3.2 — Repeated Games

ECON 316 • Game Theory • Fall 2021

Ryan Safner

Assistant Professor of Economics

safner@hood.edu

ryansafner/gameF21

gameF21.classes.ryansafner.com

Outline

[When Pure Strategies Won't Work]

[MSNE in Constant Sum Games]

[Coordination Games: PSNE and MSNE]

Prisoners' Dilemma, Reprise

Prisoners' Dilemma, Reprise

Prisoners' Dilemma, Reprise

- Not technically a Prisoners' Dilemma!
  - Game affected by Joker's threat to blow both of them up at midnight if nobody acts
- Both players have a *weakly*-dominant strategy to Detonate
- What is/are the Nash equilibrium/equilibria?

		Prisoners	
		Detonate	Nothing
Civilians	Detonate	-100 -100	0 -100
	Nothing	-100 0	-100 -100

Prisoners' Dilemma, Reprise

- A true prisoners' dilemma:

$$a > b > c > d$$

- Each player's preferences:
    - 1st best: you Defect, they Coop. ("temptation payoff")
    - 2nd best: you both Coop.
    - 3rd best: you both Defect
    - 4th best: you Coop., they Defect ("sucker's payoff")
- Nash equilibrium: (Defect, Defect)
    - (Coop., Coop.) an unstable Pareto improvement

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	b_1 b_2	d_1 a_2
	Defect	a_1 d_2	c_1 c_2

Prisoners' Dilemma: How to Sustain Cooperation?

- We'll stick with these specific payoffs for this lesson
- How can we sustain cooperation in Prisoners' Dilemma?

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Repeated Games

Repeated Games: Finite and Infinite

- Analysis of games can change when players encounter each other *more than once*

- **Repeated games**: the same players play the same game multiple times, two types:

- Players know the *history* of the game with each other

- **Finitely-repeated game**: has a known final round

- **Infinitely-repeated game**: has no (or an unknown) final round

Finitely-Repeated Games

Finitely-Repeated Prisoners' Dilemma

- Suppose a prisoners' dilemma is played for 2 rounds
- Apply **backwards induction:**
    - What should each player do in the final round?

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Finitely-Repeated Prisoners' Dilemma

- Suppose a prisoners' dilemma is played for 2 rounds
- Apply **backwards induction:**
  - What should each player do in the final round?
  - Play dominant strategy: **Defect**
  - Knowing each player will Defect in round 2/2, what should they do in round 1?

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Finitely-Repeated Prisoners' Dilemma

- Suppose a prisoners' dilemma is played for 2 rounds

- Apply **backwards induction:**
    - What should each player do in the final round?
    - Play dominant strategy: **Defect**
    - Knowing each player will Defect in round 2/2, what should they do in round 1?
        - No benefit to playing Cooperate
        - No threat punish Defection!

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Finitely-Repeated Prisoners' Dilemma

- Suppose a prisoners' dilemma is played for 2 rounds

- Apply **backwards induction:**

- Both **Defect** in round 1 (and round 2)

- No value in cooperation over time!

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	3 3	1 4
	Defect	4 1	2 2

Finitely-Repeated Prisoners' Dilemma

- For any game with a unique PSNE in a one-shot game, as long as there is a known, finite end, Nash equilibrium is the same

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Finitely-Repeated Prisoners' Dilemma

- In experimental settings, we tend to see people cooperate in early rounds, but close to the final round (if not the actual final round), defect on each other

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Infinitely-Repeated Games

Infinitely-Repeated Games

- Finitely-repeated games are interesting, but rare
    - How often do we know for certain when a game/relationship we are in will end?
- Some predictions for finitely-repeated games don't hold up well in reality
    - Ultimatum game, prisoners' dilemma
- We often play games or are in relationships that are **indefinitely repeated** (have no *known* end), we call them **infinitely-repeated games**

Infinitely-Repeated Games

- There are two nearly identical interpretations of infinitely repeated games:
  1. Players play *forever*, but discount (payoffs in) the future by a constant factor
  2. Each round the game might end with some constant probability

First Intepretation: Discounting the Future

- Since we are dealing with payoffs in the future, we have to consider players' **time preferences**

- Easiest to consider with monetary payoffs and the **time value of money** that underlies finance

$$PV = \frac{FV}{(1+r)^t}$$

$$FV = PV(1+r)^t$$

Present vs. Future Goods

- **Example**: what is the present value of getting \$1,000 one year from now at 5% interest?

