Int. J. Anal. Appl. (2025), 23:298

International Journal of Analysis and Applications

A Composite Efficiency Index for ASEAN Foreign Exchange Markets

Tran Trong Huynh[1,2,*], Thi Thu Hong Dinh[2]

[1]*Department of Mathematics, FPT University, Ha Noi, Vietnam*
[2]*School of Finance, University of Economics Ho Chi Minh City (UEH), Ho Chi Minh City, Vietnam*
**Corresponding author:* huynhtt4@fe.edu.vn

ABSTRACT. The Efficient Market Hypothesis (EMH) has long been a central paradigm in finance, yet mounting evidence suggests that market efficiency is neither uniform across assets nor constant over time. This study examines the dynamics of foreign exchange (FX) market efficiency in six ASEAN economies (Vietnam, Thailand, Indonesia, Malaysia, the Philippines, and Singapore) over the period January 2000 to August 2025. Using daily bilateral exchange rates against the U.S. dollar, we construct twelve sub-indices that capture serial dependence, volatility clustering, distributional anomalies, and microstructure frictions. These standardized measures are then aggregated through principal component analysis (PCA) into a Composite Efficiency Index (CEI), complemented by an equal-weighted average as a robustness check. The empirical results reveal three key findings. First, inefficiency has declined significantly over time, consistent with the Adaptive Market Hypothesis (AMH), but with pronounced spikes during global and local crises such as the Global Financial Crisis and the COVID-19 pandemic. Second, substantial heterogeneity is observed across markets: Singapore emerges as the most efficient, while Vietnam is persistently the least efficient. Third, changes in CEI predict higher-order return dependencies, though not mean returns themselves, underscoring its validity as a forward-looking measure. These results provide new insights into the evolving nature of FX efficiency, offering both academic contributions and policy relevance.

1. Introduction

The Efficient Market Hypothesis, originally formalized by Fama [1], has long served as a cornerstone of financial economics. In its strictest form, EMH posits that asset prices fully and instantaneously reflect all available information, implying that returns follow a martingale difference sequence with no predictable component. A direct corollary is that price changes should resemble a random walk (RW), with no systematic autocorrelation or profitable forecasting opportunities. The foreign exchange market has been a central testing ground for this

Received Nov. 7, 2025

2020 Mathematics Subject Classification. 91G15, 62M10, 62P05.

Key words and phrases. composite efficiency index; foreign exchange markets; AMH; EMH; random walk.

https://doi.org/10.28924/2291-8639-23-2025-298

ISSN: 2291-8639

hypothesis, particularly since the seminal work of Meese and Rogoff [2], who demonstrated that most structural and macroeconomic models fail to outperform a simple random walk in forecasting short-horizon exchange rate movements. Their conclusion that "the random walk is unbeatable" has profoundly shaped subsequent debates, establishing RW as the de facto benchmark of efficiency in currency markets.

Yet, the assumption of constant and universal efficiency has been increasingly challenged. Lo [3] argues that efficiency is not static but evolves with changing market conditions, competition, and investor behavior. In this framework, markets may display efficiency during tranquil periods but exhibit marked inefficiencies during crises or structural shifts. The AMH thus shifts the focus from binary questions of whether markets are efficient toward continuous assessments of how efficiency fluctuates across time and environments. For foreign exchange markets in particular—characterized by deep liquidity, 24-hour trading, and strong exposure to global shocks—efficiency may be highly dynamic, shaped by both domestic fundamentals and international contagion effects.

Testing the AMH requires measurement tools capable of capturing time-varying inefficiency. Early studies often focused on isolated diagnostics such as short-run autocorrelation, variance ratio tests, or long-memory estimates [4-6]. While informative, such measures capture only one dimension of efficiency. More recent contributions have advanced composite indices of market inefficiency. Le Tran and Leirvik [7], for example, introduced the Automatic Portmanteau-based Market Inefficiency Measure (AMIM), a robust single-metric approach based on autocorrelation. Mattera, Di Sciorio and Trinidad-Segovia [8] expanded this line by proposing a Composite Efficiency Index, integrating multiple inefficiency dimensions using principal component analysis. Similarly, Bock and Geissel [9] proposed the Average Area of Inefficiency (AAI), which summarizes both the depth and duration of inefficient episodes. These innovations underscore the multidimensional nature of inefficiency, encompassing serial dependence, volatility clustering, distributional deviations, and microstructure frictions.

Despite these advances, two critical gaps remain. First, many empirical studies apply composite indices to equity markets in developed economies, with far less attention to foreign exchange markets. Yet FX markets differ structurally from equities: they are globally integrated, continuously traded, and notoriously resistant to forecasting beyond the random walk benchmark. If any asset class exemplifies the AMH's claim of adaptive and episodic efficiency, it is foreign exchange. Second, existing composite measures often emphasize methodological elegance but have not systematically tested predictive validity—whether fluctuations in inefficiency indices can forecast short-horizon predictability metrics such as autocorrelation and variance-ratio deviations. Addressing these gaps is crucial both to advance empirical tests of AMH and to inform policy debates on currency market stability.

This paper focuses on ASEAN foreign exchange markets because they exhibit unique structural, institutional, and behavioral features that make them ideal for testing the Adaptive Market Hypothesis. First, the degree of capital account openness differs markedly across ASEAN economies, ranging from highly liberalized markets such as Singapore to more managed regimes such as Vietnam and Malaysia. These differences, combined with periodic capital flow management measures, have been shown to influence liquidity, transaction costs, and arbitrage opportunities, thereby affecting overall market efficiency [10]. Second, although the Efficient Market Hypothesis has been extensively examined in developed equity and currency markets, empirical evidence for emerging Asian currencies remains more limited. Studies on Asia–Pacific FX markets indicate that deviations from the random walk hypothesis are generally more pronounced in less developed and crisis-prone economies, suggesting that this region provides a natural setting for exploring time-varying efficiency and adaptive dynamics [11]. Finally, behavioral aspects further differentiate ASEAN FX markets: recent evidence points to episodic rather than persistent herd behavior, with heterogeneity across countries reflecting differences in macroeconomic fundamentals and central bank intervention regimes [12]. Collectively, these structural, institutional, and behavioral heterogeneities justify the construction of a multidimensional Composite Efficiency Index (CEI) to capture and compare efficiency patterns across ASEAN currencies.

This study contributes by constructing a novel Composite Efficiency Index for ASEAN foreign exchange markets over the period January 2000 to August 2025. Building on 12 standardized sub-indices capturing distinct aspects of inefficiency—including serial dependence, volatility persistence, distributional anomalies, and market microstructure frictions—we apply principal component analysis to extract a single composite indicator. For robustness, we also compute a simple equal-weighted average (AVE) of the twelve measures. The dataset covers six ASEAN currencies—Vietnamese dong (VND), Thai baht (THB), Indonesian rupiah (IDR), Malaysian ringgit (MYR), Philippine peso (PHP), and Singapore dollar (SGD)—each expressed against the U.S. dollar, reflecting their role as regionally significant and globally traded exchange rates. Daily OHLC (open, high, low, close) data from Investing.com are aggregated to a monthly frequency to compute all sub-indices.

Our analysis yields several key results. First, CEI values trend downward over time, consistent with markets becoming more efficient as financial liberalization, technological advances, and integration progressed—direct evidence in favor of the AMH. Second, cross-country comparisons reveal systematic differences: Singapore emerges as the most efficient market, while Vietnam persistently ranks as the least efficient. Third, CEI spikes coincide with major global and regional crises, including the Global Financial Crisis of 2008–2009 and the COVID-19 pandemic in 2020, providing strong face validity. Fourth, regressions using two-way fixed effects show that changes in CEI (ΔCEI) significantly predict next-month inefficiency metrics such as $|AC_1|$ and $|VR(2)-1|$, though not mean returns themselves. This distinction reinforces the martingale-difference property of returns while highlighting that inefficiency manifests in higher-order serial dependence rather than drift. Collectively, these results confirm that ASEAN FX markets exhibit adaptive, multidimensional, and event-driven efficiency patterns, in line with AMH predictions.

