Implementation Theory and Bargaining Problems

Based on Chapters 6 (Extensive Games with Perfect Information), 7 (Bargaining Games), 10 (Implementation Theory), and 15 (The Nash Solution) in Osborne and Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.

In this talk:

- Introduction to implementation theory: basic definitions, Nash implementation, and subgame perfect equilibrium implementation.
- Bargaining problems and their implementation.

Extensive games

Let's first recall some concepts from extensive games. Here, an extensive game will always be an extensive game with perfect information.

Definition 1. An *extensive game (with perfect information)* $\langle N, H, P, (\succsim_i) \rangle$ consists of

- A nonempty finite set N (the *players*).
- A set H of (finite or infinite) sequences satisfying the following properties.
  - The empty sequence $\varepsilon \in H$.
  - If $(a^k) \in H$, every proper prefix of (a^k) belongs to H.
  - If an infinite sequence $(a^k)_{k=1}^{\infty}$ satisfies $(a^k)_{k=1}^{L} \in H$ for each $L \in \mathbb{Z}^+$ then $(a^k)_{k=1}^{\infty} \in H$.

Each member of H is a *history*; each component of a history is an *action* taken by a player. A history $(a^k) \in H$ is *terminal* if it is infinite, or if it is finite and there is no a such that $(a^k, a) \in H$. The set of terminal histories is denoted by Z.

- A player function $P\colon H \setminus Z \to N$ ($P(h)$ is the player who takes an action after the history h).
- For each player $i \in N$, a preference relation $\succsim_i$ on Z.

If each member of H is finite, the game is said to have *finite horizon*.

After each nonterminal history h, player $P(h)$ chooses an action from the set

$$A(h) = \{a \mid (h, a) \in H\}.$$

There is a straightforward generalization that allows *chance moves*.

Strategies

Definition 2. A *strategy* of player $i \in N$ in an extensive game $\langle N, H, P, (\succsim_i) \rangle$ is a function that assign to each $h \in H \setminus Z$ for which $P(h) = i$ an action in $A(h)$.

For each strategy profile $s = (s_i)_{i \in N}$, the *outcome* $O(s)$ of s is the terminal history that results when each player follows its strategy s_i.

Definition 3. A *Nash equilibrium* of the extensive game $\langle N, H, P, (\succsim_i) \rangle$ is a strategy profile s^* such that for each player $i \in N$,

$$O(s^*_{-i}, s^*_i) \succsim_i O(s^*_{-i}, s_i)$$

for every strategy s_i of i.

Subgame perfect equilibrium

Definition 4. The *subgame* of the extensive game $\Gamma = \langle N, H, P, (\succsim_i) \rangle$ *that follows* h is the extensive game $\Gamma(h) = \langle N, H|_h, P|_h, (\succsim_i|_h) \rangle$, where

- $H|_h = \{h' \mid (h, h') \in H\}$,
- $P|_h(h') = P(h, h')$ for all $h' \in H|_h$, and
- $h' \succsim_i|_h h'' \iff (h, h') \succsim_i (h, h'')$.

A strategy s_i for Γ induces a strategy $s_i|_h$ for $\Gamma(h)$: $s_i|_h(h') = s_i(h, h')$. Finally, we let O_h denote the outcome function of $\Gamma(h)$.

Definition 5. A *subgame perfect equilibrium* of an extensive game $\Gamma = \langle N, H, P, (\succsim_i) \rangle$ is a strategy profile s^* such that for each player $i \in N$ and every $h \in H \setminus Z$ for which $P(h) = i$, we have

$$O_h(s^*_{-i}|_h, s^*_i|_h) \succsim_i|_h O_h(s^*_{-i}|_h, s_i)$$

for every strategy s_i of player i in the subgame $\Gamma(h)$.

That is, $s^*|_h$ is a Nash equilibrium of the subgame $\Gamma(h)$ for each $h \in H \setminus Z$.

The following result will be needed later.

Lemma 6 (The one deviation property). Let Γ be a finite horizon extensive game with or without chance moves. The strategy profile s^* is a subgame perfect equilibrium of Γ if and only if for every player $i \in N$ and history $h \in H$ for which $P(h) = i$ we have

$$O_h(s^*_{-i}|_h, s^*_i|_h) \succsim_i|_h O_h(s^*_{-i}|_h, s_i)$$

for every strategy s_i of player i in $\Gamma(h)$ that differs from $s^*_i|_h$ only in the action it prescribes after the initial history of $\Gamma(h)$.

