Beyond Returns: A Candlestick-Based Approach to Spot Covariance Estimation*

Yasin Simsek

Department of Economics, Duke University

This version: October 17, 2025

[Click here for the latest version]

Abstract

Spot covariance estimation is commonly based on high-frequency open-to-close return data over short time windows, but such approaches face a trade-off between statistical accuracy and localization. In this paper, I introduce a new estimation framework using high-frequency candlestick data, which include open, high, low, and close prices, effectively addressing this trade-off. By exploiting the information contained in candlesticks, the proposed method improves estimation accuracy relative to the return-based approach while preserving local structure. I further develop a test for spot covariance inference based on candlesticks that demonstrates reasonable size control and a notable increase in power, particularly in small samples, compared to the test built on returns. Motivated by recent work in the finance literature, I empirically test the market neutrality of the iShares Bitcoin Trust ETF (IBIT) using 1-minute candlestick data for the full year of 2024. The results show systematic deviations from market neutrality, especially in periods of market stress. An event study around FOMC announcements further illustrates the new method's ability to detect subtle shifts in relatively mild information events.

Keywords: Covariance estimation, high-frequency data, candlesticks, spot beta.

JEL Codes: C13, C32, C58, G11.

*I am very grateful to Andrew J. Patton, Tim Bollerslev and Anna Bykhovskaya for their guidance and support. I would also like to thank Anna Cieslak, Mehmet Caner, Peter R. Hansen, Ilze Kalnina, Raul Guarini Riva, Adam Rosen, George Tauchen, Christopher Walker and seminar participants at Duke University, 2025 Triangle Econometrics Conference, and 2025 Annual Meeting of SoFiE for their helpful comments and suggestions. Email: yasin.simsek@duke.edu.

1 Introduction

The covariance matrix of asset returns is fundamental to many financial applications such as portfolio selection, risk management, and asset pricing, which has led to a large literature on its estimation. A key finding of this literature is that covariances vary substantially over time. High-frequency intraday returns provide an effective means to capture this variation and have enabled more precise estimation over shorter horizons, such as a day or a month (e.g., Aït-Sahalia and Jacod (2014), Bollerslev et al. (2018a) and references therein). Recent evidence shows that covariances also vary significantly within a trading day (e.g., Bibinger et al. (2019), Andersen et al. (2021)), creating a demand for spot covariance estimates at specific points in time. Constructing such spot estimates, however, relies on returns sampled from narrow time intervals at moderate frequencies. Therefore, it suffers from limited local information and in turn generates a trade-off between statistical precision and localization. Sampling returns at finer frequencies is a natural way to address this trade-off, but doing so introduces market microstructure noise, requiring additional modeling assumptions to form an estimator.

Building on this background, I develop a new estimation framework for spot covariances using an additional source of information to alleviate the localization-precision trade-off. Specifically, my approach leverages the richer information in high-frequency candlesticks that contain the open, high, low, and close prices within each sampling interval, moving beyond the conventional methods that rely solely on (open-to-close) returns. Importantly, this work provides a practical way to improve estimates and enable reliable inference by avoiding intricate modeling assumptions while preserving the local structure of the estimates.

This paper makes three main contributions. First, it introduces a new candlestick-based estimator for spot covariances that exploits all available price information within each interval. This estimator is defined as a weighted sum of quadratic forms based on candlestick variables within a local window. The paper also develops an econometric framework that determines the estimator's form by minimizing a specific risk function, thereby improving efficiency relative to existing return-only method. Second, it presents a procedure for conducting inference on spot covariances using the new candlestick-based estimator, including a formal

hypothesis test with asymptotically correct size. Third, the practical value of the proposed estimator is demonstrated through an empirical application to assess the market neutrality of Bitcoin.[1] The empirical results show frequent rejections of the null hypothesis of market neutrality, challenging existing beliefs about Bitcoin's risk exposure.

The proposed candlestick-based estimator has several desirable properties. First, it is derived by minimizing an asymptotic approximation of the risk function. Here, the asymptotic approximation is performed by taking the size of the estimation window (i.e. number of observations in the estimation sample) as a fixed, possibly small constant. As such, my approach effectively accounts for the scarcity of local information. Moreover, the estimator's form is analytically tractable, resembling the least squares formula and does not depend on the specific sample at hand. This feature makes the estimator readily applicable to various settings without requiring further econometric procedures. I find that my estimator substantially reduces the asymptotic risk relative to the conventional estimators that only use return observations. Importantly, it continues to outperform even when implemented with shorter estimation windows. This makes my approach particularly useful for high-frequency event studies which require high localization for identification (e.g. Bollerslev et al. (2018b), Nakamura and Steinsson (2018)).[2]

Having established the candlestick-based spot covariance estimator, I further develop a formal procedure for inference on the spot covariance matrix. While recent studies have investigated candlestick-based inference for spot variances (e.g., Li et al. (2024); Bollerslev et al. (2024a)), inference methods for spot covariances remain underexplored. To address this gap, I develop a hypothesis test and corresponding test statistic. I show that this test statistic can be approximated by a limiting variable whose distribution can be characterized via Monte Carlo simulations under the null hypothesis. Accordingly, I determine the critical values by simulating the quantiles of the limiting variable. This inference procedure yields a test with asymptotically correct size. Moreover, a finite sample simulation study shows that

[1]The term "market neutrality" refers to the idea that an asset's returns do not comove with the overall market returns. Such assets are particularly important for risk management purposes because they can provide diversification benefits.

[2]Specifically, Nakamura and Steinsson (2018) exploit high-frequency bond returns over a short window around FOMC announcements to identify the effects of monetary policy shocks. Bollerslev et al. (2018b) estimate disagreement among investors by local jump regressions around news announcements.

the proposed test has reasonable size control and a notable increase in power, relative to the return-based test, particularly pronounced in small samples.

Motivated by the recent work of Liu and Tsyvinski (2021), I apply the candlestick-based estimation and inference framework to examine the market neutrality of Bitcoin, a popular cryptocurrency. In recent years, cryptocurrencies have become increasingly common in both institutional and retail portfolios.[3] However, their market exposure remains a subject of debate among academics and practitioners. Crypto advocates often frame Bitcoin as "digital gold" due to its potential hedging properties. In a similar vein, Liu and Tsyvinski (2021) find limited evidence of systematic exposure of crypto assets to market risk. Using the proposed candlestick-based spot estimator, I provide a more granular approach to this empirical question, offering deeper insights into the risk characteristics of crypto assets.

Specifically, I use 1-minute candlestick observations of the iShares Bitcoin Trust ETF (IBIT) and the SPDR S&P 500 ETF Trust (SPY) as proxies for Bitcoin and the overall market returns, respectively. I collect data for the full year of 2024 to construct spot covariance estimates at a 10-minute frequency.[4] From these, I compute candlestick-based spot beta estimates, defined as the ratio of covariance between IBIT and SPY to variance of SPY. I then test the null hypothesis of market neutrality (zero beta) using the test statistics developed in this paper. The null is rejected at the 5% significance level approximately 35% of the time, revealing very different pricing dynamics in the Bitcoin market compared to prior analyses. I also find that rejection rates exceed 50% in August and September 2024, coinciding with heightened financial market turmoil, precisely when the diversification and hedging benefits are most needed.

Along with this analysis, I conduct an event study around two FOMC announcements on June 12 and September 18, 2024. In September, the Federal Reserve cut interest rates by 50 basis points which was the first reduction since the 2022 tightening cycle. Thus, this decision is perceived as a strong dovish signal. Both return-based and candlestick-based estimates reveal that IBIT's spot beta exhibits no clear pattern before the announcement but increases

[3] In 2017, Bitcoin futures were introduced by the Chicago Mercantile Exchange (CME). Later, several cryptocurrency ETFs were launched in January 2024, allowing the investors to gain exposure to different cryptocurrencies such as Bitcoin.

[4] IBIT is a relatively new financial product that has been traded on the NASDAQ since January 2024.

sharply afterward, reaching approximately 1.5 and remaining statistically significant through the end of the trading day. By contrast, during the June meeting, these two methods lead to different conclusions. In this meeting, the Federal Reserve kept rates unchanged as expected but released a relatively hawkish dot plot which suggests less clear signals about future policy paths. Consequently, my results show that candlestick-based estimates reject the null that beta is zero immediately after the announcement, while return-based estimates fail to do so. This finding highlights the ability of the candlestick-based test to detect shifts in response to less informative events.

In an attempt to improve the estimation, one could alternatively employ tick-level data which provides the most granular information, as it records every single transaction in the market usually at ultra-high frequencies (e.g., milliseconds). However, using tick-level data faces important limitations. First, this data is accessible only to well-resourced researchers since it requires costly subscriptions through commercial providers such as NYSE Trade and Quotes (TAQ) or TickData. Furthermore, prices sampled at such ultra high-frequencies are invariably contaminated by market microstructure noise, which necessitates imposing additional modeling assumptions on the noise structure to obtain enhanced estimates, see for example Diebold and Strasser (2013). By contrast, candlestick data is widely accessible through many different public sources at "not-too-fine" frequencies (e.g., 1-minute or 5-minute), which naturally guards against the impact of microstructure noise. From this perspective, my paper adopts a more practical and accessible approach for improving spot covariance estimation.[5]

This paper contributes to multiple strands of the literature. From a technical perspective, this paper is closely related to the high-frequency econometrics literature on spot and integrated covariance estimation, see for example Barndorff-Nielsen and Shephard (2004a); Fan and Wang (2008); Barndorff-Nielsen et al. (2009) among many others. These papers develop estimators primarily based on high-frequency returns from open and close prices. This work complements this literature by introducing a new class of estimators for covariances

[5]Within the high-frequency econometrics literature, researchers have developed models for market microstructure noise. Notable references include Hayashi and Yoshida (2005); Aït-Sahalia et al. (2010); Christensen et al. (2010); Barndorff-Nielsen et al. (2011). These models can potentially be integrated into my framework. However, this is beyond the scope of this paper and left for future research.

that leverage candlestick observations, which expands the information set of high-frequency intervals with high and low prices. Moreover, these papers have primarily considered that the number of high-frequency observations growing asymptotically. However, my approach assumes a fixed estimation window following the recent work of Bollerslev et al. (2021).

My work is inspired by the range-based volatility estimation literature. Starting from seminal papers by Garman and Klass (1980) and Parkinson (1980), this literature highlights remarkable efficiency gains in estimating variances by extracting more information from candlestick prices. Notably, Christensen and Podolskij (2007) introduces the realized-range estimator for integrated variance, constructed by high-frequency ranges. Later, it is extended to be robust to microstructure noise (Martens and Van Dijk (2007); Christensen et al. (2009)); jumps (Christensen and Podolskij (2012)) and drifts (Li et al. (2025)). More recently, several papers study spot volatility estimation using candlesticks, see for example Li et al. (2024), Bollerslev et al. (2024a) and Bollerslev et al. (2025). The main focus of these papers is to derive optimal estimation and inference frameworks for spot volatility and its functionals. This paper complements this literature by extending these ideas to multivariate settings.

This paper is perhaps most closely related to Bollerslev et al. (2024a) which develops a decision-theoretic framework to construct optimal estimators for volatility functionals. In their approach, the risk function is asymptotically approximated using pivotal random variables, and the asymptotic risk is minimized through Monte Carlo simulations. However, in multivariate settings, the limiting distributions are generally non-pivotal, which makes direct minimization of a traditional risk function infeasible. To overcome this challenge, I integrate the asymptotic risk to obtain an average risk function which is pivotal, allowing me to determine the optimal estimators by minimizing this objective.

Prior to this study, Brandt and Diebold (2006) and Bannouh et al. (2009) introduced range-based covariance estimators that exploit triangular no-arbitrage conditions in currency markets. While effective in that setting, this is generally not applicable to equity markets. Similarly, Rogers and Zhou (2008) estimated the correlation coefficient of a bivariate Brownian motion using open, high, low and close price observations. However, their analysis is conducted under a constant volatility framework. In contrast, this paper accommodates a more general Itô semimartingale environment and provides a feasible inference procedure

with candlesticks which is not considered in the abovementioned papers.

The remainder of the paper is organized as follows. Section 2 introduces the candlestick-based estimation methodology. Section 3 develops the inference procedure. Section 4 presents simulation results, and Section 5 provides empirical applications. Section 6 concludes. Additional technical details and all proofs are provided in the appendices.

2 Spot Covariance Estimation with Candlesticks

Section 2.1 introduces the underlying price process. Section 2.2 presents the observation scheme and defines the candlestick returns. Section 2.3 proposes a class of new estimators for spot covariance. Then, Section 2.4 solves for the optimal estimator in the proposed class. Finally, Section 2.5 discusses the implications.

2.1 Price Process

The log-price vector at time t is denoted by $\boldsymbol{X}_t = [X_{1,t}, ..., X_{N,t}]^\top$. Assume that $\boldsymbol{X}_t$ follows an Itô semimartingale process defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$, given by

$$d\boldsymbol{X}_t = \boldsymbol{b}_t dt + \boldsymbol{\sigma}_t d\boldsymbol{W}_t, \tag{1}$$

where $\boldsymbol{b}_t$ is the drift process, $\boldsymbol{\sigma}_t$ is $N \times N$ stochastic volatility matrix and $\boldsymbol{W}_t$ is an N-dimensional vector of independent Brownian motions. The spot covariance at time t is defined as $\boldsymbol{c}_t \equiv \boldsymbol{\sigma}_t \boldsymbol{\sigma}_t^\top$, where $\boldsymbol{c}_t$ is an $N \times N$ matrix.