The contributions of this paper are threefold. First, it provides the first comprehensive composite measure of efficiency for ASEAN currency markets, filling an important geographic and asset-class gap in the literature. Second, by integrating 12 diverse inefficiency diagnostics into a single CEI, the study overcomes the limitations of single-measure approaches and provides a richer depiction of market dynamics. Third, by linking ΔCEI to subsequent predictability metrics, the analysis offers a novel test of predictive validity, demonstrating that efficiency indices not only describe but also forecast aspects of market behavior. These findings hold theoretical significance for the adaptive markets paradigm and practical relevance for policymakers and investors seeking to understand efficiency in globally interconnected but regionally heterogeneous FX markets.

The remainder of the paper is structured as follows. Section 2 details the construction of the twelve sub-indices and the composite efficiency index. Section 3 outlines the data and econometric methodology. Section 4 presents empirical results, including descriptive statistics, time-varying dynamics, cross-market comparisons, event-driven peaks, and predictive regressions. Section 5 discusses the implications concerning existing literature, and Section 6 concludes.

2. Measuring an Inefficient Market

We summarize foreign-exchange market inefficiency each month by a Composite Efficiency Index constructed from 12 sub-indices computed on daily OHLC data within a rolling window (12 months ending at month m). Each sub-index captures a distinct facet of deviation from the Efficient Market Hypothesis. After standardization, the sub-indices are aggregated using Principal Component Analysis to obtain a single composite measure. Lower CEI values indicate higher efficiency (closer to EMH), and higher values indicate stronger inefficiency.

Data windowing and basic notation

Let C_t, O_t, H_t, and Lt be the daily close, open, high, and low prices. Define daily log returns:

$$r_t = \ln C_t - \ln C_{t-1}, \qquad r_t^{\mathrm{ON}} = \ln O_t - \ln C_{t-1}, \qquad r_t^{\mathrm{ID}} = \ln C_t - \ln O_t$$

For each market i and month m, we collect all trading days whose calendar month belongs to the 12-month window (m−11, …, m) and compute the 12 sub-indices below. Notably, $ACF_\ell(x)$ is the sample autocorrelation at lagℓ; $1\{\cdot\}$ is the indicator function.

Twelve sub-indices

The Composite Efficiency Index integrates twelve standardized sub-indices, each designed to capture a particular dimension of market inefficiency. These measures span return dynamics, volatility persistence, distributional features, and market microstructure frictions. In line with the Efficient Market Hypothesis and the Adaptive Market Hypothesis, higher values of each sub-index indicate stronger deviations from efficiency.

Autocorrelation Index (I^{AC}). Under the Efficient Market Hypothesis, returns should exhibit no serial correlation; departures from a random walk can be detected via variance-ratio ideas Lo & MacKinlay (1988) [13] and, more practically, for emerging markets, autocorrelation-based diagnostics used in recent composite-efficiency work [7, 8]. The I^{AC} summarizes short-horizon serial dependence by pooling information from multiple lags and from the Ljung–Box omnibus test.

$$I^{AC} = \frac{1}{2}\left(\sum_{\ell=1}^{p} |ACF_\ell(r_t)| + (1 - p_{LB})\right) \tag{1}$$

with p=5 and p_{LB} denoting the Ljung–Box statistic.

Variance Ratio Index (I^{VR}): The variance ratio test evaluates whether k-period return variances equal k times the variance of one-period returns, as predicted by a random walk. I^{VR} captures absolute deviations from this benchmark. Recent applications confirm the sensitivity of variance ratios to efficiency shifts in FX and equity markets [5, 14]. In an efficient market, the variance of k-period cumulative returns should equal k times the variance of one-period returns. Deviations from unity, therefore, capture random-walk violations. We define the variance ratio as $VR(k) = \frac{Var\left(\sum_{j=0}^{k-1} r_{t-j}\right)}{k \cdot Var(r_t)}$ and construct the sub-index

$$I^{VR} = \sum_{k \in \{2,5,10\}} |VR(k) - 1| \tag{2}$$

Volatility Clustering Index (I^{VC}): Under weak-form EMH, second-moment dynamics should not be predictably persistent; in practice, financial returns display ARCH-type dependence—high-volatility days tend to cluster with high-volatility days, and low with low. I^{VC} quantifies this persistence and treats stronger clustering as greater inefficiency. Volatility persistence has recently been documented as a robust signal of lower informational efficiency in both developed and emerging markets [15]. To measure this, we compute autocorrelations of both absolute and squared returns and aggregate them into the index:

$$I^{VC} = \frac{1}{2}\left(\sum_{\ell=1}^{p} |ACF_\ell(r_t)| + \sum_{\ell=1}^{p} |ACF_\ell(r_t^2)|\right) \tag{3}$$

with p=5.

Range-Based Volatility Predictability (I^{RVP}): Intraday range estimators provide efficient volatility measures [16]. IRVP tests whether such measures exhibit serial predictability. Recent work demonstrates that range-based volatility dynamics signal inefficiencies during turbulent periods [17]. A related measure employs the Garman–Klass range-based volatility estimator, defined as: $\sigma_{GK,t}^2 = 0.5(ln(H_t/L_t))^2 - (2\,ln\,2 - 1)(ln(C_t/O_t))^2$. We compute its daily values, take square roots, and examine their serial dependence. The corresponding index is

$$I^{RVP} = \sum_{\ell=1}^{p} |\,ACF_\ell\,(\sqrt{\sigma_{GK,t}^2})| \tag{4}$$

Hurst Exponent Index (I^H): The Hurst exponent assesses long-memory behavior, with EMH implying H = 0.5. I^H measures the deviation from this benchmark. Recent applications in FX and cryptocurrencies highlight long-memory as a time-varying inefficiency [18]. Long memory in returns is captured by the Hurst exponent. A martingale difference sequence should yield H=0.5. Deviations from this value reflect persistent or anti-persistent dynamics inconsistent with efficiency. We therefore define:

$$I^H = |H - 0.5| \tag{5}$$

Skewness Index (I^{Skew}): Efficient markets should not display systematic asymmetry in return distributions. I^{Skew} captures the skewness of daily returns within each month. Persistent skewness is linked to information asymmetry and order-flow imbalances. Distributional asymmetries provide additional information on inefficiency [19]. The skewness of returns, if significantly different from zero, indicates one-sided order flow or information arrival. We therefore define:

$$I^{Skew} = |Skew(r_t)| \tag{6}$$

Kurtosis Index (I^{Kurt}): Fat-tailed distributions imply frequent extreme returns inconsistent with Gaussian benchmarks. I^{Kurt} measures excess kurtosis as an inefficiency indicator. Recent studies highlight that kurtosis spikes are linked to market stress and structural breaks [20]. We compute the excess component beyond the Gaussian benchmark,

$$I^{Kurt} = max\{\,Kurt(r_t) - 3,\ 0\} \tag{7}$$

Return Decomposition Inefficiency (I^{RDI}): I^{RDI} compares overnight returns (close-to-open) and intraday returns (open-to-close). Large and persistent differences indicate segmented price discovery. Recent literature identifies the "overnight return puzzle" as a form of inefficiency. We also account for the decomposition of returns between overnight and intraday trading [21]. Define overnight returns as $r_t^{ON} = ln\,O_t - ln\,C_{t-1}$, and intraday returns as $r_t^{ID} = ln\,C_t - ln\,O_t$. Persistent discrepancies in means or variances between the two indicate segmented price discovery. Our measure is

$$I^{RDI} = \left|\overline{r^{ON}} - \overline{r^{ID}}\right| + \left|ln\frac{Var(r^{ON})}{Var(r^{ID})}\right| \tag{8}$$

Tail Risk Index (I^{Tail}): Extreme return realizations undermine efficient pricing. I^{Tail} combines average daily ranges with the frequency of shocks exceeding two standard deviations.

