

Game Theory and Algorithms*

Lecture 10: Extensive Games: Critiques and Extensions

March 31, 2011

Summary: We discuss a game called the *centipede game*, a simple extensive game where the prediction made by backwards induction contrasts sharply with what happens in practice. Then, we briefly discuss a standard model even more general than “extensive with simultaneous moves:” games of *incomplete information*. We give some examples showing how its concepts of randomization and equilibrium are more complex than those encountered previously.

1 Imps and Centipedes

The Bottle Imp is a short story by the Scottish writer Robert Louis Stevenson, from 1891. The central object of the story is a magical bottle inhabited by an imp who bestows good fortune and wealth upon its owner. There is a catch: each owner of the bottle must sell it before they die or else they will burn in hell; and each sale price must be strictly less than the previous sale price. Legend has it that the imp was the source of power for Napoleon and Captain Cook.

In the story, a man buys it, gets a beautiful wife, then sells it, and contracts leprosy. He is able to track down the current owner of the bottle but unfortunately that owner paid only 2¢ for it. Should the man buy it for 1¢ and seal his doom, or live in a leper colony? What would you do in this situation? Is it rational to buy the bottle in the first place?


We will return to this story later, but for now let us introduce a related one, the *centipede game* (introduced by Rosenthal in 1981). In the centipede game, we have two players with a joint investment.

- The two players take alternate turns caring for the investment, which starts off with a value of \$0.
- On each turn, either the player can **Invest**, increasing the value of the investment by one dollar, or else they can **Cash In**.

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

- If a player cashes in, the game ends: the value of the investment is split equally between the two players, *and the player who cashed in gets a \$1 bonus.*
- We limit the investment's value (e.g., at \$100); the next player to move once it reaches this value is forced to Cash In.

This is a finite extensive game without simultaneous moves. Let's draw its game tree — we will use a limit of \$5 to make the tree smaller, and multiply everything by 2 to remove fractions:



The game gets its name since the shape of the tree looks a little bit like a caterpillar.

Next, we would like to see what backwards induction predicts in this game.

- If the investment reaches \$10 in value, the second player is forced to cash in, giving utilities (5, 7).
- If the investment reaches \$8 in value, the first player can either cash in to get utility 6, or invest and get utility 5. So, he would cash in.
- Likewise, if the investment reaches \$6 in value, the second player will cash in since $5 > 4$.
- Continuing in this way, the only subgame perfect equilibrium is the strategy profile (always cash in, always cash in).

So the SPE behaviour would give a payoff of \$2 to player 1 and \$0 to player 2 (no matter how large the limit is).

1.1 Experiments

In experiments, players do not do what SPEs predict (even though we are in the special class “without ties” where SPEs seemed to be very logical predictors). Here are the results of one particular study [1]:

- The experiment took place in 1989 with a maximum investment value of \$6.
- The actual cash value given to participants was $\$0.10 \cdot 2^x$ instead of x .
- Players played about 10 games, never facing the same opponent twice.

Roughly, the result was a bell curve: most of the games terminated after 3 or 4 Invest moves, with a smooth decrease to a handful of games stopping immediately or going as long as possible.

1.1.1 Discussion

One fact which is important is that the players' non-equilibrium behaviour actually gives them more money than equilibrium behaviour would. In this way it is similar to what is sometimes reported when humans play a finitely-repeated Prisoners' Dilemma: there is cooperation at the start, and back-stabbing at the end.

Nonetheless, even taking into account randomization, the game has no mixed SPE other than the pure SPE we already found. So, it means that at least one of the assumptions (greed, rationality, maximizing their income in every subgame, predicting and simulating their opponent) is invalid.

It raises some interesting questions:

- Suppose player 1 Invests on her first move. Then player 2 could deduce that player 1 is not using an SPE strategy; and given that player 1 Invested once, might they not Invest again?
- One might approximate large versions of this game by an "infinite" game; in this game (always invest, always invest) is an SPE. Is it easier to do this approximate calculation rather than reason everything through?

We also note that even the weaker notion of Nash equilibria always leads to Cashing in on the first move.

1.2 Bottle Imp Conclusion

The bottle imp game leads to basically the same analysis as the centipede game, where the game ends once we hit the minimum possible monetary value. The bottle imp game illustrates that the fact that there are 2 players is immaterial; we could equally well have a different player play at each node. And for the story...

As you may recall, our hero was faced with the decision whether to become a leper, or buy the bottle for 1¢, condemning his soul to hell since he'd have to sell it for less than 1¢. Luckily, his wife suggests they sail to a colony where they use *centimes* worth 1/5 of a cent. His wife bribes a sailor to buy it for four centimes, and she buys it afterwards. The man discovers this and bribes a drunk to buy it from her for two centimes, proposing to buy it back from the drunk for 1 centime. The drunk buys the bottle, and the man tries to buy it back, politely reminding that "the man who has that bottle goes to hell." The drunk replies "I reckon I'm going anyway, and this bottle's the best thing to go with I've struck yet. No, sir! This is my bottle now, and you can go and fish for another." The man and wife live happily ever after.

2 Games of Incomplete Information

We now explain a little bit about how the model of extensive games seen previously can be extended to include poker. The relevant model is called a *game with incomplete information*.

Such a game is represented with the same sort of game tree used for an extensive game without simultaneous moves; we also require that the edges (actions) are labelled. Next, we draw *information sets* around some of the internal nodes of the tree.


