Harvard CS 1360, Spring 2026
Instructor: Ariel Procaccia
Scribe: Taig Singh, Jakob de Raaij

Extensive-Form Games

—

Lecture 3

1 Extensive-Form Games and Subgame-Perfect Equilibria

Definition 1 (Extensive-Form Games)**.** An *extensive-form game* is represented as a tree. Each node is labeled with a player. The game starts at the root. At each node, the node's corresponding player chooses one of the edges (corresponding to the player's strategies) to move to a child node. Once the game reaches a leaf, the players get the given utilities.

In contrast to the normal-form games from the last two lectures, the players in an extensive-form game move sequentially and possible over multiple rounds, not simultaneously.

Example 1 (Basic extensive-form game)**.** Consider the following tree representation of an extensive-form game:

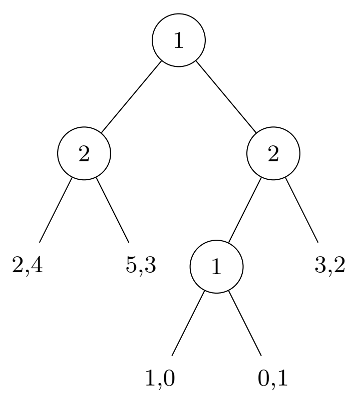

Every path from the root to a leaf in the tree represents a combination of strategies of players 1 and 2 that results in a payoff.

- If player 1 chooses LEFT and then player 2 chooses LEFT, player 1 gets payoff 2 and player 2 gets 4.
- If player 1 chooses LEFT and then player 2 chooses RIGHT, player 1 gets payoff 5 and player 2 gets 3.
- If player 1 chooses RIGHT, then player 2 chooses LEFT, and then player 1 chooses LEFT, player 1 gets payoff 1 and player 2 gets 0.
- If player 1 chooses RIGHT, then player 2 chooses LEFT, and then player 1 chooses RIGHT, player 1 gets payoff 0 and player 2 gets 1.
- If player 1 chooses RIGHT and then player 2 chooses RIGHT, player 1 gets payoff 3 and player 2 gets 2.

So how do equilibria in extensive-form games compare to normal-form games? Consider the following (very abstract) model of the Cuban Missile Crisis as a game between the US and the USSR:

Example 2 (The Cuban Missile Crisis)**.** The USSR has just positioned its missiles in Cuba. The US now has two options: It can either ignore this threat or it can respond to it. Further, if the US responds, the USSR has two options: It can either compromise or start a nuclear war. We can represent this game as an extensive-form game.

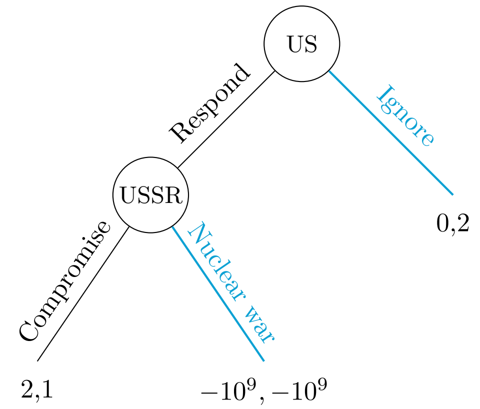

If the US chooses to ignore, the payoff is $(0, 2)$. If the US chooses to respond and the USSR chooses to compromise, the payoff is $(1, 2)$; if the US chooses to respond and the USSR chooses to engage in a nuclear war, the payoff is $(-10^9, -10^9)$.

Now, consider the equivalent normal-form representation of the game.

	Compromise	Nuclear War
Respond	$2, 1$	$-10^9, -10^9$
Ignore	$0, 2$	$0, 2$

The payoff is $(0, 2)$ if the US chooses to ignore, regardless of the USSR's strategy, as the game always ends in this leaf if the US ignores.

The strategy profile (Ignore, Nuclear War), highlighted in blue in both representations, is a Nash equilibrium of this game. However, there is a problem: The USSR's threat of a nuclear war is not credible! If the game ever ends up at the node in which the USSR must make a decision, the USSR will always choose to compromise to get a payoff of 1, avoiding the payoff of -10^9 that they will receive if they choose to engage in a nuclear war.

