Volume 2026, Issue 2 - Jun 2026

Quant Quarterly

NC STATE UNIVERSITY

Financial Mathematics

Contents

Dr. Tao Pang
Ph.D., CFA, FRM
Professor and Director

We are thrilled to share several exciting updates from our Financial Mathematics (FM) program this spring.

Student Success & Career Achievements
Our community continues to thrive. This past spring, we welcomed three new students who all achieved outstanding academic results in their first semester. Additionally, we proudly celebrate four students who successfully earned their Master's degrees in Financial Mathematics this May. Exceptionally, 100% of our graduates secured full-time job offers prior to graduation.

Innovation in the Classroom: Spring Projects
During the Spring 2026 semester, students completed 12 group projects covering cutting-edge industry topics, including: Volatility, Factor Modeling, CCAR, Tail Risk, Climate Risk, Deep Learning, and Yield Curve Modeling, etc. These projects were presented in person on April 24 to an invited audience of industry guests and alumni. Students received invaluable, first-hand feedback from seasoned professionals, many of whom have previous experience with these exact types of projects.

Alumni Engagement & Networking
Following the student presentations on April 24, we hosted our Workshop on Quantitative Finance. The event brought together 18 alumni and senior industry leaders from top firms, including Bank of America, Charles Schwab, Citi, First Citizens Bank, J.P. Morgan, Morgan Stanley, and Wells Fargo. These professionals shared their expertise through dynamic panel discussions, presentations, and valuable networking sessions with current students. It was an incredible opportunity for face-to-face mentorship, and our alumni thoroughly enjoyed reconnecting with the program that launched their careers.

Editorial Policy Update
As our program continues to expand, we are evolving our content strategy. Due to limited publication space, we will no longer publish articles from every student. Moving forward, we will feature a curated selection of standout student articles, allowing us to devote more space to program news and comprehensive student project reports.

Wishing everyone a wonderful, restful, and joyful summer!

Director of Career Services

Patrick Roberts
Director of Career Services

The Impact of AI on Quant Finance Careers

One of the most common questions asked by prospective students considering enrollment in the Financial Mathematics program and pursuing a career in quantitative finance is how will artificial intelligence (AI) and the rise in utilizing this technology impact the career outlook within this industry. With the rapid advancement and adoption of AI technologies, many students wonder whether these developments will significantly alter the career outlook for quantitative professionals.

There is no question that AI is revolutionizing quantitative finance by enabling real-time analysis of massive datasets, enhancing predictive modeling for alpha generation, and automating complex trading strategies. This technology also enhances risk management with dynamic models, utilizes techniques like natural language processing (NLP) for sentiment analysis, and accelerates approaches to model development.

Although this technology will create changes within quantitative finance, there are strategies for new professionals entering the field to understand how these technologies are being applied and develop complementary skills to thrive in an increasingly AI-driven financial ecosystem.

To begin, it is important to understand the ways that AI is currently being utilized within finance. Some key impacts of AI on quantitative finance include:

- **Alpha Generation & Data Analysis**: AI, particularly machine learning, identifies subtle, non-linear patterns in data that traditional models miss. It allows for processing alternative data, such as satellite imagery or consumer behavior, to gain a competitive edge.
- **Algorithmic Trading & Execution**: AI systems improve trade execution speed and precision, enhancing liquidity. They adapt to changing market conditions in real-time for better profitability.
- **Efficiency & Automation**: Generative AI accelerates the development lifecycle, from data cleaning to back testing, reducing time-to-market for trading strategies. It handles tedious tasks like documentation and code generation, allowing quants to focus on higher-level strategy.
- **Risk Management**: AI enhances risk assessment by analyzing vast amounts of data to predict market downturns and simulate, say, liquidity crises.

What can new quant professionals do to take advantage of this current and imminent change?

First, it is important to recognize that the change is here and course correct, focusing on a skills shift that embraces the techniques utilized within AI. The demand at financial institutions is shifting toward quants with skills in machine learning, data engineering, and AI tool utilization rather than just mathematics and coding. Several financial organizations have recently posted positions that include "AI Engineer" or "Machine Learning Engineer" dedicated to building out AI tools internally and proprietary versions of open-source large language models.

Second, a shift in focus from purely programming and technical skills to a more holistic understanding of the mathematics, financial theory, model interpretation, and overall communication skills. The ability to interpret results in a clear effective manner will be an essential proficiency. AI is not replacing human quants but transforming their role into one that focuses on interpreting results, managing risks, and developing strategies in partnership with AI.

At the end of the day, stakeholders, regulators, clients, and consumers will still prefer and expect human oversight, accountability and explanation over an AI output. Regulatory requirements further reinforce the need for human expertise. Interpretability and model transparency are critical in many regulated financial activities, meaning institutions must maintain clear oversight of automated systems. As a result, human judgment remains central to financial engineering, risk management, and model governance.

While concerns about automation replacing human expertise are understandable, AI should be viewed more accurately as a powerful productivity tool. By automating routine tasks and accelerating data analysis, AI allows quantitative professionals to devote more time to creative problem solving, innovative strategy development, and deeper financial insight. Rather than diminishing the role of quantitative professionals, AI is likely to amplify their impact, creating a new generation of quants who combine strong mathematical foundations with advanced computational tools to navigate an increasingly data-driven financial world.

Jabbir Ahmed
From Beta To Alpha: Building Regime-Aware Multifactor Strategies In The FM Program

Jacob Attia
Rethinking Model Complexity In Low-SNR Environments

Sartaz Aziz
A Journey Through Financial Mathematics: Theory, Research, And Applications

Shubham Balodi
Learning About Quantitative Risk From An Industry Professional

Aditya Bedi
Narrowing the Lens: How Graduate School Has Helped Me Refine My Career Direction

Tobi Bui
Bridging Environmental Risk and Fixed Income Analytics

Yiwei Cao
From Models to Markets to Makers: My Long-Term Career Journey

Leo Chien
Interest In Credit Risk Analysis In The Banking Industry

Olivia Dittrich
Insights From An Industry Conversation: Risk And Trading At A Commodity Trading Firm

Kristen Epperly
The Power Of Data Visualization

Clay Ewing
Simulating The Limit Order Book: A Collaborative Quantitative Project

Muxin Feng
Project Responsibilities Lead To Meaningful Growth

Ekaterina Finiutina
Building Predictive Trading Signals: Lessons From My Alpha Research Project

Paras Gupta
Portfolio Optimization: Bridging Mathematical Theory and Wall Street Reality

Chinna Gurram
Understanding Financial Markets Through The Shape Of Data

Dhrubojeet Haldar
Mean Reversion In Motion: Building An FX Statistical Arbitrage System From Scratch

Aahan Jain
Betting Against Beta

Soham Joshi
Building A CCAR-Style Capital Forecasting Framework

Michelle Kadiman
Signals Before Stress: A Panel Data View Of Credit Delinquency

Subham Khinchi
Derivatives Pricing Using The Heston Model

Burke King
From Probability to Policy: An Actuarial Career Vision

Jake Kostoryz
Where Sports, Markets, And Data Science Collide

Ted Kratt
Equity Analysis 101

Daksh Kumar
Navigating the Volatility Smile: A Unified Framework for Advanced Pricing Models

Jiayuan Li
Expanding My Path in Quantitative Risk at NC State

Lambert Li
A Conversation With Dimitri Bianco

Bedant Lohani
Building Toward Alpha

Siyuan Lu
The Practical Impact In The MFM Program

Students' Corner

Paritosh Mallick
How Networking Shaped My Journey In Financial Mathematics At NC State

Nicholas Moss
Building a Quantitative Finance Career Through Structure, Projects, and Industry Insight

Anas Mouden
From Fundamental Finance To Quantitative Finance: Understanding Risk, Drift, And Randomness

Maria Muntaha
From Curiosity To Quant Trading

Isaiah Nick
The Hidden Option In Prediction Markets

Trevor Otte
Crossroads of the Industry

Gaurav Pai
From Theory To Toolsets: Bridging Quantitative Rigor And Career Readiness

Pushkar Patil
Why Education Debt Isn't A Leveraged Buyout

Aditya Purohit
Quants In Climate Risk

Antarish Rautela
Birth Of A New Possibility

Noah Regal
When The Hypothesis Fails

Melanie Sanchez
Predicting Credit Risk

Catherine Sexton
The Impact Of NC State's MFM Program

Shreyansh Sharma
Delivery Risk In Futures Contracts

Students' Corner

Shiqin Song
Risk Management Through Hands-On Projects

Surya Sridhar
Climate, Credit, And Collateral

Abhishek Tandon
Local Volatility Surface Estimation

Vedashree Teli
Climate Risk: The Next Frontier In Quantitative Finance

Aleksey Urmanov
The Duplicity Of Market Microstructure

Anand Vadlamani
Beyond VaR: Measuring Tail Risk In A Fixed Income Portfolio

Isha Wadekar
Modeling Market Asymmetry: A GJR-GARCH Approach To Volatility Clustering

Yueyang Wang
Monte Carlo Simulation for Exotic Options

Zach Wiebe
Musings on Symmetry and Clarity

Derek Wu
A Calculated Decision: Why NC State's MFM Was The Right Fit

Joy Wu
Modeling Market Volatility Using GARCH Family Models: Evidence From The S&P 500

Shixun Wu
From Models to Decisions: My Career Goals in Financial Mathematics

Daniel Yang
Bridging Theory and Practice in the NC State Financial Mathematics Program

Peng Yang
High-Frequency Volatility Prediction Using Deep Learning

Jialin Zhang
From Models to Markets: What a Volatility Practitioner Taught Me About Financial Mathematics

Luolin Zhang
Where Theory Meets Practice: My Experience In Financial Mathematics

Jabbir Ahmed

From Beta To Alpha: Building Regime-Aware Multifactor Strategies In The FM Program

From Beta to Alpha: Building Regime-Aware Multifactor Strategies

If there is one thing this program has drilled into me, it is that elegant theory and live market data rarely agree without a fight. Two group projects this year forced me to bridge that gap head-on — one reimagining the classic Betting Against Beta (BAB) strategy around market regimes, and the other building a fully automated PCA-driven momentum trading system in QuantConnect. Together they changed how I think about factor investing.

Betting Against Beta — But Only When the Market Lets You

The BAB strategy, from Frazzini and Pedersen (2014), exploits a well-documented anomaly: the Security Market Line is flatter than theory predicts, meaning low-beta assets consistently outperform on a risk-adjusted basis. The idea is intuitive — go long a leveraged portfolio of low-beta stocks, short high-beta stocks, and collect the spread. But our team kept asking a question the original paper sidesteps: does this hold when markets are in crisis? The answer is no, and the data makes it obvious.

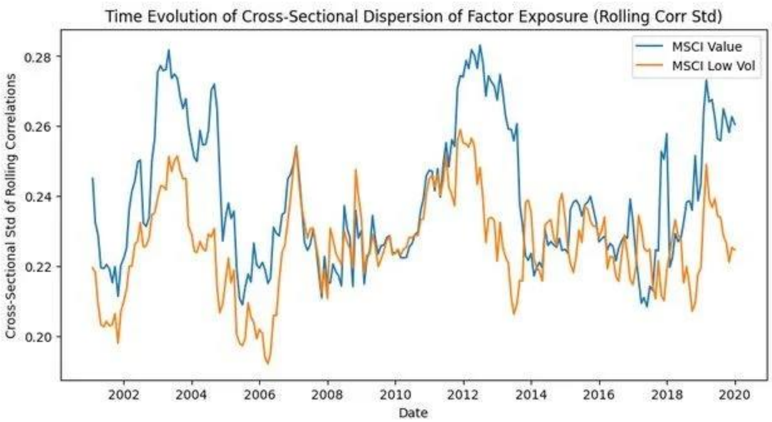

During stress periods, correlations across assets spike, low-beta stocks get dragged down alongside everything else, and the BAB signal breaks down or reverses outright as forced deleveraging unwinds positions. So we built a regime-conditioning layer on top of the standard framework.

Using the MSCI USA Index (540 constituents) as our universe, with six MSCI factor proxies — Value, Low Size, Low Volatility, High Yield, Quality, and Momentum — we classify each month as risk-on or risk-off based on the VIX. In calm regimes, the optimizer runs at full gross exposure, maximizing Sharpe under the BAB long-short constraints. In stress regimes, exposure is scaled back meaningfully. The rolling correlation structure shifts visibly across regimes, which is precisely why a static approach leaves performance on the table.

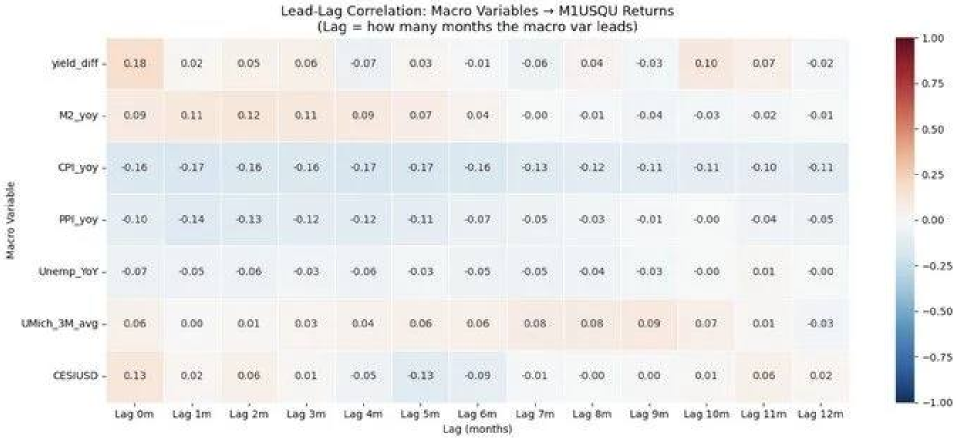

We trained the model on 1999–2019 data and tested it out-of-sample from 2020–2025, a window that includes COVID, rate shock, and the 2022 drawdown — about as adversarial a test set as you could ask for.

PCA Momentum Strategy — Making the Signal Speak for Itself

The second project tackled a different but related problem: when you have 13 correlated factors across ~800 large-cap U.S. stocks, how do you isolate a clean alpha signal? We started with a broad factor library spanning momentum (3-, 6-, 9-month), value (P/B, P/E, EV/EBITDA, EV/Sales), quality (gross margin), size, liquidity, and two volatility measures. After winsorizing outliers and z-scoring cross-sectionally, we ran PCA to extract uncorrelated principal components capturing 90% of factor variance — cutting the noise that kills out-of-sample performance.

Continue Next Page

Jabbir Ahmed

From Beta To Alpha: Building Regime-Aware Multifactor Strategies In The FM Program

Continued From Previous Page

Mispricing was then estimated using two regression methods in parallel: WLS weighted by GARCH(1,1) volatility forecasts (to downweight noisy high-vol stocks) and Lasso (to shrink irrelevant PCs to zero). The residual from each regression is the alpha signal — positive residual means undervalued, go long; negative means overvalued, go short. Positions are scaled inversely to forecasted volatility and normalized to a dollar-neutral book. A 99% VaR cap on daily leverage ensures the portfolio liquidates fully if GARCH-estimated tail risk gets too large. Backtested on monthly rebalancing from 2017–2021, the strategy produced a Sharpe ratio of 0.901 with a maximum drawdown of 11%, hit during the March 2023 stress period.

$$\hat{y}_i = w_0 + \sum_{j=1}^{m} X_{ij} w_j$$

$$J(w) = \sum_{i=1}^{n} (y_i - \hat{y}_i)^2 + \alpha \sum_{j=1}^{m} |w_j|$$

$$\|w\|^2 = \sum_{j=1}^{m} |w_j|^2$$

The win rate came in at 53%. The cumulative return path shows the strategy navigating Volmageddon in early 2018, the COVID crash, and the meme stock disruptions in 2021 without catastrophic loss — a testament to the GARCH risk module doing its job.

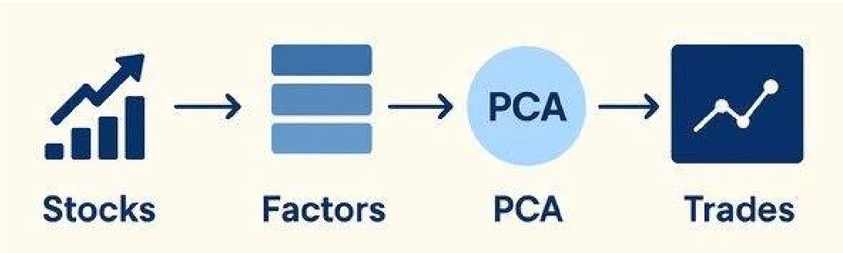

What Both Projects Taught Me

Both strategies converge on the same hard truth: factor signals are not stationary. The BAB project answers this by explicitly conditioning on market regime. The PCA project answers it by dynamically re-estimating volatility and shrinking unstable factors. Neither approach is better in isolation — but having built both from scratch, from raw data all the way to a live backtesting dashboard, I now understand intuitively why any factor model that ignores regime or vol dynamics will eventually blow up in exactly the markets it needs to perform in. That is the kind of lesson you can only learn by getting burned by your own model at 2am before a presentation.

Jacob Attia

Rethinking Model Complexity In Low-SNR Environments

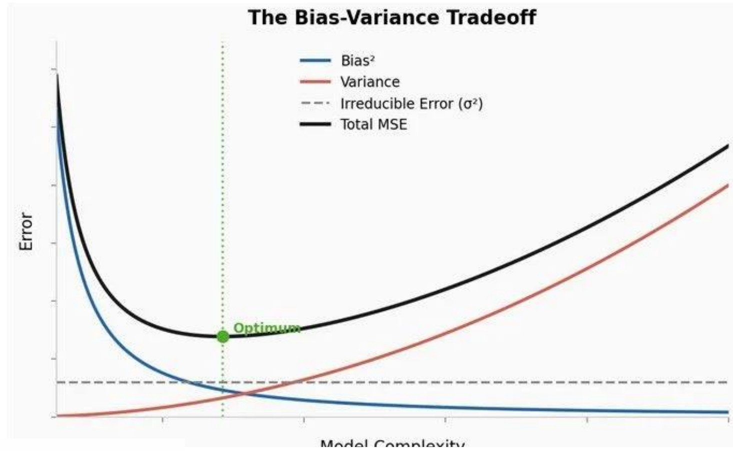

When a model performs well in-sample but falls apart out-of-sample, the instinct is often to blame the data. In quantitative finance, this reaction is almost reflexive, but the more useful diagnosis lives in a framework every statistician encounters early: the bias-variance tradeoff. Understanding it not as a mathematical curiosity but as a practical lens for model selection is one of the more consequential shifts a quant researcher can make.

The decomposition is familiar:

For a model $\hat{f}$ trained to predict a response Y, the expected test error at a point x_0 is:

$$MSE(x_0) = Bias^2(\hat{f}(x_0)) + Var(\hat{f}(x_0)) + \sigma^2$$

The irreducible noise σ^2 is the floor - model choice cannot touch it. Everything above that floor is fair game, and the tradeoff is this: as model complexity increases, bias falls but variance rises. The question is never "which is lower?" but "where does their sum bottom out?"

In most financial prediction tasks - forecasting excess returns, estimating default probabilities, classifying market regimes - the signal-to-noise ratio (SNR) is extremely low. A large fraction of σ^2 is genuine randomness baked into asset prices. This has a direct implication for where the bias-variance sum bottoms out: usually much earlier along the complexity axis than practitioners expect. A highly flexible model has the capacity to fit the data well, but in a low-SNR environment, that flexibility is mostly spent memorizing noise.

This is where model choice becomes a deliberate decision rather than a default. In the classification setting, logistic regression estimates $p(Y|X)$ directly without assumptions on the distribution of X. LDA and QDA are more restrictive - both assume Gaussian class-conditional distributions, with LDA further constraining covariance to be equal across classes. These extra assumptions increase bias relative to logistic regression, but in small-sample financial datasets that bias often pays off in reduced variance. Naive Bayes goes further, assuming conditional independence of features given the class, which is aggressive, but sometimes effective when features are many and data is scarce. In the regression setting, Ridge and Lasso introduce explicit regularization, penalizing coefficient magnitude via l_2 and l_1 norms respectively to trade a small increase in bias for a potentially large reduction in variance. Elastic net combines both. Each of these is not just a different algorithm but a different statement about where you believe the bias-variance optimum lies.

The practical implication is that in finance, the regularized linear models - Ridge, Lasso, Elastic Net, tend to outperform their unregularized counterparts not because they are inherently superior, but because the penalty term is doing real work. It is compressing the model toward a region of parameter space that generalizes better when the true signal is weak. Lasso has the added property of performing implicit variable selection, which matters when you have many candidate features and suspect most are noise - a situation any factor researcher will recognize immediately.

The practical tool for locating that bias-variance optimum is cross-validation. By repeatedly holding out subsets of the data and measuring out-of-sample error, CV traces the bias-variance tradeoff empirically - error tends to fall as complexity increases, then rises again as variance dominates. The minimum of that curve is where you want to be. In finance this is especially important because the gap between in-sample and out-of-sample performance is wider than in almost any other prediction domain. A model that looks well-specified on training data may simply have found spurious structure; CV forces the question of whether the fit generalizes.

The broader point is that model selection in a low-SNR environment should start from a tendency toward simplicity, with complexity only earned through evidence. Holdout performance, cross-validated error, out-of-sample R^2 should not be assumed. The bias-variance tradeoff gives one the language to make that argument explicitly: every increment of complexity is a bet that the reduction in bias outweighs the increase in variance. In finance, that bet loses more often than intuition suggests.

Sartaz Aziz

A Journey Through Financial Mathematics: Theory, Research, And Applications

The theoretical rigor of the Financial Mathematics program has been the most impactful aspect of my experience so far. Coming from a background that combines finance and mathematics, I was particularly drawn to the depth with which the program approaches quantitative problems in financial markets. The coursework does not simply focus on applications at a superficial level, instead, it emphasizes the mathematical foundations that make modern quantitative finance possible. This approach has significantly strengthened my ability to think about financial problems in a structured and analytical way.

During this semester, I have taken several core quantitative courses that have deepened my understanding of topics such as stochastic processes, statistical inference, and computational methods used in financial modeling. The learning experience in these courses has been intellectually demanding but extremely rewarding. Rather than merely applying formulas or software tools, the emphasis is on understanding the theoretical principles that drive these methods. This level of rigor has helped me appreciate how mathematical theory translates into practical tools used in the financial industry.

Another component of the program that has been particularly valuable is the Financial Mathematics Seminar. Over time, it naturally builds the habit of applying to jobs on a regular basis, which I think is really helpful.

My current career goal is to break into the trading space, particularly in roles related to quantitative research or systematic trading. I am aware that this path is highly competitive and demanding. However, I see that challenge as a motivation rather than an obstacle. The combination of rigorous coursework, technical projects, and independent study is helping me gradually build the skill set required for such roles. I am actively working on strengthening my programming abilities, statistical intuition, and understanding of financial markets so that I can contribute meaningfully in a quantitative trading environment.

One of the most exciting parts of my academic journey this semester has been working on a research project with my teammates. In this project, we are exploring the use of stereographic projection as a mathematical tool to model tail risk in financial returns. Tail risk, which refers to the probability of extreme market events, remains one of the most challenging aspects of financial modeling. Traditional models often struggle to capture the complex behavior of extreme returns, especially during periods of market stress.

Our project approaches this problem from a more geometric and theoretical perspective. By using stereographic projection, we aim to transform the space in which extreme observations are represented, allowing for a potentially more stable and interpretable representation of tail events. Although the framework we are developing is quite theoretical, it opens the door to several practical financial applications.

Recently, I proposed a potential extension of our framework. In addition to modeling tail behavior, the theory we have developed could potentially be used to generate more stable trading signals for mean reversion strategies. Mean reversion strategies rely heavily on identifying deviations from equilibrium levels, but such signals can often be noisy and unstable, especially during volatile market conditions. If the transformation we are studying can help stabilize the representation of extreme movements, it may also improve the reliability of signals used in systematic trading strategies. This idea is still in an early stage of exploration, but it highlights how theoretical mathematical tools can eventually lead to practical innovations in trading.

Beyond coursework and projects, I have also been actively seeking guidance from professionals working in the industry. Recently, I had the opportunity to speak with Rahul Jindal, a quantitative researcher at PGIM. During our conversation, I asked for advice regarding the job search process and preparation for full-time quantitative roles. He shared valuable insights about the skills that firms typically look for, the importance of strong statistical and programming foundations, and how to approach technical interview preparation effectively.

Overall, the Financial Mathematics program has been a deeply enriching experience. The theoretical rigor of the coursework, the collaborative research environment, and the exposure to industry perspectives have all played an important role in shaping my development. While the path toward a career in quantitative trading may be challenging, the training and opportunities provided by the program are helping me steadily move toward that goal.

