



Are Financial Markets Efficient?

Professor Raghavendra Rau

10 June 2024

Introduction

My name is Raghavendra Rau and I'm a professor at the University of Cambridge. This is the final lecture in a series of Gresham lectures this year on the big ideas of finance. The lectures this year are drawn from my textbook A short introduction to corporate finance, published by Cambridge University Press.

As I said before, there are only six major ideas in finance, five of which have won their originators Nobel prizes in economics. What are these ideas?

1. Net Present Value (NPV): NPV is the key principle in investment decision-making, where the objective is to maximise the present value of future payoffs. It involves three steps: computing cash flows, discounting those flows to a single present value using a discount rate, and deciding the financing method, which affects taxes.

2. Portfolio Theory and the Capital Asset Pricing Model (CAPM)

- The interest rate or discount rate in NPV is determined by investors, based on available investment opportunities. Markowitz and Sharpe, Nobel laureates, proposed that individual investments are parts of portfolios. They combined portfolios with risk-free assets to determine the market portfolio. The discount rate is determined using the CAPM formula.

3. Capital Structure Theory:

- Capital structure explains how the discount rate changes based on the firm's financing decision – whether to go with debt or equity. Modigliani and Miller, Nobel winners, posited that in a perfect world, financing form does not affect firm value. But with real-world imperfections (like taxes), it does matter.

4. Option Pricing Theory:

- This theory discusses how to price options, which are contracts that give rights (but not obligations) to buy or sell assets. Black, Scholes, and Merton, with the latter two winning a Nobel, provided a solution based on the no-free lunch principle. They matched the cost of a portfolio replicating an option's payoff to the option's cost.

5. Asymmetric Information:

- This lecture deals with information imbalances in transactions, where one party has more information than the other. Akerlof, Spence, and Stiglitz, Nobel laureates, developed the key concepts in this area, illustrating how information imbalances affect markets from used cars to financial policies.

6. Market Efficiency:

- This is the idea we are discussing in this lecture. It discusses how markets reflect all available information. The debate lies in the relationship between market prices and NPV. Three Nobel winners, Kahneman, Fama, and Shiller, contributed pioneering ideas on this topic, discussing market behaviour and efficiency.

In essence, corporate finance revolves around six central ideas, with five of them recognized by Nobel Prizes.

What is market efficiency?

In the first four lectures in this series, we have made two basic assumptions. The first is that everyone has symmetric access to all the information necessary to make a decision. The second is that each person rationally uses all this information to decide.

In the last two lectures, we relax these assumptions. In the previous lecture on asymmetric information, we explore what happens if we do not have the same amount of information as the people on the other side of the trade. In this lecture, we examine the concept of market efficiency.

Let's start with understanding why prices are so important in a modern economy. Prices signal how resources should be allocated efficiently. When prices rise due to higher demand or lower supply, it indicates producers should increase production. When prices fall due to lower demand or oversupply, producers should cut back. This mechanism directs resources to their most valued uses. For example, during an oil shortage, rising prices signal producers to extract more oil while also nudging consumers to reduce usage or find alternatives. This helps balance supply and demand over time.

Prices also provide incentives for production and consumption decisions. Higher prices incentivize businesses to boost supply to earn more profits, while lower prices encourage consumers to buy more. This dynamic helps maintain market equilibrium. Analysing price data helps forecast future market conditions too. Businesses use price forecasting to predict demand changes and plan production levels. Investors forecast prices to guide investment decisions. Governments forecast prices to shape policies like interest rates, taxes, and subsidies to stabilise the economy.

If prices are set arbitrarily and markets are not efficient, it leads to inefficient allocation of resources and market failures. One major issue is the misallocation of resources, where supply and demand are out of balance, resulting in overproduction or underproduction of goods and services. Resources end up being used inefficiently instead of where they are most valued. Market distortions also arise when prices do not reflect true values, causing excess supply or demand imbalances. This leads to shortages or surpluses. Arbitrary pricing causes productive inefficiency, with goods not produced at lowest costs, and allocative inefficiency, with goods mismatched to consumer preferences. It can also enable negative externalities like environmental harm if prices do not account for external costs. Monopolies may arbitrarily distort prices for their benefit, reducing consumer welfare. Information asymmetry worsens when one party lacks full information for optimal decisions. Public goods face undersupply due to the free-rider problem, where people consume without paying their share under inefficient arbitrary pricing systems.

