14.126 GAME THEORY

MIHAI MANEA

Department of Economics, MIT,

1. Sequential Equilibrium

In multi-stage games with incomplete information, say where payoffs depend on initial moves by nature, the only proper subgame is the original game, even if players observe one another's actions at the end of each period. Thus the refinement of Nash equilibrium to subgame perfect equilibrium has no bite. Since players do not know the others' types, the start of a period can only be analyzed as a separate subgame when the players' posterior beliefs are specified. The concept of sequential equilibrium proposes a way to derive plausible beliefs at every information set. Based on the beliefs, one can test whether the continuation strategies form a Nash equilibrium.

The complications that incomplete information causes are easiest to see in "signaling games"—leader-follower games in which only the leader has private information. The leader moves first; the follower observes the leader's action, but not the leader's type, before choosing his own action. One example is Spence's (1974) model of the job market. In that model, the leader is a worker who knows her productivity and must choose a level of education; the follower, a firm (or number of firms), observes the worker's education level, but not her productivity, and then decides what wage to offer her. In the spirit of subgame perfection, the optimal wage should depend on the firm's beliefs about the worker's productivity given the observed education. An equilibrium needs to specify not only contingent actions, but also beliefs. At information sets that are reached with positive probability in equilibrium, beliefs should be derived using Bayes' rule. However, there are some theoretical issues about belief update following zero-probability events.

Date: March 2, 2010.

Refer for more motivation to the example in FT, figure 8.1 (p. 322). The strategy profile (L, A) is a Nash equilibrium, which is subgame perfect as player 2's information set does not initiate a proper subgame. However, it is not a very plausible equilibrium, since player 2 prefers playing B rather than A at his information set, regardless of whether player 1 has chosen M or R. So, a good equilibrium concept should rule out the solution (L, A) in this example and ensure that 2 always plays B. The problem with the considered equilibrium is that player 2 does not play a best response to any possible belief at his information set.

For most definitions, we focus on extensive form games of prefect recall with finite sets of decision nodes. We use some of the notation introduced earlier.

A **sequential equilibrium** (Kreps and Wilson 1982) is an assessment (σ, μ), where σ is a (behavior) strategy profile and μ is a **system of beliefs**. The latter component consists of a belief specification $\mu(h)$ over the nodes at each information set h. The definition of sequential equilibrium is based on the concepts of sequential rationality and consistency. **Sequential rationality** requires that conditional on every information set h, the strategy $\sigma_{i(h)}$ be a best response to $\sigma_{-i(h)}$ given the beliefs $\mu(h)$. Formally,

$$u_{i(h)}(\sigma_{i(h)}, \sigma_{-i(h)}|h, \mu(h)) \geq u_{i(h)}(\sigma'_{i(h)}, \sigma_{-i(h)}|h, \mu(h))$$

for all information sets h and alternative strategies σ'.

Beliefs need to be **consistent** with strategies in the following sense. For any fully mixed strategy profile $\tilde{\sigma}$—that is, one where each action is played with positive probability at every information set—all information sets are reached with positive probability and Bayes' rule leads to a unique system of beliefs $\mu^{\tilde{\sigma}}$. The assessment (σ, μ) is consistent if there exist a sequence of fully mixed strategy profiles $(\sigma^m)_{m \geq 0}$ converging to σ such that the associated beliefs μ^{σ^m} converge to μ as $m \to \infty$.

Definition 1. *A sequential equilibrium is an assessment which is sequentially rational and consistent.*

The definition of sequential equilibrium rules out the strange equilibrium in the earlier example (FT figure 8.1). Since player 1 chooses L under the proposed equilibrium strategies, consistency does not pin down player 2's beliefs at his information set. However, sequential rationality requires that player 2 have some beliefs and best-respond to them, which ensures that A is not played.

Consistency imposes more restrictions than Bayes' rule alone. Consider figure 8.3 in FL (p. 339). The information set h_1 of player 1 consists of two nodes x, x'. Player 1 can take an action D leading to y, y' respectively. Player 2 cannot distinguish between y and y' at the information set h_2. If 1 never plays D in equilibrium, then Bayes' rule does not pin down beliefs at h_2. However, consistency implies that $\mu_2(y|h_2) = \mu_1(x|h_1)$. The idea is that since 1 cannot distinguish between x and x', he is equally likely to tremble at either node. Hence trembles ensure that players' beliefs respect the information structure.

