Game Theory

David C. Parkes

parkes@eecs.harvard.edu

Electronic Market Design: Spring, 2003

Basics

Agents I, strategy $s_i \in S_i$.

Def. A *strategy*, s_i, is a complete contingent plan; defines the action an agent will take in all states of the world.

Essentially, the *joint strategies* of agents, $s = (s_1, \ldots, s_N)$ define the outcome of the game, and the utility of agents.

Def. The payoff function, $u_i(s)$, defines agent i's utility for strategy profile $s = (s_1, \ldots, s_N)$.

Wrapped up in the term "utility" is an implicit assumption that *rational* agents behave as *expected utility maximizers*.

Def. A strategic-form game $G = (S_1, \ldots, S_N; u_1, \ldots, u_N)$.

Example: Ascending-price Auction

State of the world, (p, x), defines the ask price $p \geq 0$ and whether agent is holding the item $x \in \{0, 1\}$.

Strategy defines the bid $b_i(p, x)$ that agent i will take for every state (p, x):

$$b_{\mathrm{BR}}(p, x) = \begin{cases} p & \text{, if } x = 0 \text{ and } p < v_i \\ \text{no bid} & \text{, otherwise.} \end{cases}$$

Example 1

[Prisoner's Dilemma]

Two people are arrested for a crime. If a suspect testifies, and the other does not testify (DC), then released and receives a reward. If neither testifies (CC), both released. If both testify (DD), go to prison and collect reward.

	C	D
C	1,1	-1,2
D	2,-1	0,0

Def. Strategy profile s^* is a (weak) dominant-strategy equilibrium of a game if, for all i,

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}), \qquad \forall s_i \in S_i, \forall s_{-i} \in S_{-i}$$

Example: SPSB

Agents I, values v_i. Bidding strategy $b_i(v_i) \in [0, \infty)$.

$$u_i(b_i, b_{-i}, v_i) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{, if } b_i > \max_{j \neq i} b_j \\ 0 & \text{, otherwise} \end{cases}$$

Given value v_i, strategy $b_i^*(v_i) = v_i$ is a (weakly) dominant strategy, for all b_{-i}.

Let $b' = \max_{j \neq i} b_j$. If $b' < v_i$ then any bid $b_i(v_i) > b'$ is optimal. If $b' \geq v_i$, then any bid $b_i(v_i) \leq b'$ is optimal. Bid $b_i(v_i) = v_i$ solves both cases.

Example 2

	L	M	R
U	4,3	5,1	6,2
M	2,1	8,4	3,6
D	3,0	9,6	2,8

Example 3

	L	M	R
U	0,4	4,0	5,3
M	4,0	0,4	5,3
B	3,5	3,5	6,6

Nash Equilibrium

Def. A pure-strategy profile s^* is a *Nash equilibrium* if

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*), \qquad \forall s_i \in S_i, \forall i$$

Iterated Elimination of Dominated Strategies

Maintain a removed set of strategies, $R_i \subseteq S_i$, for each agent i. Initially, $R_i = \emptyset$.

- Choose an agent, i, and a strictly dominated strategy, $s_i \in (S_i \setminus R_i)$; i.e., such that some $s'_i \in (S_i \setminus R_i)$ satisfies:

$$u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i}), \quad \forall s_{-i} \in S_{-i} \setminus R_{-i}$$

- Add s_i to "remove set", R_i.
- Continue.

Thm. If a unique profile, s^*, survives, then it is the unique Nash equilibrium of the game.

Thm. If a profile, s^*, is a Nash eq., then it *must* survive iterated elimination.

Best-response Correspondences

[Brute force Nash eq. computation!]

First, apply iterated elimination of strictly dominated strategies, to get a candidate set $C_i \subseteq S_i$ of strategies for each agent.

Def. The best-response correspondence, $\mathrm{br}_i(s_{-i}) \subseteq S_i$, for agent i, computes the set of utility-maximizing strategies, given strategies s_{-i} from other agents.