$$PV = \frac{FV}{(1+r)^n}$$

$$PV = \frac{1000}{(1+0.05)^1}$$

$$PV = \frac{1000}{1.05}$$

$$PV = \$952.38$$

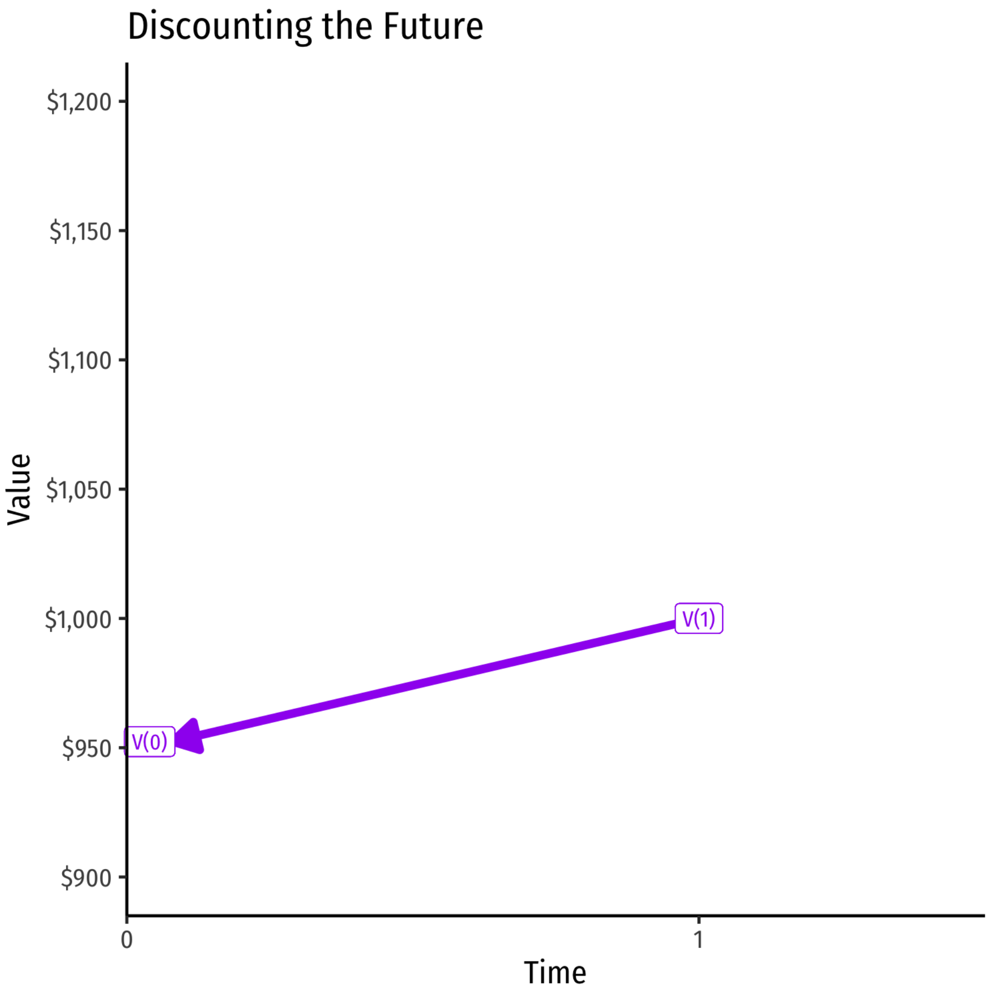

Present vs. Future Goods

- **Example**: what is the *future* value of $1,000 lent for one year at 5% interest?

