Game Theory 101: The Complete Textbook
A User-Friendly Introduction to Game Theory
by William Spaniel
(excerpts/notes compiled by Mark Zabaro)

Contents

Lesson 1.1. The Prisoner's Dilemma and Strict Dominance **4**
1.1.1: Solving the Prisoner's Dilemma . . . 4
1.1.2: The Meaning of the Numbers and the Role of Game Theory . . . 4
1.1.3: Applications of the Prisoner's Dilemma . . . 4
1.1.4: Deadlock . . . 4
1.1.5: Strict Dominance in Asymmetric Games . . . 5
Conclusion . . . 5

Lesson 1.2: Iterated Elimination of Strictly Dominated Strategies **5**
1.2.1: Using Iterated Elimination of Strictly Dominated Strategies . . . 5
1.2.2: Duopolistic Competition . . . 5
1.2.3: Does Order Matter? . . . 5
1.2.4: Weak Dominance . . . 5

Lesson 1.3: The Stag Hunt, Pure Strategy Nash Equilibrium, and Best Responses **6**
1.3.2: Safety in Numbers and Best Responses . . . 6
1.3.3: The Stoplight Game . . . 6

Lesson 1.4: Dominance and Nash Equilibrium **7**
1.4.1: Nash Equilibrium and Iterated Elimination of Strictly Dominated Strategies . . . 7
1.4.2: When IESDS Leaves Multiple Strategies . . . 7
1.4.3: Nash Equilibrium and Iterated Elimination of Weakly Dominated Strategies . . . 7
1.4.4: Simultaneous Strict and Weak Dominance . . . 7

Lesson 1.5: Matching Pennies and Mixed Strategy Nash Equilibrium **7**
1.5.1: What Is a Mixed Strategy? . . . 7
1.5.2: The Mixed Strategy Algorithm . . . 8
1.5.3: How NOT to Write a Mixed Strategy Nash Equilibrium . . . 8
1.5.4: Mixed Strategies in the Stag Hunt . . . 8
1.5.5: How Changing Payoffs Affects Mixed Strategy Nash Equilibria . . . 8
1.5.6: Invalid Mixed Strategies . . . 8
1.5.7: Mixing and Dominance . . . 8

Lesson 1.6: Calculating Payoffds **8**
1.6.1: Chicken . . . 8
1.6.2: Battle of the Sexes . . . 9
1.6.3: Pure Coordination . . . 9
1.6.4: A Shortcut for Zero Sum Games . . . 9
1.6.5: Checking Your Answer . . . 9

Lesson 1.7: Strict Dominance in Mixed Strategies **10**
1.7.1: Mixed Dominance and IESDS . . . 10

Lesson 1.8: The Odd Rule and Infinitely Many Equilibria **10**
1.8.1: Infinitely Many Equilibria . . . 10
1.8.2: Take or Share? . . . 10

Lesson 2.1: Game Trees and Subgame Perfect Equilibrium **10**
2.1.1: The Meaning of the Numbers . . . 11
2.1.2: Games with Simultaneous Moves . . . 11

2.1.3: Constructing Games with Simultaneous Moves . . . 11
2.1.4: Why We Like Game Trees . . . 12

Lesson 2.2: Backward Induction **12**
2.2.1: How Not to Write a Subgame Perfect Equilibrium . . . 12
2.2.2: Practice with Backward Induction . . . 12

Lesson 2.3: Multiple Subgame Perfect Equilibria **13**
2.3.1: The Ultimatum Game . . . 13
2.3.2: Multiple Equilibria, Same Outcome . . . 13
2.3.3: When There Must Be a Unique SPE . . . 13
2.3.4: Multiple Equilibria with Simultaneous Moves . . . 13

Lesson 2.4: Making Threats Credible **13**
2.4.1: Burning Bridges . . . 13
2.4.2: Tying Hands . . . 13

Lesson 2.5: Commitment Problems **13**
2.5.1: Civil War . . . 13
2.5.2: Contracts . . . 14

Lesson 2.6: Backward Induction without a Game Tree **14**
2.6.1: Pirates! . . . 14
2.6.2: Nim . . . 14

Lesson 2.7: Problems with Backward Induction **14**
2.7.1: Mistake-Free Games . . . 14
2.7.2: Complete Information: The Chain Store Paradox . . . 14
2.7.3: Feigning Irrationality: The Centipede Game . . . 15

Lesson 2.8: Forward Induction **15**
2.8.1: Pub Hunt . . . 15
2.8.2: Defenestrated Chicken . . . 15
2.8.3: Costly Defenestration . . . 16
2.8.4: Burned Battle of the Sexes . . . 16

Lesson 3.1: Probability Distributions **16**
3.1.1: The Golden Rules of Probability Distributions . . . 16
3.1.2: Testing the Validity of a Probability . . . 16

Lesson 3.2: Mixed Strategy Nash Equilibria in Generalized Games **16**
3.2.1: Generalized Battle of the Sexes . . . 16
3.2.2: Generalized Prisoner's Dilemma . . . 17
3.2.3: Generalized Deadlock . . . 17

Lesson 3.3: Knife-Edge Equilibria **17**
3.3.1: The Hawk-Dove Game . . . 17
3.3.2: Why Are Knife-Edge Equilibria Unrealistic? . . . 18
3.3.3: When Knife-Edge Conditions Are Important . . . 18

Lesson 3.4: Comparative Statics **18**
3.4.1: Penalty Kicks . . . 19
3.4.2: The Volunteer's Dilemma . . . 19
3.4.3: Comparative Statics of the Hawk-Dove Game . . . 20
3.4.4: Curveballs with a Runner on Third Base . . . 20
3.4.5: Comparative Statics of Take or Share (or Lack Thereof) . . . 20

Lesson 3.5: Generalizing Mixed Strategy Nash Equilibrium **20**

3.5.1: The Support of a Mixed Strategy . . . 20
3.5.2: A Necessary but not Sufficient Condition . . . 20
3.5.3: A Trick with Weak Dominance . . . 20

Lesson 3.6: Rock-Paper-Scissors 20
3.6.1: A Trick with Symmetric, Zero Sum Games . . . 20
3.6.2: Generalized Rock-Paper-Scissors . . . 21
3.6.3: Mixed Strategies as Population Parameters . . . 21

Lesson 4.1: Infinite Strategy Spaces, Second Price Auctions, and the Median Voter Theorem 22
4.1.1: A Simple Game . . . 22
4.1.2: A Game with No Equilibria . . . 22
4.1.3: Hotelling's Game and the Median Voter Theorem . . . 22
4.1.4: A Duel . . . 22
4.1.5: Cournot Duopolistic Competition . . . 22
4.1.6: Second Price Auctions . . . 22

Lesson 1.1. The Prisoner's Dilemma and Strict Dominance

"game theory is the study of strategic interdependence–that is, situations where my actions affect both my welfare and your welfare and vice versa."