Recent research links tail events to inefficiency episodes during crises [22]. To capture tail risk, we consider both the average relative intraday range $\rho_t = (H_t - L_t)/C_{t-1}$ and the frequency of extreme daily returns exceeding twice the unconditional standard deviation. The combined tail index is

$$I^{Tail} = \overline{\rho_t} + \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\{|r_t| > 2\sigma_r\} \quad (9)$$

Closing Location Value Index (I^{CLV}): I^{CLV} reflects the closing price's position within the daily high–low range. Persistent biases indicate order-flow imbalances not arbitraged away. While CLV is an established microstructure measure, recent evidence shows its utility for efficiency diagnostics in equity markets [23]. Market microstructure frictions are reflected in the location of the closing price within the daily range. We calculate the Close-Location Value (CLV) as

$$CLV_t = \frac{C_t - O_t}{max(H_t - L_t, \varepsilon)}$$

with a small ε>0 to avoid division by zero. An imbalance in its distribution, or persistence across time, signals inefficiency. Thus,

$$I^{CLV} = \frac{1}{2}\left(|\overline{CLV_t}| + \sum_{\ell=1}^{p} |ACF_\ell(CLV_t)|\right) \quad (10)$$

Gap Index (IGap): I^{Gap} measures systematic differences between prior-day closing and next-day opening prices. Predictable gaps imply delayed overnight information incorporation. Recent work shows that gap dynamics remain a significant inefficiency source in FX and equity. In a similar vein, opening gaps between the close of the previous day and the following open provides further evidence [24]. Define

$$Gap_t = \frac{|O_t - C_{t-1}|}{max(H_t - L_t, \varepsilon)}$$

and construct the index

$$I^{Gap} = \frac{1}{2}\left(\overline{Gap_t} + \sum_{\ell=1}^{p} |ACF_\ell(Gap_t)|\right) \quad (11)$$

Relative Close-to-Close vs Range Volatility Mismatch (I^{RCVM}): I^{RCVM} compares realized volatility from intraday ranges with volatility estimated from close-to-close returns. Persistent discrepancies signal incomplete incorporation of intraday information into closing prices. Andersen & Bollerslev (1998) [25] provided the initial framework, with recent applications confirming its relevance for efficiency testing. We compare volatility estimates from intraday ranges against close-to-close returns [26]. The monthly range-based volatility is $\sigma_{range} = \sqrt{\overline{\sigma^2_{GK,t}}}$, while the conventional measure is $\sigma_{cc} = sd(r_t)$. Persistent differences indicate incomplete incorporation of information. The index is

$$I^{RCVM} = \left|ln\, \sigma_{range} - ln\, \sigma_{cc}\right| \tag{12}$$

These twelve sub-indices span the domains of serial dependence, volatility dynamics, distributional properties, intraday–overnight segmentation, tail risk, and microstructure imbalance. They provide a rich and complementary set of signals from which a composite measure of inefficiency can be derived.

Standardization and composite construction

Once the twelve sub-indices have been computed for each market and month, the next step is to bring them to a common scale and combine them into a single composite measure. Because the raw sub-indices differ in units, magnitudes, and statistical distributions, direct aggregation would be misleading. We therefore apply a two–stage transformation to ensure comparability across indices and robustness to outliers.

First, each sub-index is standardized into a z-score relative to its own time-series distribution within a given market. Specifically, for sub-index j in market ii and month mm, we compute

$$z_{i,m}^{(j)} = \frac{I_{i,m}^{(j)} - \mu_i^{(j)}}{\sigma_i^{(j)}}$$

where $\mu_i^{(j)}$ and $\sigma_i^{(j)}$ denote the historical mean and standard deviation of that sub-index for market i. This transformation ensures that all components are expressed in standard deviation units, thus rendering them comparable across measures and across markets. To mitigate the influence of extreme values, we additionally truncate each standardized score to the interval [−3,3]. This clipping procedure preserves the relative ranking of observations while preventing a small number of outliers from dominating the aggregate measure.

Second, we combine the standardized scores into a single composite index. The principal method employs Principal Component Analysis (PCA). Let $\tilde{\boldsymbol{z}}_{i,m} = (z_{i,m}^{(1)}, \dots, z_{i,m}^{(12)})^\top$ denote the vector of clipped standardized scores. We estimate PCA on the pooled panel of markets and periods, and extract the loadings of the first principal component $\boldsymbol{w}$. The Composite Efficiency Index for market ii at month mm is then given by

$$CEI_{i,m} = \boldsymbol{w}^\top \tilde{\boldsymbol{z}}_{i,m} \tag{13}$$

The sign of $\boldsymbol{w}$ is normalized such that higher CEI values correspond consistently to greater market inefficiency.

As a robustness check, we also construct a simple equal–weighted benchmark, defined as the arithmetic mean of the twelve standardized sub-indices:

$$AVE_{i,m} = \frac{1}{12}\sum_{j=1}^{12} z_{i,m}^{(j)} \tag{14}$$

Consistent with recent composite approaches, higher CEI and AVE values denote stronger deviations from EMH. The correlation between CEI and AVE is found to be consistently high across markets, confirming that the composite measure does not depend critically on the specific weighting scheme provided by PCA.

In practical implementation, PCA is applied to the set of observations with complete data across the twelve sub-indices, while variables with near-zero variance are excluded to prevent numerical instability. Missing values at the scoring stage are imputed with training-set means, ensuring that no forward-looking information is introduced. To maintain stability of the composite index, we also enforce a consistent orientation of the first principal component across rolling estimation windows by aligning its sign with the cross-sectional mean of standardized scores.

Through this standardization and aggregation process, the CEI transforms a wide range of inefficiency indicators into a parsimonious monthly time series. This single index serves as the basis for subsequent empirical analyses of time-varying efficiency, cross-market comparisons, and predictive validity.

3. Method

Data and Variables

The empirical analysis is conducted using daily foreign exchange rates obtained from Investing.com for the period January 2000 to August 2025. The dataset covers six ASEAN currencies, each expressed in terms of their bilateral exchange rate against the United States dollar (USD). These six series represent the most liquid and widely monitored currency markets in the region, providing a consistent benchmark for evaluating efficiency across countries.

For each market, we collect the standard OHLC quotes: the daily closing rate (C_t), opening rate (O_t), highest intraday rate (H_t), and lowest intraday rate (L_t). Based on these four series, we compute logarithmic returns and derive a wide range of measures designed to capture distinct dimensions of potential market inefficiency. The daily indicators are then aggregated to the monthly frequency in order to form the twelve sub-indices described in Section 2. Examples include autocorrelation of returns to measure short-term predictability, volatility clustering derived from absolute and squared returns, and tail risk proxied by the frequency of extreme return realizations.

To ensure comparability, each sub-index is standardized and subsequently combined using Principal Component Analysis to produce the Composite Efficiency Index. As a robustness check, we also compute a simple equal-weighted average (AVE) of the standardized sub-indices.

Table 1 provides detailed definitions of all variables used in the analysis. For each sub-index, we report its formula and the theoretical rationale linking the measure to market efficiency. The table also introduces the CEI and AVE as composite indicators.

