AlphaManager: A Data-Driven-Robust-Control Approach to Corporate Finance

Murillo Campello[1] **Lin William Cong**[2] **Luofeng Zhou**[3]

[1]UF Warrington [2]Cornell Johnson [3]NYU Stern

ABFER Corporate Finance, May 19, 2025

Corporate Finance Challenges and AI to the Rescue?

- Graham (2022, AFA Presi-Address): Corporate finance and reality
  - CF models limited ability for explaining/predicting outcomes
  - around 10% of R^2 in-sample, worse out-of-sample
  - call for models closer to the reality

Corporate Finance Challenges and AI to the Rescue?

- Graham (2022, AFA Presi-Address): Corporate finance and reality
  - CF models limited ability for explaining/predicting outcomes
  - around 10% of R^2 in-sample, worse out-of-sample
  - call for models closer to the reality
- Mitton (2022, RFS): Methodological Variation in Empirical CF
  - p-hacking and theory-fitting in empirical CF
  - call for unified definition and framework for universal analysis

Corporate Finance Challenges and AI to the Rescue?

- Graham (2022, AFA Presi-Address): Corporate finance and reality
  - CF models limited ability for explaining/predicting outcomes
  - around 10% of R^2 in-sample, worse out-of-sample
  - call for models closer to the reality
- Mitton (2022, RFS): Methodological Variation in Empirical CF
  - p-hacking and theory-fitting in empirical CF
  - call for unified definition and framework for universal analysis
- Spiegel (2023, Financial Review): For corporate finance to truly advance we need more genuinely testable models
  - CF models are often static
  - lack of interplay between firms and financial markets
  - call for more dynamic and testable models

Corporate Finance Challenges and AI to the Rescue?

- Graham (2022, AFA Presi-Address): Corporate finance and reality
- Mitton (2022, RFS): Methodological Variation in Empirical CF
- Spiegel (2023, Financial Review): For corporate finance to truly advance we need more genuinely testable models
- Challenge: many states and controls with endogeneous and nonlinear interactions

Corporate Finance Challenges and AI to the Rescue?

- Graham (2022, AFA Presi-Address): Corporate finance and reality
- Mitton (2022, RFS): Methodological Variation in Empirical CF
- Spiegel (2023, Financial Review): For corporate finance to truly advance we need more genuinely testable models
- Challenge: many states and controls with endogeneous and nonlinear interactions
- AI to the rescue?
  - big data for firms and financial markets
  - more flexible and efficient algorithms
  - more powerful computation
  - advancement of large models applied to finance

A Data-Driven-Robust-Control Approach to Corporate Finance

Understanding Underlying Economic Mechanisms ⟵ ⟶ Empirically Explaining Outcomes & Guiding Actions

Reduced Form	**Structural Estimation**	?
transparent easy to estimate locally causal	holistic theory-based counterfactual	data-driven unrestricted modeling counterfactual

- CF fundamentally a stochastic control problem

$$\max_{\{u_{t_0},\dots,u_{t_0+T}\}} \mathbb{E}_{t_0} \sum_{t=t_0}^{t_0+T} r(X_t, u_t) \qquad s.t.\ \Delta X_{t+1} = f(X_t, u_t) + \varepsilon_{t+1}$$

  - e.g., a manager as an economic agent trying to maximize shareholder's equity by making managerial decisions
  - X_t: state
  - u_t: control
  - f: mean law of motion function
  - r: reward function (instantaneous utility function)

A Data-Driven-Robust-Control Approach to Corporate Finance

Understanding Underlying Economic Mechanisms	**Reduced Form**	**Structural Estimation**	?	Empirically Explaining Outcomes & Guiding Actions
	transparent easy to estimate locally causal	holistic theory-based counterfactual	data-driven unrestricted modeling counterfactual	

- CF fundamentally a stochastic control problem

$$\max_{\{u_{t_0},\ldots,u_{t_0+T}\}} \mathbb{E}_{t_0} \sum_{t=t_0}^{t_0+T} r(X_t, u_t) \qquad s.t.\ \Delta X_{t+1} = f(X_t, u_t) + \varepsilon_{t+1}$$

