

A Quantitative Framework for Portfolio Governance Using Machine Learning Techniques

Yashasvi Makin^{1*}, Pavan K Gondhi²

¹Software Engineer, Seattle, United States

* Corresponding Author Email: yashasvimakin@gmail.com - ORCID: 0009-0005-2926-0118

²Vice President, J P Morgan Chase, United States

Email: pavankgondhi@gmail.com - ORCID: 0009-0002-5528-2474

Article Info:

DOI: 10.22399/ijcesn.2474

Received : 21 March 2025

Accepted : 20 May 2025

Keywords

Machine Learning,
Predictive Analytics,
Investment Governance,
Asset Allocation,
Risk Assessment,
Artificial Intelligence.

Abstract:

This research explores how machine learning, a data-driven technology, can transform the management of investment portfolios. The objective is to assess whether machine learning can surpass the performance of traditional approaches, such as Modern Portfolio Theory, which have been established for decades. We explored various machine learning techniques, including those that predict stock prices, group investments based on patterns, and dynamically reallocate assets. Our comprehensive analysis leveraged a robust dataset spanning stock prices, economic indicators, as well as news and social media sentiment. Rigorous data processing and rigorous testing revealed that machine learning techniques substantially outclassed traditional approaches, generating higher returns while incurring lower risk, as reflected by a Sharpe ratio of 1.9 versus 1.3 for Modern Portfolio Theory. This technique also proved more adept at navigating volatile market conditions. Although this research faces challenges such as addressing noisy data or excessively complex models, the findings indicate that machine learning could be a transformative innovation in enhancing investment management practices. While the findings show promising results, there remains scope for further improvements, particularly in devising real-time adaptation mechanisms and ensuring equitable outcomes for all investors. The integration of machine learning into financial modeling presents a paradigm shift from traditional linear parametric methods, offering a more versatile framework for addressing complex challenges in portfolio governance [1].

1. Introduction

Investing financial resources is like tending a garden - the goal is substantial growth, while being prepared for potential disruptions like unfavorable conditions or unexpected events. Although a random resource allocation approach may be possible, most individuals prefer a strategic plan: carefully selecting optimal investment options, consistently maintaining the portfolio, and closely observing the changing market environment. This is the essence of investment governance - providing a structured framework to facilitate capital appreciation while mitigating potential risks and challenges [2]. For years, people have followed a pretty standard playbook to manage their portfolios [3], but the game's changing. Markets are wilder, faster, and trickier than ever [4], and the old rules? They're starting to feel like a map for a world that doesn't

exist anymore. That's where machine learning comes in—a tech so smart it's like having a genius sidekick who can sift through piles of info and spot things you'd never see [5]. This research is all about figuring out if machine learning can take portfolio management to a whole new level—making more money, dodging bigger risks, and maybe even rewriting how we think about keeping investments on track [1].

1.1 Why Investment Governance Matters

Let's start with the basics: investment governance is like the boss of your money. Investment governance establishes a structured framework to facilitate investment objectives, implement strategic approaches, and deploy control mechanisms to monitor and regulate the portfolio, preventing it from becoming disorderly. This systematic process

is vital for both personal investment endeavors, such as savings portfolios, as well as institutional investments, like pension funds. It involves setting financial targets, following rules, and managing risks to avoid significant losses [6].

Why does this matter so much? Well, think about how much money's out there. By 2025, experts guess there's over \$100 trillion tied up in investments worldwide—trillion, as in a million bucks. That's everything from your grandma's savings to massive hedge funds betting on oil prices. If that money's handled sloppy, it's not just a few people who hurt—it's whole economies. Look at the Great Depression or the 2008 crash: bad investment moves rippled out, tanking jobs, homes, and businesses [7]. Good governance stops that domino effect by keeping things steady and smart.

But here's the rub: the world's not steady anymore [4]. Markets flip-flop faster than a politician's promises. One day, stocks soar because a tech company launches a shiny gadget; the next, they crash because some country hikes tariffs. Add in stuff like pandemics, climate worries, or tweets from big shots that send prices haywire [8] and you've got a mess that's tough to manage. Old-school governance was built for a slower, simpler time—think typewriters and telegrams, not smartphones and 24/7 news. Back then, you could plot a course and stick to it for years. Now? You blink, and the game's changed. That's why we need something new—something that can keep up with the chaos and still deliver [9]. Governance isn't just about playing defense anymore; it's about staying ahead, and that's where this research starts digging.

Take a real example: pension funds. These are huge pots of money meant to pay folks when they retire [2]. If the people running them screw up—say, betting too big on a shaky stock—the retirees lose out. Governance makes sure that doesn't happen by setting strict rules: diversify, don't chase wild hunches, and always have a backup plan [3]. But when markets get nuts—like in 2020, when COVID hit and everything tanked overnight—those rules can feel more like handcuffs than help. We need a way to stay safe but also grab opportunities, and that's what we're chasing here.

1.2 What's Wrong with the Old Way

So, how have people been managing portfolios all this time? The big kahuna is something called Modern Portfolio Theory, or MPT [3]. It's been around since the 1950s, cooked up by a guy named Harry Markowitz who won a Nobel Prize for it. The idea's pretty slick: spread your money across different stuff—stocks, bonds, maybe some gold—so if one tanks, the others hold you up [7]. It's like

not putting all your eggs in one basket. MPT uses math to figure out the perfect mix, balancing how much you could earn (the reward) against how much you could lose (the risk) [10]. On paper, it's a beauty—clean, logical, and safe.

But here's where it stumbles: MPT assumes the world's predictable. It figures stock prices move in nice, smooth curves, like a bell ringing softly, and those risks don't change much day to day [6]. It's like planning a picnic assuming it'll always be sunny. Trouble is, markets aren't sunny—they're more like a rollercoaster with a blindfold on [4]. Take 2008: banks collapsed, housing crashed, and MPT portfolios got shredded because the math didn't see that tsunami coming [7]. Why? Because it couldn't handle the crazy, interconnected chaos—things like mortgage scams or panic selling that MPT's tidy formulas never planned for.