The lack of credibility of the USSR's threat is due to the sequential nature of the game. This is only captured in the extensive-form representation of the game, not the normal-form representation. Although (Ignore, Nuclear War) is a Nash equilibrium of this game, the USSR would never choose to engage in a nuclear war after the US has chosen to respond. This leads us to believe that the current notion of a Nash equilibrium does not work well in this setting.

Definition 2 (Subgame-Perfect Equilibrium)**.** In an extensive-form game, each subtree forms a *subgame*. A set of strategies is a *subgame-perfect equilibrium* if it is a Nash equilibrium in each subgame.

Let's apply this to the current example. There is one subtree rooted at the USSR node , corresponding to the subgame where the USSR chooses either to compromise or to engage in nuclear war. In this game, choosing compromise is the only Nash equilibrium as this is a one-player game where compromising is the dominant strategy. Now, considering the subgame corresponding to the entire game. We know that in any subgame-perfect equilibrium, the USSR will always compromise if the US chooses to respond. However, knowing that, it is the dominant strategy for the US to respond. Thus, (Respond, Compromise) is the unique subgame-perfect equilibrium.

An interesting phenomenon of subgame-perfect equilibria in extensive-form games is that agents may be able to improve their equilibrium payoff by eliminating some of their own strategies from the game. In our example, the USSR can improve their subgame-perfect equilibrium outcome by eliminating their Compromise strategy from the game. In this case, the US will always choose to ignore, resulting in a payoff of 0 for the US and 2 for the USSR, since they know that responding will result in a certain nuclear war with a payoff of -10^9. Thus, by removing their option to compromise in the case that the US respond, the USSR can increase their equilibrium payoff from 1 to 2.

Luckily, subgame-perfect equilibria are a lot easier to find than Nash Equilibria by using a technique called *backward induction*. We start at nodes that only have edges to leaves of the tree. The subgames corresponding to the subtrees rooted at these notes are 1-player games, so we know that any edge that maximizes the payoff for the player of this node is a (Nash) equilibrium in this subgame. Now, knowing what strategies the players in nodes that only have edges to leaves are going to play, we can look at all nodes that only have edges to leaves or other nodes for which we already determined the equilibrium strategy, and solve for the optimal strategy of the players corresponding to those nodes. We iteratively work our way up to the root, where we know the strategies for all nodes in the subgame-perfect equilibrium. We illustrate this approach with an example.

Example 3 (Backward induction)**.** Consider the extensive-form game from Example 1. In order to solve for the subgame-perfect equilibrium, we start with the subgame represented by the subtree rooted at the node highlighted in Figure 1a. Player 1 needs to choose between LEFT and RIGHT, receiving a payoff of 1 or 0 respectively. If the game ever reaches this subgame, player 1's optimal strategy will be to choose LEFT. Thus, we can replace the subtree rooted at the highlighted node with the payoff that the game will result in if the game ever reaches that subgame. This is shown in Figure 1b.

Next, we consider the subgame represented by the subtree rooted at the node highlighted in Figure 1b. Player 2 must choose between LEFT and RIGHT, receiving a payoff of 4 or 3, respectively. Thus, if the game ever reaches this subgame, player 2's optimal strategy is to choose LEFT and receive a payoff of 4. Thus, we can replace the subtree rooted at this highlighted node with the payoff that the game will result in if the game ever reaches that subgame, shown in Figure 1c.