The Itô-semimartingale representation in Equation (1) can be motivated by no-arbitrage conditions and therefore serves as a workhorse framework in the continuous-time finance literature, see for example Back (2010). Consequently, it has become the fundamental model for analyzing high-frequency asset prices; see Jacod and Protter (2012) and Aït-Sahalia and Jacod (2014) for further discussions. For simplicity, I exclude jumps from the price process.[6]

[6]The estimators discussed in this paper are generally robust to Poisson-type jumps as they occur at vanishing probability, see Theorem 13.3.3 Jacod and Protter (2012). Nevertheless, if jumps are a major concern, one can employ truncation (Mancini (2009)) or bi-power variation (Barndorff-Nielsen and Shephard (2004b)) methods to explicitly account for price discontinuities. I do not pursue this direction to present the novelty of the paper with minimal technical complexity.

2.2 Observation Scheme and Candlesticks

Suppose that the price process $\boldsymbol{X}_t$ is sampled on a regular time grid, $\{i\Delta_n : i = 0, 1, ..., n\}$ over a fixed time span $[0, T]$. Here, $\Delta_n = T/n$ refers to the sampling frequency and n is the number of observations, assumed to be an integer. High-frequency intervals are denoted by $\mathcal{T}_i \equiv [(i-1)\Delta_n, i\Delta_n]$ for each $i \in \{1, ..., n\}$. Following the standard practice in the high-frequency financial econometrics literature, I consider an in-fill asymptotic framework where $\Delta_n \to 0$ asymptotically, see for example Aït-Sahalia and Jacod (2014).[7]

A typical candlestick over the interval $\mathcal{T}_i$ consists of four observed prices:

$$\boldsymbol{X}_{(i-1)\Delta_n}, \quad \sup_{t\in\mathcal{T}_i} \boldsymbol{X}_t, \quad \inf_{t\in\mathcal{T}_i} \boldsymbol{X}_t, \quad \boldsymbol{X}_{i\Delta_n}$$

which are called the open, high, low, and close prices, respectively. From these prices, one can construct the following normalized variables:

$$\begin{aligned} \boldsymbol{r}_i &\equiv \frac{\boldsymbol{X}_{i\Delta_n} - \boldsymbol{X}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}}, \\ \boldsymbol{h}_i &\equiv \frac{\sup_{t\in\mathcal{T}_i} \boldsymbol{X}_t - \boldsymbol{X}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}}, \\ \boldsymbol{\ell}_i &\equiv \frac{\inf_{t\in\mathcal{T}_i} \boldsymbol{X}_t - \boldsymbol{X}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \end{aligned} \tag{2}$$

where sup and inf operators are applied element-wise. The first line defines the standard high-frequency (open-to-close) return, commonly employed in the literature. The second and third lines indicate the high-open and low-open returns, respectively. These returns stand as a new source of information in this framework. Thus, the bundle $(\boldsymbol{r}_i, \boldsymbol{h}_i, \boldsymbol{\ell}_i)$ summarizes the price dynamics within the interval $\mathcal{T}_{n,i}$ through the lens of candlesticks.

Following Bollerslev et al. (2024a), I also define the range and asymmetry variables:

$$\begin{aligned} \boldsymbol{w}_i &\equiv \boldsymbol{h}_i - \boldsymbol{\ell}_i, \\ \boldsymbol{a}_i &\equiv \boldsymbol{h}_i + \boldsymbol{\ell}_i - \boldsymbol{r}_i. \end{aligned} \tag{3}$$

[7] The choice of Δ_n is usually guided by volatility signature plots introduced in Andersen et al. (2000). Following the standard practice in high-frequency econometrics literature, I adopt moderate sampling frequencies like $\Delta_n = 1, 5, 10$-min in my practical implementations, thereby avoiding ultra high-frequency observations that could be contaminated by microstructure effects.

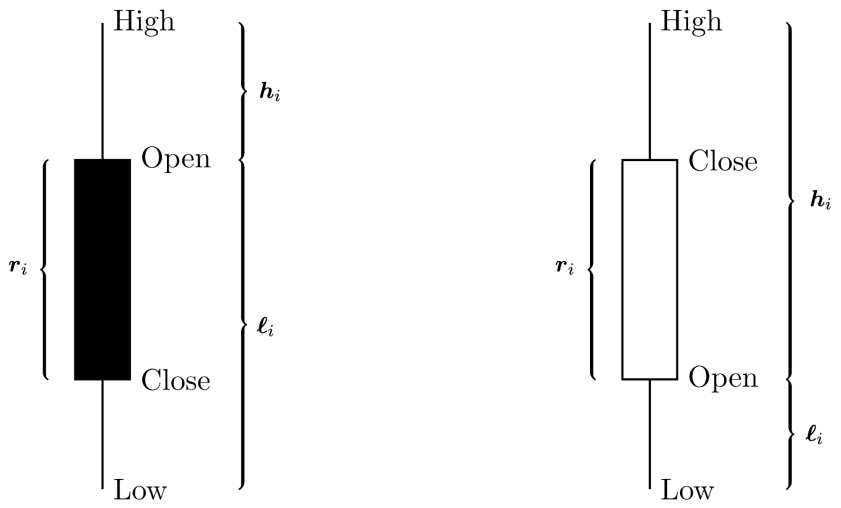

Figure 1: **Illustration of candlestick returns:** The figure shows examples of bearish (left panel) and bullish (right panel) candlesticks, illustrating the open, high, low, and close prices. Also it highlights the open-close return $\boldsymbol{r}_i$, high return $\boldsymbol{h}_i$, and low return $\boldsymbol{\ell}_i$ as defined in Equation (2).

Figure 1 presents the graphical representation of the candlestick returns on a typical candlestick chart. Looking at this figure, the range $\boldsymbol{w}_i$ is shown as the vertical distance between the high and low prices, while the return $\boldsymbol{r}_i$ reflects the length of the thick body. The asymmetry $\boldsymbol{a}_i$ indicates the position of the returns (or the thick body) within that range. For example, if $\boldsymbol{a}_i = 0$ then the thick body of the candlestick is exactly centered between the high and low prices. On the other hand, if $\boldsymbol{a}_i > 0$ then the body is skewed towards the low price, and vice versa.

One can show that $(\boldsymbol{r}_i, \boldsymbol{a}_i, \boldsymbol{w}_i)$ is a rotation (i.e., a linear transformation) of $(\boldsymbol{r}_i, \boldsymbol{h}_i, \boldsymbol{\ell}_i)$, implying that both vectors carry the same information. In this paper, I primarily work with the rotated set $(\boldsymbol{r}_i, \boldsymbol{a}_i, \boldsymbol{w}_i)$ to facilitate the analysis, but the proposed methodology can be easily adapted to the original set.

2.3 Candlestick-based Estimator

The spot covariance matrix $\boldsymbol{c}_t$ can be simply estimated as the local average of the outer products of high-frequency (open-to-close) returns. For any $t \in [0, T]$, this can be formally

expressed as:

$$\frac{1}{k}\sum_{i\in\mathcal{I}_{n,t}} \boldsymbol{r}_i\boldsymbol{r}_i^\top. \tag{4}$$

where $\mathcal{I}_{n,t} = \{\lceil \frac{t}{\Delta_n}\rceil + 1, \ldots, \lceil \frac{t}{\Delta_n}\rceil + k\}$ is the local estimation window and k stands for the size of that window.[8] This estimator represents a widely adopted method in the literature; see, among others, Fan and Wang (2008).[9]

The estimator in Equation (4) relies solely on open-to-close returns and therefore ignores potentially valuable information embedded in other candlestick returns. Thus, I introduce a flexible class of estimators for $\boldsymbol{c}_t$ that incorporates all candlestick returns $(\boldsymbol{r}_i, \boldsymbol{a}_i, \boldsymbol{w}_i)$ and combines them through a weighted sum of quadratic forms over a local window. Formally, the estimator is defined as:

$$\widehat{\boldsymbol{c}}_{n,t}(\lambda) = \frac{1}{k}\sum_{i\in\mathcal{I}_{n,t}} \left\{\lambda_1 \boldsymbol{r}_i\boldsymbol{r}_i^\top + \lambda_2 \boldsymbol{a}_i\boldsymbol{a}_i^\top + \lambda_3 \boldsymbol{w}_i\boldsymbol{w}_i^\top\right\} \tag{5}$$

where $\lambda \equiv (\lambda_1, \lambda_2, \lambda_3) \in \mathbb{R}^3$ denote the weights assigned to each component.[10] The resulting estimator is always symmetric and positive semi-definite whenever $\lambda_j \geq 0$ for all $j \in \{1, 2, 3\}$. The optimal choice of λ will be determined in subsequent sections according to specific optimality criteria. Importantly, I assume that the size of the estimation window, denoted by k, is fixed.[11]

The class of estimators defined in Equation (5) is highly flexible. For example, setting $\lambda = (1, 0, 0)$ recovers the standard spot covariance estimator in Equation (4). Moreover,

[8] Here, $\lceil\cdot\rceil$ denotes the ceiling function, which maps a real number to the smallest following integer. Thus, $\mathcal{I}_{n,t}$ can be interpreted as the right-sided local window that contains k observations after time t. Alternatively, left-sided or symmetric windows can be considered without changing the main results of the paper.

[9] The form of the estimator in Equation (4) is slightly different from that of Fan and Wang (2008). The standard approach is to let $k = k_n$ grow with n. However, I consider k as a constant number, not dependent on n.

[10] One could also include cross-product terms such as $\boldsymbol{r}_i\boldsymbol{a}_i^\top$. Within the framework of this paper, the optimal weights on such terms turn out to be zero. This is because this set of returns $(\boldsymbol{r}_i, \boldsymbol{a}_i, \boldsymbol{w}_i)$ are asymptotically orthogonal to each other. Hence, to simplify the notation, I do not include them in the estimator.

[11] This corresponds to the fixed-k asymptotic framework introduced by Bollerslev et al. (2021) and recently extended to candlestick-based spot volatility estimation and inference by Li et al. (2024). Fixed-k asymptotics build on coupling (or strong approximation) arguments; see Jacod et al. (2021). The key idea is that, over short intervals, an Itô semimartingale can be locally approximated by a scaled Brownian motion. This feature allows for determining the (approximate) finite-sample distribution of the estimator under diverging Δ_n. Bollerslev et al. (2021) show that fixed-k asymptotics deliver confidence intervals with more accurate coverage than conventional large-k approximations, especially in small samples, where the latter often suffer from nontrivial size distortions.

when $N = 1$, the estimator encompasses several well-known candlestick-based spot volatility estimators (e.g., Parkinson (1980), Garman and Klass (1980), Li et al. (2024)), which have been shown to deliver substantial efficiency gains, yielding more precise estimates and reliable inference.[12] From this perspective, the proposed estimator can be viewed as a natural multivariate extension of these univariate methods, thereby generalizing their efficiency benefits to the multivariate setting.

At a first glance, it might be surprising that moments of high and low returns are included in the covariance estimator just like the regular returns. However, there exist a tight connection between the moments of candlestick returns and the underlying correlation structure. To illustrate this, consider a price process $\boldsymbol{X}_t$ generated by a bivariate Brownian motion with unit volatilities and constant correlation coefficient ρ.[13] Denote the correlations between the candletick returns by ρ_r, ρ_a, ρ_w. Clearly, it holds that $\rho_r = \rho$. As shown by Rogers and Shepp (2006), the correlation of range and asymmetry variables (ρ_a, ρ_w) are smooth and nonlinear functions of ρ. Figure 2 illustrates this relationship. Notably, ρ_w is symmetric around zero, while ρ_a closely follows 45-degree line. This figure thus suggests that the cross moments of candlestick returns may contain valuable information about the underlying correlation structure.

2.4 Finding the "Optimal" Weights

The primary objective of this paper is to construct an estimator that efficiently exploits the information contained in candlestick features. To this end, the optimal weighting scheme λ is determined by minimizing a suitably defined risk function, which serves to quantify the estimation error. In what follows, I first provide a detailed discussion of the loss criterion underlying this risk function.

[12] For instance, Garman and Klass (1980) propose the following estimator for spot variance, among others: $\widehat{c}_{n,t} = -0.3925 r_i^2 + 0.0095 a_i^2 + 0.5015 w_i^2$ which implies $\lambda = (-0.3925, 0.0095, 0.5015)$.

[13] This special case considers $N = 2$ assets and implies following specifications for Ito-semimartingale process in Equation (1): $\boldsymbol{b}_t = 0$ and $\boldsymbol{c}_t$ is unit diagonal matrix with off-diagonal elements given by ρ.

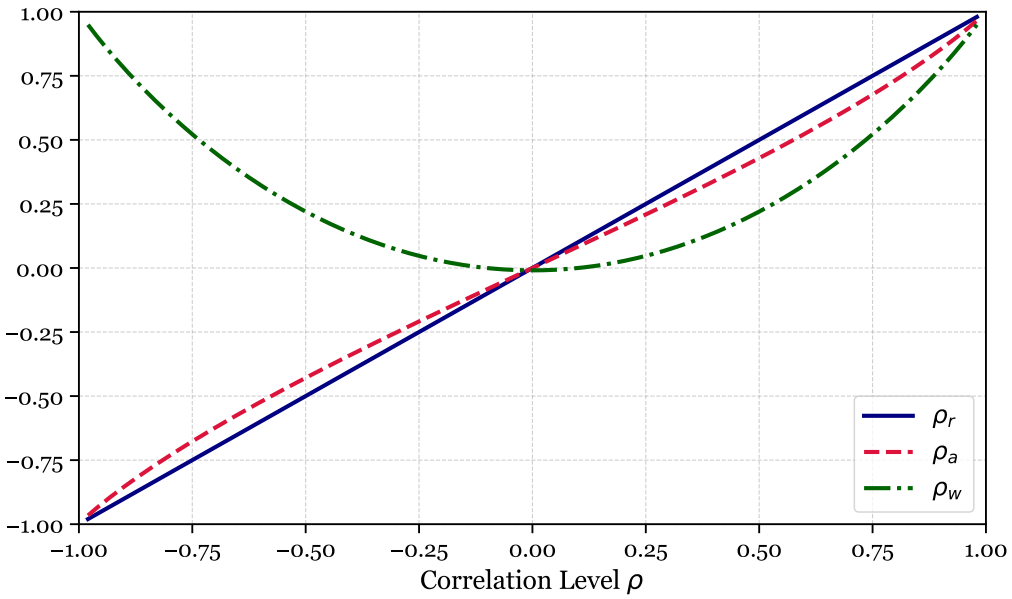

Figure 2: **Cross moments of candlestick returns:** The figure illustrates the relationship between the correlation of open-close returns ρ_r (blue) and that of asymmetry ρ_a (red) and range ρ_w (green) variables, respectively. The underlying price process is assumed to be a bivariate Brownian motion with unit volatilities and constant correlation coefficient ρ.

