TALOS

ANALYSIS

Hourly Market Invariants for Price Simulations in Digital-Asset Markets

February 2026

TALOS

Hourly Market Invariants for Price Simulations in Digital-Asset Markets

Marco Marchioro
https://www.talos.com/
Version 0.99

February 2026*

Abstract

We test whether 1-hour log-returns of major digital assets can serve as *market invariants* in the sense of Meucci (2005) for 24/7 markets. Using 18 months of hourly prices for six liquid tokens, we find broadly stable marginal distributions across three 6-month windows, weak linear serial dependence in returns (in contrast to prices), and approximate horizon consistency after square-root-of-time rescaling from 1 hour to 1 day. These diagnostics support a practical scenario engine: map hourly innovations to a daily-equivalent return and exponentiate to generate 1-day token price scenarios for portfolio repricing and daily-horizon risk measurement without relying on an end-of-day close.

Contents

1 Introduction 1

2 Market invariants for digital-asset returns 2

3 Data and preprocessing 4
3.1 Distribution of log-returns 4

4 Empirical diagnostics for invariance 5
4.1 Stationarity across three subsamples 5
4.2 Lag plots and autocorrelation function 7
4.3 Aggregation and square-root-of-time scaling 9

5 Conclusion 11

1 Introduction

Digital assets such as Bitcoin (BTC), Ether (ETH) and a growing set of alternative layer-1 and payment tokens have become an important asset class for both retail and institutional investors. These markets trade continuously, exhibit substantial volatility and heavy tails, and show stylized facts that differ in magnitude from those of traditional assets. As a result, risk management and portfolio construction for digital-asset exposures require models that are both flexible and empirically grounded.

*Latest review February 5, 2026

Symbol	Asset name
BTC	Bitcoin
DOGE	Dogecoin
ETH	Ethereum
SOL	Solana
XLM	Stellar (Lumens)
XRP	XRP (Ripple)

Table 1: Sample universe of liquid digital assets used in the empirical diagnostics.

In the framework of Meucci [1], a central step in risk and portfolio analysis is the identification of *market invariants*: transformations of observable prices that are approximately independent and identically distributed (i.i.d.) at a chosen sampling interval. Once such invariants are identified, their joint distribution can be estimated and used as an input to scenario generation and historical simulation engines.

Our contribution is to provide an empirically validated bridge from 24/7 hourly data to the daily horizon used in risk reporting. Concretely, we show that for a set of major liquid tokens, 1-hour log-returns behave as approximate market invariants and can therefore be used to generate daily-horizon *token price scenarios* via a simple daily-equivalent mapping and exponential repricing. Empirically, we document (i) approximate stability of the marginal distribution over consecutive 6-month windows, (ii) weak linear dependence at short lags, and (iii) an approximately consistent scaling behavior from 1 hour to 1 day under a daily-equivalent normalization. Methodologically, we frame these diagnostics within the market-invariants perspective to support scenario-based repricing and daily-horizon risk measurement. In short: hourly token returns provide a larger, 24/7-native scenario set that can be projected to daily price scenarios for risk measurement.

In this paper we adapt this notion of invariance to 24/7 digital-asset markets. Because there is no natural end-of-day close, we work with equally spaced *hourly* log-returns as candidate invariants and later relate them to the daily horizon typically used in risk reporting. Our empirical study uses 18 months of hourly prices for a number of liquid tokens, listed in Table 1. We investigate whether hourly returns for these assets exhibit the basic properties required of market invariants and discuss how they can be employed in historical simulations of portfolios of digital assets.

2 Market invariants for digital-asset returns

Invariants in the sense of Meucci

Following [1], we fix an *estimation interval* Δt and a corresponding equally spaced time grid $\{t_0, t_0 + \Delta t, t_0 + 2\Delta t, \dots\}$. A sequence of random variables $\{X_t\}$ observed at these times is called a set of *market invariants* if:

1. X_t is fully observable at time t;
2. the sequence $\{X_t\}$ is (approximately) independent and identically distributed in t; and
3. the common distribution of X_t does not depend on the calendar location of the estimation window.

In the context of traditional asset classes, natural invariants include non-overlapping log-returns for equities and exchange rates, changes in yields-to-maturity for fixed-income instruments, and changes in implied volatility for options. For digital assets the 24/7 trading schedule makes hourly log-returns a natural choice.