1.1.1: Solving the Prisoner's Dilemma

	Quiet	**Confess**
Quiet	-1, -1	-12, 0
Confess	0, -12	-8, -8

Rows denote strategies for Player 1 (he/him); columns denote strategies for Player 2 (she/her). Each cell represents an outcome and its respective payoffs, with Player 1's outcomes always first and Player 2's outcomes always second.

"Looking at the game matrix, people see that the <quiet, quiet> outcome leaves both players better off than the <confess, confess> outcome. They then wonder why the players cannot coordinate on keeping quiet. But as we saw, promises to remain silent are unsustainable. Player 1 *wants* player 2 to keep quiet so when he confesses he walks away free. The same goes for player 2. As a result, the <quiet, quiet> outcome is inherently unstable. Ultimately, the players finish in the inferior (but sustainable) <confess, confess> outcome."

1.1.2: The Meaning of the Numbers and the Role of Game Theory

"the cardinal values of the numbers are irrelevant to the outcome of the prisoner's dilemma"

strict dominance: "a strategy x strictly dominates strategy y for a player if strategy x provides a greater payoff for that player than strategy y regardless of what the other players do."

1.1.3: Applications of the Prisoner's Dilemma

pre-emptive attacks; arms races; levying tariffs

whether to advertise or not (if one competitor does, you must also)

"The Public Health Cigarette Smoking Act is a noteworthy application of the advertising game. In 1970, Richard Nixon signed the law, which removed cigarette ads from television. Tobacco companies actually *benefited* from this law in a perverse way–the law forced them to cooperating with each other. [...] The law simultaneously satisfied politicians, as it made targeting children more difficult for all tobacco companies."

1.1.4: Deadlock

"The 2012 Summer Olympics badminton tournament provides an interesting case study of strategic manipulation."

From the quarterfinals onward, it was single-elimination. But the second-best team in the world lost, guaranteeing they would end up in the lower half of the seeding.

"Teams who had already clinched a quarterfinal spot now had incentive to *lose* their remaining games. After all, a higher seeding meant a greater likelihood of facing the world's second-best team earlier in the elimination rounds. Matches turned into contests to see who could lose most efficiently!"

"This badminton example is a slight modification of a generic game called *deadlock*. It gets its name because the players cannot improve the quality of their outcomes unless the opponent chooses his or strategy incorrectly. Here are the generic game's payoffs:"

	Left	Right
Up	3, 3	4, 1
Down	1, 4	2, 2

"unlike in the prisoner's dilemma, no alternative outcome exists that is simultaneously better for both players. [...] As such, deadlock may be more intuitive, ut also tends to be substantively less interesting."

1.1.5: Strict Dominance in Asymmetric Games

Conclusion

Lesson 1.2: Iterated Elimination of Strictly Dominated Strategies

1.2.1: Using Iterated Elimination of Strictly Dominated Strategies

"In the previous lesson, we discussed why players ought to never play strictly dominated strategies. If players are intelligent, they should infer how others will *not* act and tailor their strategies accordingly."

Iterated Elimination of Strictly Dominated Strategies (IESDS)

1.2.2: Duopolistic Competition

Suppose fixed demand, and prices based on how much supply there is (in relation to that fixed demand), with only two producers in existence.

"the quantity produced is a function of both firms' strategic decisions. A single firm cannot control the other firm's production quantity, which in turn means it cannot unilaterally determine ultimate market price. As such, we might wonder if the firms have an optimal production strategy."

"Although the firms would like to collude to reduce production quantities and in turn artificially inflate market prices, neither firm can credibly commit to that course of action. After all, if one firm reduces its quantity produced, market prices go up, and it becomes more tempting for the other firm to break the agreement."

1.2.3: Does Order Matter?

"Suppose we had a game that started with two strictly dominated strategies. A natural question is whether we will end up with a diferent answer depending on which one we eliminate first. In fact, our first choice is irrelevant."

"When you are solving complex games and you find a strictly dominated strategy, *eliminate it immediately*. Although there may be more strategies you could eliminate in the first step, these strategies will be strictly dominated in the next step. It will also be easier to find them, as there is less information to consider in the remaining game."

1.2.4: Weak Dominance

weak dominance: "a strategy x weakly dominates a strategy y for a player if x provides *at least* as great of a payoff for that player regardless of what the other players do and there is at least one set of opposing strategies for which x pays greater than y."

"Eliminating weakly dominated strategies and analyzing the remaining game is called iterated elimination of weakly dominated stragies (IEWDS). Depending on the game, IEWDS sometimes produces sensible answers and sometimes does not."

“Depending on the order of elimination, IEWDS produces two separate answers. The problem is that iterated elimination of weakly dominated strategies gives us no guidance about which is correct or if both are.”

Lesson 1.3: The Stag Hunt, Pure Strategy Nash Equilibrium, and Best Responses

Two hunters choose independently whether to hunt stag or hare for the day. If both hunt stag, they share it and get the best possible payoff. If one hunts stag and one hunts hare, the latter gets all the hares and the former does not catch any stag. If they both hunt hare, they do a lot of work to get not much meat.

	Stag	Hare
Stag	3, 3	0, 2
Hare	2, 0	1, 1

Nash equilibrium: “A Nash equilibrium is a set of strategies, one for each player, such that no player has incentive to change his or her strategy given what the other players are doing.”

Pure strategy Nash equilibrium (PSNE): “both players are playing deterministic strategies [...] (In contrast, Lesson 1.5 covers mixed strategy Nash equilibrium, or MSNE, in which players randomize between their strategies.)”

“The stag hunt has two pure strategy Nash equilibria: <stag, stag> and <hare, hare>”

“Unlike the prisoner’s dilemma, the stag hunt illustrates game theory’s power to analyze interdependent decision making. In the prisoner’s dilemma, each player could effectively ignore what the other one planned on doing since confess generated a strictly gerater payoff regardless of the other prisoner’s choice. That is not the case with the stag hunt. Here, each player wants to do what the other is doing. That is, each player’s individually optimal strategy is a function of the other player’s choice.”

“Nash equilibrium has a “no regrets” property. If players play according to a Nash equilibrium, then they do not regret their choices once they have realized their payoffs.”

1.3.2: Safety in Numbers and Best Responses

“A **best response** is simply the optimal strategy for a particular player given what everyone else is doing.”

“For bookkeeping purposes, we mark a player’s best responses with an asterisk over his or her payoffs.”

“To find all of the game’s pure strategy Nash equilibria, we only need to check which outcomes have asterisks next to both players’ payoffs.”

“an alternative definition of a Nash equilibrium is a mutual best response.”

1.3.3: The Stoplight Game

“One interpretation is that a Nash equilibrium is a law that everyone would want to follow even in the absence of an effective police force.”