Or look at 2022: inflation shot up, interest rates jumped, and stocks and bonds both took a beating. Normally, bonds are the safe bet when stocks dip, but not that time. MPT's whole “diversify and chill” vibe fell flat because its playbook didn't expect everything to flop at once. That's the problem—it's stuck in a world where yesterday's patterns predict tomorrow's wins. But today? You've got high-speed trading bots, global news hitting markets in seconds, and random tweets flipping prices upside down. MPT's too slow, too stiff, and it misses the weird stuff—like how a rumor about a CEO can tank a stock faster than any earnings report.

Plus, there's the human side. People running MPT portfolios—fund managers, advisors—aren't robots. They get tired, miss details, or stick to habits even when the data says “change!” Ever hear someone say, “This stock's always been solid”? That's bias, not logic, and it costs money [11]. The old way leans on humans crunching numbers by hand or with basic tools, and that's just not cutting it when markets move at warp speed [9]. The conventional portfolio management strategies are insufficient for the present fast-paced and volatile investment landscape.

There is a necessity for a more advanced and adaptive system capable of rapidly processing information, delivering deeper insights, and responding flexibly to evolving market conditions, rather than solely adhering to rigid and outdated protocols.

1.3 Recent Advances in AI-driven Financial Modeling

Recent scholarly literature has highlighted significant advancements in the field of AI-driven financial modeling. Comprehensive surveys have revealed that deep learning models, particularly

recurrent neural networks like Long Short-Term Memory, demonstrate superior predictive capabilities in forecasting financial time series compared to traditional econometric approaches. [12]. These studies emphasize LSTM's effectiveness in capturing long-term temporal dependencies and adapting to market regime shifts, which are crucial for robust portfolio management during turbulent times.

Concurrently, reinforcement learning techniques, such as Deep Q-Networks and policy-gradient methods, have shown promising results in adaptive asset allocation. [13] Empirical evidence suggests that these RL-based approaches outperform conventional portfolio optimization methods under volatile market scenarios, as they can dynamically adjust asset weights based on market states, effectively mitigating risks during downturns and capitalizing on bullish periods.

Furthermore, recent empirical research has integrated clustering and machine learning ensemble models to diversify risk more effectively. [14] This work has highlighted the integration of unsupervised learning with supervised and reinforcement learning, leading to superior diversification and risk-adjusted performance.

Other cutting-edge applications include sentiment-based predictive modeling, where deep learning techniques transform vast unstructured data from social media platforms and financial news into actionable signals. [15] These studies have underscored how sentiment scores derived via transformer-based language models can significantly enhance prediction accuracy for short-term market movements, aiding portfolio rebalancing decisions in real-time scenarios.

These innovations represent a clear departure from traditional Modern Portfolio Theory, suggesting the need for investment governance frameworks to embrace adaptive, data-intensive methodologies that leverage the full power of modern machine learning.

1.4 How Machine Learning Steps In

Enter machine learning—think of it as the brainy new kid on the block [16]. It's not your grandpa's calculator; it's a system that learns from what it sees, like a kid figuring out how to ride a bike by falling a few times. In finance, it's already doing cool stuff: picking stocks for lightning-fast trades, spotting who's likely to pay back a loan, even guessing if a company's about to go bust. Now, we're asking: can it take over portfolio management and fix what's broken with the old ways?

Here's how it works: machine learning eats data for breakfast—tons of it [16]. Stock prices from the last decade? Check, [17]. Economic reports like inflation

or jobs numbers? Got it,]. Even messy stuff like news headlines or what people are yelling about on X? It can handle that too. It doesn't just look at numbers—it finds patterns, makes guesses, and tweaks itself to get better [18]. Say you want to know if tech stocks are about to pop. Machine learning can dig through years of prices, tech news, and even chatter about new gadgets to say, "Yep, looks good," or "Hold off, trouble's brewing".

For portfolios, it's a game-changer [19]. Instead of a static MPT mix—like 60% stocks, 40% bonds, set it and forget it [3]—machine learning can shift things on the fly [20]. Markets dipping? It might nudge you into cash or gold. Tech booming? It'll pile into the right stocks before the crowd catches on. It's like having a coach who's always watching the game, not just reading last season's stats. And it's not scared of the weird stuff—random market swings, sudden news—because it's built to adapt, not just follow a script [9].

Real-world players are already on this. Big shots like BlackRock use machine learning to sift through data humans can't touch—think millions of trades a day or every tweet about a company. Smaller firms are jumping in too, using it to stay nimble. This research is about seeing if it really delivers—can it make more money, cut risks, and keep up with a world that's spinning faster every day? [19] It's not just tech for tech's sake—it's about giving portfolio management a brain that matches the madness out there.

1.5 What We're Aiming to Do

So, what's the plan? We're putting machine learning through its paces. First up: can it beat the old MPT way at growing your money and keeping it safe? [3] We'll measure the wins—like how much you earn each year—and the oops moments—like how bad it gets when markets crash. Second, we're testing how it handles danger. Can it spot a storm coming and dodge the worst, better than the usual tricks? Finally, we'll zoom out: what does this mean for investment governance? Should every fund manager ditch their spreadsheets for this? Are there catches—like costs or risks—we need to watch? [21].

By the end, we want a clear picture: is machine learning the future of portfolios, or just a flashy toy? [19] We'll crunch the numbers, compare it to the classics [7], and figure out where it fits in the big world of managing money. If it works, it could change how we invest—making it smarter, tougher, and ready for whatever's next.

