figure 2, connect directly to the model of Section 3. Quotes arrive irregularly: most updates follow the previous one within seconds, with a right tail of quiet intervals and occasional halts, which is what makes the previous-tick interpolation necessary before any equidistant estimator can be computed. And intraday returns are strongly non-normal at fine sampling steps: standardized returns at $\Delta = 1$ second are sharply peaked with tails far above the Gaussian benchmark, and the densities drift toward the normal as Δ widens. The latter makes the choice of Δ concrete: finer sampling uses more of the data but are dominated by noise, and it is the Newey-West correction that allows $\Delta = 10$ seconds for the intraday measure, while the level μ_0 is fixed at the nearly noise-free step of 15 minutes.

4.1. Model Estimates and Discussion

The output of the estimation is, for each instrument, a daily series of the optimal whole-day variance measure RV_t^* together with the weights that define it. Figure 3 plots the resulting series, on a logarithmic scale, for AAPL and DIA over the full sample. Both display the familiar features of realized volatility: pronounced clustering, with quiet stretches interrupted by persistent high-volatility episodes, most prominently at the COVID-19 outbreak of March 2020 and again in the spring of 2025; the former delivers the largest whole-day variance realizations for the index fund, the latter for Apple.

4.2. Full Cross-Section

Table 3 reports, for each of the 33 instruments, the estimated moments, the importance factor $\hat{\varphi}$, and the optimal weights $\hat{\omega}_1^*$ and $\hat{\omega}_2^*$ that define the optimal measure of daily volatility. The number of trading days retained in the estimation, n, lies between 1996 and 2000, after the removal of the days under the outlier rule.

Several observations emerge from the table. First, the column of $\hat{\mu}_0$ reveals substantial variation in the level of daily volatility across instruments: the estimates range from 0.67 for DIA to 5.55 for NVDA, in the percent-squared units of the table, corresponding to annualized volatilities of roughly 13% and 37%, respectively. The ratio of "active" to "inactive" volatility, $\hat{\mu}_2/\hat{\mu}_1$, is far more stable: for most individual equities it lies between 2 and 3.6, with a median of about 2.3, which implies that roughly 30% of daily variance occurs during the close-to-open period. This overnight share is noticeably larger than the approximately 20% (a ratio of about 4) documented by Hansen & Lunde (2005) for the DJIA panel of 2001 to 2004, suggesting that the inactive period has gained importance in the price process over the past two decades. The dispersion across instruments is also informative: the three ETFs sit at the bottom of the distribution, between 1.76 and 1.87, a range matched among the equities only by Nike (1.83), followed by the financial, communication, and technology equities, roughly between 2 and 2.5, whereas the defensive names, such as Procter & Gamble (3.58), Walmart (3.44), and Coca-Cola (3.21), together with Amgen (4.83), leave the smallest fraction of their variance to the trading session. The relatively low ratios of the technology equities suggest that technology stocks have a comparatively large proportion of their volatility occurring overnight.

Second, there is a great deal of dispersion in the volatility-of-volatility estimates, $\hat{\eta}_1^2$ and $\hat{\eta}_2^2$, and in their ratio. The correlations $\hat{\eta}_{12}/(\hat{\eta}_1\hat{\eta}_2)$, in contrast, are remarkably similar across instruments, lying between 0.22 and 0.43 and mostly in the neighborhood of 0.3, the same range reported in the original study.

Third, the importance factor $\hat{\varphi}$ is close to one for all individual equities, between 0.82 and 0.99, reflecting that the intraday measure $RV_{2,t}$ is far more precise. Consequently, the optimal estimator scales the intraday component by a weight $\hat{\omega}_2^*$ near unity, from 0.63 to 1.06, while the squared overnight return is scaled down heavily, with $\hat{\omega}_1^*$ between 0.03 and 0.45: the optimal combination places disproportionately more weight on the intraday measure. The three ETFs stand apart from this pattern. Their importance factors, 0.71 (SPY) to 0.79 (QQQ), fall below those of every individual equity, and their overnight weights, $\hat{\omega}_1^*$ between 0.45 and 0.60, are the largest in the panel, with only Procter & Gamble (0.45) reaching the bottom of that range. The economic interpretation is a diversification effect: at the index level, idiosyncratic overnight news might be washed out, so the squared overnight return of a broad portfolio is a much less noisy volatility measure, relative to its intraday counterpart, than that of any single stock, and it accordingly earns a weight several times larger.

