ALPHAMANAGER: A DATA-DRIVEN-ROBUST-CONTROL APPROACH TO CORPORATE FINANCE

- by Murillo Campello, Lin William Cong, and Luofeng Zhou

Discussion by Rik Sen (University of Georgia)

AN OVERVIEW OF WHAT THE PAPER DOES

- A novel framework to enhance corporate finance decision-making with two core modules:
  - ***Predictive Environment Module (PEM)***: Uses supervised deep learning to predict future firm states based on current states and managerial decision variables (leverage, cash, investments etc.), trained on a large dataset (Compustat/CRSP 1976–2023)
  - ***Decision-Making Module (DMM)***: Employs offline reinforcement learning (RL) to recommend optimal managerial actions that maximize objectives like firm value, incorporating robust control to address model uncertainty.
- AlphaManager outperforms historical managerial decisions (e.g., 10.1% quarterly outperformance for short-term market cap growth)
- Complements traditional methods by addressing the complexity of high-dimensional, nonlinear, and dynamic corporate environments

FINDINGS AND IMPLICATIONS

Findings:

- PEM achieves good predictive accuracy (e.g., 64.7% R^2 for book asset growth, 3.2% for market cap out-of-sample)
- DMM outperforms historical decisions: 10.1% quarterly gain for short-term market cap, 8.7% for long-term
- Short-term focus boosts immediate returns but increases ambiguity; long-term strategies are more stable
- Heterogeneous performance: Strongest in manufacturing, trade, and small/illiquid firms during high-risk periods

Implications:

- Complements traditional models by identifying where theory/causal analysis is needed (high ambiguity scenarios)
- Offers actionable policy recommendations
- Opens new research avenues: learning managerial preferences via inverse RL.

MY OVERALL TAKE AND COMMENTS

- *Innovative and groundbreaking approach using tools largely missing from corporate finance research*

- A caveat to my comments before getting into them: overweights a reduced-form empiricists' point of view

MY OVERALL TAKE AND COMMENTS

- ***Innovative and groundbreaking approach using tools largely missing from corporate finance research***

- A caveat to my comments before getting into them: overweights a reduced-form empiricists' point of view

- Overarching comment: Clarify <u>identification assumptions</u>
  - Will help with quicker adoption by the profession of the framework presented in this paper

