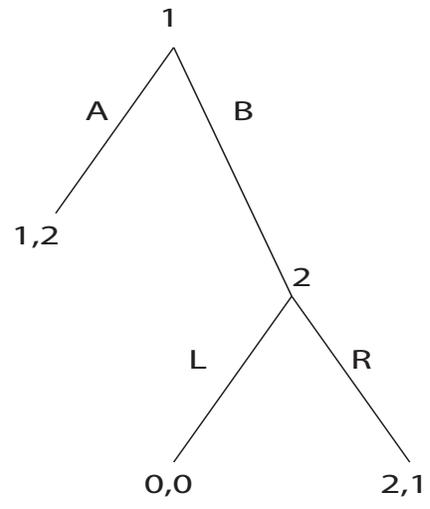


Backward Induction and Subgame Perfection

In extensive-form games, we can have a Nash equilibrium profile of strategies where player 2's strategy is a best response to player 1's strategy, but where she will not want to carry out her plan at some nodes of the game tree.

For example, consider the following game, given in both normal-form and extensive-form.

		player 2	
		L	R
player 1	A	1, 2	1, 2
	B	0, 0	2, 1



This game has two Nash equilibria: (A,L) and (B,R) . If we think of the players as selecting plans before the game starts, then the NE profile (A,L) makes some sense. If player 2 is absolutely convinced that player 1 will play A, then playing L is a best response.

While player 2 might plan to play L before the game starts, if the game proceeds and it is actually her turn to move, then L is not rational. Player 1 should realize this and play B.

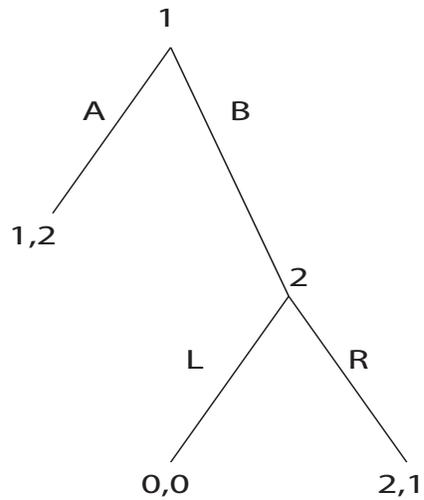
Thus, L might be a reasonable plan for player 2 if she thinks player 1 will play A, but it is not rational to carry out the plan after player 1 plays B. It is not **sequentially rational** for player 2 to play L.

Definition: *A player's strategy exhibits sequential rationality if it maximizes his or her expected payoff, conditional on every information set at which he or she has the move. That is, player i 's strategy should specify an optimal action at each of player i 's information sets, even those that player i does not believe will be reached in the game.*

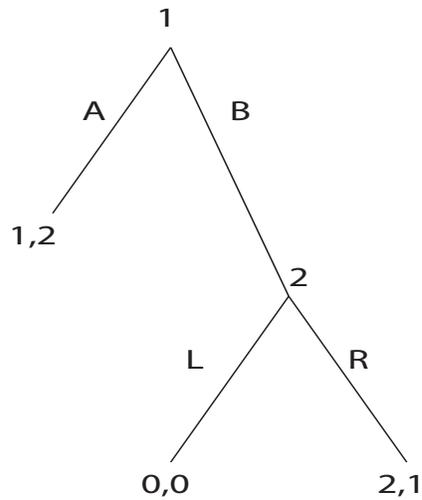
If sequential rationality is common knowledge, then players will look ahead and realize that players will not take actions down the tree that are not rational at that point.

For games of perfect information (all singleton information sets), **backward induction** is the process of "looking ahead and working backwards" to solve a game based on common knowledge of sequential rationality:

1. Start at each node that is an immediate predecessor of only terminal nodes, find the optimal action for the player who moves at that node, and change that node into a terminal node with the payoffs from the optimal action.
2. Apply step 1 to smaller and smaller games until we can assign payoffs to the initial node of the game.



Looking ahead to player 2's decision node, her optimal choice is R, so we can convert her decision node into a terminal node with payoffs (2,1).



In the smaller game, player 1 can either choose A and reach the terminal node with payoffs (1,2), or B and reach the terminal node with payoffs (2,1). Player 1's optimal choice is B. Backward induction leads to the strategy profile (B,R), with payoffs (2,1).

