Economics 262: Market Design

Spring 2020

Overview

Market design is broadly about designing interventions in economic systems so as to enhance their performance. The power and potential of market design has recently entered a new era of promise with the rise of Artificial Intelligence. Artificial Intelligence is concerned with the design of intelligent autonomous systems. Such systems are rapidly transforming our society and economy, and have been enabled by the rise of an Internet of Things (IoT) connected to public cloud computing via fast, low-latency network telecommunications.

Yet underlying the outward technological features of AI are fundamental economic and econometric principles which are central to their design and implementation. In short, to perform well, an AI system must "think like economists" - it must

1. Have a "loss" or "utility" function to evaluate the consequences of various actions it can take,
2. Form predictions and probabilistic assesments about its environment
3. Make decisions and learn the optimal plan of action in the face of uncertainty, which is a cycle that repeats and iteratively improves.

Although much of established/production AI being deployed by firms and organizations have largely been focused on achieving (1) and (2), the future trend is towards AI increasingly encompassing (3) in the form of Reinfocment Learning (RL).

Yet each of these areas is one where the human mind and resulting judgments/decisions can experience significant and prevalent departures from the rational actor model of behavior assumed in standard economic models. This suggests powerful complementarity between the design of AI systems that improve/correct the inefficiencies due to the cognitive decision biases of human actors in a way that can create immense societal value.

This course is about understanding/examining the conceptual backdrop to this core idea via class lectures, and directly engaging with the design problem in a concrete way through a team project.

Course Content

The primary goals of this course are twofold. The first is to explore the advances in behavioral economics that has documented a variety of systematic departures from this fully rational actor model that is presumed in neoclassical economics. Our second goal is to show how these very same economic principles play a role in the modern development of AI and are being encoded

in AI technologies. This juxtaposition between machine and human decision making enables us to see how the development of AI can be leveraged to enhance the efficiency of human decisions and thereby act as a complements (as opposed to substitutes) to each other.

Towards the end the course will be structured around three main themes that connect economics with design of AI solutions.

(Hu)Man vs Machine Decisions

Artificial intelligence technologies augment the capacities of human decision makers. This complementarity exists precisely because human decisions exhibit certain systematic departures from the rational actor model that dictates how economic decisions are made under classical economic theory. Documenting these biases and understanding their implications for the performance of economic systems has been the purview of the "Heuristics and Biases" paradigm in Behavioral Economics, which has resulted in 3 Nobel Prizes in Economics. In probabilistic reasoning, statistical prediction, and decision planning, there are systematic patterns whereby human behavior is naturally complemented through machine intelligence for enhancing efficiency of economic decisions.

Prediction needs Parsimony.

Much of current generation AI is in fact concerned with prediction. The key reason AI has found resurgence is that many AI problems have been successfully recast as prediction problems. These prediction problems are then solved through automated machine learning. This has spawned a whole industry of "predictive analytics" which have become rampant in society and embedded in nearly any digital product we nowadays consume. However, machine learning predictions require models to generate predictions, and as the famous statistician George E.P. Box once said "All models are wrong; some models are useful." ML techniques internalize Box's wisdom through the creation of parsimonious models of decision environment, which we operationalize through Bayesian Graphs and the "Bias-Variance" tradeoff.

Decisions are Dynamic.

Economic decisions are rarely a one-shot event. Instead the actions are linked in a sequence that unfolds over time with recurring feedback about the performance of previous decisions. Such decision problems dramatically change the action space, allowing for contingent strategies in decision making (e.g., if X happens do A, if Y happens do B). We will examine a class of these problems known as *Markov Decision Problems* which are widely prevalent in economic environments where AI technologies are being applied. Solving such decision problems can be enormously complicated in full generality due to their combinatorial complexity. Reinforcement learning is an alternative algorithmic strategy for learning the optimal action to undertake inside the context of MDP's. Central to the design of RL algorithms is recognition of the fundamental dynamic incentive of "exploitation vs exploration" - it may be advantageous to make mistakes early in the process of learning (the exploration phase), with the intention of

speeding the transition to having better information against which to optimize our decision (the exploitation phase).