Compute the BR correspondence for every agent, and search for the *fixed point*, s^* such that $s^* \in \mathrm{br}(s^*)$.

Example: Simple Competition Model

Consider a "Cournot" competition model. Two suppliers, produce quantity q_1, q_2 of a homogeneous good. Unit price is $p(Q) = [a - Q]^+$, for $Q = q_1 + q_2$; unit cost $c \geq 0$.

Utility:

$$u_1(q_1, q_2) = q_1(a - (q_1 + q_2)) - q_1 c$$

Best-response:

$$q_1^*(q_2) = \max_{0 \leq q_1 \leq \infty} q_1(a - (q_1 + q_2) - c)$$

Compute first order conditions, and solve.

Example 4

[Matching Pennies]

Each player has a penny and must choose whether to display it with heads facing up or down. If the pennies match, agent 1 wins, if the pennies do not match, agent 2 wins.

	H	T
H	1,-1	-1,1
T	-1,1	1,-1

Nash Equilibrium: Mixed

Def. A mixed strategy, $\sigma_i \in \Sigma_i$, defines a prob., $\sigma(s_i)$, for each strategy $s_i \in S_i$.

Exp. utility $u_i(\sigma) = \sum_{s \in S} (\Pi_j \sigma_j(s_j))\, u_i(s)$

Def. The support is strategies $\{s_i : \sigma(s_i) > 0\}$.

Def. A mixed-strategy profile σ^* is a *Nash equilibrium* if

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*), \qquad \forall \sigma_i \in \Sigma_i, \forall i$$

Thm. (Nash 1950) Every finite strategic-form game has a mixed-strategy equilibrium.

Note. All strategies in the *support* have the same exp. util. in any Nash eq.

Example: Matching Pennies

	H	T
H	1,-1	-1,1
T	-1,1	1,-1

Let $(r, 1-r)$ denote the mixed strategy of player 1, and $(q, 1-q)$ denote the mixed strategy of player 2. Compute BR correspondences $r^*(q)$ and $q^*(r)$.

$$r(q) = r[q-(1-q)] + (1-r)[-q+(1-q)]$$
$$= r[2q-1] + (1-r)[1-2q] = (1-2q) + r(4q-2)$$

Consider $q < 1/2, q = 1/2, q > 1/2$.

$$r^*(q) = \begin{cases} 0 & \text{, if } q < 1/2 \\ [0,1] & \text{, if } q = 1/2 \\ 1 & \text{, if } q > 1/2 \end{cases}$$

Construct $q^*(r)$ by symmetry, and solve set. of eqns.