Table 1. Variable description and definitions

Variable	Formula	Description
IAC	Eq. (1)	If markets are efficient, returns should be serially uncorrelated. This index aggregates short-horizon autocorrelations of daily returns and the Ljung–Box test statistic. Higher values imply short-run predictability and thus stronger deviations from EMH.
IVR	Eq. (2)	Under a random walk, the variance of kk-period returns should equal kk times the variance of one-period returns. This index sums absolute deviations of variance ratios from unity, indicating violations of the random walk property.
IVC	Eq. (3)	EMH implies that volatility innovations are not predictable. This measure captures persistence in absolute and squared returns; significant clustering of volatility indicates inefficient information absorption.
IRVP	Eq. (4)	Based on the Garman–Klass estimator from intraday ranges, this measure tests whether range-based volatility follows a predictable pattern. Predictability of volatility contradicts EMH.
IH	Eq. (5)	For efficient markets, the Hurst exponent equals 0.5. Values significantly above or below 0.5 imply long memory or anti-persistence in returns, indicating deviation from martingale behavior.
ISkew	Eq. (6)	In efficient markets, return distributions should not exhibit systematic asymmetry. Large skewness reflects one-sided order flow or asymmetric information, consistent with temporary inefficiency.
IKurt	Eq. (7)	A normal distribution has a kurtosis of 3. Excess kurtosis (fat tails) reflects frequent extreme returns and suggests that prices are not fully incorporating information smoothly.
IRDI	Eq. (8)	Compares overnight returns (rON) with intraday returns (rID). Large differences in means or variances imply segmented price discovery and delayed information processing, contrary to EMH.
ITail	Eq. (9)	Combines the average relative daily range with the frequency of extreme returns exceeding two standard deviations. Frequent extreme movements are inconsistent with efficient pricing.
ICLV	Eq. (10)	The closing price’s position within the daily high–low range reflects intraday order flow. Persistent imbalance or autocorrelation in CLV suggests that trading pressure is not instantly arbitraged away.
IGap	Eq. (11)	Measures systematic gaps between the previous close and the next opening price. Persistent or predictable gaps imply that overnight information is not efficiently incorporated at the open.
IRCVM	Eq. (12)	Compares volatility estimated from intraday ranges with volatility from close-to-close returns. Persistent discrepancies indicate incomplete incorporation of intraday information into closing prices.
CEI	Eq. (13)	First principal component of the standardized sub-indices. Serves as the main aggregate measure of inefficiency; higher CEI corresponds to greater deviations from EMH.
AVE	Eq. (14)	Simple arithmetic mean of the twelve standardized sub-indices. Used as a robustness check to confirm that results are not driven by PCA weighting.

The variables span multiple aspects of return dynamics, volatility persistence, distributional characteristics, and market microstructure. By construction, higher values of each sub-index indicate stronger deviations from the Efficient Market Hypothesis (i.e., lower efficiency). This property allows us to assess the evolution of efficiency over time, compare differences across ASEAN foreign exchange markets, and examine how efficiency responds to global and regional shocks.

Construction of Inefficiency Indices

To capture deviations from the weak-form Efficient Market Hypothesis, we construct twelve sub-indices that summarize distinct dimensions of return dynamics, volatility persistence, distributional properties, and market microstructure frictions. Each sub-index is designed to reflect a theoretical benchmark under efficiency and to increase in magnitude as inefficiency becomes more pronounced. The sub-indices are first computed at the monthly frequency by aggregating relevant statistics from daily data. For instance, serial dependence is measured by autocorrelations of daily returns within each month, volatility clustering is evaluated from the persistence of squared and absolute returns, while distributional anomalies are captured by skewness, kurtosis, and tail risk indicators. Market microstructure effects such as close–open gaps and intraday order-flow imbalances are incorporated to reflect information processing efficiency. Formally, the construction follows three steps. First, each sub-index is standardized into a z-score relative to its full-sample distribution, and extreme values are clipped at ±3 to mitigate outlier effects. Second, the standardized indicators are pooled into a monthly panel of twelve dimensions for each of the six markets. Third, we apply Principal Component Analysis to extract the Composite Efficiency Index. The first principal component (PC1) is retained and oriented to align positively with the simple average of the standardized sub-indices, ensuring interpretability. By construction, higher values of CEI represent greater market inefficiency.

As a robustness check, we compute an alternative measure (AVE), which is the unweighted arithmetic mean of the twelve standardized sub-indices. AVE provides a benchmark that does not rely on PCA weighting and allows us to evaluate whether results are sensitive to the choice of aggregation method. In practice, CEI and AVE are highly correlated, supporting the stability of the composite index. This composite framework has two key advantages. First, it condenses multiple, potentially noisy indicators of market behavior into a single index that reflects the dominant dimension of inefficiency. Second, it provides a consistent metric for comparing efficiency across time and across different ASEAN currency markets. In the subsequent analysis, CEI serves as the primary measure of market inefficiency, while AVE is used to validate robustness.

Regression Design and Validation

To validate the Composite Efficiency Index and evaluate its implications for foreign exchange market behavior, we employ a combination of descriptive checks and regression-based analyses. The approach is structured to align with two central empirical questions: (i) whether efficiency improves over time, and (ii) whether changes in inefficiency help forecast short-horizon predictability of returns.

Diagnostic checks. We first assess face validity by comparing peaks in CEI with major global and regional crises (e.g., the Global Financial Crisis of 2008–2009 and the COVID-19 pandemic in 2020). Convergent validity is verified through correlations between CEI and the simple average (AVE) of the twelve standardized sub-indices. Discriminant validity is assessed by checking whether CEI is distinct from volatility proxies such as realized volatility and the Parkinson range estimator. In addition, cross-market differences are analyzed using ANOVA and post-hoc pairwise tests, providing evidence on whether efficiency levels differ significantly across ASEAN markets.

Time-trend regressions. To examine whether markets have become more efficient over time, we estimate regressions of CEI on a linear time trend:

$$CEI_{i,t} = \alpha + \beta t + \mu_i + \varepsilon_{i,t} \tag{15}$$

where t t denotes the time index in months and $\alpha\ i$ α i captures market-specific effects. Standard errors are clustered by both market and time, and Driscoll–Kraay corrections are applied to account for cross-sectional dependence. A significantly negative β indicates that market inefficiency decreases over time, consistent with the Adaptive Market Hypothesis.

Predictability regressions. To test whether short-run fluctuations in inefficiency matter for return dynamics, we analyze the predictive role of the first difference of CEI (ΔCEI). Specifically, we estimate two types of models:

Predictability metrics: $Y_{i,t+1} = \alpha_i + \gamma_t + \beta \times \Delta CEI_{i,t} + \varepsilon_{i,t+1}$

where $Y_{i,t+1}$ is either the absolute first-order autocorrelation of daily returns within the next month ($|AC_1|$) or the variance-ratio deviation at horizon 2, $|VR(2)-1|$. A positive β implies that higher inefficiency today is followed by greater return predictability tomorrow.

Returns with state dependence: $ret_{i,t+1} = \alpha_i + \gamma_t + \beta_1 \Delta CEI_{i,t} + \beta_2 \left(\Delta CEI_{i,t} \times Stress_{i,t}\right) + \varepsilon_{i,t+1}$

where $ret_{i,t+1}$ is the next-month log return and $Stress_{i,t}$ is a dummy equal to one if the Parkinson volatility measure falls in the top quartile. The interaction term allows us to test whether the predictive effect of ΔCEI is amplified or attenuated under stress conditions, consistent with the AMH view that predictability is episodic and state-dependent.

4. Results

Descriptive Statistics and Correlation Matrix

The six ASEAN foreign exchange markets analyzed in this study represent a diverse group of emerging markets with heterogeneous levels of financial development, liquidity, and institutional maturity. Despite their integration with global capital markets, these economies still exhibit characteristics typical of emerging markets, including higher volatility, stronger sensitivity to external shocks, and episodes of structural inefficiency. Examining their return dynamics and inefficiency indicators provides valuable insights into the adaptive nature of market efficiency in the region. Table 2 presents the descriptive statistics of monthly exchange rate returns against the U.S. dollar, along with the standardized inefficiency sub-indices introduced earlier. Across markets, the mean monthly return is close to zero, consistent with the prediction of the Efficient Market Hypothesis that no systematic drift should be present. However, the volatility profile differs considerably. The Indonesian rupiah (JKSE) exhibits the highest return variability, with a standard deviation of 3.0% per month, followed by the Thai baht (SET) at 2.0%. By contrast, the Vietnamese dong (VNI) shows a much lower standard deviation of only 0.8%, although this reflects a shorter sample period beginning in 2008 and episodes of exchange rate stabilization. The Philippine peso (PSE) and Malaysian ringgit (KLCI) fall in the middle of the volatility spectrum, while the Singapore dollar (STI) exhibits both low average returns and modest volatility, consistent with Singapore's reputation as the most stable market in the region. Extreme values, such as monthly depreciations exceeding 17% in Indonesia, highlight the exposure of these markets to global and domestic shocks.