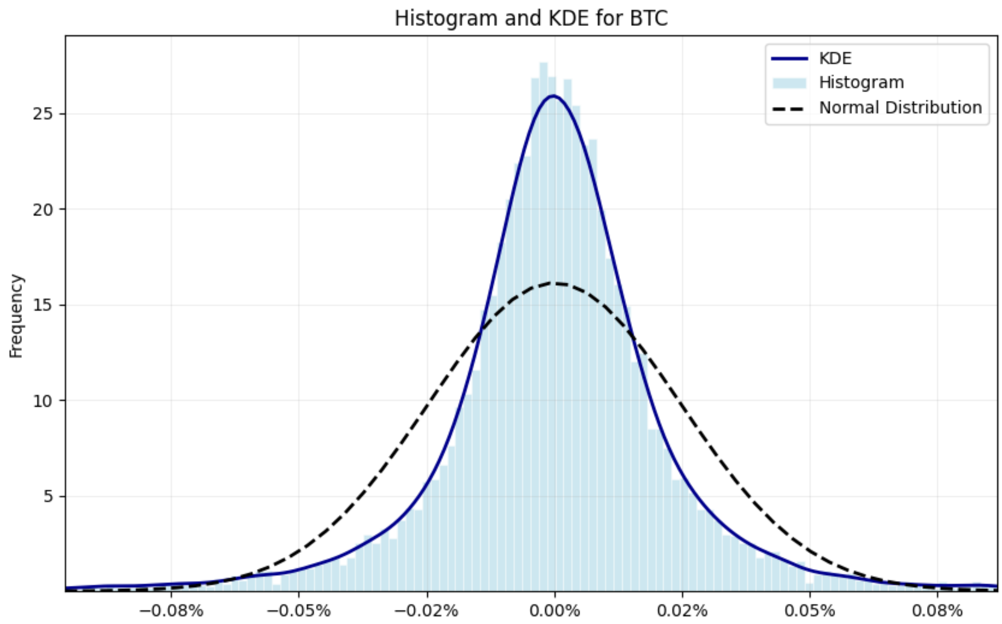

Figure 1: Empirical distribution of 1-hour BTC log-returns. Light blue bars: histogram; thick dark-blue line: kernel density estimate (KDE); black dashed line: Gaussian density with the same sample mean and standard deviation. The vertical axis is on a linear scale, highlighting the sharp central peak.

In this study we measure time in discrete hours and set the estimation interval to $\Delta t = 1$ (hour). Let $t \in \{0, 1, \ldots, T\}$ index hourly timestamps in UTC, and let P_t denote the last observed transaction price recorded in the interval $(t-1, t]$. We define the corresponding 1-hour log-return as

$$X_t = \log\left(\frac{P_t}{P_{t-1}}\right), \qquad t = 1, \ldots, T. \tag{1}$$

For each token, the collection $\{X_t\}_{t=1}^{T}$ is our candidate set of market invariants at the 1-hour horizon.

The diagnostics in this paper can be viewed as an implementation of the first step in Meucci's broader *checklist* (also known as *The Prayer*, see reference [2]), which organizes risk and portfolio modeling into a sequence of steps starting from the quest for invariance and ending with scenario-based evaluation and decision making. Within that checklist, validating approximate invariance at the chosen sampling frequency is a prerequisite for meaningful scenario generation and risk measurement at the investment horizon.

Daily-horizon token price scenario generation from hourly invariants

In reference [3], we define risk measures such as Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) and gain measures such as Gain-at-Risk (GaR) and Conditional Gain-at-Risk (CGaR) for a generic distribution of returns.

Risk measures are often reported at a 1-day horizon, while the candidate invariants in this paper are sampled hourly. There are two conceptually distinct ways to bridge these horizons: (i) explicitly aggregate 24 hourly innovations to form a 24-hour return, or (ii) apply an approximate scaling rule that maps 1-hour returns to a daily-equivalent scale. In what follows we focus on the second approach and later assess empirically whether the resulting scaled distributions are consistent across horizons.

Specifically, we define a *daily-equivalent* (pseudo-daily) log-return as

$$Y_t \equiv X_t \cdot \sqrt{24}. \tag{2}$$

This transformation matches the 1-day variance implied by i.i.d. hourly innovations, but it does not, in general, preserve the exact return distribution under aggregation. Therefore we treat Y_t as a modeling device whose adequacy must be validated empirically (see subsection 4.3).

Given a reference price P_r for the given asset, we can simulate the daily-price scenarios $\{P_t\}$ as:

$$P_t = P_r \cdot \exp(Y_t) = P_r \cdot \exp(X_t \cdot \sqrt{24}) \,. \tag{3}$$

The collection of daily-price scenarios $\{P_t\}$ can then be used in a fully repricing simulation of a portfolio of digital assets. Note that the simulated Y_t returns are not observed in the market; accordingly, the simulated P_t prices should be interpreted as model-based scenarios rather than as actual samples in a historical simulation of the given asset.

3 Data and preprocessing

We work with 18 months of hourly price data for the digital assets as listed in Table 1. For each asset we construct an equally spaced time grid of hourly timestamps in Coordinated Universal Time (UTC). When multiple trades occur within the same hour, we use the last observed transaction price to represent that hour. From these hourly prices we compute log-returns at various horizons. The primary object of interest is the sequence of 1-hour log-returns $\{X_t^{1h}\}$. For the scaling analysis we also construct returns over 2, 4, 8, and 24 hours, using non-overlapping windows to avoid artificial serial dependence.

3.1 Distribution of log-returns

Before presenting invariance diagnostics, we summarize the empirical distribution of 1-hour log-returns for a representative asset (Bitcoin). We compare the empirical density to a Gaussian benchmark with the same sample mean and standard deviation, as a compact way to visualize tail thickness and central concentration.

Figure 1 shows the histogram of the scaled BTC log-returns (light blue bars) together with a *kernel density estimate* (KDE, thick dark-blue line) and the density of the corresponding normal distribution (black dashed line), all plotted on a linear scale for the vertical axis.

The KDE provides a smooth, nonparametric estimate of the probability density function. We prefer the KDE to the raw histogram when comparing distributions across assets and horizons, because it is less sensitive to arbitrary choices of bin edges and yields a smoother representation of the underlying density, while still tracking the histogram closely in the regions where data is abundant.

On the linear scale in figure 1, the most striking feature is the pronounced central peak of BTC returns: the KDE towers above the Gaussian density around zero, indicating that small returns occur more frequently than under the normal model (leptokurtosis). The KDE also approximates the histogram well in the center of the distribution. However, in this representation it is difficult to visually distinguish the behavior of the empirical and Gaussian densities in the far tails, where probabilities are small.

