Introduction to Games of Incomplete Information

George J. Mailath

University of Pennsylvania
and
Australian National University

February 2019

Games of Complete Information

- All the previous examples are games of complete information:
    - All players know the nature of other players' information and
    - beliefs are known (being determined by the probability distribution of nature's moves).

An Example without Complete Information

	L	R
T	$1, 1$	$0, 0$
B	$0, 1$	$1, 0$

	L	R
T	$1, 0$	$0, 1$
B	$0, 0$	$1, 1$

- Suppose *I* does not which of the payoff matrices is true; *II* knows matrix, so both players know their own payoff.
- *II* plays *L* in left matrix, and *R* in right matrix.
- How will *I* choose in the presence of uncertainty?

Decision Making with Randomness

- How does a decision maker choose when faced with randomness?
  - Risk: choices have random consequences, but probability distribution of these consequences is known (lottery tickets, roulette wheel).
  - Uncertainty: choices have random consequences, and the probability distribution of these consequences is unknown (horse race). (Is probability even defined in these circumstances?)

Decision Making Under Risk

- Von Neumann and Morgenstern (1944): Primitive is a preference order $\succeq$ over lotteries (i.e., a specified objective probability distribution) over outcomes X.

Decision Making Under Risk

- Von Neumann and Morgenstern (1944): Primitive is a preference order $\succeq$ over lotteries (i.e., a specified objective probability distribution) over outcomes X.
- If the preference order $\succeq$ satisfies some “reasonable rationality” axioms, then there is a utility function $u : X \to \mathbb{R}$ such that

$$p \succeq q \iff \sum_x p(x)u(x) \geq \sum_x q(x)u(x).$$

Decision Making Under Risk

- Von Neumann and Morgenstern (1944): Primitive is a preference order $\succeq$ over lotteries (i.e., a specified objective probability distribution) over outcomes X.
- If the preference order $\succeq$ satisfies some "reasonable rationality" axioms, then there is a utility function $u : X \to \mathbb{R}$ such that

$$p \succeq q \iff \sum_x p(x)u(x) \geq \sum_x q(x)u(x).$$

- Moreover, the utility function is unique up to positive affine transformations: If v is another utility function also representing the preference order, then there exists two constants, $a > 0$ and b such that for all $x \in X$,

$$v(x) = au(x) + b.$$

Decision Making Under Uncertainty

- Savage (1972): Primitives are
  - a state space Ω, where a state resolves all uncertainty,
  - acts (bets), $f : \Omega \to X$, where X is (as in vNM) the space of outcomes, and
  - a preference order $\succeq$ over the space of acts.

Decision Making Under Uncertainty

- Savage (1972): Primitives are
  - a state space Ω, where a state resolves all uncertainty,
  - acts (bets), $f : \Omega \to X$, where X is (as in vNM) the space of outcomes, and
  - a preference order $\succeq$ over the space of acts.
- If the preference order $\succeq$ satisfies some "reasonable rationality" axioms, then there is a utility function $u : X \to \mathbb{R}$ and a subjective finitely-additive probability measure μ such that

$$f \succeq g \iff \int u(f(\omega))d\mu(\omega) \geq \int u(g(\omega))d\mu(\omega).$$

Decision Making Under Uncertainty

- Savage (1972): Primitives are
  - a state space Ω, where a state resolves all uncertainty,
  - acts (bets), $f : \Omega \to X$, where X is (as in vNM) the space of outcomes, and
  - a preference order $\succeq$ over the space of acts.
- If the preference order $\succeq$ satisfies some "reasonable rationality" axioms, then there is a utility function $u : X \to \mathbb{R}$ and a subjective finitely-additive probability measure μ such that

$$f \succeq g \iff \int u(f(\omega))d\mu(\omega) \geq \int u(g(\omega))d\mu(\omega).$$

- Moreover, the utility function is unique up to positive affine transformations and μ is unique.

Decision Making Under Uncertainty

- Savage (with the extensions to accommodate countably additive beliefs and objective probabilities as well) is the standard (classical) approach to dealing with uncertainty.

Objective beliefs are determined by nature.
Subjective beliefs are not and reflect the agent's speculations or theories.

An Example without Complete Information

	L	R
T	$1,1$	$0,0$
B	$0,1$	$1,0$

	L	R
T	$1,0$	$0,1$
B	$0,0$	$1,1$

- Suppose *I* does not which of the payoff matrices is true; *II* knows matrix, so both players know their own payoff.
- *II* plays *L* in left matrix, and *R* in right matrix.
- *I* assigns probability α to left matrix. Then, *I* plays *T* if $\alpha > \frac{1}{2}$, and plays *B* if $\alpha < \frac{1}{2}$ (and is indifferent if $\alpha = \frac{1}{2}$).

