Perfect Bayesian Equilibrium

When players move sequentially and have private information, some of the Bayesian Nash equilibria may involve strategies that are not sequentially rational.

The problem is that there are usually no proper subgames. That means that all BNE are subgame perfect. We need to modify the idea of subgame perfection so that we are able to evaluate sequential rationality at all information sets.

The following version of the Gift Game is a good illustration. Here, player 2 prefers the gift to be coming from a friend, but she would rather accept a gift from an enemy than to refuse the gift.

Gift Game 2

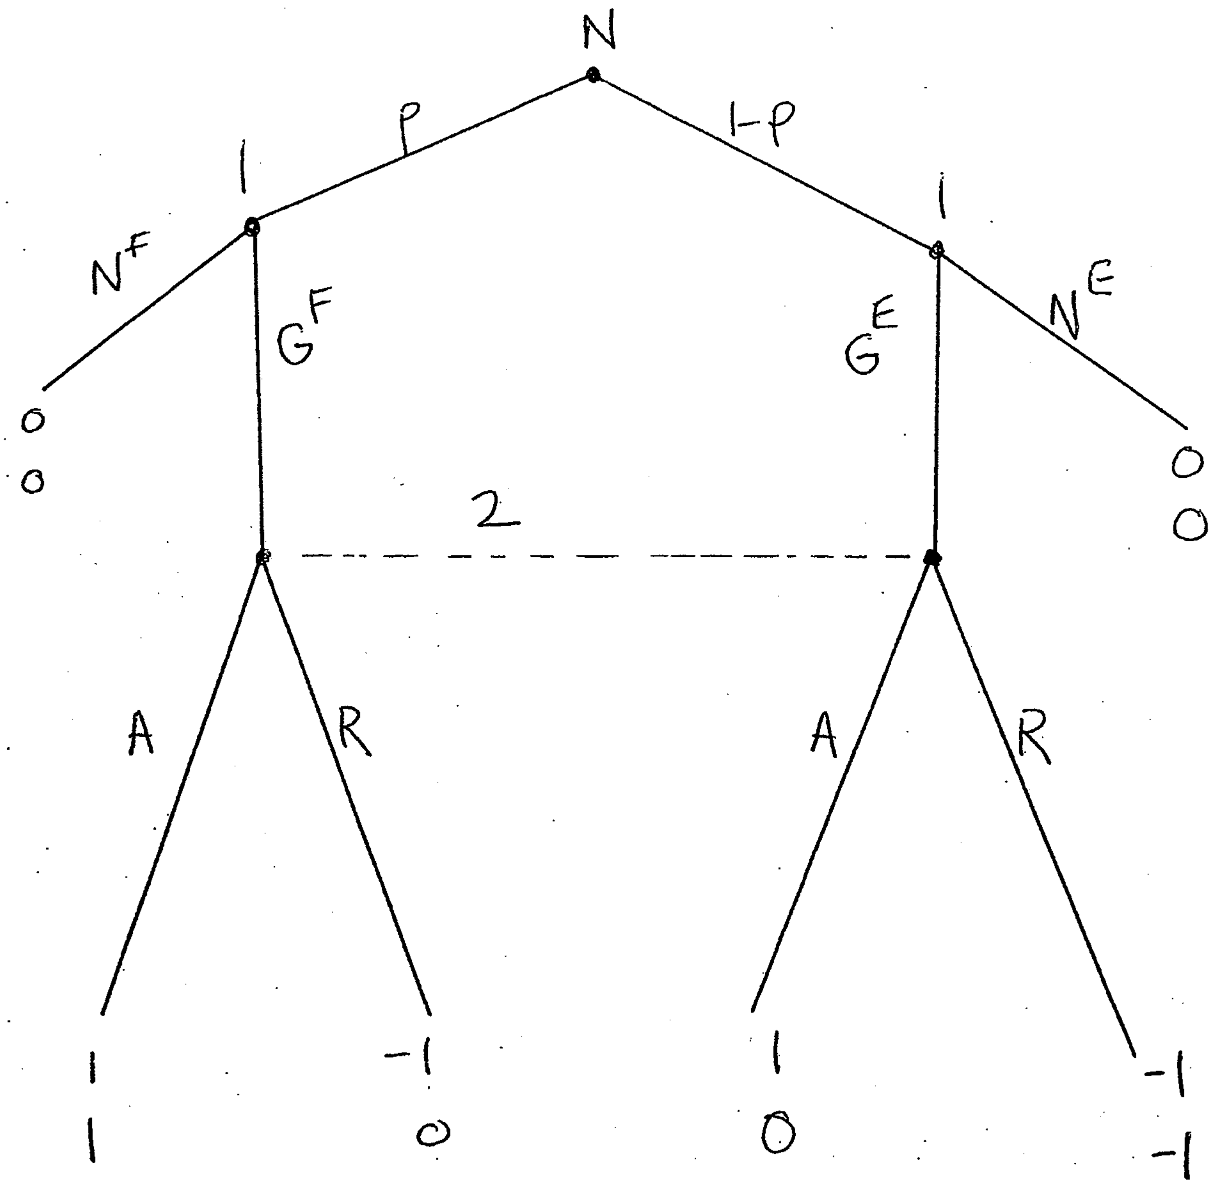

For this game, $(N^F N^E, R)$ is a BNE. Since there are no proper subgames, this is subgame perfect.

Notice that it is clearly irrational for player 2 to refuse a gift **once it is offered**, because her payoff from accepting is always greater than her payoff from refusing.

This strategy profile is not ruled out as a SPNE, because player 2's information set does not start a subgame, and switching to A in the full game does not improve her payoff, since her information set is not reached.

But how to reject this equilibrium? The concept of *Perfect Bayesian Equilibrium* (PBE) addresses this problem. A PBE combines a strategy profile and *conditional beliefs* that players have about the other players' types at every information set. This will allow us to evaluate sequential rationality by computing the expected payoff of every continuation strategy at every information set.

Conditional Beliefs about Other Players' Types

In Gift Game 2, player 2 has initial beliefs that player 1 is type F with probability p and type E with probability $1 - p$.

Notice that these are beliefs about type, *not* beliefs about player 1's strategy that we talked about earlier in the course.

Player 2's belief *conditional* on reaching her information set depends on Nature's probabilities and player 1's strategy.

For example, if player 1's strategy is $N^F G^E$, then being offered a gift causes her to *update* her beliefs about player 1's type. Her conditional beliefs are that player 1 is type E with probability one.

