When Are Stock Markets Predictable? Evidence Across Time, Firm Size, and Market Development

Harshi Pannala*

Abstract

Stock market efficiency is often tested by asking whether past returns help predict future returns. The Adaptive Markets Hypothesis suggests that this predictability may change as investors respond to new conditions. This article uses variance ratio tests to examine how short-term return predictability varies across decades, firm sizes, major market events, and developed, emerging, and frontier markets. Decade-by-decade analysis of large-cap, medium cap and small-cap U.S. stocks from the 1960s through the 2020s assess whether predictability follows a steady trend over time. A lead–lag analysis tests whether large-cap returns help predict subsequent small-cap returns. The international comparison examines whether markets differ in their responses to major shocks and how quickly their earlier patterns return.

Introduction

The Efficient Market Hypothesis (EMH) suggests that market prices reflect the information available to investors, making it difficult to earn higher returns without taking additional risk [2]. Market efficiency is generally divided into three forms: weak, semi-strong (or intermediate), and strong. Weak-form efficiency means that prices already reflect information from past prices, returns, and trading activity. Semi-strong efficiency includes all publicly available information, such as company reports and economic news. Strong-form efficiency goes further by suggesting that prices reflect both public and private information.

This paper focuses only on weak-form efficiency because it examines whether past returns contain patterns that can help predict future returns. If a market is weak-form efficient, historical price movements should not provide a reliable basis for earning abnormal returns. The variance-ratio test examines this idea by testing whether the return variance increases proportionally with time, as expected under a random walk [8]. The Adaptive Markets Hypothesis (AMH) offers a more flexible view, suggesting that market efficiency changes as investors learn and respond to new economic conditions [7]. Drawing on evolutionary biology, AMH compares financial markets to ecosystems, where investors compete and adapt until profitable opportunities become increasingly scarce. Just as an environmental disruption can favor species with advantageous traits, a shift in market conditions can create new inefficiencies that allow certain investors to gain a competitive advantage and generate alpha until others adapt. From this perspective, predictable patterns may appear during certain market regimes and disappear once conditions change or investors begin exploiting them, providing the motivation for testing algorithmic trading strategies across different regimes.

Previous empirical studies have provided evidence that market efficiency may change over time rather than remain constant. Kim et al. [5] examined more than a century of U.S. stock market data and found that return predictability varied with economic conditions, consistent with the Adaptive Markets Hypothesis. Similarly, Lim and Brooks [6], in a comprehensive review of the literature, highlighted growing evidence that market efficiency evolves over time. These findings raise an important question: are stock markets becoming progressively more efficient, or does

*American High School

efficiency fluctuate in response to economic crises, technological developments, and changes in investor behavior? This paper investigates this question by comparing variance ratios across historical periods and decades, while examining the sensitivity of the results to alternative estimation windows and statistical adjustments.

Market efficiency may also differ across firm sizes and levels of economic development. Lo and MacKinlay [9] found evidence that returns of large firms tend to lead those of smaller firms, suggesting that information may not be incorporated into all stock prices simultaneously. Hou [4] further demonstrated that the delayed diffusion of industry information contributes to this lead–lag relationship. Internationally, Bekaert et al. [1] showed that liquidity plays an important role in expected returns in emerging markets. Motivated by these findings, this study extends its analysis beyond the S&P 500 to compare large-cap and small-cap stocks, investigate whether large-cap returns predict subsequent small-cap returns, and examine sector-level efficiency across developed, emerging, and frontier markets. By considering time, firm size, and market development together, the paper aims to provide a broader understanding of how stock return predictability varies under different market conditions.

The remainder of this paper is organized as follows. The Methods section introduces the variance-ratio methodology and the statistical framework used to evaluate return predictability. The Results section first examines historical market events and decadal changes in efficiency, followed by robustness tests of the results. It then investigates differences between large-cap and small-cap stocks, tests the lead–lag relationship between their returns, and compares sector-level efficiency across developed, emerging, and frontier markets. Finally, the Conclusion summarizes the main findings and discusses their implications for the Adaptive Markets Hypothesis.

Methods

To investigate how weak-form market efficiency changes over time, I analyzed historical stock market returns using the variance-ratio test developed by Lo and MacKinlay [8, 3]. The primary analysis uses daily S&P 500 price data to examine market behavior across decades and major economic events. Additional analyses compare large-cap and small-cap stocks and examine differences across developed, emerging, and frontier markets.

