“Adaptive market hypothesis and intraday dynamics: High-frequency evidence from the Casablanca Stock Exchange”

AUTHORS	Dina Sittel Achraf Louati

ARTICLE INFO	Dina Sittel and Achraf Louati (2026). Adaptive market hypothesis and intraday dynamics: High-frequency evidence from the Casablanca Stock Exchange. *Investment Management and Financial Innovations*, *23*(4), 218–234. doi:10.21511/imfi.23(4).2026.14
DOI	http://dx.doi.org/10.21511/imfi.23(4).2026.14
RELEASED ON	Tuesday, 03 November 2026
RECEIVED ON	Friday, 08 May 2026
ACCEPTED ON	Wednesday, 21 October 2026

LICENSE	This work is licensed under a Creative Commons Attribution 4.0 International License

JOURNAL	"Investment Management and Financial Innovations"
ISSN PRINT	1810-4967
ISSN ONLINE	1812-9358
PUBLISHER	LLC “Consulting Publishing Company “Business Perspectives”
FOUNDER	LLC “Consulting Publishing Company “Business Perspectives”

NUMBER OF REFERENCES	NUMBER OF FIGURES	NUMBER OF TABLES
38	2	8

© The author(s) 2026. This publication is an open access article.

BUSINESS PERSPECTIVES

LLC "CPC "Business Perspectives"
Hryhorii Skovoroda lane, 10,
Sumy, 40022, Ukraine

www.businessperspectives.org

Dina Sittel (Morocco), **Achraf Louati** (Morocco)

ADAPTIVE MARKET HYPOTHESIS AND INTRADAY DYNAMICS: HIGH-FREQUENCY EVIDENCE FROM THE CASABLANCA STOCK EXCHANGE

Abstract

Market efficiency remains a central issue in financial economics, particularly in increasingly automated financial markets. The Adaptive Market Hypothesis suggests that efficiency evolves as market conditions change, making high-frequency analysis useful for capturing short-lived changes in return behavior. This study examines whether the Moroccan stock market exhibits time-varying short-horizon return dynamics using Moroccan All Shares Index data from August 2016 to May 2025.

The final database contains 1,222,039 unique observations. Because timestamps are irregular, we align prices to fixed intervals within continuous trading before constructing returns. The analysis combines rolling multiple variance-ratio tests with wild-bootstrap inference, a bootstrap-centered measure of random-walk departure, regression analysis, and non-Bayesian time-varying autoregressive modeling.

The results show that departures from the random-walk benchmark vary markedly over time, with similar patterns across 5-, 10-, and 15-minute frequencies. In the main 10-minute specification, the joint random-walk null is rejected in 373 of 2,814 rolling windows (13.3%). Intraday differences are modest, but formal tests do not establish statistically significant differences across the three periods. Realized volatility is positively associated with larger contemporaneous departures, whereas associations with selected calendar and event conditions are mixed and specification-dependent. Formal parameter-constancy and slope-stability tests reject stability (slope-bootstrap $p = 0.0091$), while estimated linear predictability remains below the pointwise 95% bootstrap upper bound throughout the sample. The findings are consistent with the Adaptive Market Hypothesis and show that changing return dynamics do not necessarily generate persistent predictability.

Keywords Adaptive Market Hypothesis, market efficiency, return predictability, realized volatility, variance ratio, time-varying autoregression, emerging stock market, Morocco

JEL Classification G14, G15, G12, C58

Type of the article: Research Article

Received on: 8th of May, 2026
Accepted on: 21st of October, 2026
Published on: 3rd of November, 2026

© Dina Sittel, Achraf Louati, 2026

Dina Sittel, Ph.D. Student, Faculty of Legal, Economic and Social Sciences, Mohammed V University in Rabat, Morocco. (Corresponding author)

Achraf Louati, Full Professor, Faculty of Legal, Economic and Social Sciences, Mohammed V University in Rabat, Morocco.

This is an Open Access article, distributed under the terms of the Creative Commons Attribution 4.0 International license, which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

Conflict of interest statement:
Author(s) reported no conflict of interest

INTRODUCTION

Market efficiency remains a central question in financial literature. The Efficient Market Hypothesis (EMH) defines market efficiency as the extent to which asset prices reflect available information. The assumptions underlying this theory have been consistently challenged empirically, paving the way for the introduction of the Adaptive Market Hypothesis (AMH), which posits that market efficiency is dynamic and adaptive.

This adaptive perspective is especially relevant to intraday financial markets, where price behavior can be observed at increasingly fine time intervals. Hence, high-frequency data provide an opportunity to examine short-lived departures from random-walk behavior that may remain hidden at lower sampling frequencies. However, high-frequen-

http://dx.doi.org/10.21511/imfi.23(4).2026.14

cy index observations do not directly identify the activities of high-frequency traders, and observed return dependence may also reflect market-microstructure effects.

The Casablanca Stock Exchange (CSE) provides a relevant setting for examining these dynamics in an emerging equity market. Existing research has examined market efficiency in Morocco, but evidence based on high-frequency intraday index returns remains limited. Examining departures from the random-walk benchmark alongside the stability of autoregressive relationships may therefore provide a more comprehensive understanding of short-horizon return dynamics than relying on a single measure of market efficiency. Consequently, an important question remains unresolved: to what extent do departures from the random-walk benchmark vary over time and across the trading day, and are these changes accompanied by persistent, statistically detectable linear predictability?

1. LITERATURE REVIEW

One key dimension of market quality is pricing efficiency, commonly associated with the Efficient Market Hypothesis (EMH), which implies that prices incorporate available information (Fama, 1970). Although the EMH has been the dominant theory of financial markets for decades, its assumptions have been challenged both theoretically and empirically. Grossman and Stiglitz (1980) argue that fully informative prices are incompatible with incentives to acquire costly information, while Lo and MacKinlay (1988) document departures from random-walk behavior. The random-walk benchmark and informational efficiency are nevertheless not equivalent concepts: statistical return dependence does not, by itself, establish exploitable informational inefficiency, and the absence of detectable dependence cannot establish complete efficiency.

The Adaptive Market Hypothesis (AMH) offers a dynamic alternative by combining market efficiency with evolutionary adaptation (Lo, 2004). Rather than treating efficiency as a fixed state, the AMH views return predictability and arbitrage opportunities as time-varying outcomes shaped by changing market conditions, competition, and investor adaptation (Lo, 2012). The evolution of stock-market efficiency has been widely discussed in the empirical literature (Lim & Brooks, 2011). However, time variation in predictability does not uniquely identify an adaptive mechanism. Hřebačka (2025) argues that relationships between return predictability, business cycles, and risk aversion can also be interpreted within the EMH. These perspectives place the interpretation of changing return dependence, rather than its observation alone, at the center of the empirical debate.

Empirical studies have documented changes in return predictability across developed and emerging equity markets, providing evidence consistent with some implications of the AMH (Kim et al., 2011; Smith, 2012; Urquhart & McGroarty, 2016). Yet the results are not uniform across markets, time periods, or forms of dependence. Urquhart and Hudson (2013) distinguish variation in linear predictability from more persistent nonlinear dependence, while Urquhart and McGroarty (2016) report that relationships between predictability and market conditions differ across major stock indices. The literature therefore points to a changing and multidimensional phenomenon rather than a simple alternation between efficient and inefficient market states. In particular, evidence on linear return dependence does not establish whether other forms of predictability are absent or whether any statistical departure could be profitably exploited.