$$FV = PV(1+r)^n$$
$$FV = 1000(1+0.05)^1$$
$$FV = 1000(1.05)$$
$$FV = \$1050$$

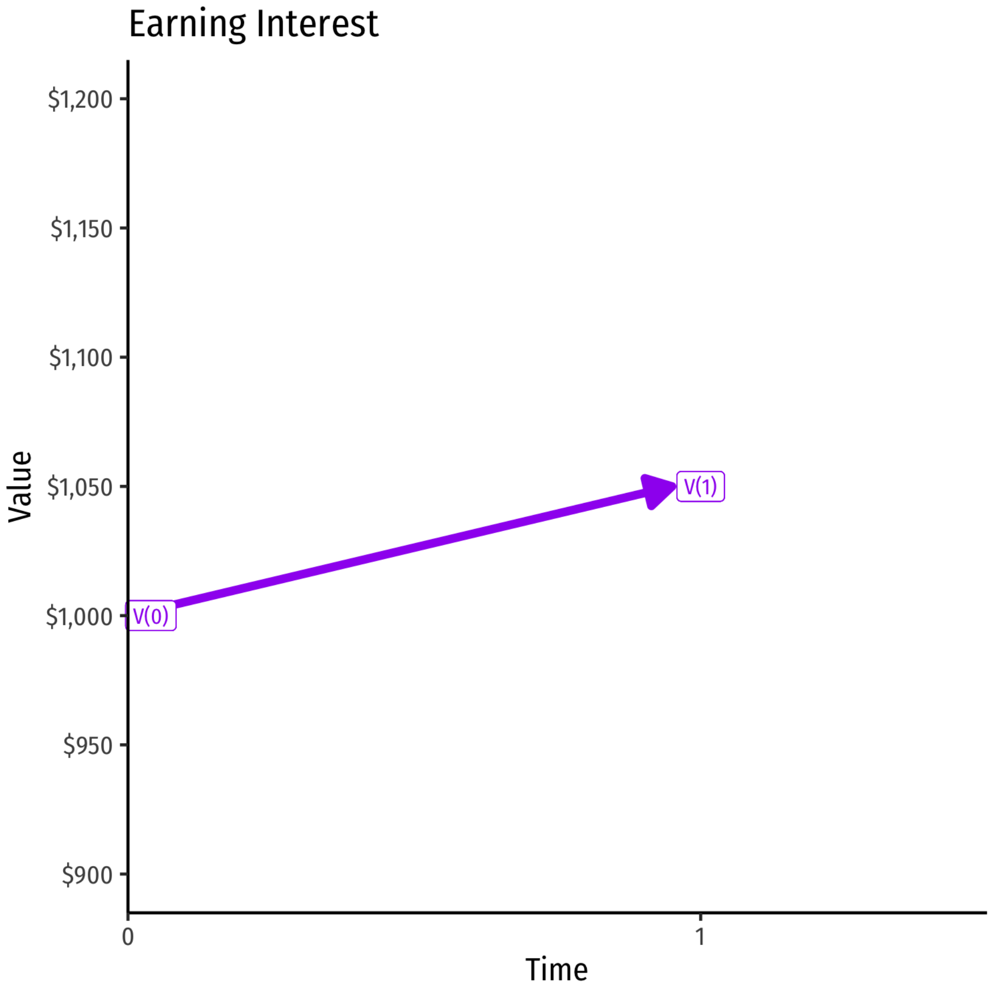

Discounting the Future

- Suppose a player values \$1 now as being equivalent to some amount with interest $1(1+r)$ *one period later*
    - i.e. \$1 with an r% interest rate over that period
- The **"discount factor"** is $\delta = \frac{1}{1+r}$, the ratio that future value must be multiplied to equal present value

Discounting the Future

$$\$1 \text{ now} = \delta\,\$1 \text{ later}$$

- If δ is low (r is high)
    - Players regard future money as worth much less than present money, **very impatient**
    - **Example**: $\delta = 0.20$, future money is worth 20% of present money
- If δ is high (r is low)
    - Players regard future money almost the same as present money, **more patient**
    - **Example**: $\delta = 0.80$, future money is worth 80% of present money

Discounting the Future

Example: Suppose you are indifferent between having \$1 today and \$1.10 next period

$$\begin{aligned}\$1 \text{ today} &= \delta \$1.10 \text{ next period}\\ \frac{\$1}{\$1.10} &= \delta\\ 0.91 &\approx \delta\end{aligned}$$

- There is an implied interest rate of $r = 0.10$
- \$1 at 10% interest yields \$1.10 next period

$$\begin{aligned}\delta &= \frac{1}{1+r}\\ \delta &= \frac{1}{1.10} \approx 0.91\end{aligned}$$

Discounting the Future

- Now consider an infinitely repeated game
    - If a player receives payoff p in every future round, the **present value** of this infinite payoff stream is

$$p(\delta + \delta^2 + \delta^3 + \cdots)$$

- This is due to compounding interest over time
    - This infinite sum converges to:

$$\sum_{t=1}^{\infty} = \frac{p}{1 - \delta}$$

- Thus, the present discounted value of receiving p in every future round is $\left(\frac{p}{1-\delta}\right)$

