Week 8: Games of Incomplete Information

Dr Daniel Sgroi

Reading: Osborne chapter 9 (and 10).

With thanks to Peter J. Hammond.

Introduction

So far we have always made an important assumption:
that the game played is common knowledge.

In particular, the players know:

1. who is playing;
2. what actions or "strategies" are possible for each player;
3. how outcomes translate into payoffs.

Furthermore, this knowledge of the game is itself common knowledge.

Common knowledge underlies solution concepts such as:

- dominant strategy equilibrium;
- iterated elimination of dominated strategies;
- Nash equilibrium;
- even subgame perfect equilibrium in an extensive form.

Relaxing Common Knowledge

This common knowledge ideal excludes many interesting and more realistic models of strategic interaction.

The Cournot and Bertrand duopoly models, for instance, each lead to a clear, precise, easily understood outcome.
But is each firm's cost function really known to the opponent, let alone common knowledge?
Clearly, it is more convincing to allow some imprecision into each firm's beliefs about the opponent's costs.

Similarly, a reasonable variant of prisoner's dilemma allows the possibility that each prisoner may be subject to omertà (the Mafia code of silence or connivance).
Or whether in BoS, either player really prefers football to the opera.

How then does one model situations where players have "incomplete information" about each other's characteristics?

Extended Beliefs

A key insight is that a similar issue
arises in any simultaneous move game:
players do not know each others' actions,
but choose best responses to (probabilistic) beliefs.
Furthermore, equilibrium requires these beliefs,
and the appropriate best responses, to be consistent.

Realizing this similarity, John C. Harsanyi (1967) developed
a very elegant and practically useful way to model beliefs
not only over other players' actions,
but also over these players' other characteristics or types.
Harsanyi described a game as having incomplete information
when the players are uncertain about each other's types.

As with games of complete information,
equilibrium analysis requires these beliefs to be consistent,
meaning there is common knowledge of players' possible types,
and of the likelihood of each type profile.

Common Knowledge of Player and Action Spaces

Like a game of complete information,
a normal or strategic form game of incomplete information
retains the following two "physical" components
of an n-player simultaneous move game:

Players: the set of players $N = \{1, 2, ..., n\}$;

Actions: the action space A_i of each player $i \in N$.

We still assume common knowledge of both these items.

But we allow uncertainty about players' preferences.

Complete vs. Incomplete Information

- In a game of complete information
  each player has a single utility function
  (more precisely, one equivalence class of utility functions)
  that maps action or strategy profiles into payoffs.
- In a game of incomplete information
  each player may have one of many possible utility functions.

Players' preferences may not be common knowledge.

Player i's type determines i's preferences/payoffs.

Randomized Types

Harsanyi (1967) captured this idea by assuming that the game begins with a move by Chance which selects the different players' preferences — or more generally, their types.

The type profile determines not only players' payoffs but also their beliefs about other players' types.

Alternatively, one may think of a family of games, and which one is played is determined by chance.

An Entry Game

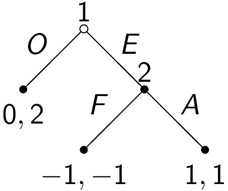

Consider a simple "entry game".

Firm 1 is a potential entrant to an industry,
which decides whether or to enter a market (E) or stay out (O).

If player 1 does enter, the incumbent firm, player 2,
responds by either fighting (F) or accommodating (A) entry.

The specified payoffs imply
that (E, A) is a unique subgame perfect equilibrium.

(Use a bimatrix game to convince yourself that (O, F) is another pure strategy Nash equilibrium that is subgame imperfect.)

An Entry Game with Two Types of Incumbent

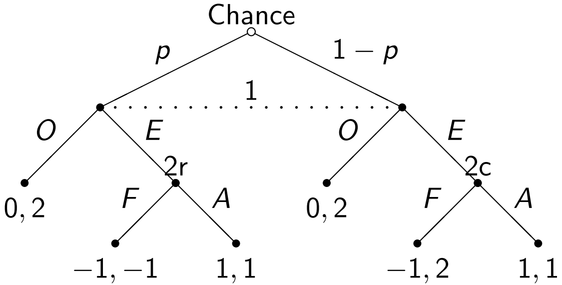

Now imagine that, whereas there is only one type of player 1, there are two types of player 2.

One "rational" type labelled $2r$ has the payoffs originally specified; the other "crazy" type labelled $2c$ enjoys fighting for its own sake, so its payoff from (Entry, Fight) is 2 instead of -1.

