Game Theory 101: The Complete Textbook

William Spaniel

Available on Amazon

Lesson 1.1: The Prisoner's Dilemma and Strict Dominance

At its core, game theory is the study of strategic interdependence—that is, situations where my actions affect both my welfare and your welfare and vice versa. Strategic interdependence is tricky, as actors need to anticipate, act, and react. Blissful ignorance will not cut it.

The prisoner's dilemma is the oldest and most studied model in game theory, and its solution concept is also the simplest. As such, we will start with it. Two thieves plan to rob an electronics store. As they approach the backdoor, the police arrest them for trespassing. The cops suspect that the pair planned to break in but lack the evidence to support such an accusation. They therefore require a confession to charge the suspects with the greater crime.

Having studied game theory in college, the interrogator throws them into the prisoner's dilemma. He individually sequesters both robbers and tells each of them the following:

We are currently charging you with trespassing, which implies a one month jail sentence. I know you were planning on robbing the store, but right now I cannot prove it—I need your testimony. In exchange for your cooperation, I will dismiss your trespassing charge, and your partner will be charged to the fullest extent of the law: a twelve month jail sentence.

I am offering your partner the same deal. If both of you confess, your individual testimony is no longer as valuable, and your jail sentence will be eight months each.

If both criminals are self-interested and only care about minimizing their jail time, should they take the interrogator's deal?

1.1.1: Solving the Prisoner's Dilemma

The story contains a lot of information. Luckily, we can condense everything we need to know into a simple matrix:

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

We will use this type of game matrix regularly, so it is important to understand how to interpret it. There are two players in this game. The first player's strategies (keep "quiet" and "confess") are in the rows, and the second player's strategies are in the columns. The first player's payoffs are listed first for each outcome, and the second player’s are listed second. For example, if the first player keeps quiet and the second player confesses, then the game ends in the top right set of payoffs; the first player receives twelve months of jail time and the second player receives zero. Finally, as a matter of convention, we refer to the first player as a man and the second player as a woman; this will allow us to utilize pronouns like “he” and “she” instead of endlessly repeating “player 1” and “player 2.”

Which strategy should each player choose? To see the answer, we must look at each move in isolation. Consider the game from player 1’s perspective. Suppose he knew player 2 will keep quiet. How should he respond?

Let’s focus on the important information in that context. Since player 1 only cares about his time in jail, we can block out player 2’s payoffs with question marks:

	Quiet
Quiet	-1, ?
Confess	0, ?

Player 1 should confess. If he keeps quiet, he will spend one month in jail. But if he confesses, he walks away. Since he prefers less jail time to more jail time, confession produces his best outcome.

Note that player 2’s payoffs are completely irrelevant to player 1’s decision in this context—if he knows that she will keep quiet, then he only needs to look at his own payoffs to decide which strategy to pick. Thus, the question marks could be any number at all, and player 1’s optimal decision given player 2’s move will remain the same.

On the other hand, suppose player 1 knew that player 2 will confess. What should he do? Again, the answer is easier to see if we only look at the relevant information:

	Confess
Quiet	-12, ?
Confess	-8, ?

Confession wins a second time: confessing leads to eight months of jail time, whereas silence buys twelve. So player 1 would want to confess if player 2 confesses.

Putting these two pieces of information together, we reach an important conclusion—player 1 is better off confessing regardless of player 2’s strategy! Thus, player 1 can effectively ignore whatever he thinks player 2 will do, since confessing gives him less jail time in either scenario.

Let’s switch over to player 2’s perspective. Suppose she knew that player 1 will keep quiet, even though we realize he should not. Here is her situation:

	Quiet	Confess
Quiet	?, -1	?, 0

As before, player 2 should confess, as she will shave a month off her jail sentence if she does so.

Finally, suppose she knew player 1 will confess. How should she respond?

	Quiet	Confess
Confess	?, -12	?, -8

Unsurprisingly, she should confess and spend four fewer months in jail.

Once more, player 2 prefers confessing regardless of what player 1 does. Thus, we have reached a solution: both players confess, and both players spend eight months in jail. The justice system has triumphed, thanks to the interrogator’s savviness.

This outcome perplexes a lot of people new to the field of game theory. Compare the <quiet, quiet> outcome to the <confess, confess> outcome:

	Quiet	Confess
Quiet	-1, -1	?, ?
Confess	?, ?	-8, -8

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 just 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

Although a large branch of game theory is devoted to the study of expected utility, we generally consider each player's payoffs as a ranking of his most preferred outcome to his least preferred outcome. In the prisoner's dilemma, we *assumed* that players only wanted to minimize their jail time. Game theory does *not* force players to have these preferences, as critics frequently claim. Instead, game theory analyzes what should happen given what players desire. So if players only want to minimize jail time, we could use the negative number of months spent in jail as their payoffs. This preserves their individual orderings over outcomes, as the most preferred outcome is worth 0, the least preferred outcome is -12, and everything else logically follows in between.