Prisoners' Dilemma, Infinitely Repeated

- With these payoffs, the value of both **cooperating** forever is $\left(\frac{3}{1-\delta}\right)$
- Value of both **defecting** forever is $\left(\frac{2}{1-\delta}\right)$

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Alternatively: Game Continues Probabilistically

- **Alternate interpretation**: game continues with some (commonly known among the players) probability θ each round

- Assume this probability is independent between rounds (i.e. one round continuing has no influence on the probability of the *next* round continuing, etc)

Alternatively: Game Continues Probabilistically

- Then the probability the game is played T rounds from now is θ^T
- A payoff of p in every future round has a present value of

$$p(\theta + \theta^2 + \theta^3 + \cdots) = \left(\frac{p}{1-\theta}\right)$$

- This is similar to discounting of future payoffs; equivalent if $\theta = \delta$

Strategies in Infinitely Repeated Games

- Recall, a **strategy** is a complete plan of action that describes how you will react under all possible circumstances (i.e. moves by other players)

  - i.e. "if other player plays x, I'll play a, if they play y, I'll play b, if, ..., etc"
  - think about it as a(n infinitely-branching) game tree, **“what will I do at each node where it is my turn?”**

- For an infinitely-repeated game, **an infinite number of possible strategies exist!**

- We will examine a specific set of **contingent** or **trigger strategies**

Trigger Strategies

- Consider one (the most important) trigger strategy for an infinitely-repeated prisoners' dilemma, the **"Grim Trigger" strategy**:
    - **On round 1**: Cooperate
    - **Every future round:** so long as the history of play has been (Coop, Coop) in every round, play Cooperate. Otherwise, play Defect *forever.*
- "**Grim**" trigger strategy leaves no room for forgiveness: one deviation triggers *infinite punishment*, like the sword of Damocles

Payoffs in Grim Trigger Strategy

- If you are playing the **Grim Trigger strategy**, consider your opponent's incentives:
    - If you both *Cooperate* forever, you receive an infinite payoff stream of 3 per round

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	3 3	1 4
	Defect	4 1	2 2

$$3 + 3\delta + 3\delta^2 + 3\delta^3 + \cdots + 3\delta^\infty = \frac{3}{1 - \delta}$$

Payoffs in Grim Trigger Strategy

- This strategy is a Nash equilibrium as long there's no incentive to deviate:

$$\text{Payoff to cooperation} > \text{Payoff to one-time defection}$$

$$\frac{3}{1-\delta} > 4 + \frac{2\delta}{1-\delta}$$

$$\delta > 0.5$$

- If $\delta > 0.5$, then player will cooperate and not defect

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Payoffs in Grim Trigger Strategy

- $\delta > 0.5$ is sufficient to sustain cooperation under the grim trigger strategy
    - This is the most extreme strategy with the strongest threat

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Payoffs in Grim Trigger Strategy

- Two interpretations of $\delta > 0.5$ as a sufficient condition for cooperation:

1. δ as **sufficiently high discount rate**
    - Players are patient enough and care about the future (reputation, etc), will not defect
2. δ as **sufficiently high probability of repeat interaction**
    - Players expect to encounter each other again and play future games together

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

Other Trigger Strategies

- "Grim Trigger" strategy is, well, grim: a single defection causes infinite punishment with no hope of redemption
    - *Very useful* in game theory for understanding the “worst case scenario” or the *bare minimum* needed to sustain cooperation!
    - Empirically, most people aren't playing this strategy in life
    - Social cooperation hangs on by a thread: what if the other player makes a *mistake*? Or *you* mistakenly think they Defected?
- There are “nicer” trigger strategies

"Nicer" Strategies

- Consider a **"Forgiving Trigger" strategy**:
  - On round 1: Cooperate
  - Every future round: so long as the history of play has been (Coop, Coop) in every round, play Cooperate. Otherwise, play Defect for 3 rounds
    - Punishment, but lasts for 3 rounds, then reverts to Cooperation

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

"Nicer" Strategies

- Consider the **"Tit for Tat" strategy**:
    - On round 1: Cooperate
    - Every future round: Play the strategy that the other player played last round
        - Example: if they Cooperated, play Cooperate; if they Defected, play Defect

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

"Nicer" Strategies

- Consider the **"Tit for 2 Tats" strategy**:
  - On round 1: Cooperate
  - Every future round: Cooperate, unless the other player has played Defect twice, then play Defect