To highlight tail behavior, figure 2 repeats the same comparison but with a logarithmic scale on the vertical axis. On this log scale, the differences in the tails become clear: both the histogram and the KDE lie systematically above the normal density for large positive and negative returns. In other words, extreme BTC returns occur more often than would be predicted by a Gaussian

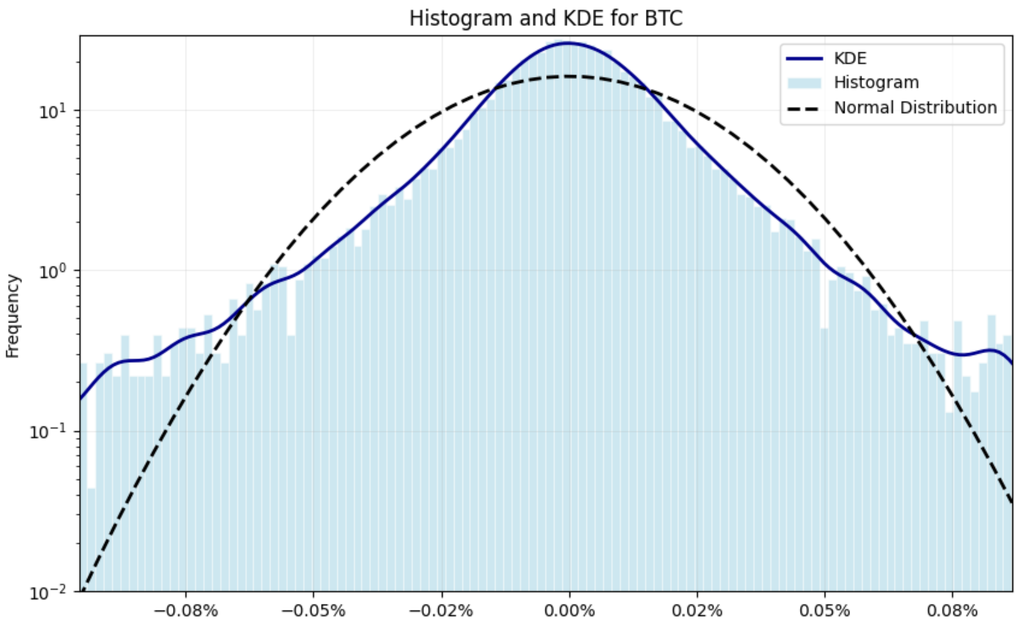

Figure 2: Same comparison as figure 1, but with a logarithmic vertical axis to emphasize tail behavior. The empirical density exceeds the Gaussian benchmark in both tails, consistent with heavy-tailed returns.

distribution with the same mean and variance, a hallmark of *fat tails*. This pattern is qualitatively similar for the other tokens in our universe and motivates the use of flexible, nonparametric tools when modeling the distribution of digital-asset returns.

4 Empirical diagnostics for invariance

In this section we perform empirical diagnostics to assess the invariance of the scaled log-returns. We examine the three aspects of invariance: time-homogeneity, weak serial dependence, and scaling across horizons.

4.1 Stationarity across three subsamples

To assess the time-homogeneity of the distribution of the scaled log-returns, we split the sample for each asset into three contiguous subsamples of equal size. With 18 months of data, splitting into three contiguous subsamples yields three 6-month windows per asset (approximately $6 \times 30 \times 24 \approx 4{,}300$ hourly returns per window). For each window we estimate a KDE using a Gaussian kernel and a fixed bandwidth choice applied consistently across subsamples, so that differences reflect the data rather than smoothing artifacts. Notice that there is no further rescaling of the returns in the subsamples, so that the KDEs are on the same scale.

In figure 3, we plot the KDEs for the three subsamples in blue, red, and green. The pooled KDE is overlaid as a black dashed line. Across all tokens, we observe that:

- the three subsample densities are close to each other in terms of location and dispersion;
- the tails of the distributions decay in a broadly similar fashion across subsamples; and

