AI for Finance: Generative Modeling Beyond Words

Lin William Cong[1]

[1] NTU (NBS, CCDS & GIFTS), ABFER, CEPR, & NBER, AI4S2@SLAI

April 2026

Core Themes of Modern AI (Beyond Basic ML)

- AI (McCarthy 1955): technologies allowing machines to perform complex tasks that typically require human intelligence.
- Intelligence (HAI): the ability to learn and perform suitable techniques to solve problems and **achieve goals**, in an uncertain, ever-varying world.
- An economist's perspective of modern AI:
  1. Instruction-driven automation/prediction $\rightarrow$ end-to-end goal-oriented optimization in enlarged modeling space.
  2. Expertise/theory-driven or reduced-form discriminative modeling $\rightarrow$ data-driven generative modeling enabling pre-trained foundational models and AI agents.
- The Tao/Ways of AI:
  1. Larger models, greater computation, general intelligence.
  2. Bigger data and associated computation; pre-training, self-learning.
  3. Goal-oriented algorithms (e.g., RL) in large spaces.
  4. Deep generative modeling for equilibrium/counterfactual analyses.
  5. AI agents for experimentation and as new/representative subjects/species.

Core Themes of Modern AI (Beyond Basic ML)

- AI (McCarthy 1955): technologies allowing machines to perform complex tasks that typically require human intelligence.
- Intelligence (HAI): the ability to learn and perform suitable techniques to solve problems and **achieve goals**, in an uncertain, ever-varying world.
- An economist's perspective of modern AI:
  1. Instruction-driven automation/prediction $\rightarrow$ end-to-end goal-oriented optimization in enlarged modeling space.
  2. Expertise/theory-driven or reduced-form discriminative modeling $\rightarrow$ data-driven generative modeling enabling pre-trained foundational models and AI agents.
- The Tao/Ways of AI:
  1. Larger models, greater computation, general intelligence.
  2. Bigger data and associated computation; pre-training, self-learning.
  3. Goal-oriented algorithms (e.g., RL) in large spaces.
  4. Deep generative modeling for equilibrium/counterfactual analyses.
  5. AI agents for experimentation and as new/representative subjects/species.
- AI for Social Sciences (e.g., Finance): (off-the-shelf) ML + text analytics
  - Return prediction, dim. redux, alt. data, etc., trained through examples.
  - Text analytics/LLMs; customized language models (Hoberg & Manela, 2025).
  - Causal ML, Causal Human+Machine Learning (e.g., Cong & Yang, 2025)

AI for Economics and Finance

Economic analysis traditionally studies three components:

- **Agents**: preferences, beliefs, action sets, and optimization under an environment.
- **Environments**: institutions, markets, policy, constraints, information structure, and rules of the game.
- **Interactions**: strategic behavior among agents that jointly shape environments and determine outcomes.

Modern AI expands/revolutionize each layer:

- Agents (consumers, voters, investors, CEOs, regulators, etc.):
  1. Large action/solution spaces;
  2. Data-driven optimization under potentially learned preferences;
  3. End-to-end, goal-oriented optimization of economic objectives.
- Environments (markets, organizations, social networks, political sys):
  1. Large & realistic environment modules & large computation;
  2. Learned from data, not all expert-specified;
  3. Potentially robust against overfitting/data shift/ambiguity.
- Interactions:
  1. Large-scale games with rich heterogeneity & interactions among agents;
  2. Learned from data with data-driven equilibrium analysis;
  3. Emergence phenomenon.

AI for Economics/Social Sciences

- What is special about social science research, such as finance research?

AI for Economics/Social Sciences

- What is special about social science research, such as finance research?
- I. Complexity, Heterogeneity, Learning, and Purpose-Driven
  1. Probabilistic, context-specific (heterogeneous), mergent interactions.
  2. Low signal-to-noise, non-stationary, variable, adaptive, etc.
  3. Value-laden with high ethical and systemic risks.
  4. (Hetergoeneous) goal-oriented; stakeholder interactions and agencies.
- II. Methods and Approach
  1. Multi-modal, qualitative, unstructured, not strictly ordered, non-systematic.
  2. Interpretation and causality.
  3. What-if questions. The future of economics lies not in explaining the past or predicting the world, but in re-imagining and generating the future.
- III. Data limitations:
  1. Scarce (mostly observational), missing, etc.
  2. Privacy-sensitive.
  3. Endogenous and time-varying data generation.
- Limitations of extant generative modeling in economics or social sciences (e.g., DSGE models, Computable General Equilibrium Models, Structural Equation Models, Agent-Based models).