An Example without Complete Information

	L	R
T	$1,1$	$0,0$
B	$0,1$	$1,0$

	L	R
T	$1,2$	$0,1$
B	$0,0$	$1,1$

- Should *II* still feel comfortable playing *R* if the payoffs are the right matrix?
- Optimality of *II*'s action choice of *R* depends on believing that *I* will play *B*, which only occurs if $\alpha \leq \frac{1}{2}$.
- But suppose *II* does not know *I*'s beliefs α. Then, *II* has beliefs over *I*'s beliefs and so *II* finds *R* optimal if he assigns probability at least $\frac{1}{2}$ to *I* assigning probability at least $\frac{1}{2}$ to the right matrix.
- But, how confident is *I* that *II* will play *R* in the right matrix?....

Games of Incomplete Information

Definition (Harsanyi (1967, 1968a,b))

A game of incomplete information or Bayesian game is the collection $\{(A_i, T_i, p_i, u_i)_{i=1}^n\}$, where

- A_i is i's action space,
- T_i is i's type space,
- $p_i : T_i \to \Delta\left(\prod_{j \neq i} T_j\right)$ is i's subjective beliefs about the other players' types, given i's type and
- $u_i : \prod_j A_j \times \prod_j T_j \to \mathbb{R}$ is i's payoff function.

- A player's type t_i describes everything that i knows that is not common knowledge (including player i 's beliefs).

Bayes-Nash Equilibrium I

A strategy for i is

$$s_i : T_i \to A_i.$$

Let $s(t) := (s_1(t_1), \dots, s_n(t_n))$, etc.

Definition

The profile $(\hat{s}_1, \dots, \hat{s}_n)$ is a Bayes-Nash (or Bayesian-Nash) equilibrium if, for all i and all $t_i \in T_i$,

$$E_{t_{-i}}[u_i(\hat{s}(t), t)] \geq E_{t_{-i}}[u_i(a_i, \hat{s}_{-i}(t_{-i}), t)], \quad \forall a_i \in A_i,$$

where the expectation over t_{-i} is taken with respect to $p_i(t_i)$.

Bayes-Nash Equilibrium II

If the type spaces are finite, then the probability i assigns to the vector $t_{-i} \in \prod_{j \neq i} T_j =: T_{-i}$ when his type is t_i can be denoted $p_i(t_{-i}; t_i)$, and the profile $(\hat{s}_1, \ldots, \hat{s}_n)$ is a Bayes-Nash (or Bayesian-Nash) equilibrium if, for all i and all $t_i \in T_i$,

$$\sum_{t_{-i}} u_i(\hat{s}(t), t) p_i(t_{-i}; t_i) \geq \sum_{t_{-i}} u_i(a_i, \hat{s}_{-i}(t_{-i})) p_i(t_{-i}; t_i), \quad \forall a_i \in A_i.$$

Return to Cournot duopoly example

- Firm 1's costs are private information, while firm 2's are public.
- Nature determines the costs of firm 1 at the beginning of the game, with $\Pr(c_1 = c_L) = \theta \in (0, 1)$.
- $A_i = \mathbb{R}_+$, firm 1's type space is $T_1 = \{t_1^L, t_1^H\}$, firm 2's is $T_2 = \{t_2\}$.
- Belief mapping p_1 for firm 1 is trivial: both types assign prob. 1 to t_2.
- The belief mapping for firm 2 is

$$p_2(t_2) = \theta \circ t_1^L + (1 - \theta) \circ t_1^H \in \Delta(T_1).$$

- Finally, payoffs are

$$u_1(q_1, q_2, t_1, t_2) = \begin{cases} [(a - q_1 - q_2) - c_L]q_1, & \text{if } t_1 = t_1^L, \\ [(a - q_1 - q_2) - c_H]q_1, & \text{if } t_1 = t_1^H, \end{cases}$$

$$u_2(q_1, q_2, t_1, t_2) = [(a - q_1 - q_2) - c_2]q_2.$$

Return to Cournot Duopoly Example with a twist

Firm 2 may know that firm 1 has low costs, c_L

- $T_1 = \{t_1^L, t_1^H\} = \{c_L, c_H\}$, $T_2 = \{t_2^l, t_2^U\} = \{t_l, t_U\}$. The prior distribution is

$$\Pr(t_1, t_2) = \begin{cases} 1 - p' - p'', & \text{if } (t_1, t_2) = (t_1^L, t_2^l), \\ p', & \text{if } (t_1, t_2) = (t_1^L, t_2^U), \\ p'', & \text{if } (t_1, t_2) = (t_1^H, t_2^U). \end{cases}$$

- The belief mappings are

$$p_1(t_1) = \begin{cases} (1-\alpha) \circ t_2^l + \alpha \circ t_2^U, & t_1 = t_1^L, \\ 1 \circ t_2^U, & t_1 = t_1^H, \end{cases}$$

$$p_2(t_2) = \begin{cases} 1 \circ t_1^L, & t_2 = t_2^l, \\ \theta \circ t_1^L + (1-\theta) \circ t_1^H, & t_2 = t_2^U. \end{cases}$$

Interim perspective and the role of priors

- The perspective of a game of incomplete information is interim: the beliefs of player i are specified type by type.