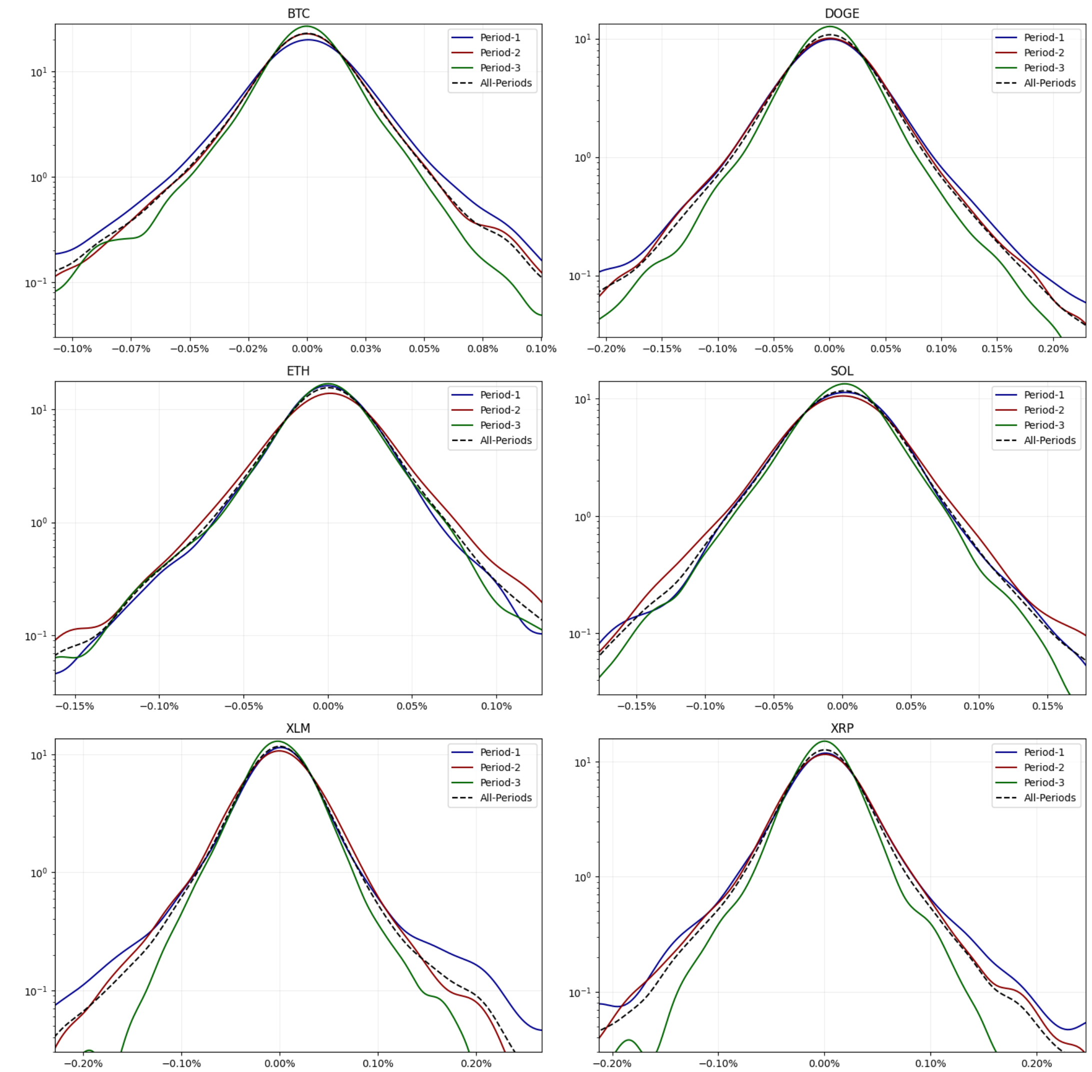

Figure 3: Time-homogeneity diagnostic for 1-hour log-returns. For each token, KDEs are computed on three consecutive 6-month subsamples (colored lines) and compared to the pooled-sample KDE (black dashed line). The close overlap of the subsample KDEs with the pooled KDE indicates approximate time-homogeneity of the marginal distribution over the sample window.

- the pooled KDE essentially interpolates the three individual curves.

These features indicate that the marginal distribution of hourly returns is relatively stable over the sample window, with no obvious regime shifts or large-scale structural breaks. From the perspective of Meucci's framework, this supports the interpretation of hourly returns as approximately time-homogeneous market invariants, at least over medium-term windows of roughly eighteen months.

A systematic pattern remains: in high-volatility market regimes the subsample densities widen for all tokens, and in low-volatility regimes they narrow. A standard remedy is to *volatility-filter* returns by dividing by an estimated conditional volatility (e.g. using GARCH models), estimate the innovation distribution on the filtered series, and then re-scale by the current volatility level when generating scenarios (see, e.g., reference [4]).

We do not implement this volatility filtering here; instead, we use the pooled KDE as an unconditional benchmark and leave conditional extensions to future works.

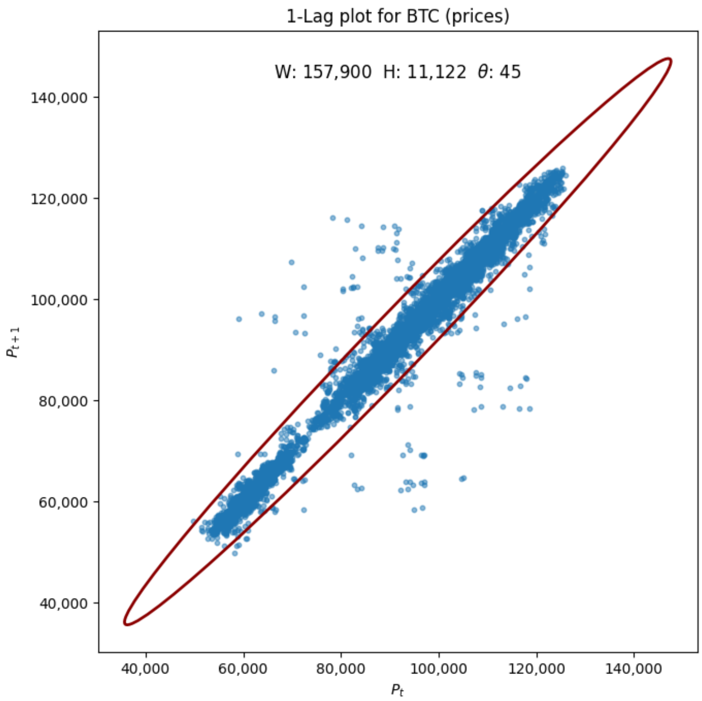

Figure 4: Lag-1 scatter plot for hourly BTC prices: points (P_t, P_{t+1}). The strong diagonal elongation indicates high persistence in price levels, illustrating why prices are not suitable market invariants. The point cloud *cannot* be inscribed in a near-circular ellipse, confirming the presence of serial dependence.

4.2 Lag plots and autocorrelation function

Lag plots provide a visual diagnostic for *serial dependence*. For a time series $\{x_t\}$, the lag-1 plot shows pairs (x_t, x_{t+1}). If there is strong linear dependence, the point cloud becomes elongated along a diagonal direction; if linear dependence is weak, the cloud shows little preferred slope. While lag plots focus on a specific lag, the autocorrelation function provides a more comprehensive view of serial dependence across many lags.

Lag plots for prices

As an example of serial dependence we consider the lag plots for the prices of Bitcoin (BTC). In figure 4 we plot the lag-1 plot for the prices of BTC. We observe a very strong positive association between P_{t+1} and P_t, reflected in a narrow, elongated ellipse along the main diagonal. This is a graphical manifestation of the persistence in price levels: large movements in the underlying price occur over many small consecutive increments rather than instantaneously.

