6.254 : Game Theory with Engineering Applications
Lecture 17: Games with Incomplete Information: Bayesian Nash Equilibria

Asu Ozdaglar
MIT

April 15, 2010

Outline

- Incomplete information.
- Bayes rule and Bayesian inference.
- Bayesian Nash Equilibria.
- Auctions.

- Reading:
- Fudenberg and Tirole, Sections 6.1-6.5.
- Krishna, Chapters 1-5.

Incomplete Information

- In many game theoretic situations, one agent is unsure about the payoffs or preferences of others.
- Incomplete information introduces additional strategic interactions and also raises questions related to "learning".
- Examples:
    - *Bargaining* (how much the other party is willing to pay is generally unknown to you)
    - *Auctions* (how much should you bid for an object that you want, knowing that others will also compete against you?)
    - *Market competition* (firms generally do not know the exact cost of their competitors)
    - *Signaling games* (how should you infer the information of others from the signals they send)
    - *Social learning* (how can you leverage the decisions of others in order to make better decisions)

Example: Incomplete Information Battle of the Sexes

- Recall the battle of the sexes game, which was a complete information "coordination" game.
- Both parties want to meet, but they have different preferences on "Ballet" and "Football".

	B	F
B	$(2,1)$	$(0,0)$
F	$(0,0)$	$(1,2)$

- In this game there are two pure strategy equilibria (one of them better for player 1 and the other one better for player 2), and a mixed strategy equilibrium.
- Now imagine that player 1 does not know whether player 2 wishes to meet or wishes to avoid player 1. Therefore, this is a situation of **incomplete information**—also sometimes called *asymmetric information*.

Example (continued)

- We represent this by thinking of player 2 having two different **types**, one type that wishes to meet player 1 and the other wishes to avoid him.

- More explicitly, suppose that these two types have probability $1/2$ each. Then the game takes the form one of the following two with probability $1/2$.

	B	F
B	$(2,1)$	$(0,0)$
F	$(0,0)$	$(1,2)$

	B	F
B	$(2,0)$	$(0,2)$
F	$(0,1)$	$(1,0)$

- Crucially, player 2 knows which game it is (she knows the **state of the world**), but player 1 does not.

- What are strategies in this game?

Example (continued)

- Most importantly, from player 1's point of view, player 2 has two possible types (or equivalently, the world has two possible states each with 1/2 probability and only player 2 knows the actual state).

- How do we reason about equilibria here?
- **Idea:** Use Nash Equilibrium concept in an expanded game, where each different type of player 2 has a different strategy
  - Or equivalently, form conjectures about other player's actions in each state and act optimally given these conjectures.

Example (continued)

- Let us consider the following strategy profile $(B, (B, F))$, which means that player 1 will play B, and while in state 1, player 2 will also play B (when she wants to meet player 1) and in state 2, player 2 will play F (when she wants to avoid player 1).
- Clearly, given the play of B by player 1, the strategy of player 2 is a best response.
- Let us now check that player 1 is also playing a best response.
- Since both states are equally likely, the expected payoff of player 1 is

$$\mathbb{E}[B, (B, F)] = \frac{1}{2} \times 2 + \frac{1}{2} \times 0 = 1.$$

- If, instead, he deviates and plays F, his expected payoff is

$$\mathbb{E}[F, (B, F)] = \frac{1}{2} \times 0 + \frac{1}{2} \times 1 = \frac{1}{2}.$$

- Therefore, the strategy profile $(B, (B, F))$ is a (Bayesian) Nash equilibrium.

Example (continued)

- Interestingly, meeting at Football, which is the preferable outcome for player 2 is no longer a Nash equilibrium. Why not?
- Suppose that the two players will meet at Football when they want to meet. Then the relevant strategy profile is $(F, (F, B))$ and

$$\mathbb{E}[F, (F, B)] = \frac{1}{2} \times 1 + \frac{1}{2} \times 0 = \frac{1}{2}.$$

- If, instead, player 1 deviates and plays B, his expected payoff is

$$\mathbb{E}[B, (F, B)] = \frac{1}{2} \times 0 + \frac{1}{2} \times 2 = 1.$$

- Therefore, the strategy profile $(F, (F, B))$ is not a (Bayesian) Nash equilibrium.