More generally, consistency imposes common beliefs following deviations from equilibrium behavior. There are criticisms of this requirement—why should different players have the same theory about something that was not supposed to happen? A contra-argument is that consistency matches the spirit of equilibrium analysis, which normally assumes that players agree in their beliefs about other players' strategies (namely, players share correct conjectures about each other's strategies).

2. Properties of Sequential Equilibrium

Theorem 1. *A sequential equilibrium exists for every finite extensive-form game.*

This is a consequence of the existence of perfect equilibria, which we prove later.

Proposition 1. *The sequential equilibrium correspondence is upper hemi-continuous with respect to payoffs.*

Proof. Let $u^k \to u$ be a convergent sequence of payoff functions and $(\sigma^k, \mu^k) \to (\sigma, \mu)$ be a convergent sequence of sequential equilibria of the games with corresponding payoffs u^k. We need to show that (σ, μ) is a sequential equilibrium for the game with payoffs given by u. Sequential rationality of (σ, μ) is straightforward because the expected payoffs conditional on reaching any information set are continuous in the payoff functions and beliefs.

We also have to check consistency of (σ, μ). As (σ^k, μ^k) is a sequential equilibrium of the game with payoff function u^k, there exists a sequence of completely mixed strategies $(\sigma^{m,k})_m \to \sigma^k$, with corresponding induced beliefs given by $(\mu^{m,k})_m \to \mu^k$. For every k, we can find a sufficiently large m_k so that each component of σ^{k,m_k} and μ^{m_k} are within

$1/k$ from the corresponding one under σ^k and μ^k. Since $\sigma^k \to \sigma, \mu^k \to \mu$, it must be that $\sigma^{m_k,k} \to \sigma, \mu^{m_k,k} \to \mu$. Thus we have obtained a sequence of fully mixed strategies converging to σ, which induces beliefs converging to μ. □

Kreps and Wilson show that in generic games (i.e., a space of payoff functions such that the closure of its complement has measure zero), the set of sequential equilibrium outcome distributions is finite. Nevertheless, it is not generally true that the set of sequential equilibria is finite, as there may be infinitely many belief specifications for off-path information sets that support some equilibrium strategies. We provide an illustration in the context of the beer-or-quiche signaling game of Cho and Kreps (1987).

See figure 11.6 in FT (p. 450). Player 1 is wimpy or surly, with respective probabilities 0.1 or 0.9. Player 2 is a bully who would like to fight the wimpy type but not the surly one. Player 1 orders breakfast and 2 decides whether to fight him after observing his breakfast choice. Player 1 gets a utility of 1 from having his favorite breakfast—beer if surly, quiche if weak—but a disutility of 2 from fighting. When player 1 is weak, player 2's utility is 1 if he fights and 0 otherwise; when 1 is surly, the payoffs to the two actions are reversed. One can show that there are two classes of sequential equilibria, corresponding to two distinct outcomes. In one set of sequential equilibria, both types of player 1 drink beer, while in the other both types of player 1 eat quiche. In both cases, player 2 must fight with probability at least 1/2 when observing the out-of-equilbrium breakfast in order to make the mismatched type of player 1 endure gastronomic horror. Note that either type of equilibrium can be supported with any belief for player 2 placing a probability weight of at least 1/2 on player 1 being wimpy following the out-of-equilbrium breakfast. Hence there is an infinity of sequential equilibrium assessments.

Kohlberg and Mertens (1986) criticized sequential equilibrium for allowing "strategically neutral" changes in the game tree to affect the equilibrium. Compare, for instance, the two games in FT figure 8.6 (p. 343). The game on the right is identical to the one on the left, except that player 1's first move is split into two moves in a seemingly irrelevant way. Whereas (A, L_2) can be supported as a sequential equilibrium for the game on the left, the strategy A is not part of a sequential equilibrium for the one on the right. For the latter game, in the simultaneous-move subgame following NA, the only Nash equilibrium is (R_1, R_2), as L_1 is strictly dominated by R_1 for player 1. Hence the unique sequential equilibrium strategies for the right-hand game are (NA, R_1, R_2).

