EFMA Doctoral Session on AI and Finance

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

Murillo Campello Lin William Cong Luofeng Zhou

June 26, 2024

ML/AI for in Finance

- Financial big data: high-dim, non-linearity, interactions, low signal-to-noise, non-stationarity/heteroskedasticity, multi-sequence panel.
- Economic motivation and guidance; interpretability (Karolyi & Van Nieuwerburgh, 2020)?
- Asset pricing and investments
  - Off-the-shelf ML/AI emphasize prediction over inference/testing (Gu, Kelly, & Xiu; Karolyi & Van Nieuwerburgh, 2020).
  - Extant literature on supervised learning for recovering pricing kernel.
- Corporate finance
  - Textual: Hoberg & Phillips (2016); Hoberg & Hanley (2019); Li et al (2020)...
  - Supervised learning: Erel et al. (2021); Lyonnet & Stern (2022).
- Learning/training through examples (supervised learning) or relying only on statistical properties of data (e.g., unsupervised learning).

ML/AI for in Finance

- Financial big data: high-dim, non-linearity, interactions, low signal-to-noise, non-stationarity/heteroskedasticity, multi-sequence panel.
- Economic motivation and guidance; interpretability (Karolyi & Van Nieuwerburgh, 2020)?
- Asset pricing and investments
  - Off-the-shelf ML/AI emphasize prediction over inference/testing (Gu, Kelly, & Xiu; Karolyi & Van Nieuwerburgh, 2020).
  - Extant literature on supervised learning for recovering pricing kernel.
- Corporate finance
  - Textual: Hoberg & Phillips (2016); Hoberg & Hanley (2019); Li et al (2020)...
  - Supervised learning: Erel et al. (2021); Lyonnet & Stern (2022).
- Learning/training through examples (supervised learning) or relying only on statistical properties of data (e.g., unsupervised learning).

ML/AI for in Finance

- Financial big data: high-dim, non-linearity, interactions, low signal-to-noise, non-stationarity/heteroskedasticity, multi-sequence panel.
- Economic motivation and guidance; interpretability (Karolyi & Van Nieuwerburgh, 2020)?
- Asset pricing and investments
  - Off-the-shelf ML/AI emphasize prediction over inference/testing (Gu, Kelly, & Xiu; Karolyi & Van Nieuwerburgh, 2020).
  - Extant literature on supervised learning for recovering pricing kernel.
- Corporate finance
  - Textual: Hoberg & Phillips (2016); Hoberg & Hanley (2019); Li et al (2020)...
  - Supervised learning: Erel et al. (2021); Lyonnet & Stern (2022).
- Learning/training through examples (supervised learning) or relying only on statistical properties of data (e.g., unsupervised learning).

ML/AI for in Finance

- Financial big data: high-dim, non-linearity, interactions, low signal-to-noise, non-stationarity/heteroskedasticity, multi-sequence panel.
- Economic motivation and guidance; interpretability (Karolyi & Van Nieuwerburgh, 2020)?
- Asset pricing and investments
  - Off-the-shelf ML/AI emphasize prediction over inference/testing (Gu, Kelly, & Xiu; Karolyi & Van Nieuwerburgh, 2020).
  - Extant literature on supervised learning for recovering pricing kernel.
- Corporate finance
  - Textual: Hoberg & Phillips (2016); Hoberg & Hanley (2019); Li et al (2020)...
  - Supervised learning: Erel et al. (2021); Lyonnet & Stern (2022).
- Learning/training through examples (supervised learning) or relying only on statistical properties of data (e.g., unsupervised learning).

ML/AI for in Finance

- Financial big data: high-dim, non-linearity, interactions, low signal-to-noise, non-stationarity/heteroskedasticity, multi-sequence panel.
- Economic motivation and guidance; interpretability (Karolyi & Van Nieuwerburgh, 2020)?
- Asset pricing and investments
  - Off-the-shelf ML/AI emphasize prediction over inference/testing (Gu, Kelly, & Xiu; Karolyi & Van Nieuwerburgh, 2020).
  - Extant literature on supervised learning for recovering pricing kernel.
- Corporate finance
  - Textual: Hoberg & Phillips (2016); Hoberg & Hanley (2019); Li et al (2020)...
  - Supervised learning: Erel et al. (2021); Lyonnet & Stern (2022).
- Learning/training through examples (supervised learning) or relying only on statistical properties of data (e.g., unsupervised learning).

