AlphaManager:

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

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

[1]University of Florida & NBER
[2]Nanyang Technological University, ABFER, CEPR, & NBER
[3]New York University

AFA 2026

Corporate Finance Research: Challenges and New Techniques

- Graham (2022, AFA President Address): CF research & reality
  - CF models limited ability for explaining/predicting outcomes
  - around 10% of R^2 in-sample, worse out-of-sample
  - calls for models closer to the reality and new approaches (e.g., surveys)

Corporate Finance Research: Challenges and New Techniques

- Graham (2022, AFA President Address): CF research & reality
  - CF models limited ability for explaining/predicting outcomes
  - around 10% of R^2 in-sample, worse out-of-sample
  - calls for models closer to the reality and new approaches (e.g., surveys)
- 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 Research: Challenges and New Techniques

- Graham (2022, AFA President Address): CF research & reality
  - CF models limited ability for explaining/predicting outcomes
  - around 10% of R^2 in-sample, worse out-of-sample
  - calls for models closer to the reality and new approaches (e.g., surveys)
- 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 Research: Challenges and New Techniques

- Graham (2022, AFA President Address): CF research & 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
- Challenges:
  - many states and controls with endogeneous and nonlinear interactions
  - fragmented and misspecified models

Corporate Finance Research: Challenges and New Techniques

- Graham (2022, AFA President Address): CF research & 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
- Challenges:
  - many states and controls with endogeneous and nonlinear interactions
  - fragmented and misspecified models
- 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

- Firm decisions fundamentally a stochastic control problem

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

  - 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 ⟷ Empirically Explaining Outcomes & Guiding Actions

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

- Firm decisions 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: expertise-driven and ad hoc
  - identify local causality $\Rightarrow$ counterfactuals
  - High internal validity
  - fragmented knowledge

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	

- Firm decisions 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: expertise-driven and ad hoc
- Structural approaches: search within theoretically tractable space
  - 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 ⟷ Empirically Explaining Outcomes & Guiding Actions

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

- Firm decisions 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: expertise-driven and ad hoc
- Structural approaches: search within theoretically tractable space
- Our approach: model the whole system and simulate the environment
  - predictive environment module: supervised learning to estimate the law-of-motions of states and the model uncertainty
  - decision-making module: reinforcement learning for high-dimensional stochastic control approximation
  - enhance internal validity using transfer learning and external validity via robust control

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
- AI, especially GenAI, for Finance:
  - Goal-oriented algorithms in large spaces (Cong et al., 2020, 2022, 2023)
  - Model-based offline reinforcement learning (RL, empirical)
  - Simulation-based studies: ABM and large environment models
  - 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 Architecture: Baseline Modules

Predictive Environment Module (PEM), 11 aux, DL, 3 x 300 (40 epochs):

- $\sim$30% training. 1976Q1 - 1991Q4; Transformer-based; quarterly rolling.

Decision-Making Module (DMM), RL/policy-gradient, 4 x 256 (200 epochs):

$$\max_{g} \ \mathbb{E}_{t_0} \sum_{t=t_0}^{t_0+T} \{r(X_t, u_t) - \lambda \cdot \mathit{BoostingError}(X_t, u_t)\}$$

$$\text{s.t. } \Delta X_{t+1} = f(X_t, u_t) + \varepsilon_{t+1}, \ \ u_t = g(X_t)$$

Robust Control and Ambiguity (Hansen and Sargent, 2023):

- Overfit and data shifts.
- Misspecification - limited power of the model class
- Risk - in-model stochastic innovation
- Ambiguity - uncertainty about detailed modeling choice
- Inspiration from climate finance (Barnett, Brock, and Hansen, 2020): max-min $+$ relative-entropy punishment $+$ probability adjustment.
- Pick the “worst” environment and punish/avoid ambiguity

Validation: Holmstrom-Tirole (1997 QJE) Moral Hazard Model

- An ideal setting for validation:
  (i) analytical solutions for data-generation and optimal policy,
  (ii) simple and convincing channels/mechanisms demonstration, and
  (iii) sufficiently flexible to generate distributional shifts:
  - An entrepreneur has an investment opportunity, can exert unobservable effort to increase success rate, but cannot afford required initial investment
  - The entrepreneur writes a contract to share a proportion of return (if successful) to external investor in exchange of initial investment
  - An external investor chooses to take or reject proposed contract
  - Moral hazard: the investor cannot observe the entrepreneur's effort
- Formulation in stochastic control setting
  - States: entrepreneur's initial wealth (A), entrepreneur's utility of not exerting effort (B), change in probability of success Δp, the successful rate when no effort is exerted (p_L), total return of the investment opportunity (R)
  - Controls: contractual return to the entrepreneur R_b, effort $e \in \{0, 1\}$
  - Environment: given states and controls, whether the external investor takes or rejects the proposed contract $q \in [0, 1]$

