# Games of Incomplete Information

Jonathan Levin

February 2002

## 1 Introduction

We now start to explore models of incomplete information. Informally, a game of *incomplete information* is a game where the players do not have common knowledge of the game being played. This idea is tremendously important in capturing many economic situations, where a variety of features of the environment may not be commonly known. Among the aspects of the game that the players might not have common knowledge of are:

- Payoffs
- Who the other players are
- What moves are possible
- How outcome depends on the action.
- What opponent knows, and what he knows I know....

To take a couple of simple examples: (1) in price or quantity competition, firms might know their own costs, but not the costs of their rivals; (2) firms investing in R&D might know how their project is coming along, but have no idea who else is working on the same problem; (3) the government may design the tax code not envisioning what ploys people will come up with to avoid taxs; (4) countries may negotiate climate change agreements having different beliefs about the costs and benefits of global climate change; (5) plaintiffs may offer settlements to defendants not knowing what sort of case the defendant will be able to bring to court, or what sort of case the defendant thinks the plaintiff will be able to bring.

# 2 Examples

## 2.1 Entry with a "Tough" Incumbent

Recall our canonical entry model where Firm 2 (the entrant) must decide whether or not to enter a market, and Firm 1 (the incumbent) must decide whether to Fight or Accomodate entry. Let's modify this game by assuming that with probability $1-p$, the incumbent is "tough," while with probability $p$, the incumbent is normal. The payoffs in the game depend on whether the incumbent is tough or normal:

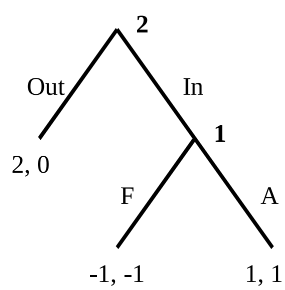


**Entry w/ Normal Incumbent**

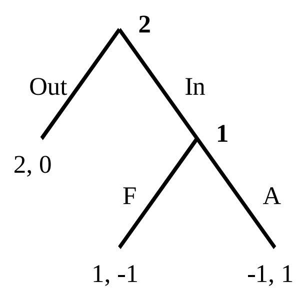


**Entry w/ Tough Incumbent**

## 2.2 Sealed Bid Auction

Two bidders are trying to purchase the same item in a sealed bid auction. The bidders simultaneously submit bids $b_1$ and $b_2$ and the auction is sold to the highest bidder at his bid price (this is called a "first price" auction). If there is a tie, there is a coin flip to determine the winner. Suppose the players utilities are:

$$u_i(b_i, b_{-i}) = \begin{cases} v_i - b_i & \text{if } b_i > b_{-i} \\ \frac{1}{2}(v_i - b_i) & \text{if } b_i = b_{-i} \\ 0 & \text{if } b_i < b_{-i} \end{cases}$$

The key informational feature is the each player knows his own value for the item (i.e. bidder $i$ knows $v_i$), but does not know the valuation of his rival. Instead, we assume that each bidder had a prior belief that his rival's valuation is a draw from a uniform distribution on $[0, 1]$, and that these prior beliefs are common knowledge.

### 2.3 Public Good Provision

Two faculty members in the economics department both want to recruit a top graduate student to their department. Either faculty member can ensure the student will accept the offer by getting on the phone and shamelessly promoting the graduate program. However, there is some cost to making this call. Assume the payoffs can be represented as follows:

| | Call | Don't |
|---|---|---|
| Call | $1-c_1, 1-c_2$ | $1-c_2, 1$ |
| Don't | $1, 1-c_1$ | $0, 0$ |

Assume the faculty members choose their actions simultaneously (or at least without learning the other's action) and the faculty have private information about their costs of making the call. That is, faculty member $i$ knows $c_i$, and believes that $c_j$ is a random draw from a uniform distribution on $[\underline{c}, \overline{c}]$. Faculty member $i$'s belief about $c_j$ is commonly known.