Shubham Balodi

Learning About Quantitative Risk From An Industry Professional

One of the most valuable experiences during my time in the Master of Financial Mathematics (MFM) program at North Carolina State University was a recent conversation I had with Abhinav Singh, Senior Assistant Vice President and Quantitative Analytics Specialist at Wells Fargo. As someone interested in pursuing a career in quantitative risk management, I was eager to learn more about how the skills we develop in the program translate into real-world applications within the financial industry.

During our conversation, Mr. Singh shared insights into his work developing and validating quantitative models used to measure and manage financial risk. He explained how these models play a crucial role in helping financial institutions understand credit exposure, perform regulatory stress testing, and evaluate portfolio risk under different economic scenarios. What stood out to me was how quantitative analysis goes beyond simply building models; it also involves interpreting results and communicating them effectively to decision-makers who may not have a technical background.

Another key takeaway from our discussion was the importance of strong programming and data analysis skills. Mr. Singh emphasized that tools such as Python, and SQL are essential for working with large datasets and implementing risk models. He also encouraged me to continue developing practical experience through projects that involve real financial data, as hands-on work helps bridge the gap between theoretical concepts and industry applications.

I also asked him for advice about selecting electives within the MFM program. He suggested focusing on courses that strengthen core quantitative skills, particularly time series analysis, stochastic processes, and machine learning applications in finance. According to him, these areas are especially valuable for roles involving risk modeling and quantitative analytics. His guidance helped me think more strategically about how to structure my coursework to build both technical depth and practical expertise.

Overall, this conversation reinforced my interest in quantitative risk management and provided valuable perspective on how the skills developed in the MFM program can translate into impactful work in the financial industry. Speaking with someone actively working in the field gave me a clearer understanding of the technical, analytical, and communication skills required to succeed as a quantitative professional.

Aditya Bedi

Narrowing the Lens: How Graduate School Has Helped Me Refine My Career Direction

During my first semester at NC State's Master of Financial Mathematics Program, particularly around late October when the pressure and workload noticeably increased compared to August and September, I had moments of questioning my interest in Quantitative Finance and Financial Mathematics altogether. Amidst all the pressure, I found myself not being as deeply connected to or interested in the material that had once deeply resonated with me. For example, I was quite interested in a September guest lecture from a Citadel professional, who had previously worked at Goldman Sachs, on the applications of machine learning in quantitative finance, but by late October, I was finding myself less interested in similar things or less motivated to go out and look for them the way I had been even weeks earlier.

Incidentally, it was starting out my technical training assignment due in a month in early November that really helped me start to refine my career direction, much as I did not realize it then. To complete the FIM 500 technical training assignment, I started off with a Forage simulation for Capital Markets that was related to Citi. It was then that I realized my interest in capital markets, in the macroeconomic effects of inflation and how it could have been predicted as an impact of, among other things, the then interest rate policy of the Federal Reserve in response to inflation around September 2021. As the simulation progressed into how different types of swaps affect clients taking certain positions, while I initially misunderstood the intuition, I found it quite interesting. That laid the groundwork for my realizing the full extent of my interest in capital markets and financial intuition, even if I knew I was interested in it on paper.

Fast forward to January 2026, the start of the new Spring 2026 semester. It was on the first day in Dr. Ellson's FIM 549 class that the spark of interest once lighted in me by the simple simulation was reignited. Though I admittedly had my ups and downs of interest even in FIM 549, I was quite interested in a lot of related topics in that class such as mutual funds, ETFs and how they really worked in depth, as well as the mathematical mechanics behind swaps that I had gained basic familiarity with before. An interest I especially recall was how in depth the class went into what caused the 2008 Financial Crisis. I had been interested in the 2008 Financial Crisis since I was a junior in high school, so when Dr. Ellson went in depth about topics like securitization, senior substructures and collateralized debt obligations and how the combination of that and incentives created a storm, I was quite deeply interested.

Reflecting on these experiences, I realize that graduate school has not only been a place to learn more about quantitative finance but also a place to refine my interests. While I do not know what the future holds, I increasingly look forward to a career hopefully in capital markets, with renewed perspective from my experiences at NC State.

Tobi Bui

Bridging Environmental Risk and Fixed Income Analytics

As we venture into the Spring 2026 cycle, the Fixed Income landscape is undergoing a profound transformation. While the sector has traditionally relied on macroeconomic indicators such as interest rate trends and employment data, a new, formidable variable has entered the equation: Climate Change. Our current initiative represents a significant expansion of our analytical horizons, integrating high-fidelity climate elements into the rigorous world of mortgage-backed securities (MBS) and government-sponsored enterprise (GSE) risk management.

The Intersection of Climate Volatility and Housing Finance

The financial stakes of environmental shifts are no longer theoretical. Natural disasters—ranging from catastrophic hurricanes in coastal corridors to the relentless expansion of floodplains and the increasing frequency of wildfires—extract a staggering economic toll, often measured in the hundreds of billions of dollars. For lenders and GSEs, the primary challenge lies in anticipation rather than reaction.

Our project focuses on a critical, often overlooked dimension of risk: the lead-lag relationship. Traditionally, losses are quantified through the "rearview mirror" of insurance payouts after a disaster has already struck. Our approach pivots toward a proactive methodology. We are modeling how climate exposure influences mortgage behavior before the catalyst event occurs. By synthesizing publicly available climate data into proprietary input features, we are building a predictive framework that identifies vulnerabilities in the housing market long before the first storm surge or ember.

Methodology: From Zip Codes to Securitization

A particularly compelling aspect of this research is the shift in market dynamics. Current data suggests that mortgages originating in climate-prone regions are increasingly being securitized. This "green shifting" of risk onto the broader market makes our work from a GSE standpoint more vital than ever. To address this, our team has undertaken an exhaustive Exploratory Data Analysis (EDA) that transcends traditional financial boundaries.

We recognized early on that financial services are fundamentally about people, not just abstract profit-and-loss statements. To capture the human element of risk, we have:

- Conducted deep-dive research into ESG (Environmental, Social, and Governance) policy frameworks and the evolving news cycle surrounding the US housing market.
- Granularized our data to the Zip Code level across diverse counties, allowing us to "flag" specific properties for potential physical damage.
- Modeled the ripple effect: We track how property damage leads to fluctuating house prices, which in turn spikes the probability of default (PD). This is often driven by the dual burden of renovation costs and localized spikes in the unemployment rate following a disaster.

Redefining Risk Assessment for RMBS

The ultimate goal of this journey is to influence the way Residential Mortgage-Backed Securities (RMBS) are modeled. By embedding climate risk into our proprietary models, we provide companies with a more sophisticated lens for risk assessment. We can now simulate potential scenarios to determine how many months after a disaster a borrower is likely to default, providing a "stress test" that accounts for the reality of the 21st-century environment.

This project is an ongoing learning experience that blends the precision of quantitative finance with the unpredictability of meteorology. We are committed to refining our analysis, ensuring that as the climate evolves, our models remain a step ahead, protecting both the stability of the market and the homeowners within it.

Yiwei Cao

From Models to Markets to Makers: My Long-Term Career Journey

My long-term career goal is to build a career that moves from deep analytical work in finance to broader investment judgment and, eventually, to supporting innovation at its earliest stages. The path I envision is clear: first becoming a bank quant, then transitioning to the buy side, and ultimately moving into venture capital. Although these stages are different, they are connected by one common theme: using rigorous analysis to make better decisions under uncertainty.

My first goal is to work as a quantitative analyst in banking. I see this as the ideal starting point because it offers strong technical training and exposure to real financial markets. As a bank quant, I would be able to apply mathematics, statistics, and programming to problems such as pricing, risk modeling, derivatives analysis, and market behavior. This role would allow me to strengthen both my theoretical foundation and my practical problem-solving skills. More importantly, it would train me to think carefully, work with complex data, and understand how financial institutions manage risk in fast-moving environments. I believe this stage will give me the discipline and technical depth necessary for every step that follows.

After building this quantitative foundation, my next goal is to move to the buy side. While banking would teach me how to model markets, the buy side would teach me how to make investment decisions. This transition is important to me because I want to go beyond analyzing financial products and begin evaluating opportunities from the perspective of capital allocation. On the buy side, I hope to develop a stronger sense of market intuition, portfolio construction, and investment strategy. I am especially interested in learning how investors combine quantitative evidence with judgment, timing, and risk appetite. In this stage of my career, I want to become someone who not only understands models, but also knows when and how to act on them.

My ultimate goal is to enter venture capital. In the long run, I want to work with innovative companies and help identify businesses with the potential to create lasting impact. Venture capital appeals to me because it combines finance, strategy, technology, and entrepreneurship. Compared with public markets, it requires a different kind of thinking: instead of focusing only on current performance, investors must evaluate vision, founders, market potential, and long-term scalability. I am drawn to the idea of supporting companies at an early stage and helping them grow. After gaining technical discipline in banking and investment judgment on the buy side, I hope to bring both skills into venture capital, where I can assess startups with both analytical rigor and strategic perspective.

What excites me most about this career path is that each stage builds naturally on the previous one. Banking will sharpen my quantitative skills. The buy side will expand my investment mindset. Venture capital will allow me to apply both in a more forward-looking and entrepreneurial setting. I do not see these goals as separate jobs, but as a progression toward becoming a well-rounded investor with strong technical ability, sound judgment, and a long-term view of innovation.

In the end, my career ambition is not only to succeed in finance, but also to grow from a specialist into an investor who can understand numbers, markets, and ideas at the same time. That is the direction I want my career to take, and it is the path I am determined to pursue.

Leo Chien
Interest In Credit Risk Analysis In The Banking Industry

My current career goal is to work in credit risk at a large financial institution. I am particularly interested in roles such as Credit Risk Analyst or Credit Risk Model Analyst at major banks like JPMorgan Chase, Citi, or Bank of America. These roles focus on evaluating the risk that borrowers may default on their obligations and helping banks manage their exposure to credit losses.

I became interested in credit risk through my studies in the Financial Mathematics program at NC State. One of the courses that influenced me the most was Financial Risk Analysis, where we studied how financial institutions measure and manage different types of risk. In that class, we learned about important risk measures such as Value-at-Risk (VaR) and Expected Shortfall, which are commonly used to estimate potential losses in financial portfolios.

Although VaR is often discussed in the context of market risk, learning about these risk measures helped me understand how quantitative tools are used more broadly in risk management. Financial institutions rely on models to estimate potential losses and determine how much capital they need to hold against those risks. This process is especially important for large banks that manage very large portfolios of loans and financial assets.

Credit risk analysis focuses on understanding the likelihood that borrowers may fail to repay their loans. Analysts often use statistical models and historical data to estimate the probability of default and potential losses associated with different types of borrowers. These models help banks evaluate lending decisions and manage the overall risk of their portfolios.

One aspect of credit risk that I find particularly interesting is how quantitative models are used to support real business decisions. Risk analysts must analyze large datasets, test model assumptions, and interpret the results carefully. Small changes in model assumptions can sometimes lead to very different estimates of risk, which means analysts need both strong technical skills and good judgment.

The Financial Mathematics program has helped me develop many of the skills needed for this type of work. Courses in statistics, econometrics, and financial modeling provide the theoretical background for understanding risk models, while programming assignments help us work with financial datasets and implement quantitative analysis.

Looking ahead, I hope to continue strengthening my skills in statistical modeling and programming, especially using Python. These tools are widely used by risk teams in large financial institutions. My goal is to apply these quantitative skills in a credit risk role where I can help banks better understand and manage the risks associated with lending and financial markets.

Olivia Dittrich

Insights From An Industry Conversation: Risk And Trading At A Commodity Trading Firm

One of the most valuable parts of my Master of Financial Mathematics program at NC State has been the opportunity to connect with professionals working in the industries I hope to enter. Recently, I had the chance to speak with Martijn Bron, a former Global Head of Trading for Cargill's cocoa and chocolate division and the current co-host of the Strong Source commodities podcast. I originally discovered his podcast while trying to learn more about commodity markets and trading houses such as Cargill, where I will be working as a Risk Management Intern this summer.

After listening to several episodes of the podcast, I reached out to Martijn on LinkedIn. He generously agreed to speak with me about the relationship between risk management and trading inside commodity firms and what someone beginning their career at a company like Cargill should understand.

One of the most interesting parts of our conversation was his discussion of the relationship between risk managers and traders. From the outside, it is easy to assume that both groups simply work toward the same goal of managing financial risk within the firm. However, Martijn explained that in practice the relationship can be more complicated. Traders are responsible for generating profits and managing positions in volatile markets, while risk managers are responsible for monitoring exposures, verifying positions, and ensuring that activity remains within the firm's risk limits. Because of this difference in responsibilities, the relationship between the two functions is not always perfectly aligned.

Martijn was very candid in explaining that risk and trading teams do not always have an easy relationship. Risk managers sometimes need to question positions or profit-and-loss reports, which can naturally create tension. However, he emphasized that the best risk professionals approach this responsibility carefully and thoughtfully.

One piece of advice he gave me that I will not forget was about verifying P&L and trade data. When reviewing a trader's positions or daily P&L, he advised that someone working in risk should thoroughly check all possible sources of error before assuming that a trader has made a mistake. Market prices, trade capture systems, timing differences, and valuation models can all affect reported results. In other words, discrepancies in P&L are not always the result of incorrect trades. A good risk manager investigates the full set of possible explanations before raising concerns.

This advice highlighted an important principle about risk management in trading organizations: credibility and trust matter. If a risk manager frequently questions traders without fully understanding the data or systems behind the positions, the working relationship can quickly deteriorate. On the other hand, if risk professionals demonstrate that they have carefully checked the information and understand the mechanics behind the trades, their concerns will be taken much more seriously.

Beyond this specific advice, the conversation gave me a clearer understanding of how risk management functions within a large commodity trading organization. Risk teams play a critical role in ensuring that traders operate within defined limits and that exposures are properly measured. At the same time, they must maintain productive working relationships with trading teams in order to do their jobs effectively.

As someone preparing to begin an internship in risk management at Cargill, this perspective was extremely valuable. Hearing directly from someone who spent decades in commodity trading helped me understand not only the technical aspects of the role but also the interpersonal dynamics that shape how trading organizations operate.

Conversations like this are a reminder that learning in a financial mathematics program does not happen only in the classroom. Speaking with industry professionals provides context for the models and concepts we study and helps connect quantitative methods with the realities of financial markets.

Kristen Epperly

The Power Of Data Visualization

Throughout my academic and professional career, I have been inspired by the ability to transform data into valuable insights through analysis and visualization. My internship experiences in Business Intelligence and Financial Planning and Analysis have shaped my technical skills as well as my future career goals. These roles allowed me to see how data visualization can improve decision making and the understanding of complex ideas. Moving forward, my goal is to pursue a career in Financial Planning and Analysis (FP&A), ideally in the insurance industry where I can continue to develop my knowledge of financial concepts and use data visualization to drive impactful decisions.

In my role as a Business Intelligence intern for Coca Cola Consolidated, I gained foundational knowledge of how data is used by companies to optimize performance. There, I created a forecasted headcount based on deliveries and created data visualizations using Tableau to demonstrate my results. I learned that when presented clearly, data can reveal patterns that would otherwise be difficult to convey. By telling a story using the data, I was able to highlight trends and identify areas of improvement. This is the role that opened my eyes to the possibilities of understanding raw data through data visualization.

My internship in Financial Planning and Analysis at Enact Mortgage Insurance deepened my interest in the financial side of business. While my business intelligence role focused on data analysis and reporting, FP&A introduced me to how financial performance is monitored, strategic planning is developed, and how the finance teams partner with other departments to guide in their workflow.

In this role, I leveraged Tableau to create an interactive dashboard to be used for premiums reporting. The idea was to improve the timeliness and accuracy of the premiums reporting process. By consolidating the data into one clean and easy to use Tableau dashboard, employees were able to navigate the filters to investigate varying trends in the data. This made for a faster reporting process and deeper understanding of what metrics were driving business changes. This experience reinforced my belief that data visualization is a powerful advantage for the finance industry.

Working for Enact Mortgage Insurance also sparked my interest in furthering my career in the insurance sector. My goal is to pursue a career in FP&A as well as to continue implementing data visualization tools in finance to discover new ways to analyze data. I am passionate about the opportunity to improve efficiency and accuracy through strategic data investigation and visualization.

Clay Ewing

Simulating The Limit Order Book: A Collaborative Quantitative Project

Modern financial markets rely on a mechanism that most investors never see directly: the limit order book. Behind every trade lies a constantly evolving system of buy and sell orders that determine prices, liquidity, and overall market dynamics. Last semester, I had the opportunity to explore this system through a collaborative project where our group built a simulated limit order book using Python. The project gave me a deeper appreciation for how markets function at their most fundamental level while also highlighting the importance of teamwork in solving complex quantitative problems.

Our goal was to create a simplified simulation of a limit order book and observe how order flow affects market behavior. Using Python, we constructed a model that tracked buy and sell orders across different price levels. We modeled order arrivals for both noise traders and informed traders using a Poisson-based process, allowing for randomness in the flow of orders. As market and limit orders entered the system, the simulation updated the order book dynamically, allowing us to observe how trades occurred and how liquidity evolved over time.

Even within this simplified framework, the model demonstrated how price movements emerge from the interaction between supply and demand. We were able to observe how small changes in order flow could significantly impact spreads, market depth, and trade execution. The graphs below illustrate two different market dynamics. In the first, short-term volatility leads to potential price jumps and increased downside risk. In the second, many buyers are waiting at lower prices while very few sellers are offering shares near the current price. This imbalance pushes the displayed price upward more than the underlying liquidity might suggest.

One of the most valuable aspects of the project was the collaborative environment in which it was completed. Our group consisted of five members, guided by a group leader who was a year ahead in the program. Having someone with additional experience helped us navigate both the technical and conceptual challenges of the project. Each member contributed in different ways. Some focused on coding the simulation structure, while others concentrated on analyzing the output and interpreting the results. This division of responsibilities allowed us to move forward efficiently while still collaborating closely on key decisions.

Regular discussions were essential throughout the project. As we refined the simulation, we frequently reviewed each other’s work and debated how best to represent different aspects of order flow. These conversations helped us better understand both the technical implementation and the economic intuition behind the model.

This project ultimately gave us a deeper appreciation for the complexity of market microstructure. While investors often think about markets in terms of prices and returns, those outcomes are generated by the continuous interaction of orders within the limit order book. The experience reinforced the idea that understanding the mechanics behind trading is essential for anyone pursuing a career in quantitative finance.

Projects like this are a major reason I chose to pursue the MFM program at NC State. They provide an opportunity to apply mathematical and computational tools to real financial systems while learning from peers with different perspectives. This type of collaboration is something I believe will remain extremely important as we all continue developing our careers. Through this experience, I developed a stronger interest in market modeling and quantitative analysis. As I continue through the program, I hope to build on this foundation and further explore how models like these can be used to better understand trading behavior, liquidity risk, and financial market dynamics.

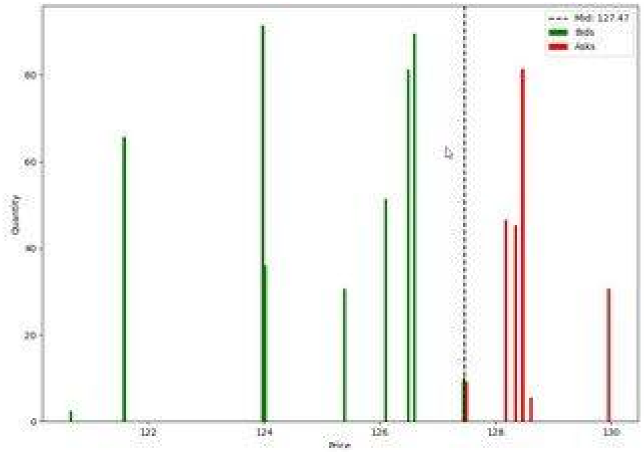

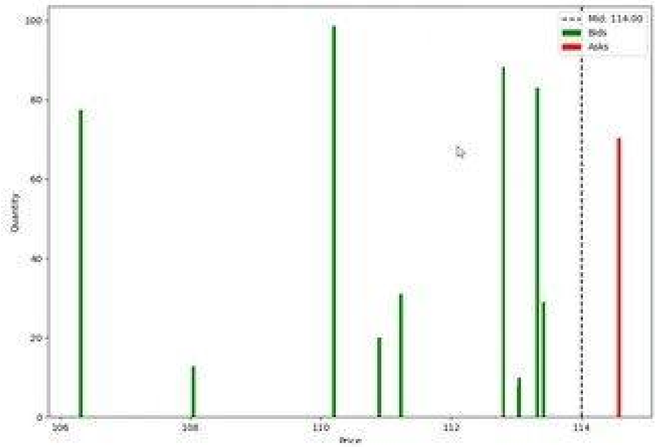

Muxin Feng

Project Responsibilities Lead To Meaningful Growth

One of the most impactful aspects of the Master of Financial Mathematics (MFM) program at NC State University so far has been the opportunity to work on applied projects that closely reflect real world financial risk management practices. While many quantitative finance concepts can initially seem theoretical, the program emphasizes connecting mathematical models with practical financial applications. Through coursework and collaborative projects, I have begun to see how statistical models, economic scenarios, and financial data are combined to support decision making in the banking industry.

A recent group project in the program focuses on building a simplified framework that replicates parts of the Comprehensive Capital Analysis and Review (CCAR) process used by large financial institutions. The goal of the project is to understand how banks evaluate their financial resilience under different economic stress scenarios. CCAR is an important regulatory framework in the United States that requires banks to demonstrate that they have sufficient capital to withstand adverse economic conditions. By studying this process, we are able to see how quantitative models are used in practice to forecast revenue, estimate losses, and project capital levels during economic downturns.

Within our group project, my primary responsibility has been working on the Pre-Provision Net Revenue (PPNR) forecasting component of the framework. PPNR represents the revenue generated from a bank's core operations before accounting for loan loss provisions and is an important input in stress testing and capital planning. In the context of CCAR, forecasting PPNR helps estimate how much income a bank may generate to absorb potential losses during periods of economic stress.

My role has involved collecting and organizing historical financial data and exploring ways to construct a model that can capture patterns in revenue over time. This process has required careful attention to data preparation, since financial statement data must be structured consistently before it can be used for modeling. Through this work, I have gained practical experience in cleaning datasets, identifying relevant financial variables, and thinking about how historical information can be used to forecast future performance.

Even though the modeling process is still in development, working on this component has helped me better understand the challenges involved in financial forecasting. Financial data often contains structural changes and economic influences that can affect model performance. As a result, building a reliable model requires both technical knowledge and financial intuition. This experience has helped me appreciate that quantitative finance is not only about applying formulas but also about understanding the economic context behind the data.

Beyond the technical work, the collaborative nature of the project has also been a valuable part of the experience. Each team member is responsible for different parts of the stress testing framework, which means that the final model depends on integrating several different components together. Working in this structure has helped me understand how quantitative teams in financial institutions often collaborate across different modeling areas, such as LGD models as well as PD models. It has also strengthened my communication skills when discussing technical ideas and the visualization within a group.

Experiences like this project have helped clarify my career goals in quantitative finance, particularly in the area of financial risk modeling and model validation. I am especially interested in roles where mathematical models are used to assess financial stability and evaluate risk exposure. Stress testing frameworks such as CCAR demonstrate how quantitative tools can help financial institutions prepare for periods of economic uncertainty. Contributing to these types of models is meaningful to me because they sit at the intersection of computational skills, quantitative analysis, and real world financial decision making.

Overall, one of the most valuable aspects of the MFM program so far has been the opportunity to work on projects that bridge the gap between theory and industry practice. By applying mathematical and statistical techniques to real financial problems, I am beginning to see how quantitative finance is used in the real world to support risk management and strategic decision making. As I continue progressing through the program, I look forward to developing stronger modeling skills and gaining a deeper understanding of how financial institutions use quantitative methods to navigate complex and uncertain markets.

Ekaterina Finiutina

Building Predictive Trading Signals: Lessons From My Alpha Research Project

Working with real financial data is very different from working with textbook examples. That was probably the biggest thing I took away from my alpha research project this semester, where I had the chance to build and test predictive signals for currency markets using high-frequency trading data.

The project centered on foreign exchange, and I worked with 1-minute price data across five currency pairs over a two-year period. The dataset had over 3 million data points. My job was to engineer features from raw price and volume data that could predict future price movements, then rigorously test whether those features actually held up on data the model had never seen.
On paper, that sounds straightforward. In practice, it was anything but.