2.4.1 Loss Function

I adopt the following quadratic loss function:

$$L(\lambda; \boldsymbol{c}_t) = \left\| \boldsymbol{\sigma}_t^{-1} \widehat{\boldsymbol{c}}_{n,t}(\lambda) \boldsymbol{\sigma}_t^{-1\top} - \mathbf{I} \right\|^2 \tag{6}$$

where $||\cdot||$ denotes the Frobenius norm and $\mathbf{I}$ is the identity matrix. Recall that $\boldsymbol{c}_t$ and $\boldsymbol{\sigma}_t$ denote the true spot covariance and volatility matrices at a fixed time point t, respectively, and $\widehat{\boldsymbol{c}}_{n,t}$ is the estimator defined in Equation (5).

The term $\boldsymbol{\sigma}_t^{-1} \widehat{\boldsymbol{c}}_{n,t} \boldsymbol{\sigma}_t^{-1\top}$ captures the multiplicative estimation error of the spot covariance matrix relative to the identity, naturally inducing a scale-invariant loss function. This property is particularly desirable, as covariance matrices represent scale parameters. In the univariate setting ($N = 1$), the loss reduces to

$$\left(\frac{\widehat{c}_{n,t}}{c_t} - 1 \right)^2,$$

which corresponds to the relative squared error that has been previously analyzed in the

context of volatility estimation with candlestick data by Li et al. (2024).

Accordingly, one can naturally define the risk of the estimator $\widehat{\boldsymbol{c}}_{n,t}$ as the conditional expectation of the loss function:

$$R(\lambda; \boldsymbol{c}_t) \equiv \mathbb{E}\left[L(\lambda; \boldsymbol{c}_t) \mid \mathcal{F}_t\right]. \tag{7}$$

This quantifies the exact finite sample risk of an estimator. For a given λ, the risk depends on the joint distribution of $(\boldsymbol{r}_i, \boldsymbol{a}_i, \boldsymbol{w}_i)$ which in turn is determined by the law of $(\boldsymbol{b}, \boldsymbol{\sigma}, \boldsymbol{W})$. Since this law is unknown, the estimation of λ with direct minimization of $R(\lambda; \boldsymbol{\sigma}_t)$ is infeasible in practice.

To simplify the analysis, I first establish, in the next subsection, an asymptotic approximation of the multiplicative estimation error $\boldsymbol{\sigma}_t^{-1}\widehat{\boldsymbol{c}}_{n,t}\boldsymbol{\sigma}_t^{-1\top}$ using coupling arguments (see, e.g., Jacod et al. (2021)). This characterization is subsequently extended to the loss function and, in turn, to the risk. The resulting asymptotic risk functional is then used to determine the optimal weights $\boldsymbol{\Lambda}$.

2.4.2 Asymptotic Approximation for the Risk Function

The assumption below gathers a set of regularity conditions, required for the asymptotic approximation results.

Assumption 1. *Suppose that $\boldsymbol{X}_t$ has the form in Equation (1) and there exists a sequence $(T_m)_{m\geq 1}$ of stopping times increasing to infinity and the following conditions hold for each $m \geq 1$:*

(i) $\|\boldsymbol{b}_t\| + \|\boldsymbol{\sigma}_t\| + \|\boldsymbol{\sigma}_t^{-1}\| \leq K_m$ *for some constant* K_m *for all* $t \in [0, T_m]$*;*

(ii) $\mathbb{E}\left[\|\boldsymbol{\sigma}_{t\wedge T_m} - \boldsymbol{\sigma}_{s\wedge T_m}\|^2\right] \leq K_m|t-s|$ *for all* $t, s \in [0, T_m]$.

These conditions are quite standard in the high-frequency econometrics literature; see, for example, Jacod and Protter (2012) for further details. The first part of the Assumption 1 implies local boundedness of the drift, volatility, and inverse volatility processes while the second part ensures a degree of smoothness in the volatility process, specifically called as locally 1/2-Hölder continuous. These assumptions are sufficiently general to accommodate a wide

range of volatility dynamics, including volatility jumps, leverage effects, and intraday seasonality. For example, they are satisfied if the volatility process is itself an Itô semimartingale or long-memory process driven by fractional Brownian motion.

The following proposition describes the key approximation result for the estimation error $\boldsymbol{\sigma}_t^{-1}\widehat{\boldsymbol{c}}_{n,t}\boldsymbol{\sigma}_t^{-1\top}$. The proof is deferred to Appendix A.1.

Proposition 1. *Suppose that Assumption 1 holds. Fix any $t \in [0,T]$. For any $k \geq 1$ and $\boldsymbol{\Lambda}$, the following holds as $\Delta_n \to 0$:*

$$\left\| \boldsymbol{\sigma}_t^{-1}\widehat{\boldsymbol{c}}_{n,t}(\boldsymbol{\Lambda})\boldsymbol{\sigma}_t^{-\top} - U_{n,t}(\lambda) \right\| = o_p(1) \tag{8}$$

where $U_{n,t}(\lambda) = \frac{1}{k}\sum_{i\in\mathcal{I}_{n,t}} \left\{ \lambda_1 \boldsymbol{\zeta}_{i,r}\boldsymbol{\zeta}_{i,r}^\top + \lambda_2 \boldsymbol{\zeta}_{i,a}\boldsymbol{\zeta}_{i,a}^\top + \lambda_3 \boldsymbol{\zeta}_{i,w}\boldsymbol{\zeta}_{i,w}^\top \right\}$ and, for any $i \in \mathcal{I}_{n,t}$,

$$\begin{aligned}
\boldsymbol{\zeta}_{i,r} &\equiv \frac{\boldsymbol{W}_{i\Delta_n} - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \\
\boldsymbol{\zeta}_{i,a} &\equiv \boldsymbol{\varrho}_t^{-1} \sup_{\tau\in\mathcal{T}_i} \boldsymbol{\varrho}_t \left(\frac{\boldsymbol{W}_\tau - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) + \boldsymbol{\varrho}_t^{-1} \inf_{\tau\in\mathcal{T}_i} \boldsymbol{\varrho}_t \left(\frac{\boldsymbol{W}_\tau - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) - \left(\frac{\boldsymbol{W}_{i\Delta_n} - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) \\
\boldsymbol{\zeta}_{i,w} &\equiv \boldsymbol{\varrho}_t^{-1} \sup_{\tau\in\mathcal{T}_i} \boldsymbol{\varrho}_t \left(\frac{\boldsymbol{W}_\tau - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) - \boldsymbol{\varrho}_t^{-1} \inf_{\tau\in\mathcal{T}_i} \boldsymbol{\varrho}_t \left(\frac{\boldsymbol{W}_\tau - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right)
\end{aligned} \tag{9}$$

with $\boldsymbol{\varrho}_t$ being the square root of spot correlation matrix $\boldsymbol{\rho}_t$, i.e., $\boldsymbol{\rho}_t = \boldsymbol{\varrho}_t\boldsymbol{\varrho}_t^\top$.[14] *All* inf *and* sup *operators are applied element-wise.*

Proposition 1 establishes that the multiplicative estimation error of the spot covariance estimator $\widehat{\boldsymbol{c}}_{n,t}(\lambda)$ is asymptotically approximated by the random matrix $U_{n,t}$ in probability as $\Delta_n \to 0$. The structure of $U_{n,t}$ mirrors that of the original estimator, with the candlestick returns replaced by the variables $\boldsymbol{\zeta}_i$s. These variables are functions of Brownian motion $\boldsymbol{W}$ and square root of spot correlation matrix $\boldsymbol{\varrho}_t$. In addition, the approximation error ($o_p(1)$ term) captures nonparametric biases arising from stochastic volatility and drift. Put differently, in the limiting case of constant volatility and vanishing drift, the relationship in Proposition 1 holds exactly rather than approximately.

Using Proposition 1 and the conditions stated therein, analogous approximations apply to the loss and risk functions. Since the loss function is continuous, the continuous mapping

[14] In explicit terms, $\boldsymbol{\rho}_t = \mathrm{diag}(\boldsymbol{c}_t)^{-\frac{1}{2}}\boldsymbol{c}_t\,\mathrm{diag}(\boldsymbol{c}_t)^{-\frac{1}{2}}$ and $\boldsymbol{\varrho}_t = \mathrm{diag}(\boldsymbol{c}_t)^{-\frac{1}{2}}\boldsymbol{\sigma}_t$ where $\mathrm{diag}(\boldsymbol{c}_t)$ is a diagonal matrix with the same diagonal elements as $\boldsymbol{c}_t$.

theorem implies

$$L(\lambda; \boldsymbol{c}_t) = \|U_{n,t}(\lambda) - \mathbf{I}\|^2 + o_p(1). \tag{10}$$

This relation can similarly be extended to the risk function by taking the conditional expectation of both sides given $\mathcal{F}_t$:

$$R(\lambda; \boldsymbol{c}_t) = \underbrace{\mathbb{E}\left[\|U_{n,t}(\lambda) - \boldsymbol{I}\|^2 \mid \mathcal{F}_t\right]}_{\widetilde{R}(\lambda; \boldsymbol{\rho}_t)} + o_p(1) \tag{11}$$

where $\widetilde{R}(\lambda; \boldsymbol{\rho}_t)$ is the asymptotic risk of the estimator $\widehat{\boldsymbol{c}}_{n,t}$.[15] Importantly, I switch to notation $\widetilde{R}(\lambda; \boldsymbol{\rho}_t)$ to emphasize that the asymptotic risk depends on the spot correlation matrix $\boldsymbol{\rho}_t$ rather than the spot covariance matrix $\boldsymbol{c}_t$. This is because the risk formula involves a term, $U_{n,t}$, which depends on $\boldsymbol{\varrho}_t$, i.e., the square root of the spot correlation matrix. The use of the Frobenius norm in the final calculation effectively considers the product $\boldsymbol{\varrho}_t \boldsymbol{\varrho}_t^\top$, making the risk solely dependent on $\boldsymbol{\rho}_t$.

The above expressions indicate that the asymptotic risk of an estimator is determined by the $\mathcal{F}_t$-conditional distribution of the term $U_{n,t}$. In the univariate case, i.e. $N = 1$, this conditional distribution is pivotal, see Theorem 1 in Li et al. (2024). In contrast, this does not hold in the multivariate setting, as $U_{n,t}$ depends on $\boldsymbol{\rho}_t$, a population quantity that is unknown to the researcher.

To build intuition, consider the structure of $U_{n,t}$, which is derived from the variables $(\boldsymbol{\zeta}_{i,r}, \boldsymbol{\zeta}_{i,a}, \boldsymbol{\zeta}_{i,w})$. By the scaling property of Brownian motion, $\boldsymbol{\zeta}_{i,r}$ is standard normally distributed and therefore naturally pivotal. However, the remaining variables, $\boldsymbol{\zeta}_{i,a}$ and $\boldsymbol{\zeta}_{i,w}$, depend on $\boldsymbol{\varrho}_t$ in a nontrivial manner. This dependence arises because the supremum and infimum operators are nonlinear, preventing the scaling by $\boldsymbol{\varrho}_t$ and its inverse from canceling. Consequently, the distribution of $U_{n,t}$ is generally not free from $\boldsymbol{\varrho}_t$. In the univariate case, by contrast, the terms $\boldsymbol{\zeta}_{i,a}$ and $\boldsymbol{\zeta}_{i,w}$ reduce to functions of standard Brownian motion increments. This is because the scaling by $\boldsymbol{\varrho}_t$ and its inverse cancels, as they become scalars, allowing the distribution of these terms to be fully characterized through simulations of Brownian motion functionals. Consequently, the asymptotic risk can be computed directly from this simulated

[15] This step requires an additional uniform integrability condition to interchange the limit and expectation operators. This condition is satisfied under the assumptions stated in Assumption 1.

distribution. Such pivotality considerably simplifies the derivation of optimal estimators via risk minimization; see Li et al. (2024) and Bollerslev et al. (2024a) for further discussion. In this sense, the multivariate volatility estimation problem studied here departs from the standard univariate frameworks and requires a more intricate treatment.

2.4.3 Average Risk Minimization

To address the non-pivotal nature of the asymptotic risk in the multivariate setting, I propose marginalizing over the unobserved quantities and working with the resulting *average risk*. This approach is directly inspired by the classical decision theory literature, see for example Lehmann and Casella (2006), which advocates the use of integrated risk functions to handle nuisance parameters.