Bayesian Games

- More formally, we can define Bayesian games, or "incomplete information games" as follows.

Definition

A Bayesian game consists of

- A set of players $\mathcal{I}$;
- A set of actions (pure strategies) for each player i: S_i;
- A set of types for each player i: $\theta_i \in \Theta_i$;
- A payoff function for each player i: $u_i(s_1, \ldots, s_I, \theta_1, \ldots, \theta_I)$;
- A (joint) probability distribution $p(\theta_1, \ldots, \theta_I)$ over types (or $P(\theta_1, \ldots, \theta_I)$ when types are not finite).

- More generally, one could also allow for a signal for each player, so that the signal is correlated with the underlying type vector.

Bayesian Games (continued)

- Importantly, throughout in Bayesian games, the strategy spaces, the payoff functions, possible types, and the prior probability distribution are assumed to be **common knowledge**.
    - Very strong assumption.
    - But very convenient, because any private information is included in the description of the type and others can form beliefs about this type and each player understands others' beliefs about his or her own type, and so on, and so on.

Definition

A (pure) strategy for player i is a map $s_i : \Theta_i \to S_i$ prescribing an action for each possible type of player i.

Bayesian Games (continued)

- Recall that player types are drawn from some prior probability distribution $p(\theta_1, \ldots, \theta_I)$.

- Given $p(\theta_1, \ldots, \theta_I)$ we can compute the conditional distribution $p(\theta_{-i} \mid \theta_i)$ using Bayes rule.
  - Hence the label "Bayesian games".
  - Equivalently, when types are not finite, we can compute the conditional distribution $P(\theta_{-i} \mid \theta_i)$ given $P(\theta_1, \ldots, \theta_I)$.

- Player i knows her own type and evaluates her expected payoffs according to the **conditional distribution** $p(\theta_{-i} \mid \theta_i)$, where $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_I)$.

Bayesian Games (continued)

- Since the payoff functions, possible types, and the prior probability distribution are common knowledge, we can compute expected payoffs of player i of type θ_i as

$$U\left(s_i', s_{-i}, \theta_i\right) = \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})$$

when types are finite

$$= \int u_i(s_i', s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) P(d\theta_{-i} \mid \theta_i)$$

when types are not finite.

Bayesian Nash Equilibria

Definition

(Bayesian Nash Equilibrium) The strategy profile $s(\cdot)$ is a (pure strategy) Bayesian Nash equilibrium if for all $i \in \mathcal{I}$ and for all $\theta_i \in \Theta_i$, we have that

$$s_i(\theta_i) \in \arg \max_{s'_i \in S_i} \sum_{\theta_{-i}} p(\theta_{-i} \mid \theta_i) u_i(s'_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}),$$

or in the non-finite case,

$$s_i(\theta_i) \in \arg \max_{s'_i \in S_i} \int u_i(s'_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i}) P(d\theta_{-i} \mid \theta_i) .$$

- Hence a Bayesian Nash equilibrium is a Nash equilibrium of the "expanded game" in which each player i's space of pure strategies is the set of maps from Θ_i to S_i.

Existence of Bayesian Nash Equilibria

Theorem

Consider a finite incomplete information (Bayesian) game. Then a mixed strategy Bayesian Nash equilibrium exists.

Theorem

Consider a Bayesian game with continuous strategy spaces and continuous types. If strategy sets and type sets are compact, payoff functions are continuous and concave in own strategies, then a pure strategy Bayesian Nash equilibrium exists.

- The ideas underlying these theorems and proofs are identical to those for the existence of equilibria in (complete information) strategic form games.

Example: Incomplete Information Cournot

- Suppose that two firms both produce at constant marginal cost.
- Demand is given by $P(Q)$ as in the usual Cournot game.
- Firm 1 has marginal cost equal to C (and this is common knowledge).
- Firm 2's marginal cost is private information. It is equal to C_L with probability θ and to C_H with probability $(1-\theta)$, where $C_L < C_H$.