If player 1's strategy is $G^F N^E$, then her conditional beliefs are that player 1 is type F with probability one.

If player 1's strategy is $G^F G^E$, then there is no new information revealed by having a gift offered. Her conditional beliefs are that player 1 is type F with probability p and type E with probability $1-p$.

If player 1's strategy is $N^F N^E$, being offered a gift is a "surprise," but player 2 still should have some beliefs conditional on the "surprise" offer of a gift. More on that later.

Sequential Rationality

Specifying beliefs about the other players' types, *conditional on reaching each information set*, allows us to evaluate each player's best response to the strategy profile of the other players at every information set, even at information sets that are not reached given those strategies.

Equivalently, beliefs can be about the various nodes in an information set, conditional on reaching that information set.

Consider again Gift Game 2. Let q denote player 2's probability assessment (belief) that player 1 is type F, conditional on reaching the information set in which she is offered a gift. Then player 2's belief that player 1 is type E is $1 - q$.

It is easy to see that player 2's payoff from A is greater than her payoff from R, for any value of q. Thus, the sequentially rational action is A, no matter what her beliefs. This rules out the $(N^F N^E, R)$ BNE.

In the original Gift Game, the sequentially rational strategy for player 2 now depends on her beliefs about player 1's type.

Again let q denote player 2's belief that player 1 is type F, conditional on reaching the information set in which she is offered a gift. If she rejects the offer, her payoff is 0.

If player 2 accepts the offer, her expected payoff is $q \cdot 1 + (1-q) \cdot (-1) = 2q - 1$. This expected payoff is greater than zero if and only if $q > \frac{1}{2}$ (that is, she believes that player 1 is more likely to be a friend).

Thus, the sequentially rational action is A if and only if $q > \frac{1}{2}$.

We can convert the game from a game of incomplete information into a game of imperfect information by modeling Nature as a player who selects player 1's type. Here is the game in Bayesian extensive form:

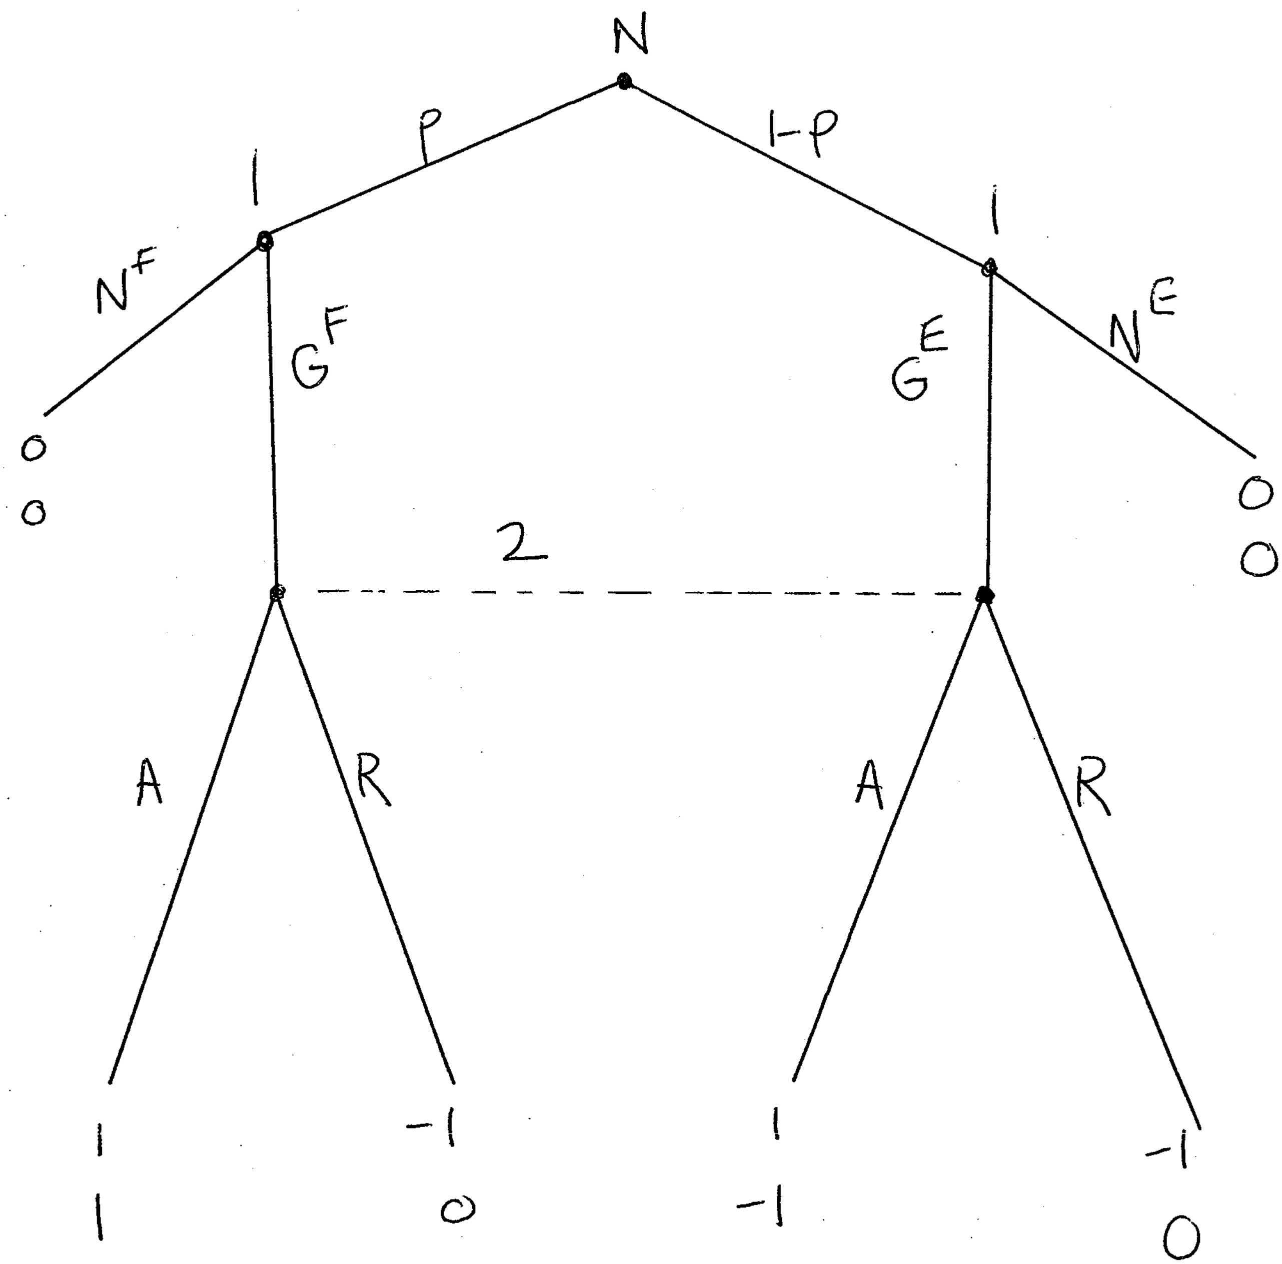

Consistency of Beliefs

For a given strategy profile, are all possible beliefs consistent with rational play? No, rational players use Bayes' rule "whenever possible."