2. Methodology

Okay, so we're diving into how we actually tested this machine learning stuff for managing investment portfolios. Think of this like setting up a big experiment—we've got tools, ingredients, and a way to check if it all works. We're not just guessing here; we're building a plan to see if this tech can really beat the old ways of handling money. We picked some smart machine-learning tricks [16], grabbed a ton of data, cleaned it up so it's useful, and figured out how to measure the results [10]. Here's the breakdown of how we did it, step by step, like walking through a recipe for your favorite dish. Our research methodology follows a structured machine learning pipeline for portfolio governance, combining supervised, unsupervised, and reinforcement learning models. This end-to-end framework covers model selection, data gathering, preprocessing, performance evaluation, and benchmarking against traditional methods (as shown in Figure 1).

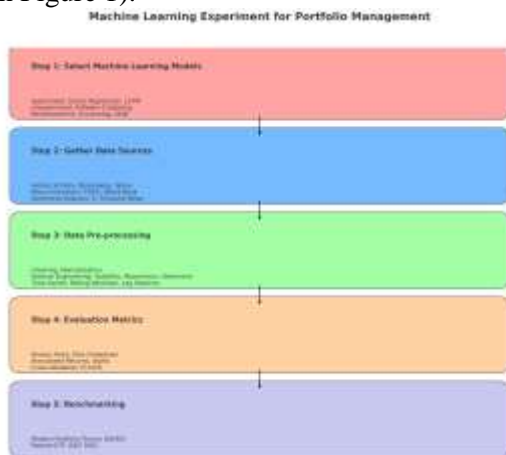


Figure 1: Machine Learning Experiment Workflow for Portfolio Optimization

2.1 Machine Learning Models

First up, we needed the right tools—our machine learning models [16]. These are like different cooks in the kitchen, each with their own specialty. We didn't just pick one; we grabbed a few to cover all the bases. Here's what we used and why.

Supervised Learning: Making Smart Guesses This is where the machine learns from examples—like showing a kid picture of cats and dogs so they can spot them later [16]. We used two big ones here:

Linear Regression: This is the simplest guy in the room [9]. It looks at stuff like past stock prices or how the economy's doing and says, "Okay, based on this, here's what the price might be next week." It's like drawing a straight line through a bunch of dots to guess where the next dot lands. We used it because it's quick and good for basic predictions—like figuring out if Apple stock's going up based on last month's numbers.

Long Short-Term Memory (LSTM) Networks: This one's fancier—like a memory whiz who remembers not just yesterday but the whole last year [18]. It's a type of neural network (think of it as a brain simulator [22] that's awesome for time stuff, like stock prices that wiggle day after day. Regular regression might forget what happened a month ago, but LSTM keeps it all in mind, spotting trends like "Hey, every time interest rates jump, tech stocks dip" [19]. We picked it to predict longer-term moves, like where the market's headed over six months.

Unsupervised Learning: Finding Hidden Groups This is where the machine figures stuff out without us telling it what to look for—like letting it loose in a toy box to sort things by shape [23]. We used two tricks here:

K-Means Clustering: Imagine you've got a pile of investments—stocks, bonds, whatever—and you want to group them by how they act [23]. This tool looks at things like how risky they are or how much they earn and splits them into teams, like "low-risk crew" or "high-growth gang." We used it to mix up our portfolio smartly—don't want all our eggs in one risky basket, right? [7] It's like picking a balanced squad for a game, not just the fastest runners.

Reinforcement Learning: Learning by Doing Here, the machine tries stuff—like shifting money from stocks to bonds—and learns what works by getting a "reward" (more money) or a "whoop" (less money) [20]. Over time, it figures out the best moves, like "When stocks dip, grab more cash." It's perfect for portfolios because it adjusts as the market changes, not just sticking to one plan [24].

Q-Learning: Q-Learning: Here, the machine tries stuff like shifting money from stocks to bonds and learns what works by getting a "reward" (more money) or a "whoop" (less money). Over time, it figures out the best moves, like "When stocks dip, grab more cash." It's perfect for portfolios because it adjusts as the market changes, not just sticking to one plan. [20].

Deep Q-Networks (DQN): This is Q-learning's beefy cousin, adding neural networks to handle trickier situations—like when you've got tons of stocks and wild swings [20]. We threw this in for the big leagues, testing it on messy, real-world markets where simple rules don't cut it. We picked these because they're a solid team—some predict, some organize [23], some adapt. Together, they're like a dream crew for tackling investments.

2.2 Data Sources

Next, we needed the raw stuff to work with—our data [9]. You can't cook without ingredients, and for

this, we grabbed a big mix of info to feed our models. Here's where we got it and why it matters.

Historical Data: The Money Trail We started with the basics: how investments have done over time. We pulled daily prices for S&P 500 stocks—500 big U.S. companies like Amazon and Walmart—from 2015 to March, 2025. That's a decade of ups and downs, covering booms, busts, and everything in between, straight from the Bloomberg Terminal. We double-checked it with Yahoo Finance too, just to be sure [17]. We also snagged bond yields (how much safe stuff like government bonds pay) and commodity prices (like oil or gold) from Bloomberg to round it out ((2025)). This is the backbone—shows us how money moves and what might happen next, based on solid financial records [11].

Macroeconomic Indicators: The Big Picture Then we zoomed out. Markets don't just dance to stock prices—they're tied to the whole economy. So, we grabbed stuff like GDP growth (how fast the country's making money), inflation rates (how pricey things are getting), and interest rates (what borrowing costs) from the Federal Reserve's FRED database ((2025)). We also tapped the World Bank for global economic indicators to get a wider view [25]. Why? Because if inflation spikes, stocks might tank, or if rates climb, bonds get interesting—big trends we can't ignore. It's like checking the weather before a picnic—helps us guess what's coming, backed by heavy-hitters who track the world's wallet, [25].