“consider the role of stoplights in society. Imagine two cars are approaching an intersection at 40 miles per hour from perpendicular directions. If both continue full speed, they will crash spectacularly. But if both stop, they waste time deciding who should go through the intersection first. Both drivers benefit if one continues without stopping while the other momentarily brakes to allow the other to pass.”

	Go	Stop
Go	-5, -5	1*, 0*
Stop	0*, 1*	-1, -1

“How can the players resolve their dilemma [between the two available Nash equilibria]? Stoplights provide a solution. The stoplight tells one driver to go with a green light, while it orders the other to stop with a red light. The players have no incentive to deviate from the stoplight’s suggestion. If the driver at the red light goes, he causes an accident. If the driver at the green light stops, he unnecessarily wastes some time. Thus, the stoplight instructs the drivers to play a Nash equilibrium.”

Lesson 1.4: Dominance and Nash Equilibrium

“if iterated elimination of strictly dominated strategies reduces the game to a single outcome, that outcome is a Nash equilibrium and it is the *only* Nash equilibrium of that game. Meanwhile, iterated elimination of weakly dominated strategies is not as kind: although any solution found through IEWDS is a Nash equilibrium, the IEWDS process sometimes eliminates other Nash equilibria.”

1.4.1: Nash Equilibrium and Iterated Elimination of Strictly Dominated Strategies

1.4.2: When IESDS Leaves Multiple Strategies

1.4.3: Nash Equilibrium and Iterated Elimination of Weakly Dominated Strategies

1.4.4: Simultaneous Strict and Weak Dominance

Lesson 1.5: Matching Pennies and Mixed Strategy Nash Equilibrium

“You and I each have a penny. Simultaneously, we choose whether to put our penny on the table with heads or tails facing up. If both of the pennies show heads or both of the pennies show tails (that is, if they match), then you have to pay me a dollar. But if one shows heads and the other shows tails (that is, they do not match), then I have to pay you a dollar.”

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

“Matching pennies is an example of a strictly competitive (or **zero sum**) game. In the prisoner’s dilemma and stag hunt, the players had incentive to cooperate with each other to acheive mutually beneficial outcomes. Here, however, the players actively want to see the other perform poorly; player 1 wins whatever player 2 loses, and vice versa.”

“As it turns out *every* finite game has at least one Nash equilibrium. [...] A **game is finite** if the number of players is finite and the number of pure strategies each player has is finite.”

1.5.1: What Is a Mixed Strategy?

mixed strategy: “randomizing over multiple strategies rather than playing a single “pure” strategy.”

1.5.2: The Mixed Strategy Algorithm

1.5.3: How NOT to Write a Mixed Strategy Nash Equilibrium

1.5.4: Mixed Strategies in the Stag Hunt

1.5.5: How Changing Payoffs Affects Mixed Strategy Nash Equilibria

"we must recalculate mixed strategy Nash equilibria every time a payoff changes."

1.5.6: Invalid Mixed Strategies

"Not all games have a mixed strategy Nash equilibrium. Deadlock, for example, does not."

When the Mixed Strategy for a player gives a contradiction (two constants are equal; a probability <0; a probability >1) then that player has no MSNE. Other players still may.

"if we show that one player cannot mix in such a manner, we still cannot eliminate the possibility that no MSNE exists. In particular, there are games where one player plays a pure strategy while the other mixes. Some differentiate these as "partially mixed strategy Nash equilibria" because one player mixes and the other does not."

1.5.7: Mixing and Dominance

"a strictly dominated strategy cannot be played with positive probability in a MSNE."

Lesson 1.6: Calculating Payoffds

1.6.1: Chicken

"Two [drivers] are on opposite ends of a one lane street and begin driving full speed toward one another. AT the last possible moment, they must decide whether to swerve or continue going straight. If one continues while the other swerves, the one who swerves is a "chicken" while the other has proven his or her bravery. If both swerve, then neither can claim superiority. But if they both continue, they crash straight into each other in an epich conflagration, which is the worst possible outcome for both players."

	Continue	Swerve
Continue	-10, -10	2*, -2*
Swerve	-2*, 2*	0, 0

""Snowdrift" is an alternative title for this game. Under that framework, two drivers are stuck on the opposite ends of a snowy road, and they simultaneously decide whether to stay in their cars or shovel a passageway. [...] We still stick to chicken because it allows for the possibility of a fiery explosion and does not involve any depressing winter weather."

in the MSNE, player 1 plays Continue with probability 1/5 and player 2 plays Continue with probability 1/5.

"As claimed, player 2's payoff is $\frac{-2}{5}$, just like player 1's. Now that we know each player's expected utility in the MSNE, we can see the <swerve, swerve> outcome leaves both players better off, as 0 is greater than $\frac{-2}{5}$. However, as we saw with the prisoner's dilemma, such an outcome is inherently unstable, as one of the players could profitably deviate to continuing."

1.6.2: Battle of the Sexes

A man and a woman want to go on a date, each has a preference between going to the balley and going to a fight. They don't have a way to coordinate on which to go to, and they'd rather do their less preferred option than spend the night apart.

	Ballet	Fight
Ballet	1, 2	0, 0
Fight	0, 0	2, 1

In the MSNE, each player goes their favorite entertainment option with probability $\frac{2}{3}$, giving expected utility $\frac{2}{3}$.

"Why is the mixed strategy Nash equilibrium so bizarre? Both the <ballet, fight> and <fight, ballet> outcomes represent coordination failure. They both occur with positive probability in the MSNE, accounting for $\frac{5}{9}$ of the outcomes. That means the couple go on their date less than half of the time if they mix, which drags down their payoffds. Indeed, each would be better off agreeing to meet their *lesser* preferred form of entertainment; the 1 they earn from that outcome beats the $\frac{2}{3}$ they earn in the MSNE."

1.6.3: Pure Coordination

	Up	Down
Up	1, 1	0, 0
Down	0, 0	1, 1

Here an MSNE exists where each player plays each strategy half the time, giving expected utility fo reach of $\frac{1}{2}$.

"In pure coordination, they only care about being together. One simple interpretation of this is choosing which side of the street to drive on. It does not really matter whether we all drive on the left side or all drive on the right side, as long as ome of us do not drive on the left while others drive on the right."

"One way to escape the inefficient mixed strategy Nash equilibrium in pure coordination and battle of the sexes is to follow social norms and laws. Driving on the road in the United States is very easy because a law tells us to drive on the right side, and that is an efficient Nash equilibrium. In battle of the sexes, perhaps the couple had a rule of thumb that the man chooses where to go on Fridays and the woman chooses where to go on Saturdays. If that were the case, they would only need to look at a calendar to coordinate even if they could not directly communicate. Thus, these strange MSNE help us interpret the usefulness of these types of coordination rules."

