1 The General Definition

A *dynamic game* (or extensive game, or game in extensive form) consists of:

- A set of *players* N;
- A set H of sequences (called *histories*) such that:
  1. $\emptyset \in H$;
  2. If $(a^1, \ldots, a^T) \in H$ then $(a^1, \ldots, a^S) \in H$ for all $S < T$ (T may be ∞);
  3. If $(a^1, \ldots, a^T) \in H$ for all T, then $(a^t)_{t=1,\ldots,\infty} \in H$.

  The elements of a history are actions played by the players. A history $(a^1, \ldots, a^T)$ is terminal if $T = \infty$ or if there is no a^{T+1} with $(a^1, \ldots, a^{T+1}) \in H$. Let Z denote the set of terminal histories.
- A *player function* P such that for each nonterminal history $h \in H$, $P(h)$ denotes the player that moves after history h. If $P(h) = c$, then we say that "Nature" (or "chance") moves after history h. For any nonterminal history h, the *action set* $A(h)$ denotes the set of actions available to $P(h)$.
- For each nonterminal history h with $P(h) = c$, a probability distribution $\mathbb{P}(\cdot|h)$ on the set of actions $A(h)$. (All such distributions are independent of each other.)
- For each player i, a partition $\mathcal{I}_i$ of $\{h \in H : P(h) = i\}$ (i.e., each history occurs in exactly one of the sets in $\mathcal{I}_i$), such that if $h, h' \in I_i$, then $A(h) = A(h')$.

  The collection $\mathcal{I}_i$ is the *information partition* of player i, and a member $I_i \in \mathcal{I}_i$ is *an information set* of player i.

  For any information set I, let $P(I)$ and $A(I)$ denote the player and action set corresponding to the information set. For any history h, let $I(h)$ denote the information set containing h.
- For each player i, a *payoff function* Π_i such that for each terminal history $h \in Z$, the payoff to player i is $\Pi_i(h)$.

(Our presentation is inspired by Osborne and Rubinstein [3].)

Remarks:

- Note that this is just a generalization of game trees – there need not be a single root, so the game graph in general is an *arborescence* (a collection of trees). With the graphical interpretation, a node in the game graph is always uniquely identified with a history.

- When the information sets are singletons everywhere, we say the game is one of *perfect information*. In all other cases, we say the game is one of *imperfect information*.

- Note that the union $\bigcup_i \mathcal{I}_i$ is a partition of the entire set of histories H.

- We will focus our attention on games of *perfect recall*: players don't forget the past. For a history h, let $X_i(h)$ denote the sequence consisting of the information sets seen by player i, and the actions played by i in those information sets. This contains all information that player i has about his past experience in the game. Perfect recall means that for two histories h, h' with $P(h) = P(h') = i$ and $I(h) = I(h')$, we have $X_i(h) = X_i(h')$; i.e., for two histories in the same information set, the player's experience in the past must have been the same.

- A (pure) *strategy* (or behavioral strategy) s_i for a player i is a complete contingent plan that specifies an action $s_i(h) \in A(h)$ for each history h where $P(h) = i$. (In a mixed strategy, each $s_i(h)$ is a distribution over the elements of $A(h)$.)

- A *subgame* after the history h is an extensive game $(N, H', Z', P, A, \mathcal{I}, \mathbf{\Pi})$ defined as follows. First, h must be a history such that the information set $I(h)$ is a singleton. The histories H' are *all* histories h' of the original game that can be written as $h' = (h, h'')$. Z' are the terminal histories of the original game that are in H'. Finally, the key requirement is that "no information sets are cut": if h is a history in H', then $I(h) \subset H'$; i.e., all other histories in the same information set are also in the subgame history set H'.

  In game tree terminology, this means a subgame is a subtree beginning at a single node that does not cut any information sets.

2 Subgame Perfect Equilibrium

Typically, when we discuss equilibrium, we always assume players are rational; i.e., they act to maximize their payoffs. The first notion of equilibrium is *Nash equilibrium*: $\boldsymbol{s}$ is a Nash equilibrium if s_i is a best response to $\boldsymbol{s}_{-i}$ for all players i. The actual history h that is realized under $\boldsymbol{s}$ is called the *equilibrium path*. (Note that with randomization, an equilibrium path is chosen according to a probability distribution over the set of histories H.)