Generative Modeling

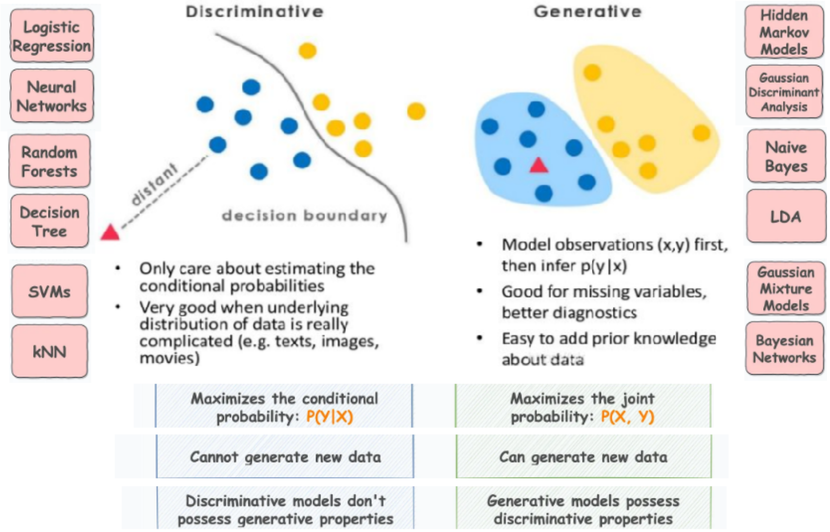

Reinforcement Learning as Efficient Heuristic Search

- Economically guided heuristic search in a large decision/action space.
- People learn by interacting with the environment in an active and sequential way, to optimize some **rewards**.
- Direct construction of portfolios with flexible objectives:
  - "**AlphaPortfolio: Goal-Oriented Investment Management Through Deep Reinforcement Learning**" (Cong, Tang, & Wang, 2024); JFE R&R.
- An alternative framework to studying CF using DL, RL & IRL:
  - "**AlphaManager: A Data-Driven-Robust-Control Approach to Corporate Finance**" (Campello, Cong, & Zhou, 2022).

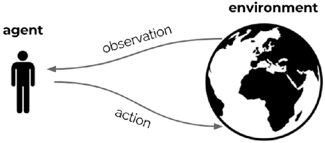

- At each step t the agent:
  - Receives observation O_t (and reward R_t)
  - Executes action A_t
- The environment:
  - Receives action A_t
  - Emits observation O_{t+1} (and reward R_{t+1})

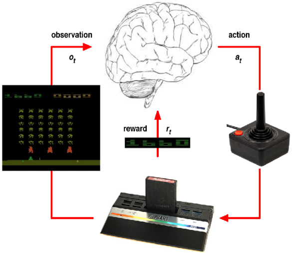

Goal-Oriented Portfolio Management Through Transformer-Based RL

- **AlphaPortfolio: Cong, Tang, & Wang (2019)**
  - First "large" model in finance (a few million parameters) with Transformer/attention/offline DRL in 2019.
  - End-to-end, any objective; not just terminal wealth or additive utilities.
  - Action-space interactions, interpretation through distillations.
- Architecture Innovation: Transformer + Cross-Asset Attention Network + RL

