R. Gibbons
A Primer in Game Theory

Chapter 3

Static Games of Incomplete Information

This chapter begins our study of games of *incomplete information*, also called *Bayesian games*. Recall that in a game of complete information the players' payoff functions are common knowledge. In a game of incomplete information, in contrast, at least one player is uncertain about another player's payoff function. One common example of a static game of incomplete information is a sealed-bid auction: each bidder knows his or her own valuation for the good being sold but does not know any other bidder's valuation; bids are submitted in sealed envelopes, so the players' moves can be thought of as simultaneous. Most economically interesting Bayesian games, however, are dynamic. As we will see in Chapter 4, the existence of private information leads naturally to attempts by informed parties to communicate (or mislead) and to attempts by uninformed parties to learn and respond. These are inherently dynamic issues.

In Section 3.1 we define the normal-form representation of a static Bayesian game and a Bayesian Nash equilibrium in such a game. Since these definitions are abstract and a bit complex, we introduce the main ideas with a simple example—Cournot competition under asymmetric information.

In Section 3.2 we consider three applications. First, we provide a formal discussion of the interpretation of a mixed strategy given in Chapter 1: player j's mixed strategy represents player i's uncertainty about j's choice of a pure strategy, and j's choice depends on the realization of a small amount of private information. Second, we analyze a sealed-bid auction in which the bidders' valuations are private information but the seller's valuation is known. Finally, we consider the case in which a buyer and a seller each have private information about their valuations (as when a firm knows a worker's marginal product and the worker knows his or her outside opportunity). We analyze a trading game called a double auction: the seller names an asking price and the buyer simultaneously names an offer price; trade occurs at the average of the two prices if the latter exceeds the former.

In Section 3.3 we state and prove the *Revelation Principle*, and briefly suggest how it can be applied in designing games when the players have private information.

3.1 Theory: Static Bayesian Games and Bayesian Nash Equilibrium

3.1.A An Example: Cournot Competition under Asymmetric Information

Consider a Cournot duopoly model with inverse demand given by $P(Q) = a - Q$, where $Q = q_1 + q_2$ is the aggregate quantity on the market. Firm 1's cost function is $C_1(q_1) = cq_1$. Firm 2's cost function, however, is $C_2(q_2) = c_H q_2$ with probability θ and $C_2(q_2) = c_L q_2$ with probability $1 - \theta$, where $c_L < c_H$. Furthermore, information is asymmetric: firm 2 knows its cost function and firm 1's, but firm 1 knows its cost function and only that firm 2's marginal cost is c_H with probability θ and c_L with probability $1-\theta$. (Firm 2 could be a new entrant to the industry, or could have just invented a new technology.) All of this is common knowledge: firm 1 knows that firm 2 has superior information, firm 2 knows that firm 1 knows this, and so on.

Naturally, firm 2 may want to choose a different (and presumably lower) quantity if its marginal cost is high than if it is low. Firm 1, for its part, should anticipate that firm 2 may tailor its quantity to its cost in this way. Let $q_2^*(c_H)$ and $q_2^*(c_L)$ denote firm 2's quantity choices as a function of its cost, and let q_1^* denote firm 1's single quantity choice. If firm 2's cost is high, it will

choose $q_2^*(c_H)$ to solve

$$\max_{q_2} \; [(a - q_1^* - q_2) - c_H]q_2.$$

Similarly, if firm 2's cost is low, $q_2^*(c_L)$ will solve

$$\max_{q_2} \; [(a - q_1^* - q_2) - c_L]q_2.$$

Finally, firm 1 knows that firm 2's cost is high with probability θ and should anticipate that firm 2's quantity choice will be $q_2^*(c_H)$ or $q_2^*(c_L)$, depending on firm 2's cost. Thus, firm 1 chooses q_1^* to solve

$$\max_{q_1} \; \theta[(a - q_1 - q_2^*(c_H)) - c]q_1 + (1-\theta)[(a - q_1 - q_2^*(c_L)) - c]q_1$$

so as to maximize expected profit.