Stock market prices generally change over time, making it difficult to compare price movements directly across different periods. Instead, I calculated the daily logarithmic return, defined as

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right).$$

Variance-Ratio Test

The variance-ratio test evaluates whether stock returns behave as expected under a random walk. In a random walk, price changes are unpredictable from previous changes, and returns have no serial correlation. Consequently, the variance of returns should increase proportionally with the holding period. For example, the variance of a five-day return should be approximately five times the variance of a one-day return.

To test this relationship, I calculated the variance ratio for holding periods of $q = 2, 5, 10$, and 20 trading days:

$$VR(q) = \frac{\mathrm{Var}(r_t + r_{t-1} + \cdots + r_{t-q+1})}{q\,\mathrm{Var}(r_t)}.$$

Under the null hypothesis of random-walk, $VR(q) = 1$, while values greater than 1 indicate positive or less than 1 indicate negative serial correlation.

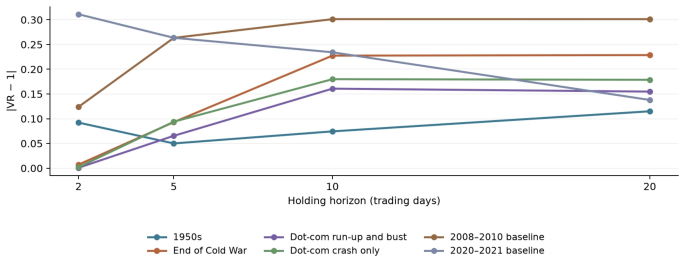


Figure 1: Magnitude of deviation of S&P 500 variance ratios from unity ($|VR - 1|$) across six historical periods at holding horizons of 2, 5, 10, and 20 trading days, illustrating time-varying departures from market efficiency.

Historical and 10-Year Comparisons

Variance ratios were calculated separately for different historical periods, including the 1950s, the end of the Cold War, the dot-com boom and crash, the 2008 financial crisis, and the COVID-19 pandemic. To compare the magnitude of deviations across periods, I also calculated $|VR(q) - 1|$, which measures the absolute distance from the random-walk benchmark without considering the direction of the deviation.

Firm Size and Lead–Lag Analysis

To examine whether market efficiency depends on company size, I compared variance ratios for large-cap and small-cap stocks across decades. I also investigated whether returns of larger companies contain information about subsequent returns of smaller companies. This relationship is commonly called the lead–lag effect.

The relationship can be estimated using the regression

$$r_{t+1}^{S} = \alpha_d + \beta_d r_t^{L} + \varepsilon_{t+1}, \tag{1}$$

where r_t^L is the large-cap return on day t, r_{t+1}^S is the small-cap return on the following trading day, and β_d measures their relationship within decade d. A positive β_d indicates that higher large-cap returns tend to be followed by higher small-cap returns. A coefficient close to zero suggests little linear predictive relationship.

International Market Comparisons

Finally, I compared sector-level market efficiency across developed, emerging, and frontier markets during three periods: 2008–2010, 2020–2021, and 2022–2025. This analysis investigates whether industries and markets at different stages of economic development exhibit different levels of return predictability, particularly during periods of financial disruption.

Finally, I compared sector-level market efficiency across developed, developing (MSCI Emerging proxy), and frontier markets during three periods: 2008–2010, 2020–2021, and 2022–2025. The analysis used monthly, value-weighted GICS sector portfolio returns and variance-ratio horizons of 2, 3, and 6 months. This analysis investigates whether industries and markets at different stages of economic development exhibit different levels of return predictability, particularly during periods of financial disruption. Country scores were averaged equally within each sector and market group. The resulting scores are descriptive and do not formally test whether one market group is significantly more efficient than another.

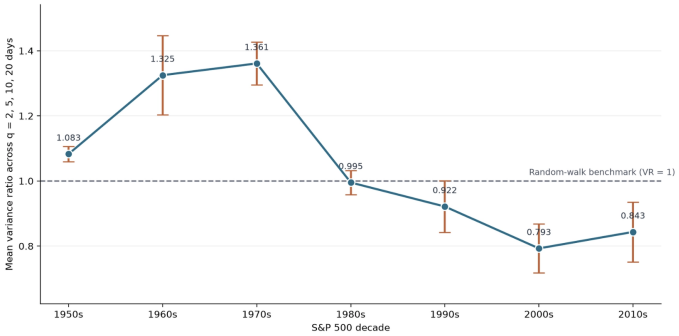


Figure 2: Mean variance ratios of the S&P 500 over 10 years, averaged across holding horizons of 2, 5, 10, and 20 trading days from the 1950s to the 2010s. The dashed line represents the random-walk baseline ($VR = 1$), highlighting changes in market efficiency across decades.