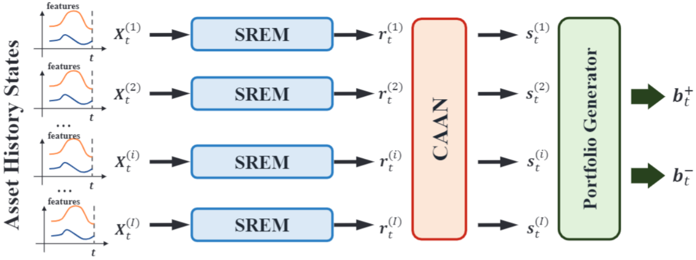

Application to U.S. Equities

- CRSP, Compustat, EDGAR, etc.
- July 1965-June 2016: pre-1990, 1990-2016, post 2001
- 51 characteristics/signals with 1-12 month lags, similar to (Freyberger et.al., 2019).
  - Past-return based.
  - Investment-related characteristics: annual percentage change in total assets (**Investment**),change in inventory over total assets (**IVC**),
  - Profitability-related characteristics such as gross profitability over the book-value of equity (**Prof**) or return on operating assets (**ROA**),
  - Intangibles such as operating accruals (**OA**) and tangibility (**Tan**)
  - Value-related characteristics such as the book-to-market ratio (**BEME**) and earnings-to-price (**E2P**)
  - Trading frictions such as the average daily bid-ask spread (**Spread**) and standard unexplained volume (**SUV**).
- Macro variables and alternative data: MD&A (10-K & 10-Q), Risk Factor Discussions (1A of 10-K), Analyst Reports, etc.

AlphaPortfolio Performance on Test Samples

	AP Performance				AP Excess Alpha					
	(1)	(2)	(3)		(4)	(5)	(6)	(7)	(8)	(9)
Firms	All	$> q_{10}$	$> q_{20}$	Factor Models	All $\alpha(\%)$	All R^2	$> q_{10}$ $\alpha(\%)$	$> q_{10}$ R^2	$> q_{20}$ $\alpha(\%)$	$> q_{20}$ 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

	AP Performance				AP Excess Alpha					
	(1)	(2)	(3)		(4)	(5)	(6)	(7)	(8)	(9)
Firms	All	$> q_{10}$	$> q_{20}$	Factor Models	All $\alpha(\%)$	All R^2	$> q_{10}$ $\alpha(\%)$	$> q_{10}$ R^2	$> q_{20}$ $\alpha(\%)$	$> q_{20}$ R^2
Return(%)	11.11	17.86	14.08	CAPM	9.64***	0.041	19.01***	0.025	15.51***	0.049
Std.Dev.(%)	6.82	6.74	6.00	FFC	10.91***	0.093	19.29***	0.316	15.48***	0.326
Sharpe	1.63	2.65	2.35	FFC+PS	10.86***	0.093	18.45***	0.352	14.33***	0.413
Skewness	0.12	0.42	-0.19	FF5	10.73***	0.134	19.21***	0.401	15.85***	0.342
Kurtosis	0.38	0.76	0.43	FF6	10.65***	0.135	18.94***	0.413	15.56***	0.357
Turnover	0.20	0.24	0.24	SY	-	-	-	-	-	-
MDD	0.04	0.02	0.02	Q4	11.67***	0.234	21.90***	0.308	17.45***	0.232

Panel Trees as Interpretable & Effective Greedy Search

- Interpretable and not-so-large AI models (not supervised/unsupervised).
- Human intelligence and *The Art of War*: divide and conquer.
- Conventional trees not designed for panel data, economic objectives, and interpretability comes at the cost of overfitting.
- Panel trees: panel data analytics with global split criteria using economic goals instead of local split criteria based on statistical error minimization.

Panel Trees as Interpretable & Effective Greedy Search

- Interpretable and not-so-large AI models (not supervised/unsupervised).
- Human intelligence and *The Art of War*: divide and conquer.
- Conventional trees not designed for panel data, economic objectives, and interpretability comes at the cost of overfitting.
- Panel trees: panel data analytics with global split criteria using economic goals instead of local split criteria based on statistical error minimization.
- Creating test assets and latent factor models:
  - "**Growing the Efficient Frontier on Panel Trees**" (Cong, Feng, He, & He, 2021); JFE forthcoming.
- Uncommon factors and macro regimes for asset pricing:
  - "**Sparse Modeling Under Grouped Heterogeneity with an Application to Asset Pricing**" (Cong, Feng, He, & Li, 2022).
- Heterogeneous predictability and trading:
  - "**Mosaics of Predictability**" (Cong, Feng, He, & Wang, 2024).