Interim perspective and the role of priors

- The perspective of a game of incomplete information is interim: the beliefs of player i are specified type by type.
- Suppose the type spaces are finite or countably infinite. Let $\hat{q}_i$ be an arbitrary full support distribution on T_i, i.e., $\hat{q}_i \in \Delta(T_i)$. Then defining

$$q_i(t) := \hat{q}_i(t_i) p_i(t_{-i}; t_i) \qquad \forall t$$

  generates a prior $q_i \in \Delta(T)$ for player i with the property that $p_i(\cdot; t_i) \in \Delta(T_{-i})$ is the belief on t_{-i} conditional on t_i.
- There are many priors consistent with the subjective beliefs (since $\hat{q}_i$ is arbitrary).

Common Prior Assumption

Definition

The subjective beliefs are consistent or satisfy the Common Prior Assumption (CPA) if there exists a single probability distribution $p \in \Delta\left(\prod_i T_i\right)$ such that, for each i, $p_i(t_i)$ is the probability distribution on T_{-i} conditional on t_i implied by p.

If the type spaces are finite, this is equivalent to the existence of a distributio p over type profiles such that

$$p_i(t_{-i}; t_i) = p(t_{-i}|t_i) = \frac{p\left(t\right)}{\sum_{t'_{-i}} p\left(t'_{-i}, t_i\right)}.$$

- If beliefs are consistent, the Bayesian game can be interpreted as having an initial move by nature, which selects $t \in T$ according to p.
- Suppose type spaces are finite. Viewed as a game of complete information, a profile $\hat{s}$ is a Nash equilibrium if, for all i, for all $s_i : T_i \to A_i$,

$$\sum_t u_i(\hat{s}(t), t)p(t) \geq \sum_t u_i(s_i(t_i), \hat{s}_{-i}(t_{-i}), t)p(t).$$

- If beliefs are consistent, the Bayesian game can be interpreted as having an initial move by nature, which selects $t \in T$ according to p.
- Suppose type spaces are finite. Viewed as a game of complete information, a profile $\hat{s}$ is a Nash equilibrium if, for all i, for all $s_i : T_i \to A_i$,

$$\sum_t u_i(\hat{s}(t), t)p(t) \geq \sum_t u_i(s_i(t_i), \hat{s}_{-i}(t_{-i}), t)p(t).$$

- This inequality can be rewritten as (where $p_i^*(t_i) := \sum_{t_{-i}} p(t_{-i}, t_i)$)

$$\sum_{t_i} \left\{ \sum_{t_{-i}} u_i(\hat{s}(t), t) p_i(t_{-i}; t_i) \right\} p_i^*(t_i) \geq \sum_{t_i} \left\{ \sum_{t_{-i}} u_i(s_i(t_i), \hat{s}_{-i}(t_{-i}), t) p_i(t_{-i}; t_i) \right\} p_i^*(t_i).$$

Global Games

Carlsson and van Damme (1993)

	A	B
A	θ, θ	$\theta - 9, 5$
B	$5, \theta - 9$	$7, 7$

- The parameter θ is uniformly distributed on the interval $[0, 20]$.
- For $\theta < 5$, B is strictly dominant, while for $\theta > 16$, A is strictly dominant.
- Each player i receives a signal x_i, with x_1 and x_2 independently and uniformly drawn from the interval $[\theta - \varepsilon,\ \theta + \varepsilon]$ for $\varepsilon > 0$.
- A pure strategy for player i is a function

$$s_i : [-\varepsilon,\ 20 + \varepsilon] \to \{A, B\}.$$

- For $x_i \in [\varepsilon,\ 20-\varepsilon]$, player i's posterior on θ is uniform on $[x_i-\varepsilon,\ x_i+\varepsilon]$.
- For $x_i \in [\varepsilon,\ 20-\varepsilon]$, player i's posterior on x_j is symmetric around x_i with support $[x_i-2\varepsilon,\ x_i+2\varepsilon]$. Hence,

$$\Pr\{x_j > x_i \mid x_i\} = \Pr\{x_j < x_i \mid x_i\} = \frac{1}{2}.$$

- For $x_i \in [\varepsilon,\ 20-\varepsilon]$, player i's posterior on θ is uniform on $[x_i-\varepsilon,\ x_i+\varepsilon]$.
- For $x_i \in [\varepsilon,\ 20-\varepsilon]$, player i's posterior on x_j is symmetric around x_i with support $[x_i-2\varepsilon,\ x_i+2\varepsilon]$. Hence,

$$\Pr\{x_j > x_i \mid x_i\} = \Pr\{x_j < x_i \mid x_i\} = \frac{1}{2}.$$

Lemma

For $\varepsilon < \frac{5}{2}$, the game has an essentially unique Nash equilibrium (s_1^*, s_2^*), given by

$$s_i^*(x_i) = \begin{cases} A, & \textit{if } x_i \geq 10\frac{1}{2}, \\ B, & \textit{if } x_i < 10\frac{1}{2}. \end{cases}$$

Proof

- Suppose $x_i < 5 - \varepsilon$, so that $\theta < 5$ (so that B is strictly dominant).