Turning to the inefficiency sub-indices, the lower panel of Table 2 confirms that all twelve measures are standardized with a mean of zero and unit variance by construction. Nonetheless, their ranges are economically meaningful. For example, the autocorrelation index (I^{AC}) varies between -2.41 and +2.43, signaling that while many months show near-random return behavior, others exhibit strong serial correlation inconsistent with market efficiency. Similarly, the volatility clustering measure (I^{VC}) reaches values above 5.6, indicating prolonged periods of predictable volatility dynamics. The Hurst exponent deviation (I^{H}) peaks at 9.1, reflecting long-memory characteristics in certain episodes. Measures such as tail risk (I^{Tail}) and close–location imbalance (I^{CLV}) also attain extreme values, consistent with markets experiencing sharp intraday swings or persistent order-flow pressures. These results underscore that ASEAN currency markets, despite an overall trend toward greater integration, still display episodes of pronounced inefficiency.

Table 2: Descriptive statistics for monthly returns and sub-indices.

Index	n	Mean	SD	Min	Median	Max
JKSE	309	0.0025	0.0300	-0.1753	0.0018	0.1449
KLCI	309	0.0003	0.0188	-0.0656	0.0000	0.0850
PSE	309	0.0012	0.0171	-0.0525	0.0002	0.1007
SET	289	-0.0009	0.0198	-0.0722	-0.0016	0.0861
STI	309	-0.0010	0.0155	-0.0678	-0.0012	0.0821
VNI	309	0.0018	0.0084	-0.0370	0.0006	0.0687
I^{AC}	1726	0.0000	0.9985	-2.4148	0.1878	2.4255
I^{VR}	1726	0.0000	0.9985	-0.4288	-0.2328	5.8622
I^{VC}	1726	0.0000	0.9985	-0.6293	-0.3190	5.6163
I^{RVP}	1726	0.0000	0.9985	-1.5673	-0.2017	5.8668
I^{H}	1726	0.0000	0.9985	-0.5897	-0.2695	9.1158
I^{Skew}	1726	0.0000	0.9985	-1.8479	-0.3227	4.7440
I^{Kurt}	1726	0.0000	0.9985	-2.5290	0.0862	3.1448
I^{RDI}	1726	0.0000	0.9985	-1.7389	-0.1815	3.3040
I^{Tail}	1726	0.0000	0.9985	-1.1179	-0.2854	9.3865
I^{CLV}	1726	0.0000	0.9985	-3.2419	0.0246	4.6123
I^{Gap}	1726	0.0000	0.9985	-1.6485	-0.2391	4.0893
I^{RCVM}	1726	0.0000	0.9985	-1.3212	-0.2921	6.4840

Table 3 reports the correlation matrix among the twelve sub-indices. Correlations are generally modest, with most coefficients below 0.40. This pattern suggests that the sub-indices capture distinct facets of inefficiency rather than overlapping dimensions. As expected, some clusters emerge: autocorrelation (I^{AC}) and variance ratio deviation (I^{VR}) are strongly related (0.60), reflecting their common link to return predictability. Volatility clustering (I^{VC}) also correlates with the range–volatility predictability measure (I^{RVP}) at 0.54, as both emphasize persistence in volatility dynamics. By contrast, return decomposition imbalance (I^{RDI}) and tail risk (I^{Tail}) show weak or even negative correlations with most other measures, underscoring their role in capturing unique microstructure frictions and rare-event risk. The low overall correlations provide empirical justification for the use of principal component analysis, since combining relatively independent sources of inefficiency enhances the interpretability of the composite index.

Table 3: Correlation matrix of the twelve sub-indices.

	I^{AC}	I^{VR}	I^{VC}	I^{RVP}	I^{H}	I^{Skew}	I^{Kurt}	I^{RDI}	I^{Tail}	I^{CLV}	I^{Gap}	I^{RCVM}
I^{AC}	1.0000	0.5995	0.3950	0.2311	0.2948	0.1049	0.1949	-0.1743	-0.1270	-0.0619	-0.1143	0.0653
I^{VR}	0.5995	1.0000	0.3533	0.2556	0.3517	0.0938	0.2646	-0.1834	-0.0879	-0.0658	-0.0897	0.1903
I^{VC}	0.3950	0.3533	1.0000	0.5414	0.1385	0.0491	0.0152	-0.0277	0.1056	-0.0231	-0.0280	0.0399
I^{RVP}	0.2311	0.2556	0.5414	1.0000	0.0330	0.2208	0.1983	-0.0597	0.1279	-0.0304	0.0828	0.0631
I^{H}	0.2948	0.3517	0.1385	0.0330	1.0000	0.0007	-0.0246	-0.0260	-0.0618	-0.0848	-0.1005	0.0835
I^{Skew}	0.1049	0.0938	0.0491	0.2208	0.0007	1.0000	0.6579	0.0477	-0.3793	-0.0408	-0.0064	0.0165
I^{Kurt}	0.1949	0.2646	0.0152	0.1983	-0.0246	0.6579	1.0000	-0.1546	-0.4199	0.0005	0.0036	0.1772
I^{RDI}	-0.1743	-0.1834	-0.0277	-0.0597	-0.0260	0.0477	-0.1546	1.0000	0.1428	0.0607	0.0411	-0.3803
I^{Tail}	-0.1270	-0.0879	0.1056	0.1279	-0.0618	-0.3793	-0.4199	0.1428	1.0000	-0.0706	-0.0445	-0.0434
I^{CLV}	-0.0619	-0.0658	-0.0231	-0.0304	-0.0848	-0.0408	0.0005	0.0607	-0.0706	1.0000	0.4430	-0.1385
I^{Gap}	-0.1143	-0.0897	-0.0280	0.0828	-0.1005	-0.0064	0.0036	0.0411	-0.0445	0.4430	1.0000	-0.1183
I^{RCVM}	0.0653	0.1903	0.0399	0.0631	0.0835	0.0165	0.1772	-0.3803	-0.0434	-0.1385	-0.1183	1.0000

Finally, Table 4 examines how the composite indices (CEI and AVE) relate to the individual sub-indices. Both composite measures load heavily on core dimensions of inefficiency. The CEI correlates strongly with autocorrelation (0.64), variance ratio deviation (0.63), and volatility clustering (0.46), confirming that return predictability and volatility persistence are central to inefficiency in these markets. The AVE, being an equal-weighted measure, also displays high correlations with these indicators, especially variance ratio deviation (0.64) and volatility clustering (0.60). Interestingly, some sub-indices exhibit negative correlations with the composites. For instance, I^{RDI} and I^{Tail} correlate negatively with CEI (-0.33 and -0.25, respectively), implying that when inefficiency manifests primarily through overnight–intraday imbalances or extreme returns, it does not necessarily coincide with the broader inefficiency dynamics captured by PCA. These differences highlight the complementary nature of the sub-indices: while some align closely with the composite inefficiency dimension, others provide independent information about specific anomalies.

Table 4: Correlation between CEI/AVE and each sub-index.

	CEI	**AVE**
I^{AC}	0.6405	0.6081
I^{VR}	0.6341	0.6429
I^{VC}	0.4632	0.5955
I^{RVP}	0.3120	0.5384
I^{H}	0.3641	0.4182
I^{Skew}	0.1885	0.2502
I^{Kurt}	0.3309	0.3111
I^{RDI}	-0.3285	-0.0564
I^{Tail}	-0.2523	0.0454
I^{CLV}	-0.0472	-0.0253
I^{Gap}	-0.0916	-0.0567
I^{RCVM}	0.2335	0.2031

The descriptive statistics confirm that ASEAN foreign exchange markets are characterized by low average returns but diverse volatility profiles, reflecting their emerging market status. The inefficiency sub-indices display rich variation, capturing both common and idiosyncratic deviations from EMH. The correlation structure demonstrates that no single measure dominates; instead, the CEI effectively synthesizes multiple weakly correlated dimensions into a coherent indicator of time-varying efficiency.