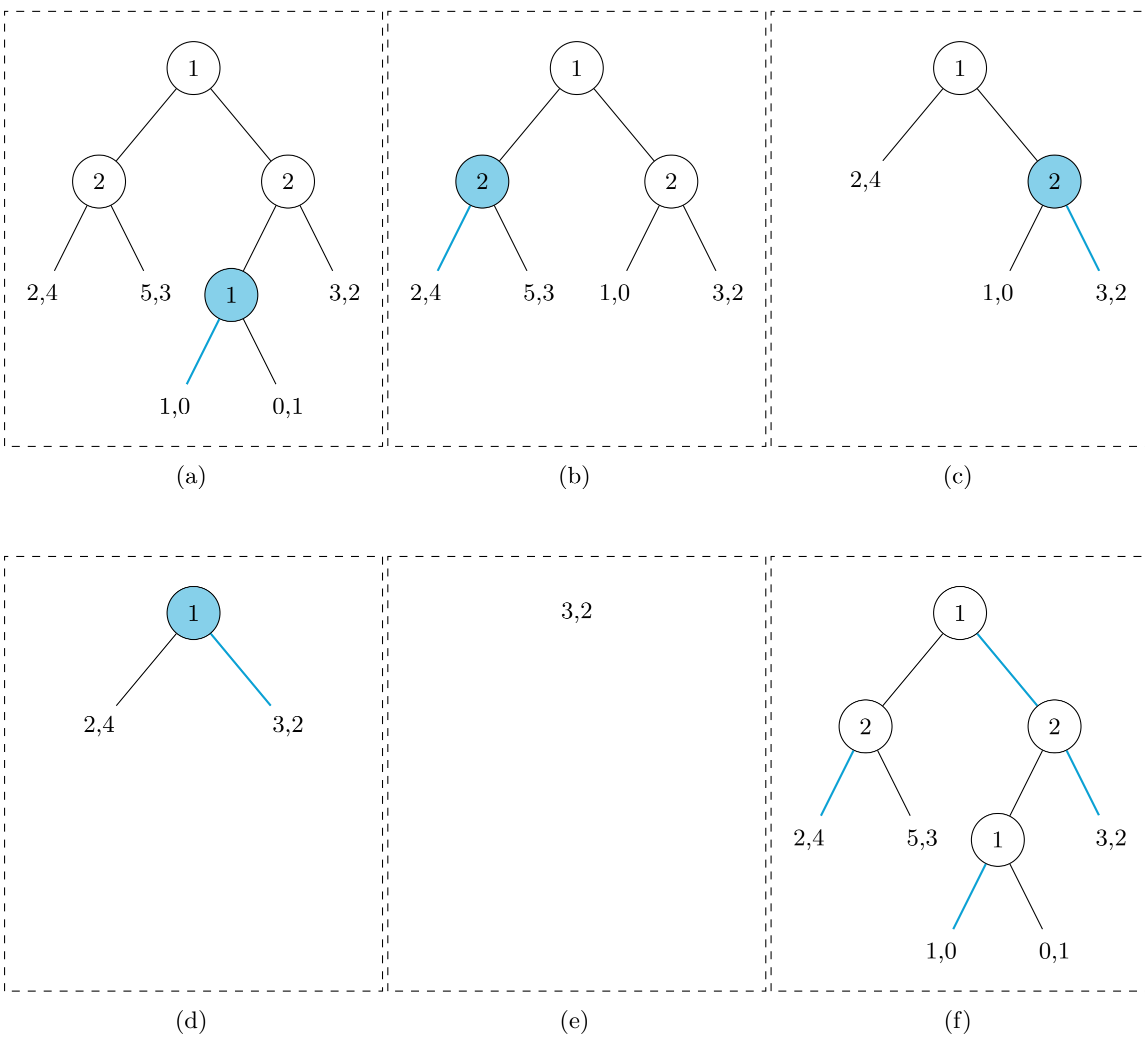

Figure 1: Backward induction for the extensive-form game from Example 1.

Continuing this process, we consider the subgame represented by the subtree rooted at the node highlighted in Figure 1c. Player 2 will always choose RIGHT in this subgame, so we replace this subtree with $3, 2$ in Figure 1d. Finally, we consider the subgame represented by the subtree rooted at the node highlighted in Figure 1d, the root of the original game. Player 1 will always choose RIGHT in this subgame, so we can replace this tree with a $3, 2$ as depicted in Figure 1e. This is the payoff of the players in the subgame-perfect equilibrium.

To get the strategy profile for this subgame-perfect equilibrium, we keep track of the strategies we found for the two players in each subgame, as shown in Figure 1f.

While we did not formally prove that backward induction successfully identifies a subgame-perfect equilibrium, the intuition for why this is the case is clear. In the strategy profile we construct, at every point in the game, a player is doing what is best for them, given that all following players will also do what is best for themselves.