Let $\mathcal{P}$ denote the parameter space for positive semidefinite correlation matrices.[16] Then define the average risk $\bar{R}(\lambda)$ by integrating $\widetilde{R}$ over $\mathcal{P}$:

$$\bar{R}(\lambda) \equiv \int_{\mathcal{P}} \widetilde{R}(\lambda; \boldsymbol{\rho}) d\boldsymbol{\rho}. \tag{12}$$

For any fixed λ, the mapping $\boldsymbol{\rho} \mapsto \widetilde{R}(\lambda; \boldsymbol{\rho})$ can be computed via Monte Carlo simulation. This in turn allows for the determination of the weights λ by solving the following optimization problem:

$$\lambda^* = \arg\min_{\lambda} \bar{R}(\lambda). \tag{13}$$

The solution to this optimization problem is available in closed form. To see this, rewrite the explicit form of the average risk $\bar{R}(\lambda)$ by applying the half-vectorization operator:

$$\int_{\mathcal{P}} \mathbb{E}\left[\left\| \frac{1}{k} \sum_{i \in \mathcal{I}_{n,t}} \left\{ \lambda_1 \operatorname{vech}(\boldsymbol{\zeta}_{i,r} \boldsymbol{\zeta}_{i,r}^\top) + \lambda_2 \operatorname{vech}(\boldsymbol{\zeta}_{i,h} \boldsymbol{\zeta}_{i,a}^\top) + \lambda_3 \operatorname{vech}(\boldsymbol{\zeta}_{i,w} \boldsymbol{\zeta}_{i,w}^\top) \right\} - \operatorname{vech}(\mathbf{I}) \right\|^2 | \mathcal{F}_t \right] d\boldsymbol{\varrho} \tag{14}$$

[16] When $N = 2$, $\mathcal{P}$ is simply all possible 2×2 correlation matrices.

Next, stack the vectorized terms into a single matrix:

$$\mathbf{\Pi} \equiv \begin{bmatrix} \vdots & \vdots & \vdots \\ \frac{1}{k}\sum_i \text{vech}(\boldsymbol{\zeta}_{i,r}\boldsymbol{\zeta}_{i,r}^\top) & \frac{1}{k}\sum_i \text{vech}(\boldsymbol{\zeta}_{i,a}\boldsymbol{\zeta}_{i,a}^\top) & \frac{1}{k}\sum_i \text{vech}(\boldsymbol{\zeta}_{i,w}\boldsymbol{\zeta}_{i,w}^\top) \\ \vdots & \vdots & \vdots \end{bmatrix} \tag{15}$$

where $\mathbf{\Pi}$ is a $\frac{N(N+1)}{2} \times 3$ matrix. Then, the average risk may be expressed as:

$$\bar{R}(\lambda) = \int_{\mathcal{P}} \mathbb{E}\left[\left\|\mathbf{\Pi}\,\lambda - \boldsymbol{y}\right\|^2 | \mathcal{F}_t\right] d\boldsymbol{\rho} \tag{16}$$

in which $\boldsymbol{y} = \text{vech}(\mathbf{I})$. This representation resembles a standard least-squares estimation problem, allowing for a closed-form solution. Consequently,

$$\lambda^* = \left(\int_{\mathcal{P}} \mathbb{E}[\mathbf{\Pi}^\top \mathbf{\Pi} | \mathcal{F}_t] d\boldsymbol{\rho}\right)^{-1} \left(\int_{\mathcal{P}} \mathbb{E}[\mathbf{\Pi}^\top \boldsymbol{y} | \mathcal{F}_t] d\boldsymbol{\rho}\right). \tag{17}$$

Finally, the numerical computation of λ^* proceeds as follows:

Step 1. Uniformly draw $\boldsymbol{\rho}$ from the parameter space $\mathcal{P}$.[17]

Step 2. For the given $\boldsymbol{\rho}$, compute the conditional expectation $\mathbb{E}[\cdot \mid \boldsymbol{\rho}]$ using Monte Carlo simulation.

Step 3. Repeat Steps 1–2 over a large number of draws and average the resulting estimates to approximate $\bar{R}(\lambda)$.

Step 4. Plug the resulting approximation of $\bar{R}(\lambda)$ into Equation (13) and solve for λ^*.

The above numerical procedure applies for all values of k, estimation window size, and N, number of assets. For clarity, I focus on the case of $N = 2$ in my practical implementations. The optimal weights λ^* for different local window sizes $k \in \{5, 10, 20\}$ are given by in Equation

[17] For $N = 2$, this corresponds to drawing the correlation coefficient ρ uniformly from $(-1, 1)$. When $N > 2$, one can use the distribution developed by Lewandowski et al. (2009) which allows for uniform sampling of correlation matrices. The implementation of this procedure is available in the programming languages like R and Python.

(18):

$$\lambda^* = \begin{cases} \big(0.4106,\ 1.4550,\ 0.0013\big)^\top & \text{if } k = 5, \\ \big(0.4725,\ 1.6280,\ 0.0002\big)^\top & \text{if } k = 10, \\ \big(0.5176,\ 1.7039,\ 0.0001\big)^\top & \text{if } k = 20. \end{cases} \tag{18}$$

2.5 Discussion

The resulting estimator $\widehat{\boldsymbol{c}}_{n,t}(\lambda^*)$ aligns with the broader motivation of candlestick-based estimators to be straightforward to implement, as the weights in the linear combination can be computed once and subsequently reused across different settings without the need for complex statistical or econometric procedures.

When defining the average risk, the integration over $\mathcal{P}$ is performed uniformly which also induces a uniform marginal distribution for the correlation parameter. This may appear simplistic, and alternative choices including more flexible distributional assumptions or those informed by the historical data and expert judgment could be certainly considered. Nevertheless, the asymptotic risk comparisons presented in Section 4.1 reveals that this approach performs nearly as well as oracle estimators, leaving limited potential improvement through alternative specifications. Moreover, solving this optimization problem for λ does not inherently depend on uniform integration. As such, the optimal weights can be obtained for any given marginal distribution of $\boldsymbol{\rho}$ following the same computational steps.

An estimator of this form, a linear combination of $\boldsymbol{r}_i\boldsymbol{r}_i^\top$ and $\boldsymbol{a}_i\boldsymbol{a}_i^\top$, was previously studied by Rogers and Zhou (2008) for the purpose of estimating the correlation between two Brownian motions. Their analysis assumes that the price process follows a scaled Brownian motion and determines the weights by minimizing the variance of the estimation error, subject to an unbiasedness constraint under the assumption of zero correlation. By contrast, my approach is formulated within a general Itô-semimartingale framework and selects the weights by minimizing the average risk of the estimation error without imposing any restrictions on the correlation structure. Moreover, an inference procedure is developed in the next section that enables hypothesis testing for spot covariances, whereas Rogers and Zhou (2008) focus exclusively on point estimation.

3 Spot Beta Estimation and Inference with Candlesticks

In this section, I study inference on spot covariances using the candlestick-based estimator introduced in the previous section. While prior work has focused on inference for spot volatilities using candlesticks (see, e.g., Li et al. (2024); Bollerslev et al. (2024a)), the problem of conducting inference on covariance terms remains unexplored. To address this gap, I propose a hypothesis testing procedure. In particular, I focus on spot betas, defined as the ratio of covariance to variance, which naturally arise as a by-product of spot covariance estimation. Focusing on betas is practically relevant as these quantities are widely employed in asset pricing and portfolio management to measure systematic risk.

3.1 Candlestick-Beta Estimator

To fix ideas, I focus on the $N = 2$ case and tailor the price process $\boldsymbol{X}_t = [X_{1,t}, X_{2,t}]^\top$ to the following regression representation:

$$\begin{aligned} dX_{1,t} &= \nu_t^{1/2} dW_{1,t} \\ dX_{2,t} &= \beta_t dX_{1,t} + \varsigma_t^{1/2} dW_{2,t}. \end{aligned} \tag{19}$$

where $\boldsymbol{W} = [W_{1,t}, W_{2,t}]^\top$ is a standard bivariate Brownian motion. This is also equivalent to assuming the following spot covariance structure:

$$\boldsymbol{c}_t = \begin{pmatrix} \nu_t & \beta_t \nu_t \\ \beta_t \nu_t & \beta_t^2 \nu_t + \varsigma_t \end{pmatrix} \quad \text{and} \quad \boldsymbol{\sigma}_t = \begin{pmatrix} \nu_t^{1/2} & 0 \\ \beta_t \nu_t^{1/2} & \varsigma_t^{1/2} \end{pmatrix}. \tag{20}$$

Through the lens of factor models in asset pricing literature (e.g., Sharpe (1964); Lintner (1965); Fama and French (1992)), one can consider the first asset as the market portfolio and the second asset as a risky asset. Then, ν_t and ς_t refer to the market and idiosyncratic variances, respectively. Finally, β_t indicates the market beta of the risky asset.

From the above, it is clear that $\beta_t = \frac{\boldsymbol{c}_{12,t}}{\boldsymbol{c}_{11,t}}$ and therefore, naturally suggesting the spot beta estimator:

$$\widehat{\beta}_{n,t}(\boldsymbol{\Lambda}) = \frac{[\widehat{\boldsymbol{c}}_{n,t}(\boldsymbol{\Lambda})]_{12}}{[\widehat{\boldsymbol{c}}_{n,t}(\boldsymbol{\Lambda})]_{11}}. \tag{21}$$

where the notation $[\cdot]_{ij}$ refers to the (i,j)-th entry of a matrix. This defines a candlestick-based spot beta estimator, hereafter referred to as the *candlestick-beta* estimator.

3.2 Testing on Spot Beta

I consider the following test statistics:

$$\widehat{T}_n = \frac{\sqrt{k-1}\left(\widehat{\beta}_{n,t} - \beta_t\right)}{\sqrt{\widehat{\varsigma}_{n,t}/\widehat{\nu}_{n,t}}}$$

where $\widehat{\nu}_{n,t} = [\widehat{\boldsymbol{c}}_{n,t}(\boldsymbol{\Lambda})]_{11}$ and $\widehat{\varsigma}_{n,t} = [\widehat{\boldsymbol{c}}_{n,t}(\boldsymbol{\Lambda})]_{22} - \frac{[\widehat{\boldsymbol{c}}_{n,t}(\boldsymbol{\Lambda})]_{12}^2}{[\widehat{\boldsymbol{c}}_{n,t}(\boldsymbol{\Lambda})]_{11}}$.

Similar to spot covariance estimation, one can establish an asymptotic approximation to $\widehat{T}_n$ using the coupling variable $U_{n,t}$ defined in Proposition 1. This result is summarized in the following proposition. Proof is provided in Appendix A.2.

Proposition 2. *Suppose that the conditions of Proposition 1 holds and $\boldsymbol{X}_t$ follows Equation (19). Then, for any fixed $k \geq 2$ and λ, the following holds as $\Delta_n \to 0$:*

$$|\widehat{T}_n - \widetilde{T}_n| = o_p(1) \quad \textit{where} \quad \widetilde{T}_n \equiv \frac{\sqrt{k-1}[U_{n,t}^{-1}]_{12}}{\sqrt{[U_{n,t}^{-1}]_{11}[U_{n,t}^{-1}]_{22} - [U_{n,t}^{-1}]_{12}^2}}.$$

This proposition shows that $\widehat{T}_n$ can be asymptotically approximated by $\widetilde{T}_n$, which is a function of the matrix $U_{n,t}$ defined in Proposition 1. Consequently, the distributional properties of the limiting variable $\widetilde{T}_n$ can be exploited to test the null hypothesis $\beta_t = \beta_0$ or to construct confidence intervals.

Accordingly, Bollerslev et al. (2024b) demonstrate that $\widehat{T}_n$ follows a Student's t distribution with $k-1$ degrees of freedom when the beta estimator is based solely on return observations, that is, when $\lambda = (1,0,0)^\top$, which makes inference straightforward. The key complication arises when the beta estimator incorporates the other candlestick returns, i.e., when $\lambda \neq (1,0,0)^\top$. In this case, as discussed in Section 2.4.2, $U_{n,t}$ and hence $\widetilde{T}_n$ depend on population quantities in a non-trivial manner, generally resulting in a non-pivotal distribution, thereby complicating inference.

Meanwhile, under the null hypothesis $H_0 : \beta_t = 0$, the limiting distribution of $\widetilde{T}_n$ becomes pivotal and hence allows for feasible inference. The reason is that under this null hypothesis,

the covariance matrix $\boldsymbol{c}_t$ defined in Equation (20) simplifies to a diagonal matrix and consequently the $\boldsymbol{\varrho}_t$ terms in the coupling returns provided in Equation (9) cancel out, leaving the limiting variable as a function of the Brownian motion functionals only. Therefore, the distribution of $\widetilde{T}_n$ can be characterized by Monte Carlo simulations.

Based on this simulated distribution, one can easily determine the critical values for the test. Specifically, for a given significance level α, I define constants B_α^+ and B_α^- such that:

$$\mathbb{P}\left(B_\alpha^- < \widetilde{T}_n < B_\alpha^+\right) = 1 - \alpha. \tag{22}$$

These constants serve as the lower and upper critical values for this test. By construction, this procedure delivers asymptotically correct size under the null hypothesis $H_0 : \beta_t = 0$.

Table 1 reports the critical values for various combinations of k and α, computed using 10,000 Monte Carlo simulations. Since infinitely many pairs can satisfy this condition, I select the highest density interval (HDI), which is the narrowest interval containing the desired probability mass. For comparison, I also report the critical values corresponding to the return-based beta estimator. These are derived analytically from the Student's t-distribution, following Bollerslev et al. (2024b). Finally, I present the interval widths, defined as $B_\alpha^+ - B_\alpha^-$, for both estimators to highlight efficiency gains achieved by the new candlestick-based method. From this table, it is evident that the critical values are tighter for the candlestick-based beta estimator compared to the return-based estimator, across all values of k and α. Specifically, the difference in interval widths becomes more pronounced as k decreases. This highlights the usefulness of my estimator, particularly in limited data scenarios.