The economic magnitude of these weights is best read as the share of the optimal measure that each component contributes on an average day. Since the weights satisfy $\omega_1\mu_1 + \omega_2\mu_2 = \mu_0$, the fraction $\omega_1\mu_1/\mu_0$ measures the contribution of the overnight component to RV_t^*. For a typical individual equity this contribution is modest, 6.6% for Apple and as little as 1.5% for Nike, so that in effect the whole-day measure is a rescaled intraday measure. For the index funds the picture is different: the overnight component accounts for 29% of the optimal measure for SPY, 24% for DIA, and 21% for QQQ. The lackings of ignoring the overnight period altogether are also easy to state in volatility units: with about 36% of close-to-close variance occuring overnight for the ETFs, an open-to-close measure alone understates whole-day volatility by roughly 20%; applied to DIA's whole-day level of 13% annualized, that means reporting a volatility near 10.5%. A further contrast with the original estimates concerns the scale of $\hat{\omega}_2^*$: in the panel of 2001 to 2004 the intraday weight typically exceeded one, because with an overnight share near 20% the intraday measure had to be scaled up to reach the whole-day level, whereas here it lies slightly below one for most instruments, a direct reflection of the larger overnight share.

4.2.1. Tests of the Identifying Assumptions

The optimally combined estimator relies critically on the fixed-split condition of Section 3, so I test its validity following the specification checks of Hansen & Lunde (2005). The logic of the first check is simple. If a fixed share of each day's variance occurs overnight, then the ratio of the two components, $r_{1,t}^2/RV_{2,t}$, moves around a constant only through the measurement errors of day t, and its logarithm should therefore be unpredictable. Conversely, if the overnight share rises and falls with market conditions, then variables that predict volatility will also predict the log-ratio. This motivates the *kitchen sink regression*

$$\log\left(r_{1,t}^2/RV_{2,t}\right) = \alpha + \beta' Z_t + u_t \tag{4.1}$$

in which a rejection of $H_0 : \beta = 0$ is evidence against the fixed split. For the test to be valid, the instruments in Z_t must be unrelated to the measurement errors, which are specific to day t; for it to be powerful, they should be good predictors of volatility. Both requirements point to the same choices: the lagged intraday measure, $\log\big(RV_{2,t-1}\big)$, a strong predictor because volatility is highly persistent, and weekday dummies for Monday through Thursday (Friday is omitted to avoid collinearity), which capture the systematic weekly pattern in average volatility. Estimation is by least squares with heteroskedasticity-robust p-values (White, 1980), on a sample that excludes the days where $r_{1,t} = 0$, for which the logarithm is undefined.

The second check targets conditional unbiasedness directly: the optimal measure should be right on average not only over the whole sample, but within any group of days selected on prior information, Monday or Friday. To make this operational, I sort the days into quintiles by the previous day's intraday variance, and separately by weekday, and compare within each group the average of RV_t^* with the average of $r_{1,t}^2 + RV_{2,t}^{(\Delta_0)}$, the nearly unbiased whole-day proxy built on the 15-minute grid. If RV_t^* is conditionally unbiased, the ratio of the two averages should be close to one in every group; a ratio above one in the turbulent quintiles, say, would reveal that the measure systematically overstates variance when markets are volatile. Because a few extreme days can dominate these averages, I compute the ratios on the outlier-truncated sample. The formal counterpart of the ratios replaces the dependent variable of Equation (4.1) with $\log\Big[RV_t^*/\Big(r_{1,t}^2 + RV_{2,t}^{(\Delta_0)}\Big)\Big]$ and runs the same regression on the full sample: under conditional unbiasedness, nothing known at the start of the day should predict this ratio either.

4.2.2. Technology Equities and ETFs

This subsection examines in closer detail the two groups that bracket the cross-section in Table 5: the technology equities, the most volatile sector, and the ETFs, the least volatile one. Technology displays the highest average level of daily variance (mean $\hat{\mu}_0$ of 2.56, pulled up by NVDA) and by far the largest volatility-of-volatility, whereas the ETFs show the lowest levels (mean $\hat{\mu}_0$ of 0.89) and the smallest active-to-inactive ratios of the panel. The two groups also sit at opposite ends in the composition of the optimal estimator: the mean overnight weight $\hat{\omega}_1^*$ is 0.15 for technology but 0.52 for the ETFs, the possible diversification effect discussed above.

Given the prominence of these two groups, they are the natural candidates on which to test the identifying assumptions that support RV_t^*. Table 4 reports the results of the specification checks described above: the p-values of the kitchen sink regression Equation (4.1) (Panel A), and the bias ratios across quintiles of lagged intraday variance and across weekdays, together with the p-values of the conditional-unbiasedness regression (Panel B).

The bias ratios of Panel B are reassuring: all ratios lie between 0.84 and 1.17, and the majority within 5% of unity, so there are no large systematic difficulties in the conditional unbiasedness of RV_t^* across volatility regimes or weekdays. The most visible pattern is a weekly seasonal one in the index funds: the Monday ratios sit below one (0.91 for DIA, 0.94 for SPY) while the Friday ratios sit above it, at 1.03 to 1.06, indicating that the optimal measure slightly underestimates the variance of days that follow a weekend and overstates that of days that precede one. This is exactly where a deviation should appear if weekend volatility does not scale proportionally with trading-session volatility.