When a player has multiple optimal options given them equal payoff at a step of backwards induction, we can break these ties arbitrarily: Any choice of tie-breaking will lead to a subgame-perfect equilibrium. In fact, also every subgame-perfect equilibrium in an extensive-form game can be obtained this way, by doing backward induction with the 'right' tie-breaking choices. This process for finding *one* subgame-perfect equilibrium is efficient (given the entire tree representation of the game) since we only need to eliminate nodes one by one. However, finding *all* subgame-perfect equilibria via backward induction is not always efficient, since the number of subgame-perfect equilibria can grow exponentially with the number of tie-breaking choices.

While subgame-perfect equilibria capture more nuances of player strategies in extensive-form games than Nash equilibria, they still may not explain the empirically observed strategies of players in certain games:

Example 4. Consider the following extensive-form game played between two players, rooted at the leftmost node.

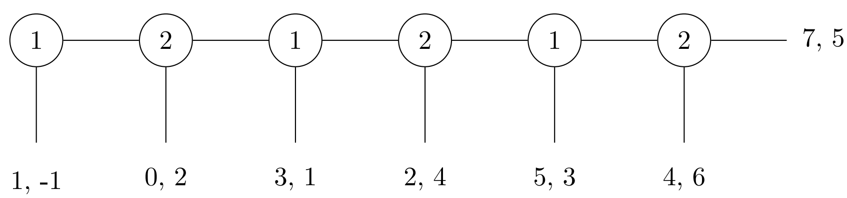

We can solve for the subgame-perfect equilibrium via backward induction. Starting at the smallest subgame, the rightmost node, we see that player 2 will always choose DOWN (with a payoff of 6) over RIGHT (with a payoff of 5), if the game ever reaches this subgame. Continuing up the tree to the subgame rooted at node just to the left, we see that player 1 will always choose DOWN (payoff of 5) over RIGHT (payoff of 4, because player 2 will choose DOWN). Continuing this process up the tree, we get that at any node the player whose turn it is will choose DOWN. In particular, when the game starts, player 1 will choose DOWN in the subgame-perfect equilibrium, leading to a payoff of $(1, -1)$.

However, because of the structure of the game, payoffs generally increase as the game continues down the tree. There are outcomes later in the tree in which both players are better off than with $(1, -1)$, that they get to by letting the game continue rather than "defecting" by choosing DOWN at some early point in the game.

If you were playing the game as player 1, when, if at all, would you choose DOWN? Although the subgame-perfect equilibrium suggests the answer to this question is to choose DOWN on the leftmost node, most people would play differently.

2 Imperfect Information Games

So far, we have assumed that the players in extensive-form games know, once it is their turn, which strategies the players preceding them played. Relaxing this assumption leads to a large class of games, that we will only cover at a high level.

Definition 3 (Imperfect Information Games)**.** A *chance node* does not belong to a player, but instead chooses one of its outgoing edges according to a known probability distribution. An *information set* is a set of nodes that a player may be in, so that the player only knows that they are in one node of the set, but do not know in which node of the set. Their strategy thus must be identical for all nodes in the information set. An imperfect information game is a extensive-form game with chance nodes and information sets.

Example 5 (Basic Imperfect Information Game)**.** Consider the following imperfect information game:

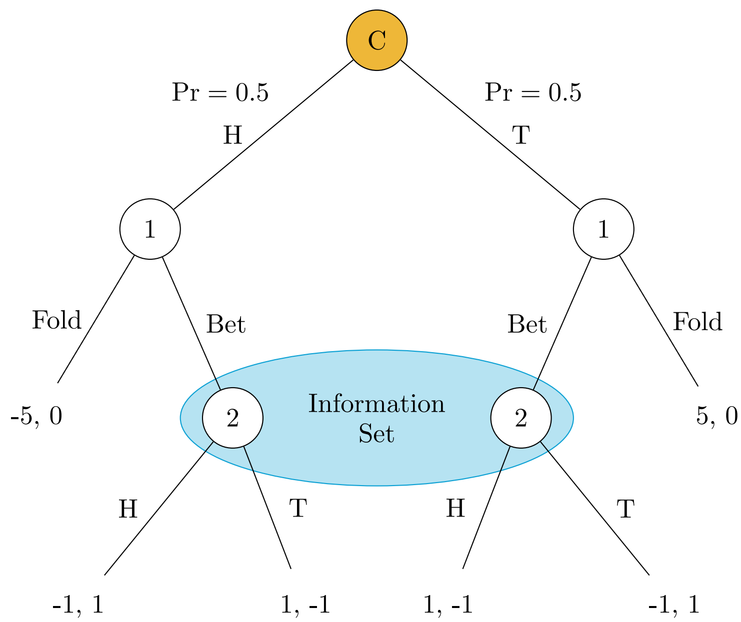

The yellow root node represents a chance node where a coin is flipped, with a 0.5 probability of H and a 0.5 probability of T. Player 1 observes the outcome of the flip before choosing their action. Player 1 can then either "Bet" or "Fold." If Player 1 folds, the game ends with payoffs of -5 for Player 1 and 0 for Player 2 if the coin flip was H, and 5 for Player 1 and 0 for Player 2 if the coin flip was T. If Player 1 bets, the next moves is up to Player 2, who does not know the outcome of the coin flip, represented in the game tree through the blue indifference set. Player 2 has to choose between H and T, with the payoff being -1,1 if player 2 guesses the coin flip correct, and 1,-1 if not.

Observe that if the coin flip is H, then player 1 will bet, and if the coin flip is T, then player 1 will fold. Therefore, player 2 will always play H: They will only ever be in the position to make this decision if the coin flip was H and player 1 thus chose to bet. If the -5 and 5 payoffs were flipped, however, then player 2 would always play T, by the same reasoning. We see that it's impossible to compute the optimal strategy of a subgame in isolation in imperfect information games, making finding optimal strategies a significantly harder task.

3 Stackelberg Games

We now consider a special case of imperfect information games with only 2-steps that we'll find to be tractable:

Definition 4 (Stackelberg Games)**.** In a Stackelberg game between 2 players is represented as a payoff matrix. The row player is the *leader* and commits to a certain (possibly mixed) strategy. The column player is the *follower*, wo observes the leader's commitment and then chooses a strategy.

Example 6 (Basic Stackelberg Game)**.** Consider the following Stackelberg game:

	LEFT	RIGHT
UP	$1, 1$	$3, 0$
DOWN	$0, 0$	$2, 1$

Playing UP is a dominant strategy for the row player. In that case, the column player is better of playing LEFT. Thus, (UP,LEFT) is the only Nash equilibrium in this game, with payoffs $(1, 1)$.

However, since this is a Stackelberg game, the row player as the leader can commit to playing DOWN. In response, it is optimal for the column player as the follower to choose RIGHT. The payoffs now are $(2, 1)$, so that the leader is strictly better off! Given that the leader announces their commitment to a pure strategy, the Stackelberg game can be rewritten as an extensive-form game of perfect information:

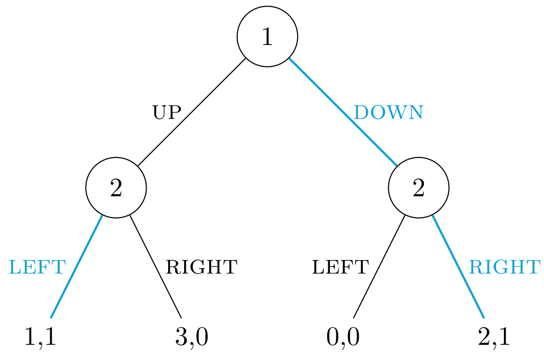

The subgame-perfect equilibrium is highlighted in blue. It is achievable when the game is played as a Stackelberg game.

We just saw that in a Stackelberg game, it may be optimal for the leader to play a pure strategy in spite of the existence of another, dominating strategy. Can the leader increase their payoff even more by committing to a mixed strategy? In particular, what is the maximum payoff the leader can get with a mixed strategy? This turns the game into the following imperfect information game:

assuming that the follower breaks their ties in favor of the leader? This turns

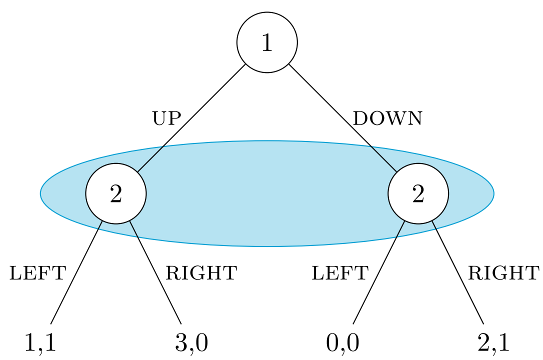

Let's assume the leader's mixed strategy is to play UP with probability p and DOWN with probability $1 - p$. Thus, when it is the follower's turn, they do not know which of the two nodes they are actually in, they just know they are in the left node (corresponding to player 1 choosing UP) with probability p and in the right node (corresponding to player 1 choosing DOWN) with probability $1 - p$. Thus, the follower knows that their expected payoff for choosing LEFT is

$$p \cdot 1 + (1 - p) \cdot 0 = p$$

and that their expected payoff for choosing RIGHT is

$$p \cdot 0 + (1 - p) \cdot 1 = 1 - p.$$

If $p > 1/2$, the follower will choose LEFT and if $p \leq 1/2$, the follower will choose RIGHT. Note that there is a tie at $p = 1/2$ and we are assuming that the follower chooses the strategy RIGHT which makes the leader better off — we'll explain why this is reasonable in a moment.

Consequently, if $p > 1/2$, the leader's expected payoff is

$$p \cdot 1 + (1-p) \cdot 0 = p$$

and if $p \leq 1/2$ the leader's expected payoff is

$$p \cdot 3 + (1-p) \cdot 2 = 2 + p$$

Therefore, the leader's expected payoff is maximized when they set $p = 1/2$ for an expected payoff of 2.5. This is greater than 2, their best payoff from a pure strategy! Thus, the leader can use randomness and the induced imperfect information to increase their payoff even further.

Let's now reconsider the assumption that the follower will break ties in favor of the leader, as we encountered it at $p = 1/2$. To avoid trusting the benevolence of the follower, the leader may choose $p = 1/2 - \varepsilon$ for an arbitrarily small $\varepsilon > 0$. In that case, it is strictly better for the follower to chose RIGHT, giving the leader a payoff of $2.5 - \varepsilon$. Since the leader can pick ε as small as they desire, their payoff can get as close to 2.5 as they desire — we say that it *is* 2.5 at the optimum. To avoid this awkward treatment of edge cases, it is customary to assume that the follower breaks ties in favor of the leader.

Definition 5 (Stackelberg Equilibrium)**.** For a mixed strategy $x_1 \in \Delta(S_1)$ of the leader, define the best response set of the follower as

$$B_2(x_1) = \arg\max_{s_2 \in S_2} u_2(x_1, s_2)$$

In a *strong Stackelberg equilibrium (SSE)* (x_1^*, s_2^*) the leader plays a mixed strategy x_1^* in

$$\arg \max_{x_1 \in \Delta(S_1)} \max_{s_2 \in B_2(x_1)} u_1(x_1, s_2)$$

and the follower plays a best-response strategy $s_2^* \in B_2(x_1^*)$.

This definition is assuming that in response to the leader committing to strategy $x_1 \in \Delta(S_1)$, player 2 responds optimally by picking a strategy in $B_2(x_1)$, but breaking ties in favor of player 1 (which is a reasonable simplifying assumption, by the same ε-argument as above). Knowing this, player 1 picks their strategy x_1 to maximize their utility $\max_{s_2 \in B_2(x_1)} u_1(x_1, s_2)$.

We will next see how an SSE can be computed via linear programming.

Example 7. Consider the following Stackelberg game where the row player is the leader and the column player is the follower:

	LEFT	RIGHT
UP	$1, 0$	$0, 2$
MIDDLE	$0, 1$	$1, 0$
DOWN	$0, 0$	$0, 0$

We write player 1's strategy as $x_1 = (p_1, p_2, 1 - p_1 - p_2)$, where p_1 is the probability of playing UP, p_2 is the probability of playing MIDDLE, and $1 - p_1 - p_2$ is the probability of playing DOWN. The follower's utility given x_1 for playing LEFT is p_2 and for playing RIGHT is $2p_1$, so their best response depends on p_1 and p_2.