- This game has 2 players, 2 states (L and H) and the possible actions of each player are $q_i \in [0, \infty)$, but firm 2 has two possible types.
- The payoff functions of the players, after quantity choices are made, are given by

$$\begin{aligned} u_1((q_1, q_2), t) &= q_1(P(q_1 + q_2) - C) \\ u_2((q_1, q_2), t) &= q_2(P(q_1 + q_2) - C_t), \end{aligned}$$

where $t \in \{L, H\}$ is the type of player 2.

Example (continued)

- A strategy profile can be represented as (q_1^*, q_L^*, q_H^*) [or equivalently as $(q_1^*, q_2^*(\theta_2))$], where q_L^* and q_H^* denote the actions of player 2 as a function of its possible types.
- We now characterize the Bayesian Nash equilibria of this game by computing the best response functions (correspondences) and finding their intersection.
- There are now three best response functions and they are are given by

$$B_1(q_L, q_H) = \arg \max_{q_1 \geq 0} \{\theta(P(q_1 + q_L) - C)q_1$$
$$+ (1 - \theta)(P(q_1 + q_H) - C)q_1\}$$
$$B_L(q_1) = \arg \max_{q_L \geq 0} \{(P(q_1 + q_L) - C_L)q_L\}$$
$$B_H(q_1) = \arg \max_{q_H \geq 0} \{(P(q_1 + q_H) - C_H)q_H\}.$$

Example (continued)

- The Bayesian Nash equilibria of this game are vectors (q_1^*, q_L^*, q_H^*)
  such that

$$B_1(q_L^*, q_H^*) = q_1^*, \quad B_L(q_1^*) = q_L^*, \quad B_H(q_1^*) = q_H^*.$$

- To simplify the algebra, let us assume that $P(Q) = \alpha - Q$, $Q \leq \alpha$.
  Then we can compute:

$$
\begin{aligned}
q_1^* &= \frac{1}{3}(\alpha - 2C + \theta C_L + (1 - \theta)C_H) \\
q_L^* &= \frac{1}{3}(\alpha - 2C_L + C) - \frac{1}{6}(1 - \theta)(C_H - C_L) \\
q_H^* &= \frac{1}{3}(\alpha - 2C_H + C) + \frac{1}{6}\theta(C_H - C_L).
\end{aligned}
$$

Example (continued)

- Note that $q_L^* > q_H^*$. This reflects the fact that with lower marginal cost, the firm will produce more.
- However, incomplete information also affects firm 2's output choice.
- Recall that, given this demand function, if both firms knew each other's marginal cost, then the unique Nash equilibrium involves output of firm i given by

$$\frac{1}{3}(\alpha - 2C_i + C_j).$$

- With incomplete information, firm 2's output is more if its cost is C_H and less if its cost is C_L. If firm 1 knew firm 2's cost is high, then it would produce more. However, its lack of information about the cost of firm 2 leads firm 1 to produce a relatively moderate level of output, which then allows firm 2 to be more "aggressive".
- Hence, in this case, firm 2 benefits from the lack of information of firm 1 and it produces more than if 1 knew his actual cost.

Auctions

- A major application of Bayesian games is to auctions, which are historically and currently common method of allocating scarce goods across individuals with different valuations for these goods.

- This corresponds to a situation of incomplete information because the valuations of different potential buyers are unknown.

- For example, if we were to announce a **mechanism**, which involves giving a particular good (for example a seat in a sports game) for free to the individual with the highest valuation, this would create an incentive for all individuals to overstate their valuations.

- In general, auctions are designed by **profit-maximizing** entities, which would like to sell the goods to raise the highest possible revenue.

Auctions (continued)

- Different types of auctions and terminology:
    - *English auctions*: ascending sequential bids.
    - *First price sealed bid auctions*: similar to English auctions, but in the form of a strategic form game; all players make a single simultaneous bid and the highest one obtains the object and pays its bid.
    - *Second price sealed bid auctions*: similar to first price auctions, except that the winner pays the second highest bid.
    - *Dutch auctions: descending sequential auctions*; the auction stops when an individual announces that she wants to buy at that price. Otherwise the price is reduced sequentially until somebody stops the auction.
    - *Combinatorial auctions*: when more than one item is auctioned, and agents value combinations of items.
    - *Private value auctions*: valuation of each agent is independent of others' valuations;
    - *Common value auctions*: the object has a potentially common value, and each individual's signal is imperfectly correlated with this common value.