The formal tests are considerably less favorable than in the original application. The kitchen sink regression of Panel A rejects the null at the 5% level for seven of the nine instruments; only Cisco (p-value of 0.106) passes this first check, with IBM at the boundary (0.049). There is one possible explanation for this: for the large technology names, following high-volatility days, a disproportionately large share of the daily variance occurs overnight, so the fixed active-to-inactive split does not hold. For the index funds the weekday dummies contribute just as strongly, while IBM's rejection rests on the weekday effects alone. The conditional-unbiasedness regressions of Panel B reject throughout, with the loading on lagged volatility now negative: relative to the nearly unbiased benchmark $r_{1,t}^2 + RV_{2,t}^{(\Delta_0)}$, the optimal measure is somewhat too high after calm days and too low after turbulent ones. This stands in contrast to Hansen & Lunde (2005), who found significant p-values for only two of thirty equities in Panel A and three in Panel B. I find two possible roots for this. First, my sample contains more than twice as many trading days ($n \approx 2030$ against 986), so the tests have substantially more power to detect possible deviations. Second, the sample considers 2018 to 2026, a period that includes the COVID-19, during which the overnight share of volatility is unlikely to have remained a fixed proportion. Taken together, the evidence suggests that RV_t^* remains a reasonable measure of daily volatility on average, but that the consistency of the optimal weights, and in particular the proportionality assumption, is just an approximation over a long modern sample than it was over the four years studied in the original article. The practical content of the rejections is that a fixed-weight RV_t^* underweights the overnight component precisely during high volatilty episodes.

4.2.3. Limitations and Practical Relevance

Three limitations qualify these results. First, the quotations originate from a single venue: for the equities whose primary market is the NYSE, the NASDAQ feed is a non-primary source, and part of the measured noise may be venue-specific rather than natural to the price process. Second, the DJIA is a price-weighted index, so the results for DIA are not a capitalization-weighted aggregate of the component results, and comparisons between the fund and its components partially reflect the weighting scheme. Third, and most consequentially, the weights are held fixed over eight years, and the specification tests show that this constancy might not be the best practice; allowing the weights to vary with the level of volatility is a natural refinement for a following research project.

5. Conclusions

This article asked how the variance of a full trading day should be measured when high-frequency prices exist only for the trading session, and answered it by extending the framework of Hansen & Lunde (2005) to one-second quotations for the DJIA equities and three index ETFs over 2018 to 2026. The mechanics of the original application survive intact: the optimal combination places disproportionate weight on the noise-corrected intraday measure, yet the squared overnight return always retains informative weight. What has changed is the object being measured. The overnight period now carries roughly 30% of daily variance, against about 20% two decades ago; the gain is most pronounced for the diversified ETFs, whose overnight returns earn weights several times larger than those of any single stock.

Beyond the limitations discussed in Section 4, the results carry an implication for how realized measures are used in practice: a framework that treats the trading session as the whole day understates daily risk materially more today than when the method was proposed, and the understatement is largest exactly where realized measures are most consumed, at the index level.

Future work follows naturally: weights that vary with the level of volatility or are re-estimated over rolling windows, weights pooled within sectors to sharpen the ticker-level estimates, and jump-robust variants that separate the noisy and time-variant components of both periods.

Bibliography

Ait-Sahalia, Y., & Jacod, J. (2014). *High-Frequency Financial Econometrics*. Princeton University Press.

Bauwens, L., & Hautsch, N. (2009). Modelling Financial High Frequency Data Using Point Processes. In *Handbook of Financial Time Series* (pp. 953–979). Springer.

Black, F. (1986). Noise. *The Journal of Finance*, *41*(3), 528–543.

Engle, R. F., & Russell, J. R. (1998). Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data. *Econometrica*, *66*(5), 1127.

Ghysels, E., & Sinko, A. (2011). Volatility forecasting and microstructure noise. *Journal of Econometrics*, *160*(1), 257–271.

Gottlieb, G., & Kalay, A. (1985). Implications of the Discreteness of Observed Stock Prices. *The Journal of Finance*, *40*(1), 135–153.

Hansen, P. R., & Lunde, A. (2005). A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data. *Journal of Financial Econometrics*, *3*(4), 525–554.

Hansen, P. R., & Lunde, A. (2006). Consistent ranking of volatility models. *Journal of Econometrics*, *131*(1–2), 97–121.

Harris, L. (1990). Statistical Properties of the Roll Serial Covariance Bid/Ask Spread Estimator. *The Journal of Finance*, *45*(2), 579–590.