Chance Types and Self Knowledge

We finish specifying the game by letting p denote the exogenously specified and commonly known probability that Chance chooses "rational" (r) as player 2's type.

This information allows us to depict the extensive form of this entry game with incomplete information.

It is natural to suppose players know their own preferences, and so can find best responses given what they assume about other players' behavior.

This explains the dashed information sets in the extensive form entry game. These show that firm 2, when making its decision, knows its own preferences.

Player 1, however, remains unsure about player 2's preferences.

Type Probabilities . . .

Players know their own preferences,
but not other players' preferences or types.

In order to use some kind of equilibrium analysis
we assume players form rational conjectures or correct beliefs
about these unknown preferences and types.

These beliefs allow players to predict other players' behaviour.

More precisely, though each player
does not necessarily know other players' actual preferences,
we assume he does know the precise way
in which Chance determines these preferences.

That is, each player knows the probability distribution over types,
and this itself is common knowledge among all the players.

... Are Commonly Known

In the extensive form of the entry game,
this is represented by player 1's information set,
and by specifying that p is common knowledge.

COMMENT: Actually, only player 1's estimate of p really matters;
player 2 has a unique dominant strategy at each information set,
so its behavior is independent of p.

An important feature of the extensive form
is that player 1 has one information set, whereas player 2 has two.

This is because player 2 knows its own type,
while player 1 only has a (correct) belief about player 2's type.
This belief is specified by the probabilities p and $1 - p$
that 2's type is "rational" (r) or "crazy" (c).

The Common Prior Assumption

Following the terminology of Bayesian statistics,
which Harsanyi himself adopted,
this is often referred to as the common prior assumption.

It means that all the players, before discovering their own types,
share a common prior probabilistic theory
of how Chance chooses the players' random type profile.

The assumption is rather strong,
but essential if we are to discuss equilibrium behaviour
in any but very special models.

Type Contingent Strategies

We can translate the entry game into its normal form.
In this form, player 1 must have four pure strategies:
at each of the two information sets
there are two actions to choose from.

A strategy of player 2 can be expressed
as a pair $xy \in \{A_r A_c, A_r F_c, F_r A_c, F_r F_c\} = \{A_r, F_r\} \times \{A_c, F_c\}$,
where $x \in \{A_r, F_r\}$ describes
what player 2 does if its type is "rational" (r),
and $y \in \{A_c, F_c\}$ what it does if its type is "crazy" (c).

More generally: when information is incomplete, a player's strategy must prescribe what each type of that player should do.

(COMMENT:
Each player's strategy is contingent on that player's type.)

Of course, player 1's strategy set is simply $\{E, O\}$.

Expected Payoffs

In this two-player game, each pair of pure strategies gives rise to a path of play that starts with Nature's choice, and follows with the simultaneous actions of both players.

In this example, suppose player 2 plays $A_r F_c$
— i.e., A if "rational" (type r), but F if "crazy" (type c).

Suppose player 1 plays E. The outcome will yield payoffs of:

- $(1, 1)$ with probability p;
- $(-1, 2)$ with probability $1 - p$.

Expected payoffs from the strategy pair $(E, A_r F_c)$ are

$$\begin{array}{rrcl} & u_1 & = & p \cdot 1 + (1-p) \cdot (-1) \;=\; 2p - 1 \\ \text{and} & u_2 & = & p \cdot 1 + (1-p) \cdot 2 \;=\; 2 - p. \end{array}$$

The Entry Game in Normal Form

When $p = \frac{2}{3}$ these payoffs will be $(u_1, u_2) = (\frac{1}{3}, \frac{4}{3})$.

Similarly, the payoffs from the other seven pure strategy pairs are as indicated in the table below:

P_1 \ P_2	$A_r A_c$	$A_r F_c$	$F_r A_c$	$F_r F_c$
O	0, 2	0, 2	0, 2	0, 2
E	1, 1	$\frac{1}{3}$, $\frac{4}{3}$	$-\frac{1}{3}$, $-\frac{1}{3}$	-1, 0

This is like any other normal form game.
Its origins as a game of incomplete information are irrelevant.

Standard tools can find best responses and Nash equilibria.

