Anveshanā • Article • July 2026

How mathematicians changed the world of finance, Forever.

by Sruthi Subramanian

Mathematics and Finance have traditionally been perceived as disjoint fields. At least in India, many students choose commerce because they do not enjoy mathematics, while those who do are often steered towards science by default. Commerce and finance are taught in a way that is largely disconnected from mathematics, at least at the school level. Formulas are used, but they are not proved, and the underlying mathematical ideas are rarely emphasized. At the same time, so-called pure mathematicians have often preferred to separate their work from practical, real-world applications and have not necessarily wanted to be associated with finance. It is therefore no surprise that when mathematicians first began walking into the world of finance, they were met with more suspicion than enthusiasm.

Many traditional investors dismissed the notion that mathematical models could rival time-tested fundamental analysis, which relied on company earnings, dividends, and financial statements to determine stock prices. Meanwhile, many academics viewed the shift of talent from pure mathematics to finance as a waste of intellectual talent. None of this deterred Jim Simons (1938 - 2024), an MIT graduate and UC Berkeley PhD, who had his eye on markets from a young age. From scratch, he built the Medallion Fund and Renaissance Technologies, producing extraordinary annual returns of 66% before fees, far surpassing even Warren Buffett. He achieved this by bringing mathematics to the heart of trading. This is the story of how Simons and a small group of mathematicians sparked the Quant Revolution.

Mathematics was always central to Jim Simons' identity. From an early age, he knew he wanted to pursue the subject. His PhD research focused on the geometry of multi-dimensional curved spaces, and later, while working with Shiing-Shen Chern, he co-developed the Chern-Simons invariants, concepts that would eventually find important applications in physics [1]. After completing his doctorate, Simons taught at MIT and conducted research at Harvard before joining the Institute for Defense Analyses (IDA) as a codebreaker, continuing his own

Sruthi Subramanian is MFin candidate at Princeton's Bendheim Center for Finance and an outgoing undergraduate from IIT Kanpur in Mathematics and Scientific Computing. She loves teaching and writing, both technical and creative, and cares deeply about making ideas accessible. Outside the classroom, she is a classical Bharatanatyam dancer and often hosts keyboard quizzes at family events. Her corresponding address is sruthisubramanian17@gmail.com.

[1]For instance, it later provided the grounds for the understanding of Topological Quantum Field Theory.

mathematical research alongside government work. The open and collaborative culture at the IDA left a lasting impression on him and later became a defining feature of Renaissance Technologies. After leaving the IDA, Simons helped build the mathematics department at Stony Brook University before eventually turning his attention fully toward finance, where he founded the Medallion Fund and Renaissance Technologies (RT).

Jim Simons at his New York office in 2007. Bill Cramer

Jim Simons always envisioned an automatic trading system, one capable of generating profits with minimal human intervention. Early on, however, the process still relied on a combination of computers and people. Medallion's first trading system, known as the 'Piggy Basket', was designed for currency markets and produced vectors indicating the proportions of various currencies to buy. Human traders then manually executed trades to match those recommendations. Over time, Simons, together with prominent figures like Henry Laufer, Peter Brown, and Robert Mercer, developed increasingly sophisticated mathematical strategies. Their work eventually led to the creation of fully automated systems driven entirely by mathematical models, bringing Simons' vision of a "money-making machine" to life.

Many of the mathematicians Simons encountered during his years at UC Berkeley, the IDA, and Stony Brook later became instrumental in building Medallion, with some even joining him as partners. James Ax, a mathematician famous for the Ax-Kochen theorem, whom Simons had persuaded to join Stony Brook, introduced trading strategies based on moving averages. Sandor Straus, a UC Berkeley PhD who worked at Stony Brook's Computer Center, focused on building clean and reliable datasets, an effort that later proved critical to the success of Henry Laufer's automated systems. Rene Carmona, then a professor at Princeton, joined the firm and experimented with kernel methods, early machine learning techniques, to identify patterns in financial data. Simons was initially cautious with these methods due to their lack of explainability. Another important figure was Elwyn Berlekamp, who had crossed paths with Simons at MIT, Berkeley, and the IDA. Berlekamp had worked with Claude Shannon (yes, Claude AI is named after this person!), the father of information theory, and himself made major contributions to combinatorial game theory. At Renaissance, he encouraged the firm to move toward short-term trading and argued that success in markets, much like success in casinos, required being correct only slightly more than half the time. Even a small statistical edge, applied consistently, could generate enormous profits.

