A Whirlwind Tour of Game Theory

(Mostly from Fudenberg & Tirole)

Players choose actions, receive rewards based on their own actions and those of the other players.

Example, the Prisoner's Dilemma:

	Cooperate	Defect
Cooperate	$+3, +3$	$0, +5$
Defect	$+5, 0$	$+1, +1$

Strategies and Nash Equilibrium

A **strategy** is a specification for how to play the game for a player. A **pure strategy** defines, for every possible choice a player could make, which action the player picks. A **mixed strategy** is a probability distribution over strategies.

A **Nash equilibrium** is a profile of strategies for all players such that each player's strategy is an optimal response to the other players' strategies. Formally, a mixed-strategy profile σ_*^i is a Nash equilibrium if for all players i:

$$u^i(\sigma_*^i, \sigma_*^{-i}) \geq u^i(s^i, \sigma_*^{-i}) \forall s^i \in S^i$$

Nash equilibrium of Prisoner's Dilemma: Both players defect!

Matching Pennies

	H	T
H	$+1, -1$	$-1, +1$
T	$-1, +1$	$+1, -1$

No pure strategy equilibria

Nash equilibrium: Both players randomize half and half between actions.

More on Equilibria

Dominated strategies: Strategy s_i (strictly) dominates strategy $s_i^{'}$ if, for all possible strategy combinations of opponents, s_i yields a (strictly) higher payoff than $s_i^{'}$ to player i.

Iterated elimination of strictly dominated strategies: Eliminate all strategies which are dominated, relative to opponents' strategies which have not yet been eliminated.

If iterated elimination of strictly dominated strategies yields a unique strategy n-tuple, then this strategy n-tuple is the unique Nash equilibrium (and it is strict).

Every Nash equilibrium survives iterated elimination of strictly dominated strategies.

Multiple Equilibria

A coordination game:

	L	R
U	$9,9$	$0,8$
D	$8,0$	$7,7$

U, L and D, R are both Nash equilibria. What would be reasonable to play? With and without coordination?

While U, L is pareto-dominant, playing D and R are "safer" for the row and column players respectively...

Existence of Equilibria

Nash's theorem, translated: every game with a finite number of actions for each player where each player's utilities are consistent with the (previously discussed) axioms of utility theory has an equilibrium in mixed strategies.

Idea 1: Reaction correspondences. Player i's reaction correspondence r_i maps each strategy profile σ to the set of mixed strategies that maximize player i's payoff when her opponents play σ_{-i}. Note that r_i depends only on σ_{-i}, so we don't really need all of σ, but it will be useful to think of it this way. Let r be the Cartesian product of all r_i. A fixed point of r is a σ such that $\sigma \in r(\sigma)$, so that for each player, $\sigma_i \in r_i(\sigma)$. Thus a fixed point of r is a Nash equilibrium.

Kakutani's FP theorem says that the following are sufficient conditions for $r : \Sigma \to \Sigma$ to have a FP.

1. Σ is a compact, convex, nonempty subset of a finite-dimensional Euclidean space.

   Satisfied, because it's a simplex

2. $r(\sigma)$ is nonempty for all σ

   Each player's playoffs are linear, and therefore continuous, in her own mixed strategy. Continuous functions on compact sets attain maxima.

3. $r(\sigma)$ is convex for all σ

   Suppose not. Then $\exists \sigma^{'}, \sigma^{''}$ such that $\lambda\sigma^{'} + (1-\lambda)\sigma^{''} \notin r(\sigma)$ But for each player i,

   $$u_i(\lambda\sigma_i^{'} + (1-\lambda)\sigma_i^{''}, \sigma_{-i}) =$$

   $$\lambda u_i(\sigma_i^{'}, \sigma_{-i}) + (1-\lambda)u_i(\sigma_i^{''}, \sigma_{-i})$$

   so that if both $\sigma^{'}$ and $\sigma^{''}$ are best responses to σ_{-i}, then so is their weighted average.