Interestingly, the cardinal values of the numbers are irrelevant to the outcome of the prisoner's dilemma. For example, suppose we changed the payoff matrix to this:

	Quiet	Confess
Quiet	3, 3	1, 4
Confess	4, 1	2, 2

Here, we have replaced the months of jail time with an ordering of most to least preferred outcomes, with 4 representing a player's most preferred outcome and 1 representing a player's least preferred outcome. In other words, player 1 would most like to reach the <confess, quiet> outcome, then the <quiet, quiet> outcome, then the <confess, confess> outcome, then the <quiet, confess> outcome.

Even with these changes, confess is still always better than keep quiet. To see this, suppose player 2 kept quiet:

	Quiet
Quiet	3, ?
Confess	4, ?

Player 1 should confess, since 4 beats 3.

Likewise, suppose player 2 confessed:

	Confess
Quiet	1, ?
Confess	2, ?

Then player 1 should still confess, as 2 beats 1.

The same is true for player 2. First, suppose player 1 kept quiet:

	Quiet	Confess
Quiet	?, 3	?, 4

Player 2 ought to confess, since 4 beats 3.

Alternatively, if player 1 confessed:

	Quiet	Confess
Confess	?, 1	?, 2

Player 2 should confess as well, as 2 is greater than 1. Thus, regardless of what the other player does, each player's best strategy is to confess.

To be clear, this preference ordering exclusively over time spent in jail is just one way the players may interpret the situation. Suppose you and a friend were actually arrested and the interrogator offered you a similar deal. The results here do not generally tell you what to do in that situation, unless you and your friend only cared about jail time. Perhaps your friendship is strong, and both of you value it more than avoiding jail time. Since confessing might destroy the friendship, you could prefer to keep quiet if your partner kept quiet, which changes the ranking of your outcomes. Your preferences here are perfectly rational. However, we do not yet have the tools to solve the corresponding game. We will reconsider these alternative sets of preferences in Lesson 1.3.

Indeed, the possibility of alternative preferences highlights game theory's role in making predictions about the world. In general, we take a three step approach:

1) Make assumptions.
2) Do some math.
3) Draw conclusions.

We do steps 1 and 3 everyday. However, absent rigorous logic, some conclusions we draw may not actually follow from our assumptions. Game theory—the math from step 2 that this book covers—provides a rigorous way of ensuring that that our conclusions follow directly from the assumptions. Thus, correct assumptions imply correct conclusions. But *incorrect* assumptions could lead to ridiculous claims. As such, we must be careful (and precise!) about the assumptions we make, and we should not be surprised if our conclusions change based on the assumptions we make.

Nevertheless, for the given payoffs in the prisoner's dilemma, we have seen an example of *strict dominance*. We say that 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. In this example, confessing strictly dominated keeping quiet for both players. Unsurprisingly, players never optimally select strictly dominated strategies—by definition, a better option *always* exists regardless of what the other players do.

1.1.3: Applications of the Prisoner's Dilemma

The prisoner's dilemma has a number of applications. Let's use the game to explore optimal strategies in a number of different contexts.

First, consider two states considering whether to go to war. The military technology available to these countries gives the side that strikes first a large advantage in the fighting. In fact, the first-strike benefit is so great that each country would prefer attacking the other state even if its rival plays a peaceful strategy. However, because war destroys property and kills people, both prefer remaining at peace to simultaneously declaring war.

Using these preferences, we can draw up the following matrix:

	Defend	Attack
Defend	3, 3	1, 4
Attack	4, 1	2, 2

From this, we can see that the states most prefer attacking while the other one plays defensively. (This is due to the first-strike advantage.) Their next best outcome is to maintain the peace through mutual defensive strategies. After that, they prefer declaring war simultaneously. Each state's worst outcome is to choose defense while the other side acts as the aggressor.

We do not need to solve this game—we already have! This is the same game from the previous section, except we have exchanged the labels “quiet” with “defend” and “confess” with “attack.” Thus, we know that both states attack in this situation even though they both prefer the <defend, defend> outcome. The first-strike advantages trap the states in a prisoner’s dilemma that leads to war.