Modern AI (Beyond Basic ML) as Goal-Oriented Search

- Intelligence $\rightarrow$ associated with humans; instruction-driven automations are not intelligent.
- AI: the ability to learn and perform suitable techniques to solve problems and achieve goals, appropriate to the context in an uncertain, ever-varying world.
- Instruction-driven $\rightarrow$ goal-oriented in complex/uncertain world.
- The way of AI:
  1. Virtue of complexity, larger models, greater computation, general intelligence.
  2. More data and more computation; pre-training, self-learning.
  3. Clever goal-oriented search and incorporation of knowledge, logic, and expertise (our focus).
- Agenda: economically guided **heuristic search** and **greedy search** in a large decision/action space to achieve goals.
  - Portfolio management with arbitrage objectives and unlabeled data.
  - Generalized security sorting to generate test assets and kernels.
  - Data-driven, high-dimensional, dynamic managerial decision-making.

(Deep) Reinforcement Learning as Efficient Heuristic Search

People learn by **interacting with the environment** in an active and sequential way, to optimize some **rewards**.

The Reward Hypothesis: *Any goal can be formalized as the outcome of maximizing a cumulative reward, when the transition and reward probabilities are initially unknown..*

1. Fly a helicopter (Reward: air time, inverse distance, . . .).
2. Make a robot walk (Reward: distance, speed, . . .)
3. Play games (football, go, etc.; Reward: win, maximize scores, . . .)
4. Manage portfolio (Reward: returns, Sharpe ratio, min var, . . .)

- Reward, Value, Policy (Actions).
- Agents: Value-based, Policy-based, Actor Critic, etc.
- Approximate dynamic programming/stochastic control; multi-arm bandit.

DRL and Attention as Core of Recent AI Development

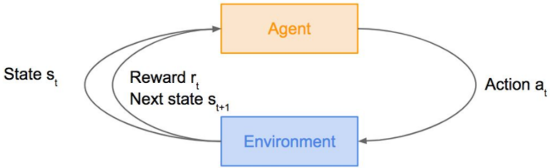

DRL and Attention as Core of Recent AI Development

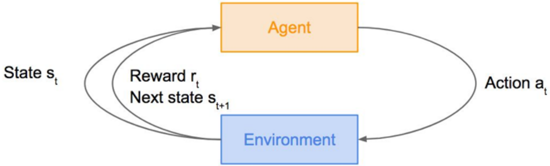

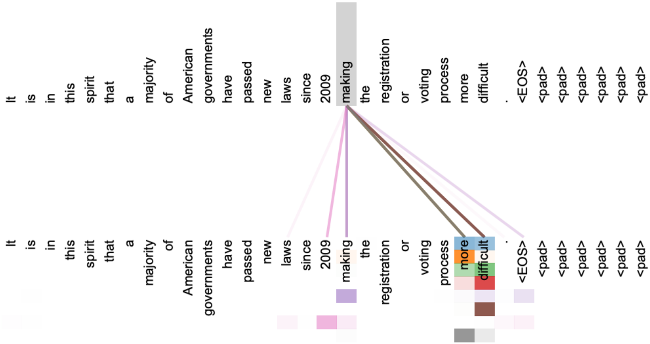

Goal-Oriented Portfolio Management Through Transformer-Based RL

- Cong, Tang, & Wang (2019, earlist Transformer/offline DRL in Finance).
  - First "large" model in finance ($>$ 10 million parameters).
  - Virtual of complexity/large models (Belkin et al., 2019, Kelly, Malamud, & Zhou, 2022, Fan et al., 2022).
  - Share the spirit of Grounded Theory ($>$ 20, 000 citations).

- Drawbacks in the conventional paradigms (e.g., Markowitz, 1952):
  - Two steps; first return distributions/premia/pricing errors.
  1. Estimation errors in the first step; traditional fixes/econometrics insufficient.
  2. Misalignment between asset pricing and general portfolio objectives.
  3. No action-state interactions (price impact, fund survival, etc.)