The first-order conditions for these three optimization problems are

$$q_2^*(c_H) = \frac{a - q_1^* - c_H}{2},$$

$$q_2^*(c_L) = \frac{a - q_1^* - c_L}{2},$$

and

$$q_1^* = \frac{\theta[a - q_2^*(c_H) - c] + (1-\theta)[(a - q_2^*(c_L) - c]}{2}.$$

Assume that these first-order conditions characterize the solutions to the earlier optimization problems. (Recall from Problem 1.6 that in a complete-information Cournot duopoly, if the firms' costs are sufficiently different then in equilibrium the high-cost firm produces nothing. As an exercise, find a sufficient condition to rule out the analogous problems here.) The solutions to the three first-order conditions are

$$q_2^*(c_H) = \frac{a - 2c_H + c}{3} + \frac{1-\theta}{6}(c_H - c_L),$$

$$q_2^*(c_L) = \frac{a - 2c_L + c}{3} - \frac{\theta}{6}(c_H - c_L),$$

and

$$q_1^* = \frac{a - 2c + \theta c_H + (1-\theta)c_L}{3}.$$

Compare $q_2^*(c_H)$, $q_2^*(c_L)$, and q_1^* to the Cournot equilibrium under *complete* information with costs c_1 and c_2. Assuming that the values of c_1 and c_2 are such that both firms' equilibrium quantities are both positive, firm i produces $q_i^* = (a - 2c_i + c_j)/3$ in this complete-information case. In the incomplete-information case, in contrast $q_2^*(c_H)$ is greater than $(a - 2c_H + c)/3$ and $q_2^*(c_L)$ is less than $(a - 2c_L + c)/3$. This occurs because firm 2 not only tailors its quantity to its cost but also responds to the fact that firm 1 cannot do so. If firm 2's cost is high, for example, it produces less because its cost is high but also produces more because it knows that firm 1 will produce a quantity that maximizes its expected profit and thus is smaller than firm 1 would produce if it knew firm 2's cost to be high. (A potentially misleading feature of this example is that q_1^* exactly equals the expectation of the Cournot quantities firm 1 would produce in the two corresponding games of complete information. This is typically not true; consider the case in which firm i's total cost is $c_i q_i^2$, for example.)

3.1.B Normal-Form Representation of Static Bayesian Games

Recall that the normal-form representation of an n-player game of *complete* information is $G = \{S_1 \ldots S_n; u_1 \ldots u_n\}$, where S_i is player i's strategy space and $u_i(s_1, \ldots, s_n)$ is player i's payoff when the players choose the strategies $(s_1, \ldots, s_n)$. As discussed in Section 2.3.B, however, in a simultaneous-move game of complete information a strategy for a player is simply an action, so we can write $G = \{A_1 \ldots A_n; u_1 \ldots u_n\}$, where A_i is player i's action space and $u_i(a_1, \ldots, a_n)$ is player i's payoff when the players choose the actions $(a_1, \ldots, a_n)$. To prepare for our description of the timing of a static game of *incomplete* information, we describe the timing of a static game of *complete* information as follows: (1) the players simultaneously choose actions (player i chooses a_i from the feasible set A_i), and then (2) payoffs $u_i(a_1, \ldots, a_n)$ are received.

We now want to develop the normal-form representation of a simultaneous-move game of incomplete information, also called a static Bayesian game. The first step is to represent the idea that each player knows his or her own payoff function but may be uncertain about the other players' payoff functions. Let player i's possible payoff functions be represented by $u_i(a_1, \ldots, a_n; t_i)$, where

t_i is called player i's *type* and belongs to a set of possible types (or *type space*) T_i. Each type t_i corresponds to a different payoff function that player i might have.