- Reduced-form approaches: need not spell out explicitly
  - identify local causality $\Rightarrow$ counterfactuals
  - fragmented knowledge
  - internal validity

A Data-Driven-Robust-Control Approach to Corporate Finance

Understanding Underlying Economic Mechanisms	**Reduced Form**	**Structural Estimation**	?	Empirically Explaining Outcomes & Guiding Actions
	transparent easy to estimate locally causal	holistic theory-based counterfactual	data-driven unrestricted modeling counterfactual	

- CF fundamentally a stochastic control problem

$$\max_{\{u_{t_0},\dots,u_{t_0+T}\}} \mathbb{E}_{t_0} \sum_{t=t_0}^{t_0+T} r(X_t, u_t) \qquad s.t.\ \Delta X_{t+1} = f(X_t, u_t) + \varepsilon_{t+1}$$

- Reduced-form approaches: need not spell out explicitly
- Structural approaches: specify a simplified version
  - limited state variables of interest (for tractability)
  - dynamics of these variables exogenously given
  - micro-founded parameters within the framework
  - balance between internal and external validity

A Data-Driven-Robust-Control Approach to Corporate Finance

Understanding Underlying Economic Mechanisms	Reduced Form	Structural Estimation	DDRC	Empirically Explaining Outcomes & Guiding Actions
←	transparent easy to estimate locally causal	holistic theory-based counterfactual	data-driven unrestricted modeling counterfactual	→

- CF fundamentally a stochastic control problem

$$\max_{\{u_{t_0},\ldots,u_{t_0+T}\}} \mathbb{E}_{t_0} \sum_{t=t_0}^{t_0+T} r(X_t, u_t) \qquad s.t.\ \Delta X_{t+1} = f(X_t, u_t) + \varepsilon_{t+1}$$

- Reduced-form approaches: need not spell out explicitly
- Structural approaches: specify a simplified version
- Our approach: specify and solve the whole problem
    - more to the external validity
    - predictive environment module (PEM): supervised learning to estimate the law-of-motions of states and the model uncertainty
    - decision-making module (DMM): reinforcement learning (RL) for high-dimensional stochastic control approximation
    - supplement the internal validity concern using transfer learning

Literature and Contribution

- Corporate Finance:
  - New DDRC overcoming limitations and unifying framework
  - Machine learning in Corporate Finance
    (i) Textual analysis, e.g., Bellstam, Bhagat, & Cookson, 2021, Li et al., 2021, Hanley and Hoberg, 2019, Cong, Liang, & Zhang, 2019, etc.;
    (ii) Supervised learning, e.g., Erel et al., 2021, Lyonnet and Stern, 2022)
  - Non-text-based “large” model tailored for CF
- Robust Control:
  - Mostly theory, focus on macro time series rather than utilizing cross-sectional info (e.g., Hansen and Sargent, 2001; Ju and Miao, 2012)
  - Application in corporate finance
  - Use ambiguity to assess the importance of causality/theory
- Artificial Intelligence:
  - Goal-oriented search (Cong et al., 2020, 2022, 2023)
  - Model-based offline RL (empirical)
  - Incorporate theory/reduced-form/structural into DDRC (transfer learning)

Data and Variables

- Data: Compustat (firm fundamentals), CRSP (market return and volatility), and Chicago Fed (macro state variables)
- From 1976 to 2023, quarterly; 20,485 different firms ranging from 1976:Q1 to 2023:Q2, with 784,460 firm-quarter observations
- State variables (built from 10 fundamental + 4 market + 4 macro)
  - Total asset, current asset, gross revenue, accounts payable, cogs, interest paid net, inventories, book current liabilities, receivables, revenue
  - Market cap, enterprise value, quarterly equity return, quarterly volatility
  - Chicago Fed indices: risk, credit, leverage, and non-financial leverage
  - Plus their History (last 4 observations) and their growth rate version
- Decision variables (9 dimensions of actions in the current quarter)
  - Leverage, acquisitions, investment, cash savings, dividend, debt issuance, equity issuance, R&D expenses, repurchases
- Total over 3M parameters; trained using A100 GPU (RedCloud)/P100 (Azure)/T4 (RedCloud) with training time $\sim 3 - 7$ days per set