Teaching Philosophy and Class Structure

This course will attempt to adopt elements of a SAIL (Structured Active In-Class Learning) format class structure. My own sympathy to the SAIL approaching to teaching the class content is very much in line with the following observation by an education researcher:

"… I point to the following unwelcome truth: much as we might dislike the implications, research is showing that didactic exposition of abstract ideas and lines of reasoning (however engaging and lucid we might try to make them) to passive listeners yields pathetically thin results in learning and understanding–except in the very small percentage of students who are specially gifted in the field." A. Arons (1997)

Although for we will not be explicitly adopting the SAIL format (e.g., I will still lecture on the reading material as the primary method of content delivery), I will make an effort to add elements of the SAIL approach to add team discussion and deliberation and diversify the classroom experience. In particular, for a handful of classes there will be a "pop" simple quiz to start the calss on the reading material to incentivize everyone to have a common baseline exposure to the topic. Then a "team" experience will take place during the class, which can take one of two forms depending on the content.

- I will start class by distributing a few discussion questions at the start of class that you will discuss and answer as a team for 10 minutes. Teams will present their perspectives and then the lecture will ensue in a fashion that is enriched by this "pre-work" to engage with the subject matter.

- I will start class with a lecture, and during class break out into teams for one or more segment on solving a problem when an important concept needs an active learning experience to reinfornce.

Finally, an integral part of the course will be the group projects which will be fulfilled through the submission of a written proposal/paper as well as a 20 minute presentation given to the class. The project is aimed at researching and examining 3 inter-related issues:

1.) Identify a real world context/setting/scenario where human bias in economic decision making is a meaningful and important problem where value from machine guided decision making can be realized.

2.) Describe how AI/ML technologies can be harnessed to build a technology that alleviates the bias and create more value out of economic decisions.

3.) Anticipate the frictions/barriers to the adoption of your solution in the setting you considered. What are those factors and why do they arise? How would you design the implementation of your approach to overcome these challenges.

The emphasis is to work with a diverse team to create a new point of view on how to connect economics and AI through an application of your choice. The submitted paper should probably be roughly 10-15 pages in length and more weight given to quality/data driven arguments rather than sheer page count (e.g., short is good).

The development of the project is what I view as the central learning experience of the course, and have prioritized it this way in the grading assessment. In order to motivate continuous attention on your part being directed torwards the goal of the project, I have slated 3 distinct weeks during the semester where instead of one class lecture that week, Edvard and I will meet separately with each team (via Skype by default for scheduling ease, but can be in-person depending on the circumstance) and offer feedback/perspective/guidance on the evolution of the projects. The structure of the meetings will be oriented around discussing prepared slides each team puts together that summarize their current state of thinking/discovery and raises challenges they are having.

Assessment

Assessment will be based on the following components:

- 20% Homework
- 30% Two Exams (equal weight)
- 40% Final Project.
- 10% Class Participation

Course Materials

The course will be based on readings from literature as well as the two books that we will study and available at the bookstore.

The first book is **"Thinking Fast and Slow"** by Daniel Kahneman, which describes the conceptual and historical background behind the core insights of the enormously influential "Heuristics and Biases" research program that Kahneman partly pioneered that isolated key departures in human decisions and judgment from the rational actor model of classical economics. This body of research led to Kahneman winning the Nobel Prize in Economics in 2002.

The second book is “Decision Making Under Uncertainty” by Mykel Kochenderfer, which in written from the perspective of an aeronautical engineer and applies the tools of the rational actor model from economics as a basis for the development of AI solutions to decision problems.

An optional supplementary book is also available at the bookstore - “Prediction Machines” by Agarwal, Gans, and Godfarb – which is an interesting and examination of the implications rapid pace of progress in AI explained in terms of lowering the price of a critical input to efficient decisions – namely predictions.

