STEVENS
INSTITUTE OF TECHNOLOGY
1870

Ph.D. Dissertation Defense

Candidate:	Cheng Lu
Degree:	Doctor of Philosophy
School/Department.:	School of Business / Financial Engineering
Date:	Monday, May 5, 2025
Time:	2:00 – 4:00 pm
Location:	Babbio 601
Title:	Advancement of Reinforcement Learning in Asset Allocation and Pricing
Chairperson:	Dr. Majeed Simaan, Finance/FE, School of Business
Committee Members:	Dr. Zachary Feinstein, Financial Engineering, School of Business
	Dr. Papa Momar Ndiaye, Financial Engineering, School of Business
	Dr. Jordan Suchow, Information Systems, School of Business

Abstract

Despite its normative appeal, portfolio theory often faces significant challenges in practical implementation due to estimation errors. These frictions manifest as model risk, the uncertainty stemming from the parameters used in the model's input -- commonly known as *estimation risk.* This dissertation introduces reinforcement learning (RL), a machine learning technique designed for sequential decision-making, as a novel approach to mitigate estimation risk in asset allocation and pricing. Overall, it highlights two key contributions of RL in finance. First, it demonstrates RL's ability to address estimation risk while preserving tractability in risk-managed portfolios. Second, it shows the value of embedding domain knowledge or equilibrium principles into both asset allocation and pricing.

The dissertation comprises two main parts: the first examines how RL can enhance portfolio selection, whereas the second part explores its implications on asset pricing. In the first part, a data-driven RL approach is proposed for improving correlation matrix estimation in high-dimensional portfolios. The proposed framework enhances the traditional shrinkage framework by incorporating text-based networks (TBN) constructed from 10-K filings, thereby relaxing restrictive assumptions on asset returns. Empirical analysis utilizing over 400 assets across two decades demonstrates that RL-based estimators yield reduced out-of-sample volatility, higher Sharpe ratios, and lower downside risk -- net of transaction costs.

The second part of the dissertation utilizes RL to emulate a non-myopic financial agent, extracting a forward-looking market valuation. In turn, the framework reverse-engineers these valuations to form expectations about the equity risk premium. Unlike traditional statistical learning models, the proposed approach embeds equilibrium mechanisms with respect to market-wide states/factors. Compared to 23 benchmarks, the proposed RL-based framework achieves the highest performance in terms of *recall*, i.e., correctly identifying true market downturns. This evidence is most pronounced during recessionary periods when hedging is most critical.