

Are Financial Markets Really Random?

by Aryan Varshney

For decades, even the greatest mathematicians and the most influential economists have been unable to predict financial markets. At the surface, movements in prices may appear chaotic and unpredictable, however, this raises a very important question: are markets truly random or do they simply appear that way?

Consider this: if markets were perfectly predictable, stock traders would be able to exploit any patterns for guaranteed profit. Guaranteed profits would forever change the outlook of financial markets and would enable several traders to be made redundant as their 'timely gambling' is no longer required. This would collapse the labour market and perhaps even the whole financial system. Yet such opportunities rarely persist in the real world, suggesting that something deeper is at the root of this.

Thankfully, the diverse and powerful subject of mathematics provides us a framework to model this uncertainty and test whether randomness truly governs markets. This essay, therefore, explores whether financial markets actually follow a random process, and where this assumption may break down.

The mathematical definition of "randomness" is not disorder, but rather unpredictability given available information. A random process is one where past outcomes cannot be used to predict future ones, and this closely resembles the concept of a "fair game" in probability theory where the expected gain from any strategy presented is zero. This reflects financial markets in a similar way where predictable profits are quickly eliminated. Does this suggest that randomness may be influenced by competition instead of purely being chance?

One simplified mathematical model of financial markets is the random walk. It assumes that price changes of 'x', let's say an asset, occurs as a sequence of random steps, forming this random walk. This can be expressed mathematically, where S_t represents the price of the asset at time t , S_{t+1} represents the price after a period $t+1$, and ϵ_t represents unpredictable news or information entering the market:

$$S_{t+1} = S_t + \epsilon_t$$

Important to note, the shocks represented by epsilon are assumed to be independent of each other, meaning past price movements or changes provide us with no information about future price changes. This therefore makes trying to predict markets using historical data impossible. A common analogy that may be referred to when describing the independence of these shocks and inability to predict, is the unpredictable nature of predicting the next outcome of a coin toss, and the ineffectiveness of past tosses to us if attempting to predict our next toss.

However, like any other model, the random walk model has its limitations too. For example, it works with time being represented as a discrete variable, however, in the real world, markets evolve

continuously. To capture this, mathematicians often refer to Brownian motion to model the dynamics of prices. This can be represented by a less simple formula:

$$dS_t = \mu S_t dt + \sigma S_t dW_t$$

This time, this formula uses rates of change (dt) to capture price changes over infinitesimally small time intervals. μ represents the average growth rate, σ measures volatility (capturing some magnitude of randomness), and the component dW_t represents a random fluctuation caused by Brownian motion. By combining the elements above, the formula provides us with a model where prices follow a structured path (though it is quite noisy), and this underpins most of quantitative finance in the modern world.

A crucial assumption in many financial models is that returns will follow a normal distribution. Using our earlier variables:

$$\epsilon_t \sim N(0, \sigma^2)$$

Though this has been proven by mathematicians where, considering the Law of Large Numbers, the outcomes (or returns) do in fact tend towards the bell-shaped curve reflected by the normal distribution, we soon find out that this underlying assumption may not actually hold true in real markets. If we accept the assumption for now, it suggests that gains and losses are equally likely due to the symmetric nature of the distribution. Moreover, price changes are small and extreme movements are very rare. In theory, this assumption simplifies various calculations and also makes models more tractable, but as I suggested, it does introduce significant limitations.

Real financial returns do not follow a perfect normal distribution and instead exhibit something called "fat tails". This is where extreme events occur more frequently than expected. This therefore suggests that real, large market crashes are more common than models predict. This can therefore be used to challenge our assumption of randomness based on normal distributions. Alternatively, what we can conclude is that markets have such underlying structures which cannot be captured by simplified models of real-world markets.

Another way in which randomness is deviated from is a phenomenon called volatility clustering. In financial markets, periods of high volatility actually tend to follow each other, and similarly, calm periods are often grouped together too. So perhaps we can accept random returns are true, but we can say with certainty that the variability is not random. Furthermore, this attacks our earlier assumption that price changes are completely independent of each other.

Some of these short-term patterns include momentum effects and mean reversion. The existence of these effects is often very short-lived and small. Traders can quickly exploit and eliminate such opportunities should they arise. This gives rise to a new concept known as the Efficient Market Hypothesis. According to this theory, prices already reflect all available information, and so only new and unpredictable information (shocks) can move the market, re-establishing the fundamental concept of randomness.

From another perspective, financial markets may only deviate from true randomness due to the role of human behaviour. An underlying assumption is that investors act rationally. In reality, however,

human behaviour introduces systematic biases. Some common examples include herd behaviour, overconfidence, and emotional decision-making where consumers (or investors) allow emotions and “gut feelings” to dictate their choices. These factors can lead to the economic concepts of bubbles and crashes (simply where the perceived value of something is well above its fundamental true value). As a result of human behaviour, markets deviate from purely mathematical predictions, highlighting the further limitations of treating markets as entirely random systems.

This information now draws us to a difficult position. Evidence suggests that markets are neither perfectly random, but nor are they fully predictable. Mathematically, our models capture important aspects of market behaviour, however in the real world, complexities introduce deviations from these models, enable markets to exhibit a form of structure randomness. They are unpredictable enough to be able to resist consistent profit strategies, but they are also structured enough to be analysed using the beautiful subject of mathematics.

Therefore, despite the answer to our posed question remaining blurred between uniformity and true unpredictability, what we can conclude is that mathematics acts as the spyglass for us to even be capable of attempting to answer this gruelling question.