Perfect Bayesian Equilibrium in Extensive-Form Games

Introduction

We have seen that Nash equilibria of extensive-form games can be undesirable because they can rely on incredible threats at off-the-equilibrium-path subgames. We were sometimes able to refine away such undesirable equilibria by strengthening our solution concept—demanding subgame perfection, which requires that the restriction of a strategy profile to any subgame be a Nash equilibrium of that subgame.

Subgame perfection will not eliminate all undesirable equilibria of extensive-form games, however. Consider the extensive-form game of Figure 1. Analysis of its strategic form quickly shows that this game has two pure-strategy Nash equilibria: (U,l) and (A,r). This game has only one subgame, viz. the entire game, so both of these Nash equilibria are also subgame perfect.

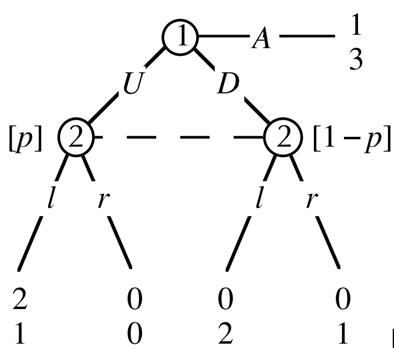

	l	r
A	1,**3**	**1,3**
U	**2,1**	0,0
D	0,**2**	0,1

□□□□□□□

Figure 1: Subgame perfection admits undesirable equilibria.

The (A,r) equilibrium is objectionable for the following reason. [Note that player 2's information set is off-the-equilibrium path with respect to the (A,r) equilibrium—i.e. it is never reached when the players conform to the equilibrium specification.] If player 2's information set were ever reached, player 2 would be uncertain about whether it was reached via player 1 having chosen U or via player 1 having chosen D. However, it doesn't matter to player 2's decision which move player 1 had chosen. No matter what player 2's beliefs about player 1's non-A choice, player 2 strictly prefers to choose l at his information set when it is reached. (If player 1 had chosen U, player 2 receives 1 from l and only 0 from r. If player 1 had chosen D, player 2 receives 2 from l and only 1 from r.) Because r is not a best response at player 2's information set for any possible beliefs which player 2 might have there, we say that r is dominated at player 2's information set. We can describe our dissatisfaction with the (A,r) equilibrium by objecting to its specification of an action at an information set which is dominated at that information set.

Let's formalize this reasoning. We begin by requiring that at every one of her information sets each player has *some* beliefs about the node at which she is located conditional on having reached that information set.

Bayes Requirement 1 For a particular strategy profile σ, we require that, for each player $i\in I$, and at each of her information sets $h_i\in H_i$, player i has beliefs $\rho_i(h_i)\in\Delta(h_i)$ about the node at which she is located conditional upon being informed that play has reached the information set h_i.

The beliefs $\rho_i(h_i)\in\Delta(h_i)$ are just a probability distribution over the nodes in the information set. Player i's beliefs in this game, then, are a specification, for each player-i information set $h_i\in H_i$, of such conditional beliefs at that information set. The n-tuple $\rho=(\rho_1,\ldots,\rho_n)$ of player beliefs is a *belief profile*.

In order to properly critique alleged equilibria we require that a candidate equilibrium be not just a strategy profile σ but be a strategy-belief profile (σ,ρ). We want to state an equilibrium requirement that would loosely say something like: For every player $i\in I$ and every player-i information set $h_i\in H_i$, player i's strategy is a best response given her beliefs $\rho_i(h_i)\in\Delta(h_i)$ at the information set h_i. However, this is too vague—at least to me!—so we must be more precise in our statement.

Recall that a subgame is formed by identifying a singleton information set and including all its successors from the original game. Information sets, actions, and payoffs for the subgame were derived from the original game by restriction. We now generalize the concept of a subgame and define a *continuation game*. A continuation game is an information set $h_i\in H_i$ for some player $i\equiv\iota(h_i)$ and all of its successor nodes from the original game. Again, information sets, actions, and payoffs in the continuation game are derived from the original game by restriction. If the designated initial information set is not a singleton, then this continuation is not a subgame.[1] And with good reason: This continuation game cannot be played as a game in its own right, because there is no initial node. So we include in the specification of the continuation game the probability distribution $\rho_i(h_i)$ over the nodes of the initial information set h_i given in the belief profile ρ. (Think of this continuation game as being preceded by a move of Nature's, where Nature chooses between the nodes of h_i according to the probability distribution $\rho_i(h_i)$.) We can restrict any strategy σ_j and any player beliefs ρ_j to this continuation game

---

1 We're assuming perfect recall.

just as we restricted a strategy to a subgame: simply throw out its specifications at information sets which don't belong to the smaller game.