Here is a more complicated example, based on the Cuban Missile Crisis.

During 1962, the Soviet Union installed nuclear missiles in Cuba. When the US found out, President Kennedy discussed the options (i) do nothing, (ii) air strike on the missiles, (iii) a naval blockade of Cuba. JFK decided on the naval blockade. Negotiations ensued, and Khrushchev threatened to escalate the situation; both sides believed that nuclear war was a possibility. Finally, the Soviet Union agreed to remove the missiles if the United States agreed not to invade Cuba. Privately, Kennedy agreed to remove some missiles based in Turkey.

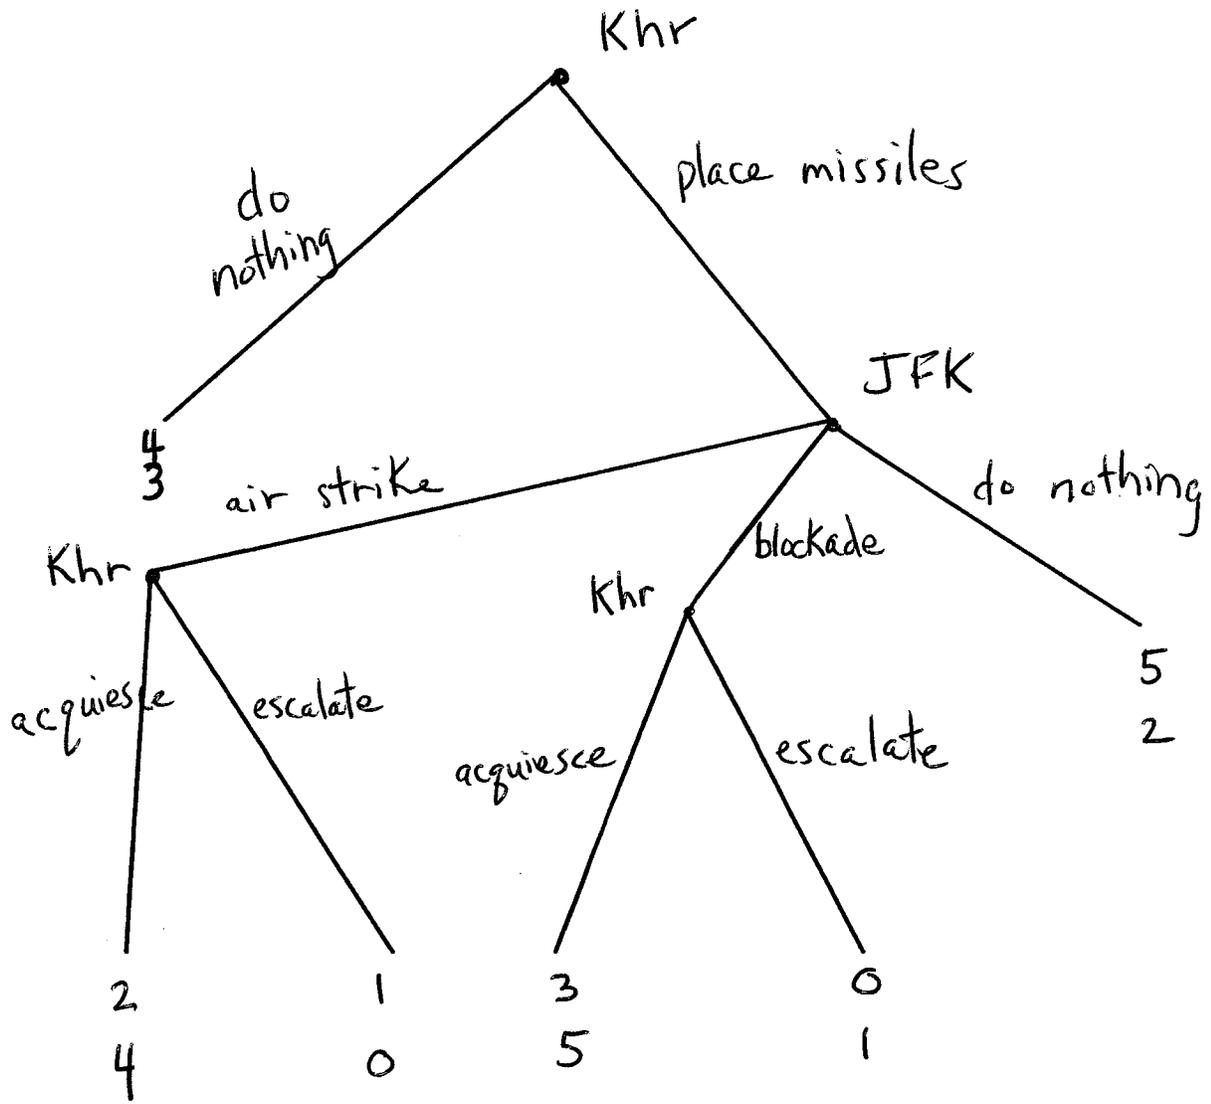
First, Khrushchev must decide whether to place the missiles in Cuba or not.