Proof

- Suppose $x_i < 5 - \varepsilon$, so that $\theta < 5$ (so that B is strictly dominant).
- Then, player i's payoff from A is less than that from B irrespective of player j's action, and so i plays B for $x_i < 5 - \varepsilon$ (as does j for $x_j < 5 - \varepsilon$).

Proof

- Suppose $x_i < 5 - \varepsilon$, so that $\theta < 5$ (so that B is strictly dominant).
- Then, player i's payoff from A is less than that from B irrespective of player j's action, and so i plays B for $x_i < 5 - \varepsilon$ (as does j for $x_j < 5 - \varepsilon$).
- But then at $x_i = 5 - \varepsilon$, since $\varepsilon < 5 - \varepsilon$, player i assigns at least probability $\frac{1}{2}$ to j playing B, and so i strictly prefers B.

Proof

- Suppose $x_i < 5 - \varepsilon$, so that $\theta < 5$ (so that B is strictly dominant).
- Then, player i's payoff from A is less than that from B irrespective of player j's action, and so i plays B for $x_i < 5 - \varepsilon$ (as does j for $x_j < 5 - \varepsilon$).
- But then at $x_i = 5 - \varepsilon$, since $\varepsilon < 5 - \varepsilon$, player i assigns at least probability $\frac{1}{2}$ to j playing B, and so i strictly prefers B.
- Define

$$x_i^* := \sup\{x_i' \mid B \text{ is implied by iterated strict dominance for all } x_i < x_i'\}.$$

Proof

- Suppose $x_i < 5 - \varepsilon$, so that $\theta < 5$ (so that B is strictly dominant).
- Then, player i's payoff from A is less than that from B irrespective of player j's action, and so i plays B for $x_i < 5 - \varepsilon$ (as does j for $x_j < 5 - \varepsilon$).
- But then at $x_i = 5 - \varepsilon$, since $\varepsilon < 5 - \varepsilon$, player i assigns at least probability $\frac{1}{2}$ to j playing B, and so i strictly prefers B.
- Define

$$x_i^* := \sup\{x_i' \mid B \text{ is implied by iterated strict dominance for all } x_i < x_i'\}.$$

- By symmetry, $x_1^* = x_2^* = x^*$. At $x_i = x^*$, player i cannot strictly prefer B to A (since $E[\theta \mid x^*] = x^*$ and $p \leq \frac{1}{2}$):

$$px^* + (1-p)(x^* - 9) = p5 + (1-p)7$$

and so $x^* \geq 10\frac{1}{2}$.

Proof

conclusion

- Define

$$x_i^{**} := \inf\{x_i'' \mid A \text{ is implied by iterated strict dominance for all } x_i > x_i''\}.$$

- Then,

$$x^{**} \leq 10\tfrac{1}{2},$$

and so

$$10\tfrac{1}{2} \leq x^{*} \leq x^{**} \leq 10\tfrac{1}{2}.$$

Mixed Strategies and Purification

The assumption that players can randomize is sometimes criticized on three grounds:

1. players don't randomize;
2. there is no reason for a player randomize with just the right probability, when the player is indifferent over all possible randomization probabilities (including 0 and 1); and
3. a randomizing player is subject to ex post regret.

Ex post regret

- A player is said to be subject to ex post regret if after all uncertainty is resolved, a player would like to change his/her decision (i.e., has regret).
- In a game with no moves of nature, no player has ex post regret in a pure (but not mixed) strategy equilibrium.
- Any pure strategy equilibrium of a game with moves of nature will typically also have ex post regret.
  - Ex post regret should not be viewed as a criticism of mixing, but rather a caution to modelers. If a player has ex post regret, then that player has an incentive to change his/her choice. Whether a player is able to do so depends upon the scenario being modeled. If the player cannot do so, then there is no issue. If, however, the player can do so, then that option should be included in the game description.

Purification

- Player i's mixed strategy σ_i of a game G is said to be purified if in an "approximating" version of G with private information (with player i's private information given by T_i), that player's behavior can be written as a pure strategy $s_i : T_i \to A_i$ such that

$$\sigma_i(a_i) \approx \Pr\{s_i(t_i) = a_i\},$$

where $\Pr$ is given by the prior distribution over T_i (and so describes player $j \neq i$ beliefs over T_i).