4. $r(\cdot)$ has a closed graph

The correspondence $r(\cdot)$ has a closed graph if the graph of $r(\cdot)$ is a closed set. Whenever the sequence $(\sigma^n, \hat{\sigma}^n) \to (\sigma, \hat{\sigma})$, with $\hat{\sigma}^n \in r(\sigma^n) \forall n$, then $\hat{\sigma} \in r(\sigma)$ (same as upper hemicontinuity)

Suppose that there is a sequence $(\sigma^n, \hat{\sigma}^n) \to (\sigma, \hat{\sigma})$ such that $\hat{\sigma}^n \in r(\sigma^n)$for every n, but $\hat{\sigma} \notin r(\sigma)$. Then there exists $\epsilon > 0$ and σ' such that

$$u_i(\sigma_i^{'}, \sigma_{-i}) > u_i(\hat{\sigma}_i, \sigma_{-i}) + 3\epsilon$$

Then, for sufficiently large n,

$$u_i(\sigma_i^{'}, \sigma_{-i}^{n}) > u_i(\sigma_i^{'}, \sigma_{-i}) - \epsilon > u_i(\hat{\sigma}_i, \sigma_{-i}) + 2\epsilon$$

$$> u_i(\hat{\sigma}_i^n, \sigma_{-i}^n) + \epsilon$$

which means that $\sigma_i^{'}$ does strictly better against σ_{-i}^n than $\hat{\sigma}_i^n$ does, contradicting our assumption.

Learning in Games*

How do players reach equilibria?

What if I don't know what payoffs my opponent will receive?

I can try to learn her actions when we play repeatedly (consider 2-player games for simplicity).

Fictitious play in two player games. Assumes stationarity of opponent's strategy, and that players do not attempt to influence each others' future play. Learn *weight functions*

$$\kappa_t^i(s^{-i}) = \kappa_{t-1}^i(s^{-i}) + \begin{cases} 1 & \text{if } s_{t-1}^{-i} = s^{-i} \\ 0 & \text{otherwise} \end{cases}$$

*Fudenberg & Levine, *The Theory of Learning in Games*, 1998

Calculate probabilities of the other player playing various moves as:

$$\gamma_t^i(s^{-i}) = \frac{\kappa_t^i(s^{-i})}{\sum_{\tilde{s}^{-i} \in S^{-i}} \kappa_t^i(\tilde{s}^{-i})}$$

Then choose the best response action.

Fictitious Play (contd.)

If fictitious play converges, it converges to a Nash equilibrium.

If the two players ever play a (strict) NE at time t, they will play it thereafter. (Proofs omitted)

If *empirical marginal distributions* converge, they converge to NE. But this doesn't mean that play is similar!

t	Player1 Action	Player2 Action	κ_T^1	κ_T^2
1	T	T	$(1.5, 3)$	$(2, 2.5)$
2	T	H	$(2.5, 3)$	$(2, 3.5)$
3	T	H	$(3.5, 3)$	$(2, 4.5)$
4	H	H	$(4.5, 3)$	$(3, 4.5)$
5	H	H	$(5.5, 3)$	$(4, 4.5)$
6	H	H	$(6.5, 3)$	$(5, 4.5)$
7	H	T	$(6.5, 4)$	$(6, 4.5)$

Cycling of actions in fictitious play in the matching pennies game

Universal Consistency

Persistent miscoordination: Players start with weights of $(1, \sqrt{2})$

	A	B
A	$0, 0$	$1, 1$
B	$1, 1$	$0, 0$

A rule ρ^i is said to be ϵ-**universally consistent** if for any ρ^{-i}

$$\lim_{T\to\infty} \sup \max_{\sigma^i} u^i(\sigma^i, \gamma_t^i) - \frac{1}{T} \sum_t u^i(\rho_t^i(h_{t-1})) \leq \epsilon$$

almost surely under the distribution generated by (ρ^i, ρ^{-i}), where h_{t-1} is the history up to time $t - 1$, available for the decision-making algorithm at time t.

Back to Experts

Bayesian learning cannot give good payoff guarantees.

- Suppose the true way your opponent's actions are being generated is not in the support of the prior – want protection from unanticipated play, which can be endogenously determined.

- The Bayesian optimal method guarantees a measure of learning something close to the true model, but provides no guarantees on received utility.

- Can use the notion of experts to bound regret!

Define *universal expertise* analogously to universal consistency, and bound regret (lost utility) with respect to the best expert, which is a strategy.

The best response function is derived by solving the optimization problem

$$\max_{\mathcal{I}^i} \mathcal{I}^i \vec{u}_t^i + \lambda v^i(\mathcal{I}^i)$$

$\vec{u}_t^i$ is the vector of average payoffs player i would receive by using each of the experts

$\mathcal{I}^i$ is a probability distribution over experts

λ is a small positive number.

Under technical conditions on v, satisfied by the entropy:

$$-\sum_s \sigma(s) \log \sigma(s)$$

we retrieve the exponential weighting scheme, and for every ϵ there is a λ such that our procedure is ϵ-universally expert.