A prominent strand of the empirical literature investigates random-walk behavior through variance-ratio testing (Lo & MacKinlay, 1988), while subsequent contributions have emphasized the importance of joint inference and finite-sample reliability (Chow & Denning, 1993; Kim, 2006). Nevertheless, a changing sequence of test outcomes cannot be interpreted independently of the inference that produces it. Hřebačka (2026) questions conclusions drawn from repeated rolling-window testing because multiple testing can create apparent changes even when individual tests are correctly conducted. This criticism highlights a distinction between observing variation in test statistics and establishing variation in the underlying return-generating process.

Related research approaches the question of changing return behavior through the stability of its autoregressive relationships. Hansen (1992) develops formal tests for parameter instability, while Ito et al. (2016) and Noda (2016) examine evolving stock-market return dependence using time-varying autoregressive models. Formal parameter-instability tests and time-varying autoregressive models examine changes in estimated return relationships, while variance-ratio tests assess departures from the random-walk benchmark. The distinction is economically important because variation in model coefficients does not necessarily imply persistent, statistically detectable linear predictability. Evidence from these approaches can consequently be informative even when their implications are not identical.

High-frequency financial data have become increasingly important in examining short-term price behavior and intraday market dynamics. Andersen et al. (2001) demonstrate how high-frequency returns can be used to study realized exchange-rate volatility. However, the use of more frequent observations also introduces challenges related to market-microstructure noise. Aït-Sahalia et al. (2005) and Hansen and Lunde (2006) examine how such noise affects the estimation of volatility using high-frequency data. These findings highlight the importance of considering potential measurement distortions when analyzing short-term financial market dynamics.

Market-microstructure considerations are particularly relevant when prices do not update synchronously. Lo and MacKinlay (1990) show that infrequent and nonsynchronous trading can affect the variance and autocorrelation properties of individual and portfolio returns. These effects are relevant to equity indices, whose constituent securities may differ in trading frequency, without implying that they account for every observed departure from a random walk. Intraday variation also has an economic foundation. Admati and Pfleiderer (1988) explain how interactions between informed and liquidity-motivated traders can generate concentrated patterns of volume and price variability. Together, these contributions suggest that short-horizon dependence can reflect both changing trading conditions and the mechanisms through which prices are observed, warranting care when interpreting high-frequency index evidence.

High-frequency trading (HFT) has become an important part of modern electronic markets, with research examining its implications for price discovery and trading costs. Carrion (2013) finds that prices incorporate information more efficiently on NASDAQ days with greater identified HFT participation, while Brogaard et al. (2014) document a role for HFT in price discovery. At the same time, Biais et al. (2015) show how investment in trading speed can generate adverse-selection externalities. These findings concern trading activity and market organization, and their interpretation should remain distinct from evidence based solely on frequently observed price indices.

Recent research has also examined changing market efficiency at high frequencies. Chu et al. (2019) report time variation in efficiency measures for high-frequency cryptocurrency returns, while Meng and Li (2021) document clustering in one-minute informational-efficiency measures for the SPY exchange-traded fund and an association with activity inferred from order-book messages. Such evidence extends the AMH discussion beyond daily equity returns but also reflects the availability of different empirical information across markets. Studies using orders or identifiable trading activity can investigate mechanisms that an index-price series alone cannot reveal.

The relationship between return predictability and market conditions remains similarly heterogeneous. Kim et al. (2011) find that predictability in long-run US equity returns is associated with volatility and economic fundamentals, but differs across crashes, crises, and bubbles. Urquhart and McGroarty (2016) likewise find that different stock markets respond differently to market conditions. In high-frequency cryptocurrency markets, Chu et al. (2019) report that news sentiment and event types are not necessarily significant determinants of measured efficiency. These findings support examining volatility and changes in the trading environment without presuming that shocks or calendar conditions have uniform effects on short-horizon return dependence.

Despite these contributions, much of the existing high-frequency equity-market literature focuses on developed and highly liquid markets, while comparatively less evidence is available on high-

http://dx.doi.org/10.21511/imfi.23(4).2026.14

frequency return dynamics in emerging markets. The literature also shows that conclusions may depend on sampling frequency and empirical approach, and that changes in return dependence do not necessarily imply persistent predictability. The Casablanca Stock Exchange therefore provides a relevant setting for examining these dynamics using high-frequency data.

The purpose of this study is to examine whether the Moroccan stock market exhibits time-varying return dynamics consistent with the Adaptive Market Hypothesis using high-frequency data and to assess how departures from the random-walk benchmark vary over time and across the trading day, and how their magnitude is associated with realized volatility and selected calendar and event conditions. The study also examines whether changes in autoregressive relationships are accompanied by persistent, statistically detectable linear predictability. A central contribution of the study is therefore to distinguish changes in short-horizon return dynamics from evidence of persistent predictability, rather than relying on a simple efficient/inefficient classification. The following hypotheses are investigated:

H1: *Departures from the random-walk benchmark vary over time in the Moroccan stock market.*

H2: *The magnitude of departures from the random-walk benchmark exhibits intraday variation across early, middle, and late continuous-trading periods.*

H3: *Realized volatility is significantly associated with the magnitude of departures from the random-walk benchmark.*

H4: *The magnitude of departures from the random-walk benchmark is associated with selected calendar and event conditions in the Moroccan stock market.*

2. METHODS

This study combines rolling multiple variance-ratio tests, regression analysis, and a non-Bayesian time-varying autoregressive model. The variance-ratio analysis examines temporal and intraday departures from random-walk behavior, with sampling-frequency robustness checks. Regression analysis then assesses their association with realized volatility, calendar conditions, and major events using heteroskedasticity and autocorrelation-consistent (HAC) standard errors, while the time-varying model examines the evolution of short-horizon return predictability.

2.1. Institutional context and market structure

The CSE follows a schedule divided into specific trading sessions, as shown in Table 1. These segments may vary in liquidity, order execution methods, and market behavior, all of which are important for evaluating intraday efficiency through the lens of the AMH.

2.2. Data and sample design

The high-frequency Moroccan All Shares Index (MASI) series was obtained from the Casablanca Stock Exchange (CSE) under a non-disclosure agreement. The initial sample contains 1,410,953 observations from August 1, 2016, to May 16, 2025. Timestamps are recorded to the second but are

Table 1. Summary of trading sessions at the Casablanca Stock Exchange

Session	Regular Days	Ramadan
Pre-Trading	08:10 – 09:00	08:10 – 09:15
Opening Auction Call	09:00 – 09:30 + T0*	09:15 – 10:00 + T0*
Continuous Trading	09:30 + T0* – 15:20	10:00 + T0* – 13:20
Closing Auction Call	15:20 – 15:30 + T1*	13:20 – 13:30 + T1*
Closing Price Publication	+1 min after auction	+1 min after auction
Closing Price Cross	+9 min after publication	+9 min after publication
Post-Close	+15 min after cross	+15 min after cross

Note: T0 and T1 are random durations determined automatically for each instrument. Their maximum permitted duration varies throughout the sample period according to the applicable CSE trading rules.

irregularly spaced and are not treated as a one-second return series. Exact timestamp duplicates were removed within each source file by retaining the last observation in source order. After merging and alignment with the CSE trading calendar, the final sample contains 1,222,039 unique timestamps across 2,188 trading dates. Continuous trading represents 86.16% of the sample (Table 2); therefore, the return analysis is restricted to this phase for comparability.

Table 2. Distribution of MASI observations by market phase

Market Phase	Observations	Share
Continuous trading	1,052,850	86.16%
Closing price phase	81,958	6.71%
Closing transition	43,727	3.58%
Opening transition	38,877	3.18%
Out-of-market	4,369	0.36%
Closing auction	241	0.02%
Post-close / late phase	17	0.00%
Total	1,222,039	100.00%

To construct returns over fixed time intervals, MASI values are sampled on a regular clock-time grid using the most recent available value within the same trading day and continuous trading session. Log returns are calculated as follows:

$$r_{d,t}^{(\Delta)} = \ln\left(P_{d,t}\right) - \ln\left(P_{d,t-\Delta}\right), \quad (1)$$

where $P_{d,t}$ denotes the MASI index value on trading day (d) at time (t), and (Δ) represents the sampling interval. Returns are not calculated across trading-day or session boundaries.