Modeling Auctions

- Model of auction:
  - a valuation structure for the bidders (i.e., private values for the case of private-value auctions),
  - a probability distribution over the valuations available to the bidders.

- Let us focus on first and second price sealed bid auctions, where bids are submitted simultaneously.

- Each of these two auction formats defines a static game of incomplete information (Bayesian game) among the bidders.

- We determine Bayesian Nash equilibria in these games and compare the equilibrium bidding behavior.

Modeling Auctions (continued)

- More explicitly, suppose that there is a single object for sale and N potential buyers are bidding for the object.
- Bidder i assigns a value v_i to the object, i.e., a utility

$$v_i - b_i,$$

  when he pays b_i for the object. He knows v_i. This implies that we have a **private value** auction (v_i is his "private information" and "private value").
- Suppose also that each v_i is *independently and identically distributed* on the interval $[0, \bar{v}]$ with cumulative distribution function F, with continuous density f and full support on $[0, \bar{v}]$.
- Bidder i knows the realization of its value v_i (or realization v_i of the random variable V_i, though we will not use this latter notation) and that other bidders' values are independently distributed according to F, i.e., all components of the model except the realized values are "common knowledge".

Modeling Auctions (continued)

- Bidders are *risk neutral*, i.e., they are interested in maximizing their expected profits.
- This model defines a Bayesian game of incomplete information, where the types of the players (bidders) are their valuations, and a pure strategy for a bidder is a map

$$\beta_i : [0, \bar{v}] \to \mathbb{R}_+.$$

- We will characterize the symmetric equilibrium strategies in the first and second price auctions.
- Once we characterize these equilibria, then we can also investigate which auction format yields a higher expected revenue to the seller at the symmetric equilibrium.

Second Price Auctions

- Second price auctions will have the structure very similar to a complete information auction discussed earlier in the lectures.
- There we saw that each player had a weakly dominant strategy. This will be true in the incomplete information version of the game and will greatly simplify the analysis.
- In the auction, each bidder submits a sealed bid of b_i, and given the vector of bids $b = (b_i, b_{-i})$ and valuation v_i of player i, its payoff is

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

- Let us also assume that if there is a tie, i.e., $b_i = \max_{j \neq i} b_j$, the object goes to each winning bidder with equal probability.
- With the reasoning similar to its counterpart with complete information, in a second-price auction, it is a weakly dominant strategy to bid *truthfully*, i.e., according to $\beta^{II}(v) = v$.

Second Price Auctions (continued)

- This can be established with the same graphical argument as the one we had for the complete information case.
- The first graph shows the payoff for bidding one's valuation, the second graph the payoff from bidding a lower amount, and the third the payoff from bidding higher amount.
- In all cases B^* denotes the highest bid excluding this player.
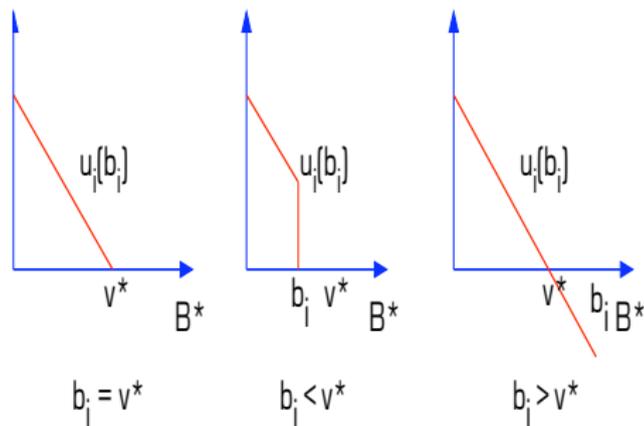
Second Price Auctions (continued)

- Moreover, now there are no other optimal strategies and thus the (Bayesian) equilibrium will be unique, since the valuation of other players are not known.
- Therefore, we have established:

Proposition
In the second price auction, there exists a unique Bayesian Nash equilibrium which involves
$$\beta^{II}(v) = v.$$

- It is straightforward to generalize the exact same result to the situation in which bidders' values are not independently and identically distributed. As in the complete information case, bidding truthfully remains a weakly dominant strategy.
- The assumption of private values is important (i.e., the valuations are known at the time of the bidding).