		Player 2	
		Cooperate	**Defect**
Player 1	**Cooperate**	3 3	1 4
	Defect	4 1	2 2

The Evolution of Cooperation

Robert Axelrod

1943—

- Research in explaining the **evolution of cooperation**
- Use prisoners' dilemma to describe human societies and evolutionary biology of animal behaviors
- Hosted a series of famous tournaments for experts to submit a strategy to play in an infinitely[1] repeated prisoners' dilemma

> “The contestants ranged from a 10-year-old computer hobbyist to professors of computer science, economics, psychology, mathematics, sociology, political science, and evolutionary biology.”

- *The Evolution of Cooperation* (1984)
- Among the most cited works in all of political science

[1] Each round had a 0.00346 probability of ending the game, ensuring on average 200 rounds of play

Axelrod, Robert, 1984, *The Evolutioon of Cooperation

The Evolution of Cooperation

Robert Axelrod

1943—

- Axelrod's discussion of successful strategies based on four properties:

  1. **Niceness**: cooperate, never be the first to defect
  2. **Be Provocable**: don't be suckered by being too nice, return defection with defection
  3. **Don't be envious**: focus on maximizing your own score, rather than ensuring your score is higher than your "partner's"
  4. **Don't be too clever**: clarity is essential for others to cooperate with you

- The winning strategy was, famously, **TIT FOR TAT**, submitted by Anatol Rapoport

Axelrod, Robert, 1984, *The Evolutioon of Cooperation

The Folk Theorem

The Folk Theorem

- Consider the **average payoff** to each player each round, depending on the strategies chosen
    - e.g. if both Cooperate forever, average payoff is (3,3) — both earn 3 every round