4 Simulations

This section examines the performance of the candlestick-based spot covariance estimator and the associated inference procedure through a series of Monte Carlo experiments. The analysis proceeds in two parts. First, I investigate the efficiency of the proposed estimator by comparing its asymptotic risk against an oracle estimator and a return-based estimator. Second, I evaluate the power of the hypothesis test for spot betas introduced in Section 3.

Table 1: **Critical values for the beta estimators**: This table reports the critical values for the hypothesis test of $H_0 : \beta_t = 0$ at significance levels $\alpha \in \{5\%, 10\%\}$ and for different local window sizes $k \in \{5, 10, 20\}$. The left panel shows the critical values for the return-based beta estimator derived from the Student's t-distribution with $k-1$ degrees of freedom. The right panel presents the critical values for the candlestick-based beta estimator computed via Monte Carlo simulations. The interval width is defined as $B_\alpha^+ - B_\alpha^-$.

	Return			**Candlestick**		
$\boldsymbol{k}$	B_α^-	B_α^+	Width	B_α^-	B_α^+	Width
Panel A: $\alpha = 5\%$						
5	-2.776	2.776	5.552	-1.657	1.415	3.072
10	-2.262	2.262	4.524	-1.430	1.460	2.890
20	-2.093	2.093	4.186	-1.429	1.383	2.812
Panel B: $\alpha = 10\%$						
5	-2.132	2.132	4.264	-1.204	1.262	2.466
10	-1.383	1.383	2.766	-1.203	1.199	2.402
20	-1.328	1.328	2.656	-1.103	1.244	2.347

4.1 Asymptotic Risk of the Estimators

To assess the effectiveness of the new approach, I compare the asymptotic risk of the estimator against two natural benchmarks. The first is an infeasible oracle estimator that minimizes the asymptotic risk under knowledge of the true correlation structure. The second benchmark is the return-based estimator, which relies solely on (open-to-close) returns. This comparison highlights the efficiency gains achieved by incorporating the additional information contained in the candlestick features.

The asymptotic risk $\widetilde{R}(\lambda; \boldsymbol{\rho}_t)$ can be computed via Monte Carlo simulations for any λ and k if the correlation structure is known. As in the previous numerical implementations, I focus on the case of $N = 2$ assets. Table 2 reports the asymptotic risk values for various configurations. I consider 3 different levels of correlation, $\rho \in \{0, 0.2, 0.6\}$.[18] Moreover, the size of the local estimation window is set to $k \in \{5, 10, 20\}$. All results are based on 10,000 Monte Carlo simulations.

[18] This numbers represent the $10, 50, 90$ percentiles of cross-section of pairwise correlations among the S&P 500 stocks.

Table 2: **Asymptotic Risk of Estimators for Different ρ Values**: The table presents the asymptotic risk of three estimators: the return-based estimator, the candlestick-based estimator, and an oracle estimator that minimizes asymptotic risk with knowledge of the true correlation structure. The results are shown for various local window sizes $k \in \{5, 10, 20\}$ and correlation levels $\rho \in \{0, 0.2, 0.6\}$. The asymptotic risk is computed via Monte Carlo simulations.

		$\rho = 0$		$\rho = 0.2$		$\rho = 0.6$	
k	Return	Candlestick	Oracle	Candlestick	Oracle	Candlestick	Oracle
5	1.212	0.427	0.384	0.429	0.384	0.440	0.391
10	0.599	0.240	0.222	0.241	0.221	0.249	0.224
20	0.297	0.128	0.119	0.128	0.118	0.134	0.119

Table 2 provides several key insights. First, my estimator achieves asymptotic risk levels that are very close to those of the oracle estimator across all values of k and ρ, implying roughly $8-9\%$ efficiency loss. This indicates that the proposed approach effectively addresses complications arising from the distribution of the estimator being non-pivotal, resulting in an estimator that performs nearly as well as the infeasible oracle estimator. Second, the proposed estimator consistently demonstrates a significant reduction in asymptotic risk compared to the return-based estimator, particularly for smaller values of k. That is, for $k = 5$, the candlestick-based estimator's risk is roughly 65% lower than that of the return-based estimator, and this difference is about 55% for $k = 20$. This suggests that my approach effectively leverages the candlestick prices and yields more efficient estimators. With only $k = 5$ observations, the proposed estimator achieves a lower risk than the return-based estimator using $k = 10$ observations, demonstrating its superior efficiency in extracting information from limited data. This advantage is particularly valuable in high-frequency event studies, where identification often relies on short time intervals and data are inherently limited.

Overall, these results show that incorporating candlestick information substantially improves efficiency, providing a powerful and practical alternative to traditional return-based methods and approaching the performance of the infeasible oracle benchmark.

4.2 Power of the Test

Next, I evaluate the properties of the hypothesis test for spot betas discussed in Section 3. To this end, I consider the data-generating process (DGP) employed in Bollerslev et al. (2024b).[19] Specifically, this DGP assume a two-factor structure for the market variance process: $\nu_t = V_{1,t} + V_{2,t}$ where $V_{1,t}$ and $V_{2,t}$ follow the processes:

$$\begin{aligned} dV_{1,t} &= 0.0128\,(0.4068 - V_{1,t})\,dt + 0.0954\sqrt{V_{1,t}}\left(\gamma dW_{1,t} + \sqrt{1-\gamma^2}dB_{1,t}\right), \\ dV_{2,t} &= 0.6930\,(0.4068 - V_{2,t})\,dt + 0.7023\sqrt{V_{2,t}}\left(\gamma dW_{1,t} + \sqrt{1-\gamma^2}dB_{2,t}\right). \end{aligned} \tag{23}$$

Here, $B_{1,t}$ and $B_{2,t}$ are independent Brownian motions that are also independent of $W_{1,t}$ and $W_{2,t}$. The parameter γ is set to -0.7, capturing the well-documented leverage effect in financial markets. The coefficients are set so that the first volatility factor is highly persistent, showing 2.5 months half-life, while the second factor reverts to mean quickly, with a half-life of a day. This allows the DGP to capture both short-term fluctuations and long-term trends in volatility.

The idiosyncratic variance ς_t and the spot beta are assumed to follow:

$$\beta_t = 1 + 0.25\sin(t)^2, \quad \varsigma_t = \left(1.5 + 0.25\sin(t)^2\right)\nu_t \tag{24}$$

where β_t fluctuates between 1 and 1.25 over time and ς_t is set to be proportional to the market variance ν_t. Finally, the price processes are generated using the representation provided in Equation (19).

I simulate the price path using an Euler scheme on a one-second grid. Then I sample the prices at one-minute frequency to construct the candlestick data, implying $\Delta_n = 1/390$ and $n = 390$ intraday observations. This setup mimics the empirical applications involving high-frequency financial data. As in the previous section, I consider three different values for the local window size, $k \in \{5, 10, 20\}$. For each configuration, I conduct 10,000 Monte Carlo simulations to compute the empirical rejection rates at significance levels of $1\%, 5\%$, and 10%. The results are summarized in Table 3. I also report the rejection rates for the

[19] This DGP is built on the univariate setup originally proposed by Bollerslev and Todorov (2011). Later, it is implemented by a number of subsequent studies (e.g., Bollerslev et al. (2021); Li et al. (2024)).

Table 3: **Power of Tests Derived from Return and Candlestick Estimators (%):** The table reports the rejection rates (in percent) for the hypothesis test of $H_0 : \beta_t = 0$ at significance levels $\alpha \in \{1\%, 5\%, 10\%\}$ and for different local window sizes $k \in \{5, 10, 20\}$ based on simulated data. The left panel shows the rejection rates for the return-based beta estimator, while the right panel presents the rejection rates for the candlestick-based beta estimator.

	Return			**Candlestick**		
$\boldsymbol{k}$	1%	5%	10%	1%	5%	10%
5	8.84	28.02	42.11	31.49	58.82	65.67
10	32.53	59.98	72.66	60.46	84.81	91.22
20	73.99	90.32	94.80	94.95	99.15	99.44

return-based beta estimator for comparison.[20]

Table 3 reports the rejection rates (in percent) for the two competing inference procedures at significance levels of 1%, 5%, and 10%. Across all scenarios, the candlestick-based test consistently rejects false null hypotheses more frequently than the return-based test. This gap widens as the local window size decreases. For instance, when $k = 20$ and $\alpha = 5\%$, the difference in rejection rates is approximately 9%, increasing to 25% for $k = 10$ and $\alpha = 5\%$. These results demonstrate a substantial reduction in false negatives, highlighting the improved power of the candlestick-based procedure.

To further illustrate the power advantage of the proposed test, Figure 3 presents the power curves for both tests as a function of the true β_t value. In this analysis, I use the same DGP as before but set β_t to be constant, varying from 0 to 2. The null hypothesis is that $\beta_t = 0$, and the empirical rejection rates are computed at a 5% significance level using $10,000$ Monte Carlo simulations. The local window size is fixed at $k = 10$.

Figure 3 displays two panels. The top panel shows the rejection rates for both tests, while the bottom panel illustrates the difference in rejection rates between the candlestick-based and return-based tests. The x-axis shows how far β_t is from the null value of 0. Looking at the leftmost point where β_t is set to zero, the return-based test appears to correctly maintain the size at approximately 5%, while the candlestick-based test slightly over-rejects at around

[20]Note that the return-based test is developed by Bollerslev et al. (2024b). As discussed in Section 3, the corresponding test statistic is shown to be Student's t distributed.

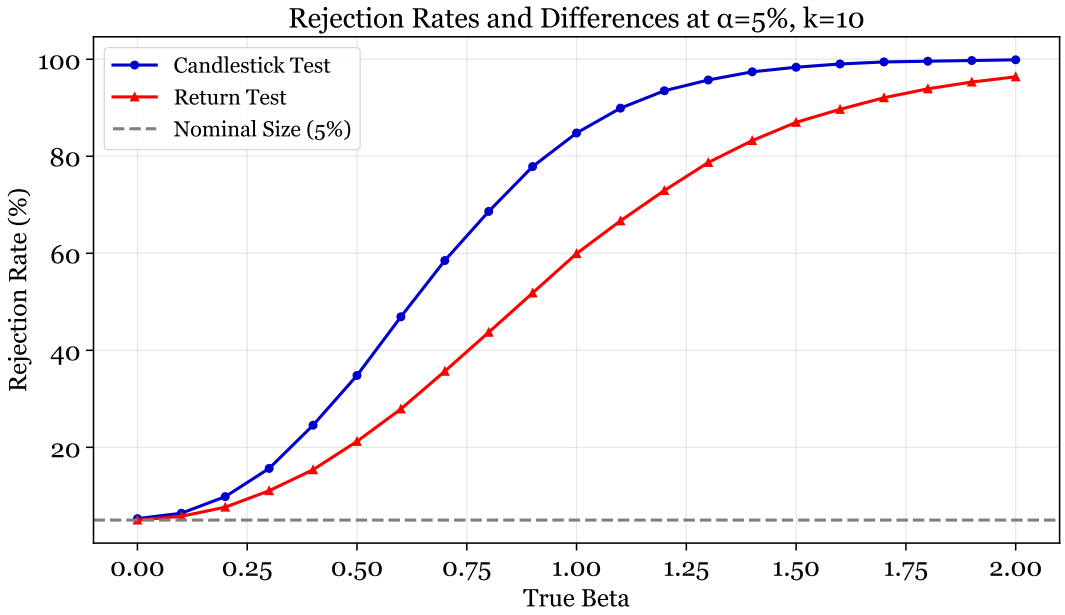

Figure 3: **Power Curves of Return-based and Candlestick-based Tests.** The figure plots the empirical power of the two tests as a function of the true β_t value (assumed to be constant), with $k = 10$ and a significance level of 5%. The power is computed using 10,000 Monte Carlo simulations for each β_t value.

6%. At the right-end where β_t is far from the null, unsurprisingly, both tests achieve a power close to 100%. As β_t deviates from zero, the power of both tests increases, with the candlestick-based test exhibiting a notably steeper increase. For example, when β_t is around 1, the candlestick-based test achieves a power of approximately 85%, while the return-based test is around 60%, indicating a 25% difference.

Overall, these results show that the candlestick-based inference for spot betas greatly improves power, especially when the true β_t is near the null, highlighting its practical value for more accurate inference on measures of the systematic risk.

5 Empirical Application

In this section, I apply the candlestick-beta estimator and the associated inference procedure to analyze Bitcoin's market exposure, i.e., its market beta. Given the growing role of crypto assets in institutional and retail portfolios, this empirical question has important implications

for risk management and portfolio selection.[21] Although crypto advocates often describe these assets as "digital gold", suggesting potential hedging benefits against aggregate market risk, empirical evidence on their risk characteristics remains limited.

To gain a deeper understanding of Bitcoin's market exposure, I estimate the spot beta and then test the null of market neutrality, i.e, zero beta, at a 5% significance level using the new candlestick-based framework proposed in Section 3. This analysis employs 1-minute price observations from two prominent ETFs: SPY and IBIT. While the former is a well-known ETF that tracks the S&P 500 index and commonly used in empirical studies, the latter is a newly launched (as of January 2024) iShares Bitcoin Trust ETF designed to track the performance of Bitcoin.[22] The sample spans the entire year 2024, covering 250 trading days and standard trading hours from $9:30$ to $16:00$. Spot beta estimates are computed using a local window of $k = 10$ one-minute candlesticks, yielding 39 spot estimates per trading day.

