The Prisoners' Dilemma

Player 1 \ Player 2	Cooperate	Defect
Cooperate	R, R	S, T
Defect	T, S	P, P

In the Prisoners' Dilemma we imagine two criminals that have been arrested and are being questioned by the police (for a detailed description of the story behind the PD see Dixit and Nalebuff p 11). This game, often called a "PD," is used frequently in political science and other social sciences. One common application is arms races between super powers. On your own, can you think of some other possible applications of the PD in international, national or local politics? What about outside politics?

Finding the Nash Equilibrium in the PD

In a strategic form game like the one above, a combination of strategies is represented by a cell in the box above. For example, the lower right hand cell represents the combination of strategies (Player 1: Defect; Player 2: Defect).
To find a Nash Equilibrium, find a cell where neither player has an incentive to leave that cell assuming the other player stays in it. So, in the above example, if both players cooperate (upper left hand cell), Player 1 will be better off defecting (assuming Player 2's move stays the same) because $T > R$. At the same time, Player 2 will be better off defecting (assuming Player 1's move stays the same) for the same reason.

$T > R > P > S$ for both players

(on your own try substituting numbers for T, R, P and S that satisfy $T > R > P > S$)

Finding Dominant Strategies: In the Prisoners' Dilemma

		Player 2
		Cooperate
Player 1	Cooperate	R R
	Defect	S T

		Player 2
		Defect
Player 1	Cooperate	T S
	Defect	P P

Finding the dominant strategy for Player 1:
First assume that Player 2 will always cooperate. Given that move by Player 2, what is the best response for Player 1? Since T > R, Player 1 is better off defecting if we know that Player 2 will cooperate.
But what if Player 2 never cooperates and instead always defects? In that case, Player 1's best response is still to defect because P > S.

Player 1's dominant strategy is to defect. Defecting is the one strategy that is always a best response regardless of what player 2 does.

Remember that for both players T > R > P > S.

To look for dominant strategies for player 2: First assume that we know that player 1 will cooperate. Is player 2 going to be better off cooperating or defecting? Since T > R, we say that player 1 is better off defecting. Now, assume that we know that player 1 will defect. Is player 2 going to be better off by defecting in this case as well? Yes. Because P > S player 2 is better off defecting. So player 2 has a dominant strategy of defecting because no matter what player 1 does, player 2 is best off defecting.

	Player 2	
	Cooperate	Defect
Player 1: Cooperate	R R	T S

	Player 2	
	Cooperate	Defect
Player 1: Defect	S T	P P

NOTE: Think back to finding the Nash Equilibrium on the first page. Can you see a relationship between the dominant strategies and the Nash Equilibrium in this game? Do you think that such a relationship exists in every game that has dominant strategies? What about every game that has a Nash Equilibrium?

Pareto Optimality And The Prisoners' Dilemma

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	R R	T S
	Defect	S T	P P

<u>Pareto Optimal:</u> An outcome (one of the quadrants in the two by two strategic form game above) is *pareto optimal* if and only if there is no other outcome that makes at least one of the players better off without hurting any of the players.

<u>Pareto Suboptimal:</u> An outcome is *pareto suboptimal* if and only if there is at least one outcome which makes at least one player better off without hurting any of the players.

<u>Pareto Improvement:</u> Outcome A is a *pareto improvement* over outcome B if outcome A is better than outcome B for at least one player without being worse for any other player than outcome B.

Let us consider the Nash Equilibrium that we found on the first page. On the first page, we found that the outcome (Defect, Defect) is the Nash Equilibrium. On the second and third pages we found that that equilibrium is produced because Defect is the dominant strategy for each player. However, is this the best outcome that both players could imagine? To answer this question, we will turn to the concept of pareto optimality. Is (Defect, Defect) *pareto optimal*? That is, are there any outcomes other than (Defect, Defect) that would make at least one player better off without hurting the other player? Compare (Defect, Defect) to (Defect, Cooperate). Player one is made better off because $T > P$. But player 2 is made worse off because $P > S$. For (Cooperate, Defect) player 2 is made better off but player 1 is made worse off. Now, compare (Defect, Defect) to the outcome (Cooperate, Cooperate). In this case player 1 is made better off because $R > P$. Player 2 is also made better off because $R > P$. So (Cooperate, Cooperate) is a *pareto improvement* over (Defect, Defect). The answer to our question is, "No, (Defect, Defect) is not the best outcome that both players could imagine."

