SOME GAME THEORY

What is game theory?

- Game theory is a framework for analysing interactions between a set of agents.
- Abstract specification of interactions.
- Describes each agent's preferences in terms of their *utility*.
  - Assume agents want to maximise utility.
- Give us a range of *solution strategies* with which we can make some predictions about how agents will/should interact.
- Game theory is *not* about being selfish.

Why game theory?

- In order to go further with the discussion of auctions, we need to get a bit more formal than we have been.
- That formality needs some game theoretic underpinning.
- Still not terribly formal :-)
- This lecture aims to give a quick overview of some of the main concepts.
- For more detail . . .

Book?

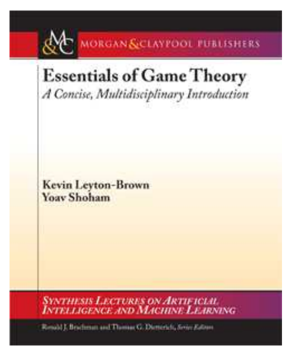

ISBN 978-159-829-5931

Congestion Game

- Using TCP — defective implementation doesn't back-off.

		i	
		defect	correct
j	defect	-3 / -3	-4 / 0
	correct	0 / -4	-1 / -1

- Agent i is the *column player*.
- Agent j is the *row player*.

Normal form games

- An n-person, finite, *normal form* game is a tuple (N, A, u), where
  - N is a finite set of players.
  - $A = A_1 \times \ldots \times A_n$ where A_i is a finite set of actions available to i. Each $a = (a_1, \ldots, a_n) \in A$ is an *action profile*.
  - $u = (u_1, \ldots, u_n)$ where $u_i : A \mapsto \Re$ is a real-valued *utility* function for i.
- Naturally represented by an n-dimensional matrix

Prisoner's Dilemma

		i	
		coop	defect
j	coop	a / a	c / b
	defect	b / c	d / d

- Any game with $c > a > d > b$ is a prisoner's dilemma.

Common payoff games

- Coordination game

	left	right
left	1 / 1	0 / 0
right	0 / 0	1 / 1

- Any game with $u_i(a) = u_j(a)$ for all $a \in A_i \times A_j$ is a common payoff game.

Constant sum games

- Matching pennies

	heads	tails
heads	-1 1	1 -1
tails	1 -1	-1 1

- Any game with $u_i(a) + u_j(a) = c$ for all $a \in A_i \times A_j$ is a constant sum game.

- Rock, paper, scissors

	rock	paper	scissors
rock	0 0	1 -1	-1 1
paper	-1 1	0 0	1 -1
scissors	1 -1	-1 1	0 0

General sum games

- Battle of the Sexes

	this	that
this	1 2	0 0
that	0 0	2 1

- Game contains elements of cooperation and competition.

Strategies

- An agent's *strategy set* is its set of available choices.
- Can just be the set of actions — *pure* strategies.
- *Mixed strategies*, probability distribution over pure strategies.
  - Strategy set S_i is set of all probability distributions over A_i.
- Set of *mixed strategy profiles* is $S_1 \times \ldots \times S_n$.
- The *support* for a mixed strategy s_i is the set of pure strategies $\{a_i | s(a_i) > 0\}$
- The payoff of a mixed strategy is the expected utility of the strategy:

$$u_i(s) = \sum_{a \in A} u_i(a) \prod_{j=1}^{n} s_j(a_j)$$

Solution concepts

- For an agent acting alone we can compute the *optimal strategy*
  - maximises the expected utility.
- In a multiagent setting this isn't very meaningful.
- Best strategy depends on what others are doing.
- *Solution concepts* identify sets of outcomes (subsets of the whole) that are interesting in some way.
- External view — Pareto optimality.
- Internal view — Nash equilibrium.

Pareto optimality

- In multiagent settings it is hard to define "best solution".
  - Can't easily handle tradeoffs between agents' utilities.
  - Which is best outcome in battle of sexes?
- But some outcomes are better than others.
- s *Pareto dominates* s' if for all i, $u_i(s) \geq u_i(s')$ and there is some j such that $u_j(s) > u_j(s')$.
- Defines a partial order over strategies.
- s is *Pareto optimal* if there is no s' such that s' Pareto dominates s.
- "Pareto optimal" is also described as "strictly Pareto efficient".