AlphaManager (AM) Architecture

Panel A: Predictive Environment Module (PEM)

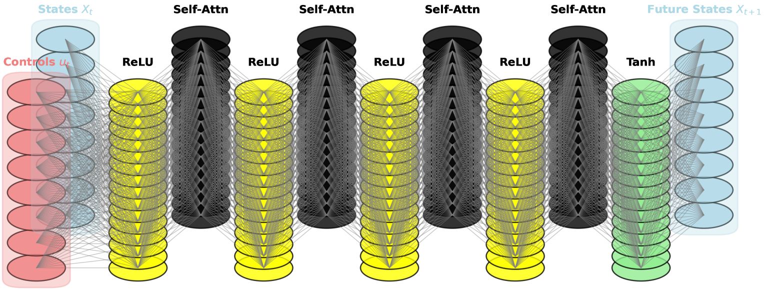

- Predictive Environment module

$$\max_{\{u_{t_0},\ldots,u_{t_0+T}\}} \mathbb{E}_{t_0} \sum_{t=t_0}^{t_0+T} r(X_t, u_t) \qquad s.t.\ \Delta X_{t+1} = f(X_t, u_t) + \varepsilon_{t+1}$$

AlphaManager (AM) Architecture

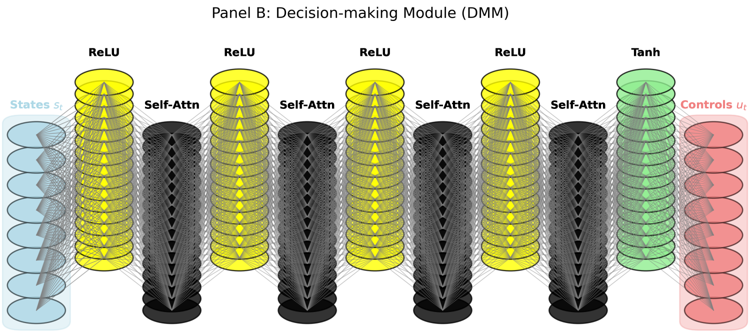

- AlphaManager Policy Module

$$\max_{\color{red}g(\cdot)} \ \mathbb{E}_{t_0} \sum_{t=t_0}^{t_0+T} r(X_t, {\color{red}g(X_t)}) \qquad s.t.\ \Delta X_{t+1} = f(X_t, {\color{red}g(X_t)}) + \varepsilon_{t+1}$$

Robust Control and Ambiguity

- Imperfection of a model: overfitting, data shifts $\Rightarrow$ model uncertainty
- Three sources of uncertainty (Hansen and Sargent, 2024):
  - Risk - in-model stochastic innovation
  - Misspecification - limited power of the model class
  - Ambiguity - uncertainty about model choice
- Inspiration from climate finance (Barnett, Brock, and Hansen, 2020): max-min + relative-entropy punishment with probability adjustment
- A bag of PEMs, indexed by $i = 1, 2, ..., I$, and ambiguity aversion
- Ambiguity aversion: maximize the minimum of reward (max-min)

$$r^i\left(X_t, g(X_t)\right) \Rightarrow \min_{i=1,2,\ldots,I} r^i\left(X_t, g(X_t)\right)$$

- Boosting error: the greatest dispersion among model predictions

$$\text{BoostingError}(X_t, u_t) = \frac{1}{D}\sum_{d=1}^{D}\left(\max_{i=1,2,\ldots,I} \hat{X}^i_{t+1,d} - \min_{i=1,2,\ldots,I} \hat{X}^i_{t+1,d}\right)^2$$

- Threshold punishment as a function of BoostingError

Empirical Results: Prediction and Ambiguity for Macroeconomic States

- Individual firms hard to influence the aggregate economy
- We train 20k NNs with past macro states as only inputs
- Prediction spread as a proxy for ambiguity
- VAR as benchmark

