Incomplete Information

We have already discussed extensive-form games with imperfect information, where a player faces an information set containing more than one node.

So far in this course, asymmetric information arises only when players do not observe the *action* choices of other players.

Not observing previous action choices is only one way that players can be asymmetrically informed. There are many other interesting sources of asymmetric information. For example, consider a game in which a house owner announces an asking price for his house and a potential buyer simultaneously announces the price she is willing to pay. A deal will be struck if the buyer is willing to pay at least as much as the seller's ask price. Also suppose that only the seller knows whether the basement leaks during a heavy rain, and only the potential buyer knows what she can afford.

A *Bayesian game of incomplete information* is a game that incorporates incomplete information. The general definition has some confusing notation, so I will just give a description here.

A Bayesian game of incomplete information consists of: (i) a set of players, (ii) a set of possible actions for each player, (iii) a set of possible types for each player, (iv) a probability distribution that specifies the probability that each player is of each of his/her possible types, (v) a function for each player that assigns payoffs, based on the action profile and the realized types of all the players $u_i(a_1, ..., a_n, t_1, ..., t_n)$.

The interpretation is that each player i observes her type before choosing her action, so her overall strategy is a function, $a_i(t_i)$. Because types are uncertain, players seek to maximize their expected payoff.

The house-selling game can be seen as a game of *incomplete information*. The actual payoffs received by the players depend on the "types" of the players as well as their actions. The type of each player is selected by "Nature" according to a probability distribution. Each player observes his or her type, but the other player does not.

For example, the seller might be the leaky-basement type with probability 0.6 and the dry-basement type with probability 0.4.

Since payoffs depend on the types of the players, the players are not sure about the payoff matrix of the game they are playing!

We will go over examples that show how to go from a Bayesian game of incomplete information to a standard normal-form game which we call a Bayesian normal-form game. (The word "Bayesian" is named after a mathematician named Bayes who worked with conditional probabilities.)

The examples will also show how to convert any extensive-form game with incomplete information into a standard extensive-form game, which we call a Bayesian extensive-form game, by including moves of "Nature" in the game tree.

By including Nature in the game tree, the players now know their payoff at every terminal node. This converts the game from one of *incomplete* information to one of complete but *imperfect* information. This is useful, because we already have solution concepts like Nash equilibrium for games of imperfect information. (This insight is due to Nobel Prize winner John Harsanyi.)

Example: An Entry Game with Cost Uncertainty

Two firms must choose whether or not to enter a market. The fixed cost of entry for firm i is c_i.

Monopoly revenues are 1, and duopoly revenues are 0. (Note–We can think of a lone entrant setting the monopoly price and if both firms enter, Bertrand competition leads to a price equal to marginal production cost.)

		firm 2 E	N
firm 1	E	$-c_1, -c_2$	$1 - c_1, 0$
	N	$0, 1 - c_2$	$0, 0$

Firm 2's cost is known to be $c_2 = \frac{1}{4}$. Firm 1's cost is uncertain. With probability one half, firm 1 is a low cost type (type L) with $c_1 = 0$, and with probability one half, firm 1 is a high cost type (type H) with $c_1 = \frac{1}{2}$.

Firm 1 learns its type before deciding whether to enter or not. Firm 2 does not know firm 1's type when it decides whether to enter, so it does not know the actual payoff matrix.

If firm 1 is type L, the payoff matrix is

		firm 2 E	N
firm 1	E^L	$0, -\frac{1}{4}$	$1, 0$
	N^L	$0, \frac{3}{4}$	$0, 0$

If firm 1 is type H, the payoff matrix is

		firm 2 E	N
firm 1	E^H	$-\frac{1}{2}, -\frac{1}{4}$	$\frac{1}{2}, 0$
	N^H	$0, \frac{3}{4}$	$0, 0$

The set of available actions for each player is $\{E, N\}$, but a strategy for firm 1 is a function specifying the action chosen for each type. Thus, firm 1 has 4 strategies: $E^L E^H$, $E^L N^H$, $N^L E^H$, and $N^L N^H$.

Firm 2 has only one type and two strategies: E and N.