Hasbrouck, J. (2007). *Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading*. Oxford University Press.

Hautsch, N. (2011). *Econometrics of Financial High-Frequency Data*. Springer Science & Business Media.

Jing, B.-Y., Kong, X.-B., & Liu, Z. (2012). Modeling High-Frequency Financial Data by Pure Jump Processes. *The Annals of Statistics*, *40*(2). https://doi.org/10.1214/12-aos977

Koopman, S. J., Jungbacker, B., & Hol, E. (2005). Forecasting Daily Variability of the S&P 100 Stock Index Using Historical, Realised and Implied Volatility Measurements. *Journal of Empirical Finance*, *12*(3), 445–475.

Li, Y., Mykland, P. A., Renault, E., Zhang, L., & Zheng, X. (2014). Realized Volatility When Sampling Times Are Possibly Endogenous. *Econometric Theory*, *30*(3), 580–605.

Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *The Journal of Business*, *36*(4), 394–419. http://www.jstor.org/stable/2350970

Newey, W. K., & West, K. D. (1986). *A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix*.

O'Hara, M. (1998). *Market Microstructure Theory*. John Wiley & Sons.

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.

Wasserfallen, W., & Zimmermann, H. (1985). The Behavior of Intraday Exchange Rates. *Journal of Banking and Finance*, *9*(1), 55–72.

White, H. (1980). A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. *Econometrica*, *48*(4), 817–838.

Appendix

Table 1: Optimal weights: population quantities and sample estimators.

Quantity	Population definition	Sample estimator
μ_1	$\mathbb{E}\big(r_{1,t}^2\big)$	$\frac{1}{n}\sum_{t=1}^{n} r_{1,t}^2$
μ_2	$\mathbb{E}\big(RV_{2,t}\big)$	$\frac{1}{n}\sum_{t=1}^{n} RV_{2,t}$
μ_0	$\mathbb{E}(IV_t)$	$\frac{1}{n}\sum_{t=1}^{n}\Big(r_{1,t}^2 + RV_{2,t}^{(\Delta_0)}\Big)$
η_1^2	$\mathrm{Var}\big(r_{1,t}^2\big)$	$\frac{1}{n}\sum_{t=1}^{n}\big(r_{1,t}^2 - \hat{\mu}_1\big)^2$
η_2^2	$\mathrm{Var}\big(RV_{2,t}\big)$	$\frac{1}{n}\sum_{t=1}^{n}\big(RV_{2,t} - \hat{\mu}_2\big)^2$
η_{12}	$\mathrm{Cov}\big(r_{1,t}^2, RV_{2,t}\big)$	$\frac{1}{n}\sum_{t=1}^{n} RV_{2,t}\big(r_{1,t}^2 - \hat{\mu}_1\big)$
φ	$\frac{\mu_2^2\eta_1^2 - \mu_1\mu_2\eta_{12}}{\mu_2^2\eta_1^2 + \mu_1^2\eta_2^2 - 2\mu_1\mu_2\eta_{12}}$	plug-in
ω_1^*	$(1-\varphi)\mu_0/\mu_1$	$(1-\hat{\varphi})\hat{\mu}_0/\hat{\mu}_1$
ω_2^*	$\varphi\mu_0/\mu_2$	$\hat{\varphi}\hat{\mu}_0/\hat{\mu}_2$

Note: Sample estimators are computed over the n trading days retained in the estimation. The step Δ_0 is set to 15 minutes, at which the microstructure bias of the plain realized variance is negligible, so that $\hat{\mu}_0$ approximates the level $\mathbb{E}(IV_t)$ without bias. Outlier rule: overnight returns tend to be small, but some are very large because prices jump between close and open in response to news, and without truncation these days dominate the moment estimates. Following Hansen & Lunde (2005), who removed the seven largest observations of $r_{1,t}^2$ and the three largest of $RV_{2,t}$ (about 1% of their sample), I classify about 2% of the trading days as outliers, the largest observations of $r_{1,t}^2$ and of $RV_{2,t}$ in a 3-to-1 proportion, and exclude them from the estimation of the moments and weights but not from the subsequent evaluation of the estimator.

Table 2: ETFs and Equities included in the empirical analysis.