QUICK COMMENT: ASSESSING PEM USING PREDICTIVE PERFORMANCE

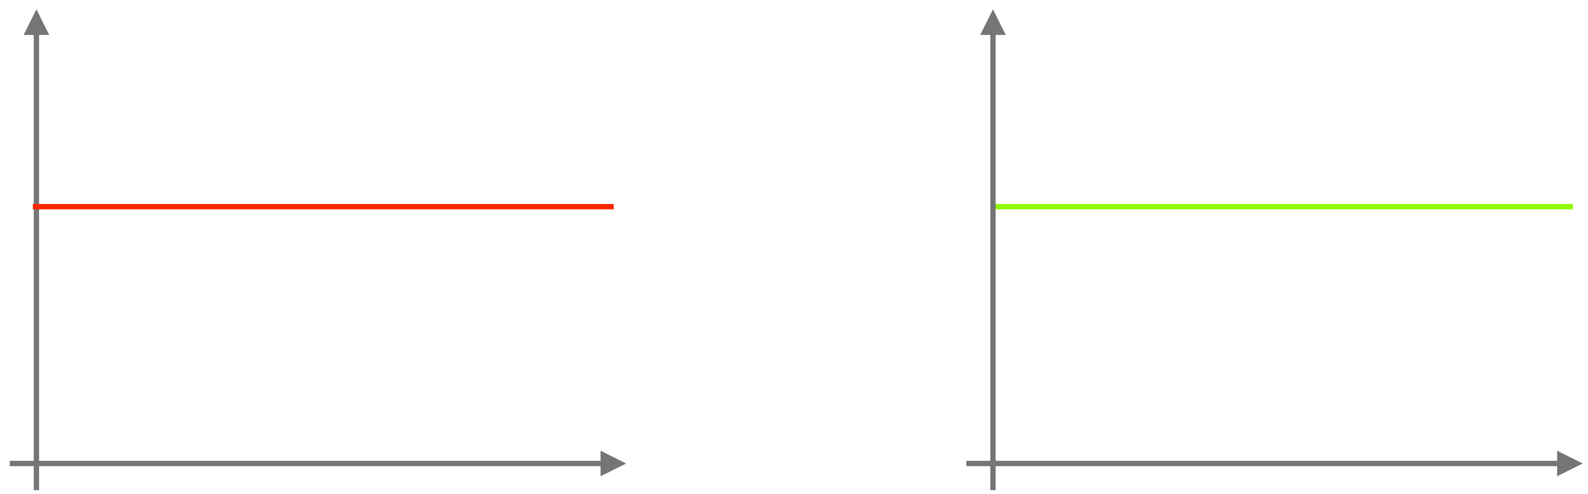

Decision variable *Objective (e.g. firm profit/value)*

- Perfect prediction of the objective possible in the above case
- Yet, we learn ***nothing*** about how decision variable affects the objective
- *The paper present predictive accuracy to assess PEM (Table 2), what about learning about the parameters of interest?*
- **Suggestion**: supplement with an appropriate ambiguity metric for causal parameter values

I AM USED TO SEEING IDENTIFICATION ASSUMPTIONS

- An instrumental variable strategy provides *causal estimates* under five assumptions: relevance, exclusion, SUTVA, exchangeability/independence, and monotonicity
- A structural model provides estimates of the parameters of interest under the assumption that the model is correct
- Under *what assumptions* will this framework uncover ***causal links*** between decision variables and states in the PEM?
  - Will help understand the strengths and weaknesses of this method relative to standard approaches
    - Currently no mention of weaknesses, which I think is suboptimal from the perspective of selling the paper

CONSIDER THE FOLLOWING DATA

Decision variable

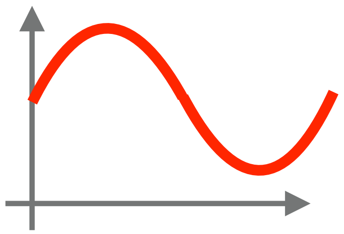

State variable X

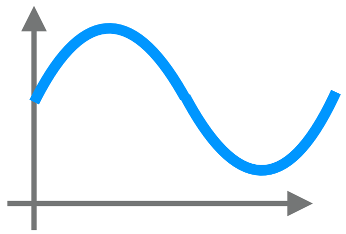

Objective

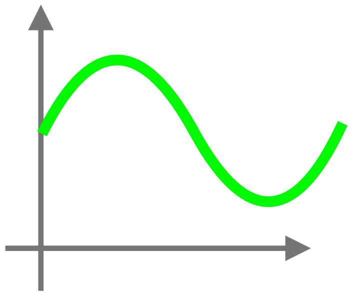

- The state variable can perfectly predict the objective
- It seems the decision variable “does not matter”