How did market efficiency develop?

Before 1960, there were two approaches to investing: Fundamental and technical analysis. Fundamental analysis involves examining a company's intrinsic value by analysing its financial statements, economic conditions, and other qualitative factors. The goal is to determine if a security is undervalued or overvalued compared to its estimated true worth. Fundamental analysts scrutinise earnings, revenue, profit margins, industry trends, and macroeconomic indicators like GDP growth and interest rates. This comprehensive assessment aims to forecast a company's future performance and potential stock price movements over an extended period.

In contrast, technical analysis almost completely disregards a company's fundamental factors and instead focuses solely on studying historical market data, primarily security prices and trading volume. Technical analysts operate under the assumption that all relevant information is already reflected in current prices. They use charts and mathematical indicators to identify patterns and trends, seeking to predict future price directions based on these visualisations. The emphasis is on short-term trading, capitalising on imminent price fluctuations rather than long-term valuations.

While fundamental analysis seeks to understand a security's intrinsic worth through rigorous examination of its financials and broader economic climate, technical analysis relies entirely on interpreting visual representations of market activity. Fundamental analysts are typically long-term investors concerned with a company's overall health and growth prospects. Conversely, technical analysts are traders focused on quickly recognizing and acting upon short-term price movements. Both methodologies claimed that they had different merits - fundamental provides depth but less timing precision, while technical offers potential exit/entry signals but lacks insight into underlying value drivers.

The concept of general equilibrium in economics laid the foundation for modern financial theories in the

1960s, culminating in Eugene Fama's Efficient Market Hypothesis (EMH). This evolution began with the early work of Adam Smith, who introduced the idea of the "invisible hand" guiding markets towards an optimal state through individual self-interest and voluntary trade. However, it was Léon Walras in the late 19th century who formalised the concept by developing a mathematical model describing how supply and demand in multiple markets interact to reach a state of simultaneous market clearance equilibrium.

In the mid-20th century, Kenneth Arrow and Gérard Debreu further advanced general equilibrium theory by proving the existence of such an equilibrium under certain conditions. Their work demonstrated that competitive markets could achieve Pareto efficiency, where no one can be made better off without making someone else worse off. This formalisation provided a robust theoretical framework for understanding how decentralised markets could lead to efficient outcomes, despite facing criticisms for its unrealistic assumptions like perfect information and rational behaviour.

Building on these principles, Eugene Fama introduced the Efficient Market Hypothesis in the 1970s, extending the idea of general equilibrium to financial markets. Fama's EMH posits that financial markets are "informationally efficient," meaning asset prices fully reflect all available information at any given time. He categorised market efficiency into three forms: weak, semi-strong, and strong, each with varying implications for the effectiveness of technical and fundamental analysis.

The weak form asserts that past price movements and volume data do not predict future prices, implying technical analysis is ineffective. The semi-strong form states all public information is reflected in stock prices, rendering both technical and fundamental analysis ineffective. The strong form claims all information, public and private, is incorporated into prices, making it impossible even for insiders to achieve abnormal returns.

Is market efficiency correct?

A straightforward example demonstrating that prices react quickly to information rather than events can be found in the study of stock market reactions to corporate announcements such as earnings announcements. When a company releases its quarterly earnings report, the stock price typically reacts almost immediately to the new information contained in the report. This reaction occurs because the earnings report provides fresh data about the company's financial performance, which investors use to update their expectations about the company's future prospects.

Studies have shown that the majority of the price response to earnings announcements happens very quickly, often within the first few trades after the announcement. For instance, research on after-trading-hours quarterly earnings announcements of NYSE and NASDAQ firms found that the price adjustment occurred during the first several post-announcement trades for NYSE stocks and was concentrated in the first post-announcement trade for NASDAQ stocks. The rapid price adjustment underscores that it is the new information about the company's earnings, rather than the mere occurrence of the earnings announcement event, which drives the price change. Investors react to the specific details in the earnings report, such as revenue, profit margins, and future guidance, which provide insights into the company's financial health and growth prospects. This example illustrates that financial markets are highly efficient in processing new information, leading to swift adjustments in stock prices based on the latest available data. The immediacy of the price reaction demonstrates that prices respond to the underlying information content rather than the events themselves, consistent with the principles of market efficiency.