Ticker	Name	Exchange	Sector	Weighting
ETFs				
DIA	SPDR DJIA ETF	NYSE	Index Fund	–
SPY	SPDR S&P500 ETF	NYSE	Index Fund	–
QQQ	Invesco QQQ ETF	NASDAQ	Index Fund	–
DJIA Components				
AAPL	Apple	NASDAQ	Technology	3.56%
AMGN	Amgen	NASDAQ	Healthcare	4.18%
AMZN	Amazon	NASDAQ	Consumer Cyclical	2.76%
AXP	American Express	NYSE	Financial Services	3.82%
BA	Boeing	NYSE	Industrials	2.55%
CAT	Caterpillar	NYSE	Industrials	10.76%
CRM	Salesforce	NYSE	Technology	1.89%
CSCO	Cisco	NASDAQ	Technology	1.29%
CVX	Chevron	NYSE	Energy	2.00%
DIS	Disney	NYSE	Communication	1.10%
GOOGL	Alphabet	NASDAQ	Communication	4.11%
GS	Goldman Sachs	NYSE	Financial Services	11.68%
HD	Home Depot	NYSE	Consumer Cyclical	3.81%
HON	Honeywell Technologies	NASDAQ	Industrials	2.50%
IBM	IBM	NYSE	Technology	3.43%
JNJ	Johnson & Johnson	NYSE	Healthcare	2.99%
JPM	JPMorgan Chase	NYSE	Financial Services	3.75%
KO	Coca-Cola	NYSE	Consumer Defensive	0.95%
MCD	McDonald's	NYSE	Consumer Cyclical	3.16%
MMM	3M	NYSE	Industrials	1.76%
MRK	Merck	NYSE	Healthcare	1.43%
MSFT	Microsoft	NASDAQ	Technology	4.35%
NKE	Nike	NYSE	Consumer Cyclical	0.49%
NVDA	Nvidia	NASDAQ	Technology	2.32%
PG	Procter & Gamble	NYSE	Consumer Defensive	1.68%
SHW	Sherwin-Williams	NYSE	Materials	3.75%
TRV	Travelers Companies, Inc.	NYSE	Financial Services	3.83%
UNH	UnitedHealth Group	NYSE	Healthcare	4.83%
V	Visa	NYSE	Financial Services	3.94%
WMT	Walmart	NASDAQ	Consumer Defensive	1.28%

The table lists the tickers, names, exchanges, and sectors of the ETFs and the 30 equities of the Dow Jones Industrial Average that are used in the empirical analysis. The DJIA components and weights are retrieved from State Street Investment Management, and are valid to the date of this article.

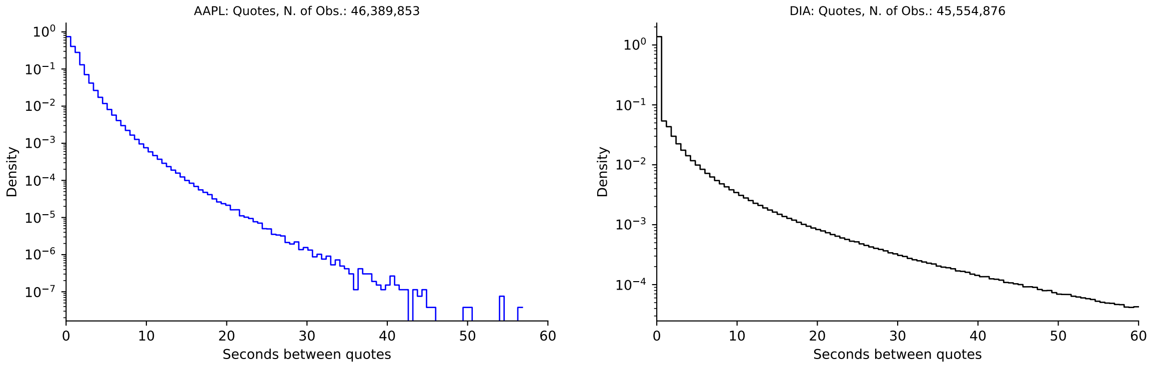

Figure 1: Time between consecutive quote updates: AAPL and DIA.

Note: The panels show the density, on a logarithmic scale, of the number of seconds elapsed between consecutive quote updates within the same trading day, over the full sample. Durations are censored at 60 seconds for display.

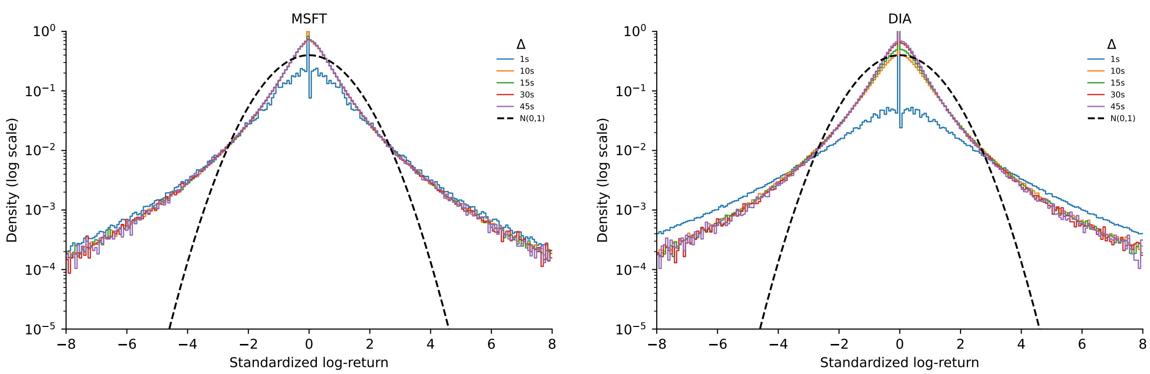

Figure 2: Log-plot of the marginal density of log-returns, at frequencies of 1, 10, 15, 30 and 45 seconds.

