Canonical game theoretic domains

Modeling Human Strategic Behavior

Kevin Leyton-Brown

University of British Columbia
Canada CIFAR AI Chair, Amii

James R. Wright

University of Alberta
Canada CIFAR AI Chair, Amii

THE UNIVERSITY OF BRITISH COLUMBIA

UNIVERSITY OF ALBERTA

Lecture Overview

Non-strategic Domains

Canonical Domains

Non-strategic Domains

Many popular games are not **strategic** in the game theoretic sense:

- Montezuma's Revenge (and many other Atari games)
  - Single-player; no need to reason about other players' incentives
  - Every action always has the same consequence
- Snakes and Ladders
  - Outcomes entirely determined by dice roll
  - No choice of actions
- War (the card game)
  - Outcomes entirely determined by shuffle order
  - No choice of actions

Lecture Overview

Non-strategic Domains

Canonical Domains

Security Games

- A “Defender” wants to prevent attacks on a set of targets
  - Airport terminals and terrorist attacks
  - Staten Island Ferries (and terrorist attacks)
  - Fare evasion on public transit
- But Defender cannot afford to guard every target all the time
  - E.g., not every ferry gets an escort
  - Not every LRT passenger gets checked
- So the Defender has to **randomize** their defenses
- But the Attacker gets to watch the Defender before attacking
  - They observe the Defender’s random distribution before acting
  - But not the Defender’s **realized actions**

Peer Grading

Peer grading:

- Gives students more feedback and exposure to others' work
- Lets us run large classes without giant teams of TAs

But graders might be **strategic** about how much effort they put into their grading

- Low-effort strategy: don't read the assignment; give 80%

How can we incentivize high-effort grading?

- Compare student grades with each other; reward agreement with other graders
- TA spotchecks: randomly grade some assignments; reward agreement with TAs

Recommendation Systems

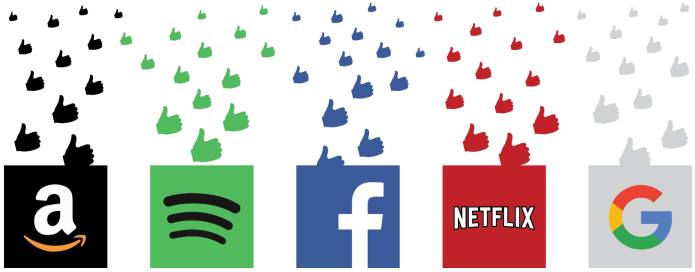

Recommendation systems are **ubiquitous**.

How do we imagine these systems as permitting a **strategic setting**?

Think of all stakeholders in an interaction:

- Users
- System Designers
- Content Providers

Spectrum Auctions

Repurpose radio spectrum from broadcast television to wireless internet

- pay broadcasters for voluntarily relinquishing their licenses
- potentially assign new channels to stations that keep broadcasting
- resell contiguous blocks of spectrum to telecoms

Many **elements of design freedom**; many **objectives**

- Participants' property rights
- Definition of goods to be traded
- Quantity of goods to trade
- Outcomes the market should seek to achieve
  efficiency; revenue; increased competition in the consumer market; bidding simplicity for unsophisticated participants
- Computational tractability

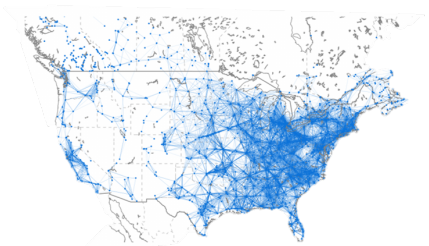

Breakout Rooms

Questions *(roughly, same for Assignment 1)*:

- **Describe the setting** and justify why a strategic model is appropriate
  - Who are the players?
  - What actions are available?
  - Where do their payoffs come from?
  - Why is the setting strategic? *(e.g., how do one agent's actions influence another's payoffs?)*
- What **game representation(s) are appropriate** to model key parts of the domain?
  - Do agents choose actions simultaneously or sequentially *(and if the latter, do they observe each other's moves)*?
  - Do agents interact once or repeatedly?
  - Do agents have knowledge of their and others' payoffs?
- Why do **human behavioral considerations** come into play in this setting?
  - How might skilled actors behave different from unskilled actors?
  - What irrational behaviors might agents exhibit?
  - What data would be useful to have?