Given the existence of hedge fund managers who earn billions of dollars annually through their trading activities, you might reasonably question how the idea of market efficiency can be correct. After all, if markets were truly efficient in pricing all available information, how could these elite investors consistently generate such massive profits year after year? So how can the market efficiency idea be correct?

Recall that the fundamental value of an asset is the present value of all cash flows that investors rationally expect to receive, basically the cash flows discounted at the rational discount rate. The one big problem with market efficiency is that it does not say how to calculate cash flows, nor does it tell us what the discount rate should be. It just says that what markets decide is correct based on the information available to them.

Testing market efficiency

Let's go back to the definition of market efficiency: In an efficient market, it is impossible to consistently

make abnormal returns on the basis of information.

The keywords here are: Impossible; Consistently; Abnormal; and Information.

Let's start with the word Abnormal. Abnormal returns are returns that, on average, beat the appropriate E[R] benchmark (e.g. CAPM). If you make returns of 10% this year, we do not know whether this is good or bad unless we have a benchmark.

Next let us move to the term Information. The EMH categorises market efficiency into three forms: weak, semi-strong, and strong, each representing a different level of information incorporation into asset prices.

Weak form efficiency asserts that *all* past trading information, such as historical prices and volumes, is already reflected in current stock prices. This means that analysing past price movements, known as technical analysis, cannot consistently yield abnormal returns. The idea is that price changes are random and not influenced by past trends, making it impossible to predict future prices based on historical data alone. For example, if an investor tries to predict future stock prices by studying past price charts and trading volumes, this strategy will not work in a weak form efficient market because all historical information is already embedded in the current price.

Semi-strong form efficiency takes it a step further by stating that all publicly available information is reflected in stock prices. This includes not only past trading data but also financial statements, news releases, and economic reports. As a result, neither technical analysis nor fundamental analysis, which involves evaluating a company's financial health, can consistently produce abnormal returns. Imagine a company announcing higher-than-expected earnings. In a semi-strong form efficient market, the stock price would quickly adjust to this new information, making it impossible for investors to profit from the news after it is released.

Strong form efficiency is the most stringent level, asserting that all information, both public and private, including insider information, is fully reflected in stock prices. According to this form, even insiders with access to confidential information cannot achieve abnormal returns because the market prices already incorporate all relevant information. Consider a CEO who knows about a groundbreaking development in their company that has not yet been made public. In a strong form efficient market, the stock price would already reflect this private information, rendering the CEO's insider knowledge useless for gaining extra profits.

An easy way to test weak form market efficiency is to use autocorrelation tests. Autocorrelation tests measure the correlation between current stock returns and past returns. If stock prices follow a random walk, there should be no significant autocorrelation, meaning past returns should not predict future returns. Most academic studies have shown that the market is pretty accurate in the weak form over short horizons.

Similarly, academics test semi-strong form market efficiency by examining how quickly and accurately stock prices adjust to new publicly available information. The primary method used for this purpose is the event study methodology. An event study analyses the impact of a specific event on the value of a security. ("Event study - Wikipedia") The goal is to determine whether the market efficiently incorporates new information into stock prices. If the market is semi-strong efficient, stock prices should adjust almost immediately to reflect the new information, leaving no opportunity for investors to earn abnormal returns based on that information.

To conduct an event study, researchers first define the event to be studied, such as earnings announcements, dividend changes, or mergers. They then establish an event window, which includes the period before and after the event date to capture the price reaction. A typical event window might be 10 days before and 10 days after the event. Next, a sample of firms or securities affected by the event is selected, such as companies that announced earnings on a specific date or those involved in a merger.