I start by assessing the monthly rejection rates (in terms of percentage) of the null hypothesis of market neutrality. The results are presented in Figure 4. The figure shows that the rejection rates are around 10% in the first two months of 2024, later increasing to 40% in mid-2024 and ending the year with a similar rate. Notably, the rejection rates are more pronounced in August and September, reaching around 60%. These months also coincide with a number of crucial economic events. Particularly, the first week of September was marked by a series of weak production and labor market data releases, raising concerns about a potential economic slowdown.[23]

Figure 5 presents the spot beta estimates and corresponding confidence intervals for September 3–6. This figure reveals that the null hypothesis of zero beta is rejected in a substantial portion of the intraday intervals, approximately 65% of the time. For instance, looking at the bottom right panel, the null is rejected in 27 out of 39 intervals on September

[21] Recent years have seen several developments that facilitated investment in crypto assets. In 2017, the Chicago Mercantile Exchange (CME) launched Bitcoin futures contracts, followed by the introduction of Bitcoin options in 2020. More recently, in January 2024, the U.S. Securities and Exchange Commission (SEC) approved the first Bitcoin exchange-traded fund (ETF).

[22] Cryptocurrency ETFs are designed to provide investors with exposure to the price movements in crypto markets. The IBIT has been the most traded one since its launch and its net asset value exceeds $70 billion as of 2025.

[23] The first week of September brought a sequence of weak data: the ISM manufacturing index on the 3rd, the ADP jobs report on the 5th, and the non-farm payrolls on the 6th, all of which came in below expectations. These developments led to a sharp sell-off in equities and a spike in volatility, as reflected in the VIX index.

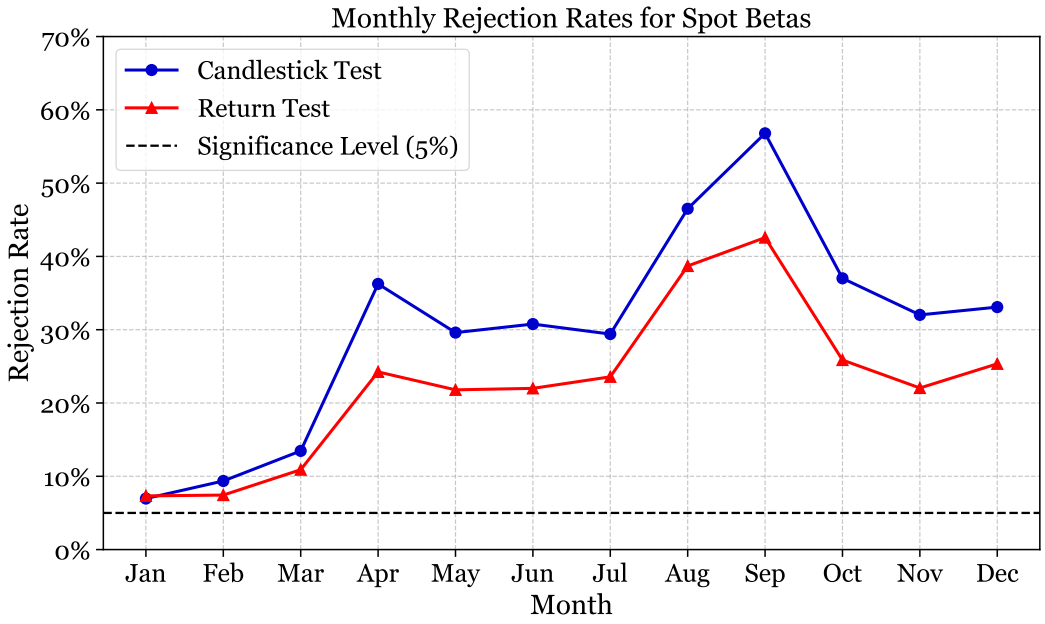

Figure 4: **Monthly Rejection Rates of the Null Hypothesis of Zero Beta for IBIT.** The figure shows the monthly rejection rates of the null hypothesis of zero beta for IBIT using candlestick-based inference framework. The sample covers the entire 2024 year, a total of 250 days, and usual trading hours from 9 : 30 to 16 : 00.

6. On that day, the NFP report was released at 8:30 AM, prior to market open, and spot beta estimates were already significant and around 1.5 at the opening. Similar patterns are observed on the other days of that week. These results indicate that the Bitcoin ETF IBIT exhibited significant positive exposure to market risk during this turbulent week, precisely the periods when such instruments may be most valuable for risk management purposes.

To illustrate the advantages of the proposed candlestick-based framework relative to conventional return-based methods, I conduct an event study analysis around two Federal Open Market Committee (FOMC) meetings on June 12, 2024, and September 18, 2024.[24] Figure 6 reports the estimation results for these two events. The top panels display the prices of IBIT and SPY, while the bottom panels present the corresponding spot beta estimates for IBIT over a two-hour window surrounding the FOMC rate decision, with time zero aligned to the announcement (indicated by the vertical dashed line). Candlestick-based estimates are depicted by blue circles, while return-based estimates are shown by red triangles, with

[24]The literature widely recognizes FOMC announcements as among the most influential scheduled events, showing immediate and significant impacts on financial markets, see for example, Bernanke and Kuttner (2005); Savor and Wilson (2014); Lucca and Moench (2015); Cieslak et al. (2019).

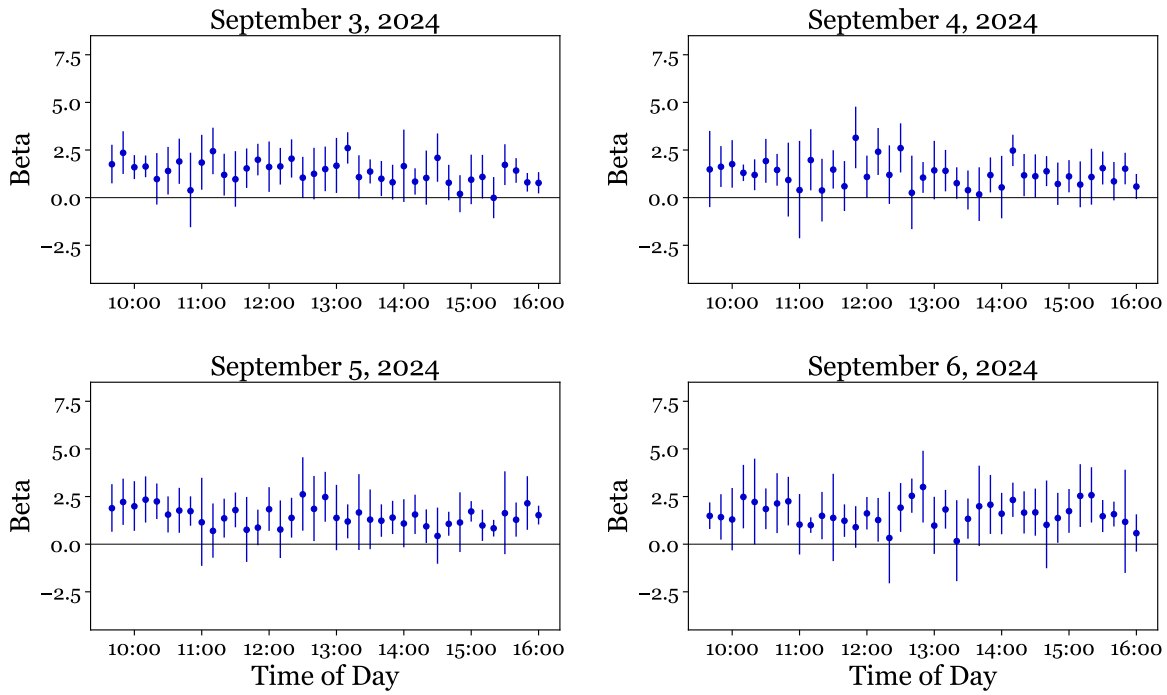

Figure 5: **Spot Beta of IBIT on September 3-6, 2024.** The figure shows the spot beta estimates of IBIT with respect to SPY using 10 $1-min$ candlestick observations in $10-min$ frequency. The vertical lines indicate the 95% confidence intervals.

vertical bars representing the 95% confidence intervals.

On September 18, 2024, the FOMC announced a 50-basis-point rate cut at 2:00 PM, which markets interpreted as a strong dovish signal, triggering increases in SPY and IBIT prices (see the top-right panel of Figure 6). Following the announcement, spot beta estimates increased from below 0.5 to around 1.5 and remained consistently positive and statistically significant for nearly an hour. In this case, the signal is strong, and both the candlestick-based and return-based methods rejected the null hypothesis, although small differences arise in the pre-announcement period.

In contrast, the June 12, 2024 FOMC meeting reflects a weaker informational shock. The Fed kept rates unchanged but released a relatively hawkish dot plot, prompting declines in SPY and IBIT prices (see the top-left panel). Here, the candlestick-based estimates still detect a significant positive exposure to market risk immediately after the release, whereas the return-based method fails to reject the null. Particularly, the return-based estimates exhibit wider confidence intervals, reflecting greater uncertainty around the estimates. This

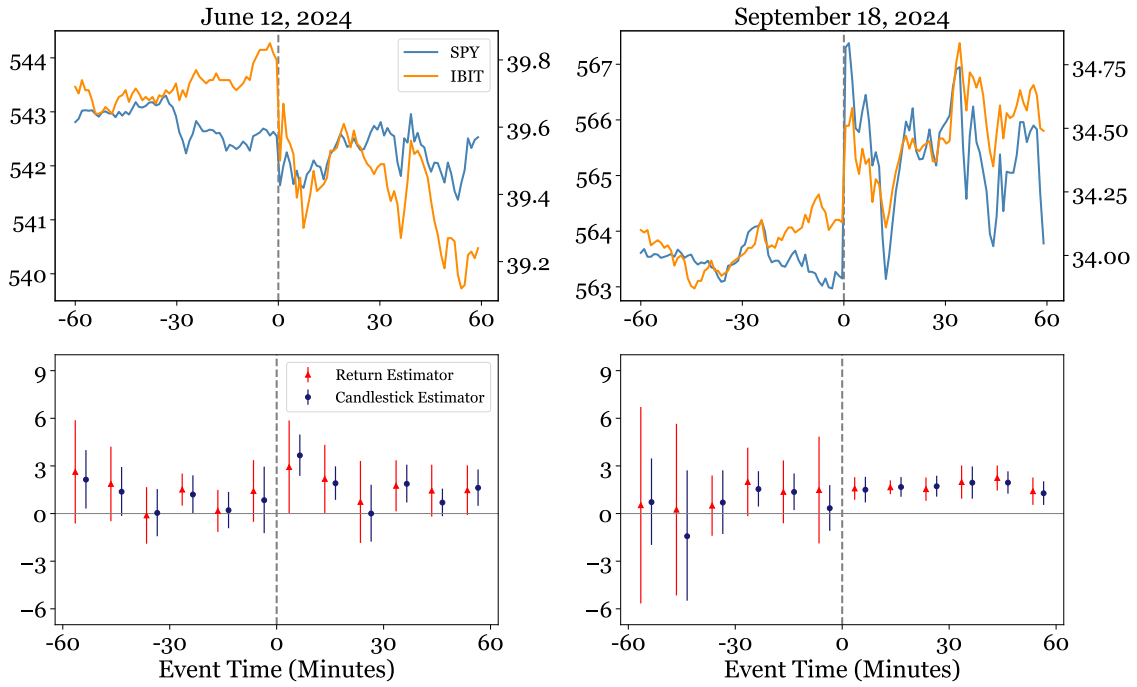

Figure 6: **Spot Beta of IBIT Around FOMC Meetings.** The figure shows the spot beta estimates of IBIT with respect to SPY using 10 1-min candlestick observations in 10-min frequency around two FOMC meetings: June 12, 2024 (left panels) and September 18, 2024 (right panels). The top panels display the price movements of IBIT and SPY, while the bottom panels present the corresponding spot beta estimates. Candlestick-based estimates are indicated by blue circles and return-based estimates by red triangles. The vertical lines indicate the 95% confidence intervals, and the vertical dashed line marks the time of the FOMC announcement.

comparison is consistent with the simulation results in Section 4, demonstrating the superior power properties of the candlestick-based approach. Consequently, these findings underscore the usefulness and robustness of the proposed approach in empirical settings.

6 Conclusion

This paper develops a new framework for estimating and conducting inference on spot covariance matrices using high-frequency candlestick data which consist of open, high, low and close prices for each interval. The approach builds on the idea of minimizing a well-defined estimation risk to optimally combine the information from all four prices. The resulting candlestick-based estimator has a simple form and is easy to implement. Monte Carlo experiments demonstrate that the candlestick-based estimator attains asymptotic risk levels

close to those of an infeasible oracle benchmark, while substantially outperforming return-based estimators, exhibiting markedly higher power (up to 25% improvement), reducing false negatives in hypothesis testing.

In an empirical study with 1-minute candlestick data of SPY (S&P 500 ETF) and IBIT (iShares Bitcoin Trust ETF), the method reveals significant positive market exposure of the Bitcoin ETF, especially during turbulent periods, challenging popular "digital gold" narrative for risk management. Moreover, in an event study around FOMC meetings, the candlestick-based approach detects significant market betas that the return-based method fails to do so. Overall, this work highlights the value of utilizing candlestick data for more accurate and powerful estimation and inference of spot covariance structures in financial markets.

References

Aït-Sahalia, Y., Fan, J., and Xiu, D. (2010). High-frequency covariance estimates with noisy and asynchronous financial data. *Journal of the American Statistical Association*, 105(492):1504–1517.

Aït-Sahalia, Y. and Jacod, J. (2014). *High-frequency financial econometrics*. Princeton University Press.

Andersen, T., Bollerslev, T., Diebold, F., and Labys, P. (2000). Great realizations. *RISK*, 13:105–108.

Andersen, T. G., Thyrsgaard, M., and Todorov, V. (2021). Recalcitrant betas: Intraday variation in the cross-sectional dispersion of systematic risk. *Quantitative Economics*, 12(2):647–682.

Back, K. (2010). *Asset pricing and portfolio choice theory*. Oxford University Press.