Bayes Requirement 2 Consider the continuation game defined by some player-i information set $h_i \in H_i$ and the conditional beliefs $\rho_i(h_i)$. The restriction of the strategy-belief profile (σ, ρ) to this continuation game must be a Nash equilibrium of the continuation game.

Definition Let (σ, ρ) be a strategy-belief profile and let $h_i \in H_i$ be an information set for player $i = \iota(h_i)$. Let $(\tilde{\sigma}, \tilde{\rho})$ be the restriction of (σ, ρ) to the continuation game which begins at the information set h_i. We say that the player-i strategy σ_i is *strictly dominated beginning at the information set* h_i if there exists another player-i strategy σ_i' such that, for all other deleted strategy profiles σ_{-i}' for the opponents, player i's expected payoff in the continuation game is strictly higher for $(\tilde{\sigma}_i', \tilde{\sigma}_{-i})$ than for $\tilde{\sigma}$.

Bayes Requirements 1 and 2 are sufficient to remove the undesirable equilibrium, viz. (A, r), in Figure 1. To see this we construct the continuation game beginning at player 2's information set for some beliefs parameterized by $p \in [0, 1]$. See Figure 2. The strategic form of this continuation game is also shown in Figure 2, from which it is clear that l is the unique Nash equilibrium in the continuation game.

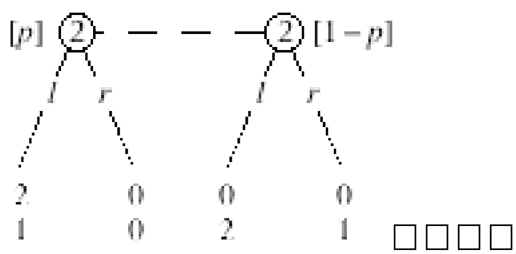

l	r
$2p, 2-p$	$0, 1-p$

Figure 2: The continuation game beginning at player 2′s information set.

More generally, Bayes Requirement 2 rejects all strategy profiles which specify at any information set an action which is dominated at that information set.

Example: Restricting a strategy-belief profile to a continuation game

Consider the strategy-belief profile $s = (U, a, d; l; p)$ for some $p \in [0, 1]$ in the extensive-form game in Figure 2a. Now consider the continuation game beginning at player 2's information set. Figure 2b depicts the restriction $\tilde{s}$ of this strategy-belief profile to the continuation game.

The expected payoff vector to the restriction $\tilde{s}$ is $p(2, 3) + (1-p)(3, 4) = (3-p, 4-p)$.

Let's evaluate whether the strategy profile s passes Bayes Requirement 2 with respect to the continuation game beginning at player 2's information set. We can construct a strategic form for the continuation game. For example, the expected payoff vector to the restricted strategy profile $(a, c; l)$ is $p(2, 3) + (1-p)(0, 1) = (2p, 2p+1)$. The expected payoff vector to $(*, *; r)$ is $p(1, 2) + (1-p)(1, 6)\ =$

$(1, 6-4p)$. Similar calculations for player 1's other strategies yield the payoff matrix in Figure 3.

In order that s passes Bayesian Requirement 2 with respect to the specified continuation game, $\tilde{s} = (a, d; l; p)$ must be a Bayesian Nash equilibrium of the continuation game. This requires that (a, d) be a best response by player 1 to l, i.e.

$$3-p \geq \max\{2p, 0, 3-3p\},$$

which is satisfied for all $p \in [0, 1]$. In order that l be a best response by player 2 to (a, d) we must have

$$4-p \geq 6-4p,$$

which is satisfied if and only if $p \in [\frac{2}{3}, 1]$. Therefore $\tilde{s}$ is a Bayesian Nash equilibrium of this continuation game if and only if $p \in [\frac{2}{3}, 1]$.