Time-Varying Dynamics of CEI

The Composite Efficiency Index (CEI) provides a synthetic measure of market inefficiency that can be tracked through time to assess how ASEAN currency markets evolve. The CEI displays pronounced time variation, with noticeable spikes around episodes of financial stress such as the Global Financial Crisis of 2008–2009 and the COVID-19 shock in 2020. These fluctuations indicate that market efficiency is not constant but adapts to external shocks, in line with the Adaptive Market Hypothesis. To formally assess these dynamics, we estimate panel regressions of CEI on a linear time trend using alternative specifications. At the same time, the longer-run pattern points to a gradual decline in inefficiency, consistent with market maturation and improvements in institutional frameworks across the region. Table 5 summarizes the results. The pooled OLS regression produces a significantly negative coefficient on time, implying that, on average, inefficiency declines by about 0.002 per month. Fixed-effects (FE) models confirm this pattern even after controlling for unobserved market heterogeneity. A two-way fixed-effects specification, which accounts for both market-specific and time-specific effects, yields similar results, reinforcing the robustness of the declining trend.

Table 5. Panel regression of CEI on time trend

Model	Coefficient on time (β)	Std. Error	t-value	p-value	R^2
Pooled OLS	-0.00215	0.00047	-4.5745	< 0.001	0.012
FE (markets)	-0.0005	0.00008	-6.25	< 0.001	0.022
FE (two-way)	-0.00056	0.00009	-6.2222	< 0.001	0.021

The coefficients are consistently negative and statistically significant across all models, providing strong evidence that ASEAN foreign exchange markets have become more efficient over the past two decades. However, the relatively low R^2 values (1–2%) indicate that while efficiency improves on average, short-term fluctuations remain large and are better explained by shocks than by a simple linear trend. In summary, the evidence supports the AMH framework: efficiency is not static but evolves adaptively. ASEAN markets have made significant progress toward greater efficiency, although episodes of severe inefficiency continue to arise during crisis periods.

Cross-Market Comparisons

While the time-series regressions demonstrate that ASEAN foreign exchange markets have generally become more efficient over time, important cross-country differences remain. Table 6 reports the average CEI for each of the six markets over the sample period. Lower values indicate greater efficiency, while higher values reflect persistent inefficiency.

Table 6. Average CEI by market

Market	Mean CEI	SD	Balanced sample (n=206)
STI (Singapore)	-1.31	1.01	Most efficient, stable
PSE (Philippines)	-0.95	0.71	Relatively efficient
SET (Thailand)	-0.69	1.18	Intermediate efficiency
KLCI (Malaysia)	-0.17	1.48	Moderate inefficiency
JKSE (Indonesia)	0.55	1.35	Persistent inefficiency
VNI (Vietnam)	1.6	1.49	Least efficient

The ranking reveals a clear gradient across ASEAN markets. Singapore (STI) and the Philippines (PSE) exhibit the lowest CEI levels, consistent with their more advanced financial infrastructures and regulatory frameworks. Thailand (SET) lies in the middle, while Malaysia (KLCI) and Indonesia (JKSE) show higher inefficiency. Vietnam (VNI) stands out with the highest CEI, reflecting its relatively recent liberalization, lower liquidity, and greater exposure to capital flow volatility.

To assess whether these differences are statistically significant, we conduct an analysis of variance. Levene's test strongly rejects the null of equal variances ($F(5, 1230) = 19.94$, $p < 0.001$), indicating heterogeneity in CEI dispersion across markets. Despite this, both classical and robust (Welch) ANOVA confirm systematic efficiency differences across ASEAN markets at the 1% level. Post-hoc pairwise comparisons (Bonferroni-adjusted t-tests) provide a sharper picture of these disparities (Table 7). Vietnam is consistently less efficient than all other markets, while Indonesia also differs significantly from all peers, underlining its persistent inefficiency. Malaysia occupies an intermediate position, being significantly different from Thailand, the Philippines, Singapore, and Indonesia. Within the more efficient group, Singapore outperforms both the Philippines and Thailand, although the difference between the Philippines and Thailand is not statistically significant.

Table 7. Pairwise comparisons of CEI means (Bonferroni-adjusted p-values)

	JKSE	KLCI	PSE	SET	STI	VNI
JKSE	–	<0.001	<0.001	<0.001	<0.001	<0.001
KLCI		–	<0.001	0.001	<0.001	<0.001
PSE			–	0.528	0.040	<0.001
SET				–	<0.001	<0.001
STI					–	<0.001
VNI						–

Note: Boldface ($p < 0.05$) indicates statistically significant differences.

The results highlight a consistent efficiency hierarchy: STI (most efficient) < PSE ≈ SET < KLCI < JKSE < VNI (least efficient). These findings reinforce the earlier time-series evidence, showing that while ASEAN markets share an overall trend of adaptation toward greater efficiency, cross-sectional heterogeneity remains large and persistent. Markets with more developed institutions and deeper liquidity, such as Singapore and the Philippines, are systematically more efficient, whereas Vietnam and Indonesia continue to exhibit significant inefficiency.

Event-Driven Peaks of Inefficiency

Beyond gradual improvements in market efficiency, the CEI also exhibits pronounced spikes during episodes of financial turmoil. These peaks serve as face-validity checks for the index: if the CEI is a meaningful measure of inefficiency, it should rise in response to major shocks that temporarily undermine market functioning. The results align well with this expectation. Table 8 lists the top event-driven peaks of inefficiency for each ASEAN market. Notably, the Global Financial Crisis (2008–2009) triggered some of the highest CEI values in Indonesia (JKSE), Thailand (SET), and Singapore (STI), reflecting severe contagion and capital flight from emerging Asia. The COVID-19 pandemic (March 2020) generated another set of extraordinary spikes, most evident in Indonesia and Singapore. Other shocks are market-specific: the Taper Tantrum of 2013 sharply affected Indonesia, the commodity price collapse of 2017–2018 elevated inefficiency in Malaysia and Vietnam, and the domestic stock market turbulence in 2022–2023 produced unusually high CEI levels in Vietnam.

Table 8. Event-driven peaks of CEI by market

Market	Peak period	CEI value	Event context
JKSE (Indonesia)	2008-11	3.01	Global Financial Crisis
JKSE (Indonesia)	2013-08	4.76	Taper Tantrum (capital outflows)
JKSE (Indonesia)	2020-03	4.79	COVID-19 pandemic shock
KLCI (Malaysia)	2001-02	4.74	Post-Asian crisis adjustment
KLCI (Malaysia)	2002-12	4.69	Early 2000s global slowdown
KLCI (Malaysia)	2017-10	4.59	Commodity/FX pressures
SET (Thailand)	2008-08 to 2008-11	3.39–3.53	Global Financial Crisis
STI (Singapore)	2001-12 / 2002-01	0.24–0.31	Dot-com bust aftermath
STI (Singapore)	2013-07	-0.04	Regional capital volatility
STI (Singapore)	2020-03	2.89	COVID-19 pandemic shock
VNI (Vietnam)	2009-12	4.57	Post-GFC adjustment
VNI (Vietnam)	2018-02 / 2018-03	4.62–4.62	Stock market turbulence
VNI (Vietnam)	2022-12 / 2023-01	4.63–4.71	Domestic market crisis

These patterns confirm that the CEI captures adaptive inefficiency, rising sharply during stress episodes and receding once markets stabilize. The timing and magnitude of the peaks correspond closely with well-documented crises, underscoring the external validity of the index. At the same time, the heterogeneity of peak events across markets illustrates how local vulnerabilities—such as reliance on commodity exports in Malaysia or nascent capital markets in Vietnam—interact with global shocks to shape efficiency dynamics.

