

Agrawal, Gans, Goldfarb Prediction, Judgment and Uncertainty

Discussant: Andrea Prat, Columbia University

The role of AI in decision making

- “AI is great but there is a lot it cannot do”
- This paper gives a name to what AI cannot do – *judgment*
- Explores how judgment and prediction interacts
- Shows judgment and prediction are likely to be complements – convincing once we see the model but not ex ante
- Analyzes the relative demand for both when: prediction quality improves, the problem becomes more complex, contracting environment, etc
- Builds on Bolton and Faure-Grimaud’s (2009) model of “thinking”
- The beginning of an exciting research program...


Comments

Example

1. Difference between state and payoff
2. Difference between reliability and accuracy
3. Complexity as number of states?
4. Endogeneity of state distribution?

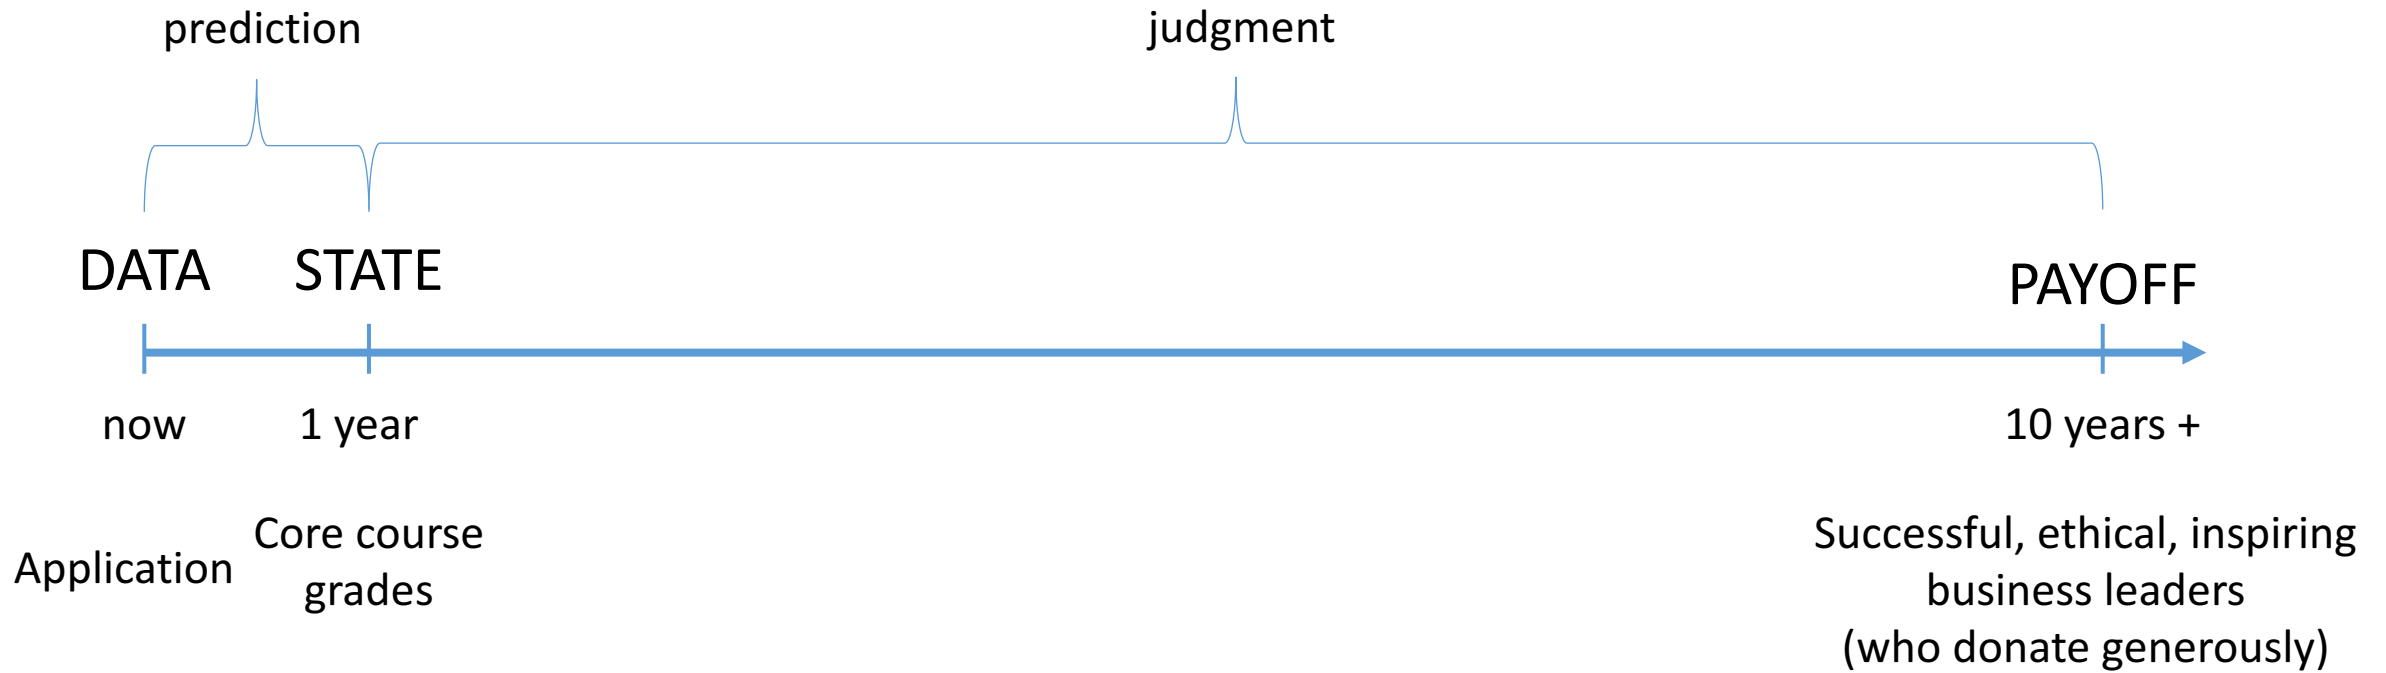
Example: College Admissions

- “Let AI run MBA admissions”
- Use all available info about successful applicants to predict course performance
 - (handle censored observations problem)
- Get a pretty accurate predictor of course performance
- So, let's do it!