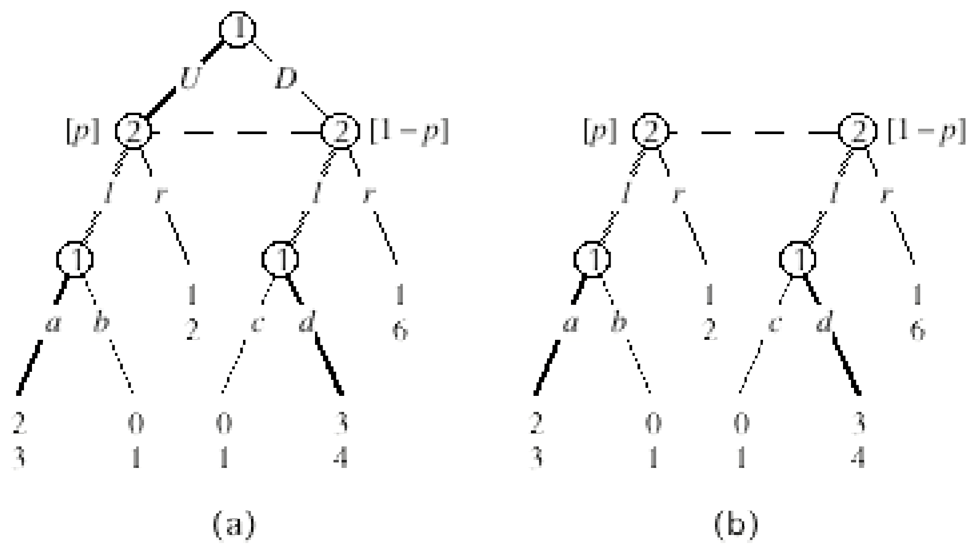

Figure 2: The restriction of the $(U, a, d; l; p)$ of (a) to the continuation game (b).

	l	r
(a,c)	$2p, 1+2p$	$1, 6-4p$
(a,d)	$3-p, 4-p$	$1, 6-4p$
(b,c)	$0, 1$	$1, 6-4p$
(b,d)	$3-3p, 4-3p$	$1, 6-4p$

Figure 3: The strategic-form matrix corresponding to the continuation game of Figure 2b.

We can also ask whether l is strictly dominated beginning at player 2's information set. This would require that

$$6-4p > \max\{1+2p, 4-p, 1, 4-3p\} = 4-p,$$

or $p < \frac{2}{3}$.

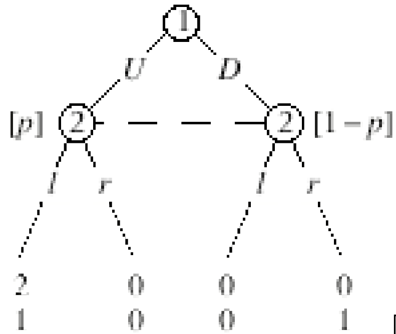

	l	r
U	**2,1**	0,0
D	0,0	**0,1**

□□□□

Figure 4: Requirements 1 and 2 don't even guarantee Nash equilibrium.

Bayes Requirements 1 and 2 are not strong enough however to generally capture even the concept of Nash equilibrium. Consider the game in Figure 4. The strategy-belief profile $(U, r; p = 0)$ satisfies Bayes Requirements 1 and 2, yet (U, r) is not even a Nash equilibrium of the game.

So we add a third requirement. It ensures that in equilibrium each player's beliefs are correct.

Bayes Requirement 3 The beliefs at any on-the-path information set must be determined from the strategy profile according to Bayes' Rule. I.e. if $h_i \in H_i$ is a player-i information set reached with positive probability when the players conform to σ, then $\rho_i(h_i) \in \Delta(h_i)$ must be computed from σ using Bayes' Rule.

This requirement eliminates the non–Nash-equilibrium profile $(U, r;\ p = 0)$ from the game of Figure 4 because, given that player 1 is choosing U, player 2's belief at his information set must put all weight on the node reached by U—i.e. we must have $p = 1$ rather than $p = 0$.

However, the addition of Bayes Requirement 3 does not guarantee that surviving strategy-belief profiles are even subgame-perfect equilibria of the game. Consider the game of Figure 5. Consider the subgame beginning with player 2's singleton information set. This subgame has a unique Nash equilibrium of (U, r). Therefore the unique subgame-perfect equilibrium of the entire game is (B, U, r). Every information set is on the path, and Bayes' Rule implies $p = 1$. This strategy-belief profile $(B, U, r; p = 1)$ satisfies Bayes Requirements 1, 2, and 3.