Nash equilibrium

- If I know how you will play the game, I can maximise. I choose my *best response*.
- i's best response to the strategy profile s_{-i} is the mixed strategy $s_i^* \in S$ such that $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$ for all s_i.
- Best response is not a solution concept since we don't, in general, know what other agents will do.
- But we build the idea of *Nash equilibrium* on top of it.
- A strategy profile $s = (s_1, \ldots, s_n)$ is a Nash equilibrium if, for all agents i, s_i is a best response to s_{-i}.
- Stable, since no agent can do better by switching strategy while everyone else sticks.
- Every game (within reason) has a (mixed strategy) Nash equilibrium.

Dominated strategies

- Let s_i and s_i' are strategies of i. S_{-i} is the set of strategy profiles of the other players.
- s_i *strictly dominates* s_i' if $u(s_i, s_{-i}) > u(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$.
- s_i *weakly dominates* s_i' if $u(s_i, s_{-i}) \geq u(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$ and $u(s_i, s_{-i}) > u(s_i', s_{-i})$ for at least one s_{-i}
- s_i *very weakly dominates* s_i' if $u(s_i, s_{-i}) \geq u(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$
- A *dominant* strategy is one that dominates all others.
- A strategy profile in which every s_i is dominant for i is a Nash equilibrium
  - equilibrium in dominant strategies

Aside — Prisoner's Dilemma redux

		i coop	defect
	coop	3	4
j		3	0
	defect	0	1
		4	1

- "defect" is a dominant strategy.
- Hence the dilemma — the dominant strategy equilibrium is the only outcome that is not Pareto optimal.

Dominated strategies II

- Game with dominated strategies

	L	C	R
U	1	1	0
	3	0	0
M	1	1	0
	1	1	5
L	1	1	0
	0	4	0

- Can eliminate the dominated strategies and simplify the game
- Remove R (dominated by L).

Dominated strategies III

- Game with dominated strategies

	L	C
U	1	1
	3	0
M	1	1
	1	1
L	1	1
	0	4

- M is now dominated by the mixed strategy that picks U and L with equal probability.
- It was not dominated before we removed R.

Dominated strategies IV

- Final game

	L	C
U	1	1
	3	0
L	1	1
	0	4

- This will not remove any Nash equilibria.
- If we only use strict dominance, the order of elimination doesn't matter.

Evolutionarily Stable Strategies

- Consider a large population of agents playing a two player game.
  - Equilibrium strategy
- Is the equilibrium strategy stable against some fraction of the population switching to a different strategy.
- A mixed strategy s is an *evolutionarily stable strategy* if for all other strategies s':
  - $u(s, s) > u(s', s)$; or
  - $u(s, s) = u(s', s)$ and $u(s, s') = u(s', s')$

- Hawk/Dove game

	hawk	dove
hawk	-2 -2	0 6
dove	6 0	3 3

- Unique symmetric Nash equilibrium, $(3/5, 2/5)$.
- Also the unique ESS.
- But, for example, $(dove, dove)$ is not an ESS, though it is Pareto optimal.

- If a mixed strategy s is an evolutionarily stable strategy, then it is a Nash equilibrium.
- Any ESS is a best response to itself, and is therefore an NE.
- The reverse does not hold — only strict Nash equilibria are ESS.
- In a two-player game, given a mixed strategy s, if (s, s) is strict Nash equilibrium, then s is an evolutionarily stable strategy.
- Interesting because we can *learn* ESS and hence NE.

Sequential Games

- In normal-form games we assume moves are simultaneous.
- Another area of game theory studies *sequential* games.
  - Players take it in turns
- We don't have time to look at this.
- Can always map the sequence of moves into a strategy, and consider this to be a very big normal form game.

Bayesian Games

- Everything we have done so far assumes agents know what game they are playing.
- Assume that:
  - Number of players
  - Set of actions
  - Payoffs

  are common knowledge across all players.
- Now look at games of *incomplete information* or *Bayesian* games.
- Represent the lack of knowledge with a probability distribution over a set of games
  - Agents' beliefs about which game they are playing.