1.6.4: A Shortcut for Zero Sum Games

"calculating payoffs in mixed strategy Nash equilibria of zero sum games is easy, since you functionally calculate both players' payoffs by finding onw player's. And second, the players' payoffs need not be equal in MSNE: "equilibrium" only refers to the stability of certain strategies, not any sort of balance in the players' payoffs."

1.6.5: Checking Your Answer

"Recall that the mixed strategy algorithm guarantees that a player earns the same payoff for selecting either of his or here pure strategies. Consequently, we can also calculate a player's payoff by calculating his or her payoff for selecting one of his or her strategies."

"although the first method is intuitively easier to grasp, the alternative method is computationally less intensive."

Lesson 1.7: Strict Dominance in Mixed Strategies

"If a mixture of two pure strategies strictly dominates a third strategy, that third strategy is strictly dominated."

1.7.1: Mixed Dominance and IESDS

"strict dominance in mixed strategies can be frustrating to work with–there are many combinations of pure strategies and an infinite range of mixtures between those strategies. Consequently, it takes effort to locate such strictly dominant mixed strategies. However, the payoff is ultimately worth it, as we can simplify games a great deal when we do find them."

Lesson 1.8: The Odd Rule and Infinitely Many Equilibria

"a 1971 paper by Robert Wilson showed almost no games have an even or infinite number of equilibria. However, some quirky games do not follow this odd rule of thumb, and our old friend weak dominance frequently claims responsibility."

1.8.1: Infinitely Many Equilibria

1.8.2: Take or Share?

"keep in mind that these equilibria are a function of the preferences of the players. In this case, we assumed that the players only wanted to maximize money. However, they may have other motivations. For example, if they are slightly vengeful [...]. Meanwhile, generous players may want to mimic the other player's strategy [...]. In any case, whenever you look at a model, you should always question the players' preferences. Are players really indifferent between sharing and taking when their opponents take in the take or share game? DO players really always want to maximize money? Does anyone have benevolent preferences?"

Lesson 2.1: Game Trees and Subgame Perfect Equilibrium

"some strategic interactions flow over time in specific steps. We call these types of games *sequential* games, since the order of play follows a sequence."

"We have seen Selten's game as a simultaneous move game. But what if the players moved sequentially? Consider the following scenario. Firm 2 currently holds a monopoly on a particular market. Firm 1 is considering whether to challenge Firm 2's monopoly status. If it enters, FIrm 2 must decide whether to accede to FIrm 1's entry or declare a price war. If Firm 2 declares a price war, all the profits go away, and both earn 0. If Firm 2 accedes, both firms can profit. Here, Firm 1 receives a payoff of 3 while Firm 2 receives a payoff of 1. If Firm 1 stays out, it saves its investment and receives a payoff of 2. Meanwhile, without the competition of Firm 1, Firm 2 can increase its payoff to 2."

"We normally express such interactions using game trees:"

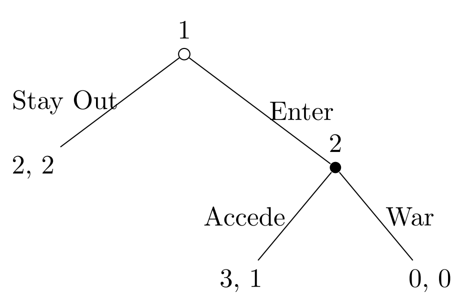

"We call this the *extensive form* of the firm entry game. The interaction begins at the open circle–called a decision node–where Firm 1 chooses whether to enter or stay out. Firm 2 selects accede or war at her decision node only if Firm 1 enters."

The 'accede' or 'war' decision of Firm 2 is called a *subgame*.

subgame perfect equilibrium (SPE) "Subgame perfection ensures that players only believe threats that others have incentive to carry out when it is time to execute those threats."

"Just as Nash equilibrium is the gold standard for simultaneous move games, subgame perfect equilibrium is the gold standard for extensive form games. As we saw in this example, all SPE are Nash equilibria, but not all Nash equilibria are SPE. As such, subgame perfection is a refinement of Nash equilibrium to ensure that players' threats are credible."

2.1.1: The Meaning of the Numbers

"In the first chapter, we discussed how the payoffs represented a player's subjective ranked ordering of possible outcomes, with the largest number representing the best outcome and the smallest outcome representing the worst. While all that remains true in extensive form games, we now assume that the payoffs represent a ranked ordering of outcomes given what has happened in the game. Consequently, the payoffs reflect a player's evaluation of fairness, distributive justice, and equality."

"game theory does not normatively tell players how to think or what their preferneces should be. Indeed, we can model scenarios where players value fairnessmore than their own financial well being. But, when we do, the payoffs *already incorporate these types of preferences*. Do not overthink the game; accept the numbers as they appear."

2.1.2: Games with Simultaneous Moves

"some extensive form games involve simultaneous moves. Here is a simple example:"

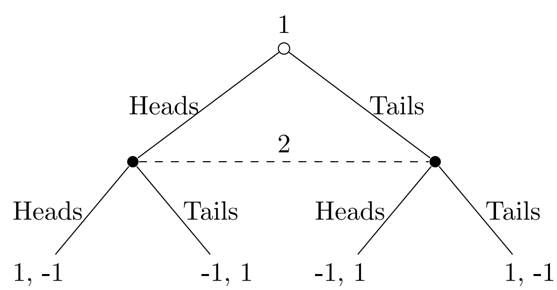

"This is matching pennies. If the coins match, player 1 earns 1 and player 2 earns -1. Otherwise, player 1 earns -1 and player 2 earns 1. Player 1 begins by choosing heads or tails. Player 2 then chooses heads or tails without seeing player 1's move. The dashed line indicates that player 2 is blind to player 1's strategy. We call this dashed line player 2's *information set*. The information it conveys is that player 1 played heads or tails, but she cannot see which."

"when we encounter simultaneous moves in extensive form games, the best thing to do is convert that game to a matrix and solve the game."

2.1.3: Constructing Games with Simultaneous Moves

"game trees must have identical strategies after simultaneous moves."

"at any information set, a player must have the same strategies available regardless of how the player arrived there."

2.1.4: Why We Like Game Trees

“When we want to analyze a strategic situation, knowing its extensive form is better than knowing its matrix. As it turns out, there is only one way to represent an extensive form game as a matrix. However, there can be multiple ways to represent a matrix in extensive form. Thus if we only have the matrix in front of us, we do not know which of its Nash equilibria will survive subgame perfection.”