The hardest constraint I had to respect was making sure no feature used future information. Every calculation at a given point in time could only draw on data available up to that exact moment. That sounds obvious, but when you are working with rolling windows, lagged variables, and forward-shifted return targets all at once, it is very easy to introduce subtle data leakage without realizing it. I spent more time auditing my feature construction pipeline than I did building it.

Beyond the technical challenge, I also had to make judgment calls about the data itself. High-frequency data is messy. Prices sometimes spike in ways that look like errors but are actually real market events. Deciding when to clean something out versus keep it requires you to understand what is actually happening in the market, not just what the numbers say.

The results were encouraging. My top features generalized across most of the currency pairs I tested, and the ensemble model produced out-of-sample performance that research literature suggests can be meaningful in real trading strategies. One pair did not perform nearly as well as the others, and that turned out to be just as instructive. It reinforced something I think is easy to overlook: markets are not uniform, and a method that works in one place will not automatically work everywhere.

What I valued most about the project was the end-to-end nature of it. I went from raw data to a validated, tested result, and every step along the way required both technical skill and practical judgment. That combination is exactly what I want to keep developing as I move toward quantitative roles after graduation.

Paras Gupta

Portfolio Optimization: Bridging Mathematical Theory and Wall Street Reality

Every day, quantitative analysts use complex mathematical models to move billions of dollars, balancing the razor's edge between maximizing profit and mitigating disastrous risk. This spring, in OR504: Introduction to Mathematical Programming, I got a firsthand look at exactly how they do it.

The course covered the width and depth of optimization techniques, from the fundamentals of Linear Programming to advanced topics like Mixed-Integer Linear Programming, Metaheuristics, and Simulated Annealing. As a Financial Mathematics student, it was a revelation to see how these theories are directly implemented in asset allocation and portfolio management using the SAS programming language.

For a hands-on learning experience, my classmates and I built an equity and bond portfolio. Our objective was clear: maximize returns while minimizing risk, subject to constraints on diversity, bond proportions, and asset allocation. The key challenge we faced was formulating the risk constraint. If we modeled risk using the standard covariance matrix:

$$\sum_{i=1}^{n}\sum_{j=1}^{n} w_i w_j \sigma_{ij} \le R^2$$

Our model would become a complex, non-linear quadratic programming problem. In the real financial world, computational speed equals money. A non-linear model might take minutes or hours to optimize across thousands of assets, whereas a linear model can be solved in milliseconds using powerful industry solvers like Gurobi or IBM CPLEX.

To keep our model linear and fast, we utilized the Capital Asset Pricing Model (CAPM) to estimate portfolio returns as a linear function of market returns. Under CAPM, the total risk of an individual asset is formulated as:

$$\sigma_i^2 = \beta_i^2 \sigma_m^2 + \sigma_{\epsilon_i}^2$$

Because we implemented a strict diversification constraint, we could confidently assume that the non-systematic (company-specific) risk (i2) would be diversified away. This allowed us to transform risk into a simple, highly efficient linear function based entirely on the portfolio's Beta:

$$\beta_p = \sum_{i=1}^{n} w_i \beta_i$$

This approach elegantly simplified our model, but we had to remain transparent about its caveats. Firstly, transforming a non-linear risk system into a linear one relies entirely on the diversification assumption. If our portfolio failed to diversify sufficiently, the model would underestimate the actual risk. Secondly, relying on individual stock betas makes the optimization extremely sensitive to how accurately those beta values are calculated.

Beyond our specific project, realizing the sheer variety of industry applications for these techniques was eye-opening:

- Linear Programming (LP): Used for Asset-Liability Matching. Pension funds and insurance companies use LP to construct bond portfolios where coupons and maturities perfectly match future payout liabilities, minimizing setup costs. Quantitative hedge funds also use LP for Arbitrage Detection, instantly identifying if a sequence of currency or bond trades will result in guaranteed profit.
- Multi-Criteria Decision Making (MCDM): Essential for ESG Investing. Modern portfolio managers must balance conflicting criteria—like maximizing returns while maintaining strict Environmental, Social, and Governance (ESG) scores. MCDM prevents "greenwashing" by mathematically forcing managers to balance these dimensions rather than just guessing.
- Discrete Optimization: Vital for Capital Budgeting (the classic Knapsack problem) and Index Tracking. For an index-driven portfolio (like tracking the S&P 500), buying all 500 stocks incurs massive transaction costs. Portfolio managers use cardinality constraints to select exactly 50 representative stocks that track the index's returns most accurately.
- Metaheuristics: Used for Stress Testing and Scenario Generation. Real-world optimization problems often become too massive and mathematically messy for traditional solvers. Nature-inspired algorithms intelligently search through millions of market shock combinations (e.g., oil prices crashing while inflation spikes) to find the specific sequence of events that causes maximum financial damage, allowing banks to successfully hedge against it.

Working on a practical project using the tools and techniques I learned in OR504 was a phenomenal experience. I now feel fully equipped and confident to apply my mathematical programming skills in the fast-paced world of modern portfolio management.

Chinna Gurram

Understanding Financial Markets Through The Shape Of Data

One aspect of the Financial Mathematics program that I have particularly enjoyed is the exposure to mathematical tools that extend beyond traditional finance methods. Many classical models in finance rely heavily on statistical assumptions and distributional properties. While these tools are powerful, financial markets are complex systems where structural changes often occur in ways that are difficult to capture with standard models.

This realization led me to explore Topological Data Analysis (TDA) as a potential tool for studying financial markets. TDA is a relatively new area that comes from algebraic topology and focuses on understanding the shape of data. Instead of analyzing only averages, variances, or correlations, TDA examines how data points form geometric structures in high-dimensional spaces.

Financial markets generate large datasets that naturally lend themselves to this kind of analysis. Asset returns, volatility measures, and macroeconomic variables can be viewed as points in a high-dimensional space. When market conditions change, the structure formed by these points also evolves. TDA provides a framework for quantifying these structural changes.

One of the most commonly used tools in TDA is persistent homology. In this approach, observations are treated as points in a point cloud. A geometric object called a simplicial complex is then constructed by connecting nearby points at different distance thresholds. As the threshold increases, clusters and loops appear and disappear within the data structure.

Persistent homology tracks the lifespan of these features and summarizes them in visual representations known as persistence diagrams or barcodes. These diagrams provide a compact way to describe the topology of the data. When markets are relatively stable, the topological structure tends to remain consistent. During periods of stress or transition, however, the structure often changes significantly.

What I find particularly interesting is how this perspective differs from traditional econometric approaches. Many models assume a specific statistical form for market dynamics. TDA, by contrast, is largely model-agnostic. It allows researchers to detect structural patterns without imposing strong parametric assumptions.

Through the Financial Mathematics program, I have been exploring how this approach could be applied to regime detection in financial markets. One possible implementation is to construct rolling windows of asset returns or factor exposures and treat each window as a point cloud. By computing persistent homology across these windows, we can track how the geometric structure of the market evolves over time.

The idea is that significant changes in topological features may serve as early indicators of regime shifts, such as transitions from low-volatility environments to periods of market stress. This approach could complement more traditional techniques such as volatility models, factor analysis, or Hidden Markov Models.

My broader career goal is to work in quantitative finance, where developing robust models for complex financial systems is essential. Exploring tools like TDA has reinforced the importance of interdisciplinary thinking. Some of the most interesting advances in quantitative finance occur when ideas from mathematics, statistics, and computer science intersect.

The Financial Mathematics program has provided an environment where exploring these connections is encouraged. Studying advanced mathematics while also applying those concepts to financial data has been one of the most rewarding aspects of the program so far. Approaches like Topological Data Analysis illustrate how mathematical ideas that initially appear abstract can ultimately provide new perspectives on real financial problems.

Dhrubojeet Haldar

Mean Reversion In Motion: Building An FX Statistical Arbitrage System From Scratch

Statistical arbitrage sits at the intersection of econometrics, signal processing, and risk management - and building one end-to-end has been one of the most technically demanding and rewarding experiences of the FM program.

The project targets the foreign exchange market. Using 24 years of tick-level data (2001–2025) across 14 major and cross currency pairs resampled to nine timeframes (from 1-minute to daily), the goal is to identify pairs whose price relationship is mean-reverting and systematically trade the deviations.

Pair Selection: Starting with Cointegration

The foundation is the Engle-Granger cointegration test. Two price series P_1 and P_2 are cointegrated if their spread of log prices $S_t = P_2 - \beta \cdot P_1$ is stationary, even though each series individually follows a random walk. The hedge ratio β is estimated via OLS regression on price levels, and stationarity is verified through ADF and KPSS tests. The Hurst exponent serves as an additional filter - values below 0.5 confirm mean-reverting behavior.

A key metric is the half-life of mean reversion, derived from the Ornstein-Uhlenbeck process: $\Delta S_t = \lambda \cdot S_{t-1} + \varepsilon_t \rightarrow \text{half-life} = -\ln(2) / \lambda$

Pairs are accepted only if the half-life falls within practical trading bounds (e.g., 2–35 bars at daily frequency). Rolling monthly calibration ensures the hedge ratio and half-life remain valid over time - a critical design choice, since static estimates deteriorate quickly in FX.

Signal Generation and Risk Management

Spread deviations are normalized as z-scores using a rolling mean and standard deviation: $z_t = (S_t - \mu) / \sigma$

Entry triggers when $|z| > 2.0$. Positions are closed when $|z| < 0.5$ (profit target), $|z| > 3.5$ (stop-loss), or the hold duration exceeds 2.5× the half-life. A consecutive-loss counter suspends trading after repeated drawdowns, preventing the system from compounding losses during regime breaks. Position sizing uses inverse-volatility weighting, clamped between 5% and 15% of notional per pair.

Extending to Kalman Filtering

The most technically involved extension replaces static rolling regression with a 2-state Kalman filter that tracks the hedge ratio β and intercept α in real time: $\log(P_{2,t}) = \beta_t \cdot \log(P_{1,t}) + \alpha_t + \varepsilon_t$

The process noise parameter δ is tuned per interval (e.g., $\delta = 5\times10^{-5}$ for 1H, $\delta = 5\times10^{-2}$ for 1D) to balance responsiveness against stability. A multi-lock state machine gates live trading behind paper-trade validation, month-long stability checks, and portfolio heat limits.

Results

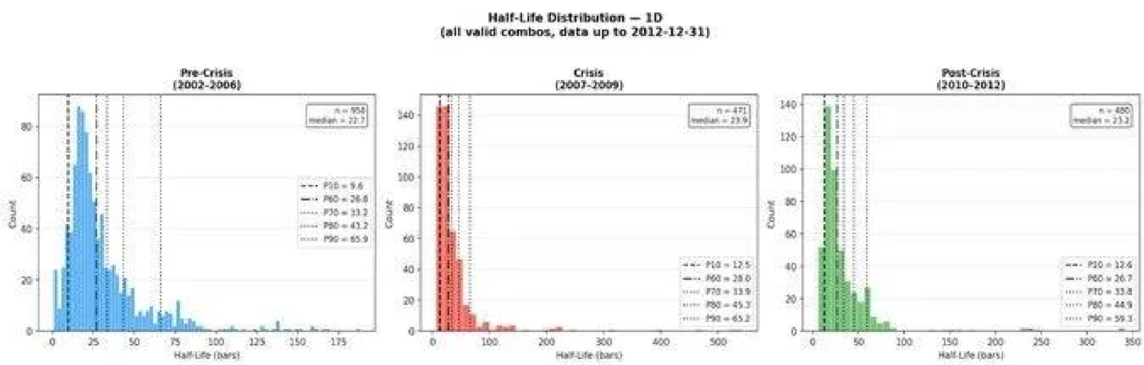

Figure 1. Half-life distributions at the 1D interval across pre-crisis (2002–2006), crisis (2007–2009), and post-crisis (2010–2012) periods. Median half-life remains stable (~21–23 bars), but the distribution widens significantly during stress.

The half-life distributions reveal a key structural property of the strategy: mean reversion speeds are consistent in stable regimes (median ~21–23 bars at daily frequency) but the distribution widens sharply during stress periods such as the 2007–2009 financial crisis. A wider distribution means some pairs revert slowly or not at all, directly degrading signal reliability. This is why static parameter calibration fails in live trading—the half-life used to set hold limits and position sizing must be re-estimated continuously as market regimes shift.

Continue Next Page

Dhrubojeet Haldar

Mean Reversion In Motion: Building An FX Statistical Arbitrage System From Scratch

Continued From Previous Page

These results highlight a recurring challenge: strategies that perform well in-sample often degrade sharply out-of-sample, particularly during the 2020 COVID-19 dislocation. The half-life distributions reveal that mean reversion speeds cluster tightly in stable regimes but widen dramatically during stress - a structural insight that directly informs parameter calibration.

Conclusion

This project forced a reckoning with the gap between academic finance and production reality. Statistical significance in-sample is not alpha. The harder engineering problems - rolling calibration, regime detection, dynamic position sizing, and state machine-based risk overlays - are where the edge (or lack thereof) actually lives. And the most important line in any backtest output is the max drawdown, not the Sharpe ratio.

Aahan Jain

Betting Against Beta

The Idea: When Low-Risk Beats High-Risk

One of the most counterintuitive findings in modern finance is that low-beta stocks — those that move less than the broad market — have historically delivered better risk-adjusted returns than their high-beta counterparts. This anomaly, first rigorously documented by Andrea Frazzini and Lasse Heje Pedersen in their landmark 2014 paper, forms the foundation of the Betting Against Beta (BAB) factor. For our team's industry project, this insight became the starting point for building a market-neutral trading strategy from scratch.

The economic intuition behind BAB is elegant. Many institutional investors face constraints — whether regulatory, mandated, or self-imposed — that prevent them from using leverage. To chase higher returns, these investors overload on high-beta assets, inflating their prices and suppressing their future returns. Meanwhile, low-beta assets sit relatively undervalued. The BAB strategy exploits this mispricing by going long leveraged low-beta assets and short de-leveraged high-beta assets, creating a portfolio that is beta-neutral by construction.

Our Methodology: Translating Theory into Practice

Turning an academic paper into a functioning strategy is harder than it sounds. Our team began by estimating rolling beta for each stock in our equity universe using a combination of one-year and five-year return windows — a dual-horizon approach recommended by Frazzini and Pedersen to balance responsiveness with stability. We then ranked stocks into deciles, constructing a long portfolio of low-beta names and a short portfolio of high-beta names.

The critical — and most technically demanding — step is the leverage and de-leverage mechanism. The long portfolio is scaled up using leverage so that its weighted beta equals one, while the short portfolio is scaled down so its weighted beta also equals one. When combined, the net portfolio beta is zero, making the strategy indifferent to broad market moves. The returns, in theory, come purely from the pricing anomaly itself.

Challenges Along the Way

The gap between theory and implementation has been one of the most instructive parts of this project. Beta is not static — it drifts considerably over time, requiring frequent portfolio rebalancing. Each rebalance introduces transaction costs that can erode the strategy's edge, forcing us to carefully calibrate rebalancing frequency against turnover costs.

We also confronted the challenge of beta estimation itself. Thin trading, non-synchronous returns, and regime changes can all distort beta calculations. Early in our backtesting, we found that naive beta estimates led to a portfolio that was far from truly market-neutral — its realized beta drifted away from zero during periods of market stress, precisely when neutrality matters most. Refining our estimation approach, incorporating Vasicek shrinkage to pull extreme betas toward the cross-sectional mean, meaningfully improved stability.

What This Project Has Taught Me

Beyond the technical skills — coding factor models, backtesting frameworks, and portfolio construction pipelines — this project has deepened my appreciation for the craft of quantitative investing. The most valuable lesson has been that a compelling academic idea and a profitable trading strategy are not the same thing. Markets are adaptive, data is messy, and implementation decisions that seem minor on paper can meaningfully alter real-world outcomes.

Working as a team has also reinforced skills I know will carry into my career: communicating technical ideas clearly, dividing complex problems into tractable workstreams, and pressure-testing each other's assumptions. The BAB project is still a work in progress, but the process of building it has already been one of the most formative experiences of my academic career.

Soham Joshi

Building A CCAR-Style Capital Forecasting Framework

One of the most impactful experiences in the Financial Mathematics (FM) program so far has been leading a team project focused on developing a CCAR-style capital forecasting and stress testing framework. The goal of the project is to simulate how a bank's capital position evolves under severe macroeconomic stress, like the methodology used in regulatory stress tests conducted by the Federal Reserve.

The first step of the project involved constructing a proxy bank portfolio using publicly available macroeconomic and financial datasets. Since real bank portfolio data is typically confidential, our objective was to create a portfolio that resembles the structure of a large bank as closely as possible using open data sources. We relied heavily on data from the Federal Reserve Economic Data (FRED) database and other public financial datasets.

A critical modeling decision was to separate the portfolio into retail and wholesale exposures. This distinction is essential because these two segments behave very differently during economic downturns. Retail portfolios—such as credit cards and consumer loans—tend to be highly sensitive to unemployment and household income shocks. Wholesale portfolios—such as corporate lending—are influenced more by broader economic conditions and business cycles. Modeling these segments separately allowed us to apply more realistic assumptions when estimating credit losses under stress.

After defining the portfolio structure, the next step was to develop Probability of Default (PD) models for the retail segment. The models were built using macroeconomic variables such as unemployment rates and economic growth indicators. One challenge we encountered was ensuring that the time-series properties of the variables were appropriate for regression modeling. To address this, we applied transformations such as differencing and logit transformations to stabilize the data and improve model interpretability.

Once the PD models were developed, we integrated them into a stress testing framework. Under different macroeconomic scenarios—such as baseline and severely adverse conditions—we forecast default rates and estimated potential credit losses over the stress horizon. These projected losses were then incorporated into a simplified capital forecasting model, allowing us to observe how the bank's capital ratio would evolve during a downturn.

Leading this project has been particularly valuable because it bridges the gap between academic quantitative modeling and real-world financial risk management. Regulatory stress testing frameworks like CCAR require institutions to combine statistical modeling, macroeconomic analysis, and financial interpretation in a cohesive way. Through this project, I gained hands-on experience in building that end-to-end workflow—from data construction and portfolio segmentation to model development and capital impact analysis.

Another key learning from this experience has been the importance of model justification and communication. In regulatory environments, building a model is only part of the process; clearly explaining assumptions, limitations, and methodological choices is equally critical. As the project lead, I focused on ensuring that our modeling decisions were transparent and supported by economic reasoning.

Overall, this project has been one of the most rewarding parts of the program so far. It allowed me to apply concepts from statistics, risk modeling, and financial economics to a realistic banking problem. More importantly, it strengthened my interest in pursuing a career in quantitative risk management and financial modeling, where rigorous analytics are used to understand and manage systemic financial risks.

Michelle Kadiman

Signals Before Stress: A Panel Data View Of Credit Delinquency

Understanding how macroeconomic conditions influence credit risk is a central question in financial risk management. During periods of economic stress, changes in unemployment, inflation, and financial market conditions often precede increases in delinquency rates. Because of this relationship, macroeconomic indicators are frequently used to help explain and anticipate movements in credit delinquency.

This analysis investigates whether macroeconomic variables help explain changes in delinquency rates. Building such a model raises several econometric questions. Which indicators meaningfully capture delinquency dynamics? Are enough variables included to limit omitted variable bias, even when multicollinearity appears low? Timing is also important, as some macroeconomic conditions influence credit performance immediately while others operate with delays. Identifying whether leading, coincident, or lagged indicators provide the strongest explanatory power therefore becomes a key step in model specification.

To address these issues, a fixed-effects panel regression framework is employed. This approach controls for unobserved entity-specific characteristics that remain constant over time. In credit risk data, institutions or portfolios often differ in structural ways that are difficult to observe directly, such as underwriting standards, borrower composition, or portfolio risk profiles. Fixed effects help isolate the impact of macroeconomic fluctuations while accounting for these persistent cross-sectional differences.

The regression results are summarized in (Image1). The model produces a within R^2 of about 0.72, indicating that the macroeconomic variables explain a large share of the time-series variation in delinquency rates within entities. At the 90% confidence level, three variables are statistically significant: CPI inflation, lagged BBB corporate bond yields, and the market volatility index. CPI inflation acts as a coincident indicator and has a negative coefficient, suggesting that higher inflation may reduce the real burden of debt and ease repayment pressure. Lagged BBB corporate bond yields are positively related to delinquency rates, indicating that tighter credit conditions are associated with greater borrower stress. The market volatility index behaves as a leading indicator and tends to rise before broader economic deterioration, which suggests that higher financial uncertainty may precede weaker credit performance.

Overall, the results suggest that macroeconomic conditions play an important role in shaping delinquency dynamics. While borrower-level characteristics remain essential for underwriting and credit assessment, macroeconomic indicators provide valuable information about the broader economic environment affecting repayment behavior. Incorporating macroeconomic signals into credit risk models can therefore enhance the ability of financial institutions and risk managers to anticipate changes in portfolio performance and monitor emerging financial stress.

```
. xtreg Deliquency_rate L2.Unemployment_rate L3.Real_GDP_growth L1.CPI_inflation_rate L4.BBB_corporate_yield Market_Volatility_Index__Level_, fe vce(robust)

Fixed-effects (within) regression               Number of obs     =        198
Group variable: entity_id                       Number of groups  =          2

R-squared:                                      Obs per group:
     Within  = 0.7163                                         min =         99
                                                              avg =       99.0
     Overall = 0.4740                                         max =         99

                                                F(1, 1)           =          .
corr(u_i, Xb) = -0.0000                         Prob > F          =          .

                                 (Std. err. adjusted for 2 clusters in entity_id)
---------------------------------------------------------------------------------
                                |               Robust
                Deliquency_rate | Coefficient  std. err.      t    P>|t|     [95% conf. interval]
--------------------------------+------------------------------------------------
              Unemployment_rate |
                            L2. |   .0007664   .0002787     2.75   0.222    -.002775    .0043079
                                |
                Real_GDP_growth |
                            L3. |  -.0000678   .0000198    -3.42   0.181   -.0003192    .0001837
                                |
             CPI_inflation_rate |
                            L1. |  -.0001682   9.45e-06   -17.79   0.036   -.0002883   -.0000481
                                |
            BBB_corporate_yield |
                            L4. |   .0057653   .0008588     6.71   0.094   -.0051469    .0166775
                                |
Market_Volatility_Index__Level_ |   .0001202   .0000136     8.84   0.072   -.0000526     .000293
                          _cons |  -.0122137   .0024369    -5.01   0.125   -.0431781    .0187507
--------------------------------+------------------------------------------------
                        sigma_u |   .0117044
                        sigma_e |  .00627844
                            rho |  .77655254   (fraction of variance due to u_i)
---------------------------------------------------------------------------------
```

Subham Khinchi

Derivatives Pricing Using The Heston Model

Introduction

Financial derivatives are priced by discounting expected payoffs under the risk-neutral measure Q. The Black-Scholes (BS) model (1973) provides an elegant closed-form solution under the assumption of constant volatility and log-normal asset returns. However, two empirical facts contradict it: (i) implied volatilities extracted from market option prices vary systematically across strikes and maturities, forming the volatility smile or skew; and (ii) realized volatility is heteroscedastic - it clusters and mean-reverts over time. These phenomena are not merely statistical curiosities; they have direct consequences for derivative hedging and risk management. Introduced by Steven Heston in 1993, the Heston model corrects both failures by making instantaneous variance a stochastic process correlated with the asset, while retaining analytical tractability through a semi-closed-form pricing formula derived via the characteristic function.

The Heston Model Dynamics

Under Q, the Heston model evolves the asset price S(t) and instantaneous variance v(t) jointly via two correlated SDEs:

SDE 1: Asset price dynamics

$$dS(t) = r\,S(t)\,dt + \sqrt{v(t)}\,S(t)\,dW_1(t)$$

μ is the risk-free drift, W_1 is a standard Brownian motion

SDE 2: Variance dynamics – mean reverting CIR process

$$dv(t) = \kappa(\theta - v(t))\,dt + \sigma\sqrt{v(t)}\,dW_2(t)$$

Correlation:

$$dW_1(t)\,dW_2(t) = \rho\,dt$$

The five parameters are: κ (mean-reversion speed), θ (long-run variance), σ (vol-of-vol), v_0 = v(0) (initial variance), and ρ (correlation). The Feller condition must hold for v(t) to remain strictly positive almost surely:

$$2\kappa\theta > \sigma^2$$

Ensures v(t) > 0 a.s.; violated in some calibrated parameter sets.