CONSIDER THE FOLLOWING DATA

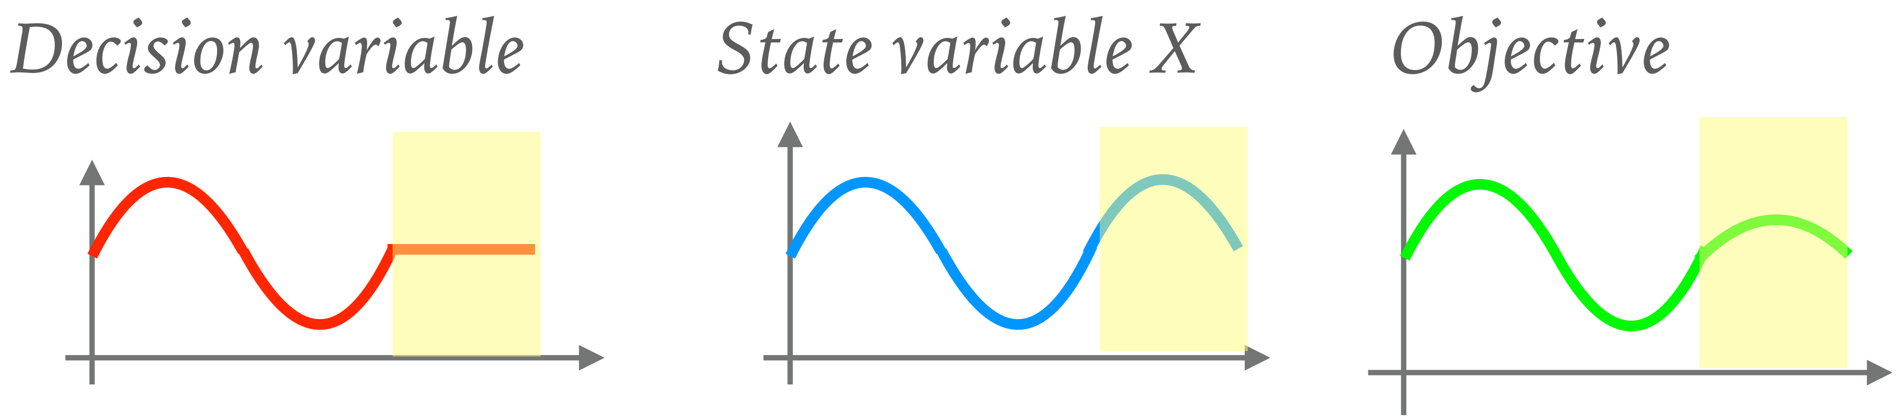

- But suppose that the decision variable counterfactually became flat, then this is what would have happened
  - E.g., state variable = productivity, decision = investment, objective = profit
- ***Takeaway:*** need suboptimal deviation of decision variables for the model to learn the parameters of interest correctly

NEED DEVIATION FROM PERFECT OPTIMAL BEHAVIOR

- If the data are generated by managers who optimize perfectly, there is no observability of cost of suboptimal decisions
- Therefore, you need the assumption: that *managers deviate from perfectly optimal behavior*
  - Noise in choosing the decision variable is useful for identification
  - Is a Bayesian learning assumption is enough? (Not sure; more on this later)
- Also, the **heterogeneity** of suboptimal deviations matters
  - If one manager is optimizing perfectly while another is not, the PEM model learning about casual effect of managerial decision on outcomes will only from the second manager