Growing the Efficient Frontier on Panel Trees: Cong et al. (2021)

P-Tree Test Assets and Factor Models

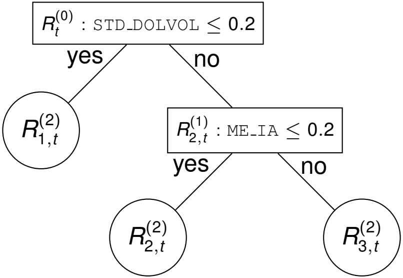

- The second split gives us three leaf basis portfolios and a updated SDF:

$$f_t^{(2)} = \widehat{\Sigma}_2^{-1} \widehat{\mu}_2 R_t^{(2)} = w_{21} R_{1,t}^{(2)} + w_{22} R_{2,t}^{(2)} + w_{23} R_{3,t}^{(2)},$$

- For the second split, the algorithm searches over all leaf nodes, characteristics, and breakpoints.

- The split criterion is calculated based on the entire cross section, thus P-Tree and its SDF are global.

A P-Tree Grown on U.S. Equities

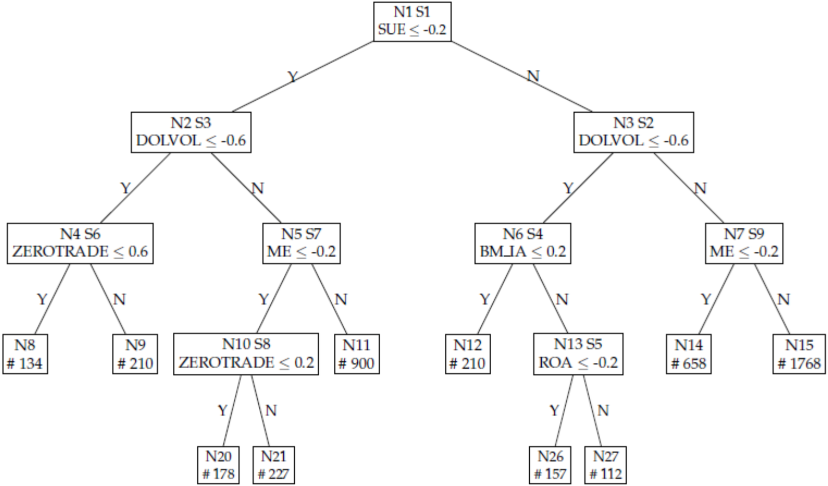

A Grown P-Tree

	SR	α_{CAPM}	α_{FF5}	α_{Q5}	α_{RP5}	α_{IP5}	SR	α_{CAPM}	α_{FF5}	α_{Q5}	α_{RP5}	α_{IP5}
	Panel B1: 20 Years In-Sample (1981-2000)						Panel B2: 20 Years Out-of-Sample (2001-2020)					
P-Tree1	7.12	1.85	1.88	1.71	1.64	1.24	3.24	1.34	1.31	1.23	1.30	0.91
P-Tree1-5	12.14	0.80	0.79	0.77	0.75	0.71	2.98	0.52	0.47	0.49	0.48	0.32
P-Tree1-10	19.16	0.64	0.65	0.62	0.61	0.57	2.57	0.36	0.29	0.33	0.34	0.19
P-Tree1-15	27.74	0.53	0.53	0.52	0.52	0.49	2.48	0.30	0.23	0.27	0.27	0.16
P-Tree1-20	41.12	0.47	0.47	0.46	0.46	0.45	2.57	0.26	0.20	0.24	0.25	0.15
	Panel C1: 20 Years In-Sample (2001-2020)						Panel C2: 20 Years Out-of-Sample (1981-2000)					
P-Tree1	5.82	1.50	1.46	1.50	1.37	1.36	4.35	1.50	1.59	1.35	1.43	1.10
P-Tree1-5	9.61	0.66	0.65	0.66	0.64	0.62	4.11	0.57	0.71	0.51	0.60	0.30
P-Tree1-10	14.24	0.46	0.46	0.46	0.45	0.45	4.07	0.40	0.49	0.35	0.44	0.23
P-Tree1-15	20.52	0.38	0.38	0.38	0.37	0.37	3.98	0.36	0.45	0.31	0.40	0.18
P-Tree1-20	26.95	0.34	0.34	0.34	0.33	0.33	3.82	0.31	0.40	0.26	0.35	0.14