An alternative formulation of this problem is that faculty member one tells everyone everything, so that everyone knows his cost $c_1 = 1/2$. However, $c_2 \in \{\underline{c}, \overline{c}\}$ is known only to player two.

Or, we could assume that player one is a senior faculty member who knows from long experience that $c_1 = 1/2$ and $c_2 = 2/3$. However, player two is a new assistant professor whose prior belief is that $c_1, c_2 \sim U[0, 2]$ and are independent.

## 3 Definitions

**Definition 1** *A game with incomplete information $G = (\Theta, S, P, u)$ consists of:*

1. *A set $\Theta = \Theta_1 \times ... \times \Theta_I$, where $\Theta_i$ is the (finite) set of possible types for player $i$.*

2. *A set $S = S_1 \times ... \times S_I$, where $S_i$ is the set of possible strategies for player $i$.*

3. *A joint probability distribution $p(\theta_1, ..., \theta_I)$ over types. For finite type space, assume that $p(\theta_i) > 0$ for all $\theta_i \in \Theta_i$.*

4. *Payoff functions $u_i : S \times \Theta \to \mathbb{R}$.*

Consider how this definition relates to each of our examples.

1. Entry: $\Theta_1 = \{\text{tough, normal}\}$; $\Theta_2 = \{\text{normal}\}$.
2. Auction: $\Theta_1 = \Theta_2 = [0, 1]$.
3. Public Good: $\Theta_1 = \Theta_2 = [\underline{c}, \overline{c}]$.

We assume that players know their own types, but do not know the types of other players.

**Remark 1** *Note that payoffs can depend not only on your own type, but on your rivals' types. If $u_i$ depends on $\theta_i$, but not on $\theta_{-i}$, we sometimes say the game has private values.*

In order to analyze these types of games, we rely on a fundamental (and Nobel-prize winning) observation by Harsanyi (1968):

> Games of incomplete information can be thought of as games of complete but imperfect information where nature makes the first move (selecting $\theta_1, ..., \theta_I$), but not everyone observes nature's move (i.e. player $i$ learns $\theta_i$ but not $\theta_{-i}$).

Consider formulating the entry model in exactly this way.

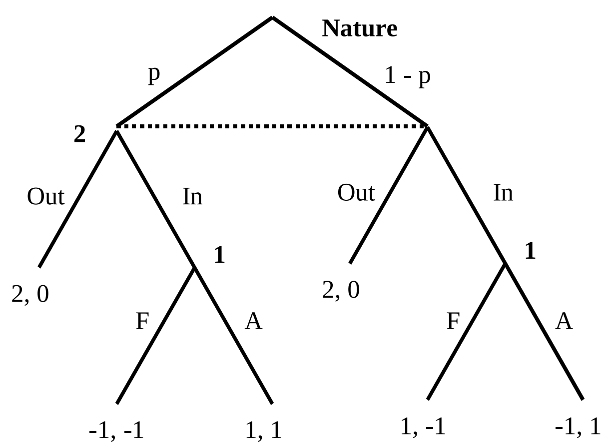


In analyzing this sort of game, we cam think of nature simply as another player. The only difference is that rather than maximizing a payoff, nature just uses a fixed mixed strategy.

This observation should make the following definitions look obvious. They just say that to analyze a game of incomplete information, we can look at the Nash Equilibrium of the game where nature is a player.