Example of Purification

	A	B
A	$9,9$	$0,5$
B	$5,0$	$7,7$

The game has two strict pure strategy Nash equilibria and one symmetric mixed strategy Nash equilibrium. Let $p = \Pr\{A\}$, then in the mixed strategy eq

$$\begin{aligned} 9p &= 5p + 7(1-p) \\ \iff 9p &= 7 - 2p \\ \iff 11p &= 7 \iff p = 7/11. \end{aligned}$$

Example of Purification (cont)

Trivial purification: Give player i payoff-irrelevant information t_i, where $t_i \sim \mathcal{U}([0,1])$, and t_1 and t_2 are independent. This is a game with private information, where player i learns t_i before choosing his or her action.

- The mixed strategy equilibrium is purified by many pure strategy equilibria in the game with private information, such as

$$s_i(t_i) = \begin{cases} B, & \text{if } t_i \leq 4/11, \\ A, & \text{if } t_i \geq 4/11. \end{cases}$$

Better Purification

Harsanyi (1973)

	A	B
A	$9 + \varepsilon t_1, 9 + \varepsilon t_2$	$0, 5$
B	$5, 0$	$7, 7$

- Player i's type $t_i \sim \mathcal{U}([0,1])$ and t_1 and t_2 are independent.
- A pure strategy for player i is $s_i : [0,1] \to \{A, B\}$. Suppose 2 is following a cutoff strategy (with $\bar{t}_2 \in (0,1)$),

$$s_2(t_2) = \begin{cases} A, & t_2 \geq \bar{t}_2, \\ B, & t_2 < \bar{t}_2. \end{cases}$$

Type t_1 expected payoff from A is

$$\begin{aligned} U_1(A, t_1, s_2) &= (9 + \varepsilon t_1) \Pr\{s_2(t_2) = A\} \\ &= (9 + \varepsilon t_1) \Pr\{t_2 \geq \bar{t}_2\} \\ &= (9 + \varepsilon t_1)(1 - \bar{t}_2), \end{aligned}$$

while from B is

$$\begin{aligned} U_1(B, t_1, s_2) &= 5 \Pr\{t_2 \geq \bar{t}_2\} + 7 \Pr\{t_2 < \bar{t}_2\} \\ &= 5(1 - \bar{t}_2) + 7\bar{t}_2 \\ &= 5 + 2\bar{t}_2. \end{aligned}$$

Thus, A is optimal if and only if

$$(9 + \varepsilon t_1)(1 - \bar{t}_2) \geq 5 + 2\bar{t}_2,$$

i.e.,

$$t_1 \geq \frac{11\bar{t}_2 - 4}{\varepsilon(1 - \bar{t}_2)}.$$

- In the symmetric equilibrium: $\bar{t}_1 = \bar{t}_2 = \bar{t}$, that is,

$$\bar{t} = \frac{11\bar{t} - 4}{\varepsilon(1 - \bar{t})},$$

  or

$$\varepsilon\bar{t}^2 + (11 - \varepsilon)\bar{t} - 4 = 0.$$

- Let $t(\varepsilon)$ denote the value of $\bar{t}$ satisfying this equality.
  - Note that $t(0) = 4/11$.
  - Write the equality as $g(\bar{t}, \varepsilon) = 0$.
  - Apply the implicit function theorem (since $\partial g/\partial \bar{t} \neq 0$ at $\varepsilon = 0$) to conclude that for $\varepsilon > 0$ but close to 0, the cutoff value of $\bar{t}$, $t(\varepsilon)$, is close to $4/11$ (the probability of the mixed strategy equilibrium in the unperturbed game).

Auctions

First-Price Sealed Bid Private-Value Auctions

- Bidder i's value for the object, v_i is known only to i.
- Nature chooses v_i, $i = 1, 2$, with v_i being independently drawn from the interval $[\underline{v}_i, \bar{v}_i]$, with distribution F_i and density f_i.
- Bidders know F_i (and so f_i).
- The set of possible bids is $\mathbb{R}_+$.
- Bidder i's ex post payoff as a function of b_1 and b_2, and values v_1 and v_2:

$$u_i(b_1, b_2, v_1, v_2) = \begin{cases} 0, & \text{if} \quad b_i < b_j, \\ \frac{1}{2}\left(v_i - b_i\right), & \text{if} \quad b_i = b_j, \\ v_i - b_i, & \text{if} \quad b_i > b_j. \end{cases}$$

- Suppose bidder 2 uses a strategy $\sigma_2 : [\underline{v}_2, \ \bar{v}_2] \to \mathbb{R}_+$.