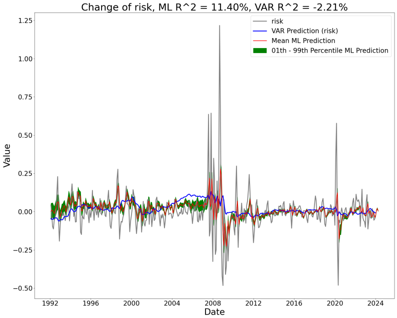

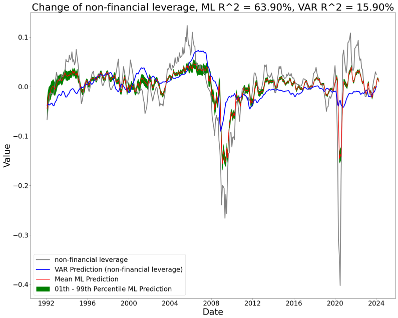

Empirical Results: PEM's Predictions of Firm Outcomes

- High-dimensional, high-fidelity OOS, reduce costly experiments.

	Ignoring Control		With Control	
State Variable	Training R^2	Test R^2	Training R^2	Test R^2
Book Asset Growth	–4.09%	–8.15%	55.44%	62.56%
Current Asset Growth	–3.58%	–7.10%	44.49%	51.21%
Gross Revenue Growth	29.54%	28.68%	31.33%	30.88%
Accounts Payable Growth	21.46%	24.43%	24.40%	27.64%
COGS Growth	25.68%	26.76%	27.00%	28.56%
Net Interest Paid Growth	73.26%	77.17%	73.36%	77.28%
Inventory Growth	12.78%	13.71%	17.04%	18.92%
Current Liability Growth	8.88%	7.72%	21.89%	22.69%
Receivables Growth	17.52%	18.77%	21.59%	23.20%
Net Income Growth	29.51%	28.59%	31.31%	30.80%
Trading Volume Growth	12.81%	16.53%	15.77%	20.75%
Log Gross Return Growth	47.90%	45.27%	50.04%	48.19%
Market Cap Growth	1.32%	–3.33%	9.32%	7.07%
Enterprise Value Growth	–0.97%	–5.73%	14.61%	13.14%

- Controls more important for some state evolution
- Consistent with known local (causal) patterns from the literature

PEM Application: Recapitalization Analysis

How does enterprise value change if a firm:

1. raises $1 more debt and put that $1 into its cash savings
2. raises $1 more equity and put that $1 into its cash savings
3. raises $1 more debt and $1 less equity
4. puts $1 cash into investment
5. raises $1 more debt and put that $1 into investment
6. raises $1 more equity and put that $1 into investment

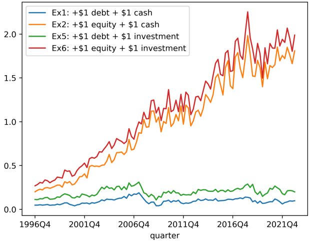

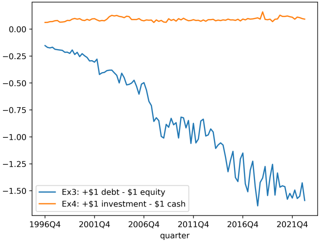

Heterogeneous PEM Performance (MSE) for System States

variable		full sample		pre-dotcom		dotcom-GFC		post-GFC	
		mean	std	mean	std	mean	std	mean	std
Total Assets	high	2.16%	8.53%	2.57%	9.42%	2.23%	8.48%	1.76%	7.68%
	low	4.40%	13.33%	4.84%	13.77%	4.30%	12.80%	4.08%	13.26%
COGS	high	2.84%	11.01%	3.30%	12.32%	2.74%	10.64%	2.50%	9.95%
	low	4.49%	13.75%	4.94%	14.67%	4.32%	13.33%	4.20%	13.15%
CurrentLiability	high	5.48%	13.98%	6.12%	14.75%	5.59%	14.20%	4.84%	13.09%
	low	6.69%	15.66%	6.83%	15.42%	6.58%	15.47%	6.64%	15.98%
MarketCap	high	7.49%	15.51%	9.77%	18.45%	7.79%	15.74%	5.30%	11.83%
	low	12.42%	21.40%	12.88%	21.69%	13.42%	22.53%	11.40%	20.34%
EnterpriseValue	high	6.11%	13.87%	8.60%	17.90%	5.99%	12.75%	4.02%	9.47%
	low	10.37%	19.01%	11.55%	20.73%	11.34%	19.67%	8.73%	16.75%
MacroRisk	high	4.81%	5.51%	6.90%	7.23%	6.24%	6.75%	5.23%	6.26%
	low	5.74%	6.57%	5.25%	5.93%	5.11%	5.92%	5.43%	6.62%