If the missiles are in place, JFK must decide on (i) nothing, (ii) air strike, or (iii) blockade.

If JFK decides on air strike or blockade, Khrushchev must decide whether to acquiesce or escalate.

Utility ranking of outcomes for Khrushchev: missiles allowed, status quo, acquiesce to blockade, acquiesce to air strike (lose resources), escalate after air strike (risk war but JFK is the initial aggressor), escalate after blockade (risk war and Khrushchev is the initial aggressor).

Utility ranking for JFK: blockade and have Khrushchev acquiesce, air strike and have Khrushchev acquiesce (be tough but use resources), status quo, allow missiles, Khrushchev escalates after blockade, Khrushchev escalates after air strike.



Backward Induction (first payoff is Khrushchev's, second is JFK's):

Khrushchev does not want to risk nuclear war and would optimally acquiesce to either a blockade or an air strike. Payoff to blockade is (3,5) and payoff to air strike is (2,4).

JFK would optimally choose blockade, leading to payoff (3,5).

Khrushchev's optimal choice is status quo (4,3), since he receives a higher payoff than placing the missiles (3,5).

Every finite game of perfect information can be solved using backward induction.

If each player has a strict preference over his possible terminal node payoffs (no ties), then backward induction gives a unique sequentially rational strategy profile.

Because no player has an incentive to deviate at any information set, the resulting strategy profile is a Nash equilibrium.

Related to backward induction is Zermelo's Theorem: In finite, two-player, "win-lose-draw" games of perfect information, then either one of the players has a strategy that guarantees a win, or both players have a strategy that guarantees a draw.

Since Chess is essentially a finite game of perfect information, then either player 1 (playing the white pieces) can guarantee a win, player 2 (playing the black pieces) can guarantee a win, or each player can guarantee a draw.

The Game of Northeast: Draw a rectangular grid of dimension $n \times k$.

Starting with player 1 on the full grid, the players take turns choosing a square remaining in the grid, removing that square and all the other squares to the northeast of that square. The player who is forced to choose the southwest square in the grid loses, and the other player wins.

For grids of equal dimensions $n \times n$, we know a winning strategy for player 1.

For grids of dimension $n \times k$, we know that player 1 has a winning strategy, even though we do not know what that winning strategy is.