Results

Figure 1 compares the variance ratios of the S&P 500 across six historical periods using holding horizons of 2, 5, 10, and 20 trading days. During the 1950s, variance ratios remain above 1 across all horizons, indicating a positive serial correlation. In contrast, the end of the Cold War (1988–1990) shows a decline from approximately 1.01 at the 2-day horizon to 0.77 at the 20-day horizon. Similar patterns appear during the dot-com boom and bust (1995–2002) and the dot-com crash (2000–2002), with variance ratios falling below 0.85 at longer horizons. The 2008–2010 financial crisis period shows the largest deviations at longer horizons, reaching approximately 0.70, while the 2020–2021 period begins with a particularly low variance ratio of 0.69 at the 2-day horizon before gradually approaching 1.

The results show that the return patterns differ across historical periods, both in the direction and in the strength of serial correlation. This is consistent with the Adaptive Markets Hypothesis, which suggests that market efficiency changes as economic conditions and investor behavior evolve. However, only the 1950s' 2-day variance ratio remains statistically significant after correcting for multiple comparisons, meaning the other differences should be treated as observed patterns rather than definitive evidence of market inefficiency.

Figure 2 hows a shift in S&P 500 return patterns across decades. Variance ratios rise from 1.083 in the 1950s to 1.361 in the 1970s, suggesting increasingly strong positive serial correlation. This trend reverses in the 1980s, when the ratio falls to 0.995, before declining further to 0.793 in the 2000s and recovering slightly to 0.843 in the 2010s. The movement from values above 1 to values below 1 suggests a shift from momentum-like to mean-reversion-like behavior rather than a steady reduction in return predictability.

These fluctuations are consistent with the Adaptive Markets Hypothesis, as the degree and direction of return dependence change across historical periods. Notably, the 1980s appear closest to the random-walk benchmark, but this apparent efficiency does not persist in subsequent decades. This raises the possibility that market efficiency is better understood as a changing characteristic influenced by economic conditions rather than a continuously improving trend. But these observations on their own cannot establish that changing market conditions caused these shifts.

Figure 3 suggests that the apparent efficiency of the S&P 500 during the 1980s depends on how the data are grouped. While the pooled decade estimate places the 1980s closest to the random-walk benchmark, averaging annual estimates produces a larger deviation, suggesting that combining returns over longer periods may hide changes in market behavior.

For example, between 1980 and 1984, variance ratios remain above 1, indicating positive serial correlation. However, during 1985–1989, the ratios fall below 1 at longer horizons, suggesting a shift toward mean-reversion-like behavior. Combining these periods produces a full-decade estimate much closer to 1, potentially masking the contrasting return patterns within the decade.

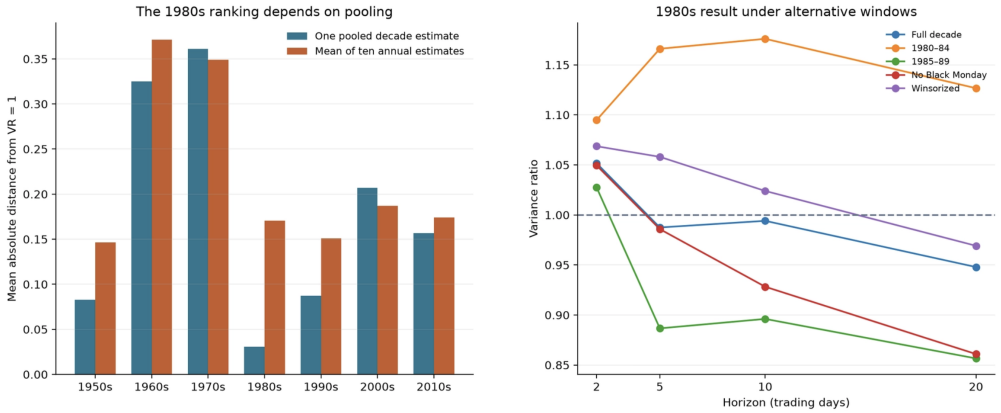
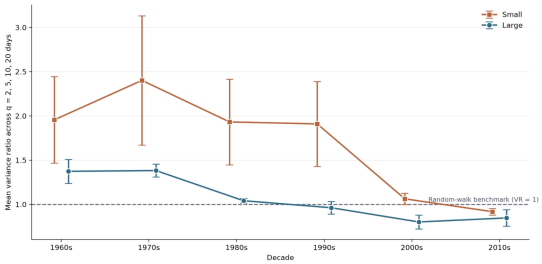
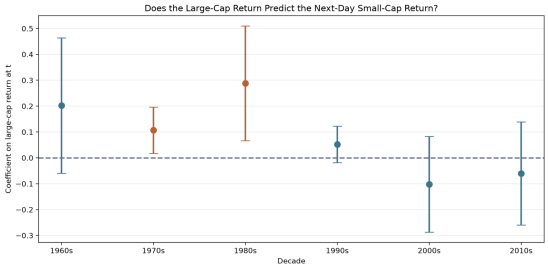