**Definition 2** *A Bayesian pure strategy for player $i$ in $G$ is a function $f_i : \Theta_i \to \Sigma_i$. Write $S^{\Theta_i}$ for the set of Bayesian pure strategies.*

**Definition 3** *A Bayesian strategy profile $(f_1, ..., f_I)$ is a* ***Bayesian Nash Equilibrium*** *if for all $i$,*

$$f_i \in \arg\max_{f_i' \in S_i^{\Theta_i}} \sum_{\theta \in \Theta} u_i \left( f_i' (\theta_i), f_{-i} (\theta_{-i}), \theta_i, \theta_{-i} \right) p (\theta_i, \theta_{-i})$$

*or alternatively, for all $i$, $\theta_i$ and $s_i$ :*

$$\begin{aligned} & \sum_{\theta_{-i} \in \Theta_{-i}} u_i (f_i (\theta_i), f_{-i} (\theta_{-i}), \theta_i, \theta_{-i}) p (\theta_{-i} | \theta_i) \\ \geq & \sum_{\theta_{-i} \in \Theta_{-i}} u_i (s_i, f_{-i} (\theta_{-i}), \theta_i, \theta_{-i}) p (\theta_{-i} | \theta_i) \end{aligned}$$

The second part of the definition just says that in order to maximize your expected payoff given that you know your types, then the strategy you choose for each type should maximize your payoff conditional on your having that type.

**Remark 2** *A Bayesian Nash Equilibrium is simply a Nash Equilibrium of the game where Nature moves first, chooses $\theta \in \Theta$ from a distribution with probability $p(\theta)$ and reveals $\theta_i$ to player $i$.*

# 4 Solving Bayesian Games

## 4.1 Public Good: version A

Consider a version of the public good game where

- Player 1 has a *known* cost $c_1 < 1/2$;
- Player 2 has cost $\underline{c}$ with probability $p$ and $\overline{c}$ with probability $1 - p$.

Assume that $0 < \underline{c} < 1 < \overline{c}$ and that $p < 1/2$.

**Proposition 1** *The unique Bayesian Nash Equilibrium is $f_1 = Call$ and $f_2(c) = Don't$ for all $c$.*

To prove this, keep in mind that each type of player must play a best response. When player 2 has type $\overline{c}$, then calling is strictly dominated:

$$u_2\left(s_1, Call; \overline{c}\right) < u_2\left(s_1, Don't; \overline{c}\right),$$

for all $s_1$. Thus, $f_2\left(\overline{c}\right) = Don't$.

Now, for player 1,

$$\begin{aligned}
u_1\left(Call, f_2; c_1\right) &= 1 - c_1 \\
u_1\left(Don't, f_2; c_1\right) &= p u_1\left(Don't, f_2\left(\underline{c}\right); c_1\right) \\
&\quad + (1-p)\, u_1\left(Don't, f_2\left(\overline{c}\right); c_1\right) \\
&\leq p \cdot 1 + (1-p) \cdot 0 = p.
\end{aligned}$$

Since $1 - c_1 > p$, then $f_1\left(c_1\right) = Call$.

But then when player 2 has type $\underline{c}$:

$$\begin{aligned}
u_2\left(f_1, Call; \underline{c}\right) &= 1 - \underline{c} \\
u_2\left(f_1, Don't; \underline{c}\right) &= 1
\end{aligned}$$

so $f_2\left(\underline{c}\right) = Don't$.

Note that this process works a bit like iterated dominance.

### 4.2 Public Good: version B

Now imagine that $c_1$ and $c_2$ are independent random draws from a uniform distribution on $[0, 2]$.

**Proposition 2** *The (essentially) unique Bayesian Nash Equilibrium is*

$$f_i\left(c_i\right) = \begin{cases} Call & \text{if } c_i \leq 2/3 \\ Don't & \text{if } c_i > 2/3 \end{cases}$$

To prove that this is actually a BNE is easy. We can just check that each player's conjectured strategy is a best response to the other's — in particular, that each player's strategy is a best response given that his opponent will call with probability 1/3 and won't call with probability 2/3.