- Subsample episodes: pre-dot com, dot com to GFC, post-GFC
- Book asset: small firms has higher prediction error and std, pre-dotcom has the highest mean and std
- COGS: both higher and lower halves have declining average MSE
- Market cap and enterprise value: lower half has higher MSE
- Macro 1 (risk): lower half is doing better in general, but not for the post-GFC period

PEM: Heterogeneous Ambiguity for System States

variable		full sample		pre-dotcom		dotcom-GFC		post-GFC	
		mean	std	mean	std	mean	std	mean	std
Total Assets	high	6.93%	3.72%	7.18%	3.74%	7.02%	3.65%	6.64%	3.72%
	low	6.82%	3.79%	6.90%	3.84%	6.84%	3.69%	6.72%	3.82%
COGS	high	5.50%	2.70%	5.46%	2.40%	5.68%	2.68%	5.42%	2.93%
	low	5.21%	2.67%	4.97%	2.33%	5.42%	2.70%	5.30%	2.89%
CurrentLiability	high	7.30%	3.67%	7.21%	3.52%	7.44%	3.66%	7.29%	3.79%
	low	6.84%	3.55%	6.57%	3.49%	6.91%	3.46%	7.02%	3.63%
MarketCap	high	4.68%	2.78%	4.51%	2.40%	5.54%	3.59%	4.28%	2.35%
	low	5.27%	3.20%	4.49%	2.31%	6.03%	3.83%	5.47%	3.29%
EnterpriseValue	high	5.85%	2.67%	5.70%	2.50%	6.57%	2.98%	5.54%	2.53%
	low	6.11%	2.91%	5.45%	2.52%	6.78%	3.15%	6.27%	2.94%
MacroRisk	high	5.63%	1.82%	5.98%	1.91%	6.33%	1.86%	5.93%	1.92%
	low	5.91%	1.94%	5.42%	1.84%	5.87%	1.87%	5.94%	2.04%

- Book asset: lower half has lower ambiguity, pre-dot com episode has the lowest mean and std
- COGS: both higher and lower halves have increasing average ambiguity
- Market cap, and enterprise value: highest ambiguity during dot com to GFC
- Macro1 (risk): lower half is doing better in general, but not for the post-GFC period

Out-Performance of AlphaManager (DMM)

- Objectives: next Q and next 8Q market cap and enterprise value.
- Next Q market cap increase (short-termist)
- Overall short-horizon outperformance: **8.44%** and 10.19%.
- Long-horizon objective: 8.73% and 4.43%

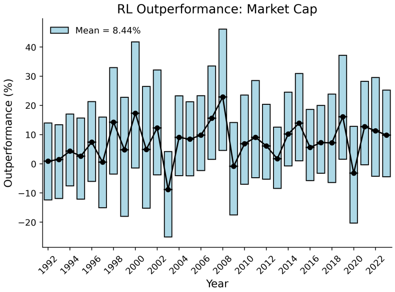

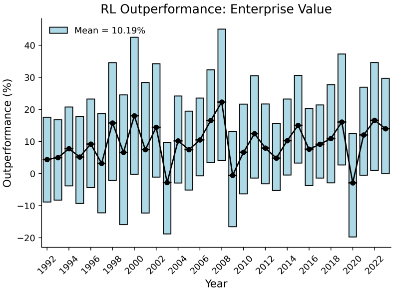

Long-term DMM Performance Under Short-Termism

- Objective: 1QTR market cap (left) or enterprise value (right) growth
- Evaluation period: 8QTRs (blue bars)
- Benchmark: firm performance in the real data (red bars)