Implementation theory

In implementation theory, we fix a set of outcomes and look for a game that yields that set of outcomes as equilibria.

The model we consider is the following. A *planner* starts with a description of the outcomes she wants to associate with each possible preference profile, and looks for a game that “implements” this correspondence.

As an example, consider a planner that wants to assign an object to one of two individuals. Assume that she wishes to give the object to the individual that values it the most, but she doesn’t know which one this is.

Her problem is then to design a game form such that for each pair of valuations, the outcome according to some solution concept is that the object is given to the individual who value it the most.

Implementation theory more formally

Definition 7. Let N be a set of individuals, C a set of feasible *outcomes*, and $\mathcal{P}$ a set of preference profiles over C. A *choice rule* is a function that assigns a subset of C to each profile in $\mathcal{P}$. A singleton-valued choice rule is called a *choice function*.

Definition 8. A *strategic game form with consequences in* C is a triple $\langle N, (A_i), g \rangle$, where A_i is the set of actions available to player $i \in N$, and $g\colon \prod_{i \in N} A_i \to C$ is an *outcome function* that associate an outcome with each action profile. A strategic game form $\langle N, (A_i), g \rangle$ and a preference profile $(\succsim_i) \in \mathcal{P}$ induce a strategic game $\langle N, (A_i), (\succsim_i') \rangle$, where $a \succsim_i' b \iff g(a) \succsim_i g(b)$.

Similarly, an *extensive game form with consequences in* C is a tuple $\langle N, H, P, g \rangle$, where H is the set of histories, $P\colon H \setminus Z \to N$ is the player function, and $g\colon Z \to C$ is an outcome function (Z is the set of terminal histories). An extensive game form and a preference profile induce an extensive game.

Definition 9. An *environment* $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ consists of

- A finite set N of players, with $|N| \geq 2$.
- A set C of outcomes.
- A set $\mathcal{P}$ of preference profiles over C.
- A set $\mathcal{G}$ of (strategic or extensive) game forms with consequences in C.

A *solution concept* for the environment $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ is a set valued function $\mathcal{S}$ with domain $\mathcal{G} \times \mathcal{P}$. If the members of $\mathcal{G}$ are strategic game forms, $\mathcal{S}$ takes values in the set of action profiles. If the members of $\mathcal{G}$ are extensive game forms, $\mathcal{S}$ takes values on the set of terminal histories.

Definition 10. Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment, and let $\mathcal{S}$ be a solution concept. The game form $G \in \mathcal{G}$ with outcome function g is said to $\mathcal{S}$-*implement* the choice rule $f : \mathcal{P} \to C$, if for each preference profile $(\succsim_i) \in \mathcal{P}$, we have $g(\mathcal{S}(G, (\succsim_i))) = f((\succsim_i))$.

Definition 11. Let $\langle N, C, \mathcal{P}, \mathcal{G}\rangle$ be an environment in which $\mathcal{G}$ is a set of strategic game forms for which the set of actions of each player $i \in N$ is a set $\mathcal{P}$ of preference profiles. Let $\mathcal{S}$ be a solution concept. The strategic game form $G = \langle N, (A_i), g\rangle \in \mathcal{G}$ is said to *truthfully $\mathcal{S}$-implement* the choice rule $f : \mathcal{P} \to C$, if for each preference profile $(\succsim_i) \in \mathcal{P}$, we have

- $a^* \in \mathcal{S}(G, (\succsim_i))$, where $a_i^* = (\succsim_i)$ for each $i \in N$ (every player reporting the true preference profile is a solution).
- $g(a^*) \in f((\succsim_i))$ (if every player reports the true preference profile, the outcome is a member of $f((\succsim_i))$).

This notation of implementation differs in in several ways form the “normal” implementation concept:

- The set of actions of each player is a set of preference profiles, and “truth telling” is always a solution.
- Non-truth telling solutions may yield outcomes that are inconsistent with the choice rule.
- There can be preference profiles for which not every outcome prescribed by the choice rule corresponds to a solution of the game.

Nash implementation

We now consider the case where the planner uses strategic game forms, and for each preference profile, the outcome of the game may be in any of its Nash equilibria.

The first result is a version of the relevation principle. The result shows that for any Nash-implementable choice rule, there is a game form in which

1. Each player has to announce a preference profile.
2. For any preference profile, truth telling is a Nash equilibrium.

The precise statement is as follows.