Second Price Auctions (continued)

- Let us next determine the expected payment by a bidder with value v, and for this, let us focus on the case in which valuations are independent and identically distributed

- Fix bidder 1 and define the random variable y_1 as the highest value among the remaining $N - 1$ bidders, i.e.,

$$y_1 = \max\{v_2, \ldots, v_N\}.$$

- Let G denote the cumulative distribution function of y_1.

- Clearly,

$$G(y) = F(y)^{N-1} \text{ for any } y \in [0, \bar{v}]$$

Second Price Auctions (continued)

- In a second price auction, the expected payment by a bidder with value v is given by

$$m^{II}(v) = \Pr(v \text{ wins}) \times \mathbb{E}[\text{second highest bid} \mid v \text{ is the highest bid}]$$
$$= \Pr(y_1 \leq v) \times \mathbb{E}[y_1 \mid y_1 \leq v]$$
$$= G(v) \times \mathbb{E}[y_1 \mid y_1 \leq v]. \qquad \textit{Payment II}$$

- Note that here and in what follows, we can use strict or weak inequalities given that the relevant random variables have continuous distributions. In other words, we have

$$\Pr(y_1 \leq v) = \Pr(y_1 < v).$$

Example: Uniform Distributions

- Suppose that there are two bidders with valuations, v_1 and v_2, distributed uniformly over $[0, 1]$.
- Then $G(v_1) = v_1$, and

$$\mathbb{E}[y_1 \mid y_1 \leq v_1] = \mathbb{E}[v_2 \mid v_2 \leq v_1] = \frac{v_1}{2}.$$

- Thus

$$m^{II}(v) = \frac{v^2}{2}.$$

- If, instead, there are N bidders with valuations distributed over $[0, 1]$,

$$G(v_1) = (v_1)^{N-1}$$
$$\mathbb{E}[y_1 \mid y_1 \leq v_1] = \frac{N-1}{N} v_1,$$

and thus

$$m^{II}(v) = \frac{N-1}{N} v^N.$$

First Price Auctions

- In a first price auction, each bidder submits a sealed bid of b_i, and given these bids, the payoffs are given by

$$U_i\left(\left(b_i, b_{-i}\right), v_i\right) = \left\{ \begin{array}{cl} v_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i > \max_{j \neq i} b_j. \end{array} \right.$$

- Tie-breaking is similar to before.
- In a first price auction, the equilibrium behavior is more complicated than in a second-price auction.
- Clearly, bidding truthfully is not optimal (why not?).
  - Trade-off between higher bids and lower bids.
- So we have to work out more complicated strategies.

First Price Auctions (continued)

- Approach: look for a **symmetric (increasing and differentiable) equilibrium**.
- Suppose that bidders $j \neq 1$ follow the symmetric and differentiable equilibrium strategy $\beta^I = \beta$, where

$$\beta_i : [0, \bar{v}] \to \mathbb{R}_+.$$

- We also assume that β is increasing.
- We will then allow player 1 to use strategy β_1 and then characterize β such that when all other players play β, β is a best response for player 1. Since player 1 was arbitrary, this will complete the characterization of equilibrium.
- Suppose that bidder 1 value is v_1 and he bids the amount b (i.e., $\beta(v_1) = b$).

First Price Auctions (continued)

- First, note that a bidder with value 0 would never submit a positive bid, so

$$\beta(0) = 0.$$

- Next, note that bidder 1 wins the auction whenever $\max_{i \neq 1} \beta(v_i) < b$.

- Since $\beta(\cdot)$ is increasing, we have

$$\max_{i \neq 1} \beta(v_i) = \beta(\max_{i \neq 1} v_i) = \beta(y_1),$$

where recall that

$$y_1 = \max\{v_2, \ldots, v_N\}.$$

- This implies that bidder 1 wins whenever $y_1 < \beta^{-1}(b)$.

First Price Auctions (continued)

- Consequently, we can find an optimal bid of bidder 1, with valuation $v_1 = v$, as the solution to the maximization problem

$$\max_{b \geq 0} G(\beta^{-1}(b))(v - b).$$

- The first-order (necessary) conditions imply

$$\frac{g(\beta^{-1}(b))}{\beta'\left(\beta^{-1}(b)\right)}(v - b) - G(\beta^{-1}(b)) = 0, \qquad (*)$$

where $g = G'$ is the probability density function of the random variable y_1. [Recall that the derivative of $\beta^{-1}(b)$ is $1/\beta'\left(\beta^{-1}(b)\right)$].