Solution Method: Graphical

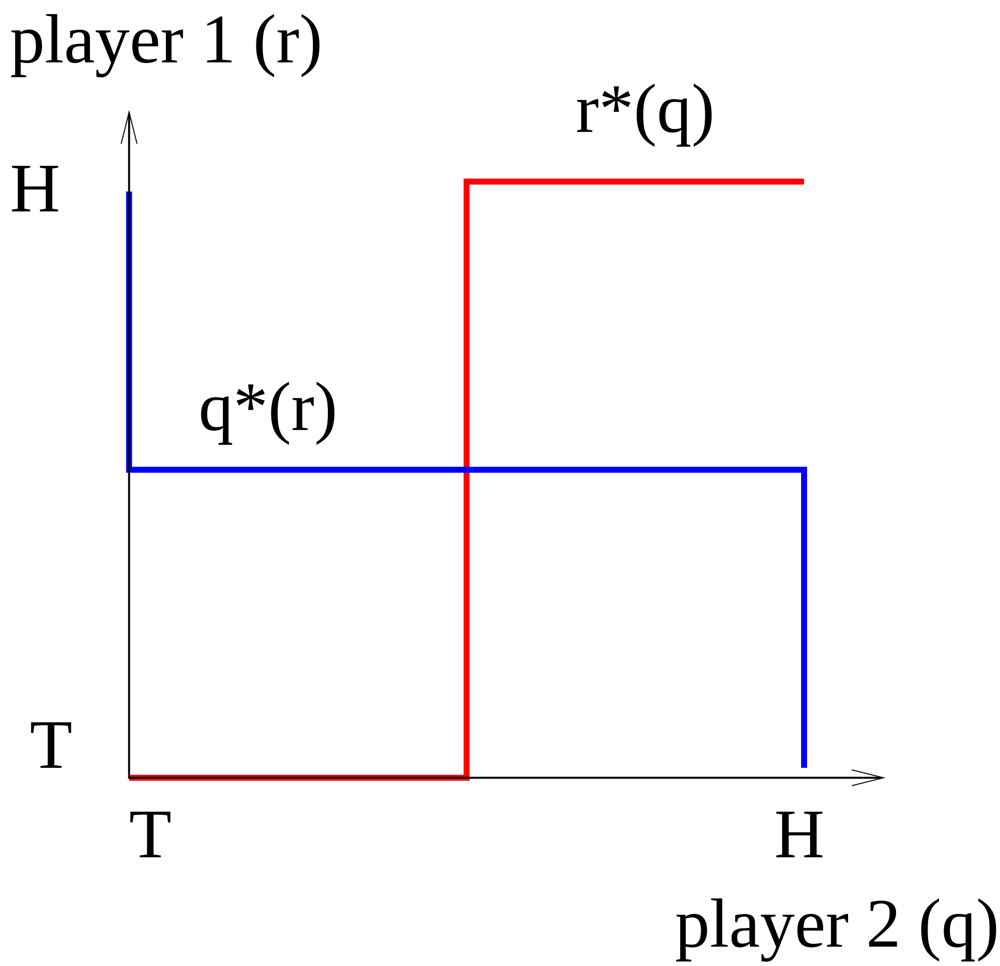

Computational Problems

Computing Nash eq. is a fundamental comput. problem, "complexity wide open". Suppose that $N = 2$ and S_1 and S_2 are finite sets. *Is there a poly. time algorithm for computing a (mixed) Nash eq. in such a game?*

"Together with factoring, the complexity of finding a Nash eq. is in my opinion the most important concrete open question on the boundary of P today." (Papadimitriou)

Specific NE problems (e.g. with maximal payoff, containing a particular subset of pure strategies) are intractable, existence of NE with expected social welfare $\geq k$ in symmetric two-player, etc. (Gilboa & Zemel 89; Conitzer & Sandholm 02).

Polynomial-time Algorithms

- two player zero-sum games (Murty'88); including algorithms that are polynomial in size of extensive-form (Koller & Megiddo 91)
- approximate every NE with a graph-theoretic representation in polynomial time in size of representation and quality, when underlying graph is a tree (Kearns et al. 01)
  – computes **all** NE, so can be used to find a NE with a useful property
  – an algorithm to compute *all* NE exactly that is exponential in # agents.
- Polynomial time algorithm to compute a **single** NE exactly for a graph with a tree-structure (Littman et al. 02)
- polynomial time algorithm to compute approximate NE in a setting with "bounded-influence" between agents (Kearns & Mansour 02)

Software Solutions

- Gambit. www.hss.caltech.edu/gambit
  – very limited, uses a normal-form representation
  – computes all Nash equilibrium
  – need to eliminate dominated strategies first
- Gala. (Koller & Pfeffer 97)
  – provides a programming language to represent games; solves extensive form without reducing to Normal form.
- Non-linear optimization; via a non-linear complementarity formulation (see McKelvey & McLennan 96)
- replicator dynamics (see Weibull 95)

Other Issues

- Multiple Nash equilibria
  – *focal points* (Schelling 60)
  – e.g. names, past experiences, Pareto-dominance, etc.
  – some help from center??
- Learning dynamics
  – adjustments, convergence, etc.
  – evolutionary game theory
- Common knowledge (knowledge “ad infinitum”)
  – of payoffs, of rationality, of beliefs, etc.
- Robustness
  – to perturbations, uncertain information, irrational agents, etc.