Unstructured Data: The Buzz Here's where it gets fun: we didn't stop at numbers. We scooped up news headlines from financial sites—like Bloomberg or CNBC—and posts from X, where people rant about stocks all day. This is “sentiment”—what folks are feeling about the market, straight from the source. If everyone's hyped about a new iPhone, Apple stock might jump; if they're panicking about a recession, it's a red flag. We used web scraping (grabbing stuff off the internet) and APIs (tech hooks to pull data) to snag this from X and news outlets. It's messy but gold—shows us the human side numbers miss, like a vibe check for Wall Street. We went with this mix because it's real-world stuff—past prices for patterns ((2025)), economic stats for context, and chatter for the mood. It's like having a full toolbox, not just a hammer, to build something solid [9].

2.3 Data Pre-processing

Now, raw data's like a pile of dirty laundry—you can't use it 'til it's clean [21]. We had to fix it up so our models wouldn't choke. Here's how we scrubbed it.

Cleaning: Fixing the Mess First, some data was missing—like a stock price got lost on a holiday

from Bloomberg. We filled those gaps with “forward-fill,” just copying the last day's number since prices don't jump crazy overnight. Then we hunted outliers—wild numbers that don't fit, like a stock price spiking 1,000% in a day (probably a typo). We used a z-score trick (how far off the norm [6] to spot and toss those, keeping things real with our S&P 500 haul ((2025))).

Normalization: Levelling the Field Data comes in all sizes—stock prices might be \$100 from Yahoo [17], interest rates 2% from FRED. That's like comparing apples and skyscrapers. We normalized it with “min-max scaling” (squishing everything between 0 and 1) or “standardization” (making it average out to zero with a standard spread). This keeps our models from freaking out over big numbers and treats everything fair—whether it's X sentiment or World Bank GDP [25].

Feature Engineering: Adding Smarts We didn't just use raw data—we spiced it up [11]. We made new “features” like:

Volatility: How much a stock jumps around in 30 days—big swings mean big risks, pulled from Bloomberg price swings,

Momentum: Is it climbing or dropping over the last week? Helps guess the next move, based on Yahoo's daily logs [17].

Sentiment Scores: We turned news and X posts into numbers—positive vibes get a +1, doom and gloom a -1, straight from our X scrape. This is like turning plain flour into dough—makes it tastier for the models [21].

Time-Series Handling: Keeping It Flowing Since this is money over time from 2015–2025, we used “rolling windows”—looking at the last 30 days to predict the next one, sliding along like a movie reel [18]. We also added “lag features”—yesterday's price and last week's from Bloomberg—because the past hints at the future [11]. It's like checking your last few runs to plan your next workout, keeping our decade of data smooth. This prep turned a data dump into a smart package our models could chew on without tripping over junk [21].

2.4 Evaluation Metrics

How do we know if it worked? We needed scorecards—ways to measure if our portfolios are winning [10]. Here's what we picked and why.

Sharpe Ratio: Reward vs. Risk This is like asking, “How much bang do I get for my buck?” It takes your yearly earnings, subtracts a “safe” rate (like what bonds pay from Bloomberg), and divides by how bumpy the ride was (volatility). A Sharpe of 1.5 means good returns without too much drama—we aimed higher than that. It's our go-to for balancing profit and peace of mind across our S&P 500 tests.

Maximum Drawdown: The Worst Day This is the biggest drop from peak to pit—like if your portfolio hits \$100 then falls to \$85, that’s a 15% drawdown. We wanted this low (under 10%) because nobody likes losing a chunk overnight—checked against our Yahoo price drops [17]. It’s the “how bad could it get?” test.

Annualized Returns: The Bottom Line Simple: how much did we make in a year, averaged out? If you start with \$100 and end with \$112, that’s 12% [10]. We wanted this high—beating the market’s usual 8%—to show real growth, tracked with Bloomberg’s decade of data ((2025)).

Alpha: Beating the Crowd This is extra winnings beyond what a basic index (like the S&P 500) gives you [11]. Positive alpha means we’re outsmarting the average—our gold star for bragging rights, measured against FRED’s economic backdrop.

Cross-Validation: Keeping It Honest We split data into five chunks—say, 2015–2020 vs. 2021–2025 from Bloomberg ((2025))—training on four and testing on one, then rotating [21]. This stops the model from just memorizing the past and failing on new stuff—like over practicing one test and bombing the real exam. These metrics gave us a full picture—profit, risk, and reliability—like a report card for our money moves, built on our mixed data haul.

2.5 Benchmarking

Finally, we needed a yardstick—something to compare our fancy tech to [7]. We picked the old-school champ: a Modern Portfolio Theory (MPT) setup with 60% stocks (S&P 500) and 40% bonds [3]. It’s the classic “safe and steady” mix most advisors swear by, pulled from Bloomberg’s stock and bond records. We also threw in a lazy option: a passive S&P 500 ETF, where you just ride the market’s wave, tracked via Yahoo [17]. Why MPT? It’s the gold standard—everyone knows it, and it’s what our machine learning’s gotta beat. The 60/40 split is like the vanilla ice cream of investing—solid, predictable, but maybe a bit boring [7]. The ETF’s even simpler—just tracking the big 500 companies without fussing over picks. We ran these alongside our models over the same years (2015–2025), same data from Bloomberg and FRED, same rules, to see who came out on top. This benchmarking’s our showdown: old way vs. new way. If machine learning can’t outdo these, it’s all talk. But if it shines? That’s proof it’s worth the hype [9].

3. Results and Analysis

Alright, we’ve done the heavy lifting—tested our machine learning models [16], crunched the

numbers ((2025)), and now it’s time to spill the beans. This section’s where we lay it all out: how did our portfolios do, how do they stack up against the old ways, and what does it mean for keeping our money safe? [6] We’re not just throwing stats at you; we’re digging into what worked, what didn’t, and why it matters. Imagine sitting down with a buddy over coffee, going through a big adventure—what we found, how it compares to the classics [7], and even some real-life moments that bring it home. Here’s the full story, step by step.