Nash Equilibria in the Entry Game when $p = \frac{2}{3}$

P_2

P_1	$A_r A_c$	$A_r F_c$	$F_r A_c$	$F_r F_c$
O	$0,\ 2_*$	$0,\ 2_*$	$0^*,\ 2_*$	$0^*,\ 2_*$
E	$1^*,\ 1$	$\frac{1}{3}^*,\ \frac{4}{3}_*$	$-\frac{1}{3},\ -\frac{1}{3}$	$-1,\ 0$

There are three Nash equilibria in pure strategies: $(O, F_r A_c)$, $(O, F_r F_c)$, and $(E, A_r F_c)$.

In the extensive form, only $(E, A_r F_c)$ is subgame perfect.
Also, $A_r F_c$ weakly dominates all player 2's other strategies.
The only mixed strategy equilibria involve 1 playing O, and 2 mixing $F_r A_c$ and $F_r F_c$ arbitrarily.
(To deter entry, firm 2 must fight if "rational"!)

Nash Equilibria in the Entry Game for General p

P_2

P_1		A_rA_c	A_rF_c	F_rA_c	F_rF_c
	O	$0,\ 2$	$0,\ 2$	$0,\ 2$	$0,\ 2$
	E	$1,\ 1$	$2p-1,\ 2-p$	$1-2p,\ 1-2p$	$-1,\ 2-p$

In case $p > \frac{1}{2}$, the best responses are as for $p = \frac{2}{3}$.
The Nash equilibria are (O, F_rA_c), (O, F_rF_c), and (E, A_rF_c), of which only (E, A_rF_c) is subgame perfect in the extensive form.

But in case $p < \frac{1}{2}$, the Nash equilibria are (O, A_rF_c) and (O, F_rF_c), of which only (O, A_rF_c) is subgame perfect in the extensive form.

"Meta" Players ...

Each type of any player $i \in N$,
given specific beliefs about other players' types and actions,
can calculate expected utility separately.

But we averaged the payoffs of all i's types
using the likelihood of each type as its weight.

This is the average payoff across the different types
of some artificial "meta" player.

Will this "meta" version of player i have optimal responses
that combine all the optimal responses of i's different types?

This is precisely where
the "type contingent" definition of strategies is so useful.

. . . and Their Strategies

For example, suppose player 1 chooses E in the entry game.

When the "meta" player 2 responds with $A_r A_c$,
player 2's "rational" type is playing a best response A,
while the "crazy" type plays the dominated strategy A.

This corresponds precisely to the fact
that $A_r A_c$ is not a best response to E, but $A_r F_c$ is.

Note that $A_r F_c$ gives one ("crazy") type a higher payoff,
and the other ("rational") type the same payoff.

Thus, the type contingent strategy $A_r F_c$
gives the "meta" player a higher expected payoff than $A_r A_c$.
Once again, an elegant part of Harsanyi's solution.

Specifying Games of Incomplete Information

Like a game of complete information,
a game of incomplete information
retains the usual physical components
— namely, the set of players and their action spaces.

These are complemented with preferences
and information components as follows:

- Chance/Nature: a probability distribution over players' types, which determine their preferences.
- Each player knows his own type, but not the other players' types.
- The probability distribution over types is common knowledge.

Bayesian Games

A normal form game of complete information is specified as $G = \langle N, (S_i)_{i=1}^n, (u_i(\cdot))_{i=1}^n \rangle$.

Recall that $N = \{1, 2, \ldots, n\}$ is the set of players,
that S_i is player i's action or strategy space,
and $u_i : S \to \mathbb{R}$ is player i's utility (payoff) function,
where $S := S_1 \times S_2 \times \cdots \times S_n$.

In a "Bayesian" game of incomplete information,
players know their own payoffs from different strategy profiles,
but may not know other players' payoffs.

To allow for this, we have introduced three new ideas.

Three New Features

1. Before the game is actually played,
   Chance or Nature chooses the different players types.
2. Each type t_i can represent information
   about player i's own payoffs, or more generally,
   about other relevant attributes of the game.
   Thus, there is a type space T_i for each player $i \in N$,
   representing the range from which Nature (or Chance)
   chooses i's type.
3. We introduce a commonly known
   common prior probability distribution $p(\cdot)$ on $\prod_{i=1}^{n} T_i$
   to describe how Chance/Nature chooses a type profile $(t_i)_{i=1}^{n}$.
   That is, every player i knows his own type $t_i \in T_i$,
   and uses this prior to form posterior beliefs
   over other profiles t_{-i} of agents' types.