A similar problem exists with arms races. Imagine states must simultaneously choose whether to develop a new military technology. Constructing weapons is expensive but provides greater security against rival states. We can draw up another matrix for this scenario:

	Pass	Build
Pass	3, 3	1, 4
Build	4, 1	2, 2

Here, the states most prefer building while the other state passes. Following that, they prefer the <pass, pass> outcome to the <build, build> outcome; the states maintain the same relative military strength in both these outcomes, but they do not waste money on weaponry if they both pass. The worst possible outcome is for the other side to build while the original side passes. Again, we already know the solution to this game. Both sides engage in the arms race and build.

Now consider international trade. Many countries place tariffs (a tax) on imported goods to protect domestic industries even though this leads to higher prices overall.

We can use the prisoner’s dilemma to explain this phenomenon. A country can levy a tariff against another country’s goods or opt for no taxes. The best outcome for a country is to tax imports while not having the other country tax its exports. This allows the domestic industries to have an advantage at home and be competitive abroad, and the country also earns revenue from the tax itself. Free trade is the next best outcome, as it allows the lowest prices for each country’s consumers. Mutual tariffs is the next best outcome, as they give each country an advantage at home but a disadvantage abroad; ultimately, this leads to higher prices than the free trade outcome. The worst possible outcome is to levy no taxes while the other country enforces a tariff, as domestic industries stand no chance against foreign rivals.

Let’s toss that information into another matrix:

	No Tax	Tax
No Tax	3, 3	1, 4
Tax	4, 1	2, 2

We know this is a prisoner's dilemma and both sides will tariff each other's goods: taxing strictly dominates not taxing in this setup.

Finally, consider two rival firms considering whether to advertise their products. Would the firms ever want the government to pass a law forbidding advertisement? Surprisingly, if advertising campaigns only persuade a consumer to buy a certain brand of product rather than the product in general, the answer is yes. If one side places ads and the other does not, the firm with the advertising campaign cuts into the other's share of the market. If they both advertise, the ads cancel each other out, but they still have to pay for the campaigns.

If we look at the corresponding matrix, we see another classic example of the prisoner's dilemma:

	No Ads	Ads
No Ads	3, 3	1, 4
Ads	4, 1	2, 2

Thus, both sides advertise to preempt the other side's campaign. The ads ultimately cancel each other out, and the firms end the game in a worse position than had they both not placed ads.

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. In terms of the game matrix, the law pushed them from the <2, 2> payoff to the mutually preferable <3, 3> payoff. The law simultaneously satisfied politicians, as it made targeting children more difficult for all tobacco companies.

These examples illustrate game theory's ability to draw parallels between seemingly dissimilar situations. We have seen models of prisoner confession, wars, arms races, taxation, and advertisements. Despite the range of examples, each had an underlying prisoner's dilemma mechanism. In this manner, game theory allows us to unify a wide-range of life decisions under a single, unified framework.

1.1.4: Deadlock

The 2012 Summer Olympics badminton tournament provides an interesting case study of strategic manipulation. The tournament featured round-robin group play with a cut to a single-elimination quarterfinals bracket. Officials determined the seeding for the quarterfinals by the win/loss records during the round-robin matches.

In the morning matches of the final day of round-robin play, the second-best team in the world lost. While their previous victories still ensured that the team would reach the quarterfinals, their defeat pushed them into the lower half of the seeding. This had an interesting impact on the afternoon matches. 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!

To untangle the twisted logic at work here, consider the following game. Two players have to choose whether to *try* or *fail*. The quality of any corresponding outcome is diagrammed in the game matrix below:

	Try	Fail
Try	0, 0	-1, 1
Fail	1, -1	0, 0

Ordinarily, we would expect trying to be a good thing and failing to be a bad thing. The reverse is true here. Each team most prefers failing while the other team tries; this ensures the team in question will lose, drop into the lower part of the quarterfinals bracket, and thus avoid the world's second best team. The worst outcome for a team is for that team to try while the other team fails; this ensures that the original team wins the match but then must face a harder path through the single elimination bracket. If both try or both fail, then neither has an inherent strategic advantage.

Like the prisoner's dilemma, we can solve this game with strict dominance alone. Here, fail dominates try for both parties. We can verify this using the same process as before. First, suppose player 2 chooses tries:

	Try
Try	0, ?
Fail	1, ?

If player 1 tries, he earns 0; if he fails, he earns 1. Since 1 beats 0, he should fail in this situation.

Now consider player 1's response to player 2 failing:

	Fail
Try	-1, ?
Fail	0, ?

Again, fail triumphs: failing nets him 0 while trying earns him -1. Because failing is better than trying for player 1 regardless of player 2's strategy, fail strictly dominates try for him.

The same story holds for player 2. Consider her response to player 1 trying:

	Try	Fail
Try	?, 0	?, 1

If player 2 tries, she earns 0; if she fails, she earns 1. Thus, she ought to fail in this situation.