Proposition 12 (Relevation principle for Nash implementation). Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment in which $\mathcal{G}$ is a set of strategic game forms. If a choice rule is Nash-implementable in the environment, it is truthfully Nash-implementable.

Proof. Let $G = \langle N, (A_i), g \rangle$ be a game form that implements the choice rule $f : \mathcal{P} \to C$, and for each $\succsim \in \mathcal{P}$, let $(a_i(\succsim))$ be a Nash equilibrium of the game $\langle G, \succsim \rangle$.

Let $G^* = \langle N, (A_i^*), g^* \rangle$, where $A_i^* = \mathcal{P}$ for each $i \in N$ and $g^*(p) = g((a_i(p_i)))$ for each $p \in \prod_{i \in N} A_i^*$ (Note that each p_i is a preference profile and that p is a profile of profiles).

The profile p^* such that $p_i^* = \succsim$ for each $i \in N$ is clearly a Nash equilibrium of $\langle G^*, \succsim \rangle$, and $g^*(p^*) \in f(\succsim)$. □

The following result gives necessary conditions for a choice rule to be Nash implementable.

Definition 13. A choice rule $f : \mathcal{P} \to C$ is *monotonic* if whenever $c \in f((\succsim_i))$ and $c \notin f((\succsim'_i))$, there is a player $i \in N$ and some outcome $b \in C$ such that $c \succsim_i b$ and $b \succ'_i c$.

Proposition 14. Let $\langle N, C, \mathcal{P}, \mathcal{G} \rangle$ be an environment in which $\mathcal{G}$ is a set of strategic game forms. If a choice rule is Nash-implementable in the environment, it is monotonic.

Proof. Suppose that the choice rule $f : \mathcal{P} \to C$ is Nash implemented by the game form $G = \langle N, (A_i), g \rangle$, $c \in f((\succsim_i))$, and $c \notin f((\succsim'_i))$. Then there is an action profile a such that $g(a) = c$, and a is a Nash equilibrium of the game $\langle G, (\succsim_i) \rangle$, but not of $\langle G, (\succsim'_i) \rangle$. Thus, there is a player j and action $a'_j \in A_j$, such that $g(a_{-j}, a'_j) \succ'_j g(a)$ and $g(a) \succsim_j g(a_{-j}, a'_j)$. □

Example 15 (Solomon's predicament). This is a classical example based on some biblical story.

Each of two women, 1 and 2, claims a baby. Each of them knows who is the true mother, but neither can prove her motherhood. Solomon tries to find the true mother by threatening to cut the baby in two relying on the fact (?) that the true mother prefers to give the baby away to see it cut in two, while the false mother rather sees the baby cut in two than gives the baby to the true mother. Solomon can give the baby to either mother, or order its execution.

Formally, let a be the outcome that the baby is given to mother 1, b that it is given to 2, and d that it is cut in two. There are two possible preference profiles:

$$\theta : a \succ_1 b \succ_1 d \text{ and } b \succ_2 d \succ_2 a \quad \text{[1 is the real mother]}$$

$$\theta' : a \succ_1' d \succ_1' b \text{ and } b \succ_2' a \succ_2' d \quad \text{[2 is the real mother]}$$

The choice rule f defined by $f(\theta) = \{a\}$ and $f(\theta') = \{b\}$ is not Nash implementable, since it is not monotonic: $a \in f(\theta)$ and $a \notin f(\theta')$, but there is no player i and outcome y such that $a \succsim_i y$ and $y \succ_i' a$.

Obviously, Solomon (or the women) didn't participate in game theory seminars.

Subgame perfect equilibrium implementation

Next, we will consider the case where the planner uses extensive game forms, and for each preference profile, the outcome of the game may be in any subgame perfect equilibria (SPE). We will restrict ourself to an illustrative example.

Example 16. The planner wants to divide an object of monetary value between two players, 1 and 2. One of the players is the legitim owner of the object, but the planner does not know which one. Suppose that the planner can give the object to any of the players, or neither of them, and that she also may impose fines on the players.

The set of outcomes is the set of triples (x, m_1, m_2), where $x = 0$ (neither player gets the object) or $x \in \{1, 2\}$ (player x gets the object), and m_i is a fine imposed on player i.

Player i's payoff if he gets the object is $v_H - m_i$ if we is the legitim owner of the object, and $v_L - m_i$ if he is not, where $v_H > v_L > 0$. If player i does not get the object, his payoff is $-m_i$.

There are two possible preference profiles, $\succsim$ in which player 1 is the legitim owner, and $\succsim'$ in which player 2 is.