Definition of Static Bayesian Game, I

Definition

The normal form representation
of an n-player static Bayesian game is

$$G^* = \langle N, (A_i)_{i=1}^n, (T_i)_{i=1}^n, ((u_i(\cdot; t_i))_{t_i \in T_i})_{i=1}^n, (p_i)_{i=1}^n \rangle.$$

As before, $N = \{1, 2, \ldots, n\}$ is the set of players,
but now A_i is player i's action set.

Also, $T_i = \{t_i^1, t_i^2, ..., t_i^{k_i}\}$ is player i's type space,
and $u_i : A \times T_i \to \mathbb{R}$ is player i's type dependent utility function,
where $A := A_1 \times A_2 \times \cdots \times A_n = \prod_{i=1}^n A_i$.

Definition of Static Bayesian Game, II

Consider again the general static Bayesian game

$$G^* = \langle N, (A_i)_{i=1}^n, (T_i)_{i=1}^n, ((u_i(\cdot; t_i))_{t_i \in T_i})_{i=1}^n, (p_i)_{i=1}^n \rangle.$$

The notation $u_i(a; t_i)$ indicates
that u_i is a utility function of the action profile a,
but it also depends on t_i as a parameter.

Finally, p_i describes player i's belief about other players' types
— that is, given that i knows his type is t_i,
the map $t_{-i} \mapsto p_i(t_{-i}|t_i)$ is the (posterior) conditional distribution
on t_{-i} (profiles of all other players' types, excluding i).

Timing of Static Bayesian Game

We assume the timing of the static Bayesian game is:

1. Nature chooses a type profile $(t_i)_{i=1}^n \in \prod_{i=1}^n T_i$.
2. Each player $i \in N$ learns his own type $t_i \in T_i$,
   which is his private information,
   and uses his prior p_i to form beliefs $p_i(t_{-i}|t_i)$
   over the other players' types t_{-i}.
3. Players simultaneously (as in a static game)
   choose actions $a_i \in A_i$ $(i \in N)$.
4. Given the players' choices $a = (a_1, a_2, \ldots, a_n)$,
   the payoffs $u_i(a; t_i)$ of each player $i \in N$ are realized.

Private Values

Note that this setup has a player's utility $u_i(a; t_i)$
depend only on all the players' actions $a \in A$
and on i's own type t_i,
but it does not depend on other players' types t_{-i}.

We call this the private values case,
since each type's payoff depends only on private information.

For example, how much are you willing to pay
for a minor artist's oil painting,
when the resale value is negligible?

Common Values

Later, some interesting examples will require us to discuss the common values case.

This allows utility functions $u_i(a_1, a_2, \ldots, a_n; t_1, t_2, \ldots, t_n)$ that depend on all players' types.

For example, how much are you willing to pay for one of Renoir's oil paintings?

It should last longer than you will,
so the resale value is important,
and depends on other connoisseur's tastes.

Bayes' Rule

Suppose two of many possible states S (sunny weather) and H (high waves) can occur exclusively or together according to some prior probability distribution $p(\cdot)$.
Thus, $p(S)$ denotes the prior probability of sunshine, $p(H)$ the prior probability of high waves, and $p(S \cap H)$ the prior probability that it will be sunny with high waves.
When you wake up, you see it is sunny; what can you infer about the probability of high waves (without seeing the sea)?
Bayes' Rule: Conditional on state S occurring, the probability that state H occurs as well is

$$\Pr\{H|S\} = \frac{p(S \cap H)}{p(S)}.$$

Intuitive Justification

Here is an intuitive justification.

The probability that S occurs
must be the sum $p(S) = p(S \cap H) + p(S \setminus H)$
because $S \cap H$ and $S \setminus H$ are mutually exclusive events.

Once S has occurred, then conditional on this knowledge,
the likelihood of H occurring
is the relative likelihood of both S and H occurring,
among all the states in which S occurs. Hence,

$$\Pr\{H|S\} = \frac{p(S \cap H)}{p(S \cap H) + p(S \setminus H)} = \frac{p(S \cap H)}{p(S)}.$$

Bayesian Game Example

Consider an example of a Bayesian game with two players, each having two possible types (say, a and b).
So t_i^k will indicate that player i is of type k.

Nature chooses pairs of players' types according to a prior joint distribution over the four possible type combinations (t_1^k, t_2^l), where $k, l \in \{a, b\}$.

This prior can be described by a joint distribution matrix:

	t_2^a	t_2^b
t_1^a	$\frac{1}{6}$	$\frac{1}{3}$
t_1^b	$\frac{1}{3}$	$\frac{1}{6}$

Table: Prior Joint Distribution

The "common prior" assumption requires everybody to take this joint distribution as given.