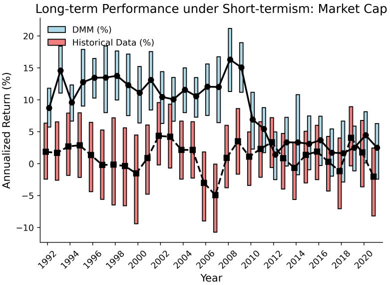

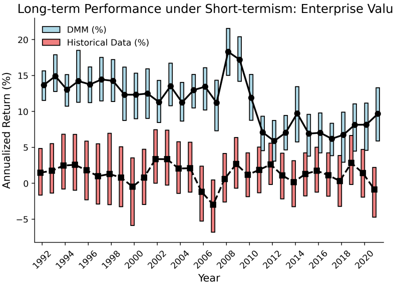

Term Structure of AM Performance Under Ambiguity

- Is long-term performance for short-term RL possibly better than that for long-term RL?
  - Intuition: LT RL is optimal under LT objective

Term Structure of AM Performance Under Ambiguity

- Is long-term performance for short-term RL possibly better than that for long-term RL?

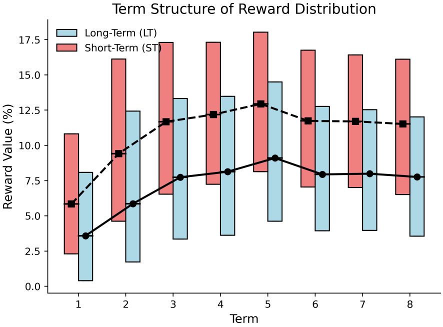

- ST RL does a better job along the term structure. Why?

Term Structure of AM Performance Under Ambiguity

- Is long-term performance for short-term RL possibly better than that for long-term RL?
- ST RL does a better job along the term structure. Why?
- ST and LT RLs have different ambiguity constraints
  - ST RL doesn't care about LT ambiguity $\Rightarrow$ looser constraint
  - Constrained instead of unconstrained optimization

Term Structure of AM Performance Under Ambiguity

- Is long-term performance for short-term RL possibly better than that for long-term RL?

- ST RL does a better job along the term structure. Why?
- ST and LT RLs have different ambiguity constraints
- Use ambiguity-adjusted reward to test

$$\text{Ambiguity Adj. Reward}_t = \frac{\text{Reward}_t}{\sqrt{\max\{1, \frac{\text{BoostingError}_t}{\text{BoostingError}_{t_0}}\}}}$$

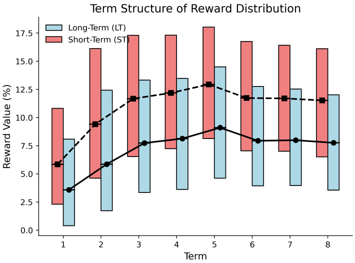

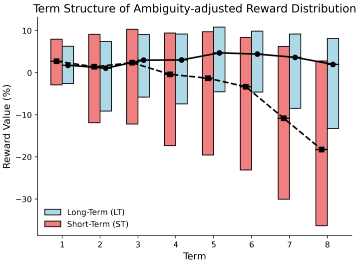

Heterogeneous Performance of AlphaManager by Value Decile

	full sample		pre-dotcom		dotcom-GFC		post-GFC	
book-to-market decile	mean	std	mean	std	mean	std	mean	std
1	3.98%	19.72%	4.22%	15.90%	-3.89%	23.80%	9.99%	17.76%
2	5.37%	17.42%	5.63%	13.74%	-1.00%	20.94%	10.27%	16.36%
3	5.87%	16.05%	6.28%	11.79%	-0.26%	19.94%	10.42%	15.07%
4	6.53%	15.47%	7.28%	11.31%	0.43%	18.94%	10.67%	14.83%
5	7.17%	15.04%	7.81%	10.90%	1.73%	18.68%	10.90%	14.42%
6	7.72%	14.69%	8.40%	10.81%	2.54%	18.10%	11.32%	14.05%
7	8.10%	14.91%	8.67%	10.72%	3.46%	18.16%	11.48%	14.84%
8	8.15%	15.10%	8.56%	10.63%	3.36%	18.55%	11.50%	15.24%
9	8.05%	15.54%	9.06%	10.18%	2.88%	19.60%	10.67%	15.73%
10	7.88%	15.38%	8.88%	10.72%	2.20%	17.65%	10.51%	16.63%