The fact that the lag-1 price plot cannot be inscribed in a near-circular ellipse emphasizes the presence of serial dependence and that raw prices are *not* suitable candidates for market invariants. Their dynamics incorporate strong serial dependence and deterministic trends, which would violate the i.i.d. assumptions.

Lag plots for log returns

We then move to the hourly log-returns $\{X_t\}$ and construct lag plots for a single-period lag. In figure 5 we plot the 1-lag plot for the log-returns for the six tokens listed in Table 1. We observe that in all cases we can inscribe an ellipse around the points that closely approximates a circle. More quantitatively measuring the width W and height H we notice that the values are very close

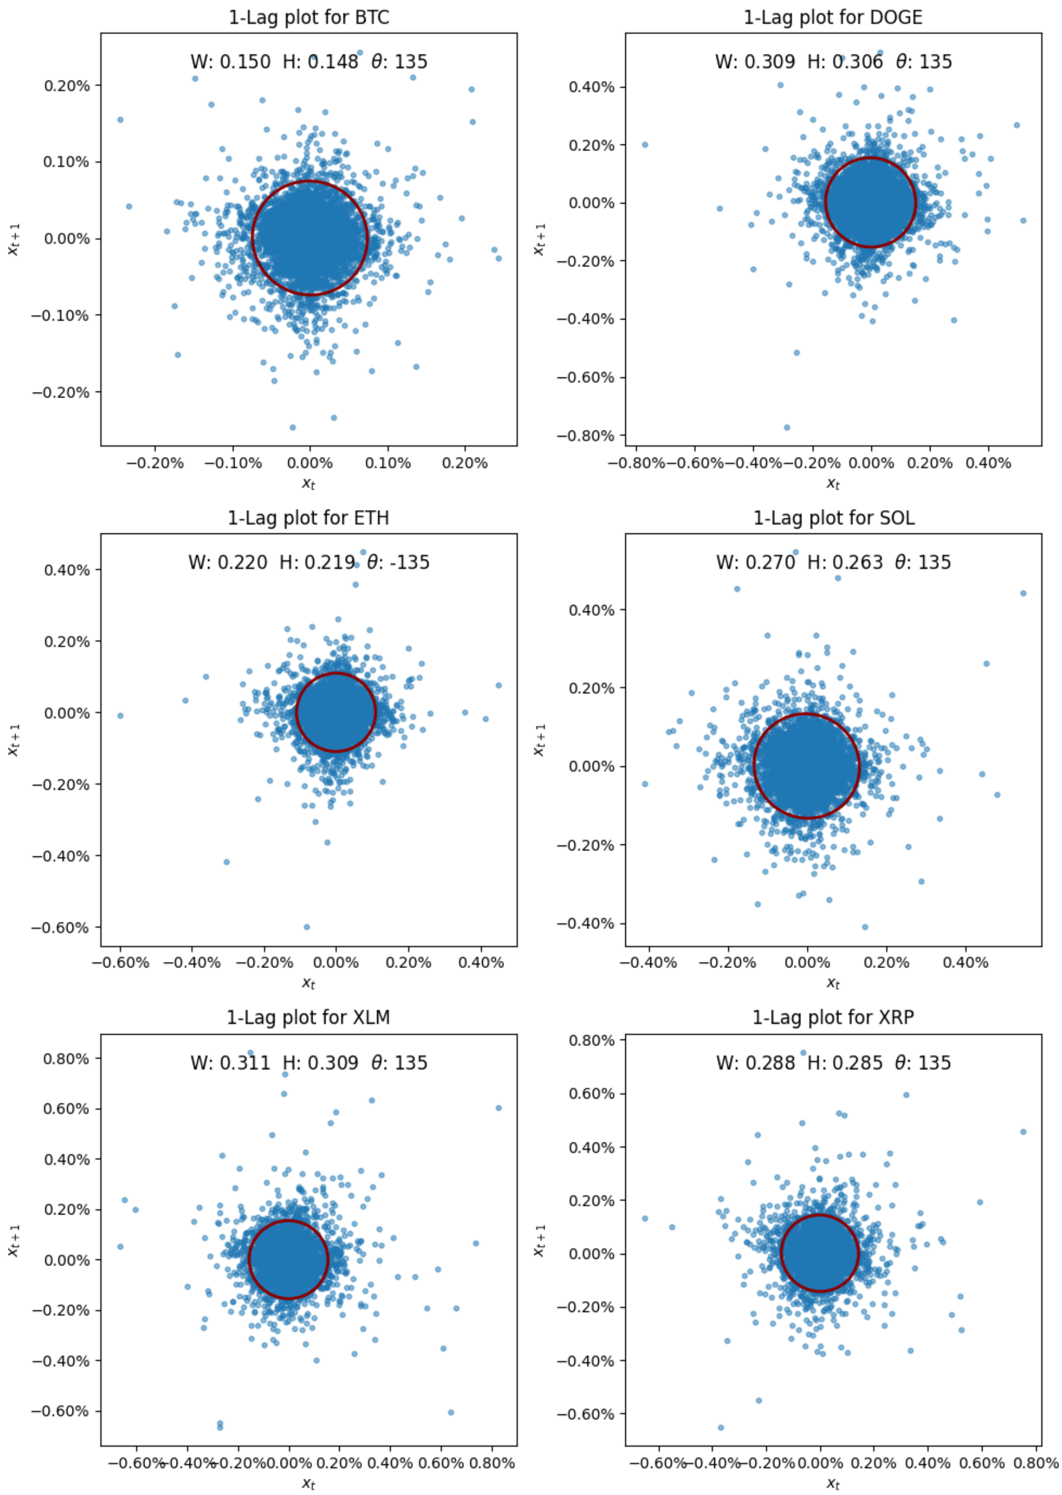

Figure 5: 1-lag scatter plots for 1-hour log-returns: points (Y_t, Y_{t+1}) for each token. The near round, slope-free clouds indicate weak linear dependence at the 1-hour horizon. The point clouds can be inscribed in near-circular ellipses, confirming the approximate independence of the log-returns.

to each other (the angle θ therefore becomes irrelevant). This is a graphical manifestation of the weak linear dependence in log-returns: when large moves occur, they tend to be concentrated within a single hour rather than spread across consecutive hours.