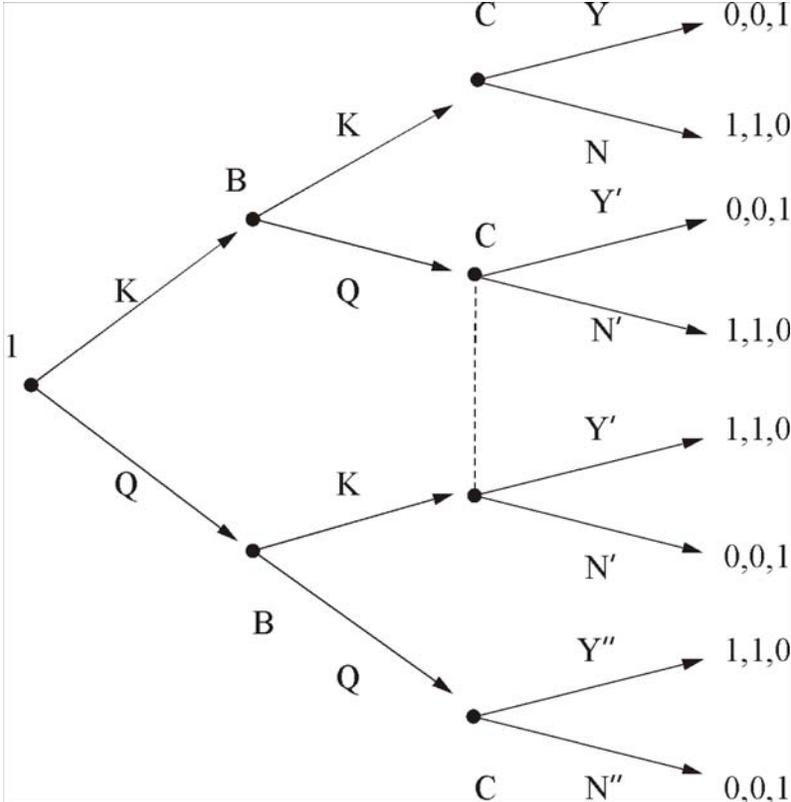
Subgame Perfection

In extensive-form games with imperfect information, backward induction can be problematic because a player's optimal action depends on which node she is at in her information set.

Still, sequential rationality can be captured by the concept of subgame perfection. First we need to define what we mean by a subgame.

Definition: Given an extensive-form game tree, a node x initiates a subgame if neither x nor any of its successors are in an information set containing nodes that are not successors of x . The tree defined by x and its successors is called a subgame.

Here is a game with 3 subgames.



Notice that:

1. Any game is a subgame of itself. Subgames other than the original game itself are called proper subgames.
2. For games of perfect information, every node other than a terminal node defines a subgame.
3. Any subgame is a game in its own right, satisfying all of our rules for game trees.
4. A strategy for the original game also defines a strategy for each of its subgames, sometimes called a continuation strategy.

Definition: *A strategy profile for an extensive-form game is a subgame perfect Nash equilibrium (SPNE) if it specifies a Nash equilibrium in each of its subgames.*

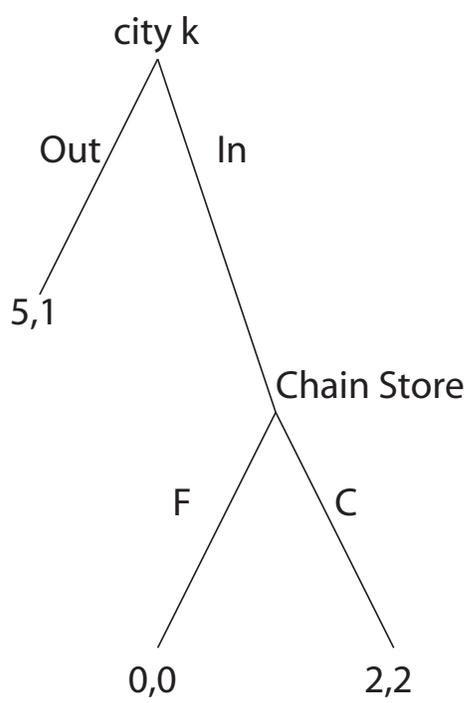
The idea behind SPNE is that even if a NE strategy profile dictates that certain subgames are not reached, we require that what the players *would do* conditional on reaching those subgames should constitute a NE. The "off-the-equilibrium-path" behavior can be important, because it affects the incentives of players to follow the equilibrium.

Notice that every SPNE must also be a NE, because the full game is also a subgame.

For finite games of perfect information, any backward induction solution is a SPNE and vice-versa. The advantage of SPNE is that it can be applied to games of imperfect information too.

The Chain Store Game

A chain store has branches in K cities, and in each city, $k = 1, \dots, K$, there is a competitor. In period k , the competitor in city k decides whether to enter the market or stay out, and if firm k enters, the chain store must decide whether to fight or cooperate. This is a game of perfect information, with the payoffs in each city given in the figure (next slide). Firm k cares only about the actions taken in its city, but the chain store's payoff is the sum of the payoffs it generates in each city.



Every path of the game in which the outcome in any period is either *out* or (in, C) is a **Nash equilibrium** outcome.

There is a unique **subgame perfect equilibrium**, where each competitor chooses *in* and the chain store always chooses C .