The main analysis uses ten-minute returns. At the 1-min frequency, 54.80% of returns are equal to zero, compared with 13.34% at 5-min, 4.04% at 10-min, and 1.49% at 15-min. The ten-minute interval therefore balances reducing stale observations with preserving intraday information. 5- and 15-minute returns are used as robustness checks.

2.3. Rolling multiple variance ratio test

To assess departures from the random-walk benchmark, variance ratios are constructed from ten-minute MASI returns following the general framework of Lo and MacKinlay (1988). Under the random-walk hypothesis, the variance of a (q)-period return should equal (q) times the variance of the one-period return. The variance ratio (VR) for horizon (q) is therefore defined as:

$$VR(q) = \frac{Var\left(\sum_{j=0}^{q-1} r_{t-j}\right)}{q \cdot Var\left(r_t\right)}, \quad (2)$$

where ($VR(q) = 1$) under the random-walk benchmark. Values above or below 1 indicate positive or negative serial dependence, respectively.

The test is implemented over rolling windows of 250 ten-minute returns, with a step of 25 returns. Within each window, variance ratios are estimated jointly for ($q = \{2,5,10\}$), corresponding to 20, 50, and 100-minute horizons. Although rolling windows may span several trading days, each (q)-period return is constructed only when all underlying observations belong to the same trading day, preventing overnight returns from entering the test.

Statistical inference is based on 10,000 wild-bootstrap replications using centered residuals and Rademacher multipliers. For each rolling window, the maximum absolute studentized statistic across the three variance ratios is compared with its bootstrap distribution. The random-walk null is rejected at the 5% level when the joint statistic exceeds the corresponding bootstrap critical value. As an additional size-calibration check, the procedure was applied to independently simulated null windows at each sampling frequency.

In addition to the joint rejection indicator, the magnitude of random-walk departure is measured using the bootstrap-centered variance-ratio deviation (BCDEV) as follows:

$$BCDEV_i = \frac{1}{3} \sum_{q \in \{2,5,10\}} \left| VR_i(q) - \overline{VR^*}_{i,0}(q) \right|. \quad (3)$$

where $\overline{VR^*}_{i,0}(q)$ denotes the mean variance ratio obtained from the bootstrap null distribution for window (i).

BCDEV is used in the subsequent regression analysis. The rolling procedure is also estimated using

http://dx.doi.org/10.21511/imfi.23(4).2026.14

5- and 15-minute returns for sampling-frequency robustness. For the intraday analysis, each continuous-trading day is divided chronologically into early, middle, and late thirds, allowing the classification to adjust to Ramadan and other shortened schedules. The intraday analysis uses q = {2,5}, as q = 10 does not provide sufficient same-segment observations. All multi-period returns remain within the same trading day and segment.

Differences in mean BCDEV and joint-rejection rates across the three intraday periods are examined using a calendar-day block bootstrap with 10,000 replications to account for dependence between overlapping windows.

2.4. Regression analysis of random-walk departures

To examine the factors associated with variations in random-walk departures, a multivariate regression model is estimated using the bootstrap-centered variance-ratio deviation defined in Equation (3). Since the distribution of this measure is right skewed, its logarithmic form is used as the dependent variable in the main specification, reducing the influence of extreme observations and allowing the estimated coefficients to be interpreted in proportional terms. The baseline model is expressed as follows:

$$\begin{aligned}&\ln\left(BCDEV_i\right)=\alpha+\beta_1\log\left(RV_i\right)\\&+\beta_2\cdot Ramadan_i+\beta_3\cdot Covid_i\\&+\sum_{y=2017}^{2025}\gamma_y Year_{y,i}+\varepsilon_i,\end{aligned}\tag{4}$$

where (*i*) denotes a rolling window. Realized volatility is calculated over the same 250 ten-minute returns used in the variance-ratio analysis, following the standard realized-volatility approach based on the sum of squared high-frequency returns (Andersen et al., 2001)

$$RV_i=\sum_{j=1}^{250}r^2_{i,j},\tag{5}$$

where $r_{i,j}$ denotes return (*j*) within window (*i*). The logarithm of realized volatility is included to capture variation in intraday market volatility. ($Ramadan_i$) represents the proportion of returns in each window occurring during Ramadan, while ($Covid_i$) measures the proportion falling within the shortened trading-hours period from March 24 to June 30, 2020. Year fixed effects are included for 2017–2025, with 2016 serving as the reference year.

An extended specification incorporates major economic and market events relevant to the Moroccan market:

$$\begin{aligned}&\ln\left(BCDEV_i\right)=\alpha+\beta_1\log\left(RV_i\right)\\&+\beta_2\cdot Ramadan_i+\beta_3\cdot Covid_i+\beta_4\cdot Boycott_i\\&+\beta_5.Earthquake_i+\beta_6\cdot Ratecut_i\\&+\beta_7\cdot Tariff_i+\sum_{y=2017}^{2025}\gamma_y Year_{y,i}+\varepsilon_i.\end{aligned}\tag{6}$$

Each event variable measures the proportion of returns within window (*i*) falling inside the corresponding event period. The consumer boycott is treated as an extended episode from April 20 to August 31, 2018. For the Marrakech earthquake, the June 2024 policy-rate cut, and the April 2025 U.S. tariff announcement, a common short-event horizon is used, corresponding to the first 10 valid CSE trading days following each event. The resulting event periods are reported in Table 3.

As a sensitivity check, the earthquake, rate-cut, and tariff event windows are also defined using 5 and 20 trading days alongside the 10-day baseline, with all three event definitions varied together in each specification.

Because consecutive rolling windows overlap substantially, coefficient inference is based on Newey-West HAC standard errors (Newey & West, 1987), with a lag length of 10. Robustness is assessed using a lag length of 20 and by re-estimating the models with the untransformed (BCDEVi) measure. As robustness checks, the regressions are re-estimated using 10 non-overlapping subsamples obtained from alternative starting offsets. Realized volatility from the preceding non-overlapping window is also examined, with contemporaneous and previous-window specifications compared to a common sample.

Multicollinearity is evaluated using variance inflation factors and condition diagnostics (Belsley et al., 1980; Kutner et al., 2005). The estimated coefficients are interpreted as conditional associations with the magnitude of random-walk departures

Table 3. Calendar-based event variables

Event Name	Variables	Event Period	Description
Consumer Boycott	boycott	April 20 – August 31, 2018	Consumers nationwide boycotted certain brands due to rising living costs
COVID-19	COVID	March 24 – June 30, 2020	CSE shortened trading hours during the COVID-19 period
Marrakech Earthquake	earthquake	September 11-22, 2023	A 6.8-magnitude earthquake near Marrakech resulted in major human and economic impacts
Central Bank Rate Cut	ratecut	June 26 – July 9, 2024	Bank Al-Maghrib reduced the policy rate from 3% to 2.75%
U.S. Tariff	us_tariff	April 3-16, 2025	The U.S. imposed a 10% tariff, affecting international trade

and not as causal effects of the corresponding market conditions or events.

2.5. Time-varying autoregressive framework (TV-AR)

The VR analysis identifies variations in random-walk departures across the rolling windows. While it serves its purpose, it remains dependent on the selected window length and does not directly assess whether the underlying return dynamics are stable over time. To complement this analysis, the non-Bayesian TV-AR framework introduced by Ito et al. (2014) was used, as it allows autoregressive coefficients to change over time and provides a continuous assessment of short-horizon return predictability consistent with the Adaptive Market Hypothesis.