As an abstract example suppose player i has two possible payoff functions. We would say that player i has two types, t_{i1} and t_{i2}, that player i's type space is $T_i = \{t_{i1} t_{i2}\}$, and that player i's two payoff functions are $u_i(a_1, \ldots, a_n; t_{i1})$ and $u_i(a_1, \ldots, a_n; t_{i2})$. We can use the idea that each of a player's types corresponds to a different payoff function the player might have to represent the possibility that the player might have different sets of feasible actions, as follows. Suppose, for example, that player i's set of feasible actions is $\{a, b\}$ with probability q and $\{a, b, c\}$ with probability $1-q$. Then we can say that i has two types (t_{i1} and t_{i2}, where the probability of t_{i1} is q) and we can define i's feasible set of actions to be $\{a, b, c\}$ for both types but define the payoff from taking action c to be $-\infty$ for type t_{i1}.

As a more concrete example, consider the Cournot game in the previous section. The firms' actions are their quantity choices, q_1 and q_2. Firm 2 has two possible cost functions and thus two possible profit or payoff functions:

$$\pi_2(q_1, q_2; c_L) = [(a - q_1 - q_2) - c_L]q_2$$

and

$$\pi_2(q_1, q_2; c_H) = [(a - q_1 - q_2) - c_H]q_2.$$

Firm 1 has only one possible payoff function:

$$\pi_1(q_1, q_2; c) = [(a - q_1 - q_2) - c]q_1.$$

We say that firm 2's type space is $T_2 = \{c_L, c_H\}$ and that firm 1's type space is $T_1 = \{c\}$.

Given this definition of a player's type, saying that player i knows his or her own payoff function is equivalent to saying that player i knows his or her type. Likewise, saying that player i may be uncertain about the other players' payoff functions is equivalent to saying that player i may be uncertain about the types of the other players, denoted by $t_{-i} = (t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n)$. We use T_{-i} to denote the set of all possible values of t_{-i}, and we use the probability distribution $p_i(t_{-i} \mid t_i)$ to denote player i's *belief* about the other players' types, t_{-i}, given player i's knowledge of his or her own type, t_i. In every application analyzed in Section 3.2

(and in most of the literature), the players' types are independent, in which case $p_i(t_{-i} \mid t_i)$ does not depend on t_i, so we can write player i's belief as $p_i(t_{-i})$. There are contexts in which the players' types are correlated, however, so we allow for this in our definition of a static Bayesian game by writing player i's belief as $p_i(t_{-i} \mid t_i)$.[1]

Joining the new concepts of types and beliefs with the familiar elements of the normal-form representation of a static game of complete information yields the normal-form representation of a static Bayesian game.

Definition *The* ***normal-form representation*** *of an n-player static Bayesian game specifies the players' action spaces $A_1, \ldots, A_n$, their type spaces $T_1, \ldots, T_n$, their beliefs $p_1, \ldots, p_n$, and their payoff functions $u_1, \ldots, u_n$. Player i's* ***type****, t_i, is privately known by player i, determines player i's payoff function, $u_i(a_1, \ldots, a_n; t_i)$, and is a member of the set of possible types, T_i. Player i's* ***belief*** *$p_i(t_{-i} \mid t_i)$ describes i's uncertainty about the $n-1$ other players' possible types, t_{-i}, given i's own type, t_i. We denote this game by $G = \{A_1, \ldots, A_n; T_1, \ldots, T_n; p_1, \ldots, p_n; u_1, \ldots, u_n\}$.*

Following Harsanyi (1967), we will assume that the timing of a static Bayesian game is as follows: (1) nature draws a type vector $t = (t_1, \ldots, t_n)$, where t_i is drawn from the set of possible types T_i; (2) nature reveals t_i to player i but not to any other player; (3) the players simultaneously choose actions, player i choosing a_i from the feasible set A_i; and then (4) payoffs $u_i(a_1, \ldots, a_n; t_i)$ are received. By introducing the fictional moves by nature in steps (1) and (2), we have described a game of *incomplete* information as a game of *imperfect* information, where by imperfect information we mean (as in Chapter 2) that at some move in the game the player with the move does not know the complete history of the game thus far. Here, because nature reveals player i's type to player i but not to player j in step (2), player j does not know the complete history of the game when actions are chosen in step (3).