Validation: Holmstrom-Tirole (1997 QJE) Moral Hazard Model

- An ideal setting for validation:
  (i) analytical solutions for data-generation and optimal policy,
  (ii) simple and convincing channels/mechanisms demonstration, and
  (iii) sufficiently flexible to generate distributional shifts:
- Formulation in stochastic control setting
  - States: entrepreneur's initial wealth (A), entrepreneur's utility of not exerting effort (B), change in probability of success Δp, the successful rate when no effort is exerted (p_L), total return of the investment opportunity (R)
  - Controls: contractual return to the entrepreneur R_b, effort $e \in \{0, 1\}$
  - Environment: given states and controls, whether the external investor takes or rejects the proposed contract $q \in [0, 1]$

$$\begin{aligned}\max_{R_b,\, e\in\{0,1\}} \quad & q \cdot \max\{A, (p_e \cdot R_b + \mathbf{1}_{e=0} \cdot B)\} + (1-q)\cdot A \\ \text{s.t.} \quad & R_b + R_l = R \\ & q = q(\Delta p, A, R_b; B, p_L) \\ & R_b \geq 0,\ R_l \geq 0,\ A < 1 < R \\ & p_0 = p_L, \quad p_1 = p_L + \Delta p\end{aligned}$$

Learning the Environment & DMM Performance in Test Sample

- Modeling the environment and the role of robust control
  1. Entrepreneur has perfect info on $q()$, standard Q-learning suffices, RL is perfect.
  2. With unknown $q()$, entrepreneur learns from data, distributional shifts makes environment model less accurate.
  3. Robust control overcome label shift in states and covariate shift in states and actions.

- In Holmstrom-Tirole, state distributional shift leads to higher ambiguity

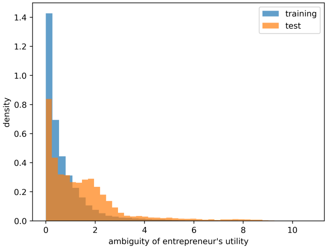

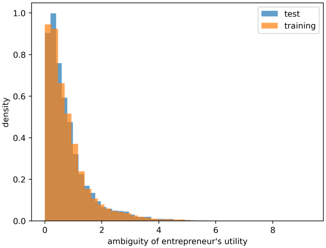

- Under ambiguity aversion, the correlation between empirical DMM controls and analytical solutions is consistently over 93%.

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/managerial actions more important for some state evolution
- Consistent with known local patterns from the literature

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%

- Subsample episodes: pre-dotcom, dotcom 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

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%

- Book asset and current liability: lower half has lower ambiguity, pre-dotcom episode has the lowest mean and std
- COGS: both higher and lower halves have increasing average ambiguity
- Market cap, and enterprise value: highest ambiguity in the middle period

PEM Power Illustrated in Traditional/Low-Dimensional settings

Recapitalization: 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?

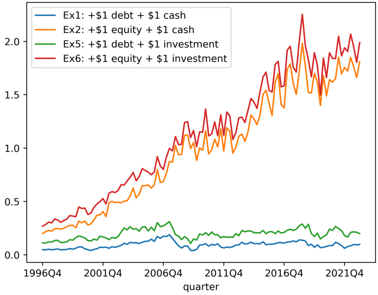

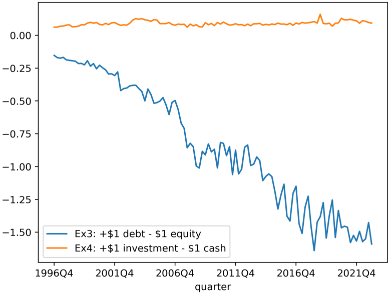

DMM & Out-Performance of AlphaManager

- 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%
- Heterogeneity: mainly driven by value firms

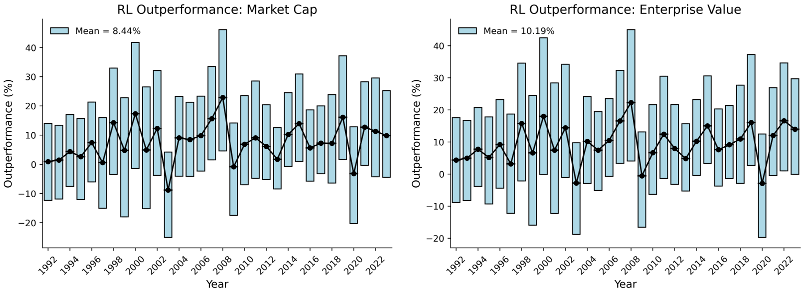

Short-Termism VS. Long-Term DMM

- Short-termism doesn't hurt LT performance much, at least not post financial crisis.
  - Counter-intuitive: LT RL should be optimal in LT

Short-Termism VS. Long-Term DMM

- Short-termism doesn't hurt LT performance much, at least not post financial crisis.

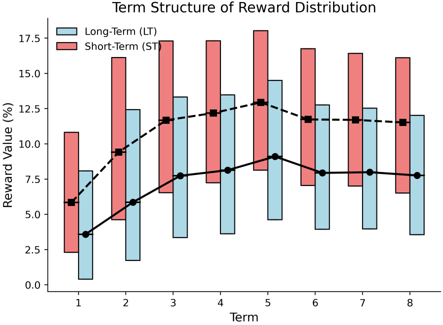

- ST RL does a better job along the term structure. Why?