For any profile of strategies, we can compute each player's expected payoff, where the expectation is taken before firm 1 learns its type. For example, the profile $(E^L E^H, E)$ yields firm 1 an expected payoff of $(\frac{1}{2})(0) + (\frac{1}{2})(-\frac{1}{2}) = -\frac{1}{4}$, and it yields firm 2 an expected payoff of $(\frac{1}{2})(-\frac{1}{4}) + (\frac{1}{2})(-\frac{1}{4}) = -\frac{1}{4}$.

By similar calculations, we can construct the entire payoff matrix corresponding to the Bayesian normal form:

		firm 2 E	N
	$E^L E^H$	$-\frac{1}{4}, -\frac{1}{4}$	$\frac{3}{4}, 0$
firm 1	$E^L N^H$	$0, \frac{1}{4}$	$\frac{1}{2}, 0$
	$N^L E^H$	$-\frac{1}{4}, \frac{1}{4}$	$\frac{1}{4}, 0$
	$N^L N^H$	$0, \frac{3}{4}$	$0, 0$

Notice what we have done. Starting with a 2×2 game of incomplete information where firm 2 does not know the correct payoff matrix, we considered firm 1's strategy set to be the set of functions from types into actions.

By computing the ex ante expected payoffs, we convert the game of incomplete information into a standard normal-form game.

Let us go through the same procedure, but for the game in extensive form. We will convert the Entry Game with Cost Uncertainty from a game of incomplete information into a Bayesian extensive form game, by adding a move by Nature that determines firm 1's type.

Here are the two game trees that player 2 could be facing:

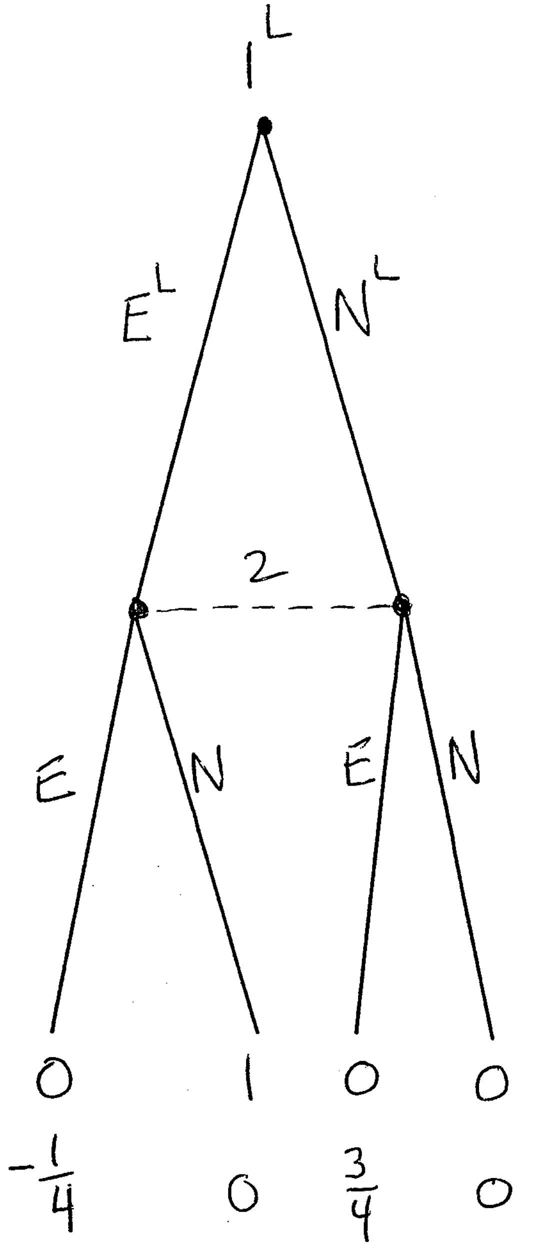
1^L
E^L
N^L
2
E
N
E
N
0
-1/4
1
0
0
3/4
0
0