Goal-Oriented Portfolio Management Through Transformer-Based RL

- Cong, Tang, & Wang (2019, earlist Transformer/offline DRL in Finance).
  - First "large" model in finance ($>$ 10 million parameters).
  - Virtual of complexity/large models (Belkin et al., 2019, Kelly, Malamud, & Zhou, 2022, Fan et al., 2022).
  - Share the spirit of Grounded Theory ($> 20,000$ citations).
- Drawbacks in the conventional paradigms (e.g., Markowitz, 1952):
  - Two steps; first return distributions/premia/pricing errors.

  1. Estimation errors in the first step; traditional fixes/econometrics insufficient.
  2. Misalignment between asset pricing and general portfolio objectives.
  3. No action-state interactions (price impact, fund survival, etc.)
- Direct optimization of portfolio weights?
  - Non-Deep RL approach: Brandt (1999), Brandt & Santa-Clara (2006), etc.
  - Other supervised-ML approach? Means to an end.
- Goal-Oriented Search
  - Feasible data-driven direct construction.
  - Powerful embedding/info extraction using self-attention.
  - General portfolio management (constraints, interactions, etc.).

AlphaPortfolio Innovation: Transformer + Cross-Asset Attention Network + RL

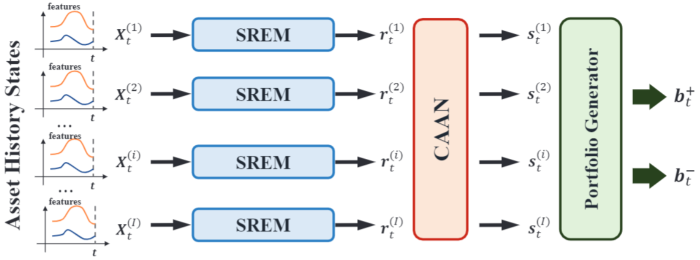

AlphaPortfolio Innovation: Transformer + Cross-Asset Attention Network + RL

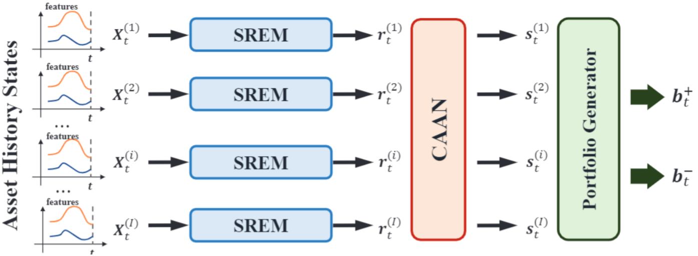

	AP Performance				AP Excess Alpha					
	(1)	(2)	(3)		(4)	(5)	(6)	(7)	(8)	(9)
Firms	All	$> q_{10}$	$> q_{20}$	Factor Models	All $\alpha(\%)$	R^2	$> q_{10}$ $\alpha(\%)$	R^2	$> q_{20}$ $\alpha(\%)$	R^2
Return (%)	17.00	17.10	18.10	CAPM	13.9***	0.005	12.2***	0.088	14.0***	0.102
Std.Dev. (%)	8.50	7.70	8.20	FFC	14.2***	0.052	13.4***	0.381	14.7***	0.465
Sharpe	2.00	2.31	2.21	FFC+PS	13.7***	0.054	12.3***	0.392	13.3***	0.480
Skewness	1.42	1.74	1.91	FF5	15.3***	0.12	13.8***	0.426	14.7***	0.435
Kurtosis	6.33	5.70	5.97	FF6	15.6***	0.128	14.5***	0.459	15.8***	0.516
Turnover	0.26	0.24	0.26	SY	17.4***	0.037	15.8***	0.332	17.0***	0.394
MDD	0.08	0.02	0.02	Q4	16.0***	0.121	15.0***	0.495	16.2***	0.521

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

(Campello, Cong, & Zhou, 2024)

Cornell University

Corporate Finance Challenges and AI to the Rescue?

- Graham (2022, JF): current CF models limited ability for explaining/predicting outcomes (around 10% in-sample, worse out-of-sample).
- Mitton (2022, RFS): P-hacking and theory-fitting; need unified definition/framework.
- Spiegel (2023): CF models static and lack interplay among firms and between firms and financial markets.

Corporate Finance Challenges and AI to the Rescue?