Posterior Beliefs

	t_2^a	t_2^b
t_1^a	$\frac{1}{6}$	$\frac{1}{3}$
t_1^b	$\frac{1}{3}$	$\frac{1}{6}$

Table: Prior Joint Distribution

Suppose player 1 learns that he is of type a.
What will be his posterior belief about player 2's type?
Using Bayes' rule we have,

$$p_1(t_2^a | t_1^a) = \frac{\Pr\{t_1^a \cap t_2^a\}}{\Pr\{t_1^a\}} = \frac{\frac{1}{6}}{\frac{1}{6} + \frac{1}{3}} = \frac{1}{3},$$

and similarly,

$$p_1(t_2^b | t_1^a) = \frac{\Pr\{t_1^a \cap t_2^b\}}{\Pr\{t_1^a\}} = \frac{\frac{1}{3}}{\frac{1}{6} + \frac{1}{3}} = \frac{2}{3}.$$

Correlated Types

	t_2^a	t_2^b
t_1^a	$\frac{1}{6}$	$\frac{1}{3}$
t_1^b	$\frac{1}{3}$	$\frac{1}{6}$

Table: Prior Joint Distribution

Note that each player has a $\frac{1}{2}$ chance of being either type a or type b.

But the types are not independent; rather, they are negatively correlated.

It is only half as likely that they will be the same as that they will differ.

Strategies and Bayesian Nash Equilibrium

In static normal form games of complete information, actions and strategies need not be distinguished, since choices are made once and for all.

For games of incomplete information, however, each player $i \in N$ can be of several different types $t_i \in T_i$, each of which may choose a different action from the set A_i.

Thus, a pure strategy for player i must specify what each of player i's types $t_i \in T_i$ will choose when Nature calls upon this type to act.

Similarly, a mixed strategy is a probability distribution over a player's pure strategies.

Strategies Are ...

Definition

Consider a static Bayesian game with private values

$$G^* = \langle N, (A_i)_{i=1}^n, (T_i)_{i=1}^n, ((u_i(\cdot; t_i)_{t_i \in T_i})_{i=1}^n, (p_i)_{i=1}^n \rangle.$$

A pure strategy for player i is a function $s_i : T_i \to A_i$
that specifies a pure action $s_i(t_i)$,
which is what i will choose when his type is t_i.

Similarly, a mixed behavioral strategy
is a function $\sigma_i : T_i \to \Delta A_i$ that specifies a lottery $\sigma_i(t_i)$
for each of i's possible types $t_i \in T_i$.

...Type Contingent

Thus, at the first stage,
even before Chance determines their types,
each player chooses a *type-contingent* strategy.

As a result of that Chance move,
players finally learn their types,
and then play according to their chosen strategy.

This offers a convenient way to model
players' beliefs over other players' strategies
when these depend on their different types.

Strategies in the Entry Game

In the Entry Game example, the incumbent firm 2 has two possible types r and c (for "rational" and "crazy"). There are two information sets, labelled $2r$ and $2c$, so $2^2 = 4$ type-dependent strategies.

Player 1 has beliefs about the type distribution that are given by Nature's probability distribution. Together with a specific strategy of player 2, these probabilities determine player 1's beliefs over different continuation paths in the game, and the resultant outcomes.

So the formulation allows players to evaluate their expected utility from their alternative choices. In the Entry Game example, let t_2^k denote player 2's type, with $k \in \{r, c\}$.

Expected Utility in the Entry Game

Suppose player 1 believes that player 2 is using the pure strategy

$$s_2(t_2) = \begin{cases} A & \text{if } t_2 = t_2^r \\ F & \text{if } t_2 = t_2^c \end{cases}$$

This simple example has only one type of player 1,
but we still denote it by t_1.

From player 1's own perspective (with some redundant notation),
the expected utility from playing E will be

$$\begin{aligned} \mathbb{E}u_1(E, s_2(\cdot); t_1) &= p_1(t_2^r|t_1)\, u_1(E, s_2(t_2^r); t_1) \\ &\quad + p_1(t_2^c|t_1)\, u_1(E, s_2(t_2^c); t_1) \\ &= p \cdot 1 + (1-p) \cdot (-1) \;=\; 2p - 1 \end{aligned}$$

When $p = \frac{2}{3}$ this yields $\mathbb{E}u_1(E, s_2(\cdot); t_1) = \frac{1}{3}$,
which is the payoff entry from the pure strategy pair $(E, A_r F_c)$
in the bimatrix that represents this game.