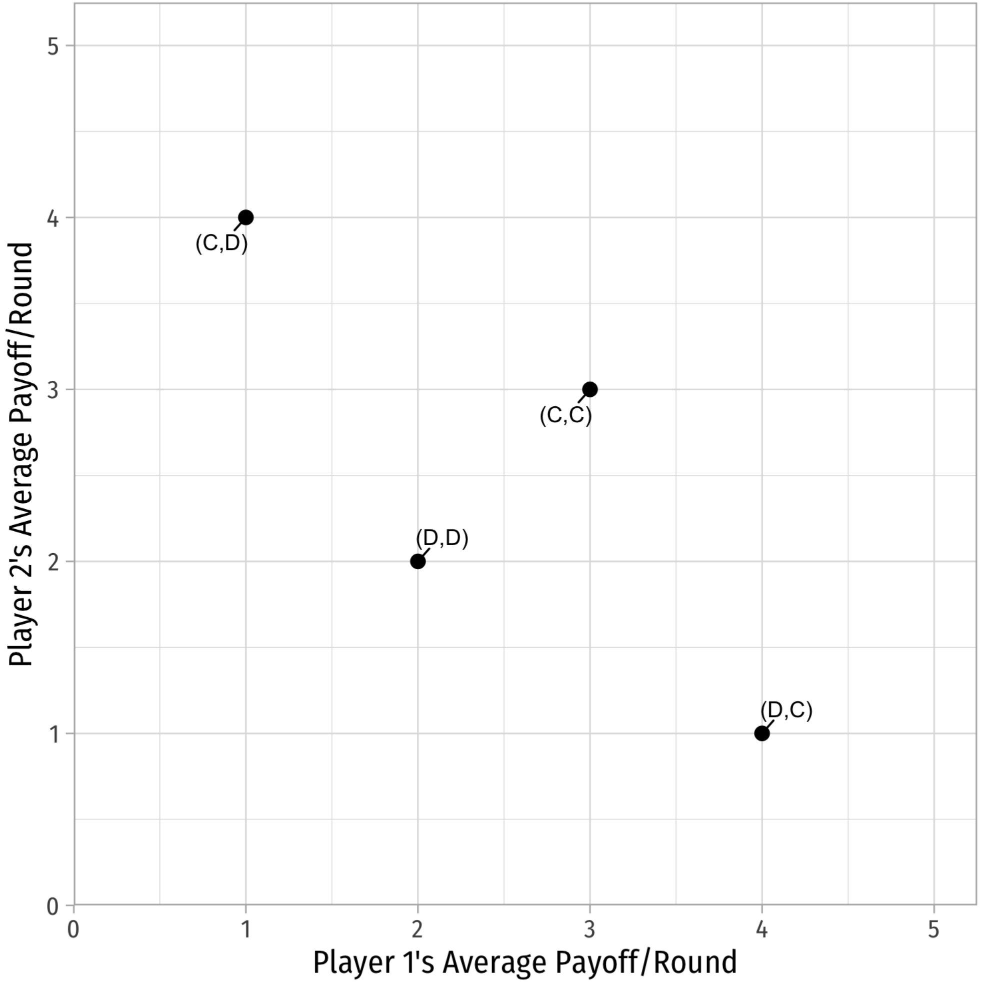

The Folk Theorem

- Consider the **average payoff** to each player each round, depending on the strategies chosen
    - e.g. if both Cooperate forever, average payoff is (3,3) — both earn 3 every round
- Consider the set of feasible average payoff
    - e.g. no way to produce average payoff of (6,6)
    - average payoff of (2.5, 2.5) is possible (players alternate between C and D each round)

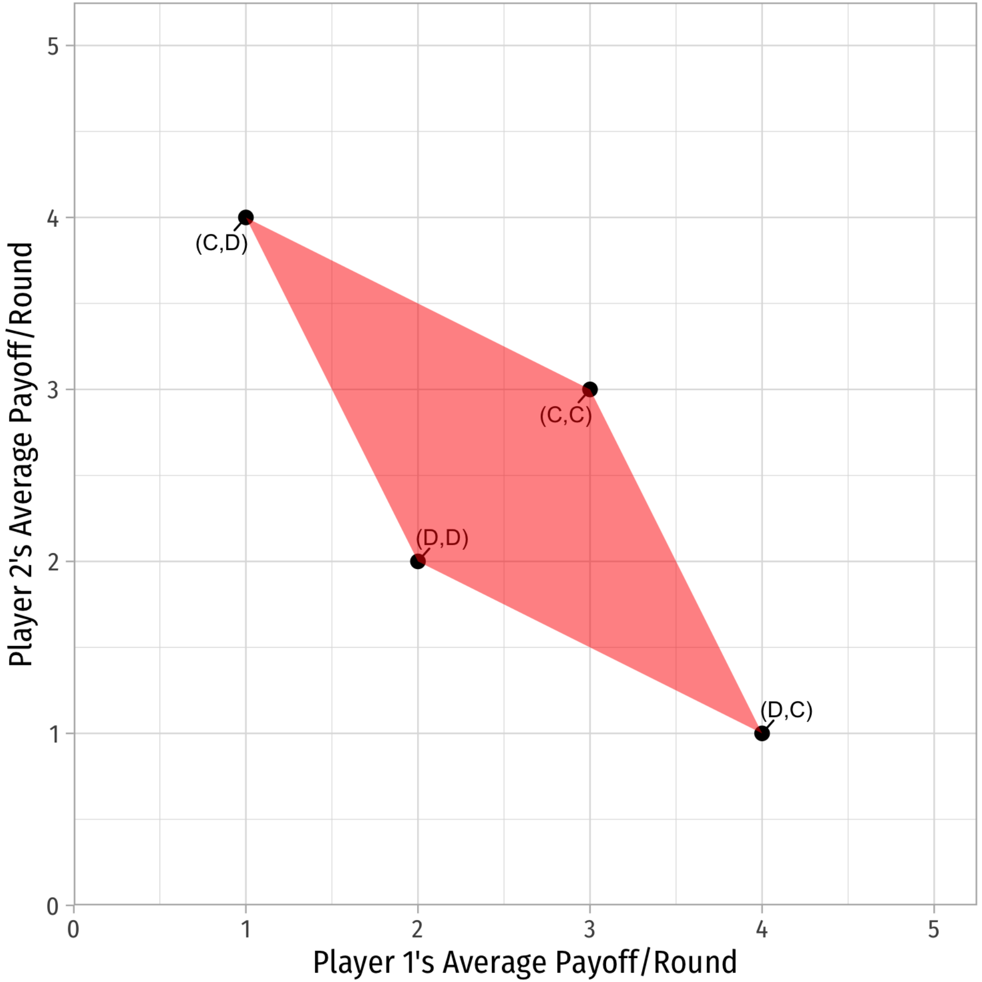

The Folk Theorem

- **Folk theorem**: any **individually rational** and feasible average payoff can be sustained with sufficiently high δ (or θ)