These findings are consistent with the well-documented empirical property that log-returns

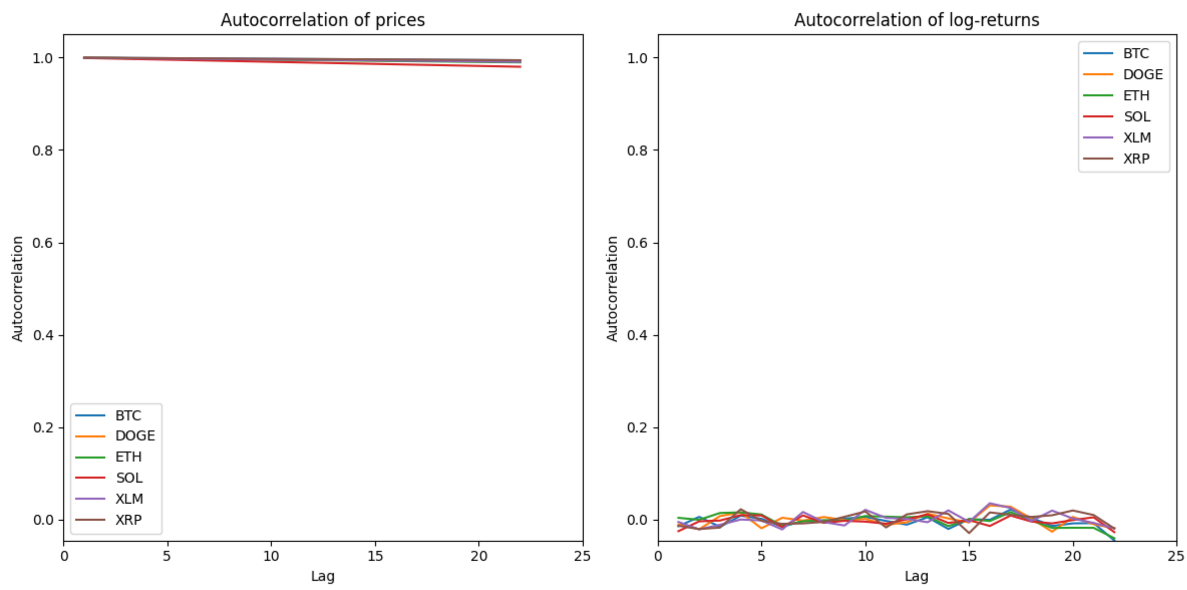

Figure 6: Autocorrelation plots from one hour to one day for both the prices (left panel) and the log-returns (right panel). The two plots are on the same scales to facilitate the comparison.

exhibit little linear autocorrelation. Therefore, we have further evidence that the log-returns are approximately independent increments.

Autocorrelation function for prices and log returns

We now consider the autocorrelation function for the prices and the log-returns. We define the autocorrelation function for the prices and the log-returns as follows:

$$\rho_P(k) = \mathrm{Corr}(P_{t+k}, P_t) \qquad \text{and} \qquad \rho_X(k) = \mathrm{Corr}(X_{t+k}, X_t).$$

A necessary condition for the log-returns to be independent increments is that the autocorrelation function is close to zero for all lags. Since, usually the autocorrelation function decreases monotonically with the lag, we limit to observe it for lags that are between one hour and one day.

In figure 6 we plot the autocorrelation function for the prices (left panel) and the log-returns (right panel) for the six tokens listed in Table 1. In order to compare the two autocorrelation functions, we plot them on the same scales for both the prices and the log-returns.

We observe that the autocorrelation function for the prices is close to one for all tokens and all lags. This is a further confirmation that the prices are not suitable market invariants. On the other hand, the autocorrelation function for the log-returns is close to zero for all tokens and all lags. This is consistent with the findings from the lag plots and give us more confidence that the log-returns are independent increments at least in an approximate sense.

4.3 Aggregation and square-root-of-time scaling

The final diagnostic concerns the behavior of returns under temporal aggregation. Starting from the hourly price series, we construct non-overlapping log-returns at horizons of 1, 2, 4, 8, and 24 hours. Denote these by