For example, in the Gift Game, we have

$$\begin{aligned} q &= pr(F|\text{gift offer}) \\ &= \frac{p \cdot pr(\text{gift offer}|F)}{p \cdot pr(\text{gift offer}|F) + (1-p) \cdot pr(\text{gift offer}|E)}. \end{aligned}$$

When player 1's strategy is $N^F G^E$, the formula becomes

$$q = \frac{p \cdot 0}{p \cdot 0 + (1-p) \cdot 1} = 0.$$

When player 1's strategy is $G^F G^E$, the formula becomes

$$q = \frac{p \cdot 1}{p \cdot 1 + (1-p) \cdot 1} = p.$$

We can use Bayes' rule to update beliefs when player 1 uses a mixed strategy. Suppose that player 1 offers a gift with probability α^F when type F, and he offers a gift with probability α^E when type E.

Then player 2's belief is

$$
\begin{aligned}
q &= pr(F|\text{gift offer}) \\
&= \frac{p \cdot pr(\text{gift offer}|F)}{p \cdot pr(\text{gift offer}|F) + (1-p) \cdot pr(\text{gift offer}|E)} \\
&= \frac{p \cdot \alpha^F}{p \cdot \alpha^F + (1-p) \cdot \alpha^E}.
\end{aligned}
$$

What should player 2 believe if player 1 never offers a gift (so $\alpha^F = \alpha^F = 0$)? The formula gives $q = \frac{0}{0}$, so that we cannot evaluate the expression. [Note: 0/0 is not zero; it is indeterminate.]

Bayes' rule cannot be applied in this case, so we consider any beliefs to be consistent.

A Perfect Bayesian Equilibrium is a strategy profile and a specification of beliefs that each player has about the other players' types.

Definition: Consider a strategy profile for all players, as well as beliefs about the other players' types at all information sets. This strategy profile and belief system is a Perfect Bayesian Equilibrium (PBE) if:

(1) sequential rationality–at each information set, each player's strategy specifies optimal actions, given her beliefs and the strategies of the other players, and

(2) consistent beliefs–given the strategy profile, the beliefs are consistent with Bayes' rule whenever possible.

Note: In Watson's definition, beliefs are about the probability of each of the nodes in an information set, conditional on reaching that information set. This is equivalent to beliefs about types.

Here is the best way to find all of the PBE of a Bayesian extensive form game.

1. Convert the game into Bayesian normal form by constructing the matrix. Find all of the Bayesian Nash equilibria from the matrix.

2. Consider the Bayesian Nash equilibria, **one at a time**. Use Bayes' rule to determine the consistent beliefs at all information sets that occur with positive probability, given the strategy profile.

3. For information sets that are never reached, given the strategy profile, find the beliefs that make the continuation strategy sequentially rational.

Signaling Games

The PBE solution is well-suited as a solution to signaling games, where player 1 observes some information (his type) and takes an action. Player 2, who does not observe player 1's type directly but observes his action, updates her beliefs about player 1's type, and takes an action herself.

If player 1 has two possible types, then pure-strategy PBE are either separating or pooling.

In a *separating equilibrium*, player 1's types choose different actions, so player 2 will be able to infer player 1's type by observing his action.

In a *pooling equilibrium*, player 1's types choose the same action, so player 2's updated beliefs about player 1's type (after observing his action) are the same as her prior beliefs.

If player 1 has more than two possible types, then "partial pooling" equilibria are also possible.

Examples of signaling games include:

1. Our two versions of the Gift Game.

2. Job-Market Signaling. Player 1 is a job applicant of either high or low productivity. Player 2 is an employer who seeks to offer the applicant a competitive wage equal to player 1's expected productivity. High productivity types may have an incentive to undertake a costly activity (get an MBA) that does not directly enhance productivity. Because the MBA is more costly for the low types than the high types, we have a separating equilibrium.

3. Advertising. Player 1 is a firm whose product quality is either high or low. A high quality firm may have an incentive to engage in costly advertising to signal its quality. Advertising gets consumers to try the product, but a high quality firm receives repeat purchases while a low quality firm does not.

4. Cheap-Talk Games. Player 1 is an expert who observes a piece of information crucial for player 2's decision. Player 1 could be an entrepreneur with an idea for a startup venture, and player 2 could be a venture capitalist deciding how much money to invest. Only player 1 knows the true value of the project. Player 1 gives some advice about how much money should be invested, and player 2 makes a decision. Because payoffs depend on the true value of the project and the money invested, the advice itself is "cheap talk." If player 1 has a bias in favor of higher investment but there is some degree of common interest, there may be PBE in which some information is credibly revealed.