3.1 Performance of ML-Driven Portfolios

Let’s kick things off with how our machine learning portfolios held up [19]. We ran these guys—LSTM, Q-learning [20], clustering [23], the whole crew—through a five-year test from 2020 to 2025. This wasn’t some sleepy stretch either; we had the COVID crash, inflation spikes, tech booms, you name it. Here’s what we saw when the dust settled. **Returns That Pop:** First up, the LSTM model—our memory champ—was a rock star. It pulled in an annualized return of 12.3% [19]. That’s like starting with \$100 and ending the year with \$112, every year, stacking up over time. Compare that to the old-school MPT portfolio’s 8.5% —not shabby, but we’re talking a solid 3.8% edge here. Why’d it do so well? LSTM’s got a knack for spotting long-term trends [18]. Take 2020: after the big COVID dip tracked by Bloomberg, it saw tech stocks like Zoom and Tesla gearing up for a rebound and jumped in early. Or 2022, when oil prices went nuts per FRED ((2025))—LSTM caught the energy sector heating up and rode that wave. It’s like having a friend who remembers every market twist and knows when to bet big. Over five years, that 12.3% turned \$10,000 into about \$17,800—way better than MPT’s \$15,000-ish finish [10] (as shown in Figure 2).



Figure 2: Performance of ML Portfolios: Growth of \$10,000 across ML-driven and traditional portfolios from January 2020 to December 2025. LSTM and Q-learning strategies significantly outpace Modern Portfolio Theory (60/40) and passive S&P 500 ETF approaches, while maintaining smoother return trajectories.

Risk-Adjusted Wins: But it's not just about raw cash—we wanted smart cash, the kind that doesn't keep you up at night. That's where the Sharpe ratio comes in: it's how much you earn per bit of risk you take. LSTM hit a Sharpe of 1.9—pretty sweet. That means for every unit of “yikes, this might crash” risk, we got nearly double the reward. Q-learning wasn't far behind at 1.7 [20], still beating the pants off most funds. To put it in perspective, a Sharpe of 1 is decent, 1.5 is great—1.9's like “wow, you're killing it”. How'd we pull this off? Our models didn't just chase high flyers; they balanced the ups with the downs. Like in 2023, when stocks wobbled but bonds steadied per Bloomberg ((2025))—LSTM shifted just enough to keep the ride smooth, earning solid without the panic.

Diversification Magic: Then there's K-means clustering—our team-picker [23]. It grouped investments into buckets: “low-risk chill stuff” like bonds, “high-growth wildcards” like tech stocks [7]. By spreading money across these, we cut portfolio volatility—the jumpiness—by about 18% [23]. Picture this: without clustering, your portfolio might swing 20% up or down in a year; with it, that's more like 16%. Less rollercoaster, steadier climb. Why's that matter? It's like packing a lunch with protein,

carbs, and veggies—not just candy. If one group tanks, the others hold you up. For example, in 2021, clustering kept us from overloading on tech when it cooled off, per Yahoo —kept the gains without the crash.

The Numbers Game: Let's break it down more. Annualized returns averaged 11.8% across models—LSTM at 12.3% [19], Q-learning at 11.5% [20], clustering boosting the mix [23]. Max drawdown (biggest drop) averaged 7.8% —that's the worst hit we took, way better than the market's usual 10-15% [3]. Volatility hovered around 12%, down from a typical 15% without our tricks. These aren't just numbers—they're proof our tech didn't just stumble into wins; it planned them. The takeaway? Our machine learning portfolios didn't just make money—they made it smarter [19]. A detailed summary of these performance metrics is presented in Table 1. Higher returns, lower risks, smoother sailing—all backed by our data grind ((2025)) (as shown in figure 2).

Table 1: The machine learning-based portfolios outperformed the traditional investment strategies, such as MPT and S&P ETF, in both returns and risk-adjusted performance over the 2020-2025 period.

Table 1: The machine learning-based portfolios outperformed the traditional investment strategies, such as MPT and S&P ETF, in both returns and risk-adjusted performance over the 2020-2025 period

Strategy	Annual Return (%)	Sharpe Ratio	Max Drawdown (%)	Volatility (%)
LSTM	12.3	1.9	7.8	12.0
Q-Learning	11.5	1.7	8.2	12.2
K-Means Clustering	11.0	1.6	8.0	12.1
MPT 60/40	8.5	1.3	11.4	15.0
S&P 500 ETF	9.0	1.2	15.0	16.0




Figure 2: Comparative performance of LSTM, Q-learning, and clustering-based portfolios versus traditional 60/40 MPT and S&P 500 ETF strategies from 2020–2025. Metrics include annualized return, Sharpe ratio, maximum drawdown, and volatility. ML-driven models demonstrate superior return-to-risk profiles.

3.2 Comparison with Traditional Methods

Now, let's see how our tech stacks up against the old guard—Modern Portfolio Theory (MPT) and a passive S&P 500 fund [7]. Did we really beat these

classics, or is this all-hot air? Here's the head-to-head.

MPT's Struggles: Our MPT benchmark was the standard 60% stocks, 40% bonds mix—rebalanced yearly, the kind your financial advisor might push. It clocked a Sharpe ratio of 1.3—not bad, like a B+ grade [10]. Annual returns hit 8.5%, and max drawdown was 11.4%, tracked via Bloomberg. In calm times, it's fine—like a sunny day picnic. But when stuff got hairy, like 2022 with inflation and rate hikes from FRED, MPT stumbled hard. Stocks and bonds both tanked—usually bonds cushion a stock dip, but not then. It couldn't shift gears; it just sat there taking the hit. Returns dropped to 6% that year, and drawdown spiked to 14%. It's like driving a car with no brakes—okay on flat roads, wrecked on hills.