- All these games have the same number of players and strategy space.
  - Not a very restrictive assumption.
  - Pad games if necessary with dominated strategies.
- Agents' beliefs are posteriors, based on a common prior conditioned on private signals.
  - Start the same, experience differs.

- Bayesian game over some familiar games

MP

2, 0	0, 2
0, 2	2, 0

$p = 0.3$

PD

2, 2	0, 3
3, 0	1, 1

$p = 0.1$

Coord

2, 2	0, 0
0, 0	1, 1

$p = 0.2$

BoS

2, 1	0, 0
0, 0	1, 2

$p = 0.4$

- Row player can only distinguish between (MP, PD) and $(Coord, Bos)$.

- Usual formal treatment uses the notion of *epistemic type*.
- Defines the payoff that a player gets from a particular outcome.
  - Private value
- A Bayesian game is a tuple (N, A, Θ, p, u) where:
  - N is a finite set of players.
  - $A = A_1 \times \ldots \times A_n$ where A_i is a finite set of actions available to i.
  - $\Theta = \Theta_1 \times \ldots \times \Theta_n$, where Θ_i is the type space of i.
  - $p : \Theta \mapsto [0, 1]$ is a common prior over types; and
  - $u = (u_1, \ldots, u_n)$ where $u_i : A \times \Theta \mapsto \Re$ is a real-valued *utility* function for i.
- Assume everyone knows the game at this level, and each agent knows its own type.

- Notion of strategy is slightly different.
- In a Bayesian game, a pure strategy is:

$$\alpha_i : \Theta_i \mapsto A_i$$

  a mapping from every type i might have to the action it would take were it to have that type.
- A mixed strategy is then just a probability distribution over these.
- $s_j(a_j|\theta_j)$ is the probability that j plays a_j in strategy s_j given that its type is θ_j.

- To analyse a game we need to say what value each strategy has to each agent.
  – Expected utility.
- *Ex post* expected utility.

$$EU_i(s, \theta) = \sum_{a \in A} \left(\prod_{j \in N} s_j(a_j|\theta_j) \right) u_i(a, \theta)$$

  where s is a strategy profile.
- We know the types of other agents, so uncertainty is just in their mixed strategies.

- But we typically don't know other agents' types.
- So, *ex-interim* expected utility:

$$EU_i(s, \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p(\theta_{-i}|\Theta_i) EU_i(s, (\theta_i, \theta_{-i}))$$

  where the embedded EU is an ex-post value and again s is a stragegy profile.
- So, we don't assume we know the types of other agents, instead we compute over all the possible types of other agents and weight each by the probability of the other agent being that type given our own type

- Finally, *ex-ante* expected utility:

$$EU_i(s) = \sum_{\theta_i \in \Theta_i} p(\theta_i) EU_i(s, \theta)$$

  where the embedded EU is an ex-post value ana again s is a strategy profile.
- Ignores any observation of type.

- Now we can define best response.
- Agent i's set of *best responses* to the mixed strategy profile set s_{-i} are:

$$BR_i(s_{-i}) = \arg\max_{s'_i \in S_i} EU_i(s'_i, s_{-i})$$

- Note that this is based on *ex-ante* EU.
- And from that we get an equilibrium definition.
- A *Bayes-Nash* equilibrium is a mixed strategy profile s that satisfies $s_i \in BR_i(s_{-i})$ for all i.
- Thus, just as before, the equilibrium is where everyone plays their best response to everyone else's best response.

- Finally, a stronger form of equilibrium.
- An *ex-post* equilibrium is a mixed strategy profile s such that for all i and θ:

$$s_i \in \arg\max_{s'_i \in S_i} EU_i(s'_i, s_{-i}, \theta)$$

- If there is an ex-post equilibrium, no agent will deviate from it even if it knew everyone's type.
  - Not the same as dominant strategy equilibrium though they often coincide.
- Said another way: don't have to believe that others have a good picture of of your type distribution.

Summary

- This lecture has given a quick tour of the main concepts of game theory.
- Concentrated on normal form games.
- Talked about the main solution concepts.
  - Nash equilibrium
- Wound up with Bayesian games.