Two slightly more technical points need to be covered to complete our discussion of normal-form representations of static Bayesian games. First there are games in which player i has private information not only about his or her own payoff function but also about another player's payoff function. In Problem 3.2, for example, the asymmetric-information Cournot model from Section 3.1.A is changed so that costs are symmetric and common knowledge but one firm knows the level of demand and the other does not. Since the level of demand affects both players' payoff functions, the informed firm's type enters the uninformed firm's payoff function. In the n-player case we capture this possibility by allowing player i's payoff to depend not only on the actions $(a_1, \ldots, a_n)$ but also on all the types $(t_1, \ldots, t_n)$. We write this payoff as $u_i(a_1, \ldots, a_n; t_1, \ldots, t_n)$.

The second technical point involves the beliefs, $p_i(t_{-i} \mid t_i)$. We will assume that it is common knowledge that in step (1) of the timing of a static Bayesian game, nature draws a type vector $t = (t_1, \ldots, t_n)$ according to the prior probability distribution $p(t)$. When nature then reveals t_i to player i, he or she can compute the belief $p_i(t_{-i} \mid t_i)$ using Bayes' rule:[2]

$$p_i(t_{-i} \mid t_i) = \frac{p(t_{-i}, t_i)}{p(t_i)} = \frac{p(t_{-i}, t_i)}{\sum_{t_{-i} \in T_{-i}} p(t_{-i}, t_i)}.$$

Furthermore, the other players can compute the various beliefs that player i might hold, depending on i's type, namely $p_i(t_{-i} \mid t_i)$ for each t_i in T_i. As already noted, we will frequently assume that the players' types are independent, in which case $p_i(t_{-i})$ does not depend on t_i but is still derived from the prior distribution $p(t)$. In this case the other players know i's belief about their types.

3.1.C Definition of Bayesian Nash Equilibrium

We now want to define an equilibrium concept for static Bayesian games. To do so, we must first define the players' strategy spaces

[1]Imagine that two firms are racing to develop a new technology. Each firm's chance of success depends in part on how difficult the technology is to develop, which is not known. Each firm knows only whether it has succeeded and not whether the other has. If firm 1 has succeeded, however, then it is more likely that the technology is easy to develop and so also more likely that firm 2 has succeeded. Thus, firm 1's belief about firm 2's type depends on firm 1's knowledge of its own type.

[2]Bayes' rule provides a formula for $P(A \mid B)$, the (conditional) probability that an event A will occur given that an event B has already occurred. Let $P(A)$, $P(B)$, and $P(A,B)$ be the (prior) probabilities (i.e., the probabilities before either A or B has had a chance to take place) that A will occur, that B will occur, and that both A and B will occur, respectively. Bayes' rule states that $P(A \mid B) = P(A,B)/P(B)$. That is, the conditional probability of A given B equals the probability that both A and B will occur, divided by the prior probability that B will occur.

are all finite sets) there exists a Bayesian Nash equilibrium, perhaps in mixed strategies. The proof closely parallels the proof of the existence of a mixed-strategy Nash equilibrium in finite games of complete information, and so is omitted here.

3.2 Applications

3.2.A Mixed Strategies Revisited

As we mentioned in Section 1.3.A, Harsanyi (1973) suggested that player j's mixed strategy represents player i's uncertainty about j's choice of a pure strategy, and that j's choice in turn depends on the realization of a small amount of private information. We can now give a more precise statement of this idea: a mixed-strategy Nash equilibrium in a game of complete information can (almost always) be interpreted as a pure-strategy Bayesian Nash equilibrium in a closely related game with a little bit of incomplete information. (We will ignore the rare cases in which such an interpretation is not possible.) Put more evocatively, the crucial feature of a mixed-strategy Nash equilibrium is not that player j chooses a strategy randomly, but rather that player i is uncertain about player j's choice; this uncertainty can arise either because of randomization or (more plausibly) because of a little incomplete information, as in the following example.