To illustrate uniqueness, let's work through how to derive the equilibrium. The first observation is the following: *if* $f_i\left(c_i\right) = Call$ *then* $f_i\left(c_i'\right) = Call$ *for all* $c_i' < c_i$. To see this, note that if $f_i\left(c_i\right) = Call$, then:

$$\mathbb{E}_{c_{-i}} u_i\left(Call, f_2\left(c_{-i}\right); c_i\right) \geq \mathbb{E}_{c_{-i}} u_i\left(Don't, f_2\left(c_{-i}\right); c_i\right)$$

which implies that:

$$1 - c_i \geq z_{-i}$$

where we let $z_{-i}$ denote the payoff to not calling. This implies that for all $c_i' < c_i$ :

$$1 - c_i' > z_{-i}$$

or equivalently

$$\mathbb{E}_{c_{-i}} u_i \left( Call, f_2 \left( c_{-i} \right); c_i' \right) \geq \mathbb{E}_{c_{-i}} u_i \left( Don't, f_2 \left( c_{-i} \right); c_i' \right).$$

There is a simple intuition: namely that calling is more attractive if the costs are lower.

In light of Observation 1, a Bayesian Nash Equilibrium must be of the form:

$$f_1 \left( c_1 \right) = \begin{cases} Call & \text{if } c_1 \leq c_1^* \\ Don't & \text{if } c_1 > c_1^* \end{cases}$$

$$f_2 \left( c_2 \right) = \begin{cases} Call & \text{if } c_2 \leq c_2^* \\ Don't & \text{if } c_2 > c_2^* \end{cases}$$

for some "cut-off" costs $c_1^*, c_2^*$. (Note: it will turn out that when $c_i$ is *exactly* equal to $c_i^*$, then agent $i$ is indifferent to calling or not. This is why the equilibrium is "essentially" unique.)

Let

$$z_j = \Pr \left[ i \text{ will call given cut-off } c_j^* \right] = \Pr \left[ c_j \leq c_j^* \right] = \frac{1}{2} c_j^*.$$

For these strategies to be a BNE, we need:

$$\begin{array}{rcll} 1 - c_i & \geq & z_{-i} & \text{for all } c_i \leq c_i^* \\ 1 - c_i & < & z_{-i} & \text{for all } c_i > c_i^* \end{array}$$

Or equivalently that $1 - c_i^* = z_{-i}$. Thus, for $i = 1, 2$,

$$1 - c_i^* = \frac{1}{2} c_i^*$$

and hence the unique equilibrium is to call whenever $c_i < 2/3$.

**Remark 3** *Note that the equilibrium outcome is inefficient in several ways. First, there is "under-investment" in the public good — it is **always** efficient for someone to call, and yet with probability* $4/9$*, no one calls. Second, there is "miscoordination" — with probability* $1/9$ *both parties call even though this is inefficient.*

## 4.3 Sealed Bid Auction

**Proposition 3** *In the first price sealed bid auction with valuations uniformly distributed on* $[0,1]$*, the unique BNE is* $f_i(v_i) = v_i/2$ *for* $i = 1, 2$.

Again, to verify that this is a BNE is relatively easy. We just show that each type of each player is using a best response. Note that:

$$\begin{aligned} \mathbb{E}_{v_2} u_1(b_1, f_2; v_1, v_2) &= (v_1 - b_1) \Pr[f_2(v_2) < b_1] \\ &\quad + \frac{1}{2}(v_1 - b_1) \Pr[f_2(v_2) = b_1]. \end{aligned}$$

We assume $b_1 \in (0, 1/2]$. No large bid makes sense given $f_2$. Then:

$$\mathbb{E}_{v_2} u_1(b_1, f_2; v_1, v_2) = (v_1 - b_1) 2b_1$$

Maximizing this by choice of $b_1$, we obtain the first order condition:

$$0 = 2v_1 - 4b_1 \qquad \Rightarrow \qquad b_1 = v_1/2.$$

To show uniqueness (or to find the equilibrium if you didn't already know it) is harder. To do it, let's consider bidder one's optimization problem given $f_2$. If we *assume* that $f_2$ is strictly increasing, then ties occur with probability zero, so bidder one's problem is:

$$\max_{b_1} (v_1 - b_1) \Pr[f_2(v_2) < b_1]$$

If we further assume that $f_2$ is increasing, then we have

$$\Pr[f_2(v_2) < b_1] = \Pr\left[v_2 < f_2^{-1}(b_1)\right] = f_2^{-1}(b_1).$$

Finally if $f_2$ is differentiable, we can solve:

$$\max_{b_1} (v_1 - b_1) f_2^{-1}(b_1)$$

to obtain the first order condition:

$$0 = -f_2^{-1}(b_1) + (v_1 - b_1) \frac{1}{f_2'\left(f_2^{-1}(b_1)\right)}.$$

For $f_1, f_2$ to be an equilibrium, it must be that:

$$b_1 = f_1(v_1)$$

so the first order condition is

$$0 = -f_2^{-1}\left(f_1\left(v_1\right)\right) + \left(v_1 - f_1\left(v_1\right)\right)\frac{1}{f_2'\left(f_2^{-1}\left(f_1\left(v_1\right)\right)\right)}.$$

Finally, assuming we can look for a symmetric equilibrium with $f_1 = f_2 = f$, the first order condition for optimality becomes:

$$0 = -v_1 + \left(v_1 - f_1\left(v_1\right)\right)\frac{1}{f'\left(v_1\right)}.$$

Solving this differential equation gives the equilibrium.

## 4.4 A Lemons Problem

Consider a seller of a used car and a potential buyer of that car. Suppose that quality of the car, $\theta$, is a uniform draw from $[0,1]$. This quality is known to the seller, but not to the buyer. Suppose that the buyer can make an offer $p \in [0,1]$ to the seller, and the seller can then decide whether to accept or reject the buyer's offer. (Note: This sounds like a dynamic game, but we can think of it as simultaneous-move if we think of the seller as announcing the set of all prices she will accept and all those she will reject.)

Payoffs are as follows:

$$u_S = \begin{cases} p & \text{if offer accepted} \\ \theta & \text{if offer rejected} \end{cases}$$

$$u_B = \begin{cases} a + b\theta - p & \text{if offer accepted} \\ 0 & \text{if offer rejected} \end{cases}$$

Assume that $a \in [0,1)$, that $b \in (0,2)$, and that $a + b > 1$. These assumptions imply that for all $\theta$, it is more efficient for the buyer to own the car.

**Proposition 4** *The (essentially) unique BNE is for the buyer to offer $p = a/\left(2-b\right)$ and the seller to accept if and only if $p \geq \theta$.*

To prove this, we first consider the seller. It is easy to see that the strategy of accepting if and only if $p \geq \theta$ is weakly dominant for the seller. Now consider the buyer's problem.

$$\begin{aligned}
\mathbb{E}_\theta u_B\left(p; S \text{ accepts if } p \geq \theta\right) &= \int_0^1 \mathbf{1}_{\{\theta < p\}}\left(a + b\theta - p\right)d\theta \\
&= \int_0^p \left(a + b\theta - p\right)d\theta \\
&= p\left(a + \frac{1}{2}bp - p\right)
\end{aligned}$$

Choosing $p$ to maximize this expression gives the first order condition

$$0 = a + (b - 2)\, p.$$

Note that the ex ante quality of the car is $1/2$. However, given an offer $p$, the expected value of the car *conditional on the seller accepting the offer* (i.e. conditional on $\theta \leq p$) is $p/2 \leq 1/2$. This is sometime's called the *winner's curse (* here it's really a *buyer's curse*). Note that what gives rise to this effect is that the seller's information is directly payoff-relevant to the buyer. Unlike our previous examples, we do not have private values.

**Remark 4** *If $a = 0$, then the buyer's curse is so strong that the unique equilibrium is for the buyer to offer a price $p = 0$. Trade never occurs despite it the fact that there are always gains from trading.*