A DIFFERENT POINT: CONSIDER THREE CASES

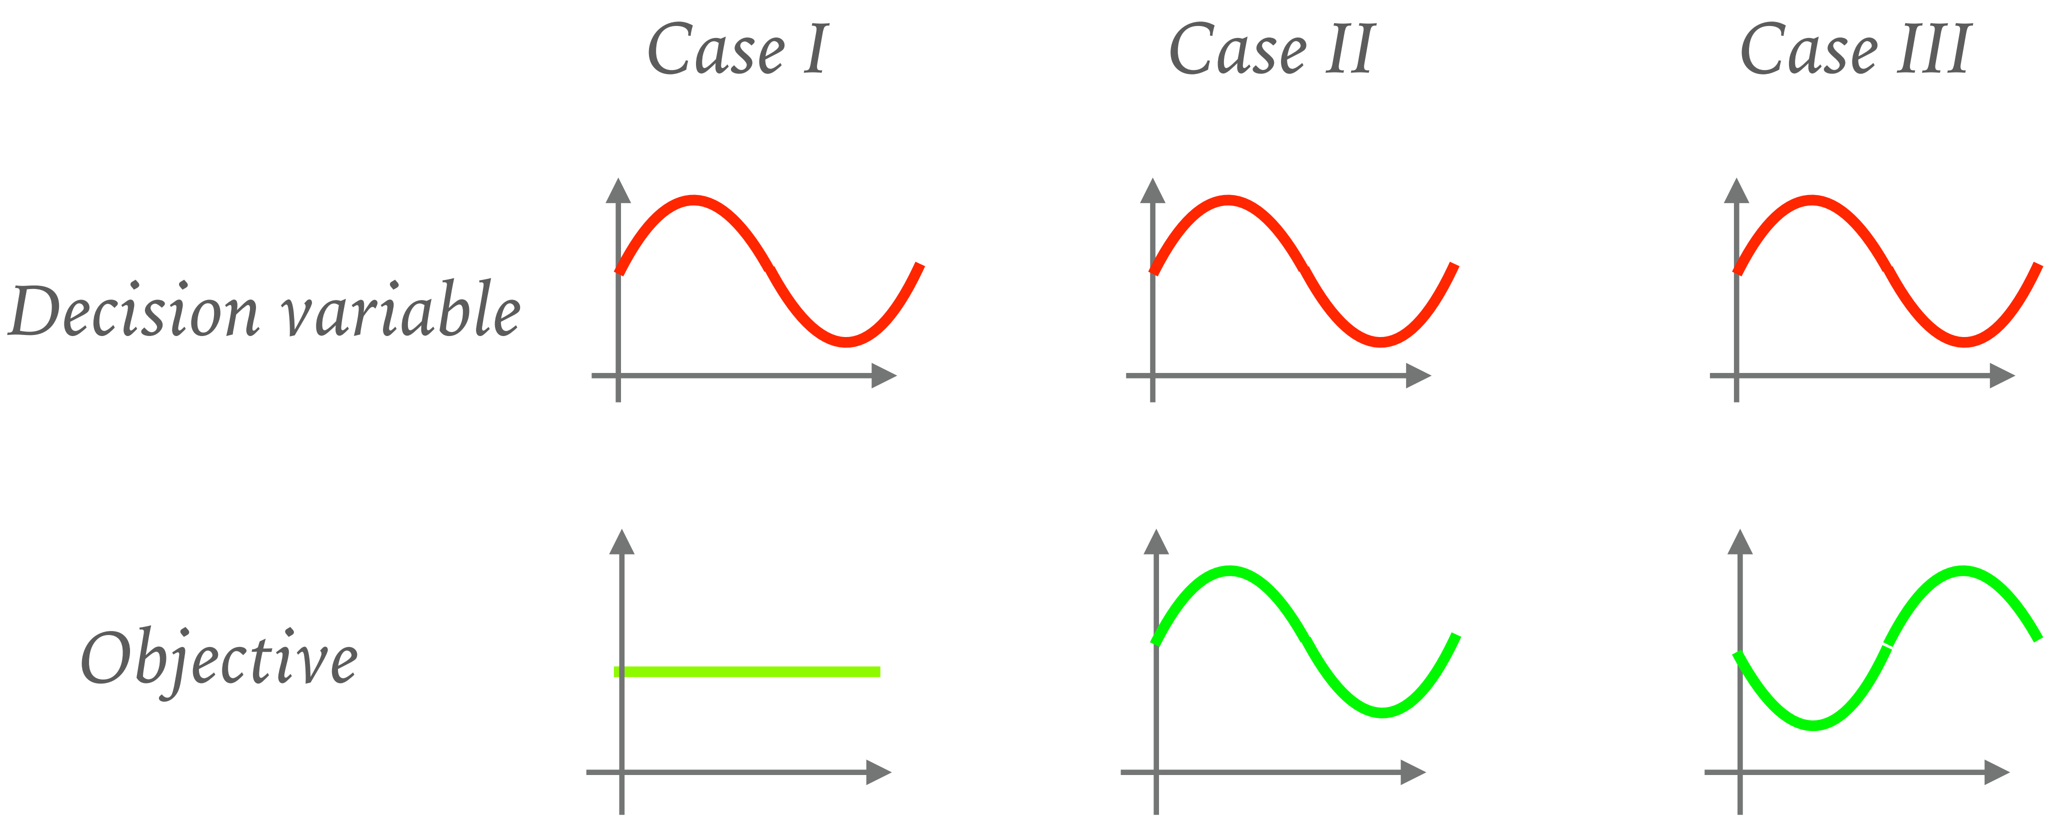

➤ How do you think the decision variable matter for outcomes?