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

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

  - 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

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

- Reduced-form approaches: expertise-driven and ad hoc
- Structural approaches: search within theoretically tractable space
- 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

AlphaManager: Campello, Cong, & Zhou (2022)

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

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

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

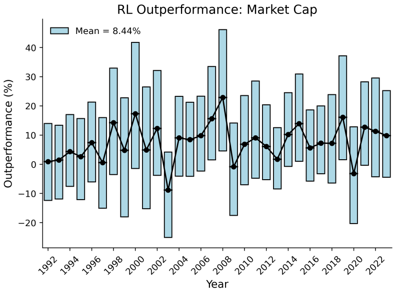

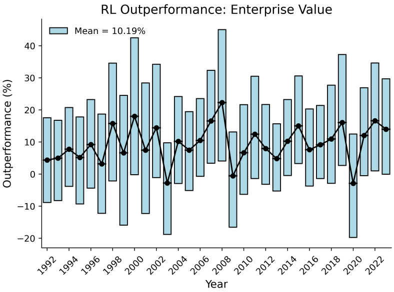

Economic Questions DDRC Enables Asking for the First Time

Piecing together CF research and answering new questions

- Ambiguity and the need for theory/reduced-form/structural models.
  - Boundary of data-driven approach.

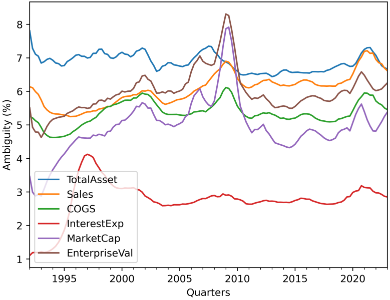

Economic Questions DDRC Enables Asking for the First Time

Piecing together CF research and answering new questions

- Ambiguity and the need for theory/reduced-form/structural models.
  - Boundary of data-driven approach.
- Ambiguity-guided transfer learning/pre-training.
  - Combining insights and predictions from other approaches/applications.
  - For example, investment-CAPM (Liu, Whited, and Zhang, 2009) links investment, sales, profitability and expected returns
  - For each observation, use investment-CAPM to synthesize counterfactual observations where sales, profitability and expected returns are impacted by a counterfactual investment ratio
  - Then, use these synthetic data to train PEM
- Holistically studying managerial objectives.
  - Actions are not additive $\rightarrow$ cannot simply study partial objectives.

Revealed Objectives and Inverse Reinforcement Learning

- Revealed preference argument and Min-Max formulation:

$$\min_{r(\cdot,\cdot)} \max_{g(\cdot)} \mathbb{E}\left[\int_{t_0}^{t_0+T} r(X_t, g(X_t)) - r(X_t, u_t^*)\, dt\right]$$
$$s.t.\ dX_t = f(X_t, g(X_t))dt + \Sigma_x dB_t. \quad (4)$$

  - u_t^* – real/optimal managerial decisions
  - r reflects real objectives $\Longleftrightarrow$ g^* yields zero difference

- Min-Max $\Rightarrow$ Min + Max + $g(X_t) = u_t^*$ or $D(g) = 0$, where

$$D(g) = \min_{r(\cdot,\cdot)} \mathbb{E}\left[\int_{t_0}^{t_0+T} r(X_t, g(X_t)) - r(X_t, u_t^*)\, dt\right] \quad (1)$$
$$s.t.\ dX_t = f(X_t, g(X_t))dt + \Sigma_x dB_t \quad (2)$$

- Generative Adversarial Networks (GANs)

AI Agents and Agent-Based Simulation

- Need to first understand the **Behavioral Economics of AI: LLM Biases and Corrections** (Bini, Cong, Huang, & Jin, 2025)

Questions that study the psychology of preferences

1. Prospect theory – diminishing sensitivity
2. Prospect theory – Loss aversion
3. Prospect theory – Probability weighting
4. Narrow framing
5. Hyperbolic discounting
6. Ambiguity aversion

Questions that study the psychology of beliefs

1. Sample size neglect (1) – binominal
2. Sample size neglect (2) – average estimation
3. Sample size neglect (3) – absolute vs. relative sample size
4. Base rate neglect
5. Conjunction fallacy
6. Gambler's fallacy
7. Confirmation bias
8. Anchoring
9. Overconfidence – over-precision
10. Overconfidence – over-estimation

Economic experiments on preferences, beliefs, and optimizations

Prospect theory – diminishing sensitivity

In addition to whatever you own, you have been given \$1,000. You now need to choose between the following two options:
A1: (\$1,000, 0.5), meaning winning \$1,000 with 0.5 probability and winning zero with 0.5 probability
B1: (\$500), meaning winning \$500 with certainty.