This is the fundamental irony of the Prisoners' Dilemma. What is the logic behind the pareto suboptimal, nash equilibrium outcome? Can you think of any ways that the two players could overcome this problem and achieve the (Cooperate, Cooperate) outcome?

The Prisoners' Dilemma in Multiple Iterations

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	R, R	S, T
	Defect	T, S	P, P

Remember that for both players T > R > P > S

The tragedy of the PD is that even though both players would be better off Cooperating with each other, they always end up Defecting. Or in more technical terms, the dominant strategies only generate one Nash Equilibrium which is for both players to Defect and that NE is pareto suboptimal (or pareto inefficient or pareto inferior).

How can we "solve" this dilemma? That is, how can we get the two players to both Cooperate with each other and avoid the pareto suboptimal equilibrium? One of the most important approaches to this problem is to suppose that the two players are playing the PD with each other over and over and over again. This is called "multiple iterations." The reasoning here is that players will cooperate now because they want their "opponents" to cooperate in the future.

Lets start by considering what happens when the number of rounds of a game are known.

WHENEVER THE NUMBER OF ROUNDS OF A MULTIPLE ITERATION PD IS KNOWN (i.e. you know when the game will end), THE DOMINANT STRATEGY REMAINS TO DEFECT ALL THE TIME.

But what if the number of rounds is not known? That is what if the game either goes on for ever or at least we don't know when it will end. This is where the two players can actually arrive at cooperation but only under certain circumstances. The requirement now is that both players have to actually CARE about the future. If neither player cares what happens in the future rounds, then it is just like those future rounds didn't exist.

How can we measure how much each player cares about the future? In game theory we use the concept of a "discount factor" to measure the value of the future. If you have had some economics you may be familiar with the concept of the discount factor. We use the symbol, δ, to indicate the discount factor.

Two Strategies In The Multi-Shot PD

Grim Trigger: Grim Trigger is a strategy one can adopt for a multi-shot PD. If player 1 is playing Grim Trigger she will cooperate in the initial round and for every round thereafter so long as her opponent cooperates. However, if player 2 should ever defect, player 1 will respond by playing defect forever thereafter.

The tables below show how two players would play if there were in an infinitely iterated PD and Player 1 played Grim Trigger and Player 2 Defected in the first round. The first table shows the strategies and the second table shows the payoffs.

	Round 1	Round 2	Round 3	Round 4	…	Round N
Player 1	Cooperate	Defect	Defect	Defect	…	Defect
Player 2	Defect	Defect	Defect	Defect	…	Defect

	Round 1	Round 2	Round 3	Round 4	…	Round N
Player 1	S	δP	$\delta^2 P$	$\delta^3 P$	…	$\delta^n P$
Player 2	T	δP	$\delta^2 P$	$\delta^3 P$	…	$\delta^n P$

The tables below show how two players would play if there were in an infinitely iterated PD and Player 1 played Grim Trigger and Player 2 Cooperated in the first round. The first table shows the strategies and the second table shows the payoffs. These tables also depict what it would look like if both players were playing Grim Trigger (because then both players would cooperate first and the keep cooperating so long as their opponent cooperated).

	Round 1	Round 2	Round 3	Round 4	…	Round N
Player 1	Cooperate	Cooperate	Cooperate	Cooperate	…	Cooperate
Player 2	Cooperate	Cooperate	Cooperate	Cooperate	…	Cooperate

	Round 1	Round 2	Round 3	Round 4	…	Round N
Player 1	R	δR	$\delta^2 R$	$\delta^3 R$	…	$\delta^n R$
Player 2	R	δR	$\delta^2 R$	$\delta^3 R$	…	$\delta^n R$

To figure out the threshold δ^*, we equate the utility from adopting both strategies and we solve for δ^*:

U (defecting against Grim): $T+\delta P+\delta^2 P+\delta^3 P\ldots+\delta^n P = T+\delta P(1+\delta+\delta^2+\delta^3...+\delta^n)$
U (cooperating with Grim): $R+\delta R+\delta^2 R+\delta^3 R\ldots+\delta^n R = R(1+\delta+\delta^2+\delta^3...+\delta^n)$

But $\mathbf{1+\delta+\delta^2+\delta^3...+\delta^n}$ **equals** $\mathbf{1/(1-\delta)}$, because if $X=1+\delta+\delta^2+\delta^3...+\delta^n$, then $\delta X=\delta+\delta^2+\delta^3+\delta^4...+\delta^{n+1}$, and $X-\delta X=1$, which means that $X(1-\delta)=1$ or $X=1+\delta+\delta^2+\delta^3...+\delta^n =1/(1-\delta)$. So…

U (defecting against Grim): $T+\delta P(1+\delta+\delta^2+\delta^3...+\delta^n) = T+\delta P/(1-\delta) = (T-T\delta+\delta P)/(1-\delta)$
U (cooperating with Grim): $R(1+\delta+\delta^2+\delta^3...+\delta^n) = R/(1-\delta)$

Equate the two to solve for δ^*: $(T-T\delta+\delta P)/(1-\delta) = R/(1-\delta)$, eliminate the common denominator to get $T-T\delta+\delta P = R$, then rearrange to get $T-R= \delta(T-P)$, or

$$\delta^* = T - R/T - P$$

IF $\delta^* > T - R/T - P$, then it makes sense to cooperate with a Grim player.
IF $\delta^* < T - R/T - P$, then it makes sense to defect on a Grim player.

The same approach leads to identical results for a threshold δ for tit-for-tat. But one defection against a TFT opponent does not terminate prospects for future cooperation. A single defection followed by cooperation at any point would make sense if **$\delta < T - R/R - S$**, because:

U (defecting once after cooperating, then cooperating again with TFT): $T+\delta S$
U (Cooperating with TFT on 2 consecutive moves after uninterrupted cooperation): $R+\delta R$

Equate the two to solve for $\delta^{\blacklozenge}$: $T+\delta S= R+\delta R$, which is equivalent to $\delta(R-S)=T-R$, or **$\delta^{\blacklozenge}=T - R/R - S$**

Note that $\delta^{\blacklozenge}$ could be impossibly high (higher than 1) to allow for continuous cooperation, thus making it impossible to have (TFT,TFT) as an equilibrium.

To sum up…

If both players are playing Grim Trigger, an outside observer would see only round after round of cooperation. The same is true when both players are playing Tit-For-Tat.

There are two ways of defecting. The first way to defect is when a player defects in the first round and goes back to cooperating thereafter (which would be suicidal against Grim Trigger but might work against Tit-For-Tat). The second way to defect from a contingent strategy is to defect at once and then forever.

The reason these two ways of defecting involve defecting in the first round is that the motive to defect is to get that Temptation payoff. The longer you wait to go for the temptation, the more it is discounted (and the less it is worth).

The first way of defecting makes sense against Tit-For-Tat when $\delta < T - R/R - S$.
Only when **$\delta > T - R/R - S$** will a player play Tit-For-Tat against Tit-For-Tat in equilibrium.

If the player defects against Tit-For-Tat by defecting in the first round and forever thereafter, it must be the case that $\delta < T - R/T - P$. Only when **$\delta > T - R/T - P$** will a player play Tit-For-Tat against Tit-For-Tat in equilibrium.

When both conditions are satisfied (**$\delta > T - R/R - S$** and **$\delta > T - R/T - P$**) then (Tit-For-Tat, Tit-For-Tat) is a Nash Equilibrium.

There is only one way to defect against Grim Trigger. That is to defect in the first round and forever thereafter. Therefore, Grim Trigger only needs one condition to support it as an equilibrium strategy against itself. When **$\delta > T - R/T - P$**, a player will play Grim Trigger against Grim Trigger in equilibrium.

For an excellent discussion of equilibrium strategies in repeated Prisoners' Dilemmas see pages 264-268 of <u>Game Theory for Political Scientists</u> by James D. Morrow. The information in this handout was derived largely from Morrow's discussion).