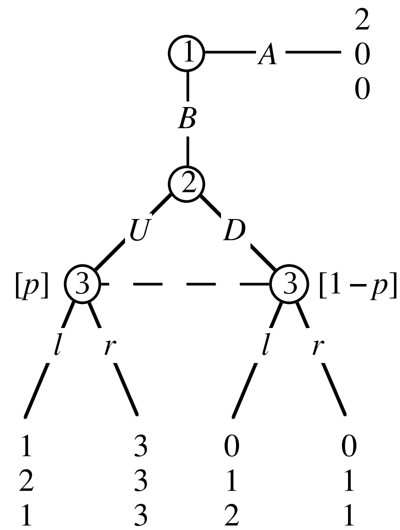

Figure 5: Requirements 1, 2, and 3 do not guarantee subgame perfection.

But now consider the strategy-belief profile $(A, U, l; p=0)$. This is a Nash equilibrium and satisfies Bayes Requirements 1, 2, and 3. (Note that player 3's information is off-the-path, and therefore Requirement 3 puts no restriction on p.) Yet this profile is not subgame perfect, because we have already seen that subgame perfection requires (U, r) by players 2 and 3. The problem with this profile can be traced to the beliefs at player 3's information set. The only way that player 3's information set could be reached is if player 1 chose B. In this case, according the strategy profile, player 2 would choose U. Therefore player 3, conditional on reaching her information set, should infer that she is located at her left-hand node and therefore should believe that $p=1$ rather than $p=0$.

We now add an additional requirement, which eliminates the non–subgame-perfect equilibrium just considered.

Bayes Requirement 4 The beliefs at any off-the-path information set must be determined from the strategy profile according to Bayes' Rule whenever possible.

Bayes Requirements 1, 2, 3, and 4 taken together constitute the definition of perfect Bayesian equilibrium in extensive-form games. (For the sender-receiver games we studied earlier, Bayes Requirement 4 has no refining power. Therefore Bayes Requirements 1, 2, and 3 constitute the definition of perfect Bayesian equilibrium in sender-receiver games.)

As we saw with sender-receiver games, perfect Bayesian equilibrium can admit equilibria which are objectionable because they rely on off-the-path beliefs which are in some sense suspect. Consider the game in Figure 6. There are two classes of pure-strategy perfect Bayesian equilibria: 1 $(U, l; p=1)$ and 2 $\{(A, r; p) : p \in [0, \frac{1}{2}]\}$.

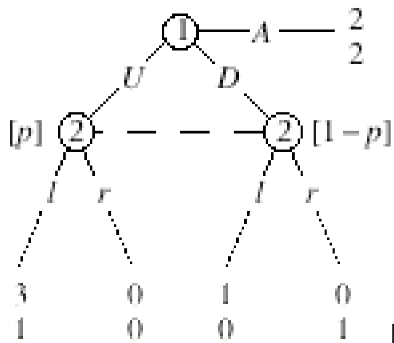

	l	r
A	2,2	2,2
U	3,1	0,0
D	1,0	0,1

□□□□

Figure 6: Perfect Bayesian equilibrium can be usefully refined.

In class 2 of these equilibria, player 2's information set is off-the-path. Player 2's beliefs at this information set are that, given that this information set was reached, there is a positive probability that the information set was reached via player 1 choosing D. But note from the strategic form that D is dominated for player 1 by A. On the other hand, U is not dominated. If player 2 observes that, contrary to his expectation, player 1 did not choose A, what should player 2 believe about player 1's actual choice? Did player 1 choose U or D?

Player 1 could never profit by playing D instead of the dominating strategy A. However, player 1 could consider playing U in hopes that player 2 would play l, giving player 1 a payoff of 3 instead of 2. Therefore we should attach zero weight to the event that a deviation by player 1 was D rather than U. Therefore we should require, when player 2's information set is off-the-path, that $p=1$.

A Refinement

If possible, each player's beliefs off the equilibrium path should put zero weight on nodes which can only be reached if another player plays a strategy that is strictly dominated beginning at some information set.