Note that the sensitivity of sequential equilibrium to the addition of "irrelevant moves" is not a direct consequence of consistency, but is rather implied by sequential rationality. In the example above, the problem arises even for subgame perfect equilibria. Kohlberg and Mertens (1986) further develop these ideas in their concept of a stable equilibrium. However, their proposition that mistakes be "conditionally optimal" is not necessarily compelling. If we take seriously the idea that players make mistakes at each information set, then it is not clear that the two extensive forms above are equivalent. In the game on the right, if player 1 makes the mistake of not playing A, he is still able to ensure that R_1 is more likely than L_1; in the game on the left, he might take either action by mistake when intending to play A.

3. Perfect Bayesian Equilibrium

Perfect Bayesian equilibrium is a concept that has been around for a while and predates sequential equilibrium. The idea is similar to sequential equilibrium but with more basic requirements about how beliefs are updated. Fudenberg & Tirole (1991) have a paper that describes various formulations of PBE. The basic requirements are that strategies should be sequentially rational and that beliefs should be derived from Bayes's rule wherever applicable, with no constraints on beliefs at information sets reached with probability zero in equilibrium.

Other properties that can be imposed:

- In a multi-stage game with independent types — i.e. exactly one move by Nature, at the beginning of the game, assigning types to players and such that types are independently distributed, with all subsequent actions of observed — beliefs about different players should remain independent at each history. (PBE is usually applied to games in which Nature moves only at the beginning and actions are observed.)
- Updating should be "consistent": given a probability-zero history h^t at time t, from which strategies do call for a positive-probability transition to history h^{t+1}, the belief at h^{t+1} should be given by updating beliefs at h^t via Bayes's rule.

- "Not signaling what you don't know": beliefs about player i at the beginning of period $t+1$ depend only on h^t and action by player i at time t, not also on other players' actions at time t.
- Two different players i, j should have the same belief about a third player k even at probability-zero histories.

All of these conditions are implied by consistency.

Anyhow, there does not seem to be a single clear definition of PBE in the literature. Different sets of conditions are imposed by different authors. For this reason, using sequential equilibrium is preferable.

4. Perfect Equilibrium

Now consider the following game:

	L	R
U	$1,1$	$0,0$
D	$0,0$	$0,0$

Both (U, L) and (D, R) are sequential equilibria (sequential equilibrium coincides with Nash equilibrium in a normal-form game). But (D, R) seems non-robust: if player 1 thinks that player 2 might make a mistake and play L with some small probability, he would rather deviate to U. This motivates the definition of **(trembling-hand) perfect equilibrium** (Selten, 1975) for normal-form games. A profile σ is a PE if there is a sequence of "trembles" $\sigma^m \to \sigma$, where each σ^m is a completely mixed strategy, such that σ_i is a best reply to σ^m_{-i} for each m.

An equivalent approach is to define a strategy profile σ^ε to be an **ε-perfect equilibrium** if there exist $\varepsilon(s_i) \in (0, \varepsilon)$ for all i, all s_i, such that σ^ε is a Nash equilibrium of the game where players are restricted to play mixed strategies where every strategy s_i has probability at least $\varepsilon(s_i)$. A PE is a profile that is a limit of some sequence of ε-perfect equilibria σ^ε as $\varepsilon \to 0$. (We will not show the equivalence here but it's not too hard.)

Theorem 2. *Every finite normal-form game has a perfect equilibrium.*

Proof. For any $\varepsilon > 0$, we can certainly find a Nash equilibrium of the modified game, where each player is restricted to play mixed strategies that place probability at least ε on every pure action. (Just apply the usual Nash existence theorem for compact strategy sets and quasiconcave payoffs.) By compactness, there is some subsequence of these strategy profiles as $\varepsilon \to 0$ that converges, and the limit point is a perfect equilibrium by definition. □

We would like to extend this definition to extensive-form games. Consider the game in Fig 8.11 (p. 353) of FT. They show an extensive-form game and its reduced normal form. There is a unique SPE (L_1L_1', L_2). But (R_1, R_2) is a PE of the reduced normal form. Thus perfection in the normal form does *not* imply subgame-perfection. The perfect equilibrium is sustained only by trembles such that, conditional on trembling to L_1 at the first node, player 1 is also much more likely to play R_1' than L_1' at his second node. This seems unreasonable — R_1' is only explainable as a tremble. Perfect equilibrium as defined so far thus has the disadvantage of allowing correlation in trembles at different information sets.