Now suppose player 1 failed instead:

	Try	Fail
Fail	?, -1	?, 0

Player 2 still ought to fail: -1 is less than 0. As a result, fail strictly dominates try for her as well. In turn, we should expect both of them to fail. Despite the absurdity of the outcome, the perverse incentives of the tournament structure make intentionally failing a sensible strategy!

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 her strategy incorrectly. Here are the generic game's payoffs:

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

Again, we can solve this game using strict dominance. Specifically, up strictly dominates down and left strictly dominates right. Let's verify this, starting with player 1's strategy:

	Left
Up	3, ?
Down	1, ?

We see that up is better than down, as 3 beats 1.

Repeating this for right, we focus on the following:

	Right
Up	4, ?
Down	2, ?

Once more, up is better than down, since 4 beats 2. So up is a better strategy than down regardless of what player 2 does.

Switching gears, suppose player 1 selected up. Then player 2 can focus on the following contingency:

	Left	Right
Up	?, 3	?, 1

Left is better than right in this case, as 3 is greater than 1.

Repeating this process a final time, player 2 now assumes player 1 will play down:

	Left	Right
Down	?, 4	?, 2

Left is still better than right, as 4 is greater than 2. Since left always beats right regardless of what player 1 does, left strictly dominates right, and therefore player 2 will play left. Thus, the outcome is <up, left>.

Thus, both players are locked into their strictly dominant strategy and will never achieve their best outcome unless the other makes a mistake. However, unlike in the prisoner's dilemma, no alternative outcome exists that is simultaneously better for both players than the <up, left> solution. As such, deadlock may be more intuitive, but it also tends to be substantively less interesting.

1.1.5: Strict Dominance in Asymmetric Games

We can use strict dominance on games even when they are not symmetric like the prisoner's dilemma or deadlock. For example, consider the arms race from earlier. Suppose that player 2 maintains her same payoffs. That is, she most prefers arming while her opponent passes and least prefers the opposite outcome. Meanwhile, she prefers neither side arming to both arming, as the balance of power remains the same but she saves on the costs of weapons. On the other hand, suppose that player 1 is a pacifist. He simply receives -1 for each party that builds weapons.

With that, we can construct the following payoff matrix:

	Pass	Build
Pass	0, 3	-1, 4
Build	-1, 1	-2, 2

Unlike before, each player has a distinct set of payoffs. But if we run through the same process as before, we will see that <pass, build> is the only reasonable solution.

Let's begin with player 1's choices. Suppose player 2 moved passes. How should player 1 respond?

	Pass
Pass	0, ?
Build	-1, ?

Recall that player 1 wants to minimize the total number of weapons. If he passes while player 2 passes, he achieves his best possible outcome. If he builds, he receives a -1. As such, he would want to pass in this situation.

Now suppose player 2 chose builds. Again, we need to find how player 1 should optimally respond:

	Build
Pass	-1, ?
Build	-2, ?

This time, player 1 cannot reach his best possible outcome. He can, however, minimize his losses by passing instead of building. Consequently, he would pass in this situation as well.

Combining the last two inferences together, we know player 1's optimal strategy: he will pass regardless of how player 2 behaves.

That leaves us to solve for player 2's strategy. Let's start she should respond to pass:

	Pass	Build
Pass	?, 3	?, 4

Player 2 can achieve her best possible outcome here by building, since she can exploit the shift in power. Since 4 beats 3, she will build in this situation.

Now suppose player 1 builds instead:

	Pass	Build
Build	?, 1	?, 2

Although player 2 can no longer reach her favorite outcome, she can at least keep pace with player 1's power by building here. As such, she would build if she knew that player 1 would build.

Despite the game's asymmetry, the game still has a solution in dominant strategies: <pass, build>. Player 2 achieves her best outcome, while player 1 must settle for a moderate result since he cannot stop her from arming.

<u>Conclusion</u>

Overall, strict dominance is a powerful tool in game theory. But while the concept is simple, applying it can be difficult. Even in matrix form, a game still has a lot of information. To successfully find dominated strategies, we must focus on one player's payoffs at a time. Above, we used question marks to isolate the relevant payoffs. When searching for strictly dominated strategies on your own, mentally block out the irrelevant payoffs and strategies in a similar manner.

Takeaway Points

1) Game theory is a mathematical method to ensure that assumptions imply conclusions.
2) Payoffs in a game matrix represent a player's preferences according to the assumptions.
3) Strategy x *strictly dominates* strategy y if it produces a higher payoff than y regardless of what all other players do.
4) Playing a strictly dominated strategy is irrational—another strategy always yields a better outcome.