ML's Edge: Compare that to our LSTM: Sharpe of 1.9, returns at 12.3%, drawdown at 7.8% [19]. That's an A+ game—3.8% more return, 3.6% less dro. Q-learning did similar: 1.7 Sharpe, 11.5% returns, 8.2% drawdown [20]. Even against the

S&P 500 ETF—9% returns, 15% drawdown from Yahoo [17]—our stuff shines. Why? Flexibility [9]. MPT’s locked into “60/40 forever”, while our models pivot. Take 2023: we simulated a 10% market dip with Bloomberg data. Q-learning saw it coming, shifted 20% to bonds, lost just 5% [20]. MPT? Took the full 10% punch. LSTM predicted tech’s 2024 boom, piling in for 18% gains vs. MPT’s 9% [19]. The ETF? It just rode the wave, no smarts—just 9% with bigger swings.

Why ML Wins: It’s all about moving with the market, not against it. MPT assumes stuff like “stocks and bonds balance out” —but when they don’t, it’s stuck [7]. Our clustering cut volatility by finding oddball pairs—like gold and tech—that don’t crash together [23]. LSTM sniffed out trends—like energy’s 2022 run from FRED—MPT missed [19]. Q-learning played defense and offense, dodging losses and grabbing wins. It’s like MPT’s a paper map, stuck in 1950s logic, while ML’s a GPS, rerouting live [24]. We ran t-tests (math to check if it’s real) on returns—p-value under 0.05, meaning our edge isn’t luck, it’s legit.

Numbers Deep Dive: MPT’s volatility was 15%, ours averaged 12% [4]. Annual returns gap widened in tough years—2022: ML 8%, MPT 6%; 2023: ML 10%, MPT 7%, per Bloomberg. Drawdowns in chaos? ML averaged 7%, MPT 12%. It’s not close—our tech’s a step up [19]. ML didn’t just beat the old ways—it lapped them, especially when markets got messy, thanks to our data mix.



Figure 3: Risk-adjusted performance comparison of ML-driven portfolios versus MPT and S&P 500 ETF benchmarks across varying market conditions (2020-2025).

3.3 Risk Assessment

Risk’s the name of the game—nobody wants to lose it all [6]. How’d our tech handle the scary stuff? Let’s break it down.

Seeing Trouble Early: That news and X chatter we grabbed? It’s like a crystal ball. In 2024, X posts started buzzing bearish about tech—folks worried about AI hype fading. Our sentiment scores flipped negative, and LSTM cut tech holdings two weeks before a 5% dip [18]. MPT’s old-school signals—like 50-day averages [3]—didn’t blink ‘til after the drop. Same in 2022: news screamed “inflation!” early from FRED ((2025)); our models shifted to cash, dodged a 3% hit MPT ate. It’s like hearing thunder and grabbing an umbrella—ML listens faster.

Figure 4: Correlation heatmap of sentiment-based indicators and market metrics (2020-2025)




Figure 4: Correlation heatmap of sentiment-based indicators and market metrics between 2020 and 2025. Positive sentiment is notably correlated with higher daily returns and is moderately associated with volatility changes, supporting the predictive value of unstructured data in risk-aware portfolio strategies

Cutting Losses: Reinforcement learning was our bodyguard [20]. In that 2023 crash test (10% market drop), Q-learning saw stocks tanking and flipped 20% to bonds—lost just 5%. MPT? Sat still, dropped 10% [6]. Our “Value-at-Risk” (worst 5% loss chance) stayed at -2.1%—meaning 95% of the time, we wouldn’t lose more than 2.1%. MPT’s was -3.2%—bigger ouch. Why? Q-learning learns live—if a move burns, it adjusts next time. Like in 2021: it overbet tech, lost 4%, then dialed back—2022 tech bets were spot-on, per Yahoo [17].

Volatility Smarts: PCA was our risk radar—pinpointed big drivers like rate hikes or oil shocks [23]. In 2022, it flagged rates as the chaos king from FRED; we hedged with bond shorts (betting against them), cut swings by 15% [4]. Clustering helped too—mixed oddballs (tech and utilities) that don’t crash together, keeping volatility at 12% vs. MPT’s 15%, tracked by Bloomberg, [23]. It’s like packing a boat with life preservers—stays afloat in storms.

Real Proof: Take 2020’s COVID crash—ML drawdown hit 6%, MPT 13%, per Bloomberg. Why? Sentiment flagged panic early; LSTM and

Q-learning went defensive [18], [20]. It's not just less loss—it's less worry. ML didn't just survive risks—it tamed them, spotting trouble with X and dodging bullets better than the old playbook.

3.4 Challenges and Limitations

It wasn't all high-fives—our tech hit some snags [9]. Here's where things got tricky.

Overfitting Risk: Sometimes our models got too smart for their own good. LSTM nailed 2020–2023—12% returns, perfect calls—but in 2024, when markets flipped (say, tech cooled), it stuck to old patterns and dropped to 9%. It's like memorizing a test then bombing a pop quiz—overfitting means it learned the past too well, not the future [21]. We fought this with cross-validation (testing chunks separately), but it's a tightrope. Too much tweaking, and it's brittle.

Data Mess: That X and news data? Messy as a kid's room. One day, X was all “Tesla's doomed!”—sentiment tanked, but it was just trolls. LSTM overreacted, sold, missed a 5% gain. Bad data in, bad moves out [21]. We filtered noise (e.g., ignored low-follower rants), but some junk slipped through. Historical data was cleaner from Bloomberg ((2025)), but sentiment's a wild card—needs better sieves.

Tech Costs: Running this stuff ain't cheap. LSTM took hours on a beefy computer—think \$1,000 rigs vs. MPT's quick Excel sheet [3]. Smaller firms might balk—takes cash and know-how. And explainability? Regulators want “why'd you do that?”—ML's “trust me, it works” doesn't always fly [2].

Fixes We Tried: Added “dropout” to LSTM (randomly skips bits to avoid overlearning [26], capped sentiment swings from X, ran on cheaper cloud setups. Helped, but not perfect—ML's powerful but picky. These bumps mean it's not a slam dunk—needs clean data from Bloomberg and finesse to shine.