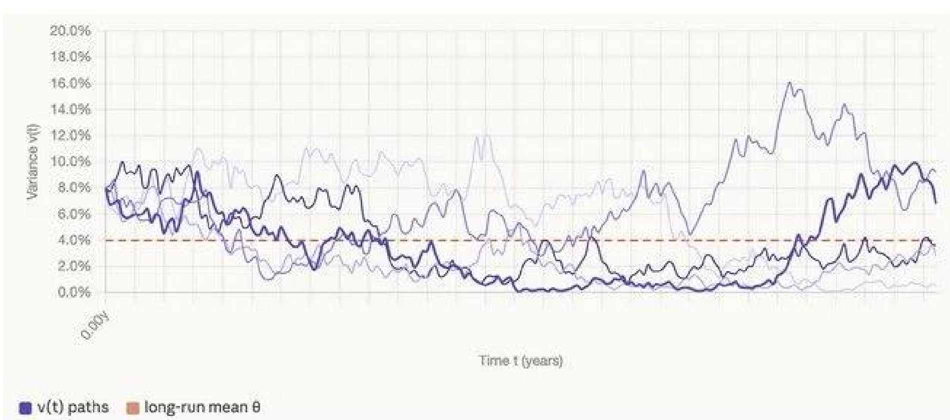

Fig 1 : Simulated variance paths of v(t)

Characteristic Function and Options Pricing

Heston's key insight is a closed-form conditional characteristic function (CF) for X = ln S(T). Letting $u \in \mathbb{R}$ and $\tau = T - t$:

$$\varphi(u;\tau) = \exp exp\left\{C(u,\tau) + D(u,\tau)\,v_0 + iu \ln ln\, S\right\}$$

Continue Next Page

Subham Khinchi

Derivatives Pricing Using The Heston Model

Continued From Previous Page

$i = \sqrt{-1}$; C and D are deterministic functions of u and τ . Where, C : Auxiliary function – drift & mean reversion contribution

$$C(u,\tau) = r\tau ui + \frac{\kappa\theta}{\sigma^2}\left[(\kappa - \rho\sigma ui - d)\tau - 2l!\left(\frac{1-g\,e^{-d\tau}}{1-g}\right)\right]$$

D : Auxiliary function – loading on current variance v0

$$D(u,\tau) = \frac{(\kappa-\rho\sigma ui-d)\left(1-e^{-d\tau}\right)}{\sigma^2\left(1-g\,e^{-d\tau}\right)}$$

d: Complex square root - Branch-cut selection in d for numerical stability

$$d = \sqrt{(\rho\sigma ui - \kappa)^2 + \sigma^2\left(ui + u^2\right)}$$

g: Ratio used in the log term of C(u,τ)

$$g = \frac{\kappa-\rho\sigma ui-d}{\kappa-\rho\sigma ui+d}$$

Given the characteristic function, European call prices follow from Gil-Pelaez Fourier inversion. With strike K, the call decomposes as:

$$C(S,K,\tau) = S\cdot P_1 - Ke^{-r\tau}\cdot P_2$$

P_1 = delta (asset-measure probability);

P_2 = risk-neutral exercise probability;

P_j : Gil-Pelaez Fourier inversion — evaluated via FFT

$$P_j = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty Re\left[\frac{e^{-iu\ln K}\varphi_j(u)}{iu}\right]du$$

Evaluated via FFT (Carr-Madan, 1999) in O(N log N) across all strikes.

Implied Volatility Surface

The Heston model generates a two-dimensional implied volatility surface over strike K and maturity T. At-the-money volatility for large τ converges to √θ, while the short-maturity skew satisfies:

$$\left.\frac{\partial\sigma_{imp}}{\partial k}\right|_{k=0} \approx \frac{\rho\sigma}{2}, \quad \tau\to 0$$

k = ln(K/F); provides a direct calibration anchor for ρ from market data

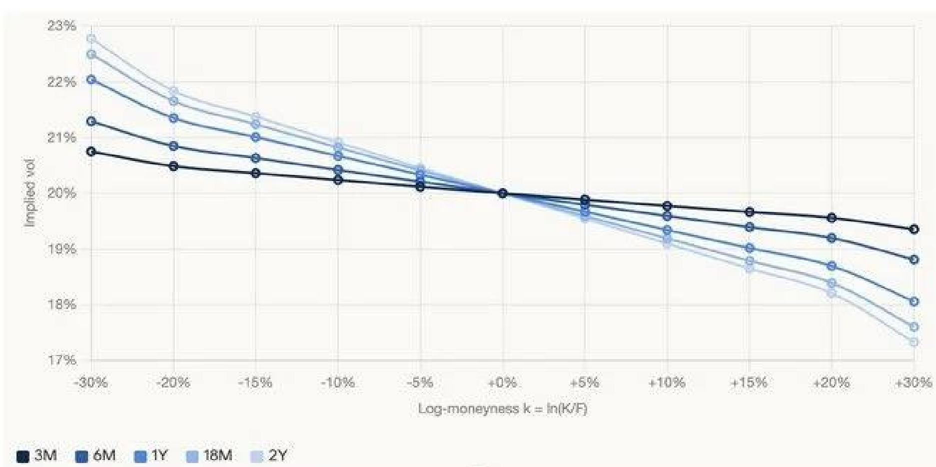

Fig 2: Heston Implied Volatility – strike x maturity

Continue Next Page

Subham Khinchi

Derivatives Pricing Using The Heston Model

Continued From Previous Page

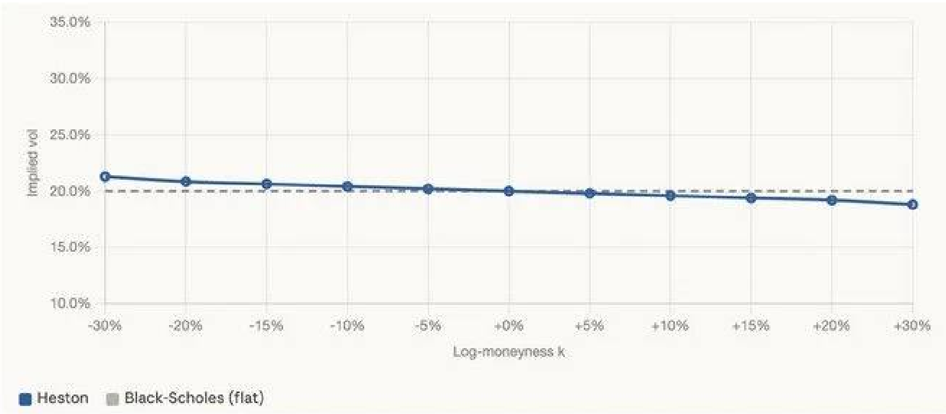

Fig3: Heston vs Black-Scholes implied vol smile

Calibration

Parameters $\Theta = \{v_0, \kappa, \theta, \sigma, \rho\}$ are calibrated by minimizing weighted squared implied-volatility errors across liquid market options:

$$\Theta^* = \arg\min_\Theta \sum_i w_i \left[\sigma_{\text{imp}}^{\text{model}}(K_i, T_i; \Theta) - \sigma_{\text{imp}}^{\text{market}}(K_i, T_i)\right]^2$$

w_i = inverse bid-ask spread; solved via Levenberg-Marquardt or differential evolution

The objective is non-convex, In practice, seed ρ from the observed skew and θ from long-run realized variance to improve convergence.

References

Heston, S. L. (1993). A Closed-Form Solution for Options with Stochastic Volatility. Review of Financial Studies, 6(2), 327–343.
Black, F. & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. Journal of Political Economy, 81(3), 637–654.
Carr, P. & Madan, D. (1999). Option Valuation Using the Fast Fourier Transform. Journal of Computational Finance, 2(4), 61–73.

Burke King

From Probability to Policy: An Actuarial Career Vision

At this moment, my ultimate career goal is to become a senior leader in the property and casualty insurance industry, holding a title such as Senior Vice President. I am motivated by the opportunity to apply quantitative analysis, financial theory, and business judgment to manage risk, guide strategic decision-making, and contribute to the long-term financial stability of an organization. I am pursuing a career as an actuary because it encompasses the intersection of mathematics, finance, and real-world applications. Actuaries play a critical role in evaluating uncertainty, pricing risk appropriately, and ensuring that insurers remain solvent while serving policyholders effectively. An actuary places emphasis on analytical rigor, ethical responsibility, and continuous learning, which align strongly with my long-term professional aspirations.

To prepare for this career path, I have made significant progress toward professional actuarial credentials. I have passed Exams P and FM, which provided a strong foundation in probability theory, financial mathematics, and risk modeling as I work towards earning my ACAS and FCAS credentials. I have completed the required VEE coursework in accounting, finance, and economics. These subjects demonstrate my commitment to the actuarial profession and my readiness to succeed in increasingly complex roles. I am particularly interested in working in the property and casualty insurance industry because of its dynamic and evolving risk landscape. Property and casualty insurers must account for catastrophe risk, climate-related volatility, inflation, regulatory changes, and emerging exposures. This environment requires actuaries to think creatively, adapt models to new data, and collaborate closely with underwriting, claims, and senior management. I am drawn to this challenge and to the opportunity to see the direct impact of actuarial analysis on pricing, reserving, and risk management decisions.

During my time in the MFM program, my goal has been to secure an actuarial position at a P&C insurance company where I can continue progressing through actuarial exams while gaining hands-on experience. I aim to develop expertise in areas such as loss reserving, pricing, capital modeling, and enterprise risk management. I also intend to build strong communication skills so that I can clearly explain technical results to non-technical stakeholders, an ability that is essential for effective actuarial leadership. I have accepted a position with Allianz Commercial as an Associate Graduate Global Pricing Actuary at their Alpharetta, Georgia location, where my hope is that I can further my experience at a property and casualty company, as I did during my 2024 summer internship with The Cincinnati Insurance Companies. During this time, I also hope to receive support as I progress further through examinations with the goal of becoming a credentialed actuary.

Over the long term, I aspire to move into executive leadership, ultimately serving as a Senior Vice President. In this role, I would seek to influence company-wide strategy, mentor younger actuaries, and help guide the organization through complex risk and financial decisions. I believe that a strong actuarial background provides an excellent foundation for executive leadership because it combines analytical discipline with a deep understanding of risk, profitability, and long-term sustainability. My career goal is to build a successful actuarial career in the property and casualty insurance industry and to grow into a senior leadership role. By continuing to develop my technical expertise, business acumen, and leadership skills, I am confident that I can make meaningful contributions to an insurance organization and to the actuarial profession.

Jake Kostoryz

Where Sports, Markets, And Data Science Collide

When I arrived at NC State to begin the MFM program, I thought I was leaving sports behind. After four years of competing as an NCAA football player at Lake Forest College, I figured the next chapter of my life would be entirely about markets, math, and quantitative finance. But over time, I realized the competitive mindset and instincts I built through sports never really went away. They just started showing up in a different place. As I got deeper into the program, I began to see a niche that brought everything together: sports, markets, and data science.

The discovery happened pretty naturally. Like a lot of people who love sports, I got into sports betting. At first it was casual, just trying to make smarter picks than my friends. But the more I dug into it, the more I realized there was real quantitative depth beneath the surface. Then a classmate in the program mentioned sports trading to me, and something clicked. This wasn't just gambling with extra steps. There are actual firms and desks out there that treat sports markets with the same rigor that traditional quant shops apply to equities or derivatives. Susquehanna International Group runs a dedicated sports trading unit called Nellie Analytics, and firms like DRW, Jump Trading, and AQR have all built or are building prediction markets and sports trading teams. On top of that, the growth of platforms like Kalshi and Polymarket is accelerating this whole trend, creating structured ways to trade on event outcomes and attracting serious quantitative talent. The infrastructure for this space is expanding fast.

What excites me the most is how naturally the skills we're developing in the MFM program translate to this space. The probability theory, stochastic modeling, and statistical methods we study every day are directly applicable to pricing sports outcomes. When you think about it, modeling the probability of a game outcome isn't all that different from pricing an option. You're estimating probabilities under uncertainty, managing risk, and looking for edge where the market has it wrong. The data science side of it is massive too. Sports generate an enormous amount of real-time, messy data, and turning that into something actionable is exactly the kind of problem I love working on.

My background as a football player is a real advantage here too. Having played the game, I understand it from the inside, and that helps me know which features and variables actually matter when I'm building a model. I know what a momentum shift looks like, how matchup dynamics change a game, and which factors drive outcomes that don't always show up in a box score. That kind of domain knowledge is critical when you're deciding what data to feed a model and what to leave out.

To be clear, my career goals extend beyond just sports trading. I'm deeply interested in quantitative finance broadly, and I've spent a lot of time this year working on projects in options pricing, volatility modeling, and commodity risk management. All of that has shaped my understanding of markets and reinforced my passion for finding quantitative edge. But sports trading represents something special to me, and I've already started exploring it hands-on. I've been working on building a trading strategy and modeling games in MLB, trying to find where the data tells a different story than the market. It's still early, but the process of digging into player stats, matchup data, and historical trends to generate predictions feels like exactly the kind of problem I want to spend my career solving.

The MFM program has been instrumental in helping me see all of these connections. Before coming to NC State, I knew I wanted to work in quantitative finance, but I didn't appreciate just how broad the field really is. Discovering that sports trading is a legitimate and growing niche has been one of the most exciting realizations of my time in the program. I don't know exactly where this path leads, but the combination of sports knowledge, quantitative skills, and genuine passion is a strong foundation. Getting to work at the intersection of sports and trading every day would be a dream.

Ted Kratt

Equity Analysis 101

Throughout my undergraduate summers, I interned at three different firms. While the companies changed, my primary duty was always answering the same core question: Should we buy this stock? After completing 20+ stock write-ups that involved dissecting different business models, foreseeing market opportunities, analyzing financial statements, and evaluating competitors, I've distilled those experiences into this tutorial on the essentials of equity analysis.

The first key to keep in mind when doing equity analysis is the audience you are presenting to. In my case, it was usually a portfolio manager who was short on time and wanted to understand things quickly. This shaped how I formatted my analyses, with the first portion being a high-level "hook" that included a summarization of their business, price, market cap, enterprise value, and important valuation ratios like price to earnings(P/E), price to earnings growth(PEG), and earnings per share(EPS). This quick and dirty breakdown gives a lot of key information to busy managers who can decide if they want to keep reading.

The next part of the firm I would analyze is the business model and segment analysis. This involves looking at revenue streams to understand where their sales are coming from, which segments are the most profitable, and at what rate they are growing. An integral part of this process is finding any "moat" that the company may possess. This could be a unique product or service that no one else in their industry provides and would not easily replicate.

Another important part of understanding how a business operates and how it will perform going forward is having a firm grasp of the market they participate in. Do they have a large total addressable market (TAM), which is the combined sales of all of the market participants, and do they have room to expand in their TAM, or do they already take up a large portion of the market? Evaluation of their TAM, paired with knowledge of future product or service releases have a huge impact on the growth potential of a company.

After understanding the growth potential of the company, the next step of the evaluation process is looking at its financial health through financial statements. This includes the income statement, balance sheet, and cash flow statement. Important metrics to focus on in the income statement are the earnings before interest, taxes, depreciation, and amortization (EBITDA) and operating margins, as well as EPS growth trends. In the balance sheet, we want to look for red flags related to ballooning debt. In the cash flow statement, we want to see strong free cash flow (FCF) generation to prove the company is growing organically. These metrics help us understand the profitability of the firm and make sure it is not drowning in debt that it cannot repay.

Next, it is important to understand how a company's competitors compare with respect to size, product availability, and growth. If there is a larger firm that is directly competing with the company you are analyzing, and they have a better product, it is probably not a good idea to invest. Another framework that helps evaluate how a company may perform is looking at Porter's five forces. This structure breaks down the power of the firm's buyers, sellers, suppliers, and the threat of substitutes. For example, a positive grade would have buyers who don't have a lot of power and products with a low threat of substitutes.

The final evaluation element is a risk assessment and rationalization of valuation. Being aware of risks in the market and risks that are directly related to the firm itself helps project future performance. Justifying the current valuation and the future valuation that is expected, given the growth catalysts we found, is also very important to have a strong conviction of future growth.

Ultimately, after completing these steps, I can confidently provide a buy, hold, or sell decision. I include this assessment in my initial "hook" and ensure every claim is backed by rigorous quantitative evidence. While an analyst may not always be right, thoroughly checking edge cases and mastering business fundamentals provides the best possible chance of a correct assessment.

Daksh Kumar

Navigating the Volatility Smile: A Unified Framework for Advanced Pricing Models

Project Overview and Objectives

As part of the Master of Financial Mathematics (MFM) program at NC State, my FIM 601 project focuses on developing a unified calibration, pricing, and evaluation framework for various volatility models. Modern options markets exhibit pronounced volatility smiles, term structure effects, and non-trivial dynamics in implied volatility. While the foundational Black-Scholes framework assumes constant volatility, empirical option prices require more expressive models to accurately capture these market phenomena.

Our central objective is not merely to fit prices, but to understand the trade-offs between static fit, dynamic realism, computational tractability, and hedging performance. We evaluate three of the most widely used frameworks: Local Volatility (LV), Stochastic Volatility (SV), and Stochastic-Local Volatility (SLV) models. By assessing these models' side-by-side on identical datasets, the project clarifies why multiple volatility models coexist in practice and dictates when each is most appropriate for real trading and risk systems.

Quantitative and Analytical Techniques

The methodology begins with constructing an arbitrage-consistent implied volatility surface using the Stochastic Volatility Inspired (SVI) parameterization. From this constructed surface, we extract a Dupire local volatility surface and implement a finite-difference numerical prices to handle the computations.

Following the local volatility implementation, we calibrate a stochastic volatility model of the Heston type to the same market data. Finally, we construct a hybrid stochastic-local volatility model. This hybrid approach is designed to combine the local calibration accuracy inherent in the LV model with the stochastic smile dynamics provided by the SV framework.

Technology, Data Sources, and Challenges

This project is implemented entirely in Python, relying on its robust ecosystem for numerical optimization, volatility surface construction, and building the finite-difference solvers. Our market data, consisting of real-world option prices and implied volatilities, is sourced directly from Bloomberg and Kaggle.

One of the primary engineering challenges we encountered was ensuring the stability of the recovered volatility surfaces while simultaneously optimizing the computational efficiency of the finite-difference solvers. Evaluating these distinct models on identical datasets requires rigorous numerical precision and careful parameter tuning to avoid unstable pricing dynamics.

Career Goals and the MFM Program

Prior to joining NC State, my background involved working as an AI Engineer and building deep learning-based portfolio optimization frameworks. While I gained substantial experience optimizing high-dimensional spaces, I realized that translating computational intuition into robust financial models requires a much deeper mathematical foundation.

I chose the MFM program to bridge the gap between advanced data science and quantitative finance. The rigorous training in the MFM program, coupled with hands-on projects like this one in FIM 601, directly supports my career aspiration of becoming a quantitative researcher. By applying my computational background to complex derivatives pricing, risk modeling, and systematic trading strategies, I am building the exact skill set needed to develop innovative, data-driven solutions in today's evolving financial markets.

Jiayuan Li

Expanding My Path in Quantitative Risk at NC State

When I first arrived at NC State, I knew I wanted to pursue a career in quantitative finance, but my understanding of “risk” was still quite general. I chose NC State because of its strong emphasis on quantitative training and practical projects that connect statistical theory with real financial applications. During my first semester, I explored credit risk through projects such as behavioral Probability of Default (PD) modeling. Entering my second semester, my focus has expanded toward market risk and volatility modeling, allowing me to apply statistical methods to real financial time series data.

One of the most valuable projects this semester involved modeling and forecasting S&P 500 volatility using GARCH-type models. The goal of the project was to understand how financial markets exhibit time-varying volatility and how these dynamics can be modeled for risk forecasting.

I began by collecting historical S&P 500 price data and transforming it into daily log returns. Exploratory analysis quickly revealed clear volatility clustering, where large market movements tend to occur in groups rather than independently. This stylized fact motivated the use of conditional heteroskedasticity models.

To model this behavior, I implemented several GJR-GARCH models with different lag structures. I estimated four candidate specifications: GJR-GARCH(1,1,1), GJR-GARCH(2,1,1), GJR-GARCH(1,1,2), and GJR-GARCH(2,1,2). Using the Bayesian Information Criterion (BIC), I selected GJR-GARCH(1,1,1) as the best model because it achieved the lowest BIC while maintaining a relatively simple structure.

The estimated parameters revealed several important characteristics of market volatility. The persistence parameter was large and highly significant, indicating that volatility shocks decay slowly over time. The asymmetric parameter was also significant, confirming the leverage effect, where negative returns lead to larger increases in volatility than positive returns of the same magnitude. In addition, modeling the residuals with a Student-t distribution allowed the model to capture the heavy tails commonly observed in financial returns.

To validate the model, I conducted several diagnostic tests. Standardized residual analysis and Ljung–Box tests indicated that the model successfully removed most serial correlation and volatility clustering from the residuals. I then implemented a rolling-window volatility forecasting framework using a 500-day estimation window to generate one-step-ahead volatility forecasts. Comparing the forecasts with realized volatility (measured using a 5-day rolling standard deviation) produced a correlation of 0.91, indicating that the model tracks changes in market risk effectively.

Working through this project helped me see how theoretical concepts from time series analysis translate into practical tools used in financial risk management. It also strengthened my programming skills in Python and deepened my understanding of how volatility forecasts support applications such as risk monitoring, portfolio management, and derivative pricing.

Outside the classroom, I have also made an effort to learn from professionals working in quantitative risk roles. Through conversations with industry practitioners, including quantitative analysts at financial institutions such as Bank of America and EY, I gained a clearer picture of how models are used in practice. These discussions highlighted that beyond predictive performance, factors such as data quality, interpretability, and the ability to communicate results to non-technical stakeholders are critical in real-world risk management.

As I continue my studies, my career goal is becoming clearer: I want to work as a quantitative risk analyst, focusing on modeling and forecasting financial risks across credit and market domains. I am particularly interested in combining traditional financial econometrics with modern machine learning methods to build models that are both accurate and interpretable for decision-making.

Through projects like this, NC State has helped me move from a general curiosity about quantitative finance toward a more focused direction in quantitative risk modeling and financial analytics.

Lambert Li

A Conversation With Dimitri Bianco

A Conversation with Dimitri Bianco

A recent conversation with Dimitri Bianco, who works in quantitative risk and model development, turned out to be very different from what most people would expect. Usually, when someone talks to an expert in this field, they expect to hear about complicated math or the advanced models. Instead, Dimitri focused on simple habits and a practical mindset. His advice made the professional world feel much less like a mystery and more like a place where being disciplined and aware of others really pays off.

The Value of Writing Things Down

One of the most important things Dimitri talked about was the habit of taking notes. Even though quant finance is a very digital world, he still keeps paper on his desk and writes things down during every meeting. This might sound like a small, old-fashioned thing to do, but it makes a huge difference in the long run. In an internship or a new job, there are always a lot of tiny details to keep track of — like which specific dates to use for a data set or exactly how a manager wants a report to look.

Dimitri pointed out that it is very easy to forget these details, and one of the biggest mistakes a new person can make is coming back to ask the same questions over and over. By writing everything down, a person shows they are prepared and that they respect the time of the people they are working with. It is a humble habit that proves someone is actually paying attention.

Getting Communication Right

When working in a remote or hybrid environment, how someone communicates says a lot about their professionalism. Dimitri explained that it is important to know which tool to use for which task. For example, internal messaging apps are great for quick, casual questions, but emails should be used for formal requests or things that need to be saved as an official record.

He also touched on the pressure people often feel when working from home. There is a common fear that if a person isn't "active" on their computer every single second, people will think they aren't working. Dimitri’s advice was to remember what a real office is like. In a normal office, people get up to get water, talk to a neighbor for a minute, or take short breaks to think. Remote work should be the same. The goal isn't to stay "green" on a chat app all day; it is to get the work done well and stay responsive when it matters.

Understanding Why the Work Matters

Finally, Dimitri shared a very grounded perspective on math and modeling. It is easy for students to focus only on making a model as accurate as possible. However, in the real world, a model is just a tool to help a business make a decision.

Dimitri believes it is more important to ask how the business is going to use the model. Does the team need something simple that they can easily explain to a client, or do they need something incredibly detailed for daily risk management? Asking these kinds of questions shows that a person is thinking about the "big picture" and not just the equations. It shows a genuine interest in being helpful to the company.

Summary

The biggest takeaway from the conversation with Dimitri was that success in a career isn't about being the smartest person in the room. It is about the simple stuff: being organized, communicating clearly, and being easy to work with. None of this advice was flashy or complicated, but it felt very real and honest. It was a reminder that growing professionally usually comes from building good, simple habits every single day.