Starting with the estimation of a stationary autoregressive model of order (p) under the assumption of time-invariant parameters, the model is specified as follows:

$$r_t = \alpha_0 + \sum_{j=1}^{p} \alpha_j r_{t-j} + u_t, \quad (7)$$

for $t = 1, 2, \ldots, T$; where (r_t) represents the ten-minute log return at time (t), (α_0) denotes the intercept, (α_j) represents the autoregressive coefficient at lag (j), and (u_t) is the error term. Autoregressive lags are constructed only within the same trading day to avoid overnight return effects. To determine the lag order (p), specifications from AR(1) to AR(12) are compared using the Akaike Information Criterion (AIC; Akaike, 1973) and the Bayesian Information Criterion (BIC) of Schwarz (1978), and residual serial-correlation diagnostics. Since the information criteria do not converge on the same lag, the smallest specification with an AIC difference of no more than two from the minimum and no significant residual autocorrelation at Ljung-Box lags 5 and 10 is selected. This procedure results in an AR(8) specification. The sensitivity of the results to lag selection is further examined by estimating an AR(12) specification and comparing it with AR(8) over a common sample.

With the potential for evolving market dynamics and structural changes, the constant AR model is generalized by allowing the autoregressive coefficients to vary over time:

$$r_t = \alpha_0 + \sum_{j=1}^{8} \alpha_{j,t} r_{t-j} + u_t, \quad (8)$$

where $(\alpha_{j,t})$ denotes the time-varying coefficient for lag (j) at time (t). Following Ito et al. (2014), the coefficients are modeled as exhibiting random-walk dynamics:

$$\alpha_{j,t} = \alpha_{j,t-1} + v_{j,t}, \quad (9)$$

for $j = 1, \ldots, 8$; where $(v_{j,t})$ represents the innovation in the time-varying coefficient. The covariance matrix of these innovations is estimated as a full covariance matrix, allowing the coefficient innovations to vary jointly over time.

The model is estimated using the non-Bayesian procedure of Ito et al. (2014), beginning with a first step least-squares estimation followed by one-step feasible generalized least squares. The degree of short-horizon linear predictability at time (t) is then measured using the following metric:

$$\zeta_t = \left| \frac{\sum_{j=1}^{8} \alpha_{j,t}}{1 - \sum_{j=1}^{8} \alpha_{j,t}} \right|, \quad (10)$$

where values of (ζ_t) closer to zero indicate smaller departures from the no-linear-predictability benchmark, while larger values indicate stronger return predictability. The series of estimated (ζ_t) values are evaluated using a residual bootstrap with 2,000 replications under the null of zero au-

http://dx.doi.org/10.21511/imfi.23(4).2026.14

toregressive predictability. The bootstrap preserves the trading-day structure and prevents overnight lags from entering the estimation. A pointwise one-sided 95% upper critical band is then constructed to identify whether the estimated degree of predictability exceeds the level expected under the null at each point in time.

Moreover, to assess whether the use of time-varying parameters is empirically justified, Hansen's (1992) joint (L_C) parameter-constancy test is applied to the constant AR(8) specification. The null hypothesis assumes that the model parameters remain constant over the sample period, while rejection of the null indicates parameter instability in at least one of the parameters included in the joint test. The stability of the autoregressive coefficients is also examined using a wild-bootstrap score test with 10,000 replications, including resampling at the trading-day level.

3. RESULTS

The empirical results are presented in four stages. The analysis begins with sampling-frequency diagnostics, followed by the temporal and intraday variance-ratio results, the regression analysis of market conditions, and the time-varying autoregressive results.

3.1. Return distribution and sampling-frequency diagnostics

Table 4 compares MASI returns across 1-, 5-, 10-, and 15-minute frequencies to identify a sampling interval that balances price variation and intraday sample size. Zero returns decline from 54.80% at 1 minute to 13.34% at 5 minutes, 4.04% at 10 minutes, and 1.49% at 15 minutes. The 10-minute series retains 70,587 observations while reducing unchanged values and is therefore used as the main specification. The 5- and 15-minute frequencies are retained for robustness.

3.2. Rolling multiple VR and intraday dynamics

The rolling multiple variance-ratio analysis addresses *H1* and *H2* by examining temporal and intraday variation in departures from the random-walk benchmark.

Using the ten-minute specification, the analysis produces 2,814 rolling windows. The mean variance ratios declined from 0.958 for q = 2 to 0.867 for q = 5 and 0.773 for q = 10, indicating that the departure from unity becomes more pronounced as the return horizon increases. The joint random-walk null is rejected in 373 windows, corresponding to 13.26% of the sample. Rejection is therefore concentrated in a minority of the rolling windows rather than occurring uniformly throughout the sample. At the same time, the mean bootstrap-centered variance-ratio deviation BCDEV is 0.169, showing that the magnitude of departure continues to vary even when the joint null is not rejected.

Figure 1 shows that the magnitude of random-walk departure varies unevenly over time, with periods of low BCDEV interrupted by shorter increases. Joint rejections are concentrated around

Table 4. Descriptive statistics of MASI returns across sampling frequencies

Freq.	N	Mean	Std. Dev.	Zero returns (%)	Min.	Max.
1 min	724,902	–0.00000036	0.000279	54.80	–0.01310	0.01176
5 min	143,631	–0.00000273	0.000619	13.34	–0.01311	0.01348
10 min	70,587	–0.00000736	0.000869	4.04	–0.01231	**0.01925**
15 min	47,696	–0.00000865	0.001071	1.49	–0.01393	0.02008

Table 5. Rolling multiple variance-ratio results (10-min specification)

Stat.	VR(2)	VR(5)	VR(10)	BCDEV	Joint test
Mean	0.9579	0.8670	0.7725	0.1692	
Median	0.9579	0.8354	0.6977	0.1448	
Std. dev.	0.0922	0.2041	0.3198	0.1191	
Rejections					373
Rejection rate (%)					13.26

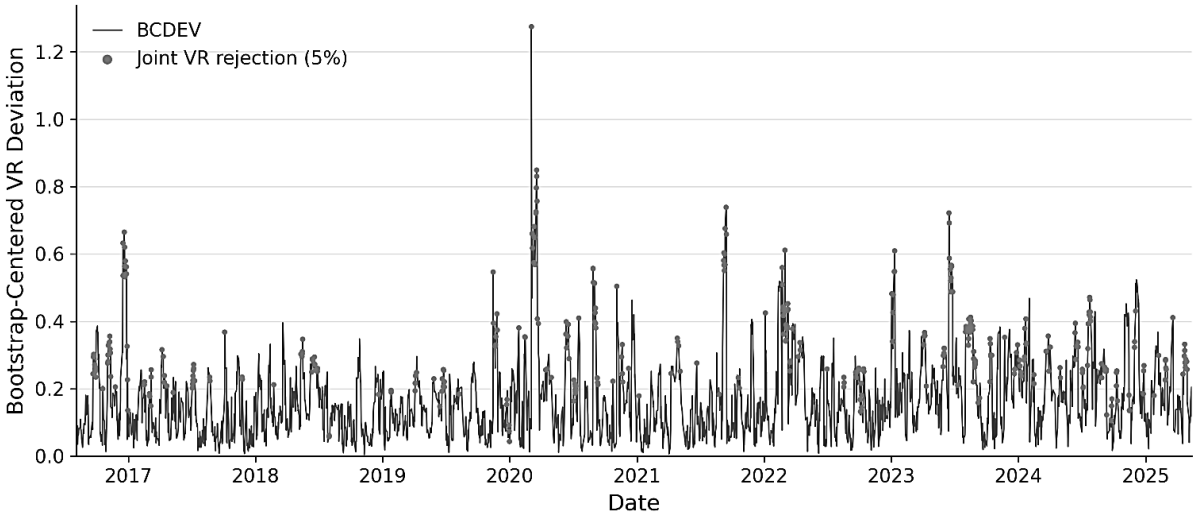


Figure 1. Evolution of the bootstrap-centered variance-ratio deviation

some of the larger deviations, while BCDEV also varies when the random-walk null is not rejected.