Note: The panels show the density, on a logarithmic scale, of intraday log-returns standardized to zero mean and unit variance, constructed by previous-tick interpolation at sampling steps Δ of 1, 10, 15, 30, and 45 seconds, against the standard normal density (dashed line).

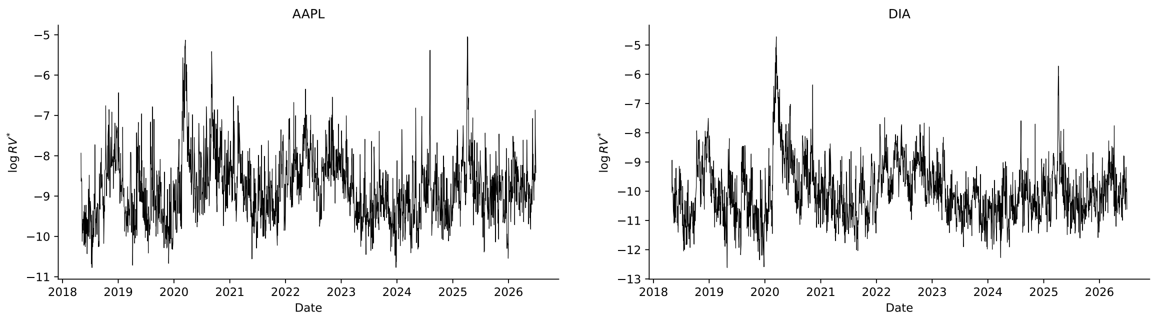

Figure 3: Optimal daily variance measure RV_t^*: AAPL and DIA.

Note: The panels plot the daily series of the optimal whole-day variance measure RV_t^* on a logarithmic scale, over the full sample from May 2018 to June 2026.

Table 3: Optimal daily volatility: empirical estimates.