In addition to whatever you own, you have been given \$2,000. You now need to choose between the following two options:
A2: (−\$1,000, 0.5), meaning losing \$1,000 with 0.5 probability and losing zero with 0.5 probability
B2: (−\$500), meaning losing \$500 with certainty

Rational responses: A1 and A2 (less risk averse) | B1 and B2 (more risk averse)
Human responses: human participants tend to choose B1 (risk averse in gain region) and A2 (risk loving in loss region)

AI Agents and Agent-Based Simulation

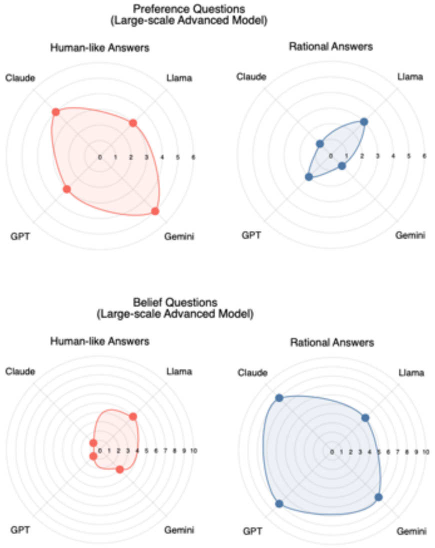

For experimental questions that study the psychology of beliefs:

- LLMs' responses become more rational for more advanced models
- LLMs' responses become more rational for models with a larger parameter scale

For experimental questions that study the psychology of preferences:

- LLMs' responses become more irrational and human-like for more advanced models
- LLMs' responses become more irrational and human-like for models with a larger parameter scale

Why (Large) Language and World Models Are Not Enough

Cong (2025) "**Economic World Models and Data-Driven Generative Equilibria**"

AI world models learn predictive or generative environment dynamics and support planning under alternative action sequences.

They are powerful, but economic analysis usually needs three additional objects:

1. **Equilibrium:** actions, beliefs, endogenous aggregates, and retraining must be jointly consistent.
2. **Information and Beliefs:** who knows what, and how beliefs update, often drives the mechanism.
3. **Interventions and Endogenous Constraints:** policies change feasible sets, information, and rules — not just rewards.

A frozen simulator evaluates outcomes under T_θ. An economic counterfactual may require T_{θ^*}, where

$$\theta^* = \mathcal{L}\big(\mathcal{D}(\pi^*, \mu^*, \theta^*; i)\big).$$

Simulation generates trajectories; equilibrium selects internally consistent ones.

Economic World Models (EWM)

An **Economic World Model (EWM)** is a data-driven generative system representing an economy or social system.

State $S_t \in \mathcal{S}$

- prices, balance sheets, networks, institutions, information
- latent economic states inferred from data

Agents

$$a_t^n = \pi_\theta^n(\mathcal{I}_t^n), \qquad n = 1, \ldots, N$$

- heterogeneous objectives and constraints
- policies may be learned from data (or represented by AI agents)

Environment dynamics

$$S_{t+1} \sim T_\theta(S_t, A_t, i)$$

- $A_t = (a_t^1, \ldots, a_t^N)$
- i denotes interventions (policies, institutional changes)
- θ includes parameters learned from data

What Economic World Models Enable

1. Equilibrium Analysis

- high-dimensional equilibrium outcomes
- policy and institutional counterfactuals
- DDGE when behavior, beliefs, and learning co-evolve

2. Transition Analysis

- feasible, adaptive, and equilibrium transitions
- shock propagation and adoption waves
- path-wise evaluation of reforms and crises

3. System Inference

- latent-state recovery
- mechanism discovery
- historical explanation and measurement

Robustness and counterfactual fragility are cross-cutting:
Lucas critique appears as regime-dependent data shift once interventions change behavior.