Researchers then calculate normal and abnormal returns. Normal returns estimate the expected return of the security in the absence of the event, often done using a market model that relates the return of the security to the return of a market index. Abnormal returns are calculated by subtracting the normal return from the actual return observed during the event window. These abnormal returns are then aggregated over the event window to assess the total impact of the event, known as the cumulative abnormal return.

Finally, statistical tests are performed to determine whether the abnormal returns are significantly different from zero. Common tests include the t-test and non-parametric tests like the rank test. If the abnormal

returns are statistically significant, it suggests that the market did not efficiently incorporate the new information into stock prices.

For example, studies have examined how stock prices react to quarterly earnings announcements. If the market is semi-strong efficient, the stock price should adjust immediately upon the announcement, reflecting the new information about the company's financial performance. Similarly, researchers have analysed stock price reactions to dividend changes and merger announcements, providing insights into the speed and accuracy of market responses.

Empirical evidence from event studies generally supports the semi-strong form of market efficiency, indicating that markets tend to quickly incorporate publicly available information into stock prices.

Another way to test semi-strong form of market efficiency is by examining the performance of mutual fund managers. Mutual fund managers claim that they use only public information to beat the market. Can they? What does the academic evidence show?

Academics test whether mutual fund managers can consistently beat the market by employing a variety of empirical methods and statistical analyses. The primary objective is to determine if mutual fund managers possess the skill to generate returns that exceed those of a passive benchmark index, after adjusting for risk and other factors.

One common approach is to analyse the historical performance of mutual funds using risk-adjusted performance measures. These measures include the Sharpe ratio, alpha, and beta, which help to account for the risk taken by the fund managers. For instance, alpha measures the excess return of a fund relative to a benchmark index, adjusted for the risk taken. A positive alpha implies that the fund has outperformed the benchmark, while a negative alpha suggests underperformance. Another benchmark uses factor models, such as the Carhart four-factor model, which includes market risk, size, value, and momentum factors. This model helps to isolate the portion of a fund's return that can be attributed to these common risk factors, allowing researchers to determine if any remaining excess return is due to managerial skill.

Academics also examine performance persistence, which refers to the ability of mutual funds to maintain their performance over time. Studies on performance persistence typically involve ranking funds based on their past performance and then tracking their future performance to see if top-performing funds continue to outperform. For example, research has shown that while some funds exhibit short-term persistence, long-term persistence is much less common, suggesting that luck rather than skill may play a significant role in short-term outperformance.

Overall, the empirical evidence on mutual fund performance suggests that while some managers may exhibit skill in the short term, consistently beating the market over the long term is exceedingly rare. This supports the semi-strong form of the EMH.

Finally, the evidence shows that markets are not strong-form efficient implying that investors can earn abnormal returns based on insider information.

So how can we explain the great investors?

There are two possibilities: luck and survivor bias.

Luck:

To illustrate why luck can explain the performance of mutual fund managers, consider a simple coin-tossing game. Imagine a scenario where 1,000 people each flip a coin. After each flip, those who get tails are eliminated, and those who get heads continue to the next round. Statistically, after about ten rounds, only one person will remain who has flipped heads every time. This person might seem extraordinarily skilled at flipping coins, but in reality, their success is purely due to luck. The large number of participants ensures that someone will inevitably achieve this unlikely outcome, even though each flip is independent and random.

This coin-tossing game parallels the performance of mutual fund managers. In the financial markets, thousands of mutual fund managers make investment decisions under conditions of uncertainty. Just as with the coin toss, some managers will inevitably experience streaks of good performance purely by chance. These streaks can create the illusion of skill, but they may not be indicative of genuine investment

skill.

Academic studies, such as those by Fama and French, have used statistical methods to distinguish between luck and skill in mutual fund performance. By analysing long-term performance data and employing techniques like bootstrap simulations, researchers can compare the actual distribution of fund returns to what would be expected by chance alone. These studies often find that while a small number of managers do exhibit skill, the majority of outperformance can be attributed to luck.

Fama and French's research demonstrated that the net returns of most actively managed funds are negative after accounting for fees and expenses. This suggests that, on average, active managers do not possess sufficient skill to consistently outperform passive benchmarks. The few managers who do achieve superior returns are often outnumbered by those who underperform, leading to an overall negative alpha in the aggregate.