5. The Beer and Quiche Game. The title of this game is based on the book in the 1980's, "Real Men Don't Eat Quiche."

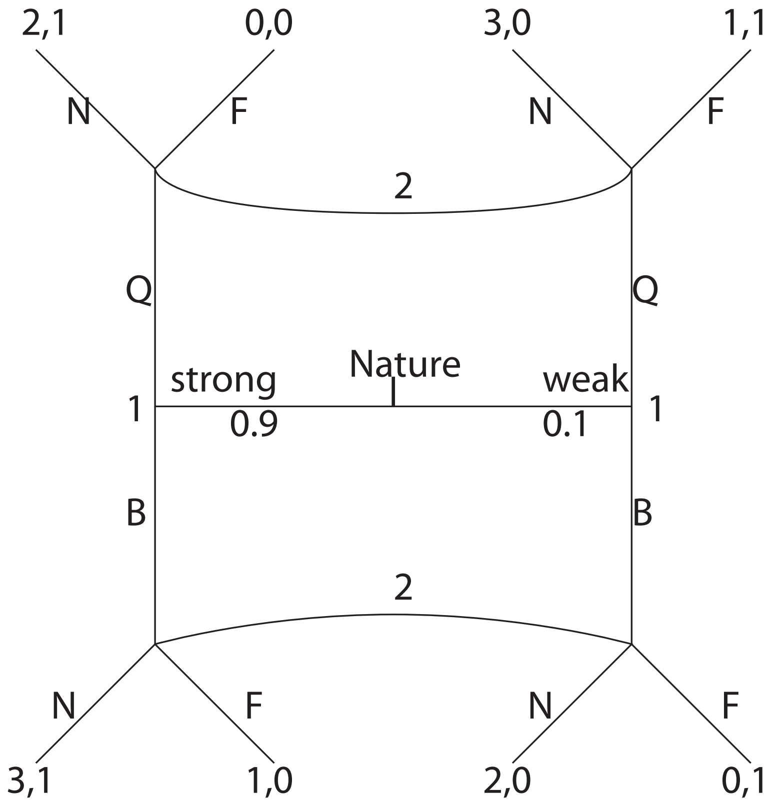

Beer and Quiche

First, notice that there cannot be a separating equilibrium. If the strong and weak types choose different actions, player 2 will infer player 1's type correctly and fight the weak type. Therefore, the weak type is not best-responding.

There are two classes of pooling PBE.

Equilibrium 1: Strong and weak types of player 1 choose B. Player 2 fights if he observes Q, but not if he observes B. Beliefs: Player 2 uses Bayes' rule if he observes B, and believes that player 1 is strong w.p. 0.9. Player 2 believes that player 1 is at least as likely to be weak as strong if he observes Q. Any belief that assigns probability of at least one-half to the weak type will do.

Given the beliefs, player 2's strategy is sequentially rational. The beliefs are consistent. Given player 2's strategy, player 1's strategy is seqentially rational. A weak player 1 receives a payoff of 2, but deviating to Q would give him a payoff of 1, due to the fact that player 2 is prepared to fight.

Equilibrium 2: Strong and weak types of player 1 choose Q. Player 2 fights if he observes B, but not if he observes Q. Beliefs: Player 2 uses Bayes' rule if he observes Q, and believes that player 1 is strong w.p. 0.9. Player 2 believes that player 1 is at least as likely to be weak as strong if he observes B.

Given the beliefs, player 2's strategy is sequentially rational. The beliefs are consistent. Given player 2's strategy, player 1's strategy is seqentially rational. A strong player 1 receives a payoff of 2, but deviating to B would give him a payoff of 1, due to the fact that player 2 is prepared to fight.

This game illustrates that our requirement that beliefs be consistent sometimes still allows some "weird" beliefs off the equilibrium path, where Bayes' rule does not apply.

Here is the Three-Card Poker game in Bayesian extensive form.