As you should recall, NE does not address the issue of credibility in information sets that are off the equilibrium path: a player can threaten action in information sets not on the equilibrium path, that would never actually be played by a rational player if those information sets were reached. This leads to the definition of *subgame perfect Nash equilibrium*: $\boldsymbol{s}$ is a subgame perfect NE if it is a NE in every *subgame* (including the original game).

3 Sequential Equilibrium

SPNE can break down as a predictive tool when there are "not enough subgames"; this can lead to a situation where any reasonable beliefs a player has in an information set are not consistent with predicted equilibrium behavior.

To correct this deficiency, Kreps and Wilson [2] suggested *sequential equilibrium*. To define SE, we first need to define *beliefs*. A *belief system* for player i is a collection of probability distributions $\{\mu_i(\cdot|I) : I \in \mathcal{I}_i\}$; each distribution $\mu_i(\cdot|I)$ is over the set of histories in the information set I. The interpretation is that $\mu_i(h|I)$ is the belief of player $P(I)$ that the prior history was h, given that he is in the information set I.

Let $P_i(\boldsymbol{s}, \mu_i|I_i)$ denote the expected payoff to player i in information set I_i, when he has belief system μ_i, and all players' strategies are $\boldsymbol{s}$. (The player computes his expected payoffs according to his beliefs.) Then *sequential rationality* requires that for all i, $I_i \in \mathcal{I}_i$, and strategies s_i':

$$P_i(\boldsymbol{s}, \mu_i|I_i) \geq P_i(s_i', \boldsymbol{s}_{-i}, \mu_i|I_i).$$

However, this is not a complete description of equilibrium, as we also require that beliefs are *consistent* with the strategies chosen. The problem is that under a given strategy profile, not all information sets will be reached with positive probability. Kreps and Wilson correct this as follows.

We define consistency for finite games (i.e., games where all relevant sets—histories, action sets, players—are finite). In this case, we say that a strategy vector $\boldsymbol{s}$ and belief system vector $\boldsymbol{\mu}$ are *consistent* if there exists a sequence $(\boldsymbol{s}^n, \boldsymbol{\mu}^n) \to (\boldsymbol{s}, \boldsymbol{\mu})$ (in the usual Euclidean sense) such that each $\boldsymbol{s}_i^n$ is a completely mixed strategy, and each belief system vector $(\mu_1^n, \ldots, \mu_N^n)$ is computed using Bayes' rule. (A completely mixed strategy chooses actions so that every possible strategy has positive probability.) The idea is that "trembles" in the strategy-belief pair should still lead to "consistent" beliefs.

A pair $(\boldsymbol{s}, \boldsymbol{\mu})$ is a *sequential equilibrium* if it is sequentially rational and consistent.

Note that any SE is a NE: the NE definition only requires that no deviations be profitable on the equilibrium path. Since information sets on the equilibrium path are reached with positive probability, the sequential rationality requirement ensures that no deviation is profitable. Also, since information sets are singletons in perfect information games, sequential rationality and SPNE coincide there—beliefs are trivial.

4 Multistage Games with Observed Actions

While compact, the definition of sequential equilibrium is troubling, because it involves a potentially complex calculation to check that the limit condition holds. Perfect Bayesian equilibrium simplifies this point, but is lengthier to state formally. Following Fudenberg and Tirole, we develop PBE here for multistage games with observed actions. We will abuse notation somewhat, but the interpretation will be clear below.

A *(Bayesian) multistage game with observable actions* is a dynamic game defined as follows. Nature moves first, and chooses a *type* $\theta_i \in \Theta_i$ for each player i. We assume types are independent; the probability player i is type θ_i is denoted $p_i(\theta_i)$. Each player i observes his own type, but not the types of the other players. From this point on, play proceeds in a sequence of stages $1, \ldots, T$. At each stage t, ever player i chooses an action a_i^t. Let $\boldsymbol{a}^t$ denote the composite action vector chosen by the players at stage t; by an abuse of notation, we let $h^t = (\boldsymbol{a}^1, \ldots, \boldsymbol{a}^{t-1})$ denote the history up to time t. "Observable actions" refers to the fact that we assume that all players observe the past history.