The coin-tossing analogy helps to underscore the importance of considering the role of luck in investment performance. Just as the last person standing in the coin-tossing game is not necessarily a skilled coin flipper, a mutual fund manager with a streak of good returns is not necessarily a skilled investor. This understanding is crucial for investors when evaluating fund performance and making informed investment decisions.

Survivorship bias:

Survivorship bias refers to the logical error of focusing on the funds that have survived over a given period while ignoring those that have failed or been closed. This bias can lead to an overly optimistic view of mutual fund performance because it excludes the poor-performing funds that no longer exist.

As an example, consider a study that aims to evaluate the performance of mutual funds over the past decade. If the study only includes funds that are still active at the end of the period, it will naturally exclude those that were closed due to poor performance. This exclusion skews the results, making the surviving funds appear more successful than they might actually be. The performance data of these surviving funds will show higher average returns because the underperforming funds, which would have dragged down the average, are not included in the analysis. For example, Elton, Gruber, and Blake conducted a study on U.S. mutual funds and found that survivorship bias significantly affected performance results. They estimated that survivorship bias could inflate the average returns of mutual funds by approximately 0.9% per year. This means that the reported performance of mutual funds is often higher than the actual performance when considering all funds, including those that failed. Survivorship bias occurs because mutual funds that perform poorly are more likely to be closed or merged with other funds. When these funds are excluded from performance studies, the remaining funds—those that have survived—tend to show better performance. This creates a misleading picture of the overall success of mutual fund managers.

Investors might be led to believe that mutual fund managers are more skilled than they actually are, simply because the data does not account for the funds that did not survive. The impact of survivorship bias is not limited to mutual funds. It can also affect other areas of finance, such as hedge funds and market indices. For instance, studies on hedge fund performance often show inflated returns because poorly performing funds are closed, and their data is excluded from the analysis. Similarly, market indices that only include surviving companies can overstate the historical performance of the market. Understanding survivorship bias is crucial for investors and researchers. It highlights the importance of considering the entire population of funds, including those that have failed, to get an accurate picture of performance. By accounting for survivorship bias, you can avoid making overly optimistic conclusions about the skill of mutual fund managers and the performance of investment strategies. In conclusion, survivorship bias explains why the performance of mutual fund managers often appears better than it truly is. By excluding failed funds from performance analyses, the data becomes skewed, leading to an inflated perception of success.

So why do hedge fund managers make so much money every year? While hedge funds do employ sophisticated strategies aimed at exploiting potential market inefficiencies, their ability to earn outsized returns stems largely from the lucrative fee structures they charge investors. Managers typically receive both a management fee, often around 2% of assets under management, as well as a coveted performance fee of 20% or more of any profits generated. With billions under management, even small percentages translate into staggering sums for the most successful funds.

Moreover, the flexibility to use leverage, take short positions, and invest across a wide range of assets allows hedge funds to potentially profit in any market environment - both bull and bear markets. This

versatility provides an advantage over traditional long-only investment vehicles.

However, despite their cutting-edge approaches, hedge funds as a group have historically struggled to consistently outperform simple index funds over the long run after accounting for their high fees. Investors like Warren Buffett have publicly challenged whether more than a handful of elite managers can truly outperform markets sustainably.

So while the existence of billionaire hedge fund managers may seem contradictory to market efficiency on the surface, their ability to generate alpha is heavily reliant on lucrative fee structures rather than market inefficiencies alone. The long-term histories of consistent outperformance remain an open debate in the investment community.

Overall, it seems that the market is pretty efficient.

© Professor Raghavendra Rau, 2024

References and Further Reading

Busse, Jeffrey A., and T. Clifton Green, 2002, Market efficiency in real time, *Journal of Financial Economics* 65, 415-437.

Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57-82.

Elton, Edwin J., Martin J. Gruber, and Christopher R. Blake, 1996, Survivorship bias and mutual fund performance, *Review of Financial Studies* 9, 1097-1120.

Fama, Eugene F., 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25, 383-417.

Rau, Raghavendra, "A short introduction to corporate finance", 2017, Cambridge University Press (Chapter 7).

© Professor Raghavendra Rau, 2024