We set up two linear programs: one maximizing player 1's utility conditioned LEFT being a best response of player 2 and another maximizing player 1's utility conditioned on RIGHT being a best response of player 2. Then, for whichever linear program achieves a higher objective function (player 1's utility), player 1's corresponding mixed strategy and player 2's best response is a SSE.

Let's first consider the case where LEFT is a best response of the follower, i.e., when $p_2 \geq 2p_1$. The leader's utility when the follower plays LEFT is p_1, so we have the following linear program:

$$\begin{aligned} &\max p_1 \\ &\text{s.t. } p_2 \geq 2p_1, \\ &\quad\; p_1 + p_2 \leq 1, \\ &\quad\; p_1, p_2 \geq 0 \end{aligned}$$

The first constraint ensures that LEFT is a best response of the follower, the other two constraints ensure that the leader's mixed strategy is a probability distribution over the leader's strategies. The solution to this LP is $p_1 = 1/3$ and $p_2 = 2/3$, with the optimal objective function value being $1/3$.

We now consider the case where RIGHT is a best response of the follower, i.e., $p_2 \leq 2p_1$. The leader's utility when the follower plays RIGHT is p_2, so we have the following linear program:

$$\begin{aligned} &\max p_2 \\ &\text{s.t. } p_2 \leq 2p_1, \\ &\quad\; p_1 + p_2 \leq 1, \\ &\quad\; p_1, p_2 \geq 0 \end{aligned}$$

The solution to this LP is also $p_1 = 1/3$ and $p_2 = 2/3$, with the optimal objective function value being $2/3$.

In this case, both linear programs gave us the same optimal mixed strategy, but the leader's utility is higher when the follower chooses RIGHT — which they will, since we are assuming they are breaking ties in favor of the leader. Thus, the SSE is $x_1 = (1/3, 2/3)$ and $s_2 =$ RIGHT.

Figure 2 illustrates the solution space of the linear programs and the position of the optimal solutions. The x-axis corresponds to p_1 and the y-axis to p_2. The solution space is limited by the constraint that $p_1 + p_2 \leq 1$ so all solutions must inside the large triangle. The blue region represents $p_2 \geq 2p_1$, where the follower plays LEFT. In the orange region, $p_2 \leq 2p_1$, where the follower plays RIGHT. Finally, note that the optimal solution exists on the indifference line at $(1/3, 2/3)$.

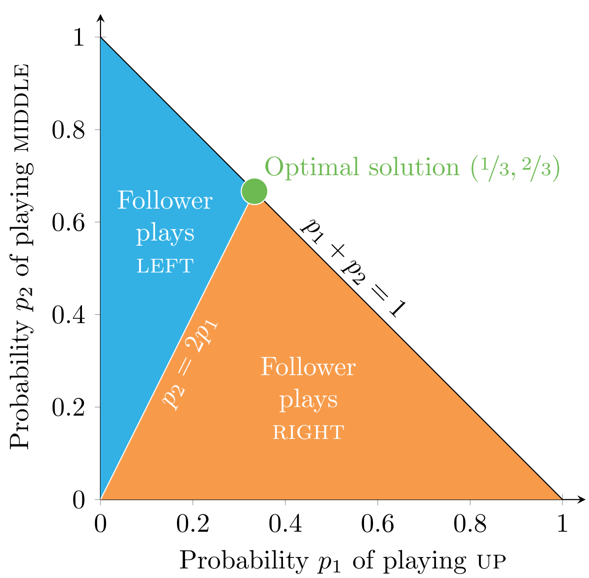

Figure 2: The solution space of the LPs for the original Stackelberg game

Now, consider a slightly modified Stackelberg game with the payoff matrix:

	LEFT	RIGHT
UP	$1, 0$	$2, 2$
MIDDLE	$0, 1$	$1, 0$
DOWN	$0, 0$	$0, 0$

This changes our second linear program (when RIGHT is a best response of the follower) to

$$\begin{aligned} \max\ & 2p_1 + p_2 \\ \text{s.t. } & p_2 \leq 2p_1, \\ & p_1 + p_2 \leq 1, \\ & p_1, p_2 \geq 0 \end{aligned}$$

because the leader now receives a payoff of 2 instead of 0 when they play UP and the follower plays RIGHT. This changes the optimal solution of the second LP to $x_1 = (1, 0)$, achieving an objective function value of 2. Since this is higher than the objective function (leader's payoff) in the first LP, corresponding to when LEFT is a best response of the follower, the SSE changes to $x_1 = (1, 0)$ and $s_2 =$ RIGHT. A depiction of the solution space of the linear programs for this modified game can be found in Figure 3. The optimal solution moved to $(1, 0)$.