For $K = 1$, subgame perfection eliminates the bad NE.

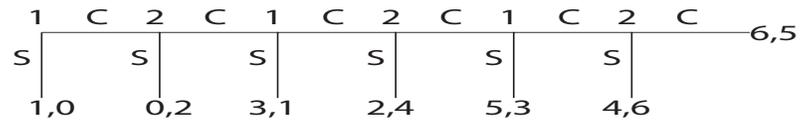
For large K , isn't it more reasonable to think that the chain store will establish a reputation for being tough? Moreover, if we see the chain store fighting the first 10 competitors, is it reasonable for the next competitor to enter? That is, if we see that the assumption of common knowledge of sequential rationality is violated, does it make sense to continue to assume it?

The Centipede Game

At each stage, a player can either stop the game, or continue the game, thereby sacrificing one dollar so that the other player can receive more than one dollar.

There is a unique subgame perfect equilibrium, where each player stops the game after every history.

There are several Nash equilibria, but all of them involve both players stopping the game at their first opportunity.



For a very long centipede, with payoffs in the hundreds, will player 1 stop immediately?

Since player 1 starting with C is not consistent with backward induction, is it reasonable for player 2 to believe that player 1 will use backward induction in the future?

Stackelberg Competition

With quantity competition, the timing of production matters. In the Stackelberg Game, firm 1 produces its output first. Then firm 2 observes q_1 before deciding its own output.

Let us adapt the Cournot model covered earlier. Marginal production cost is equal to 100, and market inverse demand is given by $p = 1000 - q_1 - q_2$. (If the expression is negative, we take the price to be zero.)

Firm 1's strategy set is

$$S_1 = \{q_1 : q_1 \geq 0\}.$$

Firm 2's strategy set is more complicated, reflecting the sequential nature of the game. Firm 2 can let q_2 depend on q_1 .

$$S_2 = \{\text{functions } q_2(q_1) : q_2(q_1) \geq 0 \text{ for each } q_1\}.$$

A Nash equilibrium is a profile of strategies, $(q_1^*, q_2^*(q_1))$.

In any SPNE, the subgame following any q_1 must be a NE. In other words, $q_2^*(q_1)$ must be a best-response to q_1 .

$$q_2^*(q_1) = BR_2(q_1) = 450 - \frac{q_1}{2}.$$

Notice the distinction: Under simultaneous-move Cournot competition, firm 2 chooses a quantity, which in the NE is a best response to firm 1's quantity q_1^* .

Under Stackelberg competition, firm 2 chooses a *function* describing firm 2's choice at each information set, and in the SPNE this function is $BR_2(q_1)$. Firm 2 must best respond to every possible output from firm 1, not just q_1^* .

Now that we have found the strategy of firm 2 satisfying sequential rationality, we can work backwards and find q_1^* . Firm 1's payoff function is

$$u_1(q_1, q_2) = [1000 - q_1 - q_2]q_1 - 100q_1.$$

In a SPNE, firm 1 chooses a best response to firm 2's strategy *function*, which we can substitute into the payoff function:

$$\begin{aligned} u_1(q_1, q_2^*(q_1)) &= [1000 - q_1 - (450 - \frac{q_1}{2})]q_1 - 100q_1 \\ &= [550 - \frac{q_1}{2}]q_1 - 100q_1 \\ &= [450 - \frac{q_1}{2}]q_1. \end{aligned}$$

Taking the derivative with respect to q_1 , setting it equal to zero, and solving for q_1 yields

$$\begin{aligned} 450 - q_1 &= 0 \\ q_1^* &= 450. \end{aligned}$$

The SPNE is $(450, 450 - \frac{q_1}{2})$.

Along the equilibrium path, we see that firm 1 chooses output of 450, and firm 2 chooses an output of 225. The corresponding price is 325. This is very different from the Cournot quantities, $(300, 300)$, and price, 400.

You can check that there is a first mover advantage. Firm 1 receives a higher payoff than firm 2, and a higher payoff than it would receive under simultaneous play.

Although the game has one SPNE, there are many NE that violate sequential rationality, for example:

$$\begin{aligned} q_1 &= 64 \\ q_2(q_1) &= 418 \quad \text{if } q_1 = 64 \\ &= 1000 \quad \text{otherwise.} \end{aligned}$$