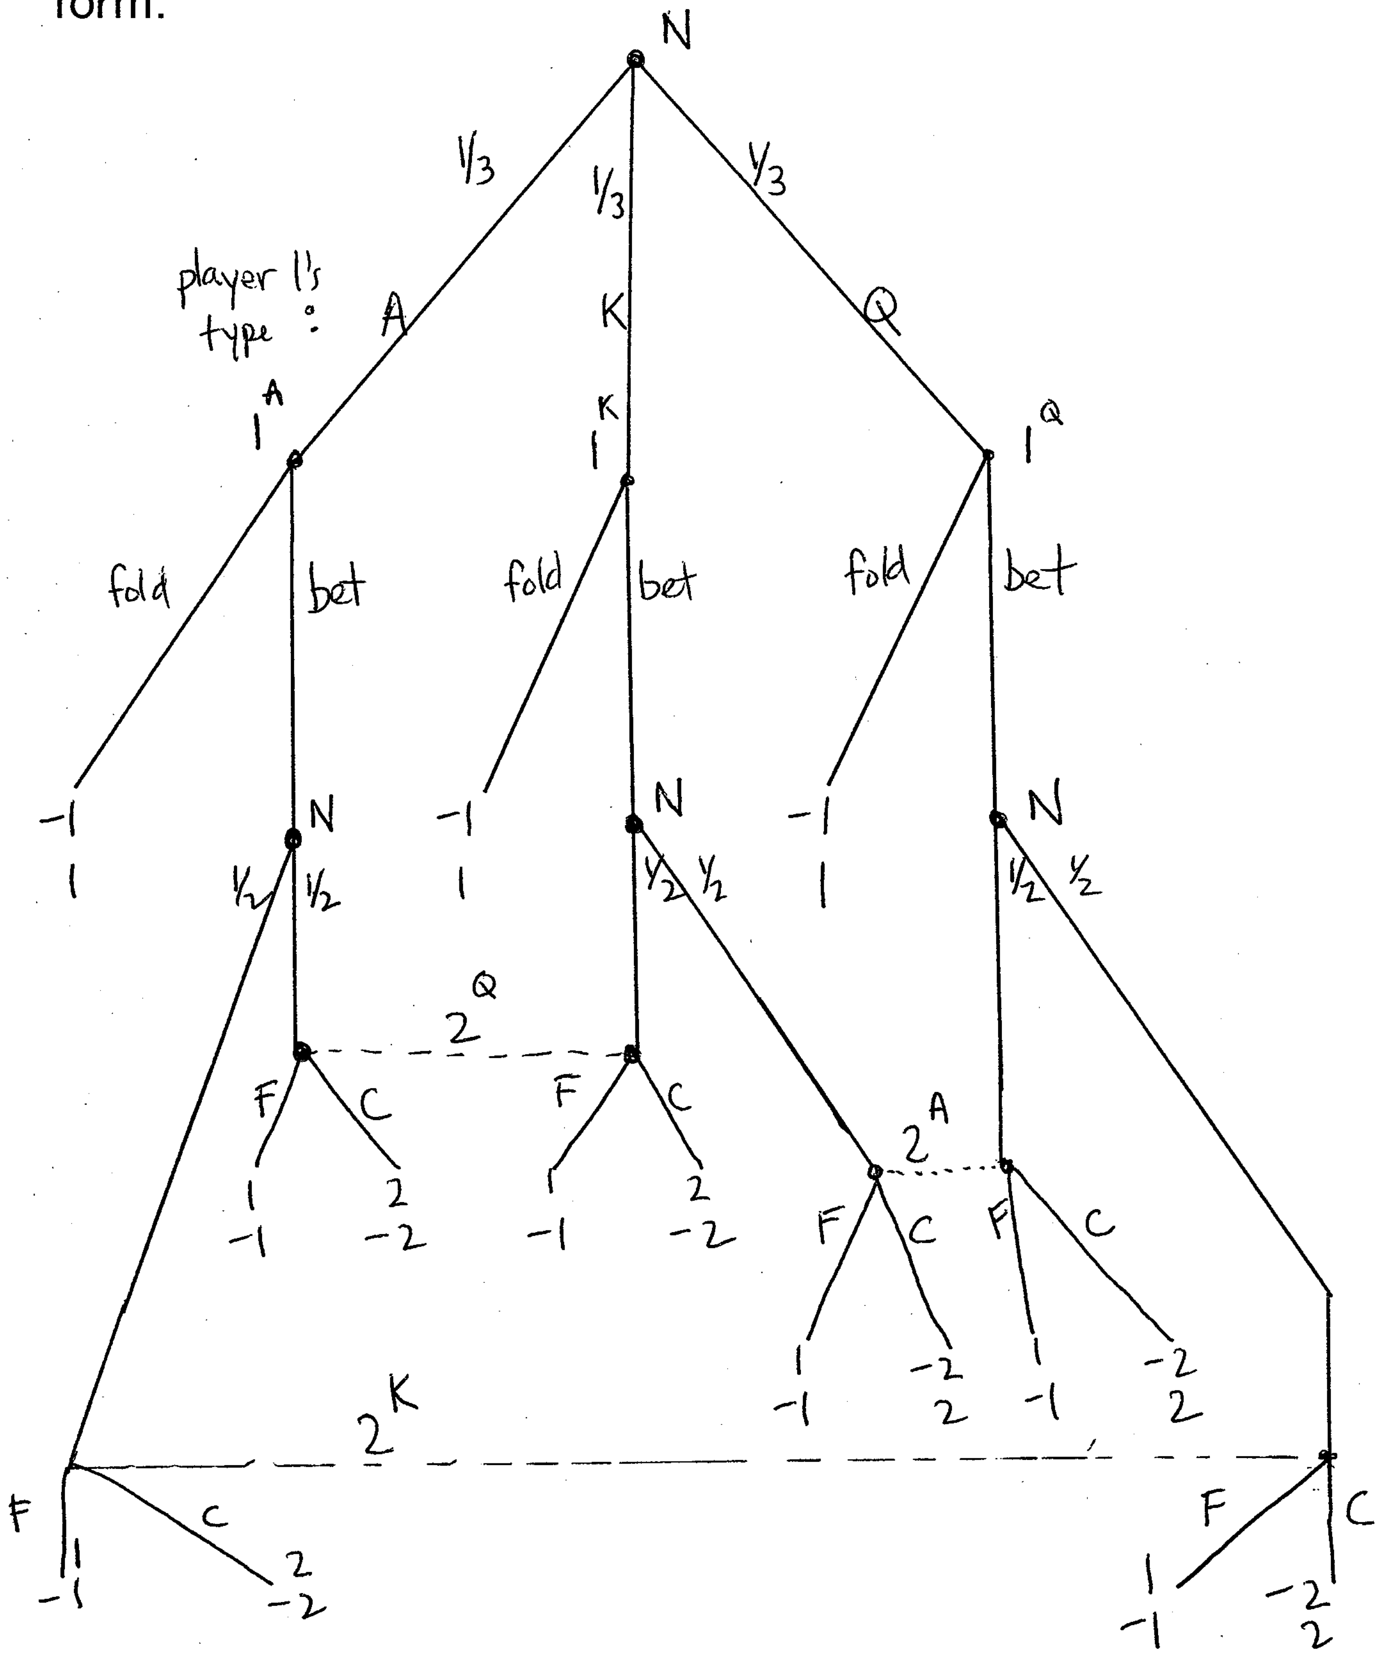

Example of PBE: Solving the 3 Card Poker Game

We will denote player 1 with an ace as type 1A, player 2 with a queen as type 2Q, etc.

We will find the PBE by figuring out which types of which players choose pure actions, and which types will be mixing. Then we can use Bayes' rule to determine beliefs and the mixing probabilities.

First, we know that player 1A will always bet his ace.

We also know that player 2A will always call with her ace, and that player 2Q will always fold with her queen.

What will player 1K choose?

Conditional on the information set in which player 1's card is a king, his beliefs about player 2's type is that she is type 2A with probability one half and type 2Q with probability one half.

Suppose player 1K bets. We know that when player 2 is type 2A, she will call and player 1 will receive -2, and when player 2 is type 2Q, she will fold and player 1 will receive $+1$. Therefore, the expected payoff from betting is

$$\frac{1}{2}(-2) + \frac{1}{2}(1) = -\frac{1}{2}.$$

Since the expected payoff from folding is -1, betting is the sequentially rational choice, given his beliefs and what we know about player 2's strategy.

What will player 1Q choose?

I will argue that in any PBE, player 1Q must mix between folding and betting.