Bedant Lohani
Building Toward Alpha

As I think about the next stage of my career, I am increasingly drawn to the intersection of quantitative research, market structure, and scalable investment systems. In the near term, my goal is to build a strong foundation in fixed income architecture. I want to work close to the models, analytics, and infrastructure that support pricing, risk, portfolio construction, and decision-making across rates and credit markets. That kind of work appeals to me because it combines rigorous mathematics with real market impact: the quality of the underlying architecture shapes how well a firm can measure risk, respond to changing conditions, and deploy capital intelligently.

Fixed income feels like the right place to start because it forces precision. It requires understanding term structures, spread dynamics, liquidity, valuation, and the way macroeconomic changes flow through portfolios. Just as important, it teaches discipline. Strong architecture is not only about building models that work in theory; it is about creating systems that are reliable, interpretable, and useful for traders, portfolio managers, and risk teams under real constraints. My upcoming summer role has reinforced this interest by introducing to me how central robust analytics and infrastructure are to institutional investing.

Over time, I hope to transition into systematic active equities. That move interests me because equities bring a different rhythm: richer cross-sectional signals, faster feedback loops, and a broader opportunity to connect data, models, and market intuition. I would like to help develop systematic frameworks that inform idea generation, portfolio tilts, factor exposures, and risk-aware implementation for clients. I see this as a valuable intermediate step because it bridges deep quantitative infrastructure work with more directly research-driven investment processes.

Long term, my ambition is to move to the buy side and focus on alpha generation. Ultimately, I want to work on building and refining strategies that convert research into repeatable investment edge. What excites me most is not just prediction, but the full chain behind it: identifying signals, stress-testing them, understanding when they fail, and embedding them into a disciplined portfolio process. I want to become the kind of investor who can move comfortably between theory and execution, someone who understands both the architecture underneath the system and the alpha it is designed to produce.

At this stage, my career goal is not to rush to the end point, but to build the right sequence of experiences. Fixed income architecture offers the technical grounding. Systematic active equities offer broader research expression and market breadth. The buy side offers the chance to own the investment outcome directly. Taken together, that path reflects the kind of career I want.

Siyuan Lu

The Practical Impact In The MFM Program

As a second-semester student in the North Carolina State University Master of Financial Mathematics program, the most impactful aspect of my experience has been the seamless integration of technical risk management with exceptional career support. The curriculum consistently bridges the gap between quantitative theory and industry application while actively preparing students for the competitive job market. Rather than treating financial mathematics purely as an abstract science, the program emphasizes how these tools function in the real business world.

During my first semester, my coursework immediately provided context for complex mathematical concepts. Classes like Machine Learning in Finance and Commercial Banking Risk Management shifted my perspective from simply calculating numbers to understanding how banks actively evaluate and mitigate risks in a highly regulated environment. I learned how financial institutions rely on probability of default models and data pipelines built in Python to manage credit risk. This foundational knowledge proved essential for moving beyond abstract mathematical theory and tackling actual business problems that companies face every day.

Now in my spring semester, this focus on practical risk application has expanded significantly. I am currently taking Enterprise Risk Management and a Seminar in Financial Mathematics alongside my core technical classes. The combination is highly effective because it forces me to apply quantitative tools to corporate strategies. For example, working on a comprehensive capital forecasting and stress testing project gives me direct exposure to the strict regulatory expectations that large banks navigate. Furthermore, analyzing the Krispy Kreme Case Study allows me to view enterprise-level risks from the perspective of an external advisor. These specific projects teach me how to interpret complex financial simulations and communicate those results clearly to stakeholders.

Beyond the rigorous classroom environment, the dedication to professional development within the program is outstanding. The career support goes far beyond basic resume reviews or interview preparation. Every week, the program invites industry professionals from diverse fields such as investment, data analytics, consulting, and risk management to host seminars on campus. These guest speakers generously share their time to engage with our cohort and provide valuable, first-hand career insights. They discuss the rapidly evolving landscape of their respective industries and highlight the precise analytical skills that drive success within their organizations. Listening to these direct experiences helps connect the quantitative techniques we learn in class to the daily realities of the financial sector. Understanding these different career trajectories has also been instrumental in helping me refine my own professional goals.

This consistent exposure to real-world practitioners, combined with a highly relevant and technical curriculum, makes the program incredibly valuable. The financial industry needs quantitative professionals who understand the broader business impact of their tools and models. By embedding practical risk management and direct industry engagement into the core academic experience, the program ensures we are ready to step into the workforce as strategic problem solvers.

Paritosh Mallick

How Networking Shaped My Journey In Financial Mathematics At NC State

When I joined the Master's in Financial Mathematics program at NC State University, I knew I was signing up for rigorous coursework like stochastic calculus, derivatives pricing, computational methods. What I didn't expect was how much of my growth over the first several months would come not from textbooks, but from conversations with people who had already walked this path.

Two events in Fall 2025 changed the way I think about my career and about myself as an emerging finance professional.

Meeting the People Behind the Profession

The FM Networking Event was my first real introduction to the community surrounding this program. Seniors, alumni, and working professionals all in one room, it sounds simple enough, but the experience was anything but ordinary.

What struck me most was how openly people shared their journeys. Nobody handed me a polished elevator pitch. Instead, I heard real stories about the interviews that didn't go well, the roles that surprised them, the skills they wished they had built earlier. For someone still figuring out where exactly they want to land, that kind of honesty is invaluable.

Several conversations gravitated toward Quantitative Finance, which has become the area I'm most drawn to. Hearing alumni talk about their day-to-day work as quant, the modeling, problem-solving, the intersection of math and markets, made the career feel tangible in a way that course descriptions never quite do.

Getting a Real Look at How Banks Hire

A short while later, I attended the Career Fair hosted by Poole Management, and it opened my eyes to something I hadn't thought deeply about before: the mechanics of how financial institutions recruit.

I had the chance to speak with representatives from Bank of America, PNC, First Citizens, Truist, UBS, and First National Bank. Each conversation was different, but a few themes kept coming up: the value of technical depth in areas like quantitative modeling and data analysis, the importance of being able to communicate complex ideas clearly, and the distinction between different kinds of finance roles within the same institution.

Before the fair, "working at a bank" felt like a vague destination. Afterward, I had a much clearer picture of what specific roles look like, what teams are hiring, and what preparation matters.

What I'm Taking Forward

Halfway through this program, I feel like I'm finally starting to connect the dots. The coursework is building my technical foundation. But it's these networking experiences that are helping me understand where that foundation can take me.

I'm now more deliberate about how I'm spending my time on which skills I'm deepening, which conversations I'm following up on, and which opportunities I'm positioning myself for. The goal of becoming a Quantitative Analyst feels less like a distant aspiration and more like a concrete next step.

If there's one thing the Master's in Financial Mathematics program at NC State has taught me so far, it's this: the technical skills get you in the room. But relationships, conversations, and the willingness to learn from others that's what shapes the kind of professional you become.

Nicholas Moss

Building a Quantitative Finance Career Through Structure, Projects, and Industry Insight

One aspect of the Financial Mathematics program that I have found particularly impactful is how well organized it is in helping students navigate both academics and career development. The program does a great job structuring coursework while also emphasizing the importance of internships and industry exposure. From course planning to career workshops, the support provided helps students focus not only on learning advanced mathematical and financial concepts but also on positioning themselves for successful careers after graduation. Having clear guidance on course sequencing and internship preparation has made it much easier to balance technical learning with professional development.

As I continue through the program, my long-term career goal is to work in New York City in either investment banking or consulting. Both fields appeal to me because they combine quantitative analysis with strategic decision-making. My academic background in mathematics and finance has prepared me to approach complex financial problems analytically, while my professional experiences have helped me understand how these analytical tools are applied in real business environments. Ultimately, I hope to work in a role where I can use quantitative methods to evaluate financial opportunities, assess risk, and help organizations make informed decisions in fast-paced markets.

One project that reflects my interest in quantitative modeling is my current research project, Modeling Tail Risk via Stereographic Projection. The goal of this project is to develop a geometric framework for analyzing extreme events in financial markets. Instead of studying probability distributions solely in their traditional form, the project maps distributions onto the unit circle using stereographic projection. This approach allows us to analyze the behavior of extreme outcomes, or “tail events,” in a new geometric context. By incorporating Hermite polynomial expansions and extending the analysis beyond Gaussian distributions to cases such as the Cauchy and Poisson distributions, the project explores how geometric transformations can help better understand financial risk and stress scenarios. Working on this project has been a valuable opportunity to apply advanced mathematical concepts to real problems in financial modeling and risk analysis.

Outside of coursework and projects, conversations with industry professionals have also helped shape my perspective on potential career paths. Recently, I had the opportunity to speak with a family friend, Jim, who previously served as Vice President and Head Controller at PVH. During our conversation, he discussed his experience managing financial reporting and overseeing large corporate accounting operations. What stood out to me most was his emphasis on understanding the business behind the numbers. While strong analytical skills are important, he explained that successful professionals also need to understand how financial decisions impact broader organizational strategy.

Our discussion gave me a clearer picture of how financial expertise develops over the course of a career. Hearing about his experiences working within a major corporation reinforced the importance of combining technical knowledge with communication and strategic thinking. It also reminded me that building relationships and learning from experienced professionals can be just as valuable as classroom learning.

Overall, the combination of a well-structured academic program, hands-on quantitative projects, and conversations with professionals in the industry has helped me develop a clearer vision for my future career. The Financial Mathematics program continues to provide both the technical foundation and professional support needed to pursue opportunities in finance and consulting.

Anas Mouden

From Fundamental Finance To Quantitative Finance: Understanding Risk, Drift, And Randomness

If I had to describe the transition from fundamental finance to quantitative finance in one sentence, I would say this: fundamental finance explains why markets move; quantitative finance explains how they move.

Before entering graduate school, my understanding of financial markets was largely rooted in fundamentals—cash flows, valuation multiples, macroeconomic drivers, and corporate performance. Risk was something we discussed in terms of volatility, drawdowns, or sector exposure. Returns were tied to growth expectations, earnings surprises, and strategic positioning. Markets moved because information changed.

Assumptions	
Tax Rate	25%
Discount Rate	12%
Perpetural Growth Rate	3%
EV/EBITDA Multiple	7.0x
Transaction Date	2017-12-31
Fiscal Year End	2018-06-30
Current Price	25.00
Shares Outstanding	20.000
Debt	30.000
Cash	239.550
Capex	15.000

Discounted Cash Flow	Entry	2018	2019	2020	2021	2022	Exit
Date	2017-12-31	2018-06-30	2019-06-30	2020-06-30	2021-06-30	2022-06-30	2022-06-30
Time Periods		0	1	2	3	4	
Year Fraction		0.50	1.00	1.00	1.00	1.00	
EBIT		47.814	51.095	55.861	58.693	63.039	
Less: Cash Taxes		11.954	12.774	13.965	14.673	15.760	
Plus: D&A		15.008	15.005	15.003	15.002	15.001	
Less: Capex		15.000	15.000	15.000	15.000	15.000	
Less: Changes in NWC		375	611	398	511	272	
Unlevered FCF		35.494	37.715	41.501	43.510	47.008	
(Entry)/Exit	(290.450)						542.129
Transaction CF	(290.450)	17.747	37.715	41.501	43.510	47.008	542.129

Terminal Value	
Perpetural Growth	537.981
EV/EBITDA	546.278
Average	542.129

Intrinsic Value	
Enterprise Value	489.747
Plus: Cash	239.550
Less: Debt	30.000
Equity Value	699.297
Equity Value/Share	34.96

Market Value	
Market Cap	500.000
Plus: Debt	30.000
Less: Cash	239.550
Enterprise Value	290.450
Equity Value/Share	25.00

Rate of Return	
Target Price Upside	40%
Internal Rate of Return (IRR)	26%

Market Value vs Intrinsic Value	
Market Value	25.00
Upside	9.96
Intrinsic Value	34.96

Source: *Corporate Finance Institute*

However, through my studies in financial mathematics, I discovered that beneath these economic narratives lies a precise probabilistic structure. Concepts such as drift, diffusion, mean reversion, and the risk-neutral measure revealed that price dynamics can be modeled mathematically—often independent of the stories we attach to them.

At the heart of modern quantitative finance lies the idea that asset prices evolve according to stochastic processes. One of the foundational models is Geometric Brownian Motion (GBM), which assumes that stock prices follow a continuous-time process with two components: a deterministic drift and a random diffusion term. The drift represents the expected rate of return, while diffusion captures uncertainty—the unpredictable shocks that continuously impact markets.

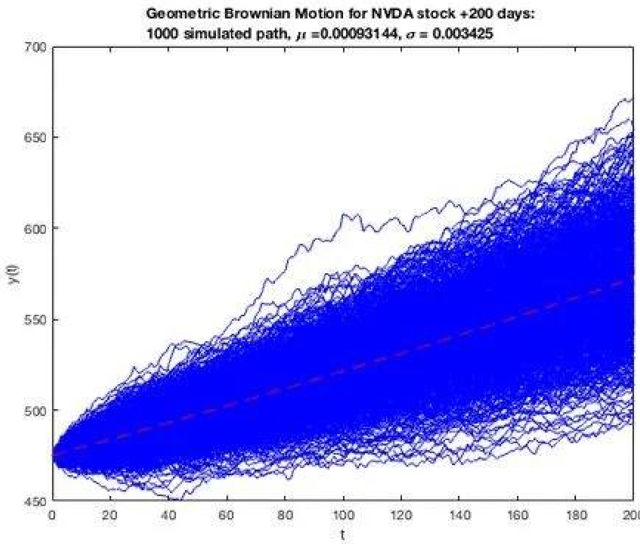

In fundamental finance, expected return reflects compensation for risk. In quantitative finance, that same idea appears as the drift term in a stochastic differential equation. But what fascinated me most was learning that when pricing derivatives, we do not use the real-world drift at all.

Instead, we transition to the risk-neutral measure. The risk-neutral framework was initially counterintuitive. Why would we assume investors are indifferent to risk? The answer lies not in behavioral assumptions, but in arbitrage logic. Under the risk-neutral measure, all assets are expected to grow at the risk-free rate. This transformation simplifies pricing because it allows us to compute the value of a derivative as the discounted expected payoff under this adjusted probability measure.

In other words, we replace economic expectations with mathematical consistency.

Continue Next Page

Anas Mouden

From Fundamental Finance To Quantitative Finance: Understanding Risk, Drift, And Randomness

Continued From Previous Page

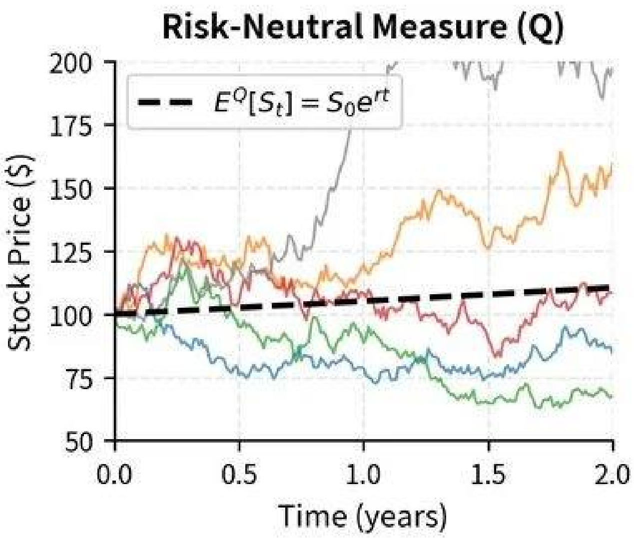

This shift is powerful. Under the real-world measure, drift reflects growth, macro outlook, and risk premia. Under the risk-neutral measure, drift becomes irrelevant for pricing; only volatility and the structure of uncertainty matter. Diffusion—the random component—remains central. This insight highlights a fundamental distinction between investing and pricing. Fundamental finance focuses on predicting returns. Quantitative finance focuses on replicating payoffs.

Another concept that deepened this bridge between the two worlds was mean reversion. In fundamental analysis, mean reversion often appears in valuation metrics—price-to-earnings ratios reverting to historical averages, margins normalizing over time, or interest rates drifting toward long-term equilibrium levels. In quantitative modeling, mean reversion is formalized through stochastic processes such as the Ornstein–Uhlenbeck model, where deviations from equilibrium are mathematically structured to decay over time.

What once felt like qualitative intuition becomes a measurable parameter. Drift, diffusion, and mean reversion together describe different philosophies of market behavior. GBM assumes proportional growth with constant volatility—appropriate for equities under the classical Black-Scholes framework. Mean-reverting processes better describe interest rates or commodities. Choosing between them is not merely a modeling decision; it reflects assumptions about how markets behave at a structural level.

Through this lens, I began to see quantitative finance not as a departure from fundamentals, but as their formalization. Economic beliefs become parameters. Market narratives become stochastic dynamics. Risk premia become transformations between probability measures.

The elegance lies in the synthesis: fundamental finance provides the economic rationale, while quantitative finance provides the mathematical machinery.

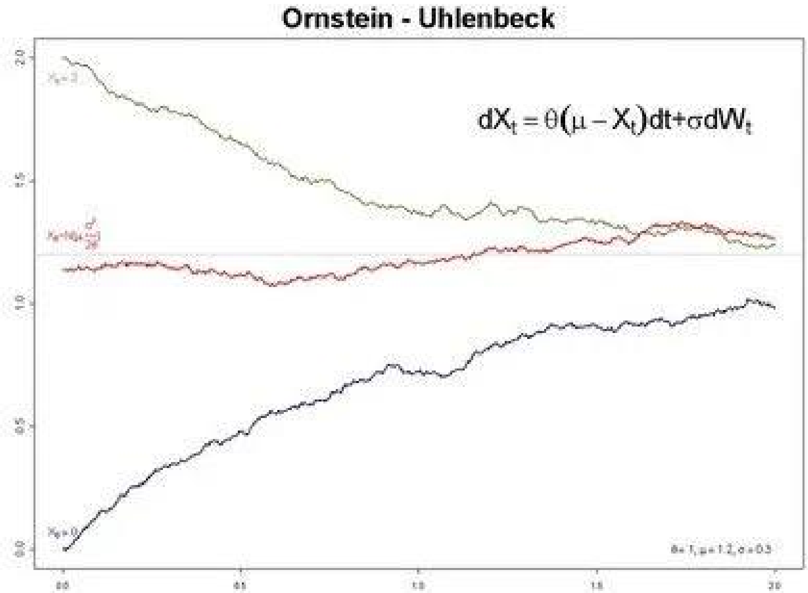

Maria Muntaha

From Curiosity To Quant Trading

When I first joined NC State University this semester, my goal was simply to improve my professional development and become more career-ready. At that time, however, I was not completely sure which direction within finance I wanted to pursue. Halfway through my first semester, that uncertainty has started to turn into a clearer goal. Through coursework and independent research, I have developed a strong interest in the trading industry, particularly quantitative trading.

One area that particularly interests me is high-frequency trading (HFT). HFT strategies rely on algorithms and extremely low-latency systems that execute trades in milliseconds. Many HFT strategies are based on small price inefficiencies or short-term signals. For example, a simplified representation of an HFT strategy could involve predicting very small price changes using statistical signals:

Price Prediction Model (Simplified): $r_t = \alpha + \beta x_t + \varepsilon_t$

Here r_t represents the short-term return, x_t is a predictive signal such as order flow or microstructure imbalance, and ε_t represents noise. If the predicted return exceeds trading costs, the algorithm executes the trade automatically.

HFT is often discussed in connection with the 2010 Flash Crash, when U.S. equity markets experienced a rapid and dramatic drop followed by a quick recovery within minutes.

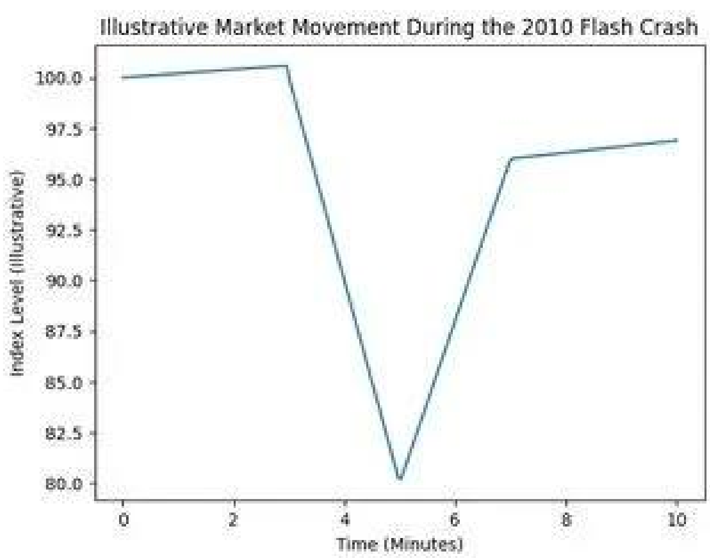

Figure 1: Illustrative market movement during the Flash Crash of May 6, 2010.

Credit: *Chart created for educational illustration based on publicly documented descriptions of the event.*

Another concept that interests me is delta hedging, a common risk-management technique used in options trading. Delta measures how sensitive the price of an option is to changes in the price of the underlying asset. What I find interesting about delta hedging is how it connects mathematical theory with real-time trading decisions. Traders must continuously adjust their hedge as market conditions change.

Option Delta: $\Delta = \partial V / \partial S$
Here V represents the option price and S represents the underlying asset price. A delta-hedged portfolio attempts to maintain neutrality to small movements in the underlying asset.

Delta-Hedged Portfolio: $\Pi = V - \Delta S$
By dynamically adjusting the position in the underlying asset, traders can maintain a near-neutral exposure to price changes while focusing on other factors such as volatility and time decay.

I am also interested in how volatility is modeled in financial markets. Returns often display volatility clustering, where periods of high volatility tend to follow each other. One widely used model that captures this behavior is the GARCH model.

GARCH(1,1) Volatility Model: $\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$

Models like this allow traders and researchers to estimate how volatility evolves over time, which is particularly useful for options pricing and volatility-based trading strategies.

Ultimately, what draws me to trading is the combination of analytical thinking, creativity, and decision-making under uncertainty. Many areas of finance contribute to understanding markets, but trading is where those ideas are ultimately tested and turned into real decisions.

As I continue my studies in Financial Mathematics at NC State University, I hope to deepen my understanding of algorithmic trading, market microstructure, and volatility modeling so that I can pursue a career as a quantitative trader.

Isaiah Nick

The Hidden Option In Prediction Markets

Every live sports position on Kalshi contains embedded optionality, but almost nobody prices it that way. When you buy a binary contract that pays $1.00 if your team wins, you are not only taking a directional view. You also retain the right to sell that position at any point before settlement, exercising only when the market bid exceeds your estimate of fair value. That early-exercise asymmetry mirrors the structure of an American-style option, and drives my college basketball live trading engine.

The foundation of this system is a continuously updating win probability for every game. As detailed last semester, an ensemble of machine learning models produces a pregame estimate. Once the game tips, that estimate gets blended with the live score through a probabilistic construction: project the pregame probability into margin space via the inverse Normal distribution, then ask how much of that expected margin remains. Remaining uncertainty scales as $\sigma\sqrt{(t/T)}$ where the square root comes directly from the random walk structure of scoring. The construction is self-consistent by design: at tip, σ cancels and the model outputs exactly the pregame probability. The single free parameter in the model σ = 17.2 controls how stubbornly the model clings to its prior, learned by minimizing the Brier score across 4,000+ games of play-by-play data.

Once a reliable fair value estimate exists, the problem changes entirely. A naive approach would be to enter when the market disagrees and exit when it converges. But that ignores something valuable: you don't have to exit. Every open position carries embedded optionality — sell when the market overprices, hold otherwise. With that in mind, the question becomes: what is this option worth, and when should I exercise it?

Quantifying that value required measuring how much the market moves around fair value. From 5.6 million in-game snapshots, the bid-versus-model residual has a standard deviation of about four cents, but conditioning on game state reveals something more interesting. Market noise follows roughly $\sigma(p) = \sqrt{(p(1 - p)}$, peaking in contested games and collapsing when the outcome is decided. That is not a modeling assumption, rather it is the empirical shape of disagreement in a prediction market, shown in Figure 1.