The solution to this is to impose perfection in the **agent-normal form**. We treat the two different nodes of player 1 as being different players, thus requiring them to tremble independently. More formally, in the agent-normal form game, we have a player corresponding to every information set. Given a strategy profile for all the players, each "player" corresponding to an information set h gets payoff given by the payoff of player $i(h)$ from the corresponding strategies in the extensive-form game. Thus, the game in figure 8.11 turns into a three-player game. The only perfect equilibrium of this game is (L_1, L_1', L_2).

More generally, a **perfect equilibrium** in an extensive-form game is defined to be a perfect equilibrium of the corresponding agent-normal form.

Theorem 3. *Every PE of a finite extensive-form game is a sequential equilibrium (for some appropriately chosen beliefs).*

Proof. Let σ be the given PE. So there exist fully mixed strategy profiles $\sigma^m \to \sigma$ which are ε-perfect equilibria of the agent-normal form game with $\varepsilon \to 0$. For each σ^m we have a well-defined belief system induced by Bayes's rule. Pick a subsequence for which these belief systems converge, to some μ. Then by definition (σ, μ) is consistent. Sequential rationality follows exactly from the fact that σ is a perfect equilibrium of the agent-normal form, using the first definition of perfect equilibrium. (More properly, this implies that there are no one-shot deviations that benefit any player; by an appropriate adaptation of the one-shot deviation principle this shows that σ is in fact fully sequentially rational at every information set.) □

The converse is not true — not every sequential equilibrium is perfect, as we already saw with the simple normal-form example above. But for generic payoffs it is true (Kreps & Wilson, 1982).

The set of perfect equilibrium outcomes is not upper-hemicontinuous (unlike sequential equilibrium or subgame-perfect equilibrium). The game has (D, R) as a perfect equilibrium

	L	R
U	$1, 1$	$0, 0$
D	$0, 0$	$1/n, 1/n$

for each $n > 0$, but in the limit where (D, R) has payoffs $(0, 0)$ it is no longer a perfect equilibrium. We can think of this as an order-of-limits problem: as $n \to \infty$ the trembles against which D and R remain best responses become smaller and smaller.

5. Proper Equilibrium

Myerson (1978) considered the notion that when a player trembles, he is still more likely to play better actions than worse ones. Myerson's notion is that a player's probability of playing the second-best action is at most ε times the probability of the best action, the probability of the third-best action is at most ε times the probability of the second-best action, and so forth. Consider the game in Fig. 8.15 of FT (p. 357). (M, M) is a perfect equilibrium, but Myerson argues that it can be supported only using unreasonable trembles, where each player has to be likely to tremble to a very bad reply rather than an almost-best reply.

Definition 2. *A ε-**proper equilibrium** is a totally mixed strategy profile σ^ε such that, if $u_i(s_i, \sigma^\varepsilon_{-i}) < u_i(s'_i, \sigma^\varepsilon_{-i})$, then $\sigma^\varepsilon_i \leq \varepsilon \sigma^\varepsilon_i(s'_i)$. A **proper equilibrium** is any limit of some ε-proper equilibria as $\varepsilon \to 0$.*

Theorem 4. *Every finite normal-form game has a proper equilibrium.*

Proof. First prove existence of ε-proper equilibria, using the usual Kakutani argument applied to the "almost-best-reply" correspondences BR_i^ε rather than the usual best-reply correspondences. ($BR_i^\varepsilon(\sigma_{-i})$ is the set of mixed strategies for player i in a suitable compact space of fully mixed strategies that satisfy the inequality in the definiton of ε-proper equilibrium.) Then use compactness to see that there exists a sequence that converges as $\varepsilon \to 0$; its limit is a proper equilibrium. □

Given an extensive-form game, a proper equilibrium of the corresponding normal form is automatically subgame-perfect; we don't need to go to the agent-normal form. We can show this by a backward-induction-type argument.

Kohlberg and Mertens (1986) showed that a proper equilibrium in a normal-form game is sequential in every extensive-form game having the given normal form. However, it will not necessarily be a trembling-hand perfect equilibrium in (the agent-normal form of) every such game. See Figure 8.16 of FT (p. 358): (Lr) is proper (and so sequential) but not perfect in the agent-normal form.