Figure 3: Robustness analysis of S&P 500 variance ratios across decades. Left: Comparison of mean absolute deviations from the random-walk benchmark ($VR = 1$) using pooled 10-year estimates and averages of annual estimates. Right: Sensitivity of the 1980s variance ratios to alternative sample windows, exclusion of Black Monday, and winsorization across holding horizons of 2, 5, 10, and 20 trading days.

Excluding Black Monday has little effect on the full-decade results, suggesting that this single event does not explain the observed pattern. Regarding the Adaptive Markets Hypothesis, a market may appear efficient over a long period even when its underlying return patterns change. Rather than reflecting consistently efficient behavior, a variance ratio near 1 may conceal opposing patterns across different market regimes, so examining shorter periods could provide a clearer understanding of how market efficiency evolves over time.



(a) Variance ratios for large- and small-cap stocks over 10 years



(b) Large-cap returns predicting next-day small-cap returns.

Figure 4: Comparison of market efficiency and return predictability across market capitalizations from the 1960s to the 2010s.

Figure 4 shows how the differences in return predictability between large- and small-cap stocks have narrowed over time. Both analyses point to a possible reduction in information gaps between market segments, as improvements in financial reporting, trading technology, and market accessibility may have made information more readily available to investors.

In Figure 4a, large-cap stocks initially had variance ratios much closer to 1 than small-cap stocks, indicating weaker return dependence. During the 1970s, small-cap stocks reached a mean variance ratio of approximately 2.40, compared to 1.38 for large-cap stocks. This gap gradually narrowed, and by the 2010s, small-cap variance ratios had fallen to approximately 0.92, while large-cap ratios reached 0.85. Interestingly, small-cap stocks ended the period closer to the random-walk benchmark, reversing the pattern seen in earlier decades.

Figure 4b can help explain this change. During the 1970s and 1980s, large-cap returns had a statistically significant positive relationship with next-day small-cap returns, with coefficients of

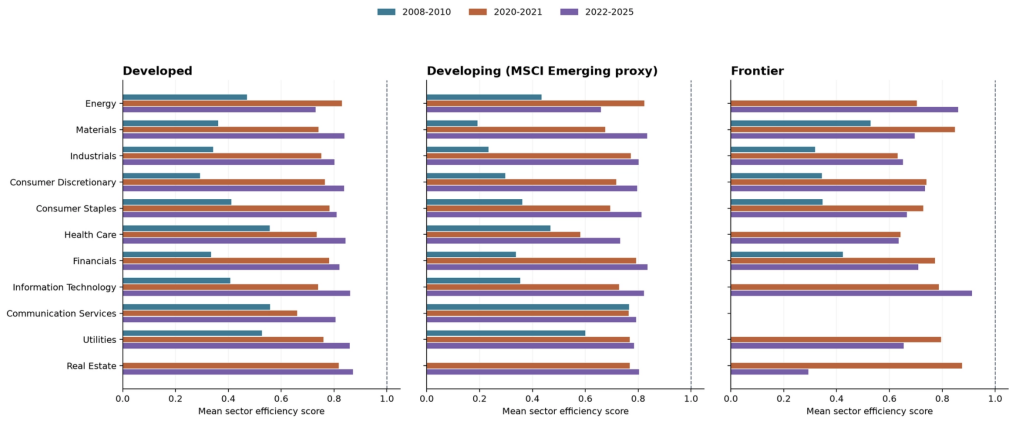


Figure 5: Comparison of sector-level market efficiency across developed, developing, and frontier markets during three periods: 2008–2010, 2020–2021, and 2022–2025. Horizontal bars represent mean sector efficiency scores, with the dashed line indicating the efficiency baseline of 1.

approximately 0.11 and 0.29, respectively. This is consistent with information reaching large-cap stocks before being reflected in smaller companies' prices. However, this relationship is no longer statistically significant from the 1990s onward, which may reflect faster information transmission between market segments.

These findings show that return predictability does not evolve uniformly across different parts of the market, consistent with AMH. While the narrowing information gap offers one possible explanation, the results also raise a broader question about how changes in market structure influence which forms of predictability persist or disappear over time.