Recall that in the Battle of the Sexes there are two pure-strategy Nash equilibria (Opera, Opera) and (Fight, Fight) and a mixed-strategy Nash equilibrium in which Chris plays Opera with probability 2/3 and Pat plays Fight with probability 2/3.

		Pat: Opera	Pat: Fight
Chris	Opera	2,1	0,0
	Fight	0,0	1,2

The Battle of the Sexes

Now suppose that, although they have known each other for quite some time, Chris and Pat are not quite sure of each other's payoffs. In particular, suppose that: Chris's payoff if both attend the Opera is $2 + t_c$, where t_c is privately known by Chris; Pat's payoff if both attend the Fight is $2 + t_p$, where t_p is privately known by Pat; and t_c and t_p are independent draws from a uniform distribution on $[0, x]$. (The choice of a uniform distribution on $[0, x]$ is not important, but we do have in mind that the values of t_c and t_p only slightly perturb the payoffs in the original game, so think of x as small.) All the other payoffs are the same. In terms of the abstract static Bayesian game in normal form $G = \{A_c, A_p; T_c, T_p; p_c, p_p; u_c, u_p\}$, the action spaces are $A_c = A_p = \{\text{Opera, Fight}\}$, the type spaces are $T_c = T_p = [0, x]$, the beliefs are $p_c(t_p) = p_p(t_c) = 1/x$ for all t_c and t_p, and the payoffs are as follows.

		Pat: Opera	Pat: Fight
Chris	Opera	$2 + t_c, 1$	0,0
	Fight	0,0	$1, 2 + t_p$

The Battle of the Sexes with Incomplete Information

We will construct a pure-strategy Bayesian Nash equilibrium of this incomplete-information version of the Battle of the Sexes in which Chris plays Opera if t_c exceeds a critical value, c, and plays Fight otherwise and Pat plays Fight if t_p exceeds a critical value, p, and plays Opera otherwise. In such an equilibrium, Chris plays Opera with probability $(x - c)/x$ and Pat plays Fight with probability $(x - p)/x$. We will show that as the incomplete information disappears (i.e., as x approaches zero), the players' behavior in this pure-strategy Bayesian Nash equilibrium approaches their behavior in the mixed-strategy Nash equilibrium in the original game of complete information. That is, both $(x - c)/x$ and $(x - p)/x$ approach 2/3 as x approaches zero.

Suppose Chris and Pat play the strategies just described. For a given value of x, we will determine values of c and p such that these strategies are a Bayesian Nash equilibrium. Given Pat's strategy, Chris's expected payoffs from playing Opera and from playing Fight are

$$\frac{p}{x}(2+t_c) + \left[1 - \frac{p}{x}\right] \cdot 0 = \frac{p}{x}(2+t_c)$$

and

$$\frac{p}{x} \cdot 0 + \left[1 - \frac{p}{x}\right] \cdot 1 = 1 - \frac{p}{x},$$

respectively. Thus playing Opera is optimal if and only if

$$t_c \geq \frac{x}{p} - 3 = c. \tag{3.2.1}$$

Similarly, given Chris's strategy, Pat's expected payoffs from playing Fight and from playing Opera are

$$\left[1 - \frac{c}{x}\right] \cdot 0 + \frac{c}{x}(2+t_p) = \frac{c}{x}(2+t_p)$$

and

$$\left[1 - \frac{c}{x}\right] \cdot 1 + \frac{c}{x} \cdot 0 = 1 - \frac{c}{x},$$

respectively. Thus, playing Fight is optimal if and only if

$$t_p \geq \frac{x}{c} - 3 = p. \tag{3.2.2}$$

Solving (3.2.1) and (3.2.2) simultaneously yields $p = c$ and $p^2 + 3p - x = 0$. Solving the quadratic then shows that the probability that Chris plays Opera, namely $(x-c)/x$, and the probability that Pat plays Fight, namely $(x-p)/x$, both equal

$$1 - \frac{-3 + \sqrt{9+4x}}{2x},$$

which approaches 2/3 as x approaches zero. Thus, as the incomplete information disappears, the players' behavior in this pure-strategy Bayesian Nash equilibrium of the incomplete-information game approaches their behavior in the mixed-strategy Nash equilibrium in the original game of complete information.