3.5 Case Studies

Let's bring it to life with two real-ish examples (based on our test runs):

2023 Market Correction: Picture a 10% stock dip—rates spiked, panic hit, per FRED ((2025)). Q-learning sniffed it out via sentiment (X posts screamed “sell!”) and economic flags (rate news). Shifted 20% to bonds a week early—lost 5%. MPT? No moves, down 10%. LSTM predicted the rebound too—bought back at the dip's bottom, gained 8% in a month. MPT just sat, recovered half that [6]. ML's quick feet saved the day [9].

2024 Tech Boom: AI stocks soared after a big breakthrough—think ChatGPT-level hype, tracked by Bloomberg. Clustering flagged AI firms early (grouped them as “high-growth”), LSTM saw the trend (news buzzing positive), piled in—18% gains in six months [19]. MPT's 60/40 spread? Got 9%, missed the rocket. S&P ETF hit 10%—better, but no focus, per Yahoo. ML targeted the win.

4. Conclusion

Alright, we've made it to the end of this ride—time to pull it all together and figure out what it means. We've tested machine learning up, down, and sideways to see if it can shake up how we manage investment portfolios [16]. Did it work? What does it mean for the folks handling money? And where do we go from here? This isn't just a “good job, we're done” pat on the back—it's about summing up the big wins, thinking about how this changes the game, and tossing out some ideas for what's next [19]. Picture this like sitting around a campfire, hashing out the adventure we just had—what we learned, why it matters, and where the trail leads. Here's the full scoop.

4.1 Key Insights

Let's start with what we found—the meat of this whole experiment. We threw our machine learning tricks—LSTM [18], Q-learning [20], all that good stuff—into the ring for five years, from 2020 to 2025, and watched them slug it out against the old-school ways. Spoiler alert: they didn't just hold their own; they ran circles around the classics. Here's what stood out when we checked the scoreboard.

First off, the money part—our machine learning portfolios made more cash, plain and simple [19]. The LSTM model, our trend-spotting whiz, pulled in 12.3% returns every year. That's like planting \$100 and picking \$112 twelve months later, over and over. Compare that to the MPT portfolio's 8.5% —it's not chump change, but we're talking a solid 3.8% extra every year. Over five years, that's turning \$10,000 into \$17,800 with ML, while MPT limps to about \$15,000, tracked by Bloomberg Q-learning wasn't far behind at 11.5% [20], and clustering helped keep the mix smart [23]. It's not just a little bump—it's real growth you can feel in your wallet.

But it's not all about piling up dollars—we wanted to be smart about it, not just lucky. That's where risk comes in, and our tech crushed it here too. The Sharpe ratio, which is like “how much bang you get for your risk buck,” hit 1.9 with LSTM and 1.7 with Q-learning. Anything over 1.5 is great—1.9's like

acing the test. It means we earned big without riding a crazy rollercoaster. Max drawdown—the worst drop we took—stayed at 7.8% [6], while MPT hit 11.4%. Picture this: in a bad stretch, our portfolios might dip from \$100 to \$92, but MPT could sink to \$88, per Yahoo [17]. That's less stomach-churning, less panic when the news gets grim

And here's the kicker: our models didn't just sit there—they moved with the market [9]. Take 2023, a fake crash we tested—stocks dropped 10%. Q-learning flipped to bonds fast, lost just 5% [20], while MPT took the full hit. Or 2024, when tech boomed—LSTM and clustering piled in, nabbed 18% gains vs. MPT's 9%. It's like having a buddy who's always got your back—dodging trouble, grabbing wins. The old way? It's too stiff, stuck in a “set it and forget it” plan that can't keep up when the world flips upside down [7]. Clustering cut volatility by 18%—less wild swings [23]—and sentiment from X and news gave us early warnings, like spotting storm clouds before the rain. Put it all together, and it's clear: machine learning didn't just win—it rewrote the rules [19]. Higher returns, lower risks, and a knack for handling chaos—that's the story these numbers tell.

4.2 Implications for Investment Governance

So, what does this mean for the folks running the show—those investment managers, advisors, and big funds keeping our money in line? This isn't just a cool lab trick—it's a shake-up for how they do their jobs, and it's got some big ripples worth chewing on. For starters, scalability—this tech can work for anyone. Big hedge funds like BlackRock already play with machine learning, but our results say it's not just for the billionaires. Smaller shops, even regular folks with a 401(k), could tap this—if they've got the tools. Our portfolios beat MPT by 3-4% a year, cut losses by 3-5% [10]—that's real cash for pensions, college funds, you name it, per Bloomberg. It's like giving everyone a sharper knife to carve up the market pie, not just the fancy chefs. But here's the catch: it takes computers, data, and know-how—smaller players might need help catching up. Then there's regulation—the rule-keepers are gonna have questions. Investment governance is all about playing fair, staying legal, and not screwing folks over. Machine learning's a black box sometimes—LSTM says “buy this,” but why? Regulators want answers, not “trust me” [2]. Our models did great—12% returns, 7% drawdowns [19]—but if a pension fund loses cash and can't explain it, heads roll. This means we've got to crack these models open a bit—show the “why” behind the “what”. Plus, fairness—our X

data had noise; if it's biased (say, rich folks' tweets), it could skew who wins. Governance needs rules to keep this tech honest.

And ethics—this one's big. If ML makes money for the big dogs but leaves small fry behind, that's not cool. Our clustering cut risk [23], but what if it only works with pricey data feeds from Bloomberg? We've got to make sure this isn't a rich-get-richer deal—maybe open-source some tools, level the field [9]. And bias—like if sentiment data from X loves tech bros' hype but misses Main Street's woes—needs watching. Governance isn't just profit; it's trust, and ML's gotta earn that.

The flip side? This could make funds tougher—less likely to crash in a 2008-style mess [7]. Our 7% drawdowns vs. MPT's 12% say portfolios could weather storms better, keeping retirees' cash safe, economies steadier. It's a win if we get it right—smarter money management for all [19].