Administrative Details

The class will meet in 3N1H DRLB (David Rittenhouse Laboratory) on Tues and Thurs from 1.30-3. Course materials such as notes, slides, HW etc will be administered through Canvass.

Amit’s office hours are 10-11 am on Wednesdays in PCPSE 622 (though feel free to get in touch outside these times if you need help).

The TA for the course is Edvard Bakhitov. His office hours are 12.30-2.30 pm on Fridays in PCPSE 500.

Homework Schedule
HW’s will be due in the following week after assigned.

Contact Details
Amit Gandhi
- email: akgandhi@upenn.edu
- Office: PCPSE 622
- Phone 215-898-7409

TA: Edvard Bakhitov
- email: bakhitov@sas.upenn.edu
- Office: PCPSE 635

Course Schedule and Readings

TFS = “Thinking Fast and Slow”
DMU = “Decision Making under Uncertainty”

Introduction to Class (Thurs, Jan 16)

Lecture 1 (Tues, Jan 21)

Topic: Introduction to AI and human decisions
Readings:

a) TFS Introduction;
b) DMU Ch 1.

Part 1: Positive Economics and Normative Opportunities

Lecture 2 (Thurs, Jan 23)

Topic: Who is Homo Economicus: Positive vs Normative Economics
Readings

a) Milton Friedman (1953), “The Methodology of Positive Economics”
b) Uskali Maki (2009), “Unrealistic assumptions and unnecessary confusions: Rereading and rewriting F53 as a realist statement **“**

Lecture 3 (Tues, Jan 28)

Topic: Positive Economics and Utility Maximization
Readings:

- Rubinstein, A., Modeling Bounded Rationality, 7-16
- Andrew Caplin & Mark Dean & Daniel Martin, 2011. "Search and Satisficing," American Economic Review, American Economic Association, vol. 101(7), pages 2899-2922, December.

Lecture 4 (Thurs, Jan 30);

Topic: Taking Friedman to the Test - Do people maximize utility?
Readings:

- Rubinstein, A., Modeling Bounded Rationality, 16-21

- Tversky, A., & Shafir, E. (1992b). Choice under conflict: the dynamics of deferred decision. Psychological Science, 3, 358-361
- Iyengar, S. S., M. R. Lepper, 2000,."When Choice is Demotivating: Can One Desire Too Much of a Good Thing?" Journal of Personality and Social Psychology, 79, 995-100
- Iyengar, Sheena S. & Kamenica, Emir, 2010. "Choice proliferation, simplicity seeking, and asset allocation," Journal of Public Economics, Elsevier, vol. 94(7-8), pages 530-539, August.
- Simonson, Itamar (1989), "Choice Based on Reasons: The Case of Attraction and Compromise Effects," Journal of Consumer Research, 16 (September), 158-73
- Marianne Bertrand, Dean Karlan, Sendhil Mullainathan, Eldar Shafir and Jonathan Zinman; "What's Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment," Quarterly Journal of Economics, 2010, 125(1), pp. 263-305.
- Amos Tversky, Paul Slovic and Daniel Kahneman "The Causes of Preference Reversal " *The American Economic Review* Vol. 80, No. 1 (Mar., 1990), pp. 204-217
- Huber, Joel, John W. Payne, and Christopher P. Puto (1982), "Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis," Journal of Consumer Research, 9 (June), 90-98

Lecture 5 (Tues, Feb 4)

Topic: Taking Friedman to the Test - Do people have utility functions?
Readings:

- TFS Ch. 25 and Ch. 27
- Kahneman, Daniel & Knetsch, Jack L & Thaler, Richard H, 1990. "Experimental Tests of the Endowment Effect and the Coase Theorem," Journal of Political Economy, vol. 98(6), pages 1325-1348
- Kahneman, Daniel & Knetsch, Jack L & Thaler, Richard H, 1991. "The Endowment Effect, Loss Aversion, and Status Quo Bias: Anomalies," Journal of Economic Perspectives, American Economic Association, vol. 5(1), pages 193-206, Winter.