Lesson 2.2: Backward Induction

The escalation game:

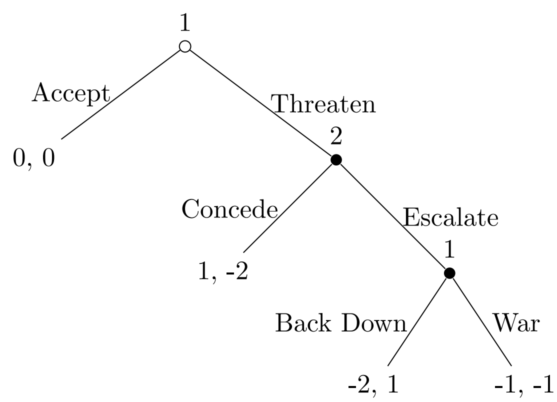

“when we use backward induction, we start at the end of the game and work our way to the beginning. Specifically, we see what players would want to do at the end of the game and take that information to the previous step to see how players should rationally respond to those future moves. After all, the smartest move today depends on what will happen tomorrow. We repeat this process until we arrive at the beginning of the game.”

2.2.1: How Not to Write a Subgame Perfect Equilibrium

“Subgame perfection, at its core, is the study of credible threats. Consequently, we want to know which threats in the escalation game are credible and which are not.”

“we say that the <(accept, war), escalate> is the SPE. This tells us that player 1 chooses accept and war at his two decision nodes, while player 2 selects escalate at hers. In contrast, merely saying that player 1 accepts the status quo does not tell us why this choice is rational for him. Since subgame perfection is the study of credible threats, we need to know that information. As such, the SPE must list the optional move at all decision noes regardless of whether the players actually reach those nodes when t hey play their equilibrium strategies.”

2.2.2: Practice with Backward Induction

Suppose the stag hunt as an extensive game, where player 1 first chooses what to hunt and player 2 chooses next (seeing player 1’s choice).

“the game’s SPE is <Stag, hare>; in the SPE, both players hunt a stag. Effectively, the sequential nature of the game solves the coordination problem. In the original stag hunt, <hare, hare> was a Nash equilibrium. But since player 1 can establish that he is hunting a stag, player 2 never has a reason to play hare in response.”

Lesson 2.3: Multiple Subgame Perfect Equilibria

2.3.1: The Ultimatum Game

"Let's start with a simple example. Player 1 has some good worth a value of 2 and has to bargain with player 2 over how to divide it. He can offer to split the good or he can attempt to take all of it. However, player 2 can reject either proposal. If she does, both receive nothing."

"This is a simple version of the ultimatum game. Player 1 begins by making an ultimatum–split or take. If he splits, player 2 accepts or rejects his division. If she accepts, they both earn 1; otherwise, they both earn 0. On the other side, if he takes, player 2 allows that to happen or spurns player 1's move. If she allows it, player 1 earns 2 while she earns 0. But if she spurns him, they both earn 0."

2.3.2: Multiple Equilibria, Same Outcome

2.3.3: When There Must Be a Unique SPE

"In a sequential game with no simultaneous moves, if an individual's payoffs are different for every outcome, and this is true for all individuals in the game, backward induction *must* yield a unique solution."

2.3.4: Multiple Equilibria with Simultaneous Moves

"If a single simultaneous decision exists in the game tree, multiple SPE may exist even if each payoff is unique."

"Backward induction requires every decision node to have a unique history."

"the solution is to utilize the *subgame* part of subgame perfect equilibrium. Rather than working from the very bottom, we work from the last decision in the game with a unique history. We call this a subgame."

Lesson 2.4: Making Threats Credible

"This lesson shows why players might want to intentionally constrain their future actions. [...] players can burn bridges–that is, make a certain future course of action impossible. [...] they can tie their hands–that is, leave a future option open but make it so extremely undesirable that they would never choose to pursue it."

2.4.1: Burning Bridges

2.4.2: Tying Hands

Lesson 2.5: Commitment Problems

"preferences matter more than words. What someone says they will do in the future may be inconsistent with what they would want to do once it is time to follow through. THat is not to say words are completely irrelevant."

2.5.1: Civil War

"Civil wars rarely end in negotiated settlements; normally, the sides fight each other until one completely militarily defeats the other. Commitment problems explain why."

2.5.2: Contracts

Lesson 2.6: Backward Induction without a Game Tree

2.6.1: Pirates!

A ship of 5 pirates (the first of which being the captain) have a scheme for the division of their spoils: the captain makes a proposal for how to split the bounty, and if at least half of the pirates agree, then it is done. If not, then the captain must walk the plank and the next pirate in succession can propose a split (which follows the same rules).

Backward induction shows that the captain need only offer a token share of the treasure to pirates 3 and 5 in order to have his split accepted.

2.6.2: Nim

Lesson 2.7: Problems with Backward Induction

"rational players may have incentive to deliberately act *irrational* so they can increase their payoffs."

2.7.1: Mistake-Free Games

"actors in our models *never* make mistakes. This is reasonable to assume when there are only a couple of players and two or three moves, but things quickly get out of hand as we increase the complexity of the game."

"We *assume* players are rational. We *assume* they know what is best for them. We *assume* they do not make mistakes. Although these assumptions are not heroic when the game is small and involves only a few players, it becomes further and further dicey with more and more additions to the game."

2.7.2: Complete Information: The Chain Store Paradox

Suppose a chain store with locations in 5 different towns. The game begins with a challenger in town 1 deciding whether to enter the market or not, and if it does, the chain store can decide whether to start a price war or not. And so on for the other four towns. While any one price war is probably not profitable, the fact that price wars are not credible threats means that a competitor would pop up in every town and the chain store would lose out bigtime. That being the case, the chain store has reason to deter would-be competitors to initiate a price war against the first competitor (even if it is irrational by our standards).

"the "chain store paradox" is that backward induction tells us that all the competitors will enter the market, yet we have a perfectly good reason to believe that no competitor will."

"in the model, we assume the chain store finds it unprofitable to start a price war against any of the challengers and that the possible challengers know this. Realistically, though, challengers might be unsure whether the chain store maximizes its profits in each individual town by starting a price war or by acceding to the challengers."

"a weak chain store has incentive to pretend that it is a strong chain store by engaging in a price war with any competitor that challenges it."

"something is still unsettling about the backward induction prediction even if we stick to the complete information story exclusively. Moreover, the chain store paradox is not the only game with this problem. Here, the chain store deliberately actived *viciously* irrational to improve its payoff."

2.7.3: Feigning Irrationality: The Centipede Game

"The centipede game is a frequent subject of laboratory experiments. Although the game endds immediately in the SPE, in practice players generally play for many rounds before someone finally takes the extra two dollars. As with the chain store paradox, game theorists have a variety of explanations for the discrepancy between subgame perfect play and play in practice."

"First, players may be irrational. They simply may be unable to work through the logic to understand that they ought to take immediately. In turn, they repeatedly add until they are close enough to understand the dilemma, which eventually causes the game to end."

"Second, note that just a single irrational player can throw a wrench into the system. Suppose you are a rational player 2, and player 1 begins the game by adding. You realize he is not rationaland wonder what would happen if you added as well. Given that he has already contributed to the pot, it stands to reason that he will do so again. [...] you may be inclined to add over your first dozen or so decision nodes. Thus, irrational play sparks further irrational play."