Robustness checks show that the temporal pattern persists across sampling frequencies. Joint rejection rates are 9.43%, 13.26%, and 18.86% at the 5-, 10-, and 15-minute frequencies, while null-size rejection rates remain close to 5% (4.4%, 5.5%, and 4.6%). These null-size rates are based on simulations of individual windows rather than the full overlapping rolling sequence. Yearly BCDEV correlations are strong between the 10-minute series and the 5-minute (Pearson 0.840; Spearman 0.818) and 15-minute series (Pearson 0.918; Spearman 0.939). Thus, rejection frequency varies with temporal resolution, but the broader time variation remains similar across specifications.

The second hypothesis concerns whether the magnitude of these departures also varies within the trading day. To examine this dimension without mixing different trading mechanisms, each continuous trading day is divided chronologically into early, middle, and late thirds. The resulting multiple variance ratio measures are reported in Table 6.

The intraday results indicate modest differences across the three periods. The early segment records the highest joint rejection rate at 10.68%, followed by the late segment at 9.05% and the middle segment at 7.84%. The ordering is different for the continuous BCDEV measure. The late segment records the highest average deviation at 0.1410, compared with 0.1310 in the early period and 0.1299 in the middle period. The mean variance ratios also remain below unity across all three segments, with the late period showing the lowest mean VR(2).

The difference between the rejection-rate ranking and the BCDEV ranking is important. It shows that the frequency with which the joint null is rejected and the continuous magnitude of departure do not provide identical information. The intraday results therefore point to some variation across the trading day, although the differences are modest compared with the temporal variation observed across the full sample.

Calendar-day block-bootstrap tests reveal no statistically significant differences in mean BCDEV ($p = 0.4765$) or joint-rejection rates ($p = 0.6067$) across the three intraday periods. Thus, while the observed temporal variation in random-walk departures is consistent with H1, the intraday results do not support *H2*. BCDEV is therefore retained as the continuous dependent variable in the fol-

Table 6. Intraday variation in multiple variance-ratio results

Intraday period	Windows	Joint rejection rate (%)	Mean VR(2)	Mean VR(5)	Mean BCDEV
Early	983	10.68	0.9738	0.8964	0.1310
Middle	906	7.84	0.9763	0.9154	0.1299
Late	906	9.05	0.9593	0.9002	0.1410

http://dx.doi.org/10.21511/imfi.23(4).2026.14

lowing regression analysis of realized volatility, calendar conditions, and major market events.

3.3. Regression analysis of random-walk departures

The regression analysis addresses *H3* and *H4* by examining whether log BCDEV is associated with realized volatility, calendar conditions, and selected market events, while controlling for year-specific effects.

Table 7. Regression results for bootstrap-centered variance-ratio deviations

Variable	(1) Baseline	(2) Extended
Log realized volatility	0.2151***	0.2175***
	(0.0653)	(0.0668)
Ramadan share	0.3567***	0.3057**
	(0.1201)	(0.1255)
COVID short hours share	−0.1528	−0.1553
	(0.2667)	(0.2673)
Boycott share	–	0.2893**
		(0.1473)
Earthquake share	–	−1.0040**
		(0.4287)
Rate-cut share	–	−0.3187
		(0.2524)
Tariff share	–	−0.7486***
		(0.2400)
Constant	−0.0290	−0.0077
	(0.6154)	(0.6285)
Year fixed effects	Yes	Yes
Observations	2,814	2,814
R^2	0.1216	0.1350
Adjusted R^2	0.1178	0.1300
Newey-West HAC lag	10	10

Note: Newey-West HAC standard errors are reported in parentheses. Both models include year fixed effects (2017–2025; 2016 omitted). ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

The results show a stable positive association between realized volatility and the magnitude of random-walk departure. The coefficient on log realized volatility is 0.215 in the baseline model and 0.217 in the extended specification, with both estimates statistically significant at the 1% level. Since both realized volatility and BCDEV enter the model in logarithmic form, the coefficient can be interpreted as an elasticity. In the extended model, a 1% increase in realized volatility is associated with a 0.217% increase in BCDEV, holding the remaining variables constant. The similarity of the coefficient across the two specifications indicates that this relationship is not materially affected by the inclusion of the event variables. This finding provides support for *H3*. Since BCDEV and realized volatility are calculated from the same return windows, this relationship reflects contemporaneous association rather than predictive evidence.

Calendar conditions produce fewer uniform results. The Ramadan share is positively associated with BCDEV in both models and remains statistically significant after the event variables are introduced. In the extended specification, a ten-percentage-point increase in the Ramadan share of a rolling window is associated with a 3.1% increase in BCDEV. By contrast, the COVID short-hours share is negative but statistically insignificant in both specifications. The results therefore indicate an association between Ramadan-period exposure and the magnitude of random-walk departure, while no comparable relationship is detected for the shortened COVID trading schedule.

The event variables also show heterogeneous associations. The boycott share is positively related to BCDEV, although its statistical significance is marginal, with a p-value of 0.0496 in the main specification. The earthquake and tariff shares are both negatively and statistically significantly associated with BCDEV. A ten-percentage-point increase in exposure to the earthquake period corresponds to a 9.6% lower BCDEV, while the same increase in tariff-period exposure corresponds to a 7.2% lower BCDEV. The rate-cut coefficient is negative but statistically insignificant. These estimates should be interpreted as conditional associations rather than causal effects of the respective events.

The robustness checks help distinguish the more stable findings from those that depend on the specification. Increasing the Newey–West lag length from 10 to 20 leaves the realized-volatility, Ramadan, earthquake, and tariff results statistically significant. The boycott coefficient becomes slightly weaker, with its p-value increasing from 0.0496 to 0.0537, while the COVID and rate-cut variables remain insignificant.

The positive association with realized volatility is also observed across all ten non-overlapping-window subsamples, although statistical significance is not uniform. By contrast, the coefficient on previous-window realized volatility is not statistically significant (β = 0.0370, p = 0.560), indicating that the positive contemporaneous association does not extend to a statistically significant relationship with previous-window volatility.

The results are also examined using BCDEV in levels rather than logarithms. Under this specification, realized volatility remains positively and significantly associated with the magnitude of departure, while the earthquake and tariff coefficients remain negative and statistically significant. The Ramadan and boycott coefficients, however, are no longer significant, and the rate-cut estimate becomes borderline. The positive volatility relationship remains significant across these alternative specifications, whereas the Ramadan and boycott results are more sensitive to the functional form of the dependent variable.

Sensitivity checks using alternative 5-, 10-, and 20-day event windows reveal some variation in statistical significance. The earthquake association remains significant under the 10- and 20-day specifications, while the tariff association is significant under the 5- and 10-day specifications. The rate-cut association is significant only under the 5-day specification. These findings suggest that the estimated relationships depend partly on the definition of event exposure.

The extended model raises R^2 from 0.122 to 0.135, indicating that the additional event variables account for some of the variation in BCDEV, although a substantial proportion of its temporal variation remains unexplained. Taken together, the regression results support *H3* through a positive contemporaneous association between realized volatility and BCDEV, while the evidence for *H4* remains mixed and sensitive to the definition of selected event periods. This heterogeneity is consistent with the idea that market behavior changes with prevailing conditions rather than following a single stable pattern.