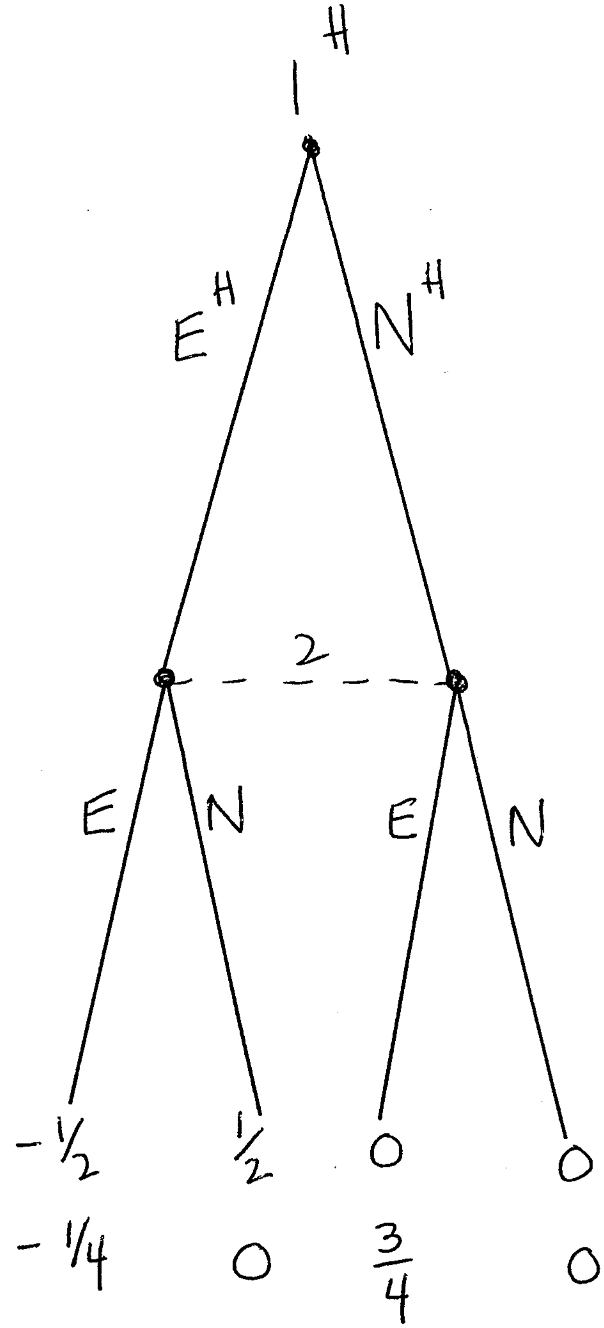
1^H
E^H
N^H
2
E
N
E
N
-1/2
-1/4
1/2
0
0
3/4
0
0

Here is the Entry game in Bayesian extensive form:

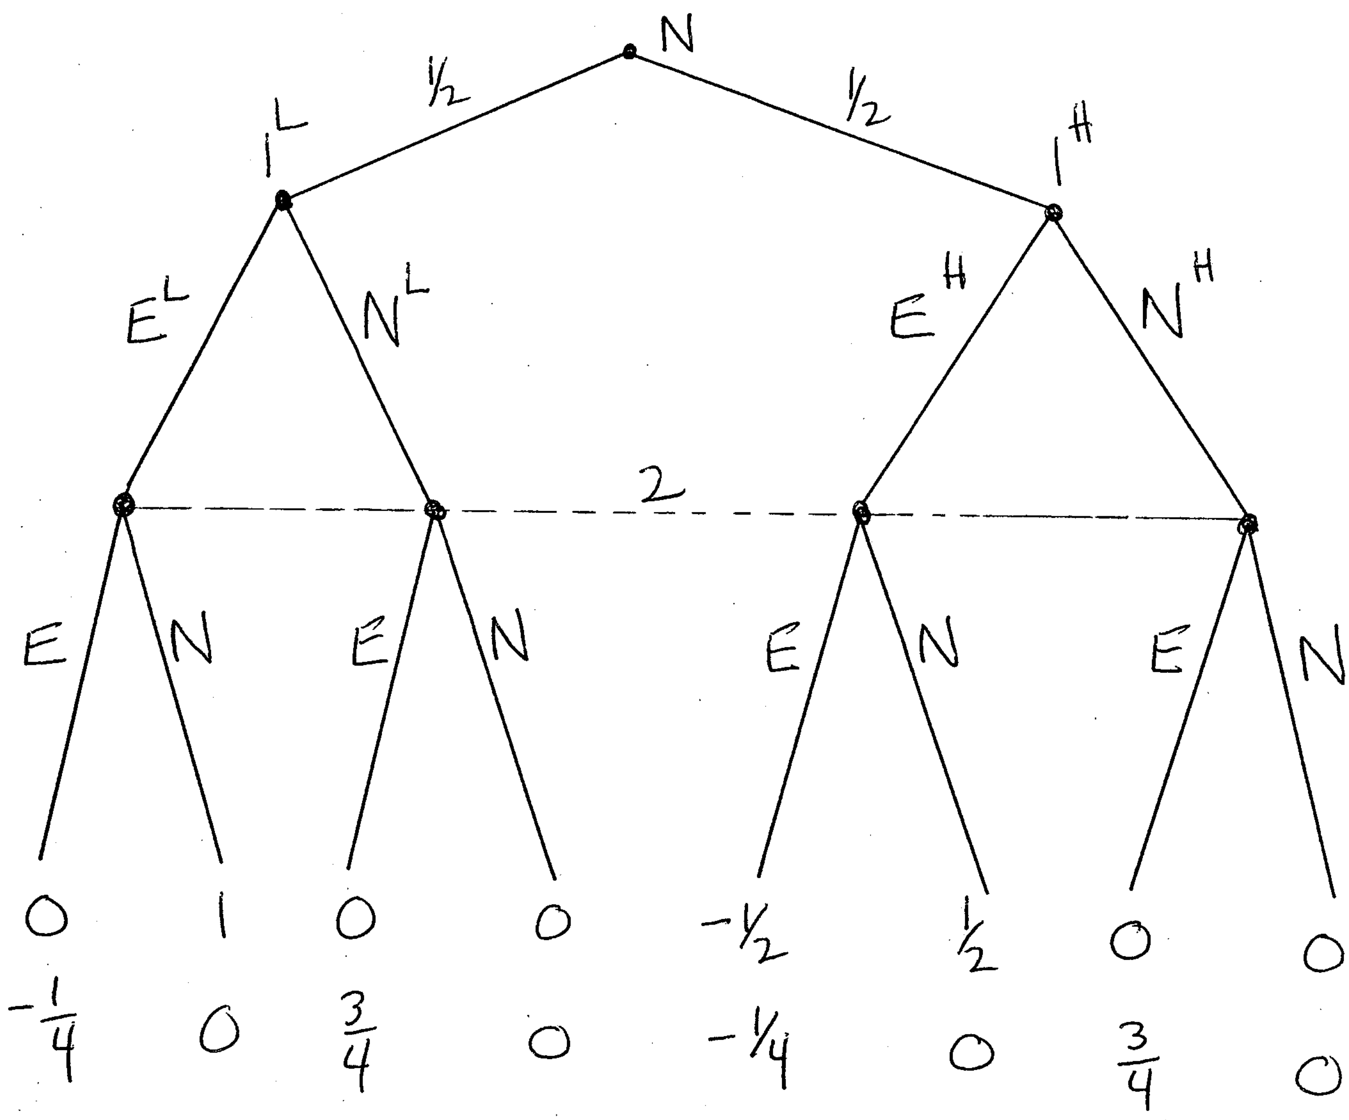

In the game tree, Nature moves first, selecting firm 1's type to be type L with probability $\frac{1}{2}$ and type H with probability $\frac{1}{2}$. These probabilities are part of the specification of the game and should be written next to the corresponding branches.

Nature is not a real player and does not have payoffs.

Based on the way the information sets are drawn, we can see that player 1 observes nature's move and player 2 does not. [Note: It is possible to have a game of perfect information with moves by nature–for example, Backgammon, Craps, or Monopoly.]

Example: The Gift Game

Player 1 decides whether or not to offer a gift to player 2.

If no gift is offered, the game ends; if player 1 offers a gift, she must decide whether to accept or reject the gift.

Player 1 either considers himself to be player 2's friend (type F) or player 2's enemy (type E), and knows his type. Player 2 does not know player 1's type, but considers the probability of type F to be p and the probability of type E to be $1 - p$. The parameter, p, is common knowledge.

Payoffs to player 2 depend on player 1's type. Think of a friend as deciding whether to offer a package of Belgian chocolates and an enemy as deciding whether to offer a package of Belgian chocolates that is rigged to explode.