"Third, this logic in turn destroys the backward induction solution when there are two rational players. Suppose you are a rational player 1 and you ignore the SPE by beginning with add. Now the rational player 2 has no idea what is going on. She may figure you are irrational and so she should continue as in the previous case. Alternatively, she may think you are rational but deliberately acting irrational in hopes that she will play irrationally as well, thus improving both of your payoffs. This time, irrational play sparked further irrational play, yet no one was irrational."

Lesson 2.8: Forward Induction

"When we solve games with backward induction, players believe all future play will be rational, and they condition their present behavior on what will occur in the future. Forward induction adds an extra layer of complexity. Here, the players believe that all *prior* play was rational as well, and they condition their present play based off what they can infer about past play."

"Although forward induction may seem like a straightforward assumption, it quickly leads to some involved inferences."

2.8.1: Pub Hunt

An extensive form of the stag hunt. In this, player 1 starts. If player 1 chooses to go to the pub instead of go hunting, player 2 sees him there and has no choice but to join him there, and each's outcome is better than it would be with any hare-hunting strategy (but is slightly worse for each player than the payoff for <stag, stag>).

In this game, if player 1 does not go to the pub, player 2 assumes player 1 is angling for a better payoff, which is only possible with <stag, stag>–i.e., if player 1's past play is rational and believes player 2 to be rational, then player 1 is hunting stag. If so, player 2's best payoff is also <stag, stag>, so they too will hunt stag.

"Although logically demanding, forward induction leads to a plausible result here. After all, in the original stag hunt, the players merely wanted to coordinate on the stag. Introducing the pub allows them to do this, even though the players never meet there."

"Unfortunately, the pub hunt is also the simplest application of forward induction. The examples grow increasingly bizarre from there."

2.8.2: Defenestrated Chicken

Consider a game of chicken where player 1 has the option of detaching their steering wheel and throwing it out the window as the game begins, and player 2 sees if player 1 has done this. This essentially ties player 1's hands to play 'Continue', leaving player 2 to decide whether they should also continue to drive straight (definitely causing a collision) or whether to swerve.

This effectively ties player 1's hands to Continue, or in other words makes player 1 playing Continue a credible (in fact, inevitable) threat. This forces player 2 to swerve.

"If both players are rational, understand backward induction, and understand forward induction, then all of those equilibria are plausible. If one of these actors is not rational, does not understand backward induction, or does not understand forward induction, then we should accept the fact that inaccurate assumptions can lead to inaccurate predictions."

2.8.3: Costly Defenestration

Consider an alternative where defenestrating the steering wheel is slightly costly, such that player 1's payoffs are slightly worse for <Defenestrate, *> (for any *) than <Continue, *>.

If player 2 assumes player 1 to be rational and she sees that player 1 has not thrown the steering wheel out the window, she must assume he is going for the only payoff that is better than <Defenestrate, swerve>, which is <Continue, swerve>, thus that he has decided to continue. If that is a given, then she must swerve, because <Continue, continue> gives her a worse payoff.

2.8.4: Burned Battle of the Sexes

Lesson 3.1: Probability Distributions

3.1.1: The Golden Rules of Probability Distributions

"A probability distribution is a set of events and the probability each event in the set occurs."

"Two golden rules of probability distributions maintain their mathematical tractability: (1) all events occur with probability no less than 0; (2) the sum of all probabilities of all events equals 1."

Four implications follow from these: (a) no probability can be greater than 1; (b) probability distributions cannot leave us wondering what else might happen; (c) the fact that the sum of the probabilities of all events equals 1 gives us a convenient way to solve for a single unknown probability; (d) probabilities can be 0 or 1.

3.1.2: Testing the Validity of a Probability

Lesson 3.2: Mixed Strategy Nash Equilibria in Generalized Games

3.2.1: Generalized Battle of the Sexes

"If we are going to encounter many different versions of battle of the sexes, it would help if we could derive a simple formula for the mixed strategy Nash equilibrium. At present, we only have the mixed strategy algorithm. That algorithm eventually finds the MSNE, but it requires a burdensome number of calculations each time we run it. Ideally, we would like to make the algorithmic calculation once and be able to apply the results every time we encounter an altered version of the game."

"All we have to do is replace the distinct numbers with *exogenous variables*:"

	Left	Right
Up	B, a	C, c
Down	C, c	A, b

“These variables are exogenous because they come from outside the game. The players do not choose them. [...] Instead, the players have these payoffs and know them going into the game, as though they were innate preferences.”

“if we create a rule that $A > B > C$ and $a > b > c$, we have the correct preference ordering for battle of the sexes.”

Finally, this gives:

	probability	EU1	EU2
Player 1	$\sigma_{up} = \frac{b-c}{a+b-2c)}$	$EU_{up} = (B)(\sigma_{left}) + (C)(1 - \sigma_{left})$	$EU_{down} = (C)(\sigma_{left}) + (A)(1 - \sigma_{left})$
Player 2	$\sigma_{left} = \frac{A-C}{A+B-2C}$		

3.2.2: Generalized Prisoner’s Dilemma

	Left	Right
Up	R, r	S, t
Down	T, s	P, p

“This time around, $T > R > P > S$ and $t > r > p > s$. You can remember the ordering like this: T is for **T**emptation, which is the payoff a player receives when he rats out his opponent and the opponent remains silent; R is for **R**eward, which is the good payoff the players receive when they both remain silent; P is for **P**unishment, which is the bad payoff both players receive when they both rat each other out; and S is for **S**ucker, which is the payoff for a player when he remains silent and the opponent rats him out. (These naming conventions come from *The Evolution of Cooperation* by Robert Axelrod, which is the seminal book on the prisoner’s dilemma. It is extremely accessible to readers new to game theory. As such, I give it my highest recommendation.”

Trying to apply the mixed strategy algorithm for either player results in a contradiction of one of our base assumptions (about probability distributions), meaning that there are no MSNE. In turn, this means that “no player can mix in this game, leaving <down, left> as the unique Nash equilibrium.”

3.2.3: Generalized Deadlock

“Let’s reinforce these same principles with deadlock. Recycling the payoffs from the prisoner’s dilemma, let $T > R > P > S$ and $t > r > p > s$. If we switch each player’s temptation payoff for the sucker’s payoff, deadlock results:”

	Left	Right
Up	R, r	T, s
Down	S, t	P, p

Lesson 3.3: Knife-Edge Equilibria

“Using exogenous variables in payoffs sometimes leads to cases where an equilibrium exists for only a single configuration of the payoffs; increasing or decreasing a single payoff by a tiny amount makes that equilibrium completely disappear. We refer to such an equilibrium as a *knife-edge equilibrium*, as they precariously rest on the skinny edge of a single number.”