CEI and Short-Horizon Predictability

We next examine whether monthly changes in market inefficiency—captured by the first difference of the Composite Efficiency Index (ΔCEI)—help forecast one-step-ahead predictability of returns rather than mean returns themselves. Two standard predictability metrics are considered at the monthly level: the absolute first-order autocorrelation of daily returns within the month ($|AC_1|$) and the variance-ratio deviation at horizon 2, $|VR(2)-1|$. Using two-way fixed effects by market and month with double clustering (id & month), we find that ΔCEI is positively associated with both measures of next-month predictability. Specifically, ΔCEI loads at 0.024 (p-value = 0.052) when the dependent variable is $|AC_1|_{t+1}$, and at 0.024 (p-value = 0.064) for $|VR(2)-1|_{t+1}$. Although the effects are modest in magnitude and only marginally significant at the 10% level, they are fully consistent with the design of CEI: months in which inefficiency increases are followed by months with more serial dependence and stronger variance-ratio departures from the random walk.

By contrast, when we regress next-month mean returns on ΔCEI (two-way fixed effects with Driscoll–Kraay/clustered SEs), the coefficients are small and statistically insignificant overall (Section 4.2), indicating that martingale-difference behavior of returns persists on average even as predictability metrics vary. A simple state-dependence check interacting ΔCEI with a "stress" dummy (top quartile of volatility by the Parkinson proxy) shows a positive baseline effect of ΔCEI ($\beta \approx 0.0019$, p-value < 0.05) that is attenuated in stress months (interaction ≈ −0.0032, ns). Taken together, the findings align with the Adaptive Market Hypothesis: predictability is episodic and tied to shifts in market conditions, with ΔCEI capturing those shifts primarily through changes in serial dependence and variance-ratio diagnostics rather than through mean-return premia.

Table 9. ΔCEI and Short-Horizon Predictability of ASEAN Equity Markets

Panel A. Predictability metrics							
Dependent variable (t+1)	ΔCEI_t coef.	Std. Error	t-value	p-value	N	FE: id	FE: month
	AC_1	t+1	0.024445	0.009614	2.543	0.052	1681
	VR(2)−1	t+1	0.024208	0.010225	2.368	0.064	1681
Panel B. Mean returns							
Dependent: ret_{t+1}	ΔCEIt_t	ΔCEIt_t × Stress	SE type	N	FE: id	FE: month	Adj. R^2
Baseline + interaction	0.001893*	−0.003210 (ns)	Cluster (id & month)	1714	Yes	Yes	0.398

*(Notes: "Stress" indicates months in the top quartile of Parkinson volatility (by market). Standard errors clustered by market and month. ***, **, * denote significance at 1%, 5%, and 10%; "ns" = not significant.)*

5. Discussion

The results of this study provide novel evidence on the adaptive and multidimensional nature of efficiency in ASEAN foreign exchange markets. Several implications emerge when juxtaposing our findings with the broader literature on market efficiency.

First, time-varying efficiency. The significant downward trend in the Composite Efficiency Index confirms that inefficiency has declined steadily since 2000, reflecting financial liberalization, regulatory reforms, and technological progress. This trajectory is consistent with Lo [3], which views efficiency as evolving rather than static. Our evidence complements earlier equity-focused research [7-9], but crucially extends the adaptive framework to FX markets—an arena historically seen as the hardest to beat due to the dominance of the random walk benchmark [2]. By showing that ASEAN currencies have grown more efficient over time, our study highlights that adaptation is not confined to equity markets but is equally observable in globally integrated currency markets.

Second, multidimensional inefficiency. The twelve sub-indices reveal that inefficiency does not arise from a single source but spans serial dependence, volatility clustering, distributional anomalies, and market microstructure frictions. Correlation analysis shows that while autocorrelation (I^{AC}), variance ratio deviations (I^{VR}), and volatility clustering (I^{VC}) load most heavily onto CEI, other indices, such as return decomposition imbalance (I^{RDI}) and tail risk (I^{Tail}), provide distinct signals. This validates the use of a composite measure rather than relying solely on traditional single-dimension tests. Prior studies such as Tran & Leirvik (2019) emphasized autocorrelation-based measures, but our findings demonstrate the importance of integrating volatility and microstructure perspectives [15, 19, 23].

Third, cross-market heterogeneity. Despite an overall regional trend toward greater efficiency, large differences persist across countries. Singapore emerges as the most efficient market, while Vietnam consistently ranks as the least efficient. This gradient mirrors differences in institutional quality, market depth, and integration. The finding parallels Bock and Geissel (2024), who show that efficiency varies systematically across European markets, and suggests that national characteristics shape efficiency trajectories even within a regionally integrated FX system [9]. For policymakers, these differences highlight the importance of strengthening domestic institutions to foster efficiency.

Fourth, event-driven inefficiency. CEI spikes during global crises (e.g., the 2008–2009 Global Financial Crisis, the 2020 COVID-19 pandemic) provide strong face validity, showing that inefficiency is episodic and tied to shocks. This aligns with the AMH view that efficiency fluctuates with market stress. Interestingly, local crises such as Vietnam's 2022–2023 stock market turbulence or Indonesia's exposure during the Taper Tantrum also triggered sharp inefficiency, underscoring the role of domestic vulnerabilities. These findings resonate efficiency deteriorates sharply during systemic stress [20, 22].

Fifth, predictive validity of CEI. The most novel contribution of this study is demonstrating that changes in CEI (ΔCEI) forecast subsequent predictability metrics such as autocorrelation and variance ratio deviations, though not mean returns themselves. This distinction is crucial: returns continue to satisfy the martingale-difference property on average, yet inefficiency manifests through higher-order dependencies. This result strengthens the argument that CEI captures genuine inefficiency dynamics, consistent with the AMH, and extends prior composite indices by validating their forward-looking content.

Collectively, our findings advance the literature in three ways. First, we provide the first comprehensive composite measure of efficiency for ASEAN FX markets, filling a geographic and asset-class gap. Second, we integrate diverse inefficiency diagnostics into a coherent framework, demonstrating that inefficiency is multidimensional and episodic. Third, we establish predictive validity, showing that inefficiency indices not only describe but also forecast short-run predictability metrics. These contributions expand the scope of AMH testing and offer practical insights for investors, risk managers, and policymakers concerned with the stability of currency markets.

Our results carry several implications for policymakers and market regulators. First, the evidence of declining inefficiency over time suggests that financial liberalization, technological progress in trading infrastructure, and stronger monetary frameworks have supported market integration and price discovery. Policymakers in less efficient markets, such as Vietnam and Indonesia, may prioritize enhancing transparency, reducing transaction costs, and strengthening institutional quality to converge toward regional benchmarks like Singapore. Second, the episodic spikes in inefficiency around global and local crises highlight the importance of macroprudential measures and credible communication strategies to dampen volatility and restore confidence during stress. Third, the predictive content of CEI implies that regulators could use such composite indices as early-warning tools for monitoring financial stability, complementing traditional volatility and liquidity indicators. Finally, from an investor's perspective, understanding efficiency dynamics across currencies can inform portfolio diversification, hedging strategies, and timing decisions, particularly under the Adaptive Market Hypothesis framework.

To ensure the robustness of our findings, we conducted several complementary exercises. First, instead of PCA, we computed an equal-weighted average (AVE) of the twelve standardized sub-indices; the results are highly correlated ($\rho \approx 0.83$) with CEI, confirming that our conclusions are not driven by PCA loadings. Second, alternative window lengths (24- and 36-month rolling estimates) for the sub-indices produced qualitatively similar trends and event-driven spikes. Third, we tested heterogeneity across markets using both ANOVA and pairwise t-tests, which consistently identified Singapore as the most efficient and Vietnam as the least efficient. Fourth, regressions were re-estimated with both clustered and Driscoll–Kraay standard errors, and the significance of time trends and predictive regressions remained intact. Finally, removing crisis periods did not materially alter the long-term declining trend of inefficiency, although the magnitudes were somewhat attenuated. Together, these robustness checks confirm that our results are stable across specifications and estimation techniques.