Short-Termism VS. Long-Term DMM

- Short-termism doesn't hurt LT performance much, at least not post financial crisis.
- 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

Short-Termism VS. Long-Term DMM

- Short-termism doesn't hurt LT performance much, at least not post financial crisis.
- 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

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

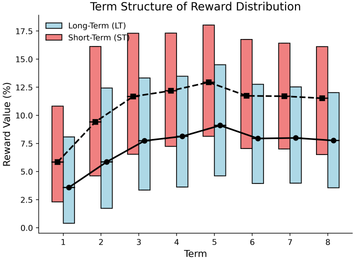

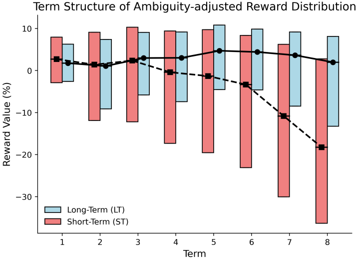

Optimal Actions Versus Historical Actions

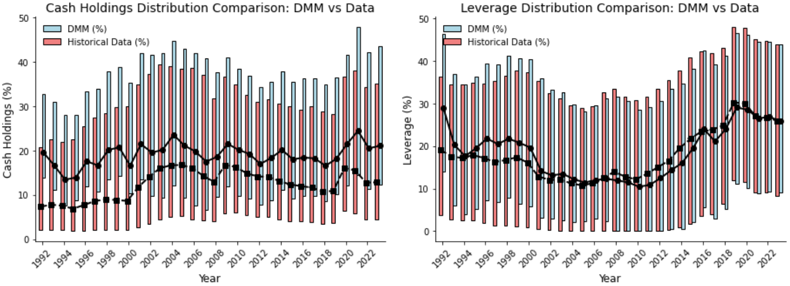

Figure: Optimal decisions (blue lines) vs real decisions (red lines): cash holdings (left) and leverage (right)

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 misspecified
  - Intuition: if managers were to be rational and the preference were correctly specified, then the expected out-performance of AM would be 0
  - Example: ESG, lobbying threat, personal achievement
  - Out-performance $\Leftrightarrow$ how misspecified 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

Understanding Managerial Objectives

- Exogeneously-specified objectives may not be accurate
- Simple projection: $R^2 \sim 10\%$; the rest nonlinearity or mis-specification

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

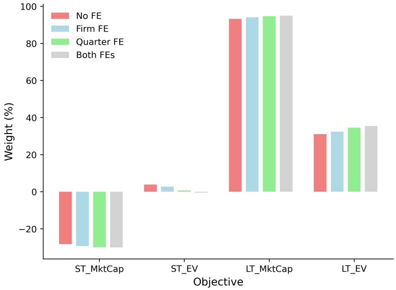

Understanding Managerial Objectives

- Exogeneously-specified objectives may not be accurate
- Simple projection: $R^2 \sim 10\%$; the rest nonlinearity or mis-specification

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

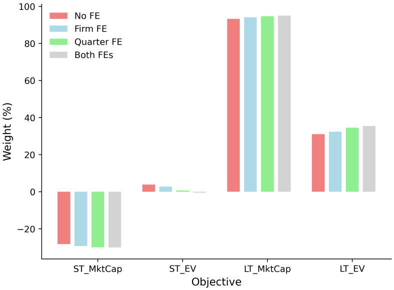

- First data-drive investigation of managerial objectives:
  - Actions not additive: no holistic study learning revealed preference,
  - A high-dim min-max problem $\rightarrow$ GANs to the rescue (companion paper).

Further Discussion and Takeaways

Informing and Piecing Together Corporate Finance Research

- Ambiguity tells the boundary of data-driven/computational approach.
  - robust control avoids those situations, but can we do better?
- Transfer learning (e.g., Cao et al., 2025; Chen et al., 2024), but ambiguity-guided
  - combining insights and predictions from other approaches and models

    $\Rightarrow$ lower ambiguity, improved internal validity

Further Discussion and Takeaways

Informing and Piecing Together Corporate Finance Research

- Ambiguity tells the boundary of data-driven/computational approach.
  - robust control avoids those situations, but can we do better?
- Transfer learning (e.g., Cao et al., 2025; Chen et al., 2024), but ambiguity-guided
  - combining insights and predictions from other approaches and models
    $\Rightarrow$ lower ambiguity, improved internal validity

Takeaways:

- DL and robust control for building "large world model" of corporate finance, the first non-text-based GenAI CF model.
- DRL as heuristic search for optimizing arbitrary managerial goals/objectives (Cong et al. 2019, 2022); AI tool for managers.
- A data-driven-robust-control approach to corporate finance: inference & predictions, understanding managerial objectives, combining approaches, defining boundaries/guiding focus, and more.