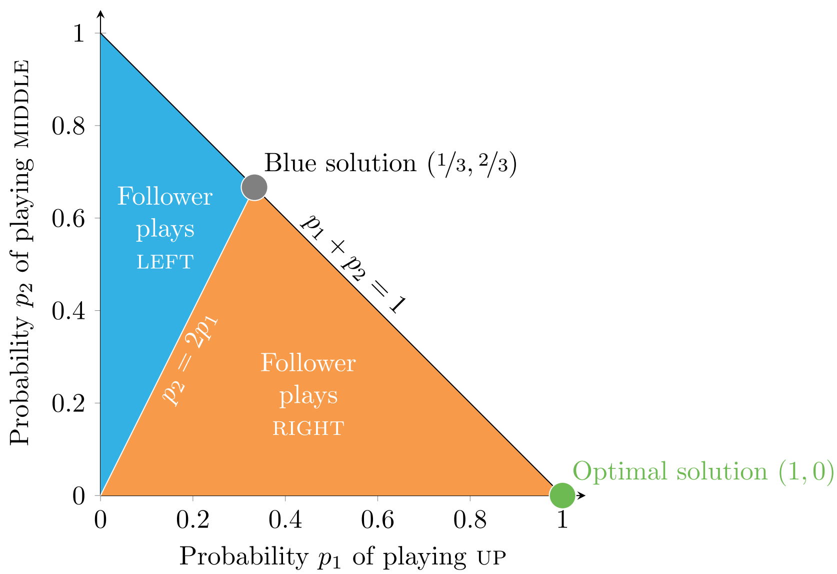

Figure 3: The solution space of the LPs for the modified Stackelberg game

We now give a general algorithm for computing SSEs using linear programming. The leader's mixed strategy is defined by variables $x(s_1)$ corresponding to the probability of playing strategy $s_1 \in S_1$. For each strategy of the follower s_2, we compute the strategy x for the leader maximizing the leaders payoff, conditioned on s_2^* being a best response for the follower. This is done via the following LP:

$$\begin{array}{ll} \max & \sum_{s_1 \in S_1} x(s_1) u_1(s_1, s_2^*) \\ \text{s.t.} & \forall s_2 \in S_2, \sum_{s_1 \in S} x(s_1) u_2(s_1, s_2^*) \geq \sum_{s_1 \in S_1} x(s_1) u_2(s_1, s_2) \\ & \sum_{s_1 \in S_1} x(s_1) = 1 \\ & \forall s_1 \in S_1, x(s_1) \geq 0 \end{array}$$

where the first set of constraints ensures that s_2^* is a best response for the follower and the other two sets of constraints ensure that the $x(s_1)$ indeed are a probability distribution. Finally, we know that the strategy x resulting from the LP corresponding to s_2^* with the largest optimal objective value is going to be an optimal strategy for player 1, and together with this s_2^* is going to be a SSE.

4 Game Theory in the Real World

We will end this lecture with a brief discussion about using these algorithms for equilibria in games in the real world. Over the last few years, we have seen many advances in computational game-playing, with a well-known example being poker. Mathematically, poker is modeled as an extensive form game with incomplete information (you do not know what cards your opponents are holding), and as such is inherently difficult to solve computationally. Nonetheless, we have reached a point where we can solve extensive-form games with incomplete information well enough to outperform humans in Poker.

However, these algorithms can also be used for more than just recreational games. We can use game-theoretic algorithms to solve high-stakes problems in the real world in fields like negotiations and (cyber)security, leading to much better outcomes in some situations. The biggest difficulty seems to be to determine which game we are playing in a real-world situation and delineating the rules, to decide which algorithms or notion of equilibrium to apply. Much of Havard Professor Milind Tambe's work has been on these topics, for example applying game theoretic algorithms in real-world situations like ensuring physical security and protecting wildlife. You can also read more about AI's need to outgrow classical games here.