Power of GenAI: Data-Driven Counterfactual/Equilibrium

- Counterfactual equilibrium, Lucas Critique, deep parameters, and a new simulation-based micro-foundation.
- Writing Quality and Soft Informationin the GenAI Age: Evidence from Online Credit Markets (Cong, Guo, Zhao, & Zhou, 2024)
  1. Writing well matters for online credit applications, depending on the dimensions; LLMs can enhance writing quality and perceived loan quality based on texts.
  2. LLM differentially more helpful for high-credit borrowers who treat good credit and quality writing as substitutes, whereas they are empirically complements for receiving funding.
  3. Our proprietary language models reveal that LLM adoption decreases soft information conveyed, as it reduces variability thus utility of loan descriptions, creating more misallocation (10-15% reduction in lender ROI).
  4. When lenders adjust their lending decisions to account for LLM adoption, misallocation is mitigated, as soft information partially "recovered" when lenders increasingly rely on hard information.
  5. In a counterfactual equilibrium, borrowers' LLM adoption level solves a fixed-point problem: with rational borrowers and low adoption costs, a unique full-adoption monotone equilibrium ensues.

Data-Driven Generative (Counterfactual) Equilibrium: Cong (2025)

1. Quality distribution can be resampled from historical cases.
2. Borrowers as goal-oriented generative agents.
3. Lenders as data-driven discriminative agents.
4. Both lenders and borrowers' actions should be optimal in equilibrium.
5. GE features fixed points (e.g., LLM/tech adoption or interest rate).

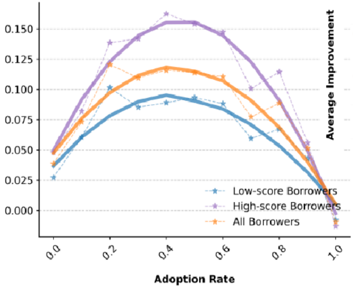

The set of agents $\mathcal{N} = \mathcal{G} \cup \mathcal{D}$ consists of:

- **Generative agents** $i \in \mathcal{G}$ with compact high-level action sets $A_i^H \subset \mathbb{R}^{k_i}$ and low-level policies $a_i^L \sim G_i(\cdot|a_i^H)$.
- **Discriminative agents** $j \in \mathcal{D}$ updating beliefs $\mu_j(\cdot|\mathbf{D}^C)$ on counterfactual data $\mathbf{D}^C = \Delta(\mathbf{a}^H, \mathbf{a}^L, \omega)$.

Each agent's payoff is:

$$u_i\left(a_i^H, a_i^L; \mathbf{a}_{-i}^H, \mathbf{a}_{-i}^L, \{a_j\}, \omega\right) \quad \text{(Generative)}$$

$$\mathbb{E}^{\mu_j(\cdot|\mathbf{D}^C)}\left[u_j(a_j; \mathbf{a}^H, \mathbf{a}^L, \omega)\right] \quad \text{(Discriminative)}$$

DDGE exists If

1. The strategy space $\prod_{i\in\mathcal{G}}\{0,1\}$ is a complete lattice,
2. The best response correspondence $\mathrm{BR}_i(\mathbf{a}_{-i}^H)$ is monotone

Sufficient Condition for Monotonicity If payoffs u_i satisfy:

$$u_i(1, a_i^L; \mathbf{a}_{-i}^H) \geq u_i(0, a_i^L; \mathbf{a}_{-i}^H) \quad \forall \mathbf{a}_{-i}^H,$$

then BR_i is monotone.

AI for Economics and Finance — EWM and DDGE

Three Messages

- Modern AI enlarges the modeling of agents, environments, and interactions.
- Economic World Models provide a common framework for system inference, transition analysis, robustness, and equilibrium analysis.
- Data-Driven Generative Equilibrium is the equilibrium concept for learned systems in which behavior, beliefs, and environments co-evolve.

Why It Matters

- moves beyond prediction in fixed environments
- enables high-dimensional, data-driven counterfactual equilibrium analysis
- contributes back to simulation and world models by adding equilibrium discipline, endogenous constraints, and institutionally meaningful interventions

AI enlarges what economics can represent and simulate;
equilibrium gives those learned systems scientific meaning.