If player 1Q always folds, then the sequentially rational choice for player 2K is to always fold, since player 2K would know that player 1 is type 1A when her information set is reached. But if player 2K always folds, player 1Q does not want to fold, since her expected payoff from betting would be $-\frac{1}{2}$ and her payoff from folding would be -1. This is a contradiction, so player 1Q cannot always fold.

If player 1Q always bets, then player 2K (from Bayes' rule) must believe that player 1 is equally likely to be type 1A or 1Q, so she receives an expected payoff of zero by calling and a payoff of -1 by folding. Therefore, the sequentially rational choice for player 2K is to call. But if player 2K always calls, then player 1Q does not want to bet. This is a contradiction, so player 1Q cannot always bet.

Thus, in the PBE, player 1Q must be indifferent between "bet" and "fold," and choose to bet with some probability, p, and fold with probability $1 - p$.

By the same argument, player 2K must be indifferent between "fold" and "call," and choose to call with some probability, q, and fold with probability $1 - q$.

Part of the equilibrium of this game involves player 1Q bluffing with what he knows is a losing hand, and it involves player 2K calling when she can only beat a bluff, to "keep player 1 honest."

We will now solve for the probabilities, p and q.

Let us compute the expected payoff for player 1Q when he bets. With probability $\frac{1}{2}$, player 2 is type 2A and always calls the bet, in which case player 1Q receives -2.