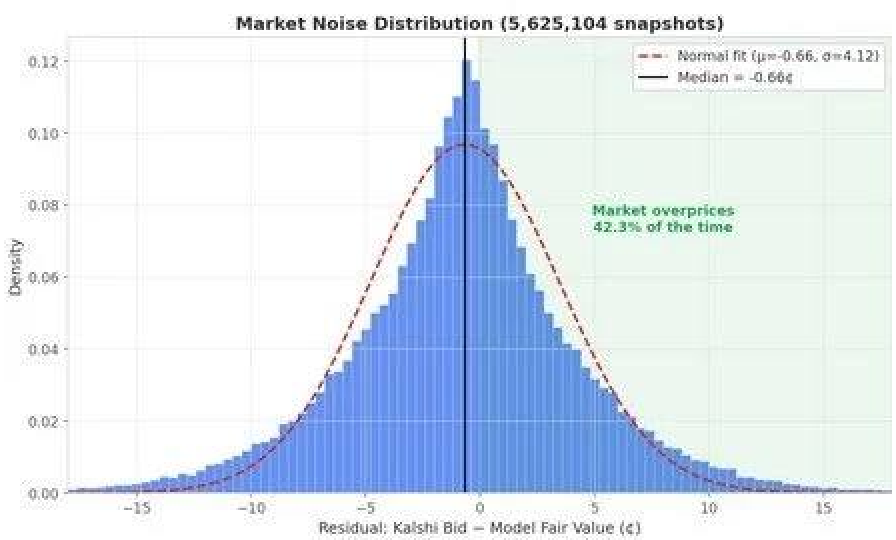

With noise characterized, option value is computed via backward induction, the same dynamic programming approach used to price American options on binomial trees. At each node, the sell premium, computed using (p), is the expected gain from selling only when the bid exceeds continuation value. But backward induction requires approximately independent timesteps, and market prices are autocorrelated. Measuring the autocorrelation structure of residuals across 970 games revealed that a 120-second interval reduces autocorrelation to roughly 0.30, or 9% shared variance, low enough to treat consecutive timesteps as independent, preventing option value inflation from correlated noise. Rather than deploying a 3-dimensional lookup table in production, a three-parameter closed-form solution was fit to 34,000 grid points at 0.31 cents RMSE.

Continue Next Page

Isaiah Nick

The Hidden Option In Prediction Markets

Continued From Previous Page

This allows us to capture the same theta and gamma with no implementation artifacts at runtime. Figure 2 shows the option value surface as fit by the closed form solution.

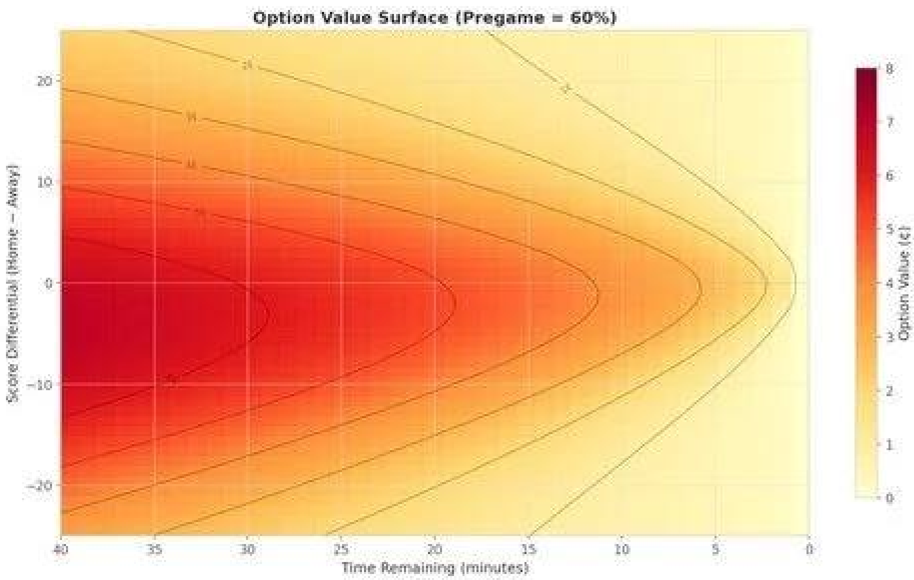

The system has been running autonomously for the last month, polling ESPN and Kalshi in parallel every second and executing via WebSocket at sub-second latency. But one blindspot required a fix: in the final seconds of a close game, microstate elements dominate, none of which the data feed captures. Market makers watching the broadcast price these elements correctly; the model cannot. Without adjustment, what looks like an edge is often just an information gap. An additive haircut on the hold value was calibrated from excess residual dispersion, raising the exit threshold precisely where the model's disadvantage is largest while remaining invisible everywhere else.

The natural validation question is simple: when the market and the model disagree, which one moves towards the other? Across 287,000+ observations, the market moves towards the model between 1.4 and 5.4 times more often than the reverse, depending on the size of disagreement and the lookahead window. This ratio decays as the window extends, consistent with information gradually being priced in, and holds above 1.0 at every threshold and horizon tested. Persistent convergence is the signature the framework is built on.

Trevor Otte

Crossroads of the Industry

From entering the Financial Mathematics graduate program at NC State, I knew the direction that I wanted to go. I wanted to apply my mathematics background in the finance industry, doing high-level work contributing to the success of a business. The more I work through my program, the more I am enjoying learning and applying the material. This has led me to want to continue my education and complete a PhD in the industry. This has been a very thought-out decision, one that has come after a lot of thought about the trajectory of my academic career.

This decision became certain for me once I had begun taking Stochastic Processes II. This course is my first introduction to PhD-level classes, and the more I immerse myself in the topic, the more I find myself wanting to learn more. This class is the precursor to SDEs, a class that I have been looking forward to taking since before I received my acceptance letter to the graduate program. I wish to continue my education in this area and, hopefully, tackle research in SDEs or SPDEs.

Therefore, after much thought and research, I will be applying to continue my studies in financial mathematics with the goal of completing research that is at the forefront of current knowledge. I am very excited for the future, and I look forward to taking SDEs and pushing the boundaries of what is currently known.

Gaurav Pai

From Theory To Toolsets: Bridging Quantitative Rigor And Career Readiness

When I started the Master of Financial Mathematics program at NC State, I expected the math to be intense. I was right. What I didn't expect was how quickly theory would turn into something I actually had to build, test, and defend in code.

Coding Is Where the Math Gets Honest

Courses like Derivative Pricing, Stochastic Calculus, Fixed Income, and Financial Risk Analysis are demanding. You can't fake understanding when you're deriving models on paper. But the real test begins when you try to implement them.

It's one thing to understand Martingales in theory. It's another to build a Python model that doesn't fall apart when exposed to real, imperfect data. My coding skills evolved from working through fundamentals to applying them in two major projects: developing a CCAR stress-testing framework and building a Market Regime Detection model using Gaussian Mixture Models and PCA.

The GMM project was especially important. Instead of relying on surface-level trend indicators, I had to construct a probabilistic system capable of identifying hidden market states such as inflationary stress. That required translating statistical theory into stable, testable code. The CCAR work pushed me to think institutionally, structuring models around risk scenarios rather than academic exercises.

The program's technical training made that progression possible. What began as learning syntax and model mechanics became the ability to design and implement systems with real financial context.

Career Readiness Isn't an Afterthought

The technical rigor is real, but the professional preparation is just as structured. Working with Career Services forced me to rethink how I present myself. Translating prior business experience into quant-relevant language wasn't automatic; it took revisions, mock interviews, and some uncomfortable feedback. Learning to clearly explain a technical project without hiding behind jargon is often harder than solving the model itself.

The mock interviews were especially useful. They exposed gaps I didn't notice: knowing something isn't the same as communicating it under pressure. Our Program Ambassadors also play a vital role here. Group projects mirror actual team dynamics, dealing with continuous updates and differing approaches. It is a constant reminder that finance isn't a solo sport.

My Objective: Integrated Risk Strategy

This program has solidified my professional objective: mastering the integrated discipline of quantitative risk management and investment strategy. While many view these as separate paths, the MFM curriculum has demonstrated that they are fundamentally the same pursuit.

Through the lens of Stochastic Calculus and Statistical Modeling, I've learned that risk is the very variable that defines the opportunity set for an investment. By applying machine learning to market regime detection, I've seen how risk management tools serve as the primary driver for proactive portfolio rotation. The MFM program has been essential in teaching me to use the "defense" of risk modeling to inform the "offense" of investment alpha. My goal is to design strategies that are not only high-performing but are built on a foundation of mathematical resilience.

The combination of technical depth and professional development has given me more than academic knowledge. It has provided a practical toolkit and a clearer understanding of how the industry functions. The math matters, and the code matters; but the ability to connect both to real institutional problems and communicate them clearly matters just as much.

Pushkar Patil

Why Education Debt Isn't A Leveraged Buyout

Every time an individual takes debt and invests in an activity seeking return, the natural expectation is a good return. Education follows the same logic, students borrow to finance studies and living expenses, hoping to secure a strong job and repay through future earnings. This mirrors private equity, where funds borrow to invest in a company with the expectation of exponential return. Education loans carry the same allure of leveraged investment.

At face value, a post-college quant job looks like a home run. Six-figure base salary. Bonus upside. Brand-name firm. The narrative writes itself: borrow $100,000, land a $150,000 compensation package, and the return feels exponential. It feels like a leveraged buyout of yourself. Debt today, explosive equity tomorrow.

But that glow fades the moment you discount the cash flows. From a quant finance perspective, what matters is not the headline salary but the present value of incremental earnings net of costs, adjusted for risk. Once you model that carefully, the comparison to leveraged buyouts breaks down sharply.

The leverage mechanics differ fundamentally. In an LBO, a PE fund acquires a company with debt. The equity slice is thin. If enterprise value grows even modestly, the residual equity value increases disproportionately because debt is fixed in nominal terms. This creates convexity. A 20% rise in enterprise value can translate into a 2–3x equity multiple. IRR is amplified through operational improvements and exit multiple expansion.

Education loans finance human capital, but human capital is not a scalable corporate asset. Income is constrained by time, market wages, and career progression. Compensation follows a stochastic but bounded process. Bonuses vary. Markets cycle. Firms restructure. The salary process resembles a drift-plus-volatility diffusion, not an exponential equity payoff.

The payoff structure is the key difference. In private equity, equity holders capture residual upside after debt repayment. The payoff is convex. In education, income growth is largely linear in experience. There is no "multiple expansion" on your own labor. You cannot 5x your output without scaling through ownership or entrepreneurship.

Discounting matters enormously. A $150,000 salary is not $150,000 today. You discount future earnings by a rate reflecting loan interest and income uncertainty. If your education loan carries 6–8% interest and your income is correlated with macroeconomic downturns, the appropriate discount rate is not trivial. The net present value looks far less explosive and that's before living costs enter the picture.

In an LBO, debt service is paid out of EBITDA after optimizing cost structures. In personal finance, your "EBITDA" is post-tax income. Rent, food, insurance, healthcare, and taxes are unavoidable. Debt service comes after survival costs. The free cash flow available to amortize student loans is far lower than gross compensation suggests.

Time horizon and liquidity differ materially. Private equity investments are structured around 5–7 year exits. Education returns unfold gradually over decades as wage differentials accumulate slowly. There is no liquidity event where you "sell" your degree. Human capital is illiquid and non-transferable.

Risk is also more concentrated. Private equity funds diversify across multiple companies. Education loans are concentrated bets on a single asset: you. If your sector contracts, income risk spikes precisely when debt obligations remain fixed. This introduces negative convexity: significant downside risk, capped upside.

Are education returns positive? Often, yes. The NPV of a strong quantitative degree can exceed tuition costs by a meaningful margin. But these returns resemble a steady, risk-adjusted wage premium and not the high-volatility, convex payoff profile of private equity. The psychological comparison persists because both involve leverage: borrow money, invest, reap future gains. But leverage alone does not create exponential returns. Exponential outcomes arise when leverage interacts with scalable assets, operational optionality, and exit dynamics. Human capital lacks that scalability.

Education is an investment. It can have positive NPV and dominate alternative career paths. But from a quant finance perspective, it is not a leveraged buyout of yourself. It is a long-duration, illiquid, concentrated exposure to a single labor-income stream with limited convexity. And that is why it does not compound like private equity.

Aditya Purohit
Quants In Climate Risk

Climate change is no longer just an environmental conversation. It has become a financial one too. As extreme weather events grow more frequent and severe, the mortgage lending industry faces a challenge that traditional credit risk models were never built to handle. This is where quantitative finance steps in, bringing data, mathematics, and modeling together to ask a question that matters deeply to lenders, regulators, and homeowners alike: how much does climate risk actually cost?

Through my work in the Financial Mathematics program at NC State, I have been part of a team tackling exactly this question. Our project focuses on quantifying climate risk in residential mortgages by integrating disaster-level data across floods, wildfires, and hurricanes with loan-level mortgage data. The goal is to understand how physical climate events influence a borrower's probability of default and ultimately reshape the risk profile of a mortgage portfolio.

The foundation of the work started with building a baseline probability of default model using traditional credit risk drivers like credit scores, loan-to-value ratios, and borrower characteristics. This model achieved an AUC of 0.83, meaning it does a solid job of distinguishing between defaulters and non-defaulters under normal conditions. But normal conditions are becoming less reliable as climate events intensify. This baseline model serves as our benchmark, and the entire next phase of the project is about measuring how much climate exposure changes that picture.

On the flood side, which is the part of the project I have focused on most deeply so far, we pulled data from FEMA's OpenFEMA database and merged it with publicly available elevation data at the zip code level. From there, we engineered a composite flood risk score ranging from one to five that accounts for elevation, proximity to flood zones, historical flood damage, recorded water levels, frequency of flood events, and total number of insured losses in a given area. What stood out immediately was the result that roughly 60 percent of properties in our dataset scored at the high-risk level of four, with elevation and proximity together accounting for about half of the predictive weight in the model. That finding alone tells a compelling story about how much geographic exposure drives financial vulnerability in ways that a credit score never could.

This kind of work sits at the intersection of data science and finance in a way that feels genuinely exciting. Feature engineering for climate risk is not like engineering features for a standard credit model. You are working with geospatial data, disaster declarations, satellite-derived burn ratios, hurricane track records going back decades, and weather API feeds. The technical side demands fluency in tools like Python and GeoPandas, while the financial side demands a clear understanding of what these features mean for loss given default and how lenders should actually price that risk.

Regulators like the Federal Reserve and international bodies like the Bank for International Settlements have started pushing financial institutions to stress-test their portfolios for physical climate risk. That pressure creates real demand for people who can bridge the gap between climate science and credit modeling.

The broader implication of this work extends beyond any one project. Mortgage markets are deeply tied to how we as a society value land and place. If climate risk starts repricing entire zip codes, the effects ripple through home values, insurance markets, local tax bases, and community stability. Quantitative tools give us a way to see those risks clearly before they crystallize into losses. That transparency is what makes this kind of work meaningful, not just technically interesting.

Being part of the Financial Mathematics program has given me the foundation to engage with problems like this seriously. The combination of statistical rigor, financial theory, and real-world data that the program emphasizes is exactly what climate risk modeling requires. I came in wanting to apply math to finance. I did not expect that one of the most important applications would be helping the financial system understand the ground beneath our feet.

Antarish Rautela

Birth Of A New Possibility

One aspect of the Financial Mathematics program that has been most impactful for me is the opportunity to enroll in PhD-level coursework in machine learning and statistics at NC State. Coming into this program, I expected rigorous finance courses but I did not fully anticipate how much statistics was available in this program.

Taking courses like statistical machine learning alongside doctoral students has pushed me to engage with material at a depth I had not previously experienced. Rather than simply applying methods, I've been challenged to interrogate why those methods work and their theoretical foundations, their failure modes, and the open questions that still surround them. That shift in perspective has been genuinely transformative. It moved me from thinking like a practitioner to thinking like a researcher.

What has struck me most is how naturally the intersection of statistics and finance generates meaningful, unsolved problems. The more deeply I've engaged with the technical side of this curriculum, the more I've come to see that the questions I find most compelling about model uncertainty, causal inference in markets, the behavior of learning algorithms under distribution shift are questions that require the kind of sustained, rigorous inquiry that a PhD is designed to produce.

During my time at NC State, I got to dive deep into the theory of machine learning and not just the rules that certain machine learning methods have. I got to understand why these rules are set into place and why these methods were formulated. For example, I am currently learning about SVM and SVC where SVM was invented from SVC to adjust for true boundaries between two classes being nonlinear and the linear algebra behind it was quite interesting.

In short, access to this level of coursework has not just strengthened my technical foundation; it has opened my mind to the real possibility of pursuing a doctorate. It gave me a concrete sense of what research culture looks and feels like from the inside, and it made me want to be part of it.

Noah Regal

When The Hypothesis Fails

Earlier this semester my partner and I worked on a research project asking whether the way a financial news headline is written affects how quickly the stock market prices it in. The idea came from a simple observation. If a headline says, "Company X misses earnings estimates," the information is clear. However, if it says, "Company X may miss earnings estimates," the language signals uncertainty. Our hypothesis was that investors would respond less completely to hedged, uncertain headlines, leaving more predictable return movement in the days that followed.

To test this we built a dataset of about 3.1 million Yahoo Finance headlines covering S&P 500 companies from 2016 through early 2026. We scored each headline for linguistic uncertainty using financial text analysis tools and an AI language model, then matched the headlines to daily stock returns to measure how much of a stock's reaction to news was delayed versus immediate.

The hypothesis failed. Linguistic hedging, whether measured through modal verbs, uncertainty wordlists, or a GPT-4o-mini scoring approach, predicted essentially nothing about delayed price incorporation. The form of headlines did not matter in the way we expected. That null result ended up being the most interesting part of the project.

While diagnosing why the linguistic features failed, we noticed that a different variable was doing something. We had included it almost as a control: the difference between a company's news sentiment on a given day and the average sentiment across all covered firms that same day. A company receiving positive news on a broadly negative day, or negative news when the broader market mood is optimistic. That variable predicted delayed price incorporation consistently and strongly. Companies whose news diverged from the prevailing market mood showed larger return corrections over the following five trading days. The relationship held across the full sample period, survived robustness checks, and held up out of sample over a seven year test period.

The intuition makes sense in hindsight. When a stock's news runs against the grain of the market's mood, investors have to separate the firm-specific signal from the broader market context. That process appears to be incomplete on the day the news arrives, with the rest of the adjustment happening over subsequent days. The effect is strongest for news that arrives after market close, when investors cannot trade on it immediately. What we took away from the project was less about the specific finding and more about how empirical research actually works. The original hypothesis was motivated by a reasonable story. Its failure was just as informative as a positive result would have been. Three separate diagnostic tests ruled out data quality and measurement problems as explanations, which meant the null was real. And the variable that ended up mattering most was one we had not originally thought of as the main event.

The paper is titled "Hedged Headlines and Slow Markets" and was submitted to the Human x AI Finance competition at UCLA Anderson. It covers many hours of work including data collection, model development, and writing. Working on it as two researchers would not have been feasible at this scale without AI tools, both in building the empirical pipeline and in the writing and revision process, which the paper documents openly.

The short version of the finding is this: it is not the uncertainty in how news is written that slows price discovery. It is whether the news itself is out of step with what the rest of the market is hearing that day.

Melanie Sanchez
Predicting Credit Risk

During my graduate studies in Financial Mathematics at North Carolina State University, I completed a project focused on building a Probability of Default (PD) model using consumer lending data from the LendingClub platform. Probability of Default models are a core component of modern credit risk management and are widely used by financial institutions to estimate the likelihood that a borrower will fail to repay a loan. These estimates influence lending decisions, interest rate pricing, and regulatory capital requirements.

The dataset consisted of historical loan records that included borrower financial characteristics, loan features, and repayment outcomes. Each observation represented an individual loan and contained variables such as credit score range, annual income, employment length, loan amount, interest rate, and debt-to-income ratio. The objective of the project was to predict whether a borrower would default or fully repay the loan, making the task a binary classification problem.

The first stage of the project focused on data exploration and preparation. This involved examining the distributions of variables, identifying missing values, and understanding how borrower characteristics differed between loans that defaulted and those that were repaid. One key challenge was the presence of class imbalance, since default events occur much less frequently than successful repayments. Addressing this imbalance is an important step in building reliable credit risk models.

To estimate default probabilities, I first implemented a logistic regression model, which is one of the most commonly used approaches in credit risk modeling due to its interpretability. Logistic regression models the probability of default as a function of borrower characteristics, producing outputs between 0 and 1 that can be interpreted as the likelihood of default. The model allowed me to examine how variables such as debt-to-income ratios or credit score categories influence the probability that a borrower defaults.

In addition to logistic regression, I also explored machine learning techniques, including gradient boosted trees (XGBoost), to capture nonlinear relationships between borrower characteristics and default risk. Machine learning models can often improve predictive accuracy by detecting interactions between variables that may not be visible in a traditional linear framework.

To evaluate model performance, I used metrics commonly applied in credit risk modeling, including ROC-AUC and classification accuracy. These metrics help determine how well the model distinguishes between borrowers who are likely to default and those who are likely to repay their loans.

This project provided valuable insight into the practical challenges of credit risk modeling. Beyond the statistical techniques themselves, building a PD model requires thoughtful data preparation, careful interpretation of model results, and an understanding of the trade-offs between model interpretability and predictive power. Overall, the experience reinforced how quantitative methods can be applied to real financial data to support risk management and lending decisions.

Catherine Sexton

The Impact Of NC State's MFM Program

I joined the Master of Financial Mathematics (MFM) program at North Carolina State University to pursue my career goals in quantitative finance. With my strong interest in the field of mathematics and my desire to apply quantitative methods to complex real-world problems, financial mathematics represented the ideal intersection of my interests. However, the field of quantitative finance is highly competitive and continues to grow in popularity. Selecting the MFM program was one of the most important decisions I have made in preparing myself for a successful career in the field.

My interest in quantitative finance began during my undergraduate studies, where I pursued degrees in both mathematics and statistics. I intentionally chose classes that aligned with my career interests, including mathematics of derivative securities, stochastic processes, and financial time series. These classes were both challenging and engaging, and they further solidified my desire to pursue graduate study in financial mathematics. My primary expectation with a master's program was to further my technical knowledge through more advanced coursework that would better prepare me for professional work in quantitative finance.

Since I have started in the MFM program, my academic expectations have certainly been met through the challenging and informative courses I am taking. However, one aspect of the program that has had the greatest impact on me has been its extensive career services.

Even before classes began, the program introduced me to the internship application process and began preparing us for the competitive recruiting cycle in quantitative finance. Once the semester starts, the career seminar course provides structured opportunities to develop the professional and interpersonal skills necessary for success in the industry. This includes tailored resume reviews, mock interviews, presentation practice, networking guidance, research projects, and industry speakers.

One of the most valuable components of the program has been the industry projects that allow students to gain experience applying their skills to real financial problems. I am currently working on a project titled Modeling Tail Risk Through Stereographic Projection. In this project, I am applying the mathematical concept of stereographic projection (mapping data points from the real line to a circle) to transform the normal distribution in order to better analyze extreme events in the tails of the distribution. Projects like this mirror the type of analytical and modeling work performed by quantitative researchers in the financial industry. This experience is not only strengthening my technical abilities but also providing concrete examples of my work that I can present to potential employers during the job search process.

Another of the many career services that has been particularly valuable is the weekly industry speakers. These presentations offer regular networking opportunities and allow students to interact with professionals working across various areas of quantitative finance. The speakers share insights about their career paths, discuss current trends in the industry, and answer student questions. These conversations have broadened my understanding of the variety of possible career paths in financial mathematics and have helped me think more strategically about my own career goals.

What makes the MFM program's career services especially impactful is the personalized support provided to each student. The faculty of the program, including the Director, Dr. Pang, the Director of Career Services, Mr. Roberts, and the professors, are all invested in the success of the program and the individual students. I have received thoughtful feedback and individualized advice on resumes, interview preparation, and career strategy. This level of mentorship has been invaluable in helping me navigate the competitive job market in quantitative finance.

The comprehensive career services and professional development opportunities have been the most impactful aspect of my experience so far in the MFM program at North Carolina State University. The program not only strengthens students' technical knowledge but also actively prepares us to translate those skills into successful careers in the financial industry.

Shreyansh Sharma

Delivery Risk In Futures Contracts

In April 2020, the NYMEX May 2020 WTI crude oil futures contract (ticker CLM20) produced one of the most dramatic dislocations ever observed in a major derivatives market. On April 20, the contract settled at −$37.63 per barrel, marking the first negative settlement price in the history of the WTI benchmark [1]. What initially appears to be an extraordinary anomaly in commodity markets becomes easier to understand once the mechanics of physical delivery are taken seriously. The episode provides a clear illustration of how delivery risk in physically settled futures contracts can dominate market outcomes when financial participants approach expiration without the operational capacity to take delivery.

The backdrop was a historic collapse in oil demand caused by the COVID-19 pandemic. Global travel restrictions and lockdowns sharply reduced transportation fuel consumption, while oil production remained elevated in early 2020 following a breakdown in coordination among major producing countries in March. U.S. refinery inputs fell to roughly 13 million barrels per day by mid-April, approximately 24% below the level observed one year earlier [2]. The resulting supply imbalance caused crude inventories to rise rapidly across the United States.