The regression analysis identifies conditions associated with changes in the magnitude of random-walk departure, but it does not establish whether the underlying autoregressive structure of returns itself remains stable over time. The following section therefore complements these results with the time-varying autoregressive framework, which directly examines the evolution of short-horizon return dynamics.

3.4. Time-varying return dynamics (TV-AR)

The time-varying autoregressive analysis complements the rolling tests by examining whether the underlying short-run return structure changes over time and whether such changes generate statistically detectable linear predictability.

Starting with the lag selection, AR orders from one to twelve were compared using a common sample of 44,331 observations. The AIC reached its minimum at (p = 12), while the BIC favored (p = 1). Following the selection rule previously described, AR(8) and AR(12) satisfied the required AIC and Ljung-Box conditions. Since AR(8) was the smallest specification meeting these conditions, it was retained for the final estimation. The final AR(8) model was then estimated using 53,083 same-day observations.

Table 8. Time-varying autoregressive and parameter-constancy results

Statistics	Result
Selected AR order	8
Observations	53,083
Mean ζ_t	0.009422
Median ζ_t	0.008011
Maximum ζ_t	0.020022
Observations above pointwise 95% band	0
Hansen L_C statistic	6.7452
1% critical value	3.05
Parameter constancy	Rejected

Note: (ζt) measures time-varying linear predictability. The bootstrap bound is a one-sided pointwise 95% upper critical band. Hansen's LC is the joint parameter-constancy test statistic.

As shown in Table 8, the joint (L_C) statistic is 6.7452, exceeding the 1% critical value of 3.05. The null of parameter constancy is therefore rejected, indicating instability in the estimated model. A separate bootstrap test also rejects the stability of the autoregressive coefficients (p = 0.0091), pro-

http://dx.doi.org/10.21511/imfi.23(4).2026.14

viding more specific evidence of changing return dynamics. This finding is consistent with the temporal variation reported in Table 5. Similar results are obtained under trading-day-clustered resampling (p = 0.0155 and 0.0182).

The table also reports the magnitude of time-varying linear predictability. The mean value of (ζ_t) is 0.00942, with a median of 0.00801 and a maximum of 0.02002. Although the autoregressive parameters vary over time, none of the 53,083 estimated (ζ_t) values exceed the pointwise 95% bootstrap upper bound. The pointwise 95% bootstrap upper bound varies only slightly over time, with a mean of 0.02436 and a range of 0.02402-0.02452. Therefore, parameter variation does not necessarily imply a large or statistically detectable level of linear predictability at each point in time.

Figure 2 provides a clearer view of this evolution throughout the sample. In the AR(8) specification,

linear predictability gradually decreases during the first years, reaches its lowest level around 2020, and then increases progressively until the end of the sample. Despite these changes, the estimated (ζ_t) series remains below the pointwise bootstrap bound over the entire period.

To assess lag-order sensitivity, AR(12) was compared with AR(8) using a common sample of 44,331 observations. The resulting predictability series are strongly correlated (Pearson r = 0.886), although their magnitudes differ substantially (median ζ_t = 0.02295 for the re-estimated AR(8) and 0.00268 for AR(12)). Thus, the magnitude of estimated predictability is sensitive to lag order. The pointwise bootstrap bound in Figure 2 applies only to the original AR(8) specification.

The results in Figure 2 should be considered together with the parameter-constancy test presented in Table 8. The slope-stability tests provide formal evidence of variation in the autoregressive coefficients, while none of the AR(8) predictability estimates exceed the pointwise 95% bootstrap bound. Taken together, these findings support H1 in the full sample, indicating that changes in short-run return dynamics do not necessarily translate into statistically detectable linear predictability.

4. DISCUSSION

Using high-frequency MASI data, this study shows that short-horizon return dynamics in the Moroccan stock market are not constant over time. The rolling multiple variance-ratio analysis reveals clear changes in the magnitude of departures from the random-walk benchmark, with broadly similar temporal patterns across the 5-, 10-, and 15-minute specifications. As reported in Table 5 and Figure 1, these departures fluctuate throughout the sample rather than remaining at a stable level. This finding aligns with the Adaptive Market Hypothesis, which views market efficiency as a changing process rather than a fixed condition (Lo, 2004, 2012). Similar evidence of time-varying return predictability has been reported by Kim et al. (2011), Lim et al. (2013), Smith (2012), and

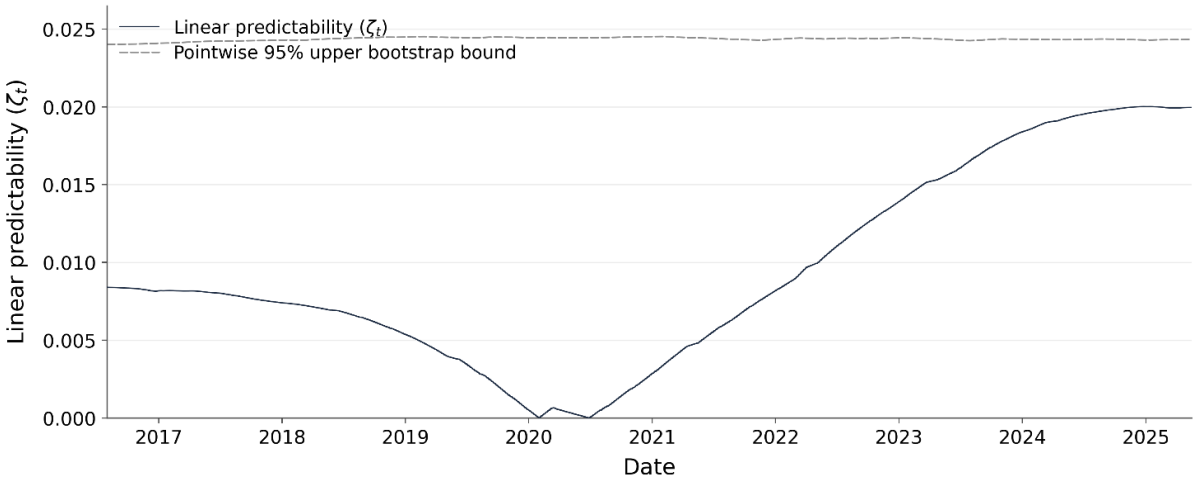


Figure 2. Time-varying linear predictability and pointwise 95% bootstrap upper bound

Urquhart and McGroarty (2016). The Moroccan evidence therefore extends this literature to an African emerging market where high-frequency evidence remains limited. Similar patterns have also been documented at much finer frequencies in cryptocurrency and electronic equity markets (Chu et al., 2019; Manahov et al., 2019; Meng & Li, 2021).

A second indication of this temporal variation comes from the evolution of the autoregressive structure itself. The joint parameter-constancy test reported in Table 8 indicates instability in the estimated model, while a separate bootstrap test provides evidence that the autoregressive coefficients also vary over time. A comparable pattern is documented by Ito et al. (2014) and Noda (2016), where market efficiency is allowed to evolve rather than being represented by a fixed parameter structure. However, an important distinction emerges from the Moroccan results. Although the slope-stability tests provide evidence of changing autoregressive coefficients, the estimated AR(8) linear predictability measure remains below the pointwise 95% bootstrap bound throughout the sample, as illustrated in Figure 2. Parameter instability therefore does not necessarily imply strong or persistent predictability. From an AMH perspective, adaptation may appear through changing return dynamics even when the magnitude of predictability remains limited.