The lack of explainability in methods like kernel models was one of the major obstacles to integrating computers into finance. Many quantitative techniques, both then and now, operate as black boxes: the system recommends trades, but no one can fully explain why those recommendations work. This naturally made investors hesitant to trust such models. Ironically, this became both an obstacle and an advantage for Simons. In the early years, potential investors were reluctant to fund Renaissance Technologies because they could not clearly understand the logic behind its strategies. Yet, when the firm eventually built its trading systems, it deliberately combined signals that were intuitively understandable with others that had no obvious explanation. Competitors tended to avoid strategies they could not rationalize, which gave Renaissance an edge. By embracing patterns that others distrusted, the firm created systems that were both highly effective and difficult for rivals to predict or replicate.

The Quant Revolution unfolded alongside rapid advances in technology and the growing influence of computer science. Improvements in computing power and data processing enabled the analysis of the enormous datasets that Sandor Straus had painstakingly assembled. Several figures connected to quantitative finance were also deeply involved in the development of modern computing and machine learning. Jeff Bezos, before founding Amazon, worked at D. E. Shaw Co., one of Renaissance's rivals, which used statistical arbitrage models for trading. Leonard Baum, Simons' first partner, together with Lloyd Welch, developed the Baum-Welch algorithm, a method applied not only in finance but also in speech recognition. Similarly, the IBM research team from which many people eventually joined RT, starting with Peter Brown and Robert Mercer, worked on speech-to-text systems and machine translation, early forms of what is now known as natural language processing. Straus' insistence on clean, structured data and the use of historical information to predict future behavior also anticipated techniques that are now central across industries, from online recommendation systems and targeted advertising to intelligence analysis. Many of the mathematical models used during the Quant Revolution relied on identifying patterns and probabilities in seemingly unpredictable systems, much as in astrophysics, which Simons also had an affinity for. One important framework was the Hidden Markov Model, which is used in conjunction with the Baum-Welch algorithm. The model below (left) denotes a Markov Model, with the transition matrix P, whose element p_{ij} gives the probability of moving from the i^{th} state to the j^{th} state at any point in time. This matrix is fixed, consistent with the assumption that future outcomes depend only on the current state rather than the entire past. The $(t+1)^{th}$ state depends only on the $(t)^{th}$ one.

Hidden Markov Models extend this idea by assuming that some underlying factors affecting the system remain unknown, forcing the model to make probabilistic inferences from partial information. As shown in the diagram above (right), we assume a Markov model is running in the background. We then condition on each hidden state to find the probability of transitioning to an observable state.

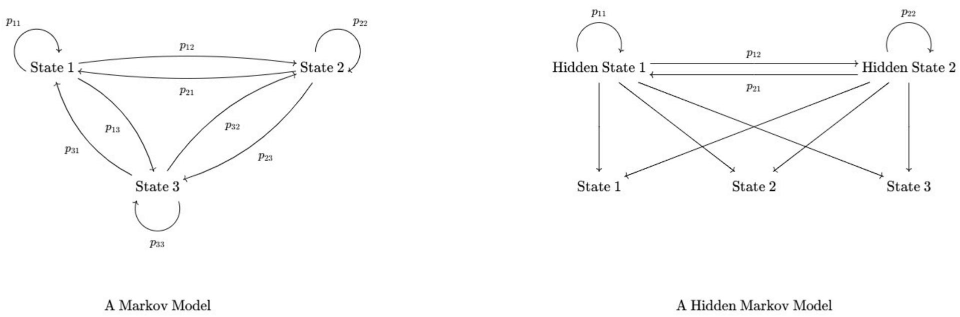

Fig: Markov Model vs Hidden Markov Model COURTESY SRUTHI SUBRAMANIAN

Bayesian inference was used extensively to update probabilities. In Bayesian inference, we begin with a prior distribution that represents our initial beliefs about the probabilities of different outcomes. As new data becomes available, we condition on this information and apply Bayes' theorem to obtain a posterior distribution, which then serves as the new prior for future updates. This sequential updating of beliefs allows the probabilities of different events to evolve over time, making the model more adaptive and reliable. By incorporating the most recent information, Bayesian methods embody the principle that additional data can improve our understanding and lead to better predictions.

> Bayesian inference updates our beliefs about an unknown parameter θ based on observed data D. The relationship is given by Bayes' theorem:
>
> $$p(\theta \mid D) = \frac{p(D \mid \theta)\, p(\theta)}{p(D)}$$
>
> - $p(\theta \mid D)$: **Posterior probability** of θ given the observed data D.
> - $p(D \mid \theta)$: **Likelihood**, the probability of observing D assuming θ.
> - $p(\theta)$: **Prior probability** of θ before observing the data.