Bannouh, K., Van Dijk, D., and Martens, M. (2009). Range-based covariance estimation using high-frequency data: The realized co-range. *Journal of Financial Econometrics*, 7(4):341–372.

Barndorff-Nielsen, O., Hansen, P. R., Lunde, A., and Shephard, N. (2009). Realized kernels in practice: Trades and quotes. *The Econometrics Journal*, 12(3):C1–C32.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2011). Multivariate realised kernels: Consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. *Journal of Econometrics*, 162(2):149–169.

Barndorff-Nielsen, O. E. and Shephard, N. (2004a). Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics. *Econometrica*, 72(3):885–925.

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

Bernanke, B. S. and Kuttner, K. N. (2005). What explains the stock market's reaction to federal reserve policy? *The Journal of Finance*, 60(3):1221–1257.

Bibinger, M., Hautsch, N., Malec, P., and Reiss, M. (2019). Estimating the spot covariation of asset prices—statistical theory and empirical evidence. *Journal of Business & Economic Statistics*, 37(3):419–435.

Bollerslev, T., Hood, B., Huss, J., and Pedersen, L. H. (2018a). Risk everywhere: Modeling and managing volatility. *Review of Financial Studies*, 31(7):2729–2773.

Bollerslev, T., Li, J., and Li, Q. (2024a). Optimal nonparametric range-based volatility estimation. *Journal of Econometrics*, 238(1):105548.

Bollerslev, T., Li, J., Li, Q., and Li, Y. (2025). Optimal candlestick-based spot volatility estimation: New tricks and feasible inference procedures. *Available at SSRN 5046917*.

Bollerslev, T., Li, J., and Liao, Z. (2021). Fixed-k inference for volatility. *Quantitative Economics*, 12(4):1053–1084.

Bollerslev, T., Li, J., and Ren, Y. (2024b). Optimal inference for spot regressions. *American Economic Review*, 114(3):678–708.

Bollerslev, T., Patton, A. J., and Quaedvlieg, R. (2018b). Modeling and forecasting (un) reliable realized covariances for more reliable financial decisions. *Journal of Econometrics*, 207(1):71–91.

Bollerslev, T. and Todorov, V. (2011). Estimation of jump tails. *Econometrica*, 79(6):1727–1783.

Brandt, M. W. and Diebold, F. X. (2006). A no-arbitrage approach to range-based estimation of return covariances and correlations. *Journal of Business*, 79(1):61–74.

Christensen, K., Kinnebrock, S., and Podolskij, M. (2010). Pre-averaging estimators of the ex-post covariance matrix in noisy diffusion models with non-synchronous data. *Journal of Econometrics*, 159(1):116–133.

Christensen, K. and Podolskij, M. (2007). Realized range-based estimation of integrated variance. *Journal of Econometrics*, 141(2):323–349.

Christensen, K. and Podolskij, M. (2012). Asymptotic theory of range-based multipower variation. *Journal of Financial Econometrics*, 10(3):417–456.

Christensen, K., Podolskij, M., and Vetter, M. (2009). Bias-correcting the realized range-based variance in the presence of market microstructure noise. *Finance and Stochastics*, 13:239–268.

Cieslak, A., Morse, A., and Vissing-Jorgensen, A. (2019). Stock returns over the fomc cycle. *The Journal of Finance*, 74(5):2201–2248.

Diebold, F. X. and Strasser, G. (2013). On the correlation structure of microstructure noise: A financial economic approach. *Review of Economic Studies*, 80(4):1304–1337.

Fama, E. F. and French, K. R. (1992). The cross-section of expected stock returns. *The Journal of Finance*, 47(2):427–465.

Fan, J. and Wang, Y. (2008). Spot volatility estimation for high-frequency data. *Statistics and its Interface*, 1(2):279–288.

Garman, M. B. and Klass, M. J. (1980). On the estimation of security price volatilities from historical data. *Journal of Business*, pages 67–78.

Hayashi, T. and Yoshida, N. (2005). On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli*, 11(2):359–379.

Jacod, J., Li, J., and Liao, Z. (2021). Volatility coupling. *The Annals of Statistics*, 49(4):1982–1998.

Jacod, J. and Protter, P. (2012). *Discretization of Processes*. Springer.

Lehmann, E. L. and Casella, G. (2006). *Theory of point estimation*. Springer Science & Business Media.

Lewandowski, D., Kurowicka, D., and Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, 100(9):1989–2001.

Li, J., Wang, D., and Zhang, Q. (2024). Reading the candlesticks: An ok estimator for volatility. *Review of Economics and Statistics*, 106(4):1114–1128.

Li, Y., Nolte, I., Nolte, S., and Yu, S. (2025). Realized candlestick wicks. *Journal of Econometrics*, 250:106014.

Lintner, J. (1965). Security prices, risk, and maximal gains from diversification. *The Journal of Finance*, 20(4):587–615.

Liu, Y. and Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *The Review of Financial Studies*, 34(6):2689–2727.

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

Mancini, C. (2009). Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps. *Scandinavian Journal of Statistics*, 36(2):270–296.

Martens, M. and Van Dijk, D. (2007). Measuring volatility with the realized range. *Journal of Econometrics*, 138(1):181–207.

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: the information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330.

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *Journal of Business*, pages 61–65.

Rogers, L. and Shepp, L. (2006). The correlation of the maxima of correlated brownian motions. *Journal of Applied Probability*, 43(3):880–883.

Rogers, L. C. and Zhou, F. (2008). Estimating correlation from high, low, opening and closing prices. *Annals of Applied Probability*, 18(2):813–823.

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

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442.

Appendix A

A.1 Proof of Proposition 1

In this section, I provide the proof of Proposition 1 which builds on the coupling techniques developed in Jacod et al. (2021) and Bollerslev et al. (2021). The former studies the approximation of the estimation error for the spot covariance estimator when k increases with n, whereas the latter focuses on coupling the spot volatility estimator in a fixed-k framework. As noted in the main text, I consider a fixed-k setup and thus my work is in the same spirit as Bollerslev et al. (2021). To ensure consistency with the existing literature, I closely follow the notation introduced in the aforementioned papers.

I rewrite the Assumption 1 and Proposition 1 here for convenience:

Assumption 1. *Suppose that $\boldsymbol{X}_t$ has the form in Equation (1) and there exists a sequence $(T_m)_{m\geq 1}$ of stopping times increasing to infinity and the following conditions hold for each $m \geq 1$:*

(i) $\|\boldsymbol{b}_t\| + \|\boldsymbol{\sigma}_t\| + \|\boldsymbol{\sigma}_t^{-1}\| \leq K_m$ *for some constant* K_m *for all* $t \in [0, T_m]$*;*

(ii) $\mathbb{E}\left[\|\boldsymbol{\sigma}_{t\wedge T_m} - \boldsymbol{\sigma}_{s\wedge T_m}\|^2\right] \leq K_m|t-s|$ *for all* $t, s \in [0, T_m]$*.*

Proposition 1. *Suppose that Assumption 1 holds. Fix any $t \in [0, T]$. For any $k \geq 1$ and $\boldsymbol{\Lambda}$, the following holds as $\Delta_n \to 0$:*

$$\left\| \boldsymbol{\sigma}_t^{-1} \widehat{\boldsymbol{c}}_{n,t}(\lambda) \boldsymbol{\sigma}_t^{-\top} - U_{n,t} \right\| = o_p(1) \tag{A.1}$$

where $U_{n,t} = \frac{1}{k}\sum_{i\in\mathcal{I}_{n,t}} \left\{ \lambda_1 \boldsymbol{\zeta}_{i,r}\boldsymbol{\zeta}_{i,r}^\top + \lambda_2 \boldsymbol{\zeta}_{i,h}\boldsymbol{\zeta}_{i,a}^\top + \lambda_3 \boldsymbol{\zeta}_{i,\ell}\boldsymbol{\zeta}_{i,w}^\top \right\}$ *and, for any* $i \in \mathcal{I}_{n,t}$*,*

$$\begin{aligned}
\boldsymbol{\zeta}_{i,r} &\equiv \frac{\boldsymbol{W}_{i\Delta_n} - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \\
\boldsymbol{\zeta}_{i,a} &\equiv \boldsymbol{\varrho}_t^{-1} \sup_{\tau\in\mathcal{T}_i} \boldsymbol{\varrho}_t \left(\frac{\boldsymbol{W}_\tau - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) + \boldsymbol{\varrho}_t^{-1} \inf_{\tau\in\mathcal{T}_i} \boldsymbol{\varrho}_t \left(\frac{\boldsymbol{W}_\tau - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) - \left(\frac{\boldsymbol{W}_{i\Delta_n} - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) \\
\boldsymbol{\zeta}_{i,w} &\equiv \boldsymbol{\varrho}_t^{-1} \sup_{\tau\in\mathcal{T}_i} \boldsymbol{\varrho}_t \left(\frac{\boldsymbol{W}_\tau - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) - \boldsymbol{\varrho}_t^{-1} \inf_{\tau\in\mathcal{T}_i} \boldsymbol{\varrho}_t \left(\frac{\boldsymbol{W}_\tau - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right)
\end{aligned} \tag{A.2}$$

with $\boldsymbol{\varrho}_t$ being the square root of spot correlation matrix $\boldsymbol{\rho}_t$, i.e., $\boldsymbol{\rho}_t = \boldsymbol{\varrho}_t\boldsymbol{\varrho}_t^\top$. In explicit terms, $\boldsymbol{\rho}_t = \operatorname{diag}(\boldsymbol{c}_t)^{-\frac{1}{2}} \boldsymbol{c}_t \operatorname{diag}(\boldsymbol{c}_t)^{-\frac{1}{2}}$ and $\boldsymbol{\varrho}_t = \operatorname{diag}(\boldsymbol{c}_t)^{-\frac{1}{2}} \boldsymbol{\sigma}_t$ where $\operatorname{diag}(\boldsymbol{c}_t)$ is a diagonal matrix with

the same diagonal elements as $\boldsymbol{c}_t$.

Proof. Fix $k \geq 1$ and λ. Let K denote a generic positive constant. As is common in the spot estimation literature, one can strengthen Assumption 1 by assuming the conditions hold with $T_m = \infty$, which can be justified by a standard localization argument (see Jacod and Protter (2012) for details).

I begin by writing out the explicit expressions for the candlestick returns:

$$\begin{aligned} \boldsymbol{r}_i &\equiv \Delta_n^{-\frac{1}{2}} \left(\textstyle\int_{(i-1)\Delta_n}^{i\Delta_n} \boldsymbol{b}_s ds + \int_{(i-1)\Delta_n}^{i\Delta_n} \boldsymbol{\sigma}_s d\boldsymbol{W}_s \right) \\ \boldsymbol{h}_i &\equiv \Delta_n^{-\frac{1}{2}} \left(\sup_{t \in \mathcal{T}_{n,i}} \left(\textstyle\int_{(i-1)\Delta_n}^{t} \boldsymbol{b}_s ds + \int_{(i-1)\Delta_n}^{t} \boldsymbol{\sigma}_s d\boldsymbol{W}_s \right) \right) \\ \boldsymbol{\ell}_i &\equiv \Delta_n^{-\frac{1}{2}} \left(\inf_{t \in \mathcal{T}_{n,i}} \left(\textstyle\int_{(i-1)\Delta_n}^{t} \boldsymbol{b}_s ds + \int_{(i-1)\Delta_n}^{t} \boldsymbol{\sigma}_s d\boldsymbol{W}_s \right) \right) \end{aligned} \tag{A.3}$$

I also introduce the following definitions:

$$\begin{aligned} \boldsymbol{r}'_i &\equiv \boldsymbol{\sigma}_{(i-1)\Delta_n} \left(\tfrac{\boldsymbol{W}_{i\Delta_n} - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) \\ \boldsymbol{h}'_i &\equiv \sup_{t \in \mathcal{T}_{n,i}} \boldsymbol{\sigma}_{(i-1)\Delta_n} \left(\tfrac{\boldsymbol{W}_t - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) \\ \boldsymbol{\ell}'_i &\equiv \inf_{t \in \mathcal{T}_{n,i}} \boldsymbol{\sigma}_{(i-1)\Delta_n} \left(\tfrac{\boldsymbol{W}_t - \boldsymbol{W}_{(i-1)\Delta_n}}{\sqrt{\Delta_n}} \right) \end{aligned} \tag{A.4}$$

which serve as the coupling variables for the candlestick returns. Similar variables can be defined for the range and asymmetry variables:

$$\begin{aligned} \boldsymbol{a}'_i &\equiv \boldsymbol{h}'_i + \boldsymbol{\ell}'_i - \boldsymbol{r}'_i \\ \boldsymbol{w}'_i &\equiv \boldsymbol{h}'_i - \boldsymbol{\ell}'_i \end{aligned} \tag{A.5}$$

The proof consists of two steps. The first step controls how well the coupling returns in Equation (A.4) approximate the candlestick returns in Equation (A.3). The second step combines these results with the continuous mapping theorem to reach the desired conclusion. Before proceeding to the first step, I derive useful intermediate results.