3.2.B An Auction

Consider the following first-price, sealed-bid auction. There are two bidders, labeled $i = 1, 2$. Bidder i has a valuation v_i for the good—that is, if bidder i gets the good and pays the price p, then i's payoff is $v_i - p$. The two bidders' valuations are independently and uniformly distributed on $[0, 1]$. Bids are constrained to be nonnegative. The bidders simultaneously submit their bids. The higher bidder wins the good and pays the price she bid; the other bidder gets and pays nothing. In case of a tie, the winner is determined by a flip of a coin. The bidders are risk-neutral. All of this is common knowledge.

In order to formulate this problem as a static Bayesian game, we must identify the action spaces, the type spaces, the beliefs, and the payoff functions. Player i's action is to submit a (nonnegative) bid, b_i, and her type is her valuation, v_i. (In terms of the abstract game $G = \{A_1, A_2; T_1, T_2; p_1, p_2; u_1, u_2\}$, the action space is $A_i = [0, \infty)$ and the type space is $T_i = [0, 1]$.) Because the valuations are independent, player i believes that v_j is uniformly distributed on $[0, 1]$, no matter what the value of v_i. Finally, player i's payoff function is

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

To derive a Bayesian Nash equilibrium of this game, we begin by constructing the players' strategy spaces. Recall that in a static Bayesian game, a strategy is a function from types to actions. Thus, a strategy for player i is a function $b_i(v_i)$ specifying the bid that each of i's types (i.e., valuations) would choose. In a Bayesian Nash equilibrium, player 1's strategy $b_1(v_1)$ is a best response to player 2's strategy $b_2(v_2)$, and vice versa. Formally, the pair of strategies $(b(v_1), b(v_2))$ is a Bayesian Nash equilibrium if for each v_i in $[0, 1]$, $b_i(v_i)$ solves

$$\max_{b_i} \ (v_i - b_i)\text{Prob}\{b_i > b_j(v_j)\} + \frac{1}{2}(v_i - b_i)\text{Prob}\{b_i = b_j(v_j)\}.$$

We simplify the exposition by looking for a linear equilibrium: $b_1(v_1) = a_1 + c_1 v_1$ and $b_2(v_2) = a_2 + c_2 v_2$. Note well that we are *not* restricting the players' strategy spaces to include only linear strategies. Rather, we allow the players to choose arbitrary strategies but ask whether there is an equilibrium that is linear. It turns out that because the players' valuations are uniformly distributed, a linear equilibrium not only exists but is unique (in a sense to be made precise). We will find that $b_i(v_i) = v_i/2$. That is, each player submits a bid equal to half her valuation. Such a bid reflects the fundamental trade-off a bidder faces in an auction: the higher the bid, the more likely the bidder is to win; the lower the bid, the larger the gain if the bidder does win.