$$X_t^{1h}, X_t^{2h}, X_t^{4h}, X_t^{8h}, X_t^{24h}. \tag{4}$$

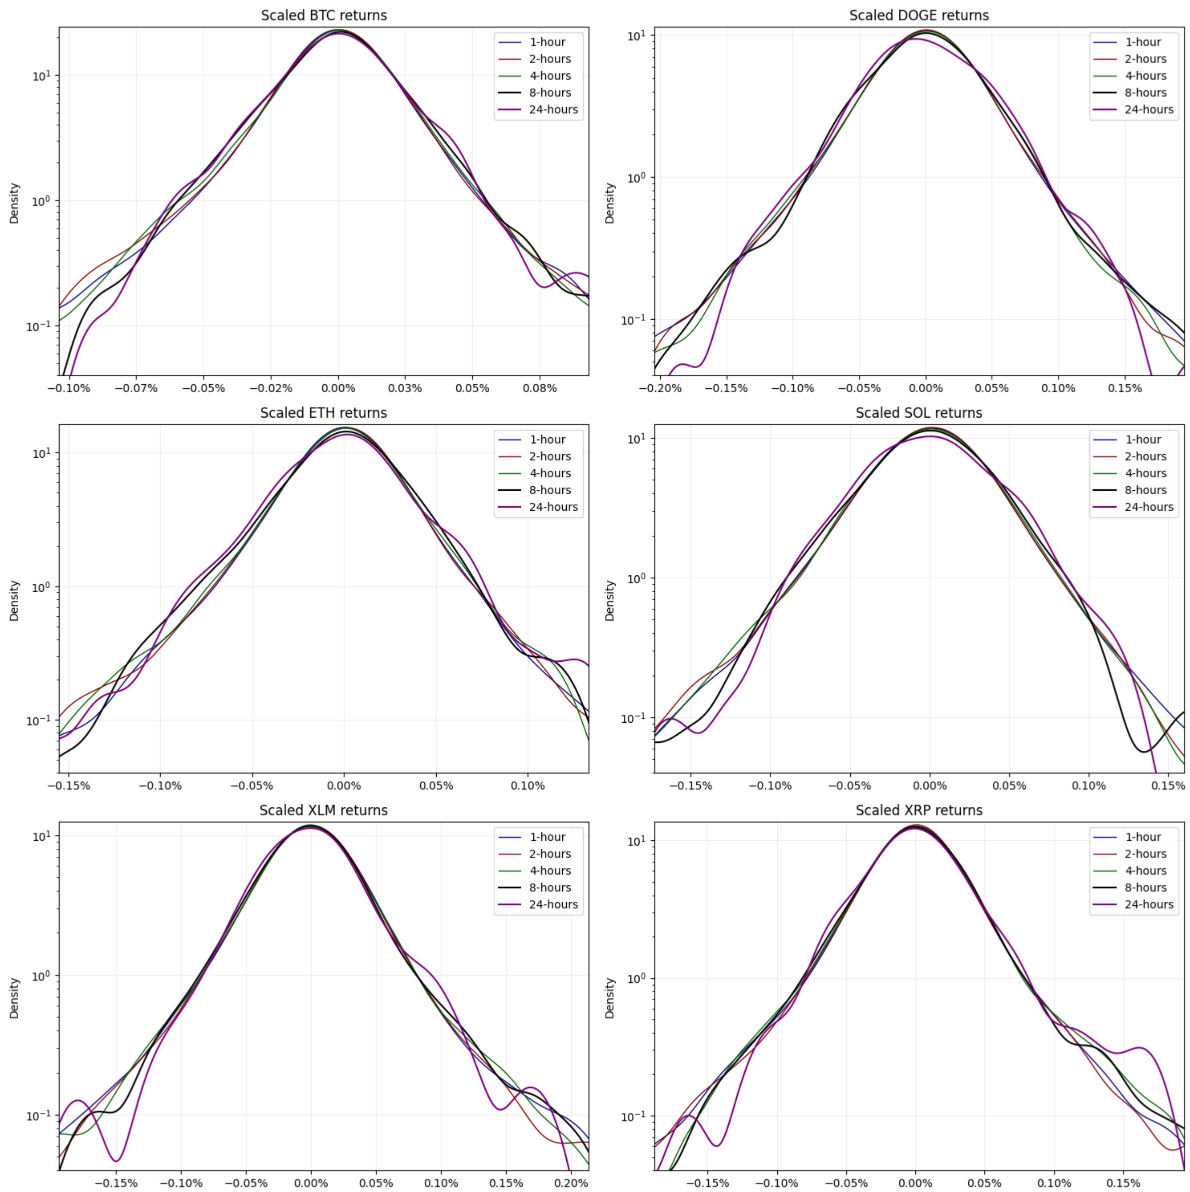

Figure 7: Scaling across horizons. KDEs of non-overlapping k-hour log-returns for $k \in \{1, 2, 4, 8, 24\}$ after rescaling by $\sqrt{24/k}$ to a daily-equivalent horizon. The overlap indicates approximate aggregation/scaling consistency over the 1 hour to 1 day horizons.

For each horizon we estimate a kernel density. To compare the shapes of these distributions on a common scale, we normalize returns by the empirical standard deviation appropriate for the horizon or, equivalently, apply the square-root-of-time scaling rule to bring them to a daily-equivalent standard deviation. In other words, we define

$$
\begin{aligned}
Y_t^{1h} &= X_t^{1h} \cdot \sqrt{24}, \\
Y_t^{2h} &= X_t^{2h} \cdot \sqrt{24/2}, \\
Y_t^{4h} &= X_t^{4h} \cdot \sqrt{24/4}, \\
Y_t^{8h} &= X_t^{8h} \cdot \sqrt{24/8}, \\
Y_t^{24h} &= X_t^{24h} \cdot \sqrt{24/24}.
\end{aligned}
$$

In figure 7 we plot the KDEs for the $\{Y_t^{1h}\}$, $\{Y_t^{2h}\}$, $\{Y_t^{4h}\}$, $\{Y_t^{8h}\}$, and $\{Y_t^{24h}\}$ for all tokens