- Graham (2022, JF): current CF models limited ability for explaining/predicting outcomes (around 10% in-sample, worse out-of-sample).
- Mitton (2022, RFS): P-hacking and theory-fitting; need unified definition/framework.
- Spiegel (2023): CF models static and lack interplay among firms and between firms and financial markets.
- Abundant big data on managerial actions, market states, & firm outcomes.
- Success of AI with more efficient algorithms and more powerful computation.
- The advancement of large models and their applications in Finance (Giannone et al., 2021; Cong et al., 2020).

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

- What are managers' optimal actions given a certain objective?
  - Counterfactual statements are essential for aiding CF decisions
    - Experiments & surveys: costly, infrequent, sometimes ethically questionable.
    - Reduced-form: restrictive models, limited performance, partial equilibrium.
    - Causal inferences: local and low-dimensional; fragmented knowledge.
    - Structural approaches: need a theory that accounts for the general dynamics.
  - Manager actions $\Leftrightarrow$ environment (e.g., Edmans, Goldstein & Jiang, 2015)
  - Environment stochastic, complex, often unknown.
- CF fundamentally a stochastic control problem:
  - Information set: internal states (accounting, audits, etc.) and market environments.
  - Dynamically learned, high-dimensional and nonlinear.
  - Action space: high-dimensional policy $\rightarrow$ need robust counterfactuals!

Literature and Contribution

- Corporate Finance:
  1. New DDRC overcoming limitations and unifying framework.
  2. ML in CF
     (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)
  3. Non-text-based "large" model tailored for CF.
- Robust Control:
  1. Mostly theory, focus on macro time series rather than utilizing cross-sectional info (e.g., Hansen and Sargent, 2001; Klibanoff et al., 2009; Ju and Miao, 2012).
  2. Application in corporate finance.
  3. 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 2020, quarterly; 19,981 firms, 887,778 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
- Trained using A100 GPU (RedCloud)/P100 (Azure)/T4 (RedCloud) with training time $\sim 3 - 7$ days per set.

AlphaManager (AM) Architecture

Predictive Environment module (PEM, 10 auxiliaries, 300+64+5):

- Warm-up 1976Q1 - 1991Q4; quarterly rolling training & test.

AlphaManager Policy Module: RL 256:64+5:

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

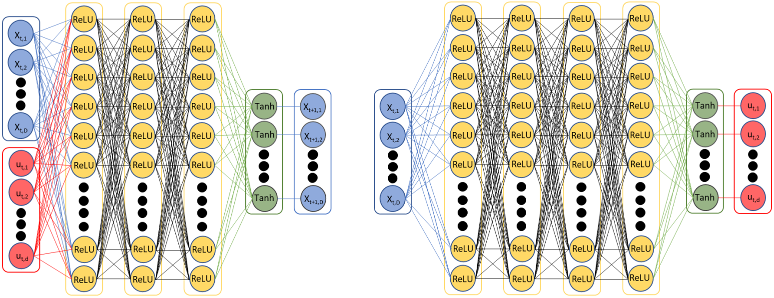

Panel A: Environment Module Panel B: Policy Module

Cornell University

Robust Control and Ambiguity

- Overfitting, data shifts, & uncertainty sources (Hansen and Sargent, 2023):
  - 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 + probability adjustment.
- A bag of PEMs, indexed by $i = 1, 2, ..., I$, and ambiguity aversion
- Pessimistic agent: 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: 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	-9.39%	-12.39%	61.28%	62.65%
Current Asset Growth	-2.16%	-5.94%	53.34%	57.01%
Gross Revenue Growth	36.43%	32.35%	38.76%	35.10%
Accounts Payable Growth	26.84%	26.17%	30.50%	30.02%
COGS Growth	34.95%	30.92%	36.45%	32.78%
Net Interest Paid Growth	81.36%	82.77%	81.58%	82.97%
Inventory Growth	16.48%	13.78%	21.08%	18.17%
Current Liability Growth	10.59%	6.93%	28.30%	25.61%
Receivables Growth	24.95%	22.86%	30.30%	28.33%
Net Income Growth	36.89%	32.88%	39.19%	35.58%
Trading Volume Growth	18.92%	15.90%	21.99%	19.92%
Log Gross Return Growth	49.86%	42.73%	52.34%	45.99%
Market Cap Growth	3.26%	-9.62%	12.98%	2.42%
Enterprise Value Growth	-2.29%	-12.83%	19.40%	10.57%

