CSCI 357: Algorithmic Game Theory

Lecture 18: Repeated Games & Incentives in P2P

Shikha Singh

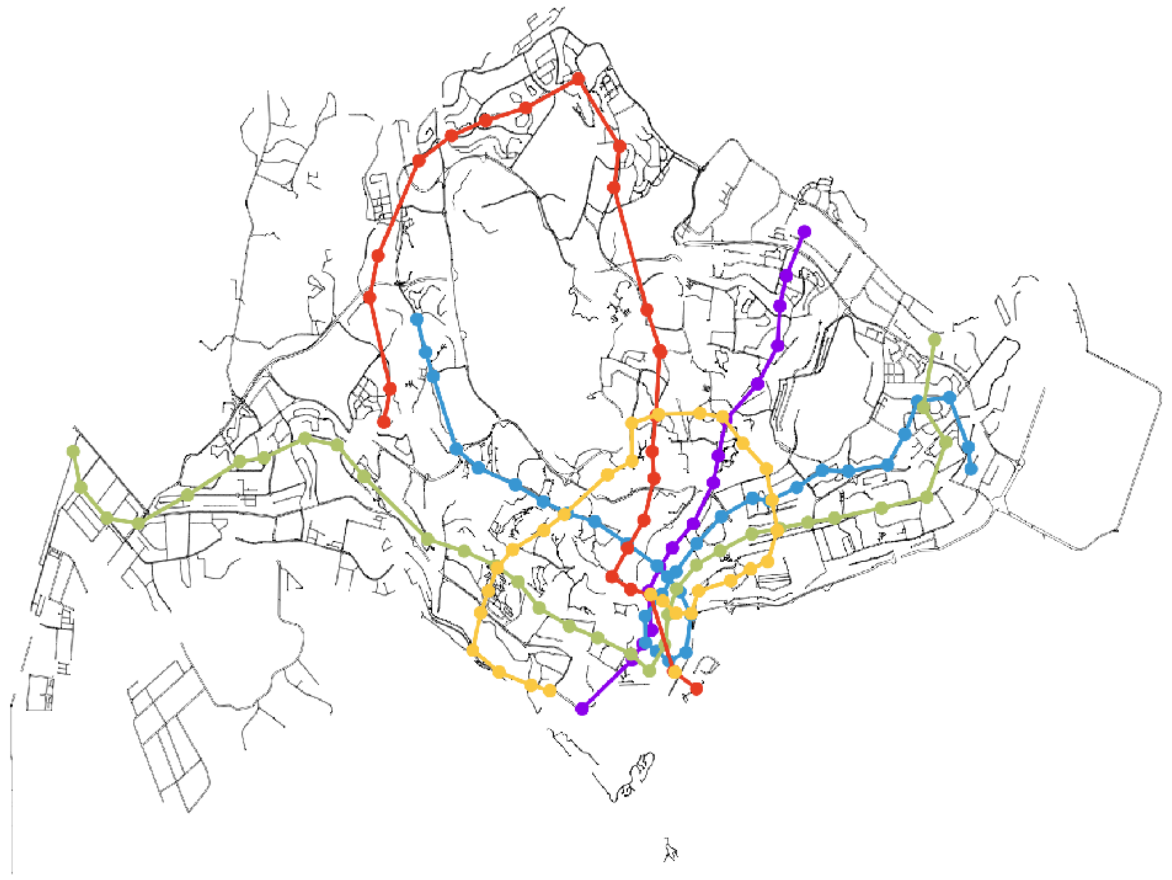

Announcements and Logistics

- HW 8 is due a day early: Wed April 20 11 pm
  - Shorter assignment: 3 questions
- Project proposal due Friday 5 pm
  - Google form that identifies partner, topic & ~150 words description of goals
- We will spend part of Thursday's lecture to discuss how to approach the project
- Drop in hours to discuss projects
  - Thursday 9 - 11 am, 4-5.30 pm
  - Friday 9 am - noon
  - Will send excel sheet so you can put down name

Questions?

Midterm 2 Logistics

- Will be a 24-hour take home exam
- Open book and notes
- Cumulative in topics with focus on second half
- LaTeX solutions and submit PDF via Gradescope
- Available to be taken during the 48 hour window:
  - Sun April 24 8am - Tues April 26 8 am
  - No lecture on Monday April 25

Questions?

Last Time

- Defined extensive-form games
  - Strategies are "plan" of action at every node in the tree where a player can called to act
- Nash equilbirium on the whole game tree can be suboptimal within subgames
- Subgame-perfect equilbrium (SPE): a strategy that is a Nash on every subgame
  - Can verify if s is a SPE using the single-deviation principle
  - Must always exist in any finite extensive form game
  - Can be found by backward induction
- Started discussion of **repeated games**

Today

- Discuss repeated games, especially repeated prisoner's dilemma
- Examples of SPE in repeated PD and how to analyze them
- How to apply the framework to repeated file sharing games in P2P systems

Repeated Games

Prisoner's Dilemma

- One shot game:
  - (D, D) is the unique DSE
- What if the players played it over and over again?
- Would cooperation emerge?

	C	D
C	4, 4	0, 5
D	5, 0	2, 2

	C	D
C	a, a	b, c
D	c, b	d, d

$c > a > d > b$

Split or Steal

- Players can choose split or steal the prize money
  - If both steal, no one gets any money
  - If one splits, other steals: the thief gets all the money
  - If both split: they share the only in half
- Weakly dominant action?
  - Steal weakly dominates Split for both players
- In both the video game and game show, the game is multi-stage and current decisions have future consequences
- Cooperation is often seen in all these situations
- https://www.youtube.com/watch?v=S0qjK3TWZE8&ab_channel=spinout3

	Split	Steal
Split	1/2, 1/2	0, 1
Steal	1, 0	0, 0

Motivation: Incentives in P2P

- P2P systems provide an intriguing case study of how a system evolves in response to incentive issues
  - Incentive properties vary widely across different protocols
- Peer-to-peer file sharing:
  - Way to distribute a file between users where they upload and download from each other
- P2P is now fundamental to blockchain platforms, such as Bitcoin and Ethereum
- **AGT view**: **do peers in a P2P system to have an incentive to cooperate**?
  - For file sharing, do users have incentive to upload while downloading from peers?

Failure of Centralization

- In the days of early internet, file sharing was done in an ad hoc way
- Napster (1999): provided a centralized, searchable directory listing which users have copies of various files (e.g. mp3s)
    - Matchmaker (matched up people who want file to people who have the file)
    - File transfer was then done directly between users
- Lawsuits against Napster for copyright infringement (2000s)
    - By RIAA, Metallica, etc
- After Napster failed to comply, it was shut down in 2001
- Napster's rise (25 million users) pointed to the demand for such systems but its failure motivated **decentralized designs**

Benefits of P2P

- Client-server model: server provider is associated with the server machines, users device is a client machine
    - These platforms need to make use of millions of distributed servers in order to cache content on machines close to users to provide low latency and maintaining this infrastructure
- In contrast, P2P systems there is no distinction between client and servers: each computer acts as both and is called a peer
- Main advantage: can scale well to large numbers of users while keeping the costs low for the initial uploader of the content
    - Provide robustness by avoiding a single point of failure
- Disadvantages: no control over content and who will download it, for how long the files will be available, etc