- This is a first-order differential equation, which we can in general solve.

First Price Auctions (continued)

- More explicitly, a symmetric equilibrium, we have $\beta(v) = b$, and therefore $(*)$ yields

$$G(v)\beta'(v) + g(v)\beta(v) = vg(v).$$

- Equivalently, the first-order differential equation is

$$\frac{d}{dv}\Big(G(v)\beta(v)\Big) = vg(v),$$

with boundary condition $\beta(0) = 0$.

- We can rewrite this as the following optimal bidding strategy

$$\beta(v) = \frac{1}{G(v)} \int_0^v yg(y)dy = \mathbb{E}[y_1 \mid y_1 < v].$$

- Note, however, that we skipped one additional step in the argument: the first-order conditions are only necessary, so one needs to show sufficiency to complete the proof that the strategy $\beta(v) = \mathbb{E}[y_1 \mid y_1 < v]$ is optimal.

First Price Auctions (continued)

- This detail notwithstanding, we have:

Proposition

In the first price auction, there exists a symmetric equilibrium given by

$$\beta^I(v) = \mathbb{E}[y_1 \mid y_1 < v].$$

First Price Auctions: Payments and Revenues

- In general, expected payment of a bidder with value v in a first price auction is given by

$$
\begin{aligned}
m^I(v) &= \Pr(v \text{ wins}) \times \beta(v) \\
&= G(v) \times \mathbb{E}[y_1 \mid y_1 < v]. \qquad \textit{Payment I}
\end{aligned}
$$

- This can be directly compared to (Payment II), which was the payment in the second price auction
  $(m^{II}(v) = G(v) \times \mathbb{E}[y_1 \mid y_1 \leq v])$.

- This establishes the somewhat surprising results that
  $m^I(v) = m^{II}(v)$, i.e., both auction formats yield the same expected revenue to the seller.

First Price Auctions: Uniform Distribution

- As an illustration, assume that values are uniformly distributed over $[0, 1]$.
- Then, we can verify that

$$\beta^I(v) = \frac{N-1}{N} v.$$

- Moreover, since $G(v_1) = (v_1)^{N-1}$, we again have

$$m^I(v) = \frac{N-1}{N} v^N.$$

Revenue Equivalence

- In fact, the previous result is a simple case of a more general theorem.
- Consider any **standard** auction, in which buyers submit bids and the object is given to the bidder with the highest bid.
- Suppose that values are independent and identically distributed and that all bidders are risk neutral. Then, we have the following theorem:

Theorem

Any symmetric and increasing equilibria of any standard auction (such that the expected payment of a bidder with value 0 is 0) yields the same expected revenue to the seller.

Sketch Proof

- Consider a standard auction A and a symmetric equilibrium β of A.
- Let $m^A(v)$ denote the equilibrium expected payment in auction A by a bidder with value v.
- Assume that β is such that $\beta(0) = 0$.
- Consider a particular bidder, say bidder 1, and suppose that other bidders are following the equilibrium strategy β.
- Consider the expected payoff of bidder 1 with value v when he bids $b = \beta(z)$ instead of $\beta(v)$,

$$U^A(z, v) = G(z)v - m^A(z).$$

- Maximizing the preceding with respect to z yields

$$\frac{\partial}{\partial z} U^A(z, v) = g(z)v - \frac{d}{dz} m^A(z) = 0.$$

Sketch Proof (continued)

- An equilibrium will involve $z = v$ (Why?) Hence,

$$\frac{d}{dy} m^A(y) = g(y)y \qquad \text{for all } y,$$

implying that

$$m^A(v) = \int_0^v yg(y)dy = G(v) \times \mathbb{E}[y_1 \mid y_1 < v],$$

(General Payment)

establishing that the expected revenue of the seller is the same independent of the particular auction format.

- (General Payment), not surprisingly, has the same form as (Payment I) and (Payment II).

6.254 Game Theory with Engineering Applications
Spring 2010