Ticker	Quotes	n	$\hat{\mu}_0$	$\hat{\mu}_1$	$\hat{\mu}_2$	$\frac{\hat{\mu}_2}{\hat{\mu}_1}$	$\hat{\eta}_1^2$	$\hat{\eta}_2^2$	$\frac{\hat{\eta}_2^2}{\hat{\eta}_1^2}$	$\frac{\hat{\eta}_{12}}{\hat{\eta}_1\hat{\eta}_2}$	$\hat{\varphi}$	$\hat{\omega}_1^*$	$\hat{\omega}_2^*$
ETFs													
SPY	47.3	1998	0.74	0.36	0.63	1.76	0.48	0.82	1.69	0.32	0.71	0.60	0.84
QQQ	47.4	2000	1.27	0.59	1.11	1.87	1.16	1.76	1.51	0.34	0.79	0.45	0.90
DIA	45.6	1998	0.67	0.32	0.58	1.82	0.44	0.68	1.53	0.31	0.76	0.51	0.88
DJIA Components													
AAPL	46.4	1999	2.04	0.96	1.94	2.03	4.01	4.15	1.03	0.39	0.93	0.14	0.98
AMGN	21.3	1997	1.33	0.37	1.80	4.83	0.68	3.56	5.27	0.33	0.92	0.27	0.68
AMZN	40.4	1997	2.65	1.22	2.52	2.07	8.15	6.12	0.75	0.36	0.97	0.06	1.02
AXP	26.0	1998	2.21	1.11	2.16	1.95	5.19	6.06	1.17	0.32	0.86	0.27	0.88
BA	27.1	1998	3.68	1.53	3.77	2.47	11.42	19.88	1.74	0.41	0.92	0.19	0.90
CAT	22.4	1997	2.47	1.16	2.48	2.13	5.28	3.96	0.75	0.30	0.95	0.10	0.95
CRM	30.2	1999	3.08	1.39	3.11	2.24	7.69	8.54	1.11	0.31	0.92	0.18	0.91
CSCO	41.1	1997	1.37	0.58	1.46	2.52	2.05	2.53	1.23	0.42	0.99	0.03	0.93
CVX	37.6	1999	2.02	0.98	2.00	2.05	3.47	5.26	1.52	0.43	0.88	0.25	0.89
DIS	34.9	1997	1.88	0.84	1.96	2.33	4.07	4.42	1.09	0.42	0.98	0.04	0.95
GOOGL	37.8	1998	2.23	1.03	2.13	2.06	5.51	3.58	0.65	0.34	0.98	0.04	1.03
GS	23.2	1997	2.20	0.98	2.27	2.31	3.45	4.14	1.20	0.37	0.94	0.13	0.91
HD	23.6	1999	1.55	0.68	1.70	2.52	1.91	2.60	1.36	0.35	0.94	0.14	0.86
HON	24.0	1997	1.34	0.57	1.46	2.58	1.46	2.42	1.65	0.34	0.91	0.20	0.84
IBM	26.5	1998	1.55	0.69	1.65	2.39	3.63	3.89	1.07	0.38	0.97	0.06	0.91
JNJ	30.9	1997	0.85	0.32	1.05	3.24	0.44	1.11	2.53	0.26	0.88	0.31	0.72
JPM	37.8	1997	1.75	0.83	1.68	2.03	2.89	3.01	1.04	0.37	0.92	0.16	0.96
KO	36.2	1999	0.76	0.28	0.91	3.21	0.42	0.78	1.86	0.36	0.97	0.08	0.81
MCD	21.8	2000	0.91	0.37	1.09	2.94	0.66	1.29	1.95	0.34	0.93	0.18	0.78
MMM	21.8	1998	1.62	0.66	1.89	2.85	1.60	3.30	2.06	0.33	0.91	0.23	0.78
MRK	31.6	1999	1.27	0.48	1.66	3.49	1.00	2.13	2.13	0.31	0.95	0.13	0.73
MSFT	45.5	1998	1.74	0.77	1.70	2.19	2.37	3.12	1.32	0.30	0.88	0.27	0.90
NKE	31.3	1996	2.24	1.13	2.08	1.83	7.21	4.17	0.58	0.38	0.98	0.03	1.06
NVDA	42.5	1999	5.55	2.43	5.50	2.26	19.08	26.10	1.37	0.33	0.90	0.24	0.90
PG	32.9	1999	0.84	0.31	1.11	3.58	0.39	1.48	3.79	0.22	0.83	0.45	0.63
SHW	13.7	1997	1.95	0.92	2.09	2.28	2.72	4.04	1.48	0.25	0.85	0.32	0.79
TRV	15.5	1997	1.65	0.84	1.76	2.09	2.25	3.19	1.42	0.25	0.82	0.35	0.77
UNH	22.3	1997	1.87	0.77	2.14	2.77	3.51	5.35	1.53	0.41	0.98	0.04	0.86
V	31.5	2000	1.44	0.65	1.53	2.36	1.95	2.88	1.48	0.32	0.89	0.24	0.84
WMT	33.2	1999	0.98	0.35	1.20	3.44	0.72	1.68	2.33	0.37	0.96	0.11	0.78

Note: Quotes is the number of best bid and offer records entering the estimation after the clean-up, in millions. n corresponds to the trading days included in the calculations after cleaning for outliers.

Table 4: Tests of the identifying assumptions: technology equities and ETFs.

	Panel A	**Panel B: Conditional Unbiasedness**										
Ticker	**p-value**	**q1**	**q2**	**q3**	**q4**	**q5**	**Mon**	**Tue**	**Wed**	**Thu**	**Fri**	**p-value**
ETFs												
SPY	**<.001**	0.99	0.99	0.98	0.98	1.02	0.94	0.98	1.01	1.03	1.03	**<.001**
QQQ	**0.021**	1.01	0.98	0.96	0.99	1.02	0.97	0.99	1.03	0.98	1.03	**0.004**
DIA	**<.001**	1.01	0.99	0.99	1.00	1.01	0.91	0.98	1.01	1.04	1.06	**<.001**
Technology												
AAPL	**<.001**	1.04	1.04	0.97	1.06	0.97	0.93	1.03	1.05	1.03	0.98	**<.001**
CRM	**0.022**	1.10	0.96	0.93	0.95	1.05	1.01	0.98	1.01	0.96	1.05	**<.001**
CSCO	0.106	0.99	1.00	0.95	1.01	1.02	0.95	1.05	1.01	0.92	1.08	**<.001**
IBM	**0.049**	1.13	1.17	0.84	0.95	1.04	0.97	1.03	0.99	1.01	0.99	**<.001**
MSFT	**<.001**	1.08	0.95	0.95	1.05	0.99	1.00	1.00	0.98	0.97	1.05	**<.001**
NVDA	**0.005**	1.00	0.96	1.01	0.98	1.02	0.99	1.00	1.00	0.99	1.02	**0.001**

Note: Panel A contains the p-values of the test of $H_0 : \beta = 0$ in the kitchen sink regression (Equation (4.1)). The first five columns of Panel B contain the bias ratios, the group average of RV_t^* over that of the proxy $r_{1,t}^2 + RV_{2,t}^{(\Delta_0)}$, across the quintiles of $RV_{2,t-1}$; the next five across weekdays; and the last column the p-values of the test of $H_0 : \beta = 0$ in the conditional-unbiasedness regression, which replaces the dependent variable of Equation (4.1) with $\log\left[RV_t^*/\left(r_{1,t}^2 + RV_{2,t}^{(\Delta_0)}\right)\right]$ on the full sample. All p-values are heteroskedasticity-robust; those significant at the 5% level are in bold.