The planner wants to implement the choice rule f for which $f(\succsim) = (1,0,0)$ and $f(\succsim') = (2,0,0)$.

This is implemented by the following extensive game form.

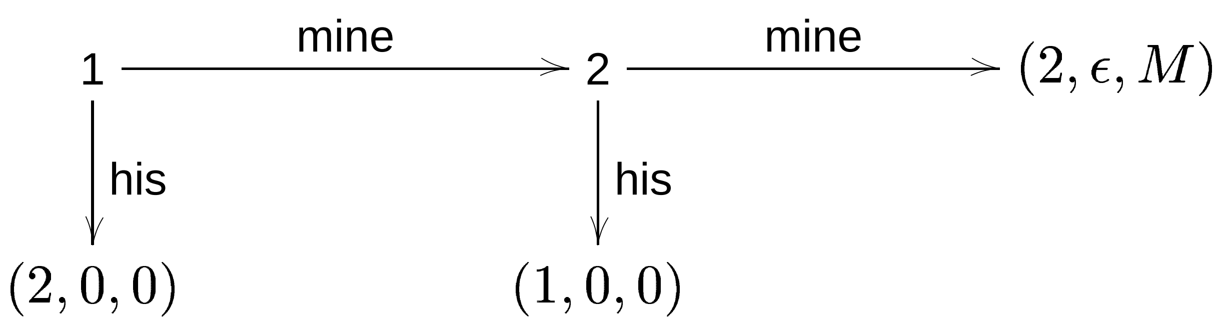

First player 1 is asked whether the object is his. If he says, “no”, the object is given to player 2. If he says “yes”, player 2 is asked if he is the owner. If player 2 answers “no”, the object is given to player 1. Otherwise, player 2 gets the object and he must pay a fine M, $v_L < M < v_H$ while player 1 has to pay a small fine $\epsilon > 0$.

It is easy to see that for each preference profile, the game has a unique SPE with outcome $(i, 0, 0)$, where i is the legitim owner. Thus, this game form SPE-implements the choice rule f.

The idea behind the game form is that in each SPE, player 2 is forced to choose truthfully. Given that player 2 always chooses truthfully, player 1 is also forced to choose truthfully in each SPE.

Bargaining problems revisited

For $p \in [0,1] \subseteq \mathbb{R}$, we let $p \cdot x \oplus (1-p) \cdot y$ denote the (discrete) probability distribution that gives x with probability p and y with probability $1-p$. Furthermore, we let $p \cdot x$ denote the distribution $p \cdot x \oplus (1-p) \cdot D$ (D is defined below).

Definition 17. A *bargaining problem* $\langle X, D, \succsim_1, \succsim_2 \rangle$ consists of

- A compact set X in a metric space (the set of *agreements*).
- An element $D \in X$ (the *disagreement* outcome).
- Two preference relations $\succsim_1, \succsim_2$ on the set of probability distributions over X satisfying $x \succsim_i D$ for all $x \in X$. The preference relations are represented by continuous utility functions $u_i\colon X \to [0,\infty) \subseteq \mathbb{R}$ such that $u_i(D) = 0$ and $x \succsim_i y \iff E[u_i(x)] \geq E[u_i(y)]$.

- The problem is non-degenerate in the sense that there is an $x \in X$ such that $x \succ_1 D$ and $x \succ_2 D$.
- (Convexity). For any $x, y \in X$ and $p \in [0, 1]$, there is an $z \in X$ such that $z \sim_i p \cdot x \oplus (1 - p) \cdot y$ for $i = 1, 2$.
- (Non-redundancy). If $x \in X$, there is no $x' \in X, x' \neq x$ such that $x \sim_i x'$ for $i = 1, 2$.

The Nash solution

Definition 18. A *bargaining solution* is a function that assigns to every bargaining problem $\langle X, D, \succsim_1, \succsim_2 \rangle$ a unique element in X.

Definition 19. The *Nash solution* is a bargaining solution that assigns to the bargaining problem $\langle X, D, \succsim_1, \succsim_2 \rangle$ an $x^* \in X$ such that

$$p \cdot x \succ_i x^*, p \in [0,1], x \in X \Longrightarrow p \cdot x^* \succsim_j x, j \neq i.$$

Proposition 20. Let $\langle X, D, \succsim_1, \succsim_2 \rangle$ be a bargaining problem. Then $x^* \in X$ is a Nash solution of the problem if and only if

$$u_1(x^*)u_2(x^*) \geq u_1(x)u_2(x), \forall x \in X.$$

Furthermore, the Nash solution is well-defined.