CONSIDER THREE CASES

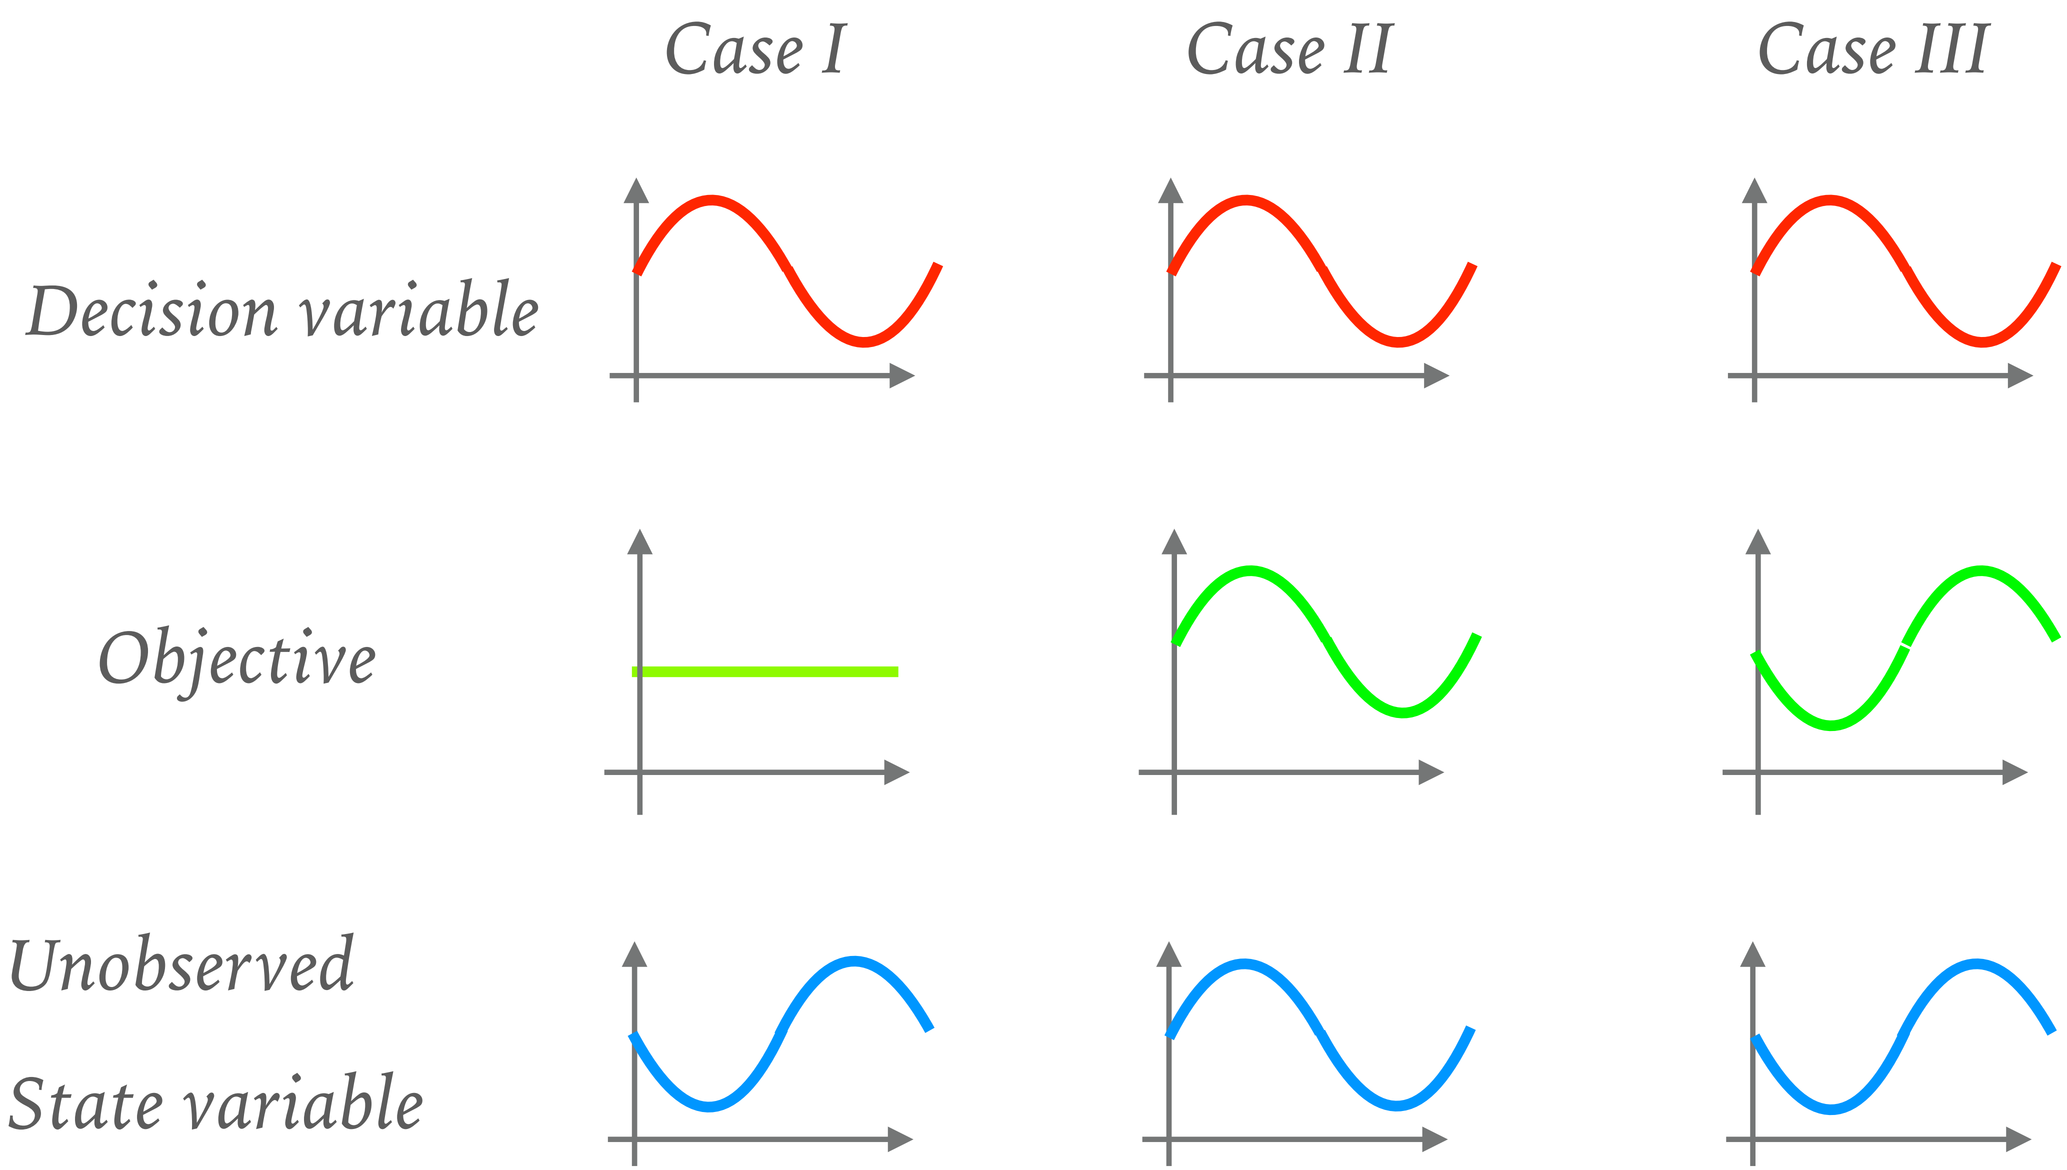

- Interpretations different now!
- Assumption 2: *All relevant state variables are observed*
  - standard omitted variable concern

IMPLICATION OF OBSERVABILITY

- This suggest that ***measurement error*** in observed state variables also matters
  - A weaker form of unobservability

- **Heterogeneity** of precision of state variable across firms also likely matters for what is being estimated

OTHER ASSUMPTIONS THAT ARE LIKELY NEEDED

- When looking at ability to predict changes in firm value, clarify assumption about *market efficiency*
  - If the markets observe managerial actions continuously and incorporate them, we should not be able to predict changes in firm value *even when managerial actions do matter for value*
  - **Heterogeneity** in market efficiency could lead to learning more about firms whose prices reflect information with a delay

- Another assumption is that the world is *(near-)Markovian*
  - E.g., actions should not take 10 periods to start to affect outcomes

