

DISSERTATION DEFENSE

Kevin Pierre Mott

“Finance-Informed Neural Networks – Deep Learning for Functional Problems in Macroeconomics and Finance”

Thursday, October 30, 2025
10:30am
Tepper 4242

Student Debt (Forgiveness) in General Equilibrium

This chapter evaluates the Biden Administration's proposed student loan forgiveness policy in a stochastic overlapping generations model with 60 periods of life and three household types differentiated by student debt-to-income ratios. The central finding is that student loan forgiveness generates minimal real economic effects, contradicting the policy's stated objective of promoting wealth accumulation among over-leveraged borrowers. Despite reduced debt burdens, borrowers allocate the transfer primarily toward consumption rather than retirement savings or productive investment, leaving aggregate capital, production, wages, and asset prices virtually unchanged. For non-borrowers, the policy delivers welfare losses driven almost entirely by higher tax obligations needed to finance the forgiveness, with negligible offsetting gains from general equilibrium spillovers. While forgiveness does provide a small welfare benefit through reduced consumption risk—acting as government-provided intergenerational risk sharing—this effect is quantitatively minor. The results suggest that student loan obligations were not the binding constraint on wealth accumulation for young highly leveraged borrowers, and the forgiveness program operates primarily as a fiscal transfer rather than a mechanism to unlock productive investment.

Real and Asset Pricing Effects of Employer Retirement Matching

This chapter asks whether employer retirement matching generates meaningful general equilibrium effects on firm investment and output. Employer matching subsidizes household equity purchases, altering savings incentives and potentially changing household intertemporal marginal rates of substitution (MRS). The firm discounts future dividends using the endogenous stochastic discount factor (SDF) arising from household MRS, so matching could affect corporate investment through this repricing channel. I integrate stochastic overlapping generations and neoclassical q-theory firm investment models, where matching enters households' Euler equations and the household MRS-implied SDF determines the firm's cost of capital. Analytically, I prove in a two-period deterministic model that matching unambiguously increases the SDF, reduces equilibrium returns, and raises capital investment regardless of if the matching is financed out of the labor or capital share of output: households tolerate lower market returns because their effective returns inclusive of the match remain attractive. Solving the full 60-period stochastic model using FINNs confirms these

predictions quantitatively. Introducing empirically realistic matching reduces equilibrium equity returns by 79 basis points, increases the aggregate capital stock by 6.1%, and raises wages by 1.7%.

Deep Learning the Term Structure for Derivatives Pricing

This chapter introduces a no-arbitrage, Monte Carlo-free approach to pricing path-dependent interest rate derivatives. The Heath-Jarrow-Morton model gives arbitrage-free contingent claims prices but is infinite-dimensional, making traditional numerical methods computationally prohibitive. To make the problem computationally tractable, I cast the stochastic pricing problem as a deterministic partial differential equation (PDE). Finance-Informed Neural Networks (FINNs) solve this PDE directly by minimizing violations of the differential equation and boundary condition, with automatic differentiation efficiently computing the exact derivatives needed to evaluate PDE terms. FINNs achieve pricing accuracy within 0.04 to 0.07 cents per dollar of contract value compared to Monte Carlo benchmarks. Once trained, FINNs price caplets in a few microseconds regardless of dimension, delivering speedups ranging from 300,000 to 4.5 million times faster than Monte Carlo simulation as the state space discretization of the forward curve grows from 10 to 150 nodes. The major Greeks—theta and curve deltas—come for free, computed automatically during PDE evaluation at zero marginal cost, whereas Monte Carlo requires complete re-simulation for each sensitivity. The framework generalizes naturally beyond caplets to other path-dependent derivatives—caps, swaptions, callable bonds—requiring only boundary condition modifications while retaining the same core PDE structure.