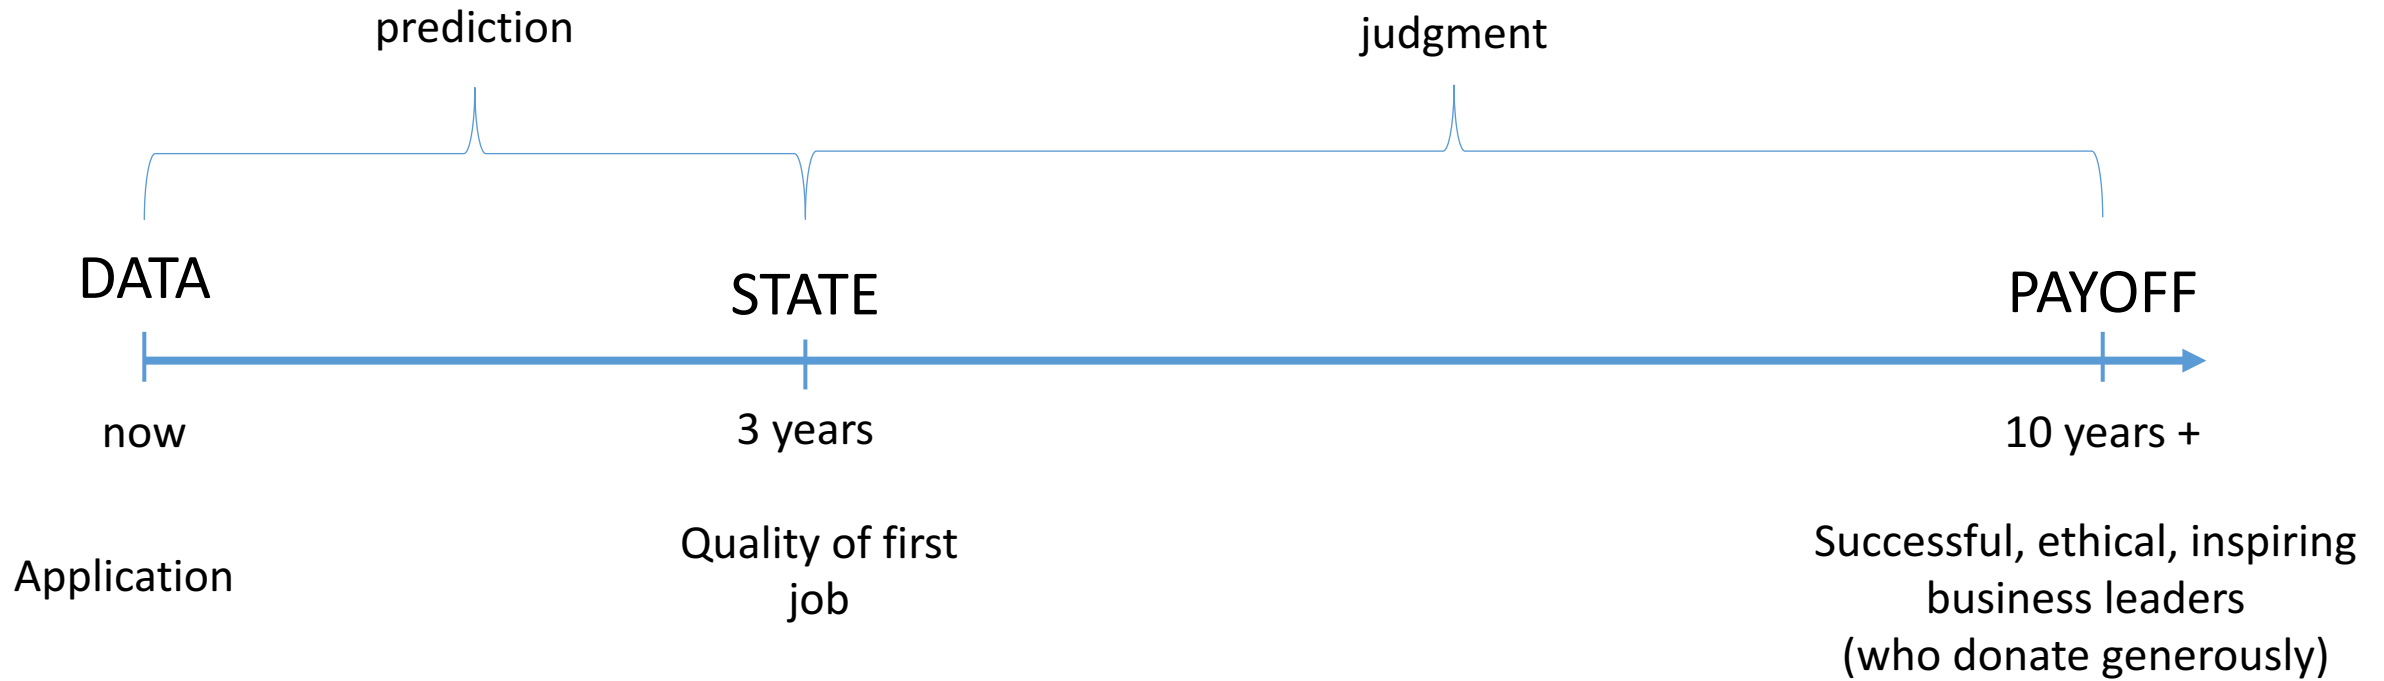


Ultimate Goal?






- The problem is that the state differs from the payoff
- Why not pick a different state?




- The problem is that the state differs from the payoff
- Why not pick a different state?
 - State = Job (3 years from now)?

prediction



- The problem is that the state differs from the payoff
- Why not pick a different state?
 - State = Job (3 years from now)?
 - State = Payoff (10 years from now)?



- The problem is that the state differs from the payoff
- Why not pick a different state?
 - State = Job (3 years from now)?
 - State = Payoff (10 years from now)?
- Lessons: (1) The state (prediction target) is a choice variable; (2) tradeoff between prediction quality and judgment quality;
- **Suggestion #1: endogenize the prediction target and choose the optimal one**

2. Reliability and Accuracy

- Accuracy: probability AI returns a prediction (e)
- Reliability: probability the prediction is correct (a)
- However, the standard output of a (binary) prediction algorithm is a number in $[0,1]$ – the probability state 1 is true
 - The distinction between a and e seems artificial
 - Every algorithm generates a conditional distribution on $[0,1]$ given the true state
- **Suggestion #2. Model $[0,1]$ case;**
 - **Precision cost = $F(\text{entropy})$**
 - **Characterize optimal signal structure + optimal cost**

3. Complexity

- Complexity = number of states N
- Only m states can be predicted
- N increases while m stays constant = prediction technology gets worse
- Does not add much to previous analysis...
- **Suggestion #3. Relate to standard notions of complexity?**
 - Size of set of possible hypotheses.
 - Tradeoff between variance and bias
 - Judgment as a way to narrow the hypothesis set?

4. Endogenous State Distribution

- In problems involving human interaction, the acceptance rule chosen by the DM affects the distribution $F(\vartheta, s)$
 - In fraud detection, the fraudsters look for weaknesses in the system
 - A whole industry around college admissions
 - ...
- AI mostly abstracts from this (huge) problem
- Need for: (i) theory of agent response (a form of judgment); (ii) field experiments – Comparative advantage of economists!!!
- **Suggestion #4 (another paper): Explore effect of endogenizing $F(\vartheta, s)$**

Thank you!