- Suppose bidder 2 uses a strategy $\sigma_2 : [\underline{v}_2,\ \bar{v}_2] \to \mathbb{R}_+$.
- Then, bidder 1's expected (or interim) payoff from bidding b_1 at v_1 is

$$
\begin{aligned}
U_1\left(b_1, v_1; \sigma_2\right) &= \int u_1\left(b_1, \sigma_2\left(v_2\right), v_1, v_2\right)\, dF_2(v_2) \\
&= \frac{1}{2}\left(v_1 - b_1\right) \Pr\left\{\sigma_2(v_2) = b_1\right\} \\
&\qquad + \int_{\{v_2 : \sigma_2(v_2) < b_1\}} \left(v_1 - b_1\right) f_2\left(v_2\right)\, dv_2.
\end{aligned}
$$

- Suppose bidder 2 uses a strategy $\sigma_2 : [\underline{v}_2, \bar{v}_2] \to \mathbb{R}_+$.
- Then, bidder 1's expected (or interim) payoff from bidding b_1 at v_1 is

$$\begin{aligned} U_1(b_1, v_1; \sigma_2) &= \int u_1(b_1, \sigma_2(v_2), v_1, v_2)\, dF_2(v_2) \\ &= \frac{1}{2}(v_1 - b_1) \Pr\{\sigma_2(v_2) = b_1\} \\ &\quad + \int_{\{v_2 : \sigma_2(v_2) < b_1\}} (v_1 - b_1) f_2(v_2)\, dv_2. \end{aligned}$$

- Player 1's ex ante payoff from the strategy σ_1 is given by

$$\int U_1(\sigma_1(v_1), v_1; \sigma_2)\, dF_1(v_1),$$

and so for an optimal strategy σ_1, the bid $b_1 = \sigma_1(v_1)$ must maximize $U_1(b_1, v_1; \sigma_2)$ for almost all v_1.

- Proceed by “guess and verify”: that is, we impose a sequence of increasingly demanding conditions on the strategy of player 2, and prove that there is a best reply for player 1 satisfying these conditions.

- Proceed by “guess and verify”: that is, we impose a sequence of increasingly demanding conditions on the strategy of player 2, and prove that there is a best reply for player 1 satisfying these conditions.
- Begin by supposing σ_2 is strictly increasing, so that $\Pr\{\sigma_2(v_2) = b_1\} = 0$.

- Proceed by "guess and verify": that is, we impose a sequence of increasingly demanding conditions on the strategy of player 2, and prove that there is a best reply for player 1 satisfying these conditions.
- Begin by supposing σ_2 is strictly increasing, so that $\Pr\{\sigma_2(v_2) = b_1\} = 0$.
- Without loss of generality, restrict attention to bids b_1 in the range of σ_2.

- Proceed by “guess and verify”: that is, we impose a sequence of increasingly demanding conditions on the strategy of player 2, and prove that there is a best reply for player 1 satisfying these conditions.
- Begin by supposing σ_2 is strictly increasing, so that $\Pr\{\sigma_2(v_2) = b_1\} = 0$.
- Without loss of generality, restrict attention to bids b_1 in the range of σ_2.
- Then,

$$\begin{aligned}
U_1(b_1, v_1; \sigma_2) &= \int_{\{v_2 : \sigma_2(v_2) < b_1\}} (v_1 - b_1) f_2(v_2)\, dv_2 \\
&= E[v_1 - b_1 \mid \text{winning}] \Pr\{\text{winning}\} \\
&= (v_1 - b_1) \Pr\{\sigma_2(v_2) < b_1\} \\
&= (v_1 - b_1) \Pr\{v_2 < \sigma_2^{-1}(b_1)\} \\
&= (v_1 - b_1) F_2(\sigma_2^{-1}(b_1)).
\end{aligned}$$

- Need to choose b_1 to max $U_1 \left(b_1, v_1; \sigma_2 \right) = (v_1 - b_1) F_2(\sigma_2^{-1}(b_1))$.

- Need to choose b_1 to max $U_1\left(b_1, v_1; \sigma_2\right) = (v_1 - b_1)F_2(\sigma_2^{-1}(b_1))$.
- Suppose σ_2 is differentiable, and that the bid $b_1 = \sigma_1(v_1)$ is an interior maximum.

- Need to choose b_1 to max $U_1\left(b_1, v_1; \sigma_2\right) = (v_1 - b_1)F_2(\sigma_2^{-1}(b_1))$.
- Suppose σ_2 is differentiable, and that the bid $b_1 = \sigma_1(v_1)$ is an interior maximum.
- The first order condition is

$$0 = -F_2\left(\sigma_2^{-1}\left(b_1\right)\right) + \left(v_1 - b_1\right) f_2\left(\sigma_2^{-1}\left(b_1\right)\right) d\sigma_2^{-1}\left(b_1\right)/db_1.$$

- Need to choose b_1 to max $U_1(b_1, v_1; \sigma_2) = (v_1 - b_1)F_2(\sigma_2^{-1}(b_1))$.
- Suppose σ_2 is differentiable, and that the bid $b_1 = \sigma_1(v_1)$ is an interior maximum.
- The first order condition is

$$0 = -F_2\left(\sigma_2^{-1}(b_1)\right) + (v_1 - b_1)\, f_2\left(\sigma_2^{-1}(b_1)\right) d\sigma_2^{-1}(b_1)/db_1.$$