Table 5: Statistical summary of empirical estimates, by sector

	$\hat{\mu}_0$	$\hat{\mu}_1$	$\hat{\mu}_2$	$\frac{\hat{\mu}_2}{\hat{\mu}_1}$	$\hat{\eta}_1^2$	$\hat{\eta}_2^2$	$\frac{\hat{\eta}_2^2}{\hat{\eta}_1^2}$	$\frac{\hat{\eta}_{12}}{\hat{\eta}_1\hat{\eta}_2}$	$\hat{\varphi}$	$\hat{\omega}_1^*$	$\hat{\omega}_2^*$
ETFs											
Mean	0.89	0.42	0.77	1.82	0.70	1.09	1.58	0.32	0.75	0.52	0.87
Min	0.67	0.32	0.58	1.76	0.44	0.68	1.51	0.31	0.71	0.45	0.84
50%	0.74	0.36	0.63	1.82	0.48	0.82	1.53	0.32	0.76	0.51	0.88
Max	1.27	0.59	1.11	1.87	1.16	1.76	1.69	0.34	0.79	0.60	0.90
Technology											
Mean	2.56	1.14	2.56	2.27	6.47	8.06	1.19	0.36	0.93	0.15	0.92
Min	1.37	0.58	1.46	2.03	2.05	2.53	1.03	0.30	0.88	0.03	0.90
50%	1.89	0.86	1.82	2.25	3.82	4.02	1.17	0.35	0.93	0.16	0.91
Max	5.55	2.43	5.50	2.52	19.08	26.10	1.37	0.42	0.99	0.27	0.98
Healthcare											
Mean	1.33	0.49	1.66	3.58	1.41	3.04	2.87	0.33	0.94	0.19	0.75
Min	0.85	0.32	1.05	2.77	0.44	1.11	1.53	0.26	0.88	0.04	0.68
50%	1.30	0.42	1.73	3.37	0.84	2.85	2.33	0.32	0.94	0.20	0.72
Max	1.87	0.77	2.14	4.83	3.51	5.35	5.27	0.41	0.98	0.31	0.86
Consumer Cyclical											
Mean	1.84	0.85	1.85	2.34	4.48	3.54	1.16	0.36	0.96	0.10	0.93
Min	0.91	0.37	1.09	1.83	0.66	1.29	0.58	0.34	0.93	0.03	0.78
50%	1.90	0.90	1.89	2.30	4.56	3.39	1.05	0.35	0.96	0.10	0.94
Max	2.65	1.22	2.52	2.94	8.15	6.12	1.95	0.38	0.98	0.18	1.06
Financial Services											
Mean	1.85	0.88	1.88	2.15	3.15	3.86	1.26	0.32	0.89	0.23	0.87
Min	1.44	0.65	1.53	1.95	1.95	2.88	1.04	0.25	0.82	0.13	0.77
50%	1.75	0.84	1.76	2.09	2.89	3.19	1.20	0.32	0.89	0.24	0.88
Max	2.21	1.11	2.27	2.36	5.19	6.06	1.48	0.37	0.94	0.35	0.96
Industrials											
Mean	2.28	0.98	2.40	2.51	4.94	7.39	1.55	0.35	0.92	0.18	0.87
Min	1.34	0.57	1.46	2.13	1.46	2.42	0.75	0.30	0.91	0.10	0.78
50%	2.05	0.91	2.18	2.52	3.44	3.63	1.70	0.34	0.92	0.20	0.87
Max	3.68	1.53	3.77	2.85	11.42	19.88	2.06	0.41	0.95	0.23	0.95
Consumer Defensive											
Mean	0.86	0.31	1.07	3.41	0.51	1.31	2.66	0.32	0.92	0.21	0.74
Min	0.76	0.28	0.91	3.21	0.39	0.78	1.86	0.22	0.83	0.08	0.63
50%	0.84	0.31	1.11	3.44	0.42	1.48	2.33	0.36	0.96	0.11	0.78
Max	0.98	0.35	1.20	3.58	0.72	1.68	3.79	0.37	0.97	0.45	0.81

Note: Sectors with three or more components. Units and scaling as in Table 3.