After the history h^t, the set of actions available to player i is $A_i(h^t)$. A strategy of a player i specifies an action $s_i(h^t, \theta_i) \in A_i(h^t)$ (or a mixed action) after each history, and given his type.

Now note that all players observe the *same* history, and it is common knowledge that players' types are independently drawn according to $\boldsymbol{p}$. It then follows that after history h^t, *all* players other than i have the same belief about player i's type; and further, i's beliefs about other players' types should be independent. Given this observation, we let $\gamma_{-i}(\theta_i|h^t)$ denote the belief of players other than i about player i's type; γ_{-i} is a conditional distribution over the types of player i.

A pair $(\boldsymbol{s}, \boldsymbol{\gamma})$ is a *Perfect Bayesian equilibrium* if:

1. *Initial beliefs*: All players' initial beliefs are consistent with $\boldsymbol{p}$: for all i and θ_i, $\gamma_{-i}(\theta_i|\emptyset) = p_i(\theta_i)$.

2. *Bayesian updating where possible*: Suppose for some θ_i' and action profile $\boldsymbol{a}^t$, $\gamma_{-i}(\theta_i'|h^t) > 0$ and $s_i(h^t, \theta_i')$ puts positive weight on action a_i^t. Then for all θ_i:

$$\gamma_{-i}(\theta_i|h^t, \boldsymbol{a}^t) = \frac{\gamma_{-i}(\theta_i|h^t)\mathbb{P}(s_i(h^t, \theta_i) = a_i^t)}{\sum_{\hat{\theta}_i} \gamma_{-i}(\hat{\theta}_i|h^t)\mathbb{P}(s_i(h^t, \hat{\theta}_i) = a_i^t)}.$$

3. *Action-determined beliefs*: Given $\boldsymbol{a}^t$ and $\hat{\boldsymbol{a}}^t$, if $a_i^t = \hat{a}_i^t$, then $\gamma_{-i}(\theta_i|h^t, \boldsymbol{a}^t) = \gamma_{-i}(\theta_i|h^t\hat{\boldsymbol{a}}^t)$.

4. *Sequential rationality*: Player i's expected payoff starting after any history h^t, $t = 1, \ldots, T$ is maximized by playing s_i. (Note that the strategy depends on both history and player i's type.)

Each of these conditions has a simple interpretation. The first says that all players start with beliefs given by the distributions $p_1, \ldots, p_N$. The second requires that, if possible, players use Bayes' rule to update their beliefs. The third says that players only use the actions of other players to update beliefs; if player i plays the same action after the same history, it should lead to the same beliefs about his type. Finally, the fourth says that players maximize their payoffs, given their beliefs.

It is possible to show that for multistage games with observable actions, *every sequential equilibrium is a perfect Bayesian equilibrium*. (To make this precise, it is important to translate from the model of beliefs γ_{-i} used here, to the model of beliefs μ_i used in our discussion of sequential equilibrium.) It need not be true, however, that any perfect Bayesian is a sequential equilibrium. It is true, however, that if either each player has at most two possible types, or the game has only two periods, then any PBE is also a sequential equilibrium. (The results of this paragraph were established by Fudenberg and Tirole [1].)

Exercise. Suppose $(\boldsymbol{s}, \boldsymbol{\mu})$ is a sequential equilibrium of the extensive form description of a Bayesian multistage game of observable actions. Show that in the extensive form description, an information set of player i at stage t corresponds to a pair (θ_i, h^t); and show there exists a collection of beliefs γ_{-i} such that:

$$\mu_i(\boldsymbol{\theta}, h^t|\theta_i, h^t) = \prod_{j \neq i} \gamma_{-j}(\theta_j|h^t).$$

Conclude that the pair $(\boldsymbol{s}, \boldsymbol{\gamma})$ is a perfect Bayesian equilibrium. (In other words, show that the joint belief over $\boldsymbol{\theta}$ in the sequential equilibrium decomposes as a product of commonly held beliefs over $\boldsymbol{\theta}_{-i}$.)

References

[1] D. Fudenberg and J. Tirole. Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53:236–260, 1991.

[2] D. Kreps and R. Wilson. Sequential equilibrium. *Econometrica*, 50:863–894, 1982.

[3] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts, 1994.