- But

$$\frac{d\sigma_2^{-1}(b_1)}{db_1} = \frac{1}{\sigma_2'(\sigma_2^{-1}(b_1))},$$

so

$$F_2\left(\sigma_2^{-1}(b_1)\right) \sigma_2'\left(\sigma_2^{-1}(b_1)\right) = (v_1 - b_1)\, f_2\left(\sigma_2^{-1}(b_1)\right),$$

i.e.,

$$\sigma_2'\left(\sigma_2^{-1}(b_1)\right) = \frac{(v_1 - b_1)\, f_2\left(\sigma_2^{-1}(b_1)\right)}{F_2\left(\sigma_2^{-1}(b_1)\right)}.$$

- Assume $F_1 = F_2$, and suppose the equilibrium is symmetric, so that $\sigma_1 = \sigma_2 = \tilde{\sigma}$, and $b_1 = \sigma_1(v)$ implies $v = \sigma_2^{-1}(b_1)$.
- Then,

$$\sigma_2'\left(\sigma_2^{-1}(b_1)\right) = \frac{(v_1 - b_1) f_2\left(\sigma_2^{-1}(b_1)\right)}{F_2\left(\sigma_2^{-1}(b_1)\right)}.$$

becomes (dropping subscripts),

$$\tilde{\sigma}'(v) = \frac{(v - \tilde{\sigma}(v))f(v)}{F(v)},$$

or

$$\tilde{\sigma}'(v)F(v) + \tilde{\sigma}(v)f(v) = vf(v).$$

- We have

$$\tilde{\sigma}'(v)F(v) + \tilde{\sigma}(v)f(v) = vf(v).$$

- We have

$$\tilde{\sigma}'(v)F(v) + \tilde{\sigma}(v)f(v) = vf(v).$$

- But

$$\frac{d}{dv}\tilde{\sigma}(v)F(v) = \tilde{\sigma}'(v)F(v) + \tilde{\sigma}(v)f(v),$$

so

$$\tilde{\sigma}(\hat{v})F(\hat{v}) = \int_{\underline{v}}^{\hat{v}} vf(v)dv + k,$$

where k is a constant of integration.

- We have
$$\tilde{\sigma}'(v)F(v) + \tilde{\sigma}(v)f(v) = vf(v).$$
- But
$$\frac{d}{dv}\tilde{\sigma}(v)F(v) = \tilde{\sigma}'(v)F(v) + \tilde{\sigma}(v)f(v),$$
so
$$\tilde{\sigma}(\hat{v})F(\hat{v}) = \int_{\underline{v}}^{\hat{v}} vf(v)dv + k,$$
where k is a constant of integration.
- Moreover, evaluating both sides at $\hat{v} = \underline{v}$ shows that $k = 0$, and so
$$\tilde{\sigma}(\hat{v}) = \frac{1}{F(\hat{v})}\int_{\underline{v}}^{\hat{v}} vf(v)dv = E[v \mid v \leq \hat{v}].$$

Summary

- Each bidder bids the expectation of the other bidder's valuation, conditional on that valuation being less than his (i.e., conditional on his value being the highest). This is not an accident.
- Summarizing the calculations till this point, we have shown that if $(\tilde{\sigma}, \tilde{\sigma})$ is a Nash equilibrium in which $\tilde{\sigma}$ is a strictly increasing and differentiable function, and $\tilde{\sigma}(v)$ is interior (which here means strictly positive), then it is given by $\tilde{\sigma}(\hat{v}) = E[v \mid v \leq \hat{v}]$.
  Note that $E[v \mid v \leq \hat{v}]$ is increasing in $\hat{v}$ and lies in the interval $[\underline{v}, Ev]$.
- It remains to verify the hypotheses. It is immediate that $\tilde{\sigma}$ is strictly increasing and differentiable. Moreover, for $v > \underline{v}$, $\tilde{\sigma}(v)$ is strictly positive. It remains to verify the optimality of bids.