Proof. Suppose first that $u_1(x^*)u_2(x^*) \geq u_1(x)u_2(x)$ for all $x \in X$. Then $u_i(x^*) > 0$ for $i = 1, 2$ (since the problem is non-degenerate). If $pu_i(x) > u_i(x^*)$ for some $p \in [0, 1]$ and $x \in X$, then $pu_i(x)u_j(x^*) > u_i(x^*)u_j(x^*) \geq u_i(x)u_j(x)$. Thus, $pu_j(x^*) > u_j(x)$. That is, $p \cdot x \succ_i x^* \implies p \cdot x^* \succsim_j x$.

Conversely, suppose that x^* is a Nash solution. By definition, we must have $u_i(x^*) > 0$ for $i = 1, 2$. Let $x \in X$ be such that $u_i(x) > 0$ for $i = 1, 2$, and $u_i(x) > u_i(x^*)$ for some i (for all other values of x, we obviously have $u_1(x^*)u_2(x^*) \geq u_1(x)u_2(x)$). If $p > u_i(x^*)/u_i(x)$ for some $p \in [0, 1]$, we have $pu_j(x^*) \geq u_j(x)$ (since x^* is a Nash solution). Hence $u_i(x^*)u_j(x^*)/u_i(x) \geq u_j(x)$ and thus $u_i(x^*)u_j(x^*) \geq u_i(x)u_j(x)$.

Finally, to show that the Nash solution is well-defined, let $U = \{(u_1(x), u_2(x)) \mid x \in X\}$. Note that x^* is a Nash solution if and only if $(v_1, v_2) = (u_1(x^*), u_2(x^*))$ maximizes $v_1 v_2$ over U. Since U is compact (u_1, u_2 are continuous), this problem has a solution. Since the function $v_1 v_2$ is strictly quasi-concave on the interior of $\mathbb{R}^2_+$ and U is convex, the solution is unique. Finally, by the non-redundancy there is a unique $x^* \in X$ that yields the pair of maximizing utilities. $\square$

Implementation of the Nash solution

Proposition 21. Fix a set X and an event $D \in X$. For all pairs $(\succsim_1, \succsim_2)$ for which $\langle X, D, \succsim_1, \succsim_2 \rangle$ is a bargaining problem, the following extensive game form (with perfect information and chance moves) SPE-implements the Nash solution.

1. Player 1 chooses $y \in X$.
2. Player 2 chooses $x \in X$ and $p \in [0, 1]$.
3. With probability $1 - p$ the game ends with outcome D. With probability p it continues.
4. Player 1 chooses either x or $p \cdot y$. This choice is the outcome.

Proof. Let x^* be the (unique) Nash solution of the bargaining problem. We claim that each SPE of the game is essentially equivalent to the following simple strategy profile.

- Step 1: Player 1 chooses $y = x^*$.
- Step 2: Player 2 chooses $x = x^*$ and $p = \inf\{p \mid u_1(x^*) \geq pu_1(y)\}$.
- Step 4: Player 1 chooses $\max(u_1(x), pu_1(y))$.

Using the one deviation property, we can easily show that this is an SPE of the game. Clearly, its outcome is x^*.

It remains to show that the SPE is unique. To this end, consider any SPE of the game.

In the last step, Player 1 is clearly forced to choose (with some abuse of notation)

$$\max(u_1(x), pu_1(y)).$$

In step 2, Player 2 is forced to choose x and p such that

$$pu_2(\max(u_1(x), pu_1(y)))$$

is maximized. In step 1, Player 1 is forced to choose y such that

$$u_1(\max_{x,p} pu_2(\max(u_1(x), pu_1(y))))$$

is maximized.

Thus, the SPE is unique ignoring ties in the $\max$ operations. By the restrictions put on u_1 and u_2, there can be no relevant ties. □

Concluding remarks

From a cryptographic point of view, the previous implementation is unsatisfactory, since it gives no method to compute the SPE unless both parties know the other party's utility function. If both parties know u_1 and u_2, they can independently compute $x^* \in X$ such that $u_1(x^*)u_2(x^*)$ is maximized without any fancy game.

Thus, it is still an open problem to designs a protocol that allows the two parties to find the value x^* without giving any non-trivial information about u_i to party j. (This is where I ran out of time . . .)

Or course, this could be done using secure function evaluation, but is this efficient enough?