6. Conclusion

This study develops and applies a Composite Efficiency Index to measure the time-varying efficiency of six ASEAN foreign exchange markets over the period 2000–2025. By aggregating twelve sub-indices that capture serial dependence, volatility clustering, distributional anomalies, and market microstructure frictions, the CEI provides a comprehensive view of deviations from the Efficient Market Hypothesis. The results reveal a significant decline in inefficiency over time, consistent with the Adaptive Market Hypothesis, while also documenting pronounced spikes during episodes of global and local crises. Cross-market comparisons show that Singapore is consistently the most efficient market, whereas Vietnam lags behind, reflecting heterogeneity in institutional development and market integration. Furthermore, changes in CEI predict higher-order dependencies such as autocorrelation and variance-ratio deviations, demonstrating the predictive validity of the index.

The findings have several implications. For policymakers, the gradual improvement in efficiency highlights the benefits of financial liberalization and technological upgrades, suggesting that further reforms can help lagging markets converge toward regional leaders. The episodic deterioration of efficiency during crises underscores the need for robust macroprudential frameworks and early-warning systems, where CEI can serve as a useful monitoring tool. For investors and market participants, the results support the AMH perspective that predictability is context-dependent and concentrated in turbulent periods, offering insights for hedging and portfolio diversification strategies.

Despite its contributions, this study has several limitations that point to future research opportunities. First, the analysis relies on monthly aggregation of daily data, which may smooth out high-frequency inefficiency patterns; extending the framework to intraday horizons could yield richer insights. Second, while PCA provides a statistically sound method of index construction, alternative machine learning techniques may capture nonlinear interactions among sub-indices more effectively. Third, the scope is limited to ASEAN currencies against the U.S. dollar; expanding to cross-rates or other emerging regions could test the generalizability of the CEI. Finally, integrating behavioral or sentiment-based measures with the CEI could further strengthen its explanatory and predictive power.

Funding: This research is funded (supported) by University of Economics Ho Chi Minh City, Vietnam (UEH)

Conflicts of Interest: The authors declare that there are no conflicts of interest regarding the publication of this paper.

References

[1] E.F. Fama, Efficient Capital Markets: A Review of Theory and Empirical Work, J. Financ. 25 (1970), 383-417. https://doi.org/10.2307/2325486.

[2] R. Meese, K. Rogoff, The Out-of-Sample Failure of Empirical Exchange Rate Models: Sampling Error or Misspecification?, in: Exchange Rates and International Macroeconomics, pp. 67-112, National Bureau of Economic Research, Inc. (1983). https://ideas.repec.org/h/nbr/nberch/11377.html.

[3] A.W. Lo, The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective, J. Portf. Manag. 30 (2004), 15-29.

[4] Y. Liu, The Short-Run and Long-Run Components of Idiosyncratic Volatility and Stock Returns, Manag. Sci. 68 (2022), 1573-1589. https://doi.org/10.1287/mnsc.2020.3884.

[5] O. Ozkan, Impact of COVID-19 on Stock Market Efficiency: Evidence from Developed Countries, Res. Int. Bus. Financ. 58 (2021), 101445. https://doi.org/10.1016/j.ribaf.2021.101445.

[6] J.H. Kim, A. Shamsuddin, Are Asian Stock Markets Efficient? Evidence from New Multiple Variance Ratio Tests, J. Empir. Financ. 15 (2008), 518-532. https://doi.org/10.1016/j.jempfin.2007.07.001.

[7] V.L. Tran, T. Leirvik, A Simple but Powerful Measure of Market Efficiency, Financ. Res. Lett. 29 (2019), 141-151. https://doi.org/10.1016/j.frl.2019.03.004.

[8] R. Mattera, F. Di Sciorio, J.E. Trinidad-Segovia, A Composite Index for Measuring Stock Market Inefficiency, Complexity 2022 (2022), 9838850. https://doi.org/10.1155/2022/9838850.

[9] J. Bock, S. Geissel, Evolution of Stock Market Efficiency in Europe: Evidence from Measuring Periods of Inefficiency, Financ. Res. Lett. 62 (2024), 105129. https://doi.org/10.1016/j.frl.2024.105129.

[10] C. Cantú, Effects of Capital Controls on Foreign Exchange Liquidity, J. Int. Money Financ. 93 (2019), 201-222. https://doi.org/10.1016/j.jimonfin.2019.01.006.

[11] A.S. Azad, Random Walk and Efficiency Tests in the Asia-Pacific Foreign Exchange Markets: Evidence from the Post-Asian Currency Crisis Data, Res. Int. Bus. Financ. 23 (2009), 322-338. https://doi.org/10.1016/j.ribaf.2008.11.001.

[12] N. Danila, P. Aggarwal, Foreign Exchange Market Herd Behaviour: Empirical Study in Asean-5 Countries, Asian J. Account. Res. 10 (2024), 63-75. https://doi.org/10.1108/ajar-11-2023-0381.

[13] A.W. Lo, A.C. MacKinlay, Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test, Rev. Financ. Stud. 1 (1988), 41-66. https://doi.org/10.1093/rfs/1.1.41.

[14] V. Hřebačka, Analyzing Market Efficiency: The Role of Business Cycles, Risk Aversion, and Occam's Razor in the Adaptive Market Hypothesis, Financ. Res. Lett. 75 (2025), 106840. https://doi.org/10.1016/j.frl.2025.106840.

[15] S. Ding, T. Cui, X. Wu, et al. Supply Chain Management Based on Volatility Clustering: The Effect of CBDC Volatility, Res. Int. Bus. Financ. 62 (2022), 101690. https://doi.org/10.1016/j.ribaf.2022.101690.

[16] M.B. Garman, M.J. Klass, On the Estimation of Security Price Volatilities from Historical Data, J. Bus. 53 (1980), 67-87. https://doi.org/10.1086/296072.

[17] D.W. Shin, Forecasting Realized Volatility: A Review, J. Korean Stat. Soc. 47 (2018), 395-404. https://doi.org/10.1016/j.jkss.2018.08.002.

[18] J. Alvarez-Ramirez, E. Rodriguez, C. Ibarra-Valdez, Long-Range Correlations and Asymmetry in the Bitcoin Market, Physica A: Stat. Mech. Appl. 492 (2018), 948-955. https://doi.org/10.1016/j.physa.2017.11.025.

[19] A. Ranaldo, F. Somogyi, Asymmetric Information Risk in FX Markets, J. Financ. Econ. 140 (2021), 391-411. https://doi.org/10.1016/j.jfineco.2020.12.007.

[20] C.M. Boya, Time-Varying Efficiency and Economic Shocks: A Rolling DFA Test in Western European Stock Markets, Int. J. Financ. Stud. 13 (2025), 157. https://doi.org/10.3390/ijfs13030157.

[21] D. Lou, C. Polk, S. Skouras, A Tug of War: Overnight Versus Intraday Expected Returns, J. Financ. Econ. 134 (2019), 192-213. https://doi.org/10.1016/j.jfineco.2019.03.011.

[22] K. Li, C. Xie, Y. Ouyang, T. Mo, Y. Feng, Tail Risk Spillovers in the Stock and Forex Markets at the Major Emergencies: Evidence from the G20 Countries, Int. Rev. Financ. Anal. 96 (2024), 103712. https://doi.org/10.1016/j.irfa.2024.103712.

[23] V. Bogousslavsky, D. Muravyev, WHO Trades at the Close? Implications for Price Discovery and Liquidity, J. Financ. Mark. 66 (2023), 100852. https://doi.org/10.1016/j.finmar.2023.100852.

[24] N. Boyarchenko, L.C. Larsen, P. Whelan, The Overnight Drift, Rev. Financ. Stud. 36 (2023), 3502-3547. https://doi.org/10.1093/rfs/hhad020.

[25] T.G. Andersen, T. Bollerslev, Deutsche Mark–Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies, J. Financ. 53 (1998), 219-265. https://doi.org/10.1111/0022-1082.85732.

[26] B. Korkusuz, D. Kambouroudis, D.G. McMillan, Do Extreme Range Estimators Improve Realized Volatility Forecasts? Evidence from G7 Stock Markets, Financ. Res. Lett. 55 (2023), 103992. https://doi.org/10.1016/j.frl.2023.103992.