3.3.1: The Hawk-Dove Game

Two animals are in conflict over some good worth v >0. They can choose whether to act like a hawk (and fight over it) or a dove (and not put up a fight). If both play the same goal they split the good in half, but there is a cost to fighting (c) that they each pay if both play hawk.

	Hawk	Dove
Hawk	$\frac{v}{2}-c$, $\frac{v}{2}-c$	v, 0
Dove	0, v	$\frac{v}{2}$, $\frac{v}{2}$

"Hawk strictly dominates dove for $\frac{v}{2} > c$. We know both players must pick hawk in equilibrium. In essence, the game turns into a prisoner's dilemma, in which "hawk" is "confess" and "dove" is "keep quiet"."

"in the case where $v < 2c$, a MSNE exists in which both players select hawk with probability $\frac{v}{2c}$ and choose dove with complementary probability."

"the partially mixed strategy Nash equilibria rely on the knife-edge condition of $\frac{v}{2} = c$. If $\frac{v}{2}$ is even slightly greater than or slightly less than c, these equilibria completely disappear."

3.3.2: Why Are Knife-Edge Equilibria Unrealistic?

Continuing with the hawk-dove game....

"if we are explicit with our assumptions, the probability that such a case exists equals *exactly* zero."

"the probability of observving a cost parameter c between any two values a and b (where $a < b$) is the integral of the probability density function between a and b. Let $f(x)$ be that probability density function: in general, $Pr[a \leq c \leq b] = \int_a^b f(x)dx$; but for the hawk-dove game in specific, $Pr[\frac{v}{2} \leq c \leq \frac{v}{2}] = \int_{\frac{v}{2}}^{\frac{v}{2}} f(x)dx$"

"Essentially, we are looking or the area under the probability density function curve at a single point. But that area has no width and hence has no area. Thus, the probability of observing a c exactly equal to $\frac{v}{2}$ is zero."

"this norm against researching knife-edge equilibria is a good thing. Knife-edge conditions induce indifference, which often leads to instances of weak dominance. Observe that hawk weakly dominates dove if $\frac{v}{2} = c$."

3.3.3: When Knife-Edge Conditions Are Important

"We can safely skip knife-edge conditions of games that occur naturally. Howevver, game masters sometimes fabricate the rules."

e.g., the take-share game from before is an instance of the knife-edge condition of the hawk-dove generalized game.

"absent a scenario where someone is actually in control of the game in this manner, we can ignore knife-edge equilibria."

"One more caveat to stress: we can only make claims about knife-edge conditions when referring to *exogenous* variables. We have seen and will see many games where an equilibrium rests on a knife-edge strategy *endogenously* selected by the players."

Lesson 3.4: Comparative Statics

"At its core, game theory is the study of altering the strategic dimensions of an environment. We want to know how subtle changes in a game affect how players behave."

"We began introducing the possibility of a fluid environment in this chapeter by adding exogenous variables to payoff matrices. In this lesson, we begin analyzing the actual change."

"The study of such changes is called **comparative statics**. In essence, we take one environment, make a slight tweak to it, and compare the outcomes of those two games. Using this method, we can discover how manipulating games affects a player's outcome or the welfare of a society."

Our method of calculating comparative statics is as follows:

1. Solve for the game's equilibria.
2. Calculate the element of interest. (This could be the probability a player selects a particular strategy, the probability the players reach a certain outcome, or a player's expected utility.)
3. Take the derivative of that element of interest with respect to the exogenous variable we want to manipulate.
4. Use that derivative to see how changing the exogenous variable affects the element of interest.

3.4.1: Penalty Kicks

Soccer penalty kicks: the striker can kick left or right, and the goalie can dive left or right, and they choose simultaneously. Suppose the goalie is a perfect blocker: if they choose correctly they always prevent a goal. Suppose also that the kicker is weaker on his left, where he has accuracy x such that $0 < x < 1$.

	D Left	D Right
K Left	0, 0	x, -x
K Right	1, -1	0, 0

So the question here is: as the kicker's accuracy on his left increases, will he kick to that side more or less frequently?

The MSNE is that the goalie dives left with probability $\frac{x}{1+x}$ while the kicker aims left with probability $\frac{1}{1+x}$. Since x represents the kicker's accuracy, we take the derivative of $\frac{x}{1+x}$ with respect to x: $f'(x) = \frac{-1}{(1+x)^2}$.

"Recall that x is bounded between 0 and 1. Thus, $\frac{-1}{(1+x)^2}$ is always negative on that interval. Therefore, the probability the kicker aims left *decreases* as his accuracy to that side improves!"

"Most people guess the opposite. After all, why would improving your abilities on one side make you want to utilize that side less frequently?"

"The critical insight is that the kicker must factor in the goalie's strategic interaction. In a world where the goalie knows nothing about the kicker's weakness, the kicker should aim toward the stronger side. However, in this version of the game, the goalie is fully aware of the kicker's weakness, and she exploits that weakness by guarding the kicker's strong side more frequently. In turn, the kicker sees value in aiming toward his weak side: although he will often miss, the goalie will not be there to stop a well-placed shot as frequently. But when teh kicker's accuracy on his left improves, the goalie can no longer camp out on the right. That makes the kicker willing to aim to the right more often, which is what the comparative static tells us will happen."

3.4.2: The Volunteer's Dilemma

Take the case of Kitty Genovese, being stabbed outside of the apartment of two of her neighbors. Each neighbor values Kitty's life at 1 so that's an incentive to call the police. But there is a cost (c, such that $0 < c < 1$) to calling the police (incurred when they have to answer all the questions from the police). So they are each best off if one or both of them calls the police, but they prefer not to be the one to do so.

"This is a **volunteer's dilemma**: each neighbor only wants to volunteer to call if he or she knows the other one will not. As such, without communication, it is unclear who should pick up the phone."

	Ignore	Call
Ignore	0, 0	1*, 1-c*
Call	1-c*, 1*	1-c, 1-c

In the MSNE, both players ignore with probability c and call with probability $1 - c$.

3.4.3: Comparative Statics of the Hawk-Dove Game

"Suppose the hawk-dove game is a model of crisis bargaining between two states, where the <hawk, hawk> outcome represents war. What can we say about the probability of war as a function of the cost of conflict?"

"if we increase the cost of war, the expected probability of war will either remain the same or decrease. Two reasons prevent us from saying it is decreasing instead of weakly decreasing. First, if we increase c but maintain $c < \frac{v}{2}$, the players will continue playing <hawk, hawk> despite paying a larger cost. And second, if we increase c in the range $c > \frac{v}{2}$, the probability of war only decreases if the players are using the MSNE. If they are playing one of the PSNE, the probability of war remains 0."

