Academic Background

I entered Boston University as a physics major but found myself drawn to the underlying mathematical structure. After taking quantum mechanics, I realized my interest lay less in the physics itself, and more in the probabilistic nature that describe these systems. From there, I shifted to applied mathematics with a physics minor, and have since focused my coursework on rigorous probability theory and stochastic processes – including graduate courses in Probability Theory (MA 779), Advanced Stochastic Processes (MA 783), and Real Analysis (MA 711).

These courses have been engaging both conceptually and in application. We study measure theory, convergence of random variables, martingales, stopping times, and more. I've already applied this material directly – using martingale representation and Feynman-Kac in my thesis work. Together, these courses have prepared me to understand and engage with research-level work in stochastic control and probability for machine learning.

Research Experience

My senior thesis under Professor Mark Kon sits at the intersection of Bayesian inference, stochastic control, reinforcement learning, and mathematical finance. I extend the continuous-time portfolio optimization framework of Wang and Zhou (2019) to a setting with unknown model parameters, where the drift is learned through Bayesian updating. I justify the Kalman-Bucy filter for the return processes and prove the optimal control is Gaussian. A belief-augmented Hamilton-Jacobi-Bellman (HJB) equation is then formulated on an extended state space that includes both wealth and posterior parameters. I then identify a bilinear wealth drift term, which breaks the classical Linear-Quadratic-Gaussian (LQG) Riccati structure. This motivates my current work on analyzing the resulting value function. (Where I am as of writing this SOP)

In summer 2024, I worked as a research intern at HKUST under Professor Yang Can, during which I was first introduced to modern time-series transformers. I later implemented PatchTST architectures for cryptocurrency volatility forecasting, benchmarking them against GARCH and Kalman filter baselines. The transformer-based model achieved 35% reduction in MSE and improved directional accuracy, demonstrating that attention-based architectures capture nonlinear and non-Gaussian patterns that classical models miss. Overall, this project taught me rigorous evaluation methods, as well as how to connect theoretical forecasting methods to high-noise financial environments.

Outside formal research appointments, I have pursued independent projects also involving reinforcement learning and portfolio optimization. I developed a Proximal Policy Optimization (PPO)-based portfolio allocator for ETFs and evaluated it against rolling Markowitz optimization. In backtests, the PPO strategy achieved a Sharpe ratio of 2.00 compared to 1.75 for the Markowitz, and ranked in the top 0.01% of 1 million Monte Carlo simulated portfolios. Building this system required me to design a reward structure that balanced risk and return, incorporating transaction costs and adequate diversity. Validation was also performed on multiple disjoint windows, with a worst-window penalty, for robustness. This experience opened my eyes to the strengths and fragility of RL-based strategies in noisy environments, and only furthers my interest in this area.

Graduate Study Goals

In graduate study, I hope to work on problems at the intersection of probability, stochastic control, and reinforcement learning. Although my work has been finance-related, broadly speaking, I am interested in algorithmic decision-making. More specifically, I am interested in: (1) Stochastic control with learning, where we estimate parameters or latent states while acting optimally; (2) Probabilistic analysis of reinforcement learning algorithms, particularly their convergence properties and stability in stochastic environments; and (3) mathematical approaches to algorithmic trading that integrate modern ML components with classical stochastic calculus. I would like my work to sit on rigorous mathematical foundations; martingale methods, HJB equations, weak convergence of controlled processes; all while being closely connected to applicable algorithms.

I am drawn to departments with active research in reinforcement learning and optimization for machine learning, stochastic control, and mathematical finance. What excites me the most about research is contributing to the theoretical frameworks used in industry. More concretely, I am motivated by the mathematical justification of algorithms used in practice by the ML and quantitative finance communities.

Career Goals

My long-term goal is to work as a quantitative researcher in finance, focusing on portfolio optimization. I see a research-oriented PhD as an 'optimal' path, allowing me to build a strong foundation in probability and stochastics. Furthermore, I will learn to formulate and solve nontrivial problems, and ideally produce valuable research. In the longer run, I hope to use my training to tackle real-world RL and financial problems, where we have expensive large-scale optimization, sample inefficiency, and convergence issues.

Throughout my education and work, I have enjoyed formulating mathematical solutions to real-world problems. I am eager to continue this path through my graduate study, both deepening my understanding of probability and contributing to the fields of learning and finance.