Optimality

- It is not optimal to bid $b_1 < \underline{v} = \tilde{\sigma}(\underline{v})$ or $b_1 > Ev = \tilde{\sigma}(\bar{v})$.
- Since $\tilde{\sigma}$ is strictly increasing and continuous, any bid in $[\underline{v}, Ev]$ is the bid of some valuation v.
- Bidding as if valuation $\hat{v}$ has valuation v' is suboptimal:

$$\begin{aligned}
U(v'; \hat{v}) &= (\hat{v} - \tilde{\sigma}(v')) \Pr(v_2 \leq \tilde{\sigma}^{-1}(\tilde{\sigma}(v'))) \\
&= (\hat{v} - E[v \mid v \leq v'])F(v') \\
&= \left(\hat{v} - \frac{1}{F(v')} \int_{\underline{v}}^{v'} vf(v)dv\right) F(v') \\
&= \int_{\underline{v}}^{v'} (\hat{v} - v)f(v)dv.
\end{aligned}$$

Common Value Auctions

- Each bidder receives a private signal about the value of the object, t_i, with $t_i \in T_i = [0, 1]$, uniformly independently distributed.
- The common (to both players) value of the object is $v = t_1 + t_2$.
- Ex post payoffs are given by

$$u_i(b_1, b_2, t_1, t_2) = \begin{cases} t_1 + t_2 - b_i, & \text{if } b_i > b_j, \\ \frac{1}{2}(t_1 + t_2 - b_i), & \text{if } b_i = b_j, \\ 0, & \text{if } b_i < b_j. \end{cases}$$

- Suppose bidder 2 uses strategy $\sigma_2 : T_2 \to \mathbb{R}_+$.
- Suppose σ_2 is strictly increasing.
- Then, t_1's expected payoff from bidding b_1 is

$$\begin{aligned}
U_1(b_1, t_1; \sigma_2) &= E[t_1 + t_2 - b_1 \mid \text{winning}] \Pr\{\text{winning}\} \\
&= E[t_1 + t_2 - b_1 \mid t_2 < \sigma_2^{-1}(b_1)] \Pr\{t_2 < \sigma_2^{-1}(b_1)\} \\
&= (t_1 - b_1)\sigma_2^{-1}(b_1) + \int_0^{\sigma_2^{-1}(b_1)} t_2 \, dt_2 \\
&= (t_1 - b_1)\sigma_2^{-1}(b_1) + (\sigma_2^{-1}(b_1))^2/2.
\end{aligned}$$

- Maximizing $U_1(b_1, t_1; \sigma_2) = (t_1 - b_1)\sigma_2^{-1}(b_1) + (\sigma_2^{-1}(b_1))^2/2$.
- If σ_2 is differentiable, the first order condition is

$$0 = -\sigma_2^{-1}(b_1) + (t_1 - b_1)d\sigma_2^{-1}(b_1)/db_1 + \sigma_2^{-1}(b_1)d\sigma_2^{-1}(b_1)/db_1,$$

and so

$$\sigma_2^{-1}(b_1)\sigma_2'(\sigma_2^{-1}(b_1)) = (t_1 + \sigma_2^{-1}(b_1) - b_1).$$

- Suppose the equilibrium is symmetric, so that $\sigma_1 = \sigma_2 = \sigma$. Then,

$$t\sigma'(t) = 2t - \sigma(t).$$

- Integrating,

$$t\sigma(t) = t^2 + k,$$

where k is a constant of integration. Evaluating at $t = 0$ shows that $k = 0$, and so

$$\sigma(t) = t.$$

Winner's Curse

- Note that this is not the profile that results from the analysis of the private value auction when $\underline{v} = 1/2$ ($E[t_1 + t_2 \mid t_1] = t_1 + 1/2$).
- In particular, letting $v' = t + \frac{1}{2}$, we have

$$\sigma_{\text{private value}}(t) = \tilde{\sigma}(v') = \frac{v' + 1/2}{2} = \frac{t+1}{2} > t = \sigma_{\text{common value}}(t).$$

- This illustrates the winner's curse: $E[v \mid t_1] > E[v|t_1, \text{ winning}]$. In particular, in the equilibrium just calculated,

$$\begin{aligned} E[v \mid t_1, \text{ winning}] &= E[t_1 + t_2 \mid t_1, \ t_2 < t_1] \\ &= t_1 + \frac{1}{t_1}\left[(t_2)^2/2\right]_0^{t_1} = \frac{3t_1}{2}, \end{aligned}$$

while $E[v \mid t_1] = t_1 + \frac{1}{2} > 3t_1/2$ (recall $t_1 \in [0,1]$).

CARLSSON, H., AND E. VAN DAMME (1993): "Global Games and Equilibrium Selection," Econometrica, 61(5), 989–1018.

HARSANYI, J. C. (1967): "Games with Incomplete Information Played by Bayesian Players. Part I," Management Science, 14(3), 159–182.

——— (1968a): "Games with Incomplete Information Played by Bayesian Players. Part II," Management Science, 14(5), 320–334.

——— (1968b): "Games with Incomplete Information Played by Bayesian Players. Part III," Management Science, 14(7), 486–502.

——— (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," International Journal of Game Theory, 2(1), 1–23.

SAVAGE, L. J. (1972): The Foundations of Statistics. Dover, New York, NY, second revised edn.

VON NEUMANN, J., AND O. MORGENSTERN (1944): Theory of games and economic behavior. Princeton University Press, 1st edn., 2nd edn 1947, 3rd edn 1953.