"Regardless, political scientists interested in interstate wars have created far more elaborate models of crisis bargaining and seen the same catch-22 result. Promoters of peace may want to limit the harm war inflicts on those unlucky enough to be fighting. However, decreasing the costs associated with conflict actually incentivizes states to fight more frequently. As such, the destructive power of nuclear weapons may ironically be better promoters of peace than even the most seasoned of diplomats."

3.4.4: Curveballs with a Runner on Third Base

3.4.5: Comparative Statics of Take or Share (or Lack Thereof)

"Not all games have interesting comparative statics. In fact, the equilibria of some games will not change at all even if you alter some of its features."

Lesson 3.5: Generalizing Mixed Strategy Nash Equilibrium

3.5.1: The Support of a Mixed Strategy

"In games with a finite number of strategies, we say a pure strategy is in the *support of a mixed strategy* if and only if the probability of playing that pure strategy inn the mixed strategy is positive."

3.5.2: A Necessary but not Sufficient Condition

"It is necessary for all pure strategies in the support of a mixed to yield the same expected utility. However, it is *not* a sufficient condition."

3.5.3: A Trick with Weak Dominance

"knowing weakly dominated strategies makes finding mixed strategy equilibria substantially less time consuming. If a player mixes among *all* of his or her strategies, in a game with a finite number of strategies, the other player cannot play a weakly dominated strategy in equilibrium."

Lesson 3.6: Rock-Paper-Scissors

3.6.1: A Trick with Symmetric, Zero Sum Games

"Rock-paper-scissors is a symmetric and zero sum game; each player has the same strategies and payoffs associated with those strategies, and the sum of both players' payoffs in each outcome is zero. Whenever a game meets these requirements, we know something about the outcome of the game: *each player's expected utility must equal zero in equilibrium.*"

3.6.2: Generalized Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0, 0	-x, x*	y*, -y
Paper	x*, -x	0, 0	-z, z*
Scissors	-y, y*	z*, -z	0, 0

"To maintain the flavor of rock-paper-scissors, constrain x, y, and z such that $x > 0$, $y > 0$, and $z > 0$. This ensures that paper still beats rock, scissors still trumps paper, and rock still destroys scissors. Howeer, by letting the exogneous variables be any strictly positive value, we can vary the lethality of each strategy against each other strategy. For example, if x is extremely large, then paper obliterates rock, and the rock player must hand a large sum of money to the paper player."

"One interpretation of this game is like regular rock-paper-scissors, except different matchups of strategies result in different amounts of dollars changing hands. A more natural interpretation is of character selection in video games, particularly two dimensional fighting games. Characters have different strengths and weaknesses, which leads to good matchups against some opposing characters and bad matchups against others. Thus, a large value for x impliese that the "paper" character has a strong matchup versus the "rock" character, and so forth."

"Rock-paper-scissors is a finite game, so it must have an equilibrium. Since said equilibrium is not in pure strategies or mixtures involving only two strategies, both players must be mixing among all three strategies."

"It is interesting to note that the main determinant of each strategy's probability has nothing to do with that strategy."

"As the matrix shows, x is the benefit a paper player receives for beating rock. It does *not* appear in any of the payoffs involving a scissors player."

"The value for x represents paper's ability to smash rocks."

"All other things being equal, x does make paper more attractive. But the players can anticipate this. In turn, scissors becomes more viable as a way to counteract paper's strength against rock. In effect, the players balance out paper's advantage by increasing their frequency of scissors."

3.6.3: Mixed Strategies as Population Parameters

"From an empirical standpoint, mixed strategies seem bizarre. [...] people do not rely on randomizing devices to make their stategic decisions even if a Nash equilibrium tells them to. [...] Are people actually playing these games rationally if they are never randomizing?"

"We could interpret a mixed strategy Nash equilibrium as the population parameters of a larger game rather than a specific strategy of an individual in a two-player game."

"Under normal circumstances, playing rock as a pure strategy does not work in a Nash equilibrium. But the player's choice can still be rationally optimal. When teh gamer logs into the online inteface for his game, he joins thousands of players on the server. If an automated matchmaking system randomly picks his opponent, what is his expected utility for the game."

"Let σ_{rock} be the portion of the population that plays rock, σ_{paper} be teh portion of the population that plays paper, and $\sigma_{scissors}$ be the portion of the population that plays scissors. From last section's indifference equations, we know that the player's expected utility is 0 if $\sigma_{rock} = \frac{z}{x+y+z}$, $\sigma_{paper} = \frac{y}{x+y+z}$, and $\sigma_{scissors} = \frac{x}{x+y+z}$. But if those are the portions of other players using each strategy, the player also has an expected utility of 0 if he plays paper or scissors. Thus, his choice to play rock as a pure strategy is rational; he cannot choose a different strategy and achieve a greater expected utility."

Lesson 4.1: Infinite Strategy Spaces, Second Price Auctions, and the Median Voter Theorem

4.1.1: A Simple Game

4.1.2: A Game with No Equilibria

“Without best responses, there cannot be mutual best responses, and in turn there cannot be a Nash equilibrium.”

4.1.3: Hotelling’s Game and the Median Voter Theorem

“Two vendors are selling identical ice cream on a beach for $2 per cone. The vendors own carts andd must therefore choose where to set up shop. Since their products and prices are identical, the location is all that matters–beachgoers will purchase from whichever vendor is closest to them, and they will split the business evenly if the vendors are in the same location.”

Both must position themselves at the midpoint. From any other position, each player has a profitable deviation (to move toward the center to get a slightly higher proportion of beachgoers’ business).

“Hotelling’s game has a unique equilibrium: the vendors occupy the same spot halfway along the beach.”

“Similarly, consider Presidential elections in the United States. Immediately after the primary season ends, both the Democratic and Republican candidates dart woward the middle of the political spectrum. Political scientists note that this is also an application of Hotelling’s game, known as the **median voter theorem**.”

4.1.4: A Duel

4.1.5: Cournot Duopolistic Competition

4.1.6: Second Price Auctions

“Second price auctions have a large number of Nash equilibria. However, we focus on one in particular: when everyone submits the maximum price they are willing to pay for the good.”

“This equilibrium is remarkable for a number of reasons. First, it is strategy-proof. [...] Players in a second price auction can be comparatively oblivious. Because submitting their maximum prices is weakl dominant, bidders will never regret having told the truth.”

“Second, submitting one’s maximum weakly dominates all other strategies.”

“Third, it is honest. The bidders simply tell the auctioneer their true value for the good.”

“Fourth, number of bidders does not matter. Whether there are two or two million, submitting the maximum price remains optimal.”

“Fifth, bidders need not nkow others’ maximum prices. [... So far] we have assumed players have *complete information*–that is, they know each others’ payoffs, they know they know each other’s payoffs, and so forth. WHile complete information can go a long way, many interesting interactions involve one or both sides being in the dark. [...] Sometimes incomplete information can drastically change the outcome of an interaction. However, for a second price auction, it does not–everyone can still safely submit their maximum price.”