- Objective: enterprise value growth in the next 2 years
- AlphaManager performance mainly driven by value firms

Optimal Actions Versus Historical Actions

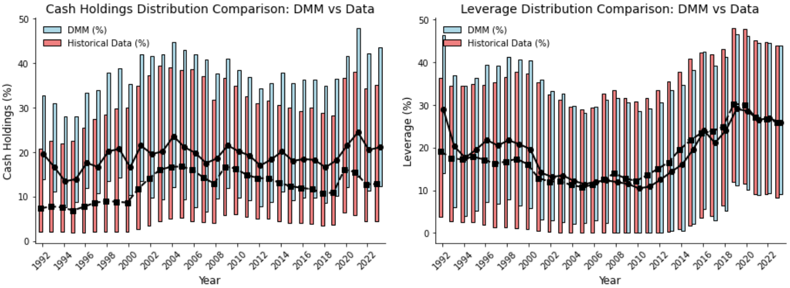

Figure: Optimal decisions (blue lines) vs real decisions (red lines): cash holdings (left) and leverage (right)

Maximizing **next Q enterprise value**: more acquisitions, increasing cash holdings more, keeping the same leverage, paying out more dividend, and increasing investment, especially in R&D, allowing more variations in investments, and more repurchases during bad times.

Three Perspectives on Out-performance of AlphaManager

- Explanation #1: managers are not skilled enough to realize their goals
  - Intuition: managerial decisions are not aligned with their preferences
  - Example: bad execution or limited information
  - Out-performance $\Leftrightarrow$ how irrational the manager is
- Explanation #2: the objective is mis-specified
  - Intuition: if managers are rational and the preference is correctly specified, then the expected out-performance of AM should be 0
  - Example: ESG, lobbying threat, personal achievement
  - Out-performance $\Leftrightarrow$ how mis-specified the preference is
- Explanataion #3: firms face unobservable constaints
  - Intuition: managers wanted to but are not able to do so
  - Example: financial constraints in borrowing, lack of investment opportunities
  - Out-performance $\Leftrightarrow$ shadow prices of binded constraints

Revealed Managerial Preferences

- Exogeneously-specified objectives are too ideal
- Real managers may care about a linear combination of them
- A projection exercise:

$$\min_{\beta_k} \left(\sum_{k=1}^{4} \beta_k \cdot u_{j,t,k} - u^*_{j,t} \right)^2 \quad \text{s.t.} \sum_{k=1}^{4} \beta_k = 1$$

- R^2s are around 10% $\Rightarrow$ the rest 90% variation because of either nonlinearity or mis-specification

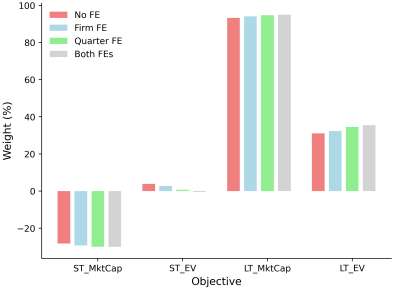

Further Discussion and Takeaways

Piecing Together Corporate Finance Research

- Ambiguity and the need for theory/reduced-form/structural models
  - Boundary of data-driven approach
- Ambiguity-guided transfer learning
  - Combining insights and predictions from other approaches to improve internal validity

Further Discussion and Takeaways

Piecing Together Corporate Finance Research

- Ambiguity and the need for theory/reduced-form/structural models
  - Boundary of data-driven approach
- Ambiguity-guided transfer learning
  - Combining insights and predictions from other approaches to improve internal validity

Takeaways:

- DL and robust control for building "world model" of corporate finance
- Deep reinforcement learning as heuristic search for optimizing arbitrary managerial goals/objectives
- A data-driven-robust-control approach to corporate finance