In Figure 5, across developed and developing markets, the efficiency scores are generally lowest during the 2008–2010 financial crisis before improving in 2020–2021 and 2022–2025. By the most recent period, the two market groups have relatively similar scores across many sectors, indicating that differences in return predictability are not necessarily determined by a country's level of economic development.

Developed markets show particularly high scores in information technology, utilities, and real estate during 2022–2025, while energy remains relatively lower. Developing markets follow a similar pattern of improvement, although the extent varies by sector. Frontier markets, however, show less consistent changes. While information technology improves considerably, real estate experiences a sharp decline in 2022–2025. This result should be interpreted cautiously because the frontier real estate category includes data from only one country, which may not represent the general trend.

These findings are consistent with AMH, as efficiency changes across economic conditions and market segments rather than following a uniform pattern. The convergence between developed and developing markets also raises questions about whether market maturity alone determines efficiency. Instead, sector-specific conditions and differences in how markets respond to economic changes may play an important role. However, the scores are descriptive and do not establish whether these differences are statistically significant.

Conclusion

I investigated how stock market efficiency changes over time, across firm sizes, and across different levels of market development. Using variance-ratio tests, I found that stock return predictability does not follow a steady trend toward greater efficiency. Instead, the strength and direction of return dependence fluctuate across decades and major economic events, consistent with the Adaptive Markets Hypothesis. An especially interesting finding was that the S&P 500

appeared closest to the random-walk benchmark during the 1980s, but further analysis revealed contrasting patterns between the early and late parts of the decade.

The analysis also revealed that differences in return predictability between large-cap and small-cap stocks have narrowed considerably over time. Large-cap returns significantly predicted next-day small-cap returns during the 1970s and 1980s, but this relationship was no longer statistically significant from the 1990s onward. This is consistent with the possibility that improvements in technology and information transmission have reduced delays in how stock prices respond to new information. International comparisons further suggested that developed and emerging markets have become more similar in their sector-level efficiency scores, while frontier markets exhibit less consistent patterns. These findings suggest that market efficiency depends not only on time but also on market structure, firm size, and economic conditions.

Future research would be to investigate whether changes in trading technology, market liquidity, and information availability can explain the declining lead–lag relationship between large-cap and small-cap stocks. The international analysis could also be expanded to include more countries, longer historical periods, and formal statistical tests of differences between market groups. Another direction is to test whether the observed return patterns can support profitable trading strategies after accounting for transaction costs would help establish whether statistical predictability translates into economically meaningful opportunities. These extensions could provide a deeper understanding of how and why market efficiency evolves. Github Link: Code

References

[1] Geert Bekaert, Campbell R. Harvey, and Christian Lundblad. "Liquidity and Expected Returns: Lessons from Emerging Markets". In: *The Review of Financial Studies* 20.6 (2007), pp. 1783–1831. DOI: 10.1093/rfs/hhm030.

[2] Eugene F. Fama. "Efficient Capital Markets: A Review of Theory and Empirical Work". In: *The Journal of Finance* (1970).

[3] Yahoo Finance. *SP 500 Historical Data*. Retrieved: Daily value-weighted size portfolio returns. 2026.

[4] Kewei Hou. "Industry Information Diffusion and the Lead-lag Effect in Stock Returns". In: *The Review of Financial Studies* 20.4 (2007), pp. 1113–1138. DOI: 10.1093/revfin/hhm003.

[5] Jae H. Kim, Abul Shamsuddin, and Kian-Ping Lim. "Stock Return Predictability and the Adaptive Markets Hypothesis: Evidence from Century-Long U.S. Data". In: *Journal of Empirical Finance* 18.5 (2011), pp. 868–879. DOI: 10.1016/j.jempfin.2011.08.002.

[6] Kian-Ping Lim and Robert Brooks. "The Evolution of Stock Market Efficiency over Time: A Survey of the Empirical Literature". In: *Journal of Economic Surveys* 25.1 (2011), pp. 69–108. DOI: 10.1111/j.1467-6419.2009.00611.x.

[7] Andrew W Lo. "The adaptive markets hypothesis: Market efficiency from an evolutionary perspective". In: *Journal of Portfolio Management, Forthcoming* (2004).

[8] Andrew W. Lo and A. Craig MacKinlay. "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test". In: *The Review of Financial Studies* (1988).

[9] Andrew W. Lo and A. Craig MacKinlay. "When Are Contrarian Profits Due to Stock Market Overreaction?" In: *The Review of Financial Studies* 3.2 (1990), pp. 175–205. DOI: 10.1093/rfs/3.2.175.