- 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

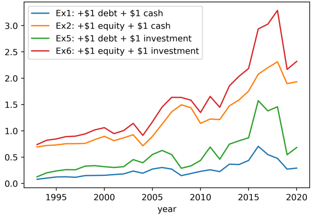

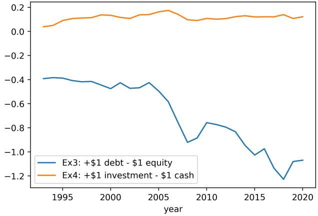

Heterogeneous PEM Performance (MSE) for System States

variable		pre-dotcom		dotcom-GFC		post-GFC	
		mean	std	mean	std	mean	std
log_book_asset	low	2.39%	8.08%	2.12%	7.37%	1.88%	7.34%
	high	1.41%	6.06%	1.31%	5.28%	0.98%	4.38%
log_cogs	low	2.88%	10.16%	3.18%	10.28%	3.54%	11.01%
	high	2.44%	9.33%	2.40%	8.94%	2.00%	8.47%
log_current_liabilities	low	3.36%	10.57%	4.06%	11.51%	4.37%	12.01%
	high	5.11%	12.04%	4.92%	11.94%	4.40%	11.28%
log_market_cap	low	8.60%	16.91%	12.67%	23.79%	8.46%	16.66%
	high	7.30%	15.85%	8.41%	16.66%	4.84%	11.12%
log_enterprise_val	low	9.39%	18.79%	11.15%	20.01%	6.92%	14.27%
	high	5.24%	14.27%	5.20%	11.05%	3.28%	7.83%
macro1 (risk)	low	3.56%	4.29%	3.84%	4.69%	3.69%	4.91%
	high	4.66%	5.67%	5.14%	5.82%	3.57%	4.46%

- Subsample episodes: pre-dot com, dot com to GFC, post-GFC
- Book asset: growth has higher prediction error and std, pre-dotcom has the highest mean and std
- COGS: higher half has declining average MSE but lower half has increasing average MSE
- Market cap, enterprise value, and macr1 (risk): highest prediction error during dotcom to GFC

PEM: Heterogeneous Ambiguity for System States

variable		pre-dotcom		dotcom-GFC		post-GFC	
		mean	std	mean	std	mean	std
log_book_asset	low	6.86%	3.00%	7.41%	3.54%	7.87%	3.87%
	high	7.15%	3.25%	8.66%	3.96%	8.76%	3.81%
log_cogs	low	6.92%	3.27%	8.20%	3.88%	8.61%	4.18%
	high	7.50%	3.52%	9.39%	4.34%	9.17%	4.10%
log_current_liabilities	low	7.50%	3.33%	9.99%	4.63%	10.78%	4.91%
	high	8.65%	3.82%	11.65%	4.80%	11.97%	4.83%
log_market_cap	low	8.57%	3.60%	18.96%	14.08%	14.74%	6.61%
	high	9.44%	4.30%	19.77%	13.11%	14.61%	6.05%
log_enterprise_val	low	8.18%	3.42%	17.41%	12.53%	13.57%	6.06%
	high	8.26%	3.84%	17.08%	11.44%	13.07%	5.43%
macro1 (risk)	low	6.95%	2.32%	9.28%	3.24%	9.84%	3.49%
	high	7.72%	2.99%	12.48%	5.35%	10.70%	3.61%

- Book asset: lower half has lower ambiguity, pre-dot com episode has the lowest mean and std
- COGS: lower half has increasing average MSE
- Market cap, and enterprise value: highest prediction error during dot com to GFC
- Macro1 (risk): lower half has lower mean and std of ambiguity

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 performance: **10.03%** and 2.71%.
- Long-horizon objective: 8.73% and 4.43%