First observation:

If player i uses a pure (type-dependent) strategy,
while Chance chooses player i's type,
for other players $j \neq i$ it is as if player i uses a mixed strategy.

Again, in the Entry Game example,
suppose player 2 uses the strategy $A_r F_c$.

For player 1 it is as if player 2
chooses A with probability p
and F with probability $1 - p$.

Second observation:

As in an extensive form game,
we effectively specify player i's action
at all its information sets, one for each type.

Thus, even though player i knows its realized type,
we specify "counterfactually" what i would have done
even for those types that have not been realized
(and never will be).

This is necessary so that players other than i
can form well defined beliefs over player i's behavior.

Namely, players $j \neq i$ need to combine
their posterior beliefs over i's types
with their beliefs over
what each type t_i of player i would want to do.

Bayesian Nash Equilibrium Defined

In the static Bayesian game

$$G^* = \langle N, (A_i)_{i=1}^n, (T_i)_{i=1}^n, ((u_i(\cdot; t_i)_{t_i \in T_i})_{i=1}^n, (p_i)_{i=1}^n \rangle,$$

a strategy profile $s^* = (s_1^*(\cdot), s_2^*(\cdot), ..., s_n^*(\cdot))$
is a *pure strategy* Bayesian Nash Equilibrium (BNE)
if for every player $i \in N$,
and for each realization $t_i \in T_i$ of player i's type,
the action $a_i = s_i^*(t_i)$ is a best response because it solves

$$\max_{a_i \in A_i} \sum\nolimits_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i)\, u_i(a_i, s_{-i}^*(t_{-i}); t_i).$$

Note that $s_{-i}^*(t_{-i}) = (s_j^*(t_j))_{j \neq i}$
is the other players' equilibrium strategy profile
when their type profile is $t_{-i} = (t_j)_{j \neq i} \in \prod_{j \neq i} T_j$.
Note too that in the maximand we could replace a_i
with any type-contingent deviation $\tilde{s}_i(\cdot)$
without affecting the result.

Conditional Expected Utility

Remark

In the definition of Bayesian Nash equilibrium, we can write the best response condition more generally and concisely as requiring each $s_i^*(t_i)$ *to solve*

$$\max_{a_i \in A_i} \mathbb{E}_{t_{-i}}[u_i(a_i, s_{-i}^*(t_{-i}); t_i)|t_i].$$

We take player i*'s conditional expectations* $\mathbb{E}_{t_{-i}}[\cdot|t_i]$
over the random realizations of t_{-i}
given that player i *knows his own type* t_i*.*
The expectation is of $u_i(a; t_i)$
when player i*'s component* a_i *of strategy profile* a *is chosen freely,*
but the other players $j \neq i$
choose their equilibrium type dependent strategies $s_j^*(t_j)$*.*

Multiple Integration

The notation $\mathbb{E}_{t_{-i}}[\cdot|t_i]$ is more general
in the sense that it applies even when one or more players
have a *continuum of possible types*.

For instance, each player may have
an interval $T_i = [\underline{t}_i, \bar{t}_i] \subseteq \mathbb{R}$ as the type space,
with the random type drawn from T_i
with cumulative distribution $F_i(t_i)$
and density function $f_i(t_i) = F_i'(t_i)$ (if F_i is differentiable).

In this case player i's expected utility will be an integral
(more precisely, a multiple integral in $n-1$ dimensions)
over the realizations of the other players' types,
and the corresponding actions specified by their strategies.

Reflections on Equilibrium

Hence, a Bayesian Nash equilibrium has each player choose a type-contingent strategy $s_i^*(\cdot)$ so that, given any one of his types $t_i \in T_i$, and his beliefs about other players' strategies $s_{-i}^*(\cdot)$, his conditional expected utility from $s_i^*(t_i)$ is no less than from any other action $a_i \in A_i$.

The conditional probability distribution is determined from the other players' type contingent strategies, and by the Chance moves that each player faces through the conditional beliefs $p_i(t_{-i}|t_i)$.

Zero Probability Events

Remark

Bayes' Rule gives an ill defined answer if there is a zero probability event that appears in the denominator of the formula for a conditional probability.

This matters little for now, but is crucial when one considers sequential rationality in games of incomplete information.