Lecture 6 (Thurs, Feb 6) HW1

Topic: Taking the test from the lab to field: the effects of expertise
Readings:

- Kevin Bryan (Oct 2019), "What randomisation can and cannot do: The 2019 Nobel Prize"
- List, J., "Does market experience eliminate market anomalies?,"The Quarterly Journal of Economics 118 (1), 41-71
- List, J., "Neoclassical theory versus prospect theory: Evidence from the marketplace," Econometrica 72 (2), 615-625

Lecture 7 (Tues, Feb 11)

Topic: What is the Alternative: Meet System 1 and System 2
Readings:

- TFS Ch. 1-6

Lecture 8 (Thurs, Feb 13);

Topic: What is the mechanism through which decision bias emerges: Heuristics and Biases
Readings:

- TFS Ch. 7-10, 11

Lecture 9 (Tues, Feb 18);

Topic: Normative Opportunity: Law of Small Numbers
Readings:

- TFS Ch. 10
- Kahneman and Tversky original "Belief in the Law of Small Numbers" paper
- Recent press: The "irony effect", and "a theory in crisis".

Lecture 10 (Thurs Feb 20);

Topic: Normative opportunity: Probability Judgement
Readings:

- TFS Ch. 13-15
- Hastie and Dawes "Rational Choice in an Uncertain World" Ch. 5 (handout)

Tues Feb 25 – Check in #1

Thurs Feb 27 – Probability and Julia/R review (Edvard to conduct)

Lecture 11 (Tues, March 3);

Topic: Normative opportunity: Regression Fallacy
Readings:

- TFS Ch. 17-18
- Hastie and Dawes "Rational Choice in an Uncertain World" Ch. 7 (handout)

Part 2: Decision Making Under Uncertainty

Lecture 12 (Thurs, March 5);

Topic: AI design: Bayesian Networks and Dimensionality Reduction
Readings:

- DMU Ch. 2.1

Lecture 13 (Tues, March 17); HW2

Topic: Normative Opportunities: Bayes rule and Bayesian Inference
Readings:

- DMU Ch. 2.2
- TFS Ch. 16

Lecture 14 (Thurs March 19);

Topic: AI design - Parameter Learning and MCMC
Readings:

- DMU Ch. 2.3

Tuesday, March 24 – Check-in#2

Thurs March 26 - Midterm

Lecture 15 (Tues, March 31);

Topic: AI design: Influence Diagrams and expert systems
Readings:

- DMU Ch. 3.1-3.20

Lecture 16 (Thurs Apr 2); HW3

Topic: Normative Opportunity: The independence axiom, Allais paradox, and Ellsberg Paradox
Readings:

- TFS Ch. 28-30, 34

Lecture 17 (Tues Apr 7)

Topic: Markov Decision Processes: Theory and Example
Readings:

- DMU Ch. 4.1-4.3

Thurs Apr 9: Check in #3

Lecture 18 (Tues April 14);

Topic: Normative opportunity: dynamic inconsistency
Readings:

- Benhabib, Jess, Bisin, Alberto and Schotter, Andrew, (2010), Present-bias, quasi-hyperbolic discounting, and fixed costs, *Games and Economic Behavior*, **69**, issue 2, p. 205-223.
- Matthew Rabin & Ted O'Donoghue, 1999. "Doing It Now or Later," American Economic Review, American Economic Association, vol. 89(1), pages 103-124, March.

Lecture 19 (Thurs April 16); HW4

Topic: Markov Decision Processes: Numerical approaches and and the curse of dimensionality
Readings:

- DMU Ch. 4.5-4.7

Lecture 20 (Tues Apr 21)

Topic: Learning over time, multiarm bandits, and Q-learning.

Readings:

- DMU Ch. 5.1, 5.2

Class Presentations (Thurs Apr 23)

Class Presentations (Tues Apr 28)

Final Exam