Out-performance of AlphaManager with Short-term Objectives

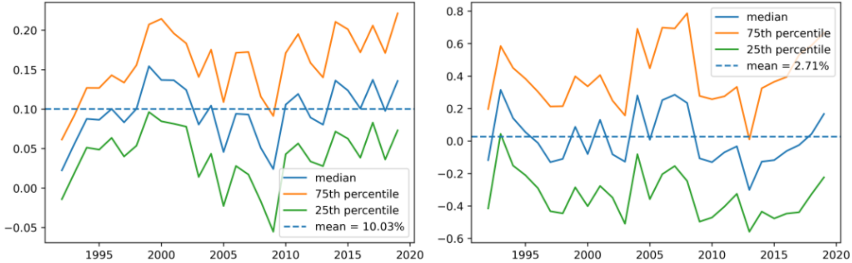

Long-term Performance Under Short-termism

- Objective: 1QTR market cap (left) or enterprise value (right) growth
- Evaluation period: 4QTRs (black lines)
- Benchmark: firm performance in the real data (blue lines)
- Takeaway: short-term AlphaManager decisions do not dominate the actual decisions in mid-term

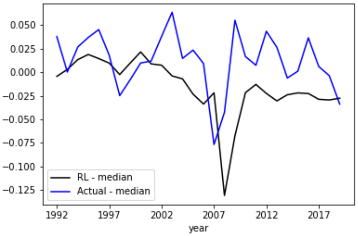

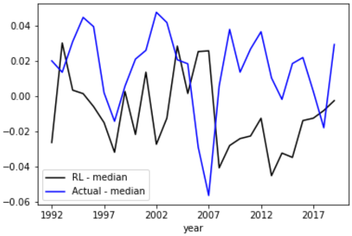

Heterogeneous Performance of AlphaManager Across Firms

variable		pre-dotcom		dotcom-GFC		post-GFC	
		mean	std	mean	std	mean	std
log_book_asset	low	8.13%	4.00%	6.48%	4.74%	6.67%	4.92%
	high	4.54%	4.02%	0.74%	5.07%	0.17%	3.50%
log_cogs	low	7.81%	3.89%	6.02%	4.90%	5.82%	5.56%
	high	4.45%	4.19%	1.27%	5.42%	1.22%	4.09%
log_current_liabilities	low	6.79%	4.35%	5.37%	5.51%	4.95%	6.35%
	high	5.21%	4.27%	1.89%	5.34%	1.99%	3.76%
log_market_cap	low	8.35%	3.65%	6.35%	4.97%	6.27%	5.37%
	high	4.18%	4.02%	1.16%	5.15%	0.88%	3.89%
log_enterprise_val	low	8.42%	3.80%	6.63%	4.70%	6.68%	4.98%
	high	4.31%	3.95%	0.74%	5.02%	0.30%	3.57%
macro1 (risk)	low	7.04%	4.31%	4.12%	5.13%	3.46%	5.43%
	high	4.99%	4.22%	3.05%	6.12%	3.35%	5.31%

- Objective: enterprise value growth in the next 2 years
- AlphaManager performance mainly driven by small caps/EVs and low-risk period within each episode (expansions)

Optimal Actions Versus Historical Actions

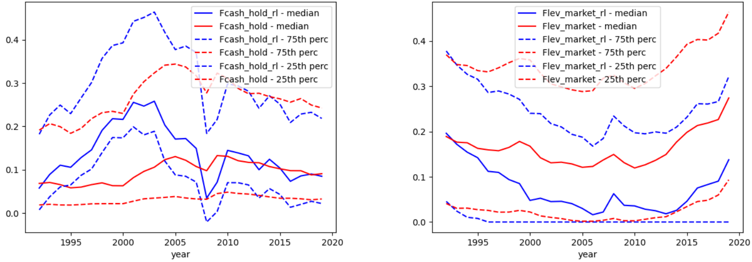

Figure: Optimal decisions (blue lines) vs real decisions (red lines): cash holdings (left) and leverage (right)

Maximizing **next Q market cap**: more acquisitions, increasing cash holdings more, issuing less debt but more equity (i.e., reducing leverage), paying out more dividend, and increasing investment, especially in R&D, allowing more variations in investments, and more repurchases during bad times.

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.
- Interpretations and consistency checks.

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.
- Interpretations and consistency checks.

Conclusion:

- Deep reinforcement learning as heuristic search for optimizing economic goals/objectives.
- Direct optimization for portfolio management with arbitrary objectives.
- A data-driven-robust-control approach to corporate finance.