- The information sets are disjoint (i.e. each internal node lies in at most one information set).
- Any two nodes in an information set must be labelled by the same player.
- Any two nodes in an information set must have the same labels on the edges coming out of them.

Then, our model is that *each player cannot distinguish nodes in the same information set*. That is to say, if nodes h, h' both lie in the same information set (say for player i), then any strategy for player i must choose the same action (label) at h and h' .

2.1 Examples

The first example shows that strategic games can be modelled as extensive games with incomplete information. We actually mentioned this in passing in the first lecture: for the taxpayer game, the citizen moves before the government, but the government can't distinguish between the taxpayer's two moves before choosing their own move.

p1 \ p2	A	DA
L	1,3	4,1
DL	2,2	3,3

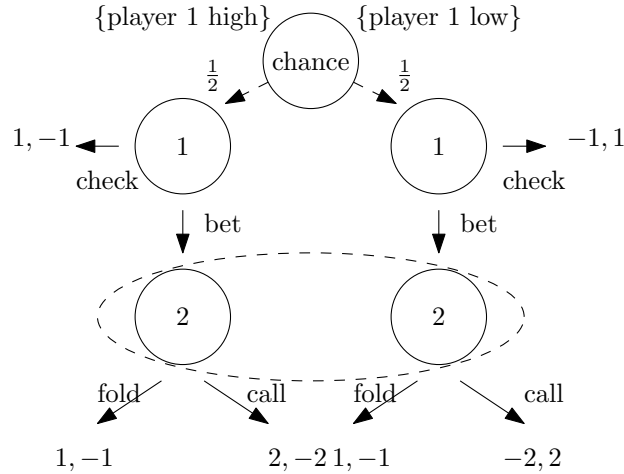


In the diagrams, we represent information sets with dashed circles.

The second example gives an example of a simplified game of poker:

- It is a zero-sum game, and the root node is controlled by external randomness.
- Both players ante \$1, and player 1 is dealt a high or low card. We imagine player 2 as always having a “middling” card.
- Player 1 looks at her card and gets to bet \$1 or check. If she checks, the game ends; she wins with a high card and loses with a low card.


- If player 1 bet, then player 2 can call or fold. Folding loses the \$1 ante immediately, while calling means that the stakes are $\pm\$2$ depending on the type of player 1's card.



For this particular game, there is a unique pair of maxminimizing strategies, as follows. Player 2 calls $2/3$ of the time. Player 1 always bets with a high card, and also bets $1/3$ of the time with a low card (intuitively, he bluffs sometimes since otherwise player 2 would never need to call).

2.2 Equilibria

Equilibrium concepts in games of incomplete information tend to be more complicated. We give one example here, adapted from Osborne & Rubinstein.



Consider the strategy profile (A, R) . Is it an equilibrium? The reasonableness of playing R for player 2 depends on her estimation of

$$\Pr[B]/(\Pr[B] + \Pr[C])$$

which is $0/0$ if player 1 is always using A .

The subgame-perfect concept of checking subgames is not so useful here, since there are no subgames. So equilibrium concepts tend to incorporate new ideas such as beliefs, or


convergent sequences of strategy profiles, or allowing errors. We mention a few names of these concepts: (weak) sequential equilibria, trembling-hand perfect equilibria, and perfect equilibria, which allow a “second most likely choice.”

2.3 Mixed Strategies

So far, we did not explicitly describe how *mixed* strategies are modelled in extensive games. There are actually two alternatives:

- A player can flip a generalized coin once at the start of the game, and then they pick a pure strategy thereafter. (Probability distribution over strategy functions.)
- At each node, a player gets to independently flip a new generalized coin, and makes just one action based on the result of that flip. (Single strategy function whose output is drawn from a distribution, called a *behavioural mixed strategy*.)

These turn out to be equivalent for all games studied before today, but are not equivalent in general. The distinction between the definitions are easiest to see with an example. It also will show that with *incomplete information*, the two concepts are different, and neither one generalizes the other.



Consider the game depicted above. Player 1 is “forgetful” whether he has already moved or not.

- If player 1 flips a coin before the game starts and thereafter picks a pure strategy A or B uniformly at random, then he gets outcome x half of the time, and outcome z half of the time.
- If player 1 flips a coin before each move and picks a uniformly random choice, then he can get to outcome x, y each a quarter of the time, and z half of the time.
- Neither model can achieve the distribution on outcomes achieved by the other model.

There is a natural sufficient condition for the equivalence of the two models: a game has *perfect recall* if for all pairs h, h' in the same information set for some player i , the sequence of prior moves made by player i are the same in h and h' . (We assume for this definition that nodes in different information sets get distinct labels for their moves.) For example, the game just pictured does not have perfect recall because of the two nodes in the same information set, one was reached by action sequence \emptyset , and the other by (A) .

One nice thing that happens with *perfect recall* is that we can work on trees in a bottom-up fashion. Moreover, we get the equivalence mentioned above.

Theorem 1. *In games with perfect recall, randomizing at the start of the game is outcome-equivalent to randomizing each choice separately.*

The proof idea is to use conditional probability.

The next topic will be mechanism design. As a segue from extensive games, we will mention *cake-cutting* games.

References

- [1] R. D. McKelvey and T. R. Palfrey. An experimental study of the centipede game. *Econometrica*, 60(4):pp. 803–836, 1992.