SUGGESTION: FIRST BUILD CONFIDENCE IN THE TOOL

- Show the performance on ***simulated data*** where we know the data-generating-process precisely
  - Consider simple models as well as models such as more involve models e.g., Henessy-Whited, that feature dynamic optimization
  - Explore importance of noise, Bayesian learning, etc., for PEM model performance
- Will help clarify the situations and questions for which the tool works well and when it does not work as well

SUGESTION: LINK TO PEARL'S VIEW OF CAUSALITY

- In finance research most causal thinking is using the potential outcomes framework (Rubin/Angrist)
- But there are others, e.g., Pearl's framework, which is likely is a better fit for the framework of this paper
- Look into style of assumptions there

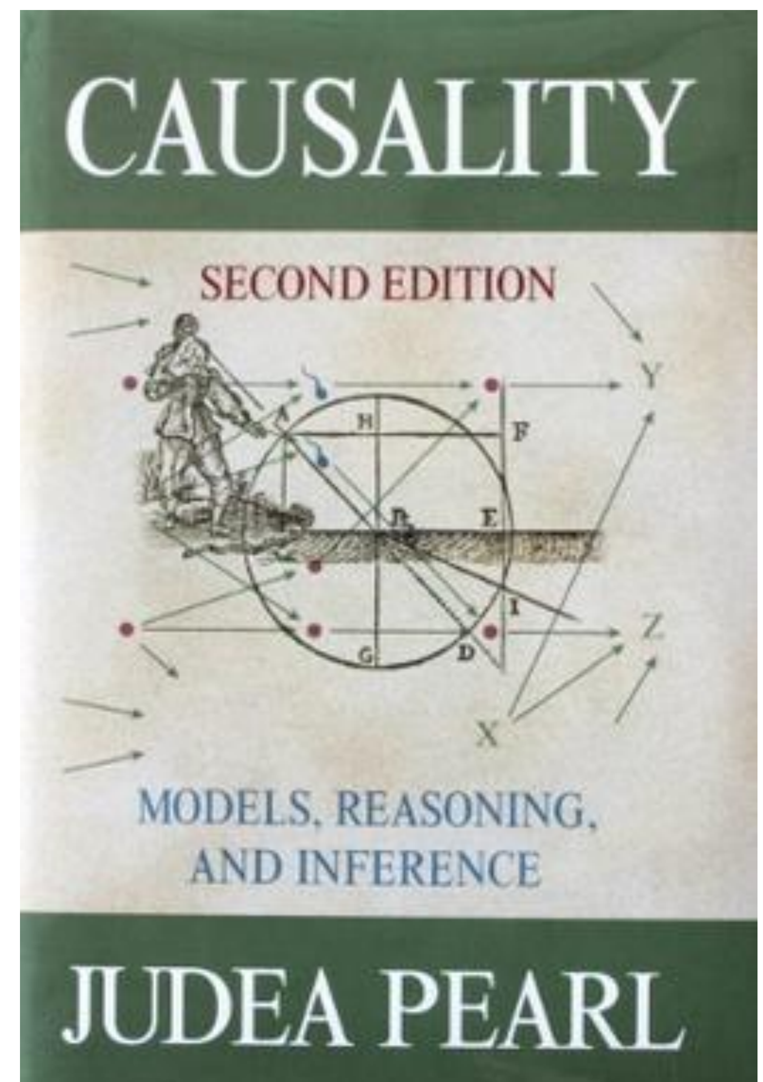

MINOR COMMENT: WHEN A QUESTION IS NOT ANSWERABLE

- In IV, we do not have a good instrument… we say sorry we do not know
- In structural model, changing the parameter has minimal effect on observables… we say parameter is not identified
- I think you make a claim that ambiguity will take care of this, but I am not sure if that is always the case
- Think a bit more about this and explain

CONCLUDING THOUGHT

- I fed this paper into an AI that relies on a similar kind of neural network used here — an LLM

- This is what it had to say:

*"**AlphaManager** is a bold step forward, blending cutting-edge AI with corporate finance to address complex decision-making challenges. Its strengths—novelty, empirical depth, and adaptability—position it as a potential game-changer. However, its technical complexity and limited economic interpretability present hurdles. By improving clarity, justifying choices, and linking to existing approaches and theory, the authors can elevate this work into a landmark contribution."*

THANK YOU!