The delivery point for NYMEX WTI futures is Cushing, Oklahoma, a major storage and pipeline hub in the U.S. crude oil network. As the supply surplus grew, available storage capacity at Cushing tightened significantly. By mid-April 2020, approximately 76% of the hub's working storage capacity of roughly 76 million barrels was already filled [2]. In practice, much of the remaining capacity was either leased or operationally constrained. This meant that traders holding long futures positions close to expiration faced a real logistical problem: unless they exited their contracts, they could be obligated to take delivery of physical crude oil with no available storage.

The May 2020 contract was scheduled to expire on April 21, leaving April 20 as the final full trading day. By that point, most market participants who did not intend to take delivery had already rolled their positions into later contracts. Open interest had fallen dramatically from approximately 635,000 contracts in early April to about 108,000 contracts by April 20 [1]. However, even this residual open interest represented a substantial quantity of crude oil that would need to be delivered if positions were not closed.

The price action on April 20 illustrates how delivery constraints can dominate market dynamics near expiration. The May contract had settled at $18.27 per barrel on April 17. During trading on April 20, the price initially declined gradually before collapsing during the afternoon session. The contract reached an intraday low of −$40.32 per barrel shortly before the settlement window and ultimately settled at −$37.63 per barrel [1].

At the same time, later-dated contracts remained positive. The June 2020 WTI futures con- tract settled at $20.43 per barrel on the same day [2]. The resulting May–June spread widened dramatically, reaching roughly −$58 per barrel by settlement. This extreme contango reflected the market's valuation of immediate storage constraints: crude oil deliverable in May effectively carried a large negative convenience value relative to oil deliverable one month later.

An important structural factor behind the dislocation was the presence of large passive financial participants. Commodity index funds and exchange-traded products obtain oil exposure through rolling futures contracts rather than taking delivery. These vehicles typically follow predetermined roll schedules that shift exposure from the expiring contract into the next maturity. In March and early April 2020, investor inflows into oil exchange-traded funds increased sharply as retail investors attempted to buy the decline in crude prices. Research by the Bank for International Settlements documents that inflows into oil ETFs surged in late March and early April, placing additional pressure on futures markets during the roll process [3].

I first came across this episode while listening to a finance podcast and ended up digging deeper to understand the delivery mechanics embedded in futures contracts. In my coursework, futures have mostly been discussed from a financial perspective, often focusing on pricing, hedging, and cash-settled contracts. The logistical side of physically settled futures, involving storage capacity, transportation constraints, and operational considerations, was something I had not previously thought about in detail. That realization reminded me of a piece of advice shared by a guest speaker in my FIM 601 class, who emphasized the value of actively listening to finance podcasts as a way to encounter real market situations beyond the classroom. This incident stood out as a good example of that idea, and it ultimately motivated me to include this anecdote in my Quant Quarterly.

References

[1] U.S. Commodity Futures Trading Commission. Interim Staff Report: Trading in NYMEX WTI Crude Oil Futures Contract Leading to the Negative Settlement Price on April 20, 2020. CFTC, 2020.
[2] U.S. Energy Information Administration. Why did WTI crude oil futures prices turn negative? Today in Energy, April 2020.
[3] Bank for International Settlements. The Oil Market Turmoil of April 2020. BIS Quarterly Review, 2021.

Shiqin Song

Risk Management Through Hands-On Projects

One aspect of the Financial Mathematics program that I have enjoyed the most is the opportunity to work on practical projects. Coming from a strong quantitative and mathematical background, I have always been interested in applying theoretical knowledge to real financial problems. The program provides many opportunities to do exactly that through project-based learning and industry connections.

In particular, I really appreciate that the program encourages students to participate in applied projects rather than focusing only on theoretical coursework. These projects allow us to work with real datasets, use programming tools such as Python, and think about how quantitative models are used in industry. Through these experiences, I have been able to better understand how concepts from statistics, finance, and mathematics can be combined to solve real problems.

One project that had a particularly strong impact on me was a credit risk modeling project where our team built a Probability of Default (PD) model. In this project, we used a large loan dataset and applied statistical methods to estimate the likelihood that borrowers would default. Initially, we implemented a logistic regression model to analyze the relationship between borrower characteristics and default probability. While this model provided a good baseline, we also explored more advanced machine learning approaches such as Xgboost to capture nonlinear relationships and interactions between variables.

Working on this project helped me understand the practical challenges of risk modeling. For example, we had to carefully select variables, evaluate model performance, and consider how the model might behave in different economic conditions. It also showed me that building a model is not only about statistical accuracy, but also about stability, interpretability, and how the results can support real decision-making in financial institutions.

This experience sparked my interest in risk management. I realized that risk analytics plays a critical role in the financial system, especially in areas such as credit risk, market risk, and risk management. Institutions rely on quantitative models to assess potential losses, allocate capital, and make informed lending or investment decisions. Being able to contribute to this process through data analysis and modeling is something I find both challenging and meaningful.

Overall, the FM program has strengthened my interest in pursuing a career in quantitative risk management. Through hands-on projects, collaborative teamwork, and industry engagement, I have gained valuable experience applying quantitative tools to real financial problems. The PD modeling project in particular helped me discover how much I enjoy working on risk analytics and reinforced my goal of building a career in this field.

Surya Sridhar

Climate, Credit, And Collateral

The nexus between climate change and financial stability has shifted from a theoretical concern for future generations to an immediate, quantifiable reality for the global banking sector. As catastrophic weather events increase in frequency and intensity, the traditional assumptions underpinning residential mortgage lending; namely, that physical assets remain stable collateral over a 30-year horizon are being fundamentally challenged. My project on climate-risk mortgages was born from this urgency: a necessity to bridge the critical gap between climate science and risk-based credit modeling.

Mortgages represent the single largest asset class for many commercial banks. When we consider the volume of capital tied to residential property, we realize that the banking sector is essentially "long" on geography. If the value of a property is physically compromised by wildfire, flood, or sea-level rise, the collateral backing the loan evaporates.

The integration of climate risk into mortgage underwriting necessitates a transition from backward-looking historical models to forward-looking, scenario-based frameworks. Traditional credit models rely on the assumption that property collateral remains stable over a 30-year duration. However, non-stationary climate patterns, such as intensifying flood hazards and chronic heat stress, introduce physical and transition risks that historical data fails to capture.

The core technical challenge in this project is quantifying the lead-lag relationship between climate hazards and financial indicators. We often treat climate risk as a "slow-moving" hazard that eventually triggers a "fast-moving" financial crisis.

- **The "Lead" (Climate Data)**: Physical indicators like rising groundwater levels, increased frequency of extreme heat days, or changes in floodplain maps. These variables are highly autocorrelated and exhibit long-term, non-stationary trends.
- **The "Lag" (Financial Indicators)**: Insurance premium spikes, property value depreciation, and, ultimately, the Probability of Default (PD) or Loss Given Default (LGD) in a mortgage book.

The challenge is that the market (the "Lag") often ignores the "Lead" data until a tipping point is reached. My model aims to quantify this "time-to-event" lag. If a property enters a new, high-risk flood zone, we need to know exactly when the insurance market adjusts, and when that adjustment cascades into a lower home equity buffer. By measuring this lag, we can stress-test portfolios before the market re-prices the risk.

The primary motivation for this project was to create a framework that translates abstract climate projections into concrete financial variables. In the climate community, data is expressed in temperature anomalies and millimeters of sea-level rise. In the finance community, data is expressed in Basis Points (bps) and Expected Loss (EL). My project focuses on the transmission channels between these worlds, specifically:

1. **Physical Risk (Direct):** Property damage caused by acute weather events.
2. **Insurance Risk (Indirect)**: The economic impact of escalating premiums or the withdrawal of insurance coverage in high-risk zones.

The vision for this project is to establish a new paradigm in underwriting. Instead of treating climate risk as a binary for lending, the model incentivizes resilience. It introduces a mechanism where mortgage terms, interest rates, or down-payment requirements are adjusted based on the specific climate-resilience characteristics of the building, such as retrofitted flood defenses or fire-resistant materials.

By internalizing the costs of climate externalities into the mortgage price, we achieve two objectives: it reduces the bank's exposure to volatile assets, and it provides a financial signal that makes climate-resilient upgrades economically viable for the borrower. Ultimately, this project is a tool for financial translation. It allows banks to acknowledge the reality of a changing climate without retreating from the mortgage market entirely. By building a robust, data-driven framework that accounts for the physical realities of the 21st century, we can ensure that the housing market remains a foundation for stability rather than a source of systemic financial fragility. The future of lending must be as dynamic as the climate itself, and this project serves as a critical step in building that necessary resilience.

Abhishek Tandon

Local Volatility Surface Estimation

Volatility is one of the most important inputs in option pricing and derivatives risk management. Classical models such as the Black–Scholes framework assume that volatility is constant across strikes and maturities. However, empirical evidence from option markets consistently contradicts this assumption. In practice, implied volatility varies across both strike price and time to maturity, forming what is commonly known as the volatility smile or volatility surface. These patterns reflect market expectations of future uncertainty and the pricing of tail risk. To address this limitation, this project focuses on estimating a local volatility surface from market option data and validating its ability to reproduce observed option prices.

The motivation for using a local volatility framework arises from the inability of constant-volatility models to capture the structure observed in real option markets. Out-of-the-money options tend to trade at higher implied volatilities than at-the-money options, and volatility also varies across maturities. These observations suggest that volatility depends on both the level of the underlying asset and time. The local volatility model, introduced by Bruno Dupire, addresses this issue by modeling volatility as a deterministic function of the underlying asset price and time, denoted as $\sigma(S,t)$. When properly calibrated, this framework can reproduce the entire implied volatility surface observed in the market.

The analysis uses SPX option data from 2023, a period characterized by relatively stable market conditions and moderate volatility levels. Using data from a lower-volatility environment allows the calibration process to focus on structural patterns in the volatility surface rather than distortions caused by extreme market events. However, raw option data often contains noise and inconsistencies, making preprocessing an essential step before calibration.

Several data quality filters were applied to ensure reliable option quotes. Observations with zero or negative bid prices were removed because such quotes are not economically meaningful. Options with ask prices less than or equal to bid prices were also discarded since they violate basic market microstructure rules. Additional filters removed incomplete observations and ensured that valid bid and ask quotes existed for both calls and puts. Call–put consistency checks were also applied to identify potential arbitrage violations. These steps help ensure that inaccurate or noisy quotes do not distort the shape of the volatility surface.

After cleaning the dataset, implied volatilities were extracted from option prices using the Black–Scholes model. Although the model assumes constant volatility, it is widely used to translate option prices into implied volatility values. Because option quotes are not evenly distributed across strikes and maturities, smoothing techniques are required to construct a stable volatility surface.

To obtain a smooth representation of the implied volatility surface, the Stochastic Volatility Inspired (SVI) parameterization was used. The SVI model represents total implied variance as a function of log-moneyness and allows the volatility smile to be captured using a small number of parameters controlling its level, slope, and curvature. The parameters were estimated by minimizing the difference between model-implied variances and observed market variances using the L-BFGS optimization algorithm.

Because the local volatility formula requires derivatives of option prices with respect to strike and maturity, additional smoothing techniques were applied to stabilize the surface and reduce noise. Methods such as isotonic regression and Savitzky–Golay filtering were used to ensure that the surface remained smooth and well behaved.

Once the implied volatility surface was constructed, the local volatility function was calculated using Dupire’s formula, which links local volatility to derivatives of option prices with respect to strike and maturity. By differentiating the smoothed option price surface, the model produces a volatility function that varies with both asset price and time while remaining consistent with observed option prices.
To validate the estimated local volatility surface, options were priced numerically using the Dupire forward partial differential equation. The PDE was solved using a Crank–Nicolson finite difference scheme, which provides a balance between numerical stability and accuracy. The resulting model prices were then compared with market prices using metrics such as mean absolute error and root mean squared error.

The results show that the local volatility model reproduces market option prices significantly more accurately than the constant-volatility Black–Scholes model. Overall, this project demonstrates a practical framework for constructing and validating a local volatility surface using market option data. The workflow combines careful data preprocessing, parametric surface calibration, and numerical methods to produce a volatility surface that is consistent with observed market prices and suitable for derivative pricing applications.

Vedashree Teli

Climate Risk: The Next Frontier In Quantitative Finance

The 2024 global insured losses from natural catastrophes came in at $140 billion, marking the fifth consecutive year that the figure has been above the 10-year average, as reported by Swiss Re. Nevertheless, financial institutions continue to rely on risk models based on historical return distributions that have zero forward-looking climate information. The disparity between physical climate exposure and its representation in balance sheet risk models is arguably the most important outstanding issue in quantitative finance today.

Regulatory pressure is accelerating. The Basel Committee on Banking Supervision, the ECB, and the Federal Reserve have each issued supervisory guidelines that require the integration of climate scenario information into stress testing models. The Network for Greening the Financial System (NGFS) has released standardized scenario pathways, including Orderly, Disorderly, and Hot House World, which financial institutions are increasingly expected to operationalize. The issue is not conceptual; rather, it is methodological in nature.

Why Existing Risk Models Fall Short

Standard VaR and Expected Shortfall models assume stationarity that return distributions are stable over time. Climate risk violates this structurally: physical risks are fat-tailed and shifting in severity, while transition risks introduce policy-driven structural breaks that historical data cannot anticipate. In my Quantitative Market Risk Engine project, stress scenarios under Basel III protocols identified a 48% higher tail loss exposure than conventional models suggested using historical shocks alone. Under climate-adjusted inputs, that divergence grows considerably larger.

Regime Detection and Transition Risk

Carbon-intensive sectors do not reprice gradually; they face sudden, policy-driven dislocations that behave like regime breaks. In my Market Regime Detection research, I built a Gaussian Mixture Model (GMM) framework across equity, volatility, and macroeconomic features, improving classification accuracy by 15% over volatility-only benchmarks. Conditioning regime states on carbon price trajectories and regulatory shock variables gives portfolio managers a probabilistic map of when transition risk materializes in asset prices, rather than treating it as a binary event.

Deep Learning and Climate-Linked Instruments

Catastrophe bonds, weather derivatives, and carbon allowances are the hardest instruments to price well; their risk drivers are physical climate variables, not financial ones, and standard models break down where volatility is structurally shifting. In my ongoing option pricing research, I am benchmarking Long Short-Term Memory networks against Black-Scholes on RMSE, with early results pointing toward meaningful accuracy gains in non-stationary environments. Better pricing means more efficient risk transfer and capital that can actually reach climate resilience.

The Road Ahead

NGFS scenarios, satellite-derived physical risk data, and a growing literature on climate-adjusted asset pricing are all available. What remains scarce is the practitioner layer, the people who can translate these inputs into production-grade systems that inform real capital allocation decisions.

That is the work I want to do. I am targeting roles in climate risk analytics within asset management, reinsurance, or financial regulation, focused on scenario analysis, physical risk scoring, and transition risk factor modeling. What drives me is seeing these tools deployed where they matter in the systems that price a future that looks less and less like the past.

Aleksey Urmanov

The Duplicity Of Market Microstructure

Everything works, everything, except common sense.

I once completed a multi-year project focused on analyzing the relationship between order flow, liquidity in the order book, and the price of an asset. My great revelation was that everything is a lot harder than it looked. This followed from a failure to understand why the price of an asset would steadily increase while the market was consistently selling in magnitudes dozens of times greater than any buying activity. This is something I observed was relatively common, in varying degrees of severity. The obvious explanation is that there is some incentive for the market to absorb the order flow while continually bidding up the price. However, that incentive is unknown to me or anyone else.

In a liquid market, the limit orderbook is the least static thing on planet earth. Every order that sits on the book is subject to change as the market price shifts around. There are bidding wars on the microsecond scale, spoofing, buffer blocks. Everything that happens in the orderbook reflects some utility function that is being maximized by some market participant. The money in the book will move around and the market will compete for price and order queue positioning even when there are no market orders. This means that the activity within the book, is a reaction to changes in the book itself as much as it is a reaction to changes in order flow. That self-reactionary behavior is what masks the relationship between order flow and liquidity and price.

The same qualities of an orderbook that seem to possess any predictive power are just as likely to be deceptive. If a participant uses the imbalance in an orderbook to predict short term movements, why wouldn’t another participant add some liquidity to one side of the market to skew the imbalance and give the appearance of pressure. For them, there is no obligation for their liquidity to transact in the market, they can move it around in perpetuity while continuing to masquerade as some form of buying or selling pressure. This exact kind of pseudo-trading activity is just another driver of difficulty in understanding the relationships between liquidity and price.

Certain autonomous market making algorithms will even perform “dances”. If you can find a market that is very stale and wait long enough, you may see the orderbook flutter and create little patterns. Imagine equalizer bars in audio equipment, flashing and wiggling up and down. It’s quite funny because you know some developer was having a slow day and decided to integrate some “dance code” into the production grade stack that is intended to provide liquidity to markets. Take what you will with this information.

Anand Vadlamani

Beyond VaR: Measuring Tail Risk In A Fixed Income Portfolio

One of the more humbling realizations in quantitative finance is how often the models we trust fail precisely when we need them most. Value at Risk (VaR) — the industry-standard measure of market risk — is a prime example. It tells you the loss you won't exceed 95% or 99% of the time. But it says nothing about what happens in that remaining 1%. During my Interest Rate Risk and Tail Loss Modeling project, I set out to stress-test that blind spot on a real bond portfolio, and the results changed how I think about risk measurement entirely.

Building the Portfolio and Quantifying Rate Sensitivity

I started by constructing a multi-maturity bond portfolio and computing the classic interest rate sensitivity measures: DV01, duration, convexity, and key rate shifts across the yield curve. This step — often treated as routine — turned out to be revealing. Duration gave a reasonable first-order picture of rate sensitivity, but convexity showed that for larger rate moves, the linear approximation breaks down meaningfully. Key rate duration analysis further exposed which points on the curve drove the most risk, which is something a single duration number simply cannot capture. These findings reinforced a key FRM Part 2 principle: interest rate risk is not one-dimensional.

Three Methods, One Question: How Bad Can It Get?

With the portfolio characterized, I estimated 95% and 99% VaR using three methodologies: historical simulation, parametric (variance-covariance), and Monte Carlo. Each method told a slightly different story. The parametric approach, which assumes normally distributed returns, consistently underestimated risk relative to the other two — a direct consequence of ignoring fat tails. Historical simulation preserved the actual empirical distribution of past rate moves, including crisis periods, and produced more conservative estimates. Monte Carlo offered the most flexibility, allowing me to simulate a wide range of interest rate scenarios including non-linear shocks. I then validated all three using backtesting, checking how often realized losses actually exceeded the predicted VaR thresholds. The parametric model failed backtesting more frequently than acceptable — a concrete reminder of why model risk is taken so seriously in the FRM curriculum.

Where VaR Ends: CVaR and Extreme Value Theory

The most intellectually rewarding part of the project was moving beyond VaR into Conditional Value at Risk (CVaR) — also known as Expected Shortfall — and Extreme Value Theory (EVT). CVaR answers the question VaR avoids: given that we are in the worst 1% of outcomes, what is the average loss? This measure is not only more informative but is now preferred under Basel III and IV frameworks precisely because it is more sensitive to tail severity. To push further, I fitted a Generalized Pareto Distribution (GPD) to the tail of the loss distribution using EVT. The GPD tail was notably heavier than what any of the three VaR models assumed, suggesting that in genuine stress scenarios — a 2008-style rates spike, a sudden credit event — the portfolio's losses could be substantially larger than standard risk models would predict. Stress testing with historically calibrated shocks confirmed this: the portfolio showed meaningful vulnerability to parallel upward yield curve shifts, particularly at the long end.

What This Project Taught Me

This project sat at the intersection of everything I find compelling about financial mathematics: rigorous theory, real data, and decisions with genuine consequences. Preparing for FRM Part 2 has deepened my appreciation for why regulators and risk managers care so much about the assumptions baked into these models. A number on a risk report is only as trustworthy as the model behind it, and this project gave me a firsthand look at where those models agree, where they diverge, and what happens in the space they all fail to capture. It is that last question — the tail — that I want to keep exploring as I move into my career in quantitative risk.

Isha Wadekar

Modeling Market Asymmetry: A GJR-GARCH Approach To Volatility Clustering

In the sphere of quantitative finance, the ability to accurately forecast volatility is the cornerstone of robust risk management and derivative pricing. As a first-year student in the Financial Mathematics program at NC State, I am currently leading a research project titled "Modeling Asymmetric Volatility Clustering in Financial Assets using GJR-GARCH". Working alongside my team and advisors, Dr. Tao Pang and Mr. Patrick Roberts, we are investigating how different asset classes react to market shocks.

The Challenge of Asymmetricity

Our research is grounded in the phenomenon of volatility clustering, where "large changes tend to be followed by large changes, of either sign". While standard GARCH models are effective at capturing this clustering, they often fall short because they assume a symmetric response to market news. Empirical evidence suggests a leverage effect: negative returns (bad news) typically increase volatility more significantly than positive returns (good news) of the same magnitude.

To capture this, our team is utilizing the GJR-GARCH model. By incorporating an indicator function, this model specifically accounts for the stronger impact of negative shocks on conditional variance:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \gamma\varepsilon_{t-1}^2\mathbf{1}_{\{\varepsilon_{t-1}<0\}} + \beta\sigma_{t-1}^2$$

Asset Analysis

We have applied this framework to a diverse set of financial assets to test its versatility:

- Index: S&P 500 (Jan 2015 – Dec 2025)
- Commodity: WTI Crude Oil (Jan 2021 – Feb 2026)
- REIT: Prologis (Jan 2009 – Dec 2023)

Assumption Testing

Before modeling, we performed rigorous assumption testing. Using the Augmented Dickey-Fuller (ADF) Test, we confirmed that the log returns for these assets are stationary. Furthermore, Ljung-Box Test results indicated significant autocorrelation, confirming that the variance is not constant and follows a pattern suitable for GARCH-family modeling.

GJR-GARCH Model Implementation

A critical development in our project is the move away from the assumption of normality. Financial return distributions frequently exhibit fat tails (excess kurtosis), meaning extreme gains or losses occur more often than a normal distribution predicts. To address this, we have successfully fitted a GJR-GARCH(1,1,1) engine using a Student's t-distribution. This allows us to simultaneously capture the leverage effect and the higher likelihood of rare market events.

Future Work

As we move into the final phases of the project, our focus shifts toward Model Validation and Risk Forecasting. Our next steps include:

- Diagnostic Checking: Using Ljung-Box diagnostics on standardized residuals to ensure the model has successfully captured all underlying market patterns.
- Conditional Volatility Forecasting: Generating forecasts to predict how current risk regimes will evolve, providing a practical tool for anticipated market shifts.

Leading this project has been a profound lesson in the necessity of model precision. It has bridged the gap between theoretical stochastic processes and the practical realities of market asymmetry. I look forward to presenting our final results and exploring how these volatility insights can be integrated into broader algorithmic trading strategies.

Yueyang Wang

Monte Carlo Simulation for Exotic Options

This semester, I had the opportunity to study Monte Carlo simulation during my second semester in the Financial Mathematics Master's program. Before taking this course, I assumed Monte Carlo was a relatively simple method: simulate many paths, observe what happens, and average the results. That intuition is not entirely incorrect, but it misses the real depth of the subject. Monte Carlo methods are not just about producing random paths. They are a flexible numerical framework for valuing complex financial products, measuring uncertainty, and approximating quantities that are difficult or impossible to compute in closed form.

Monte Carlo simulation is a method that uses repeated random sampling to estimate unknown values. Rather than producing one deterministic answer immediately, it generates many possible outcomes and studies their overall behavior. From these simulated trials, one can estimate prices, probabilities, risks, and distributions. In finance, this makes Monte Carlo especially valuable because uncertainty is built directly into asset-price models. Instead of asking for one future price, Monte Carlo asks what many possible future price paths might look like and then uses those paths to estimate the value of a contract.

This feature becomes particularly important for exotic options. For a standard European option, the payoff depends only on the asset price at maturity. In contrast, many exotic options are path-dependent, meaning that the payoff depends on the behavior of the asset price over time rather than only at expiry. As a result, knowledge of the terminal distribution alone is not sufficient for pricing these products.