By Assumption 1, it is easy to see that:

$$\left\| \int_{(i-1)\Delta_n}^{i\Delta_n} \boldsymbol{b}_s ds \right\| \leq \int_{(i-1)\Delta_n}^{i\Delta_n} \|\boldsymbol{b}_s\| ds = O_p(\Delta_n). \tag{A.6}$$

Moreover, the Burkholder-Davis-Gundy inequality and Assumption 1 imply that:

$$\begin{aligned}\mathbb{E}\left[\left\|\int_{(i-1)\Delta_n}^{i\Delta_n}(\boldsymbol{\sigma}_s-\boldsymbol{\sigma}_{(i-1)\Delta_n})d\boldsymbol{W}_s\right\|^2\right] &\leq K\Delta_n\mathbb{E}\left[\int_{(i-1)\Delta_n}^{i\Delta_n}\|\boldsymbol{\sigma}_s-\boldsymbol{\sigma}_{(i-1)\Delta_n}\|^2 ds\right]\\ &\leq K\Delta_n^2.\end{aligned} \tag{A.7}$$

Further, we can deduce that:

$$\sup_{t\in\mathcal{T}_{n,i}}\left\|\int_{(i-1)\Delta_n}^{t}(\boldsymbol{\sigma}_s-\boldsymbol{\sigma}_{(i-1)\Delta_n})d\boldsymbol{W}_s\right\| = O_p(\Delta_n). \tag{A.8}$$

Step 1: We now establish approximation results for $\boldsymbol{r}_i$, $\boldsymbol{h}_i$ and $\boldsymbol{\ell}_i$ separately. We start with the return:

$$\begin{aligned}\|\boldsymbol{r}_i-\boldsymbol{r}_i'\| &\leq \left\|\Delta_n^{-\frac{1}{2}}\int_{(i-1)\Delta_n}^{i\Delta_n}\boldsymbol{b}_s ds\right\| + \left\|\Delta_n^{-\frac{1}{2}}\int_{(i-1)\Delta_n}^{i\Delta_n}(\boldsymbol{\sigma}_s-\boldsymbol{\sigma}_{(i-1)\Delta_n})d\boldsymbol{W}_s\right\|\\ &= O_p(\Delta_n^{\frac{1}{2}}).\end{aligned} \tag{A.9}$$

where the first line directly follows from the triangle inequality and the second line uses above intermediate results in Equation (A.6) and (A.8). For the high return, we have:

$$\begin{aligned}\left\|\boldsymbol{h}_i-\boldsymbol{h}_i'\right\| &\equiv \Delta_n^{-\frac{1}{2}}\Big\|\sup_{t\in\mathcal{T}_{n,i}}\Big(\int_{(i-1)\Delta_n}^{t}\boldsymbol{b}_s ds+\int_{(i-1)\Delta_n}^{t}\boldsymbol{\sigma}_s d\boldsymbol{W}_s\Big)\\ &\qquad -\sup_{t\in\mathcal{T}_{n,i}}\boldsymbol{\sigma}_{(i-1)\Delta_n}(\boldsymbol{W}_t-\boldsymbol{W}_{(i-1)\Delta_n})\Big\|\\ &\leq \Delta_n^{-\frac{1}{2}}\sup_{t\in\mathcal{T}_{n,i}}\left\|\int_{(i-1)\Delta_n}^{t}\boldsymbol{b}_s ds+\int_{(i-1)\Delta_n}^{t}(\boldsymbol{\sigma}_s-\boldsymbol{\sigma}_{(i-1)\Delta_n})d\boldsymbol{W}_s\right\|\\ &\leq \Delta_n^{-\frac{1}{2}}\int_{(i-1)\Delta_n}^{i\Delta_n}\|\boldsymbol{b}_s\|ds+\sup_{t\in\mathcal{T}_{n,i}}\left\|\int_{(i-1)\Delta_n}^{t}(\boldsymbol{\sigma}_s-\boldsymbol{\sigma}_{(i-1)\Delta_n})d\boldsymbol{W}_s\right\|\\ &= O_p(\Delta_n^{\frac{1}{2}}).\end{aligned} \tag{A.10}$$

The first two lines are obviously implications of sup definition. Similarly, the last line follows from Equation A.6 and A.8. Finally, one can deduce the same inequality for the low return:

$$\left\|\boldsymbol{\ell}_i-\boldsymbol{\ell}_i'\right\| = O_p(\Delta_n^{\frac{1}{2}}), \tag{A.11}$$

and also for the asymmetry and range variables:

$$\begin{aligned} \|\boldsymbol{a}_i - \boldsymbol{a}_i'\| &\leq \|\boldsymbol{h}_i - \boldsymbol{h}_i'\| + \|\boldsymbol{\ell}_i - \boldsymbol{\ell}_i'\| + \|\boldsymbol{r}_i - \boldsymbol{r}_i'\| = O_p(\Delta_n^{\frac{1}{2}}) \\ \|\boldsymbol{w}_i - \boldsymbol{w}_i'\| &\leq \|\boldsymbol{h}_i - \boldsymbol{h}_i'\| + \|\boldsymbol{\ell}_i - \boldsymbol{\ell}_i'\| = O_p(\Delta_n^{\frac{1}{2}}). \end{aligned} \tag{A.12}$$

Furthermore, I claim that

$$\begin{aligned} \left\|\boldsymbol{\sigma}_t^{-1}\boldsymbol{r}_i' - \boldsymbol{\zeta}_{i,r}\right\| &= O_p(\Delta_n^{1/2}) \\ \left\|\boldsymbol{\sigma}_t^{-1}\boldsymbol{a}_i' - \boldsymbol{\zeta}_{i,a}\right\| &= O_p(\Delta_n^{1/2}) \\ \left\|\boldsymbol{\sigma}_t^{-1}\boldsymbol{w}_i' - \boldsymbol{\zeta}_{i,w}\right\| &= O_p(\Delta_n^{1/2}). \end{aligned} \tag{A.13}$$

The first line directly follows from Assumption 1. Note that $|i\Delta_n - t| \to 0$ for any $i \in \mathcal{I}_{n,t}$ as $\Delta_n \to 0$ and this implies $\|\boldsymbol{\sigma}_t - \boldsymbol{\sigma}_{(i-1)\Delta_n}\| = O_p(\Delta_n^{1/2})$. For notational convenience, I only consider the third line and the same steps can be adopted for the second line as well. Specifically, one can write:

$$\begin{aligned} \|\boldsymbol{\sigma}_t^{-1}\boldsymbol{w}_i' - \boldsymbol{\zeta}_{i,w}\| &\leq \|\boldsymbol{\sigma}_t^{-1}\| \cdot \|\boldsymbol{w}_i' - \boldsymbol{\sigma}_t\boldsymbol{\zeta}_{i,w}\| \\ &= \|\boldsymbol{\sigma}_t^{-1}\| \cdot \left\|\boldsymbol{w}_i' - \mathrm{diag}(\boldsymbol{c}_t)^{1/2} \sup_{\tau,s\in\mathcal{T}_{n,i}} \boldsymbol{\varrho}_t \left(\tfrac{\boldsymbol{W}_\tau - \boldsymbol{W}_s}{\sqrt{\Delta_n}}\right)\right\| \\ &= \|\boldsymbol{\sigma}_t^{-1}\| \cdot \left\|\boldsymbol{w}_i' - \sup_{\tau,s\in\mathcal{T}_{n,i}} \boldsymbol{\sigma}_t \left(\tfrac{\boldsymbol{W}_\tau - \boldsymbol{W}_s}{\sqrt{\Delta_n}}\right)\right\| \\ &= \|\boldsymbol{\sigma}_t^{-1}\| \cdot \left\|\sup_{\tau,s\in\mathcal{T}_{n,i}} (\boldsymbol{\sigma}_t - \boldsymbol{\sigma}_{(i-1)\Delta_n}) \left(\tfrac{\boldsymbol{W}_\tau - \boldsymbol{W}_s}{\sqrt{\Delta_n}}\right)\right\| \\ &= O_p(\Delta_n^{1/2}) \end{aligned} \tag{A.14}$$

where the first line follows from sub-multiplicative property of matrix norm, the second and third lines use the definition of $\boldsymbol{\zeta}_{i,w}$ and $\boldsymbol{\varrho}_t$, the fourth line is a direct implication of sup definition and the last line uses Assumption 1 and properties of Brownian motion.

Step 2: Rewrite the main statement of the proposition as:

$$\left\|\boldsymbol{\sigma}_t^{-1}\widehat{\boldsymbol{c}}_{n,t}(\lambda)\boldsymbol{\sigma}_t^{-\top} - U_{n,t}\right\| = \left\|\frac{1}{k}\sum_{i\in\mathcal{I}_{n,t}}\left\{\lambda_1(\boldsymbol{\sigma}_t^{-1}\boldsymbol{r}_i)(\boldsymbol{\sigma}_t^{-1}\boldsymbol{r}_i)^\top - \boldsymbol{\zeta}_{i,r}\boldsymbol{\zeta}_{i,r}^\top\right\}\right.$$
$$+\frac{1}{k}\sum_{i\in\mathcal{I}_{n,t}}\left\{\lambda_2(\boldsymbol{\sigma}_t^{-1}\boldsymbol{a}_i)(\boldsymbol{\sigma}_t^{-1}\boldsymbol{a}_i)^\top - \boldsymbol{\zeta}_{i,a}\boldsymbol{\zeta}_{i,a}^\top\right\}$$
$$\left.+\frac{1}{k}\sum_{i\in\mathcal{I}_{n,t}}\left\{\lambda_3(\boldsymbol{\sigma}_t^{-1}\boldsymbol{w}_i)(\boldsymbol{\sigma}_t^{-1}\boldsymbol{w}_i)^\top - \boldsymbol{\zeta}_{i,w}\boldsymbol{\zeta}_{i,w}^\top\right\}\right\| \tag{A.15}$$

Using the results from Step 1, it follows that the terms in curly brackets are $O_p(\Delta_n^{1/2})$ for any $i \in \mathcal{I}_{n,t}$ and fixed $\lambda = (\lambda_1, \lambda_2, \lambda_3)^\top$. Therefore, the entire expression is $O_p(\Delta_n^{1/2}) = o_p(1)$. This completes the proof.

□

A.2 Proof of Proposition 2

Now, I provide the proof of Proposition 2, which establishes the coupling result for the test statistic defined as:

$$\widehat{T}_n = \frac{\sqrt{k-1}\left(\widehat{\beta}_{n,t} - \beta_t\right)}{\sqrt{\widehat{\varsigma}_{n,t}/\widehat{\nu}_{n,t}}}.$$

The proof is based on the algebraic manipulations of the previous proposition. For convenience, I restate the proposition here:

Proposition 2. *Under the conditions of Proposition 1, for any fixed $k \geq 2$ and λ, the following holds as $\Delta_n \to 0$:*

$$|\widehat{T}_n - \widetilde{T}_n| = o_p(1) \quad \textit{where} \quad \widetilde{T}_n \equiv \frac{\sqrt{k-1}[U_{n,t}^{-1}]_{12}}{\sqrt{[U_{n,t}^{-1}]_{11}[U_{n,t}^{-1}]_{22} - [U_{n,t}^{-1}]_{12}^2}}.$$

Proof. Fix $k \geq 2$ and λ. By Proposition 1, we have:

$$\left\|\boldsymbol{\sigma}_t^{-1}\widehat{\boldsymbol{c}}_{n,t}(\lambda)\boldsymbol{\sigma}_t^{-\top} - U_{n,t}\right\| = o_p(1). \tag{A.16}$$

Note that both $\boldsymbol{\sigma}_t^{-1}\widehat{\boldsymbol{c}}_{n,t}(\lambda)\boldsymbol{\sigma}_t^{-\top}$ and $U_{n,t}$ are positive definite matrices. Therefore, by the

continuity of matrix inversion operator on the set of positive definite matrices, we have:

$$\left\|(\boldsymbol{\sigma}_t^{-1}\widehat{\boldsymbol{c}}_{n,t}(\lambda)\boldsymbol{\sigma}_t^{-1\top})^{-1} - U_{n,t}^{-1}\right\| = o_p(1). \tag{A.17}$$

This implies that each element of the matrix $(\boldsymbol{\sigma}_t^{-1}\widehat{\boldsymbol{c}}_{n,t}(\lambda)\boldsymbol{\sigma}_t^{-\top})^{-1}$ converges to the corresponding element of $U_{n,t}^{-1}$ in probability. These relations can be written in explicit forms as follows:

$$\begin{aligned}
\left|\left(\frac{\nu_t}{\widehat{\nu}_{n,t}} + \frac{\nu_t(\beta_t - \widehat{\beta}_{n,t})^2}{\widehat{\varsigma}_{n,t}}\right) - [U_{n,t}^{-1}]_{11}\right| &= o_p(1) \\
\left|\left(\frac{\nu_t^{1/2}\varsigma_t^{1/2}(\beta_t - \widehat{\beta}_{n,t})}{\widehat{\varsigma}_{n,t}}\right) - [U_{n,t}^{-1}]_{12}\right| &= o_p(1) \\
\left|\frac{\varsigma_t}{\widehat{\varsigma}_{n,t}} - [U_{n,t}^{-1}]_{22}\right| &= o_p(1).
\end{aligned}$$

Using the second and third lines, one can write:

$$\left|\frac{(\widehat{\beta}_{n,t} - \beta_t)}{\sqrt{\varsigma_{n,t}/\nu_{n,t}}} - \frac{[U_{n,t}^{-1}]_{12}}{[U_{n,t}^{-1}]_{22}}\right| = o_p(1). \tag{A.18}$$

Moreover, from all three lines, one can deduce that:

$$\left|\frac{\nu_t}{\widehat{\nu}_{n,t}} - \left([U_{n,t}^{-1}]_{11} - \frac{[U_{n,t}^{-1}]_{12}^2}{[U_{n,t}^{-1}]_{22}}\right)\right| = o_p(1). \tag{A.19}$$

Finally, combining the above equations yields:

$$\left|\frac{\sqrt{k-1}(\widehat{\beta}_{n,t} - \beta_t)}{\sqrt{\widehat{\varsigma}_{n,t}/\widehat{\nu}_{n,t}}} - \frac{\sqrt{k-1}[U_{n,t}^{-1}]_{12}}{\sqrt{[U_{n,t}^{-1}]_{11}[U_{n,t}^{-1}]_{22} - [U_{n,t}^{-1}]_{12}^2}}\right| = o_p(1). \tag{A.20}$$

as required.

□