This combination of results is particularly relevant in light of recent criticism of rolling-window tests. Hřebačka (2026) highlights the multiple-testing problem in rolling-window analyses, while Hřebačka (2025) argues that changes in return predictability may also be explained by business cycles and risk aversion without uniquely validating the AMH. Here, the broad temporal pattern remains visible across the 5-, 10-, and 15-minute return specifications, while the TV-AR model provides a complementary assessment of the same underlying question without relying on the same rolling classification. In this context, combining rolling evidence with formal slope-stability testing provides a broader basis for understanding changes in return dynamics, supporting *H1*. Nevertheless, the sensitivity of estimated predictability to model specification reinforces the need for caution when interpreting these changes as evidence of market adaptation. The findings are consistent with the Adaptive Market Hypothesis without establishing it as the only possible explanation.

Descriptive variation is also observed within the trading day, although it is more moderate than the temporal changes observed across the full sample. Table 6 shows that the Early period records the highest joint rejection rate, whereas the Late period has the highest average BCDEV, with the Middle period displaying smaller departures. The difference between these rankings is important because the frequency of rejecting the random-walk null and the magnitude of departure from that benchmark capture different aspects of return behavior. However, calendar-day block-bootstrap comparisons do not establish statistically significant differences in mean BCDEV or joint-rejection rates across the three periods. *H2* is therefore not supported by formal inference. This contrasts with Dong et al. (2017), who document systematic morning–afternoon differences in the autocorrelation of US stock returns. Chen et al. (2024) also examine intraday predictability and variance-ratio measures in Chinese stocks. These stock-level studies use different markets and empirical designs, so the comparison is informative rather than directly equivalent. In the Moroccan case, these descriptive differences appear even within continuous trading, but their lack of statistical significance does not establish identical behavior or a systematic ranking of trading-day periods.

Market conditions provide another important dimension of the results. Among the variables examined in the regression analysis, realized volatility presents the clearest and most stable relationship with the magnitude of random-walk departures. As shown in Table 7, higher realized volatility is associated with higher BCDEV across the main specifications, providing support for *H3* as a contemporaneous association. Kim et al. (2011) likewise find that stock return predictability varies with volatility and broader market conditions, although their long-horizon US setting differs from the present high-frequency index analysis. The nonsignificant previous-window volatility result and the shared returns used to construct BCDEV and realized volatility caution against a predictive interpretation. This contemporaneous rela-

http://dx.doi.org/10.21511/imfi.23(4).2026.14

tionship also fits the AMH view that return dynamics depend on prevailing market conditions rather than remaining unchanged through time. Ramadan is positively associated with BCDEV in the main logarithmic specification, but this relationship becomes weaker when the dependent variable is expressed in levels. Its effect should therefore be viewed as a possible calendar-related change in return dynamics rather than a stable structural relationship.

A less uniform pattern emerges when the analysis turns to major market events. The earthquake and U.S. tariff windows are associated with lower BCDEV in the main ten-day specification, although the results vary with event-window length. In contrast, the boycott effect is weaker and specification-dependent, while the rate-cut association is significant only under the five-day window and the COVID short-hours variable remains insignificant. H4 therefore receives only partial support. Rather than suggesting that major events necessarily increase departures from random-walk behavior, the results show that their relationship with the magnitude of random-walk departure may differ in both magnitude and direction. A similar non-uniform response to changing market conditions is reported by Kim et al. (2011), who find differing predictability patterns across crashes and crises. The Moroccan evidence therefore points to heterogeneous associations with selected calendar and event conditions rather than a common effect across all episodes.

Beyond the individual hypotheses, the Moroccan case adds evidence from a market that remains underrepresented in the high-frequency AMH literature. Most previous research has focused on developed and highly liquid electronic markets. Studies using transaction, quote, or order-book information Brogaard et al. (2014), Carrion (2013), Leone and Kwabi (2019), Manahov et al. (2019), and Meng and Li (2021) examine price discovery and trading mechanisms that cannot be observed directly using MASI index values. The present findings show that time-varying return dynamics are also observed in a North African emerging market, although the magnitude of predictability should not be overstated. By combining rolling variance-ratio tests, a continuous bootstrap-centered measure of random-walk departure, regression analysis, and TV-AR estimation, the study provides a more detailed picture of time-varying return dynamics than a simple efficient/inefficient classification.

Some limitations should nevertheless be considered when interpreting these findings. The analysis relies on the MASI index and therefore cannot determine whether the observed patterns differ across individual stocks, sectors, liquidity groups, or capitalization levels. The absence of order-book data also prevents a direct examination of trading intensity, bid-ask spreads, order flow, and high-frequency trader activity. Previous-tick sampling, infrequent trading, and nonsynchronous price updating may also contribute to measured short-horizon dependence, particularly given the high proportion of zero one-minute returns and the below-one mean variance ratios. These effects cannot be eliminated by using ten-minute returns. The tests examine departures from the random-walk benchmark and short-horizon linear predictability rather than overall informational efficiency or exploitable trading opportunities. Moreover, the event regressions identify conditional associations rather than causal effects of the events themselves. Future research could extend the analysis using firm-level intraday and order-book data and could compare the Moroccan market with other exchanges in the Middle East, North Africa, or Sub-Saharan Africa. Such extensions would help determine whether the time-varying return patterns observed at the CSE are specific to the Moroccan market or form part of a broader feature of emerging financial markets.

CONCLUSION

The purpose of this study was to examine whether the Moroccan stock market exhibits time-varying short-horizon return dynamics consistent with the Adaptive Market Hypothesis using high-frequency data from the Casablanca Stock Exchange. The analysis focused on whether short-horizon departures from the random-walk benchmark vary over time and across the trading day, and whether their magnitude is associated with realized volatility and selected calendar and event conditions.

The findings reveal that short-run return dynamics are not stable throughout the sample. Departures from the random-walk benchmark change over time, while formal slope-stability tests provide evidence that the autoregressive coefficients also vary over time. Descriptive intraday differences are modest, but formal tests do not establish statistically significant differences. Realized volatility shows the clearest contemporaneous association with larger departures. The associations with selected calendar and event conditions are mixed and specification-dependent. At the same time, the estimated AR(8) measure of linear predictability remains below the pointwise 95% bootstrap upper bound over the full sample. This shows that changing return dynamics do not necessarily imply strong or persistent predictability.

Rather than describing the market as simply efficient or inefficient, the results point to changing return dynamics that are consistent with the Adaptive Market Hypothesis, although alternative explanations remain possible. These findings concern specific patterns of short-horizon return behavior rather than overall informational efficiency.

This study adds to the existing literature by examining these dynamics at high frequency and by combining several complementary approaches within the same empirical framework. The use of rolling variance-ratio tests, a continuous measure of random-walk departure, regression analysis, and time-varying autoregressive modeling makes it possible to capture different aspects of short-horizon return behavior instead of relying on a single binary measure of efficiency. The findings also show the value of high-frequency data for identifying changes in price behavior across intraday intervals.

In this sense, the study contributes not only to the evidence on the Moroccan market, but also to the broader literature on time-varying return dynamics in emerging markets. The results provide useful insights into short-run return dependence and its association with market conditions, with potential relevance for researchers, market participants, regulators, and exchange operators.

AUTHOR CONTRIBUTIONS

Conceptualization: Dina Sittel, Achraf Louati.
Data curation: Dina Sittel, Achraf Louati.
Formal analysis: Dina Sittel, Achraf Louati.
Investigation: Dina Sittel, Achraf Louati.
Methodology: Dina Sittel, Achraf Louati.
Resources: Dina Sittel.
Software: Dina Sittel.
Supervision: Achraf Louati.
Validation: Dina Sittel, Achraf Louati.
Visualization: Dina Sittel, Achraf Louati.
Writing – original draft: Dina Sittel, Achraf Louati.
Writing – review & editing: Dina Sittel, Achraf Louati.