At Renaissance, Henry Laufer used the vast datasets assembled by Sandor Straus to search for recurring patterns in historical market data. His models identified patterns tied to specific times of day, days of the week, or recurring reactions to particular events. His strategies were tested against historical records to assess statistical significance before deployment. Often, the exact reason these patterns worked remained unclear, but if the historical evidence supported them strongly enough, then the strategies were implemented regardless. Another widely used approach was the *return-to-mean (or mean reversion) strategy*, especially in *pairs trading*, which involves two historically correlated securities. If the price relationship between the two assets deviated significantly from its historical average, traders assumed that it would eventually revert to its long-run mean. They then bought the relatively undervalued asset **(Long buy)** and simultaneously sold the relatively overvalued asset **(Short sell)**, creating a market-neutral position that minimized exposure to broader market movements. When the price relationship returned to its historical level, the **trades were reversed** — the long position was sold, and the short position was covered, allowing the trader to lock in a profit approximately equal to the magnitude of the temporary deviation. Such lower-risk strategies were central to Medallion's success and contributed to its exceptionally high Sharpe ratio (the ratio of expected return to risk), allowing it to achieve unusually high returns while taking comparatively lower risk.

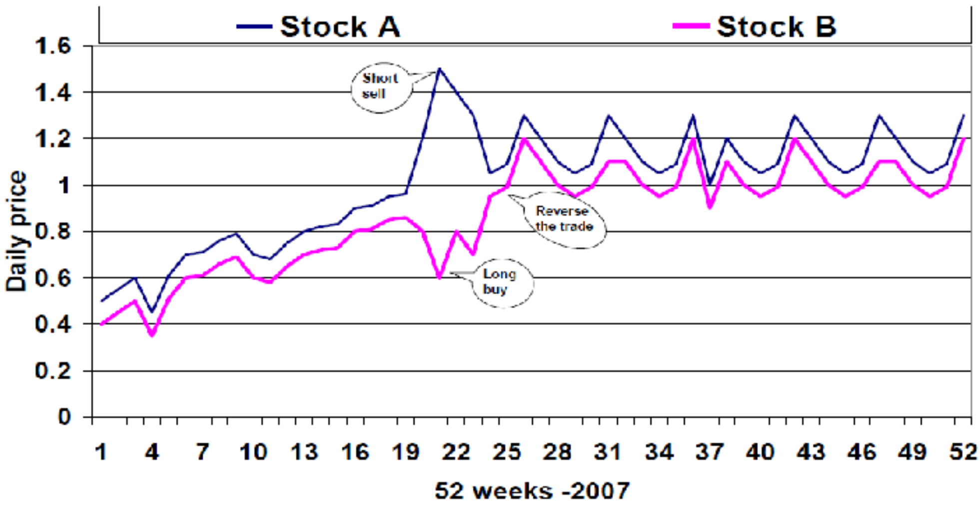

Fig: Pairs Trading Strategy Reproduced from [1] with permission.

Beginning in the 1950s, academic research witnessed a rapid growth in mathematical models designed to explain and predict financial market behaviour. Their significance is underscored by the fact that several of these contributions were later recognised with Nobel Prizes in Economics. Harry Markowitz's *Modern Portfolio Theory* introduced the concept of mean-variance optimization to determine the optimal allocation of wealth across multiple assets, balancing expected returns against risk. Building on these ideas, the *Capital Asset Pricing Model* (CAPM) of the 1960s used linear regression to relate an asset's returns to market returns, distinguishing between systematic and unsystematic risk and reinforcing the importance of diversification in reducing asset-specific risk. The 1970s then saw the emergence of the *Black-*

[1] Al-Naymat, Ghazi. "Mining pairs-trading patterns: A framework." International Journal of Database Theory and Application 6.6 (2013): 19-28.

Scholes model for option pricing, which drew upon stochastic calculus and modelled asset prices using geometric Brownian motion to derive the price of options (financial contracts that give the owner the right to buy/sell an asset at a fixed price, on a fixed future date). These developments laid the foundations of modern quantitative finance and continue to influence both academic research and industry practice.