- An average payoff is **individually rational** if it is at least as good as the one-shot Nash equilibrium (Defect, Defect), i.e. (2,2) outcome

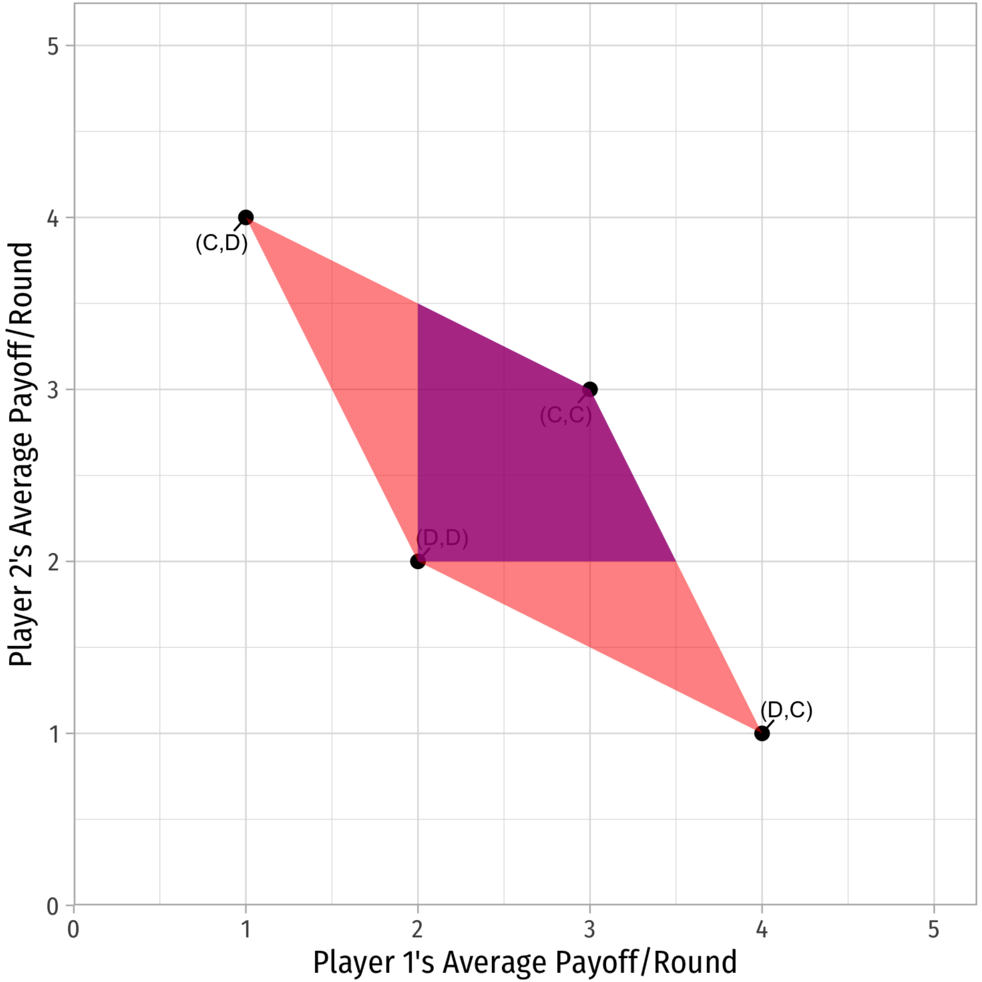

Folk Theorem: Simply Put

- **Folk theorem (simplified)**: Many strategies can sustain long-run cooperation if:
    - Each player can observe history
    - The value of future interactions must be sufficiently important to players
        - sufficiently high discount rate δ
        - sufficiently high probability of game continuing θ
- If this is true, *many* strategies can sustain long-run cooperation
    - Any in the teal set in the diagram before
    - *Grim trigger* is simply the bare minimum/worst case scenario (and, importantly, easiest to model!)

Assessing the Folk Theorem

- **The Good**: cooperation is possible, rational, and efficient!
    - Any improvement above (D,D) is a Pareto improvement for all players
- **The Bad**: lack of predictive power
    - Anything goes! Almost *any* outcome can be a sustainable equilibrium
    - This is why game theorists use the grim trigger strategy results as the *bare minimum* sufficient strategy for cooperation
- As temptation payoff increases relative to Nash equilibrium, need higher δ or θ to sustain cooperation