If player 1 is type F, the payoff matrix is

		player 2 A	R
player 1	N^F	$0, 0$	$0, 0$
	G^F	$1, 1$	$-1, 0$

If player 1 is type E, the payoff matrix is

		player 2 A	R
player 1	N^E	$0, 0$	$0, 0$
	G^E	$1, -1$	$-1, 0$

Here, the set of available actions for each type of player 1 is $\{N, G\}$, but a strategy for player 1 is a function specifying the action he chooses for each type. Thus, he has 4 strategies: $N^F N^E$, $N^F G^E$, $G^F N^E$, and $G^F G^E$.

Player 2 only has one type and therefore two possible strategies: A and R.

For any profile of strategies, we can compute each player's expected payoff, where the expectation is taken before player 1 learns his type. For example, the profile $(N^F G^E, A)$ yields player 1 an expected payoff of $p \cdot 0 + (1-p) \cdot 1 = 1-p$, and it yields player 2 an expected payoff of $p \cdot 0 + (1-p) \cdot (-1) = p-1$.

By similar calculations, we can construct the entire payoff matrix corresponding to the Bayesian normal form:

		player 2	
		A	R
	$N^F N^E$	$0, 0$	$0, 0$
player 1	$N^F G^E$	$1-p, p-1$	$p-1, 0$
	$G^F N^E$	p, p	$-p, 0$
	$G^F G^E$	$1, 2p-1$	$-1, 0$

Let us go through the same procedure, but for the game in extensive form. We will convert the Gift Game of incomplete information into a Bayesian extensive-form game, by adding a move by Nature that determines player 1's type.

Here are the two game trees that player 2 could be facing:

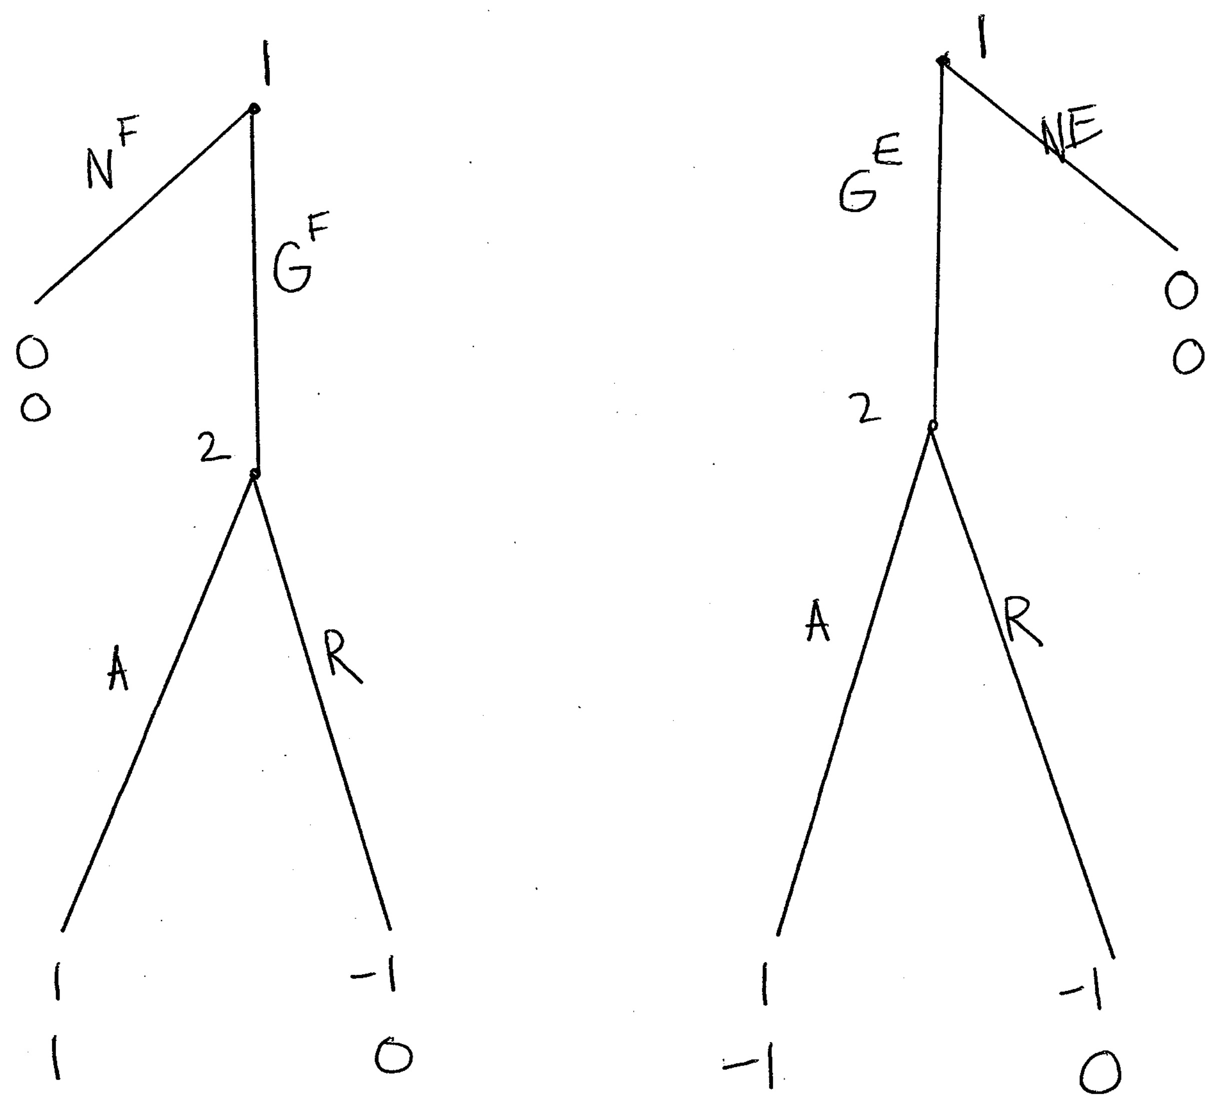

We can convert the game from a game of incomplete information into a game of imperfect information by modeling Nature as a player who selects player 1's type. Here is the game in Bayesian extensive form:

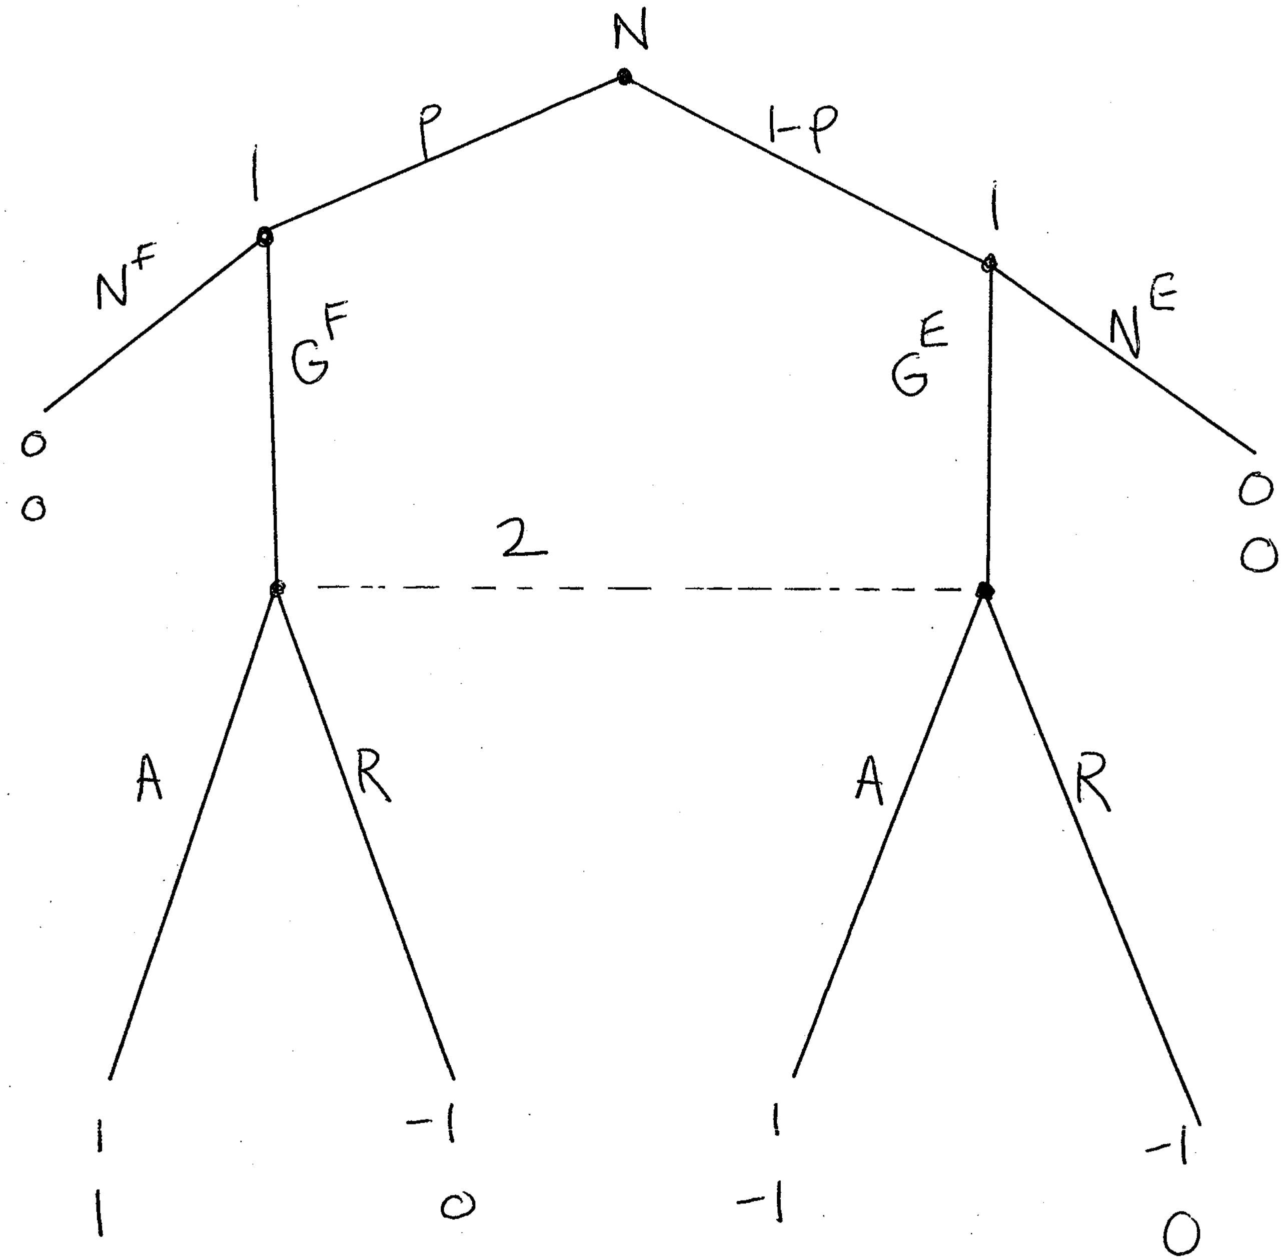

In the game tree, Nature moves first, selecting player 1's type to be type F with probability p and type E with probability $1-p$. These probabilities are part of the specification of the game and should be written next to the corresponding branches.

Based on the way the information sets are drawn, we can see that player 1 observes Nature's move and player 2 does not observe Nature's move.

Notice also that play is not simultaneous. This introduces a new element of complexity into the game, because player 2 might be able to *learn* something about player 1's type by observing his action, and issues of *sequential rationality* might come up.

In the Gift Game, a type F player 1 cannot ignore the part of the game tree in which he is type E. Because player 2 cannot observe Nature's move, whether or not she will accept the gift depends on that part of the tree as well. A rational player 1 must consider what his other types would have done.

Example: Three-Card Poker (from exercise 4 on page 333)

A deck containing an Ace, King, and Queen is shuffled and the two players each receive one card. Each player observes his/her own card but not the other player's card.

First, player 1 decides whether to fold or bet. If he folds, the game is over and payoffs are $(-1, 1)$. [Player 2 keeps her ante and wins player 1's ante of one chip.]

If player 1 bets [he puts another chip in the pot], then player 2 must decide whether to fold or call. If she folds, payoffs are $(1, -1)$. [Player 1 wins player 2's ante.]

If player 2 calls [she puts another chip in the pot], the players show their cards and the player with the higher card wins the pot; the winner receives a payoff of 2 and the loser receives a payoff of -2.

Ace is the highest ranking card, then King, then Queen.

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