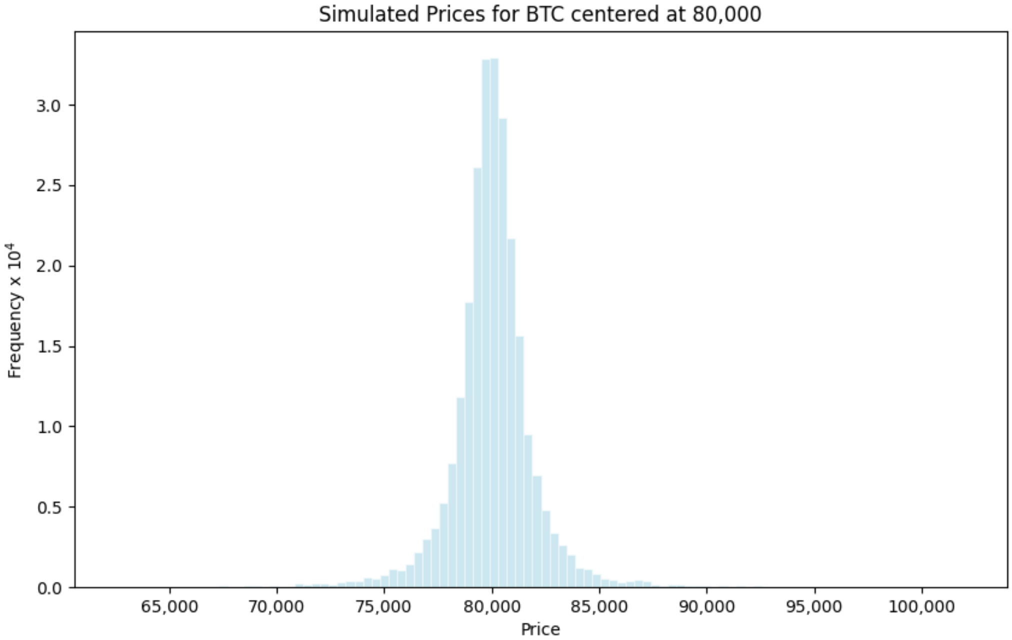

Figure 8: Model-based daily-equivalent BTC price scenarios generated from 1-hour log-returns using the scaling $Y_t = X_t\sqrt{24}$ and $P_t = P_r \exp(Y_t)$ with reference price $P_r = 80{,}000$ USD. These are model-based scenarios (not observed daily returns).

listed in Table 1.

We observe that the KDEs for the different horizons overlap closely over the central part of the distribution. The tails of the distributions also line up reasonably well, though small differences remain, particularly in the far tails where data is sparse at longer horizons. The 1-hour distribution naturally exhibits more data points in the tails simply because it is based on a larger number of observations; this improves the resolution of extreme events without necessarily implying heavier tails per se.

Overall, the scaling behavior is broadly consistent with the square-root-of-time rule between 1 hour and 1 day for the assets considered. This supports the idea that hourly invariants can be used as a flexible building block for modeling risk at daily horizons, either via explicit aggregation of multiple hourly steps or via scaling approximations.

Scaling 1-hour invariants to a daily-equivalent horizon increases the number of available scenarios by a factor of 24 relative to using one observation per day. Therefore, for liquid tokens the 1-hour sampling frequency provides a substantially richer empirical basis for estimating tail events than daily returns over the same calendar span.

5 Conclusion

In this paper we investigated whether hourly log-returns on major digital assets can be viewed as market invariants in the sense of Meucci [1]. Using 18 months of 24/7 hourly prices for six liquid tokens listed in Table 1, we examined the time-homogeneity, dependence structure, and scaling behavior of the returns.

First, by splitting the sample into three contiguous subsamples and comparing their kernel density estimates, we found that the marginal distribution of hourly returns is relatively stable across time. The three subsample KDEs largely overlap, and their tails decay in a comparable fashion, with the pooled KDE tracing their common shape closely.

Second, lag plots revealed a sharp contrast between prices and returns. Hourly prices produced an elongated diagonal ellipse in the 1-lag scatter plot (see figure 4), highlighting strong serial dependence and confirming that prices are not suitable invariants. In contrast, lag plots for hourly returns (see figure 5) produced point clouds that were well contained within near-circular ellipses, indicative of weak linear autocorrelation and approximately independent increments at these horizons. We then computed the autocorrelation function for the prices and the log-returns and observed that it is close to one for prices and close to zero for log returns. This further confirms that the prices are not suitable market invariants and that the log-returns are approximately independent increments.

Third, we studied the aggregation of returns across time by comparing KDEs at 1, 2, 4, 8, and 24-hour horizons. After appropriate square-root-of-time rescaling, the distributions overlapped closely, with the 1-hour returns naturally providing the richest information in the tails. This behavior is broadly consistent with diffusion-like scaling and supports the use of hourly invariants to model risk at the daily horizon.

For example, one may use the scaled returns to generate a large number of independent daily scenarios, as in equation (3). Figure 8 shows the histogram of the price scenarios for BTC using the hourly returns scaled by $\sqrt{24}$ applied to a reference price of $P_r = 80,000$ USD. The histogram is shown on a linear scale.

Similar scenarios could also be computed for all the other tokens in the portfolio. Once the token price scenarios are multiplied by the corresponding portfolio quantities, portfolio value scenarios can be computed. At this point we are able to compute risk and gain measures at the daily horizon as shown in reference [3].

From a risk-engineering perspective, the main implication is that hourly sampling can materially increase the effective scenario set available for nonparametric risk measurement in 24/7 markets. When hourly returns are approximately invariant, one can estimate a joint innovation distribution at the hourly frequency and project it to a daily-equivalent horizon for portfolio repricing and tail-risk measurement. For portfolio applications, scenarios should be generated jointly across tokens to preserve cross-asset dependence.

References

[1] A. Meucci. *Risk and Asset Allocation*. Springer Finance, Springer-Verlag, Berlin Heidelberg, 2005. 2, 11

[2] A. Meucci. The Prayer: Ten-Step Checklist for Advanced Risk and Portfolio Management. (Available via GARP/ARPM white paper versions), 2011. 3

[3] Marco Marchioro, *Finite-Sample Distortion Measures: Unified Risk and Gain via Scenario Weights*, Talos Quant Research, `https://www.talos.com/insights`, February 2026. 3, 12

[4] G. Barone-Adesi, K. Giannopoulos, and L. Vosper. VaR without correlations for nonlinear portfolios. *Journal of Futures Markets*, 19(5):583–602, 1999. 6

talos.com