Modern Portfolio Theory

$$\boxed{\mathbb{E}[R_p] = \sum_{i=1}^{n} w_i \cdot \mathbb{E}[R_i] = \mathbb{E}[\mathbf{w}^T R_i]}$$

- $\mathbb{E}[R_p]$: portfolio expected return
- w_i: weight of asset i
- $\mathbb{E}[R_i]$: expected return of asset i
- n: number of assets

$$\boxed{\sigma_p^2 = \sum_{i=1}^{n}\sum_{j=1}^{n} w_i w_j \sigma_i \sigma_j \rho_{i,j} = \mathbf{w}^T \Sigma \mathbf{w}}$$

- σ_p^2: portfolio variance (risk)
- σ_i: volatility of asset i
- ρ_{ij}: correlation between assets i and j
- Σ: covariance matrix of returns

$$\boxed{\text{Sharpe Ratio} = \frac{\mathbb{E}[R_p] - R_f}{\sigma_p}}$$

- R_f: risk-free rate
- The numerator is the excess return over the risk-free, and the denominator is the portfolio volatility

We can find the optimal portfolio by minimizing variance, or by maximizing the Sharpe Ratio, or by optimizing any measure relating risk and return in this setup.

Optimization Problem (maximizing Sharpe Ratio)

$$\boxed{\max_{\mathbf{w}} \frac{\mathbb{E}[R_p] - R_f}{\sigma_p} = \max_{\mathbf{w}} \frac{\mathbf{w}^\top \boldsymbol{\mu} - R_f}{\sqrt{\mathbf{w}^\top \Sigma \mathbf{w}}}}$$

- $\mathbf{w}$: vector of portfolio weights
- $\boldsymbol{\mu}$: vector of expected returns

Capital Asset Pricing Model

$$\boxed{\mathbb{E}[R_i] = R_f + \beta_i \cdot (\mathbb{E}[R_m] - R_f)}$$

- $\mathbb{E}[R_i]$: expected return on asset i
- R_f: risk-free rate
- β_i: beta of asset i
- $\mathbb{E}[R_m]$: expected market return

$$\boxed{\beta_i = \frac{\text{Cov}(R_i, R_m)}{\text{Var}(R_m)} = \frac{\sigma_{i,m}}{\sigma_m^2}}$$

- $\text{Cov}(R_i, R_m)$: asset and market returns covariance
- $\text{Var}(R_m)$: market variance

Black-Scholes Option Pricing

A *call option* is a financial contract that gives the holder the right (but not the obligation) to buy the underlying asset, usually a stock, at a fixed price, called the strike, on a fixed date, the expiry date of the options contract. A *put option* gives the holder the right to sell the underlying asset at the strike price on the expiry date.

$$\boxed{C = S_0 \cdot N(d_1) - K \cdot e^{-rT} \cdot N(d_2)} \quad \boxed{P = K \cdot e^{-rT} \cdot N(-d_2) - S_0 \cdot N(-d_1)}$$

C: call option price; P: put option price; S_0: current stock price; K: strike price; r: risk-free rate (continuous); T: time to expiration (years); σ: volatility; $N(\cdot)$: standard normal CDF, where

$$d_1 = \frac{\ln\left(\frac{S_0}{K}\right) + \left(r + \frac{\sigma^2}{2}\right) T}{\sigma\sqrt{T}}, \; d_2 = d_1 - \sigma\sqrt{T}$$

The partial derivatives of these options are taken to find out how the prices change with respect to different factors, like time or the current stock price. These derivatives are popularly known as the Greeks.

The rise of quantitative finance also coincided with the emergence of Behavioral Economics, which emphasized the importance of human psychology in shaping market behavior. Emotions, biases, and collective reactions all influence how prices move, and many argued that quantitative models were, in essence, sophisticated tools for modeling human behavior itself. By analyzing how people had reacted to past events, these systems attempted to predict how markets would respond to similar situations in the future. As long as human behavior remained relatively consistent, past data could serve as a useful guide. However, as quantitative trading becomes increasingly dominant and human-driven trading is declining, some question whether markets themselves are changing. This shift to more automated trading could make markets more stable and reduce volatility, but it also raises doubts about whether the future will continue to resemble the past. Interestingly enough, though, even the strongest advocates of quantitative methods recognized the limits of automation. When models underperformed, Simons himself occasionally intervened and adjusted the systems. Ultimately, the financial world operates through a balance between human judgment and automated systems, and maintaining that equilibrium remains crucial.

The Quant Revolution is far from over. While the first generation of quants showed that mathematics and statistics could uncover hidden patterns in markets, the current generation is exploring whether artificial intelligence can take things a step further by learning those patterns on its own. Hedge funds today increasingly rely on machine learning, natural language processing, and use data ranging from weather parameters to social media sentiment, blurring the boundaries between mathematics, finance, computer science, and data science. As algorithms continue to reshape financial markets, the same question remains as relevant as ever: Can mathematics ever fully capture the complexity of financial systems? The search for this answer continues to define the frontier of quantitative finance.