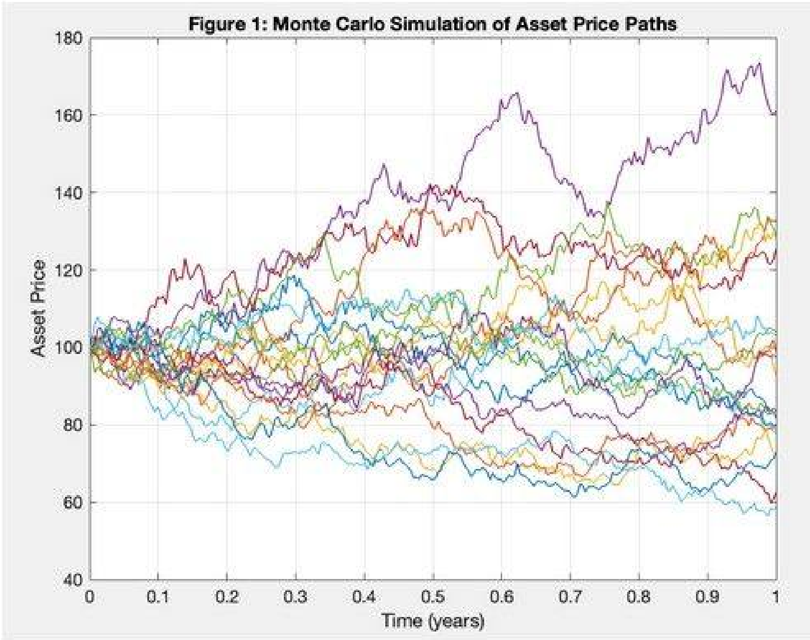

figure 1

Figure 1, Monte Carlo Simulation of Asset Price Paths, illustrates this idea by showing a set of simulated asset-price paths. Even though all of the paths start from the same initial value, they evolve very differently over time. Some paths trend upward, others decline, and others fluctuate sharply before maturity. This diagram captures the central logic of Monte Carlo simulation: instead of relying on a single forecast, the method generates many possible futures for the same underlying asset. For vanilla options, only the final point of each path may matter. For exotic options, however, the entire path can be relevant.

A clear example is the lookback option, whose payoff depends on the maximum or minimum asset price reached during the life of the option. In this case, two paths with the same terminal value may still lead to different payoffs if one path reached a higher maximum or lower minimum than the other. This is exactly why Monte Carlo simulation is so useful: it allows the pricing algorithm to track the full evolution of the asset price and record path-dependent quantities such as running maxima or minima.

Continue Next Page

Yueyang Wang

Monte Carlo Simulation for Exotic Options

Continued From Previous Page

Another important example is the barrier option, where the payoff depends on whether the asset crosses a specified barrier level before expiry. This introduces an important source of discretization error. In continuous-time models, the asset may cross the barrier between two observation dates, even if the sampled values themselves do not show a crossing. Figure 2, Barrier Crossing Missed by Discrete Monitoring, demonstrates this problem clearly. The fine simulated path rises above the barrier, but the coarse monitoring points fail to capture that event. A naive discrete-time Monte Carlo method would therefore miss the barrier crossing and may assign the wrong option value. This shows that accuracy depends not only on the number of simulated paths, but also on how finely the time interval is divided. To reduce discretization error, one may use a finer time grid or apply a Brownian bridge correction to estimate what happens between simulated dates.

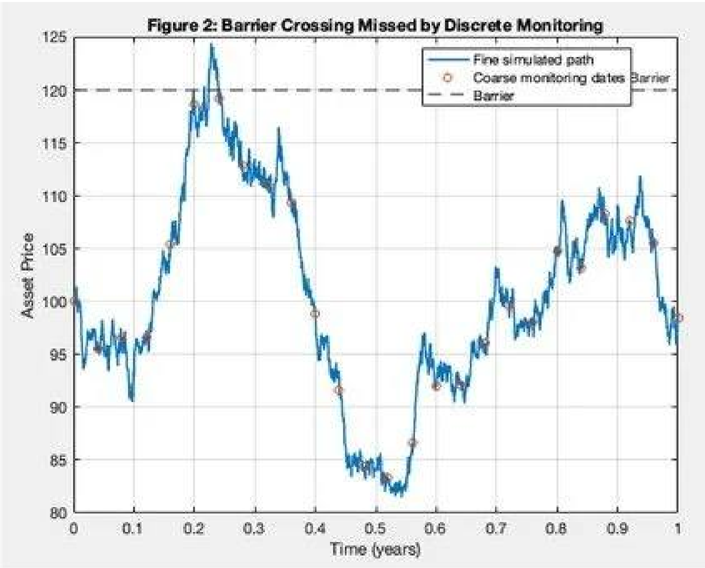

figure 2

Monte Carlo methods are also subject to statistical error. Since the option value is estimated from a finite number of random paths, the result is itself random and will vary from one simulation to another. Even if the model and time discretization are correct, too few simulated paths may produce a noisy estimate. Statistical error can be reduced by increasing the number of simulations, although that also increases computational cost. In practice, effective Monte Carlo pricing requires a balance between controlling discretization error in the time grid and controlling statistical error in the estimator.

Overall, Monte Carlo simulation changed from being, in my mind, a simple trial-and-error technique into a sophisticated and powerful method for pricing exotic options. It is valuable because financial derivatives often depend on uncertainty in a highly detailed way. As Figure 1 and Figure 2 show, the path itself matters, and numerical errors can significantly affect the result. For path-dependent exotic options, Monte Carlo provides a natural and effective framework for valuation, provided that both discretization error and statistical error are handled carefully.

Zach Wiebe

Musings on Symmetry and Clarity

This semester I was fortunate enough to be appointed an ambassador of the financial mathematics program, meaning I have had the opportunity to lead an industry project with a group of my peers.

A lifetime ago, in August 2025, I was just getting started with my fundamentals of statistical inference coursework and had been staring at normal distributions for far too long—long enough that I developed a hunch about a new symmetry hidden within the bell-curve, relating its tails to its peak. I leveraged what I remembered of my topology and non-Euclidean geometry undergraduate coursework to formalize this symmetry, and I became increasingly convinced I was closing in on a significant insight. Both frustratingly and excitingly, my research turned up no evidence of prior art.

“Now, Zach,” one might wonder, “Surely spinning normal distributions around won’t have any practical applications.” I grant that a transformation as described is largely just a change in perspective—but so is any transformation. That doesn’t stop the Fourier transform from having important signal processing applications. I would draw an analogy between this project’s transformation and a conversion between degrees and radians. There is nothing one can do with radians that cannot be done with degrees—but the intuition that working in radians provides has sparked many innovations in the 300 years they have been in use. Radians provide a framework for applying length analysis methods to angles, just as I hope to apply mean analysis methods to tails. To paraphrase Dr. Ghosh from a particularly enlightening statistical inference lecture, “Everyone knows how to deal with uncertainty around the mean. The problem is uncertainty around the tails.” Well, the goal of this project is to put the tails in terms of the means, thereby creating a way to deal with them. I whipped up a rudimentary Python program to further analyze this symmetry, but the project quickly fell through the cracks as the semester picked up the pace.

Returning to this project as an ambassador was too tempting a possibility to pass up. I had just enough of a foothold in this project that I was confident I could develop a semester-long gameplan to engage a 5-person team—there was way more to discover here than I could do alone, but not so much that the team would not have resume-worthy results by semester’s end.

Now, with Spring break approaching, the project is starting to show some intriguing results thanks to the efforts of my amazing team—Sartaz Aziz, Nick Moss, Ant Rautela, and Catherine Sexton. We’ve stumbled across plenty of difficulties that my initial investigation last semester failed to reveal, but our combined brainstorming has produced just as many solutions.

Chief among these difficulties has been the rather practical issue of taking a step back from the abstract math that we’ve been immersing ourselves in for months to communicate our project goals and methods with classmates who have spent all that time learning very different things in very different projects. We spent most of a meeting locking in terminology to reduce confusion even within the team and discussing methods to convey all of these ideas that we look forward to showing off in the upcoming final project presentations at the end of the semester. One of my go-to methods is analogy: my sister once needed help calculating volume by rotating an integrated function around the x-axis and imagining that volume as a deli sandwich crafted by stacking slices of area—meat, cheese, lettuce, bread—along the toothpick holding everything together—the x-axis—helped her a lot. We’re still searching for a suitable analogy for our project; but, even if we can’t find one, it’s sure to be a presentation to remember.

Derek Wu

A Calculated Decision: Why NC State's MFM Was The Right Fit

My pursuit of a Master of Financial Mathematics (MFM) at NC State has been an immensely impactful chapter in my life. Returning to school after several years of professional experience in an adjacent field was not an easy decision, but in doing so the program has profoundly reshaped my professional ambitions and introduced me to incredible opportunities.

At the heart of my decision to join NC State was a deep-seated love for mathematics and its ability to make sense of the financial world. Seeing the potential of quantitative methods to solve intricate financial problems during my personal studies and projects sparked a desire to go deeper. I wanted a program that could unite these interests, and I've found that perfect blend here.

What truly sets this program apart is its commitment to career development; the support feels personal from the very beginning. Before classes even started, their career services department reached out to learn about my individual story and where I hoped to go. They worked with me one-on-one to shape a resume that authentically represents my skills for the roles I was targeting, offering guidance that feels tailored, not templated.

The interview preparation has been a huge help as well; going through multiple mock interviews built my confidence and helped me find my voice. Even with years of professional experience under my belt, I realized I had never focused on developing those essential soft skills. I'm also incredibly grateful for the semester-long projects where we tackle real problems with guidance from senior students. For me, this truly bridged the gap between theory and practice.

Woven throughout it all are the industry guest speakers, who pull back the curtain on what different quantitative finance roles look like day-to-day. Their honest perspectives were invaluable in helping me envision my path forward. I have also been encouraged to participate in several career fairs and networking events that have broadened my professional network. These experiences have given me a clearer sense of how to align my skills with industry needs, particularly technical knowledge and risk analysis.

My time in NC State's MFM program has been a journey of challenge, growth, and opportunity. The rigorous coursework, practical projects, and invaluable professional guidance have fully prepared me for a future in finance. I am deeply grateful to the supportive faculty and for a curriculum that fostered my personal and professional development. This program has not only defined my career path but has also given me the confidence to navigate the dynamic world of finance.

Joy Wu

Modeling Market Volatility Using GARCH Family Models: Evidence From The S&P 500

Introduction

Volatility is a key measure of risk in financial markets and plays an important role in portfolio management, derivatives pricing, and risk control. Financial return series are known to exhibit several stylized facts, such as volatility clustering, where large price movements tend to be followed by further large movements. These characteristics make traditional constant-variance models inadequate for capturing the dynamic behavior of financial volatility.

This project focuses on modeling the volatility of the S&P 500 index using econometric models from the GARCH (Generalized Autoregressive Conditional Heteroskedasticity) family. The goal is to examine time-varying volatility dynamics and evaluate the forecasting performance of different GARCH-type models.

Data and Methodology

The dataset consists of daily closing prices of the S&P 500 index from 2023 to 2025. To analyze return dynamics, prices were transformed into log returns:

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

Where P_t denotes the closing price at time *t*.

Preliminary analysis showed that the conditional mean of returns contains little predictable structure. After testing several specifications, the mean process was modeled as an ARIMA(0,0,0), suggesting that the return series behaves approximately like white noise in the mean equation. As a result, the main focus of the analysis was on modeling the conditional variance of returns.

Volatility Modeling with the GARCH Family

My primary contribution to this project involved implementing and comparing several volatility models from the GARCH family. These models are designed to capture time-varying volatility and persistence in financial return series.

The comparison of these models highlights the importance of capturing asymmetric volatility dynamics in financial markets. Equity returns often respond differently to positive and negative shocks, a phenomenon known as the leverage effect. Models such as GJR-GARCH and EGARCH allow volatility to react more strongly to negative returns, which helps better describe the risk dynamics observed in equity markets like the S&P 500.

Implementing these models helped deepen my understanding of how volatility evolves over time and how econometric models can be used to quantify market risk.

GARCH Model

The baseline specification used in this analysis is the GARCH(1,1) model, which models conditional variance as a function of past squared shocks and past conditional variance.

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

This model captures volatility persistence commonly observed in financial markets.

GJR-GARCH

To account for asymmetric volatility responses, the GJR-GARCH model was also implemented. Financial markets often exhibit the leverage effect, where negative shocks have a stronger impact on future volatility than positive shocks.

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \gamma I_{t-1}\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

where I_{t-1} is an indicator function that equals 1 when the shock is negative and 0 otherwise.

Continue Next Page

Joy Wu

Modeling Market Volatility Using GARCH Family Models: Evidence From The S&P 500

Continued From Previous Page

EGARCH

The EGARCH model was also implemented to capture asymmetric volatility dynamics in an alternative framework. Instead of modeling variance directly, EGARCH models the logarithm of conditional variance.

$$ln(\sigma_t^2) = \omega + \beta ln(\sigma_{t-1}^2) + \alpha|z_{t-1}| + \gamma z_{t-1}$$

This specification allows the model to capture asymmetric effects while ensuring the positivity of volatility.

Forecast Evaluation

To evaluate model performance, one-step-ahead volatility forecasts were generated for each model. Forecast accuracy was assessed using several evaluation metrics, including Mean Squared Error (MSE), Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the QLIKE loss function.

Conclusion

This project demonstrates how GARCH family models can effectively capture time-varying and asymmetric volatility dynamics in financial markets. By comparing several volatility models, the analysis provides insights into how different model structures describe the behavior of market risk.

Shixun Wu

From Models to Decisions: My Career Goals in Financial Mathematics

I came to the Financial Mathematics program because I wanted a deeper toolkit, but also a clearer way to use that toolkit in real decisions. In my past roles as a quantitative analyst, I worked across the electronic, EV, and banking sectors. I enjoyed the craft of building models, cleaning data, and testing assumptions, yet I also saw how easy it is for a technically correct model to miss the point if it is not connected to a decision, a risk limit, or a business constraint.

My current career goal is to become a quantitative researcher who can bridge theory and implementation. I want to work on problems where market dynamics, risk management, and statistical learning meet, and where the final output is something a team can trust and act on. That could be research for pricing and hedging, portfolio construction, or risk measurement, but the common thread is the same: build models that are rigorous, testable, and operational.

In the near term, I am focusing on three skills that I think define strong quant work. First, I want to be faster and more disciplined with model validation. That means stress testing, sensitivity analysis, and clear diagnostics, not just a high backtest metric. Second, I want to improve my ability to translate a mathematical idea into clean, efficient code, with careful documentation and reproducibility. Third, I want to strengthen my judgment around uncertainty and model risk, especially in settings where data is limited or regimes shift.

The FM curriculum has been shaping how I frame these goals. When we work through stochastic calculus and measure changes, the value is not only in passing an exam or finishing a proof. It is in learning a structured way to ask what is being modeled, what is being assumed, and what could break. When we study numerical methods and simulation, I am reminded that implementation details matter, and that accuracy, stability, and speed are part of the same conversation. That mindset is something I want to carry into any research environment.

Looking ahead, I aim to join a team where I can contribute to research, but also learn how research becomes production. I want to participate in the full cycle: problem definition, data and feature design, modeling, calibration, validation, and deployment with monitoring. Over time, I hope to specialize in areas where my background fits well, such as systematic strategies tied to fundamental and sector signals, or risk frameworks that connect market behavior to capital and liquidity constraints.

Ultimately, I want my work to be useful in the way good quant work should be useful: it should clarify trade-offs, quantify risks, and improve decisions under uncertainty. The technical part is important, but so is communication. A model that cannot be explained clearly is hard to govern and easy to misuse. My goal is to build both: strong mathematics and clear, practical reasoning that a broader team can rely on.

Daniel Yang

Bridging Theory and Practice in the NC State Financial Mathematics Program

One of the most impactful aspects of the Financial Mathematics program at North Carolina State University has been the program's strong emphasis on connecting rigorous mathematical theory with real-world financial applications. Coming into the program, I had a solid background in mathematics and statistics, but what truly deepened my understanding was seeing how concepts like stochastic processes, optimization, and deep learning are used to model markets, price derivatives, and manage risk in practice.

In particular, the coursework that integrates quantitative theory with computational implementation has been especially valuable. Rather than stopping at derivations, we are encouraged to build simulations, implement pricing algorithms, and analyze real financial data. This hands-on approach has made topics like Monte Carlo methods, time-series modeling, and portfolio optimization much more concrete. Writing code to test ideas forced me to confront the assumptions behind models and to understand how sensitive results can be to parameter choices, estimation error, and market noise. That experience has been far more impactful than learning formulaic coursework in isolation.

Another part of the program I've found meaningful is the interdisciplinary nature of the curriculum. Financial mathematics sits at the intersection of probability, statistics, economics, and computer science, and the program reflects that. Core courses and projects draw on mathematic ideas behind machine learning and financial economics, which has broadened my perspective on how quantitative tools can be applied across domains. Furthermore, the wide variety of high level subjects (like optimal transport, efficient deep learning, nonparametric Bayesian inference) taught at the various departments of State gives students an understanding of the forefront of financial research and their implementations. This interdisciplinary approach has also aligned well with my mixed research background, allowing me to see new ways to apply industry-relevant skills in finance.

The program has also been impactful in terms of professional development. Through seminars, projects, and collaboration with classmates, I've gained a better understanding of how quantitative work is communicated in industry settings. Explaining models clearly, presenting results with appropriate caveats, and working in teams are skills that are often overlooked in purely theoretical programs but are emphasized here. These experiences have helped me feel more prepared to contribute in a professional quantitative role.

Finally, the collaborative environment within the cohort has made a lasting impression. Working with peers from diverse academic and professional backgrounds has pushed me to think differently about problems and exposed me to alternative modeling approaches. Study groups, project collaborations, and informal discussions have often been where complex concepts finally clicked, and they have reinforced how important communication and teamwork are in quantitative finance.

Overall, the most impactful aspect of the Financial Mathematics program at NC State has been its ability to bridge theory and practice. By combining rigorous mathematical training with computational implementation, interdisciplinary learning, and collaborative problem-solving, the program has strengthened both my technical foundation and my ability to apply quantitative methods to real financial problems.

Peng Yang

High-Frequency Volatility Prediction Using Deep Learning

Financial markets generate massive volumes of high-frequency data every second. Accurately forecasting short-term volatility from this data is valuable for market makers, quantitative traders, and risk management systems. In this project, I developed a deep learning framework to predict short-term stock price volatility using order book data from high-frequency trading environments.

The dataset consisted of limit order book information containing the top two levels of bid and ask orders for multiple stocks. Each observation included price and volume information for buy and sell orders, recorded at very fine time intervals. Because raw order book data is extremely granular and noisy, the first step of the project focused on building a robust data processing pipeline.

Using Python and the Polars data processing library, I implemented efficient data transformations to aggregate raw book events into meaningful features. For each stock and time interval, I computed market microstructure indicators such as bid–ask spread, total trading volume, and volume imbalance between buyers and sellers. These features capture the dynamics of supply and demand within the order book and provide signals related to short-term volatility. I also performed exploratory data analysis to verify feature distributions and detect missing values, ensuring data quality before model training.

After feature engineering, the dataset was structured as a time series problem. The objective was to predict future realized volatility based on historical order book signals. To model temporal dependencies, I experimented with transformer-based architectures, which are widely used in sequence modeling due to their ability to capture long-range relationships through attention mechanisms.

The model architecture consisted of an embedding layer to project engineered features into a higher-dimensional representation, followed by multiple transformer encoder layers. These layers allow the model to weigh the importance of different time steps in the input sequence when generating predictions. Compared to traditional models such as linear regression or simple recurrent networks, transformers can better capture complex interactions across time.

To evaluate the model, the dataset was divided into training and validation sets based on time periods to avoid information leakage. Performance was measured using root mean squared error (RMSE) between predicted and realized volatility. Preliminary experiments demonstrated that the transformer model outperformed simpler baseline models, suggesting that attention-based architectures can effectively learn patterns in high-frequency financial data.

In addition to model development, I focused on building scalable workflows for processing large datasets. High-frequency order book data can easily reach tens of millions of records, so computational efficiency was critical. By leveraging columnar data formats such as Parquet and using vectorized operations in Polars, the data pipeline remained both memory-efficient and fast.

Overall, this project provided hands-on experience with financial data engineering, feature design based on market microstructure theory, and deep learning techniques for time series forecasting. It also highlighted the challenges of working with high-frequency financial datasets, including noise, extreme data volume, and the need for careful validation strategies. Future work will involve experimenting with additional architectures, incorporating more order book levels, and optimizing the model for real-time volatility forecasting.

Jialin Zhang

From Models to Markets: What a Volatility Practitioner Taught Me About Financial Mathematics

One of the most impactful parts of the Financial Mathematics (FM) program for me so far has been how consistently it connects classroom models to the way real markets behave—especially through guest speakers who have lived inside the “messy middle” between theory and trading. A recent talk by Dennis Davitt (CIO/CEO of Millbank Dartmoor Portsmouth) crystallized that bridge for me and reshaped how I think about derivative modeling, risk, and career direction.

In class, we learn the elegant core: stochastic processes, Black–Scholes assumptions, and discrete-time approximations like binomial trees. Those tools are powerful, but it’s easy to treat them as self-contained mathematical objects. Davitt’s career story—spanning floor trading, investment banking, and asset management—highlighted that the “same” model can be used in very different environments, and the environment often determines whether a strategy is robust or fragile. On the trading floor, he emphasized liquidity and adaptation; in investment banking, leverage and balance-sheet constraints change everything; and in asset management, even being “right” can still get you fired if the path is painful or the risk story is misunderstood.

One message I found especially memorable was: **interest rates change the entire investing playbook**. The presentation framed the zero-rate era as an abnormal period that encouraged yield-seeking and option-selling behavior, while higher-rate regimes make different structures and hedges more feasible. That perspective made me re-evaluate how I interpret volatility regimes: not as abstract parameters to estimate, but as market states shaped by policy, portfolio constraints, and investor behavior.

He also discussed current derivative market themes—like the growth of zero-days-to-expiry (0DTE) options—and why option strategies can “blow up.” The takeaway wasn’t simply “options are dangerous,” but that blow-ups often share common threads: leverage and mismarked positions, plus a failure to understand how exposures evolve intraday when the market moves. For me, this connected directly to what we do in the FM program: we simulate paths, compute Greeks, and talk about hedging, but the practical question is always, “How does the exposure change when assumptions break?”

This experience also clarified my current career goals. I’m increasingly focused on roles where I can combine quantitative modeling with risk communication—such as quantitative research, risk analytics, or derivatives-focused portfolio analysis. I enjoy building models, but I’m equally motivated by the responsibility of explaining what those models mean to decision-makers: what can go wrong, how regimes shift, and how to design strategies that remain investable under stress. The FM program is helping me build that full-stack mindset—from code and math to interpretation and implementation.

Ultimately, what I’ve enjoyed most is realizing that financial mathematics is not only about finding the “right formula.” It’s about building a disciplined way to reason under uncertainty, staying humble about incentives and constraints, and translating model outputs into decisions that can survive the real world. That’s the kind of impact I was hoping for when I joined the program—and it’s why I’m excited for what comes next.

Luolin Zhang

Where Theory Meets Practice: My Experience In Financial Mathematics

One aspect of the Financial Mathematics program that I have found especially impactful is how strongly it connects rigorous quantitative training with real industry applications. The coursework is mathematically challenging, while also emphasizing how these tools are used to solve practical financial problems.

A highlight of the program for me has been the weekly seminar series. Industry professionals from a wide range of organizations share their career paths and insights into how quantitative skills are applied in their work. I particularly appreciate the diversity of speakers, from large financial institutions to startups, which provides a broader view of the financial industry.

These seminars have helped me better understand areas of finance that I previously knew little about. For example, one talk introduced how actuarial modeling is used in the insurance industry to design and price products. Seeing how quantitative models directly influence real financial decisions made the concepts we learn in class feel much more tangible.

The program also places strong emphasis on project-based learning. In one project, I worked on modeling mortgage default risk using loan-level data. More recently, I have been working on forecasting market volatility using models such as GARCH and LSTM and applying these forecasts to option pricing.

Experiences like these make the program both rigorous and practical, and they have helped me better understand how quantitative finance is applied in real-world settings.

Quantitative Finance Workshop 2026

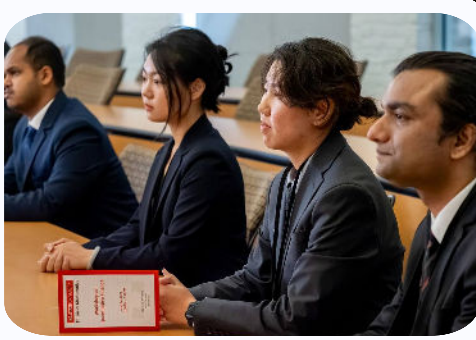

Quantitative Finance Reception

SAS Hall at NCSU's Main Campus

NC STATE UNIVERSITY

Financial Mathematics Graduate Program

Ranked # 5 by Risk.net
2025

Ranked #9 by QuantNet
2026

Financial-mathematics@ncsu.edu
https://financialmath.sciences.ncsu.edu