REFERENCES

1. Admati, A. R., & Pfleiderer, P. (1988). A theory of intraday patterns: Volume and price variability. *Review of Financial Studies, 1*(1), 3-40. https://doi.org/10.1093/rfs/1.1.3
2. Aït-Sahalia, Y., Mykland, P. A., & Zhang, L. (2005). How often to sample a continuous-time process in the presence of market microstructure noise. *Review of Financial Studies, 18*(2), 351-416. https://doi.org/10.1093/rfs/hhi016
3. Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In B. N. Petrov & F. Csáki (Eds.), *Proceedings of the Second International Symposium on Information Theory* (pp. 267-281). Akadémiai Kiadó. Retrieved from https://hero.epa.gov/reference/591/
4. Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of the American Statistical Associa-*

http://dx.doi.org/10.21511/imfi.23(4).2026.14

tion, 96(453), 42-55. https://doi.org/10.1198/016214501750332965

5. Belsley, D. A., Kuh, E., & Welsch, R. E. (1980). *Regression diagnostics: Identifying influential data and sources of collinearity*. Wiley. https://doi.org/10.1002/0471725153
6. Biais, B., Foucault, T., & Moinas, S. (2015). Equilibrium fast trading. *Journal of Financial Economics, 116*(2), 292-313. https://doi.org/10.1016/j.jfineco.2015.03.004
7. Brogaard, J., Hendershott, T., & Riordan, R. (2014). High-frequency trading and price discovery. *Review of Financial Studies, 27*(8), 2267-2306. https://doi.org/10.1093/rfs/hhu032
8. Carrion, A. (2013). Very fast money: High-frequency trading on the NASDAQ. *Journal of Financial Markets, 16*(4), 680-711. https://doi.org/10.1016/j.finmar.2013.06.005
9. Chen, J., Haboub, A., & Khan, A. (2024). Limits of arbitrage and their impact on market efficiency: Evidence from China. *Global Finance Journal, 59,* 100916. https://doi.org/10.1016/j.gfj.2023.100916
10. Chow, K. V., & Denning, K. C. (1993). A simple multiple variance ratio test. *Journal of Econometrics, 58*(3), 385-401. https://doi.org/10.1016/0304-4076(93)90051-6
11. Chu, J., Zhang, Y., & Chan, S. (2019). The adaptive market hypothesis in the high frequency cryptocurrency market. *International Review of Financial Analysis, 64,* 221-231. https://doi.org/10.1016/j.irfa.2019.05.008
12. Dong, X., Feng, S., Ling, L., & Song, P. (2017). Dynamic autocorrelation of intraday stock returns. *Finance Research Letters, 20,* 274-280. https://doi.org/10.1016/j.frl.2016.10.008
13. Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance, 25*(2), 383-417. https://doi.org/10.2307/2325486
14. Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review, 70*(3), 393-408. Retrieved from https://www.jstor.org/stable/1805228
15. Hansen, B. E. (1992). Testing for parameter instability in linear models. *Journal of Policy Modeling, 14*(4), 517-533. https://doi.org/10.1016/0161-8938(92)90019-9
16. Hansen, P. R., & Lunde, A. (2006). Realized variance and market microstructure noise. *Journal of Business & Economic Statistics, 24*(2), 127-161. https://doi.org/10.1198/073500106000000071
17. Hřebačka, V. (2025). Analyzing market efficiency: The role of business cycles, risk aversion, and Occam's razor in the Adaptive Market Hypothesis. *Finance Research Letters, 75,* 106840. https://doi.org/10.1016/j.frl.2025.106840
18. Hřebačka, V. (2026). On the testing of Adaptive Markets Hypothesis using rolling windows. *Applied Economics Letters, 33*(5), 742-746. https://doi.org/10.1080/13504851.2024.2389341
19. Ito, M., Noda, A., & Wada, T. (2014). International stock market efficiency: A non-Bayesian time-varying model approach. *Applied Economics, 46*(23), 2744-2754. https://doi.org/10.1080/00036846.2014.909579
20. Ito, M., Noda, A., & Wada, T. (2016). The evolution of stock market efficiency in the US: A non-Bayesian time-varying model approach. *Applied Economics, 48*(7), 621-635. https://doi.org/10.1080/00036846.2015.1083532
21. Kim, J. H. (2006). Wild bootstrapping variance ratio tests. *Economics Letters, 92*(1), 38-43. https://doi.org/10.1016/j.econlet.2006.01.007
22. Kim, J. H., Shamsuddin, A., & Lim, K.-P. (2011). Stock return predictability and the adaptive markets hypothesis: Evidence from century-long U.S. data. *Journal of Empirical Finance, 18*(5), 868-879. https://doi.org/10.1016/j.jempfin.2011.08.002
23. Kutner, M. H., Nachtsheim, C. J., Neter, J., & Li, W. (2005). *Applied linear statistical models* (5th ed.). McGraw-Hill Irwin.
24. Leone, V., & Kwabi, F. (2019). High frequency trading, price discovery and market efficiency in the FTSE100. *Economics Letters, 181,* 174-177. https://doi.org/10.1016/j.econlet.2019.05.022
25. Lim, K.-P., & Brooks, R. (2011). The evolution of stock market efficiency over time: A survey of the empirical literature. *Journal of Economic Surveys, 25*(1), 69-108. https://doi.org/10.1111/j.1467-6419.2009.00611.x
26. Lim, K.-P., Luo, W., & Kim, J. H. (2013). Are US stock index returns predictable? Evidence from automatic autocorrelation-based tests. *Applied Economics, 45*(8), 953-962. https://doi.org/10.1080/00036846.2011.613782
27. Lo, A. W. (2004). The adaptive markets hypothesis. *The Journal of Portfolio Management, 30*(5), 15-29. https://doi.org/10.3905/jpm.2004.442611
28. Lo, A. W. (2012). Adaptive markets and the new world order. *Financial Analysts Journal, 68*(2), 18-29. https://doi.org/10.2469/faj.v68.n2.6
29. Lo, A. W., & MacKinlay, A. C. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies, 1*(1), 41-66. https://doi.org/10.1093/rfs/1.1.41
30. Lo, A. W., & MacKinlay, A. C. (1990). An econometric analysis of nonsynchronous trading. *Journal of Econometrics, 45*(1-2), 181-211. https://doi.org/10.1016/0304-4076(90)90098-E
31. Manahov, V., Hudson, R., & Urquhart, A. (2019). High-frequency trading from an evolutionary perspective: Financial markets as adaptive systems. *International Journal of Finance & Economics, 24*(2), 943-962. https://doi.org/10.1002/ijfe.1700
32. Meng, K., & Li, S. (2021). The adaptive market hypothesis and high frequency trading. *PLOS ONE, 16*(12), e0260724. https://doi.org/10.1371/journal.pone.0260724

33. Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica, 55*(3), 703-708. https://doi.org/10.2307/1913610

34. Noda, A. (2016). A test of the adaptive market hypothesis using a time-varying AR model in Japan. *Finance Research Letters, 17,* 66-71. https://doi.org/10.1016/j.frl.2016.01.004

35. Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics, 6*(2), 461-464. https://doi.org/10.1214/aos/1176344136

36. Smith, G. (2012). The changing and relative efficiency of European emerging stock markets. *The European Journal of Finance, 18*(8), 689-708. https://doi.org/10.1080/1351847X.2011.628682

37. Urquhart, A., & Hudson, R. (2013). Efficient or adaptive markets? Evidence from major stock markets using very long run historic data. *International Review of Financial Analysis, 28,* 130-142. https://doi.org/10.1016/j.irfa.2013.03.005

38. Urquhart, A., & McGroarty, F. (2016). Are stock markets really efficient? Evidence of the adaptive market hypothesis. *International Review of Financial Analysis, 47,* 39-49. https://doi.org/10.1016/j.irfa.2016.06.011

http://dx.doi.org/10.21511/imfi.23(4).2026.14