Also with probability $\frac{1}{2}$, player 2 is type 2K and calls with probability q and folds with probability $1-q$. Then the probability of player 2 being type 2K **and** calling (player 1Q receiving a payoff of -2) is $\frac{1}{2}q$ and the probability of player 2 being type 2K **and** folding (player 1Q receiving a payoff of 1) is $\frac{1}{2}(1-q)$.

Player 1Q's expected payoff when he bets is therefore

$$\frac{1}{2}(-2) + \frac{1}{2}q(-2) + \frac{1}{2}(1-q)(1)$$

For player 1Q to be indifferent between betting and folding, we have

$$\begin{aligned} \frac{1}{2}(-2) + \frac{1}{2}q(-2) + \frac{1}{2}(1-q)(1) &= -1 \\ q &= \frac{1}{3}. \end{aligned}$$

As we just showed, player 1Q's indifference condition determines the probability that player 2K calls the bet, which is $\frac{1}{3}$ of the time.

To find the probability with which player 1Q bets, we impose the condition that player 2K is indifferent between calling and folding.

The expected payoff of player 2K when she folds is -1.

The expected payoff of player 2K when she calls the bet is

$pr(\text{1A}|\text{2K and 1 bets})(-2) + pr(\text{1Q}|\text{2K and 1 bets})(2).$

To proceed, we need to determine which beliefs about player 1's type satisfy our consistency requirement. We must use Bayes' rule to find the consistent beliefs that player 2K has about player 1's type, conditional on reaching her information set.

To find the beliefs of 2K when 1 bets, we have

$$pr(\text{1A}|\text{2K and 1 bets}) =$$

$$\frac{pr(\text{2K and 1 bets}|\text{1A})pr(\text{1A})}{pr(\text{2K and 1 bets}|\text{1A})pr(\text{1A}) + pr(\text{2K and 1 bets}|\text{1Q})pr(\text{1Q})}$$

Substitute into the above expression:

$$\begin{aligned} pr(\text{1A}) &= pr(\text{1Q}) = \frac{1}{3} \\ pr(\text{2K and 1 bets}|\text{1A}) &= \frac{1}{2} \\ pr(\text{2K and 1 bets}|\text{1Q}) &= \frac{1}{2}p, \end{aligned}$$

yielding

$$pr(\text{1A}|\text{2K and 1 bets}) = \frac{\frac{1}{2} \cdot \frac{1}{3}}{\frac{1}{2} \cdot \frac{1}{3} + \frac{1}{2}p \cdot \frac{1}{3}} = \frac{1}{1+p}.$$

The remaining probability must be that player 1 is type 1Q, so

$$pr(\text{1Q}|\text{2K and 1 bets}) = \frac{p}{1+p}.$$

Now we can compute the expected payoff of player 2K when she calls the bet,

$$\begin{aligned} & pr(\text{1A}|\text{2K and 1 bets})(-2) + pr(\text{1Q}|\text{2K and 1 bets})(2) \\ = \ & \frac{1}{1+p}(-2) + \frac{p}{1+p}(2). \end{aligned}$$

Setting this payoff equal to the payoff from folding, -1, we can solve for p to get

$$p = \frac{1}{3}.$$

Player 1Q bluffs by betting with his queen one third of the time.

Recapping, the PBE strategy profile is given by

player 1: type 1A bets, type 1K bets, type 1Q bets w.p. $\frac{1}{3}$ and folds w.p. $\frac{2}{3}$.

player 2: type 2A calls, type 2K calls w.p. $\frac{1}{3}$ and folds w.p. $\frac{2}{3}$, type 2Q folds.

The PBE beliefs are given by

player 1:

type 1A believes player 2 is type 2K w.p. $\frac{1}{2}$ and type 2Q w.p. $\frac{1}{2}$.

type 1K believes player 2 is type 2A w.p. $\frac{1}{2}$ and type 2Q w.p. $\frac{1}{2}$.

type 1Q believes player 2 is type 2A w.p. $\frac{1}{2}$ and type 2K w.p. $\frac{1}{2}$.

player 2 (at her information sets following player 1 betting):

type 2A believes player 1 is type 1K w.p. $\frac{3}{4}$ and type 1Q w.p. $\frac{1}{4}$.

type 2K believes player 1 is type 1A w.p. $\frac{3}{4}$ and type 1Q w.p. $\frac{1}{4}$.

type 2Q believes player 1 is type 1A w.p. $\frac{1}{2}$ and type 1K w.p. $\frac{1}{2}$.

Which player would you rather be? One can compute that the ex ante expected payoff for player 1 in the PBE is $-\frac{1}{9}$, and the ex ante expected payoff for player 2 is $\frac{1}{9}$.