Suppose that player j adopts the strategy $b_j(v_j) = a_j + c_j v_j$. For a given value of v_i, player i's best response solves

$$\max_{b_i} \,(v_i - b_i)\text{Prob}\{b_i > a_j + c_j v_j\},$$

where we have used the fact that $\text{Prob}\{b_i = b_j(v_j)\} = 0$ (because $b_j(v_j) = a_j + c_j v_j$ and v_j is uniformly distributed, so b_j is uniformly distributed). Since it is pointless for player i to bid below player j's minimum bid and foolish for i to bid above j's maximum bid, we have $a_j \le b_i \le a_j + c_j$, so

$$\text{Prob}\{b_i > a_j + c_j v_j\} = \text{Prob}\left\{v_j < \frac{b_i - a_j}{c_j}\right\} = \frac{b_i - a_j}{c_j}.$$

Player i's best response is therefore

$$b_i(v_i) = \begin{cases} (v_i + a_j)/2 & \text{if } v_i \ge a_j, \\ a_j & \text{if } v_i < a_j. \end{cases}$$

If $0 < a_j < 1$ then there are some values of v_i such that $v_i < a_j$, in which case $b_i(v_i)$ is not linear; rather, it is flat at first and positively sloped later. Since we are looking for a linear equilibrium, we therefore rule out $0 < a_j < 1$, focusing instead on $a_j \ge 1$ and $a_j \le 0$. But the former cannot occur in equilibrium: since it is optimal for a higher type to bid at least as much as a lower type's optimal bid, we have $c_j \ge 0$, but then $a_j \ge 1$ would imply that $b_j(v_j) \ge v_j$, which cannot be optimal. Thus, if $b_i(v_i)$ is to be linear, then we must have $a_j \le 0$, in which case $b_i(v_i) = (v_i + a_j)/2$, so $a_i = a_j/2$ and $c_i = 1/2$.

We can repeat the same analysis for player j under the assumption that player i adopts the strategy $b_i(v_i) = a_i + c_i v_i$. This yields $a_i \le 0$, $a_j = a_i/2$, and $c_j = 1/2$. Combining these two sets of results then yields $a_i = a_j = 0$ and $c_i = c_j = 1/2$. That is, $b_i(v_i) = v_i/2$, as claimed earlier.

One might wonder whether there are other Bayesian Nash equilibria of this game, and also how equilibrium bidding changes as the distribution of the bidders' valuations changes. Neither of these questions can be answered using the technique just applied (of positing linear strategies and then deriving the coefficients that make the strategies an equilibrium): it is fruitless to try to guess all the functional forms other equilibria of this game might have, and a linear equilibrium does not exist for any other distribution of valuations. In the Appendix, we derive a symmetric Bayesian Nash equilibrium,[3] again for the case of uniformly distributed valuations. Under the assumption that the players' strategies are strictly increasing and differentiable, we show that the unique symmetric Bayesian Nash equilibrium is the linear equilibrium already derived. The technique we use can easily be extended to a broad class of valuation distributions, as well as the case of n bidders.[4]

Appendix 3.2.B

Suppose player j adopts the strategy $b(\cdot)$, and assume that $b(\cdot)$ is strictly increasing and differentiable. Then for a given value of v_i, player i's optimal bid solves

$$\max_{b_i} \,(v_i - b_i)\text{Prob}\{b_i > b(v_j)\}.$$

Let $b^{-1}(b_j)$ denote the valuation that bidder j must have in order to bid b_j. That is, $b^{-1}(b_j) = v_j$ if $b_j = b(v_j)$. Since v_j is uniformly distributed on $[0,1]$, $\text{Prob}\{b_i > b(v_j)\} = \text{Prob}\{b^{-1}(b_i) > v_j\} = b^{-1}(b_i)$. The first-order condition for player i's optimization problem is therefore

$$-b^{-1}(b_i) + (v_i - b_i)\frac{d}{db_i}b^{-1}(b_i) = 0.$$

---

[3]A Bayesian Nash equilibrium is called symmetric if the players' strategies are identical. That is, in a symmetric Bayesian Nash equilibrium, there is a single function $b(v_i)$ such that player 1's strategy $b_1(v_1)$ is $b(v_1)$ and player 2's strategy $b_2(v_2)$ is $b(v_2)$, and this single strategy is a best response to itself. Of course, since the players' valuations typically will be different, their bids typically will be different, even if both use the same strategy.

[4]Skipping this appendix will not hamper one's understanding of what follows.