Bayesian-Nash equilibrium

(Harsanyi 68)

For **imperfect information** games, in which agents are uncertain of the payoffs of other agents.

Agent i has *type*, $\theta_i \in \Theta_i$, which defines payoff $u_i(s, \theta_i)$.

Agents have *common prior*, $p(\theta)$, over distr. of agent types; write the conditional prob. $p(\theta_{-i}|\theta_i)$.

Strategy spaces, payoff functions, possible types, prior distributions are *common knowledge*.

Let $\sigma_i(\theta_i) \in \Sigma_i$ denote the strategy that agent i chooses, given type θ_i; then $\sigma = (\sigma_1, \ldots, \sigma_N)$ denotes a strategy profile.

Expected utility:

$$u_i(\sigma_i(\theta_i), \sigma_{-i}(\cdot), \theta_i) = \sum_{\theta_{-i}} p(\theta_{-i}|\theta_i) u_i(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i}), \theta_i)$$

Def. Strategy-profile, $\sigma^*(\cdot)$, is in Bayesian-Nash eq. if, for every agent i, and every $\theta_i \in \Theta_i$,

$$u_i(\sigma_i^*(\theta_i), \sigma_{-i}^*(\cdot), \theta_i) \geq u_i(\sigma_i(\theta_i), \sigma_{-i}^*(\cdot), \theta_i)$$
$$\forall \sigma_i(\cdot) \neq \sigma_i^*(\cdot)$$

Each agent plays a best-response with respect to its beliefs about the types of the other agents, assuming they also play a best-response with respect to their beliefs.

Example: FPSB

Bidders 1, 2. Bidder i has value v_i. Payoff is $v_i - p$. Values $v_i \sim [0,1]$. Strategy space $S_i = [0, \infty)$. Strategy $b_i(v_i) \in S_i$, specifies a bid for each value. Payoff:

$$u_i(b_1, b_2, v_i) = \begin{cases} v_i - b_i & \text{, if } b_i > b_j \\ (v_i - b_i)/2 & \text{, if } b_i = b_j \\ 0 & \text{if } b_i < b_j \end{cases}$$

Strategies $(b_1^*(\cdot), b_2^*(\cdot))$ define a Bayesian-Nash eq. if $b_i(v_i)$ solves:

$$\max_{b_i} \left[(v_i - b_i)\mathrm{Pr}(b_i > b_j(v_j)) + 1/2(v_i - b_i)\mathrm{Pr}(b_i = b_j(v_j))\right]$$

for each $v_i \in [0,1]$.

Simplify, assume a *linear equilibrium*, with $b_i(v_i) = a_i + c_i v_i$. Solve:

$$b_i^*(v_i, a_j, c_j) = \max_{b_i}(v_i - b_i)\mathrm{Pr}(b_i > a_j + c_j v_j)$$

Clearly, $a_j \leq b_i^*(v_i, a_j, c_j) \leq a_j + c_j$, and $\mathrm{Pr}(b_i > a_j + c_j v_j) = (b_i - a_j)/c_j$.

Bayes-Nash: Comput. Issues

- existence of multiple equilibria
- common prior
  – can relax by folding beliefs about knowledge of other agents into the type space] (Mertens & Zamir 85)
- rationality assumptions
- common-knowledge assumptions
- no efficient algorithm.

Dominant strategy equilibria much more desirable!

What did we not cover?

- **Extensive-form games**
  - – multiple-stages, dynamic structure
  - – explicit order, information made explicit
  - – Subgame-perfect Nash equil.
- **Repeated games**
  - – strategic-form (or “stage game”) in each round
  - – weighted average of payoffs in each stage
  - – “Folk theorems”
- **Equilibrium refinements**
  - – perfect equil., perfect Bayesian equil, etc.