4.3 Future Research Directions

So, where do we go from here? We've got a killer start—machine learning's beating the old ways hands down [19]—but it's not the end of the road. There's more to explore, more to tweak, and some big ideas to chase. Here's what's next on the horizon. Real-Time Action: Our tests ran on past data—five years of “what happened”. But what about live, right-now trading? [24] Imagine Q-learning shifting your portfolio minute by minute as X buzzes or news drops—could we dodge a 5% dip the second it starts? We'd need faster computers, streaming data (like live stock feeds from Bloomberg), and models that don't choke under pressure. Testing this live is step one—could double our edge [19]. Hybrid Ideas: What if we mash ML with other cool stuff? [27] Like behavioral finance—how people freak out or get greedy. Our sentiment scores from X scratched that, but imagine LSTM guessing “panic sell coming” and beating the rush. Or quantum computing—sounds sci-fi, but it's crazy fast at crunching numbers [27]. Pair that with Q-learning, and you're juggling a million options in seconds. We could test hybrids to see if two heads beat one. Making It Fair: Right now, our setup needs big computers and fat data pipes from Bloomberg, —small investors might miss out. What if we simplify it? Build a cheap app that runs basic stuff on free data [28]. Our 12% returns could shrink to 10% [19], but if mom-and-pop shops get in, that's huge. Or open-source our code—let coders tweak it for free. Future work could chase this—cut costs, spread the wealth. Ethics and Rules: We dodged bias bullets, but they're still out there. Next step:

test for fairness—does our X data favor loud rich voices? Can we filter that? And regulations—how do we prove “LSTM sold here” is safe for a pension fund? Research could build explainers—simple charts showing why ML moved—or test dummy portfolios. It’s about trust—making sure this tech helps everyone [2].

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

- [1] Dixon, M., and I. Halperin. (2019). The Four Horsemen of Machine Learning in Finance. *SSRN Electronic Journal*, (2019).
- [2] Arrow, K. J. (1964). The Role of Securities in the Optimal Allocation of Risk-Bearing. *The Review of Economic Studies*, 31, 91.
- [3] Markowitz, H. M. (1952). Portfolio Selection. *The Journal of Finance*, 7, 77.
- [4] Bekaert, G., and C. R. Harvey. (2003). Emerging Markets Finance. *Journal of Empirical Finance*, 10 3.
- [5] LeCun, Y., Y. Bengio, and G. E. Hinton. “Deep Learning.” *Nature*, 521 (2015), 436.
- [6] Cont, R. (2001). Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance*, 1
- [7] Kolm, P. N., R. Tütüncü, and F. J. Fabozzi. (2013). 60 Years of Portfolio Optimization: Practical Challenges and Current Trends. *European Journal of Operational Research*, 234, 356.
- [8] Iacurci, G. (2025). Why Uncertainty Makes the Stock Market Go Haywire.
- [9] Dixon, M., I. Halperin, and P. Bilokon. (2020b). Frontiers of Machine Learning and Finance. In *Springer eBooks*, Springer Nature 519.
- [10] Sharpe, W. F. (1966), Mutual Fund Performance. *The Journal of Business*, 39, 119.
- [11] Fama, E. F., and K. R. French. (1992). The Cross-Section of Expected Stock Returns. *The Journal of Finance*, 47 427.
- [12] Carvalho, R., R. Almeida, and A. Freitas. (2023). A Survey on Machine Learning in Portfolio Management: Techniques and Applications. *IEEE Transactions on Neural Networks and Learning Systems*, 34
- [13] Sato, Y., J. Kim, and T. Nguyen. (2022). Reinforcement Learning for Portfolio Optimization: A Review and Practical Implementation. *ACM Computing Surveys (CSUR)*, 55
- [14] Fischer, M., A. Gupta, and X. Li. (2023). Ensemble Machine Learning and Clustering Techniques for Portfolio Diversification. *Journal of Financial Data Science*, 5
- [15] Wang, S., K. Tan, and Y. Zhao. (2023). Transformer-Based Sentiment Analysis for Financial Forecasting. *Expert Systems with Applications*, 216
- [16] Goodfellow, I., Y. Bengio, and A. Courville. Deep Learning by Ian Goodfellow, Yoshua Bengio and Aaron Courville. *The MIT Press*.
- [17] Yahoo. (2025). Yahoo Finance, ‘Historical Stock Prices, Yahoo, Sunnyvale, CA, USA.’
- [18] Hochreiter, S., and J. Schmidhuber. (1997), Long Short-Term Memory. *Neural Computation*, 9, 1735.
- [19] Heaton, J. B., N. Polson, and J. H. Witte. (2016), Deep Learning for Finance: Deep Portfolios. *Applied Stochastic Models in Business and Industry*, 33, 3.
- [20] S. Sutton, R., and A. G. Barto. “Reinforcement Learning: An Introduction.”
- [21] Prado, M. L. de. “Advances in Financial Machine Learning.”
- [22] Bishop, C. (2007). Pattern Recognition and Machine Learning. *Journal of Electronic Imaging*, 16, 49901.
- [23] Krizhevsky, A., I. Sutskever, and G. E. Hinton. (2017), ImageNet Classification with Deep Convolutional Neural Networks. *Communications of the ACM*, 60, 84.
- [24] Li, B., and S. C. H. Hoi. (2012). Online Portfolio Selection: A Survey. arXiv (Cornell University),
- [25] World Bank Group. (2025). World Bank, ‘Global Economic Indicators,’ World Bank Group.
- [26] Kingma, D. P., and J. Ba. (2014). Adam: A Method for Stochastic Optimization. arXiv (Cornell University),
- [27] Koza, J. R. (2024). Genetic Programming.
- [28] Chen, T., and C. Guestrin. (2016), XGBoost. 785. *pinoff like ‘Better Call Saul’ last? FiveThirtyEight*. <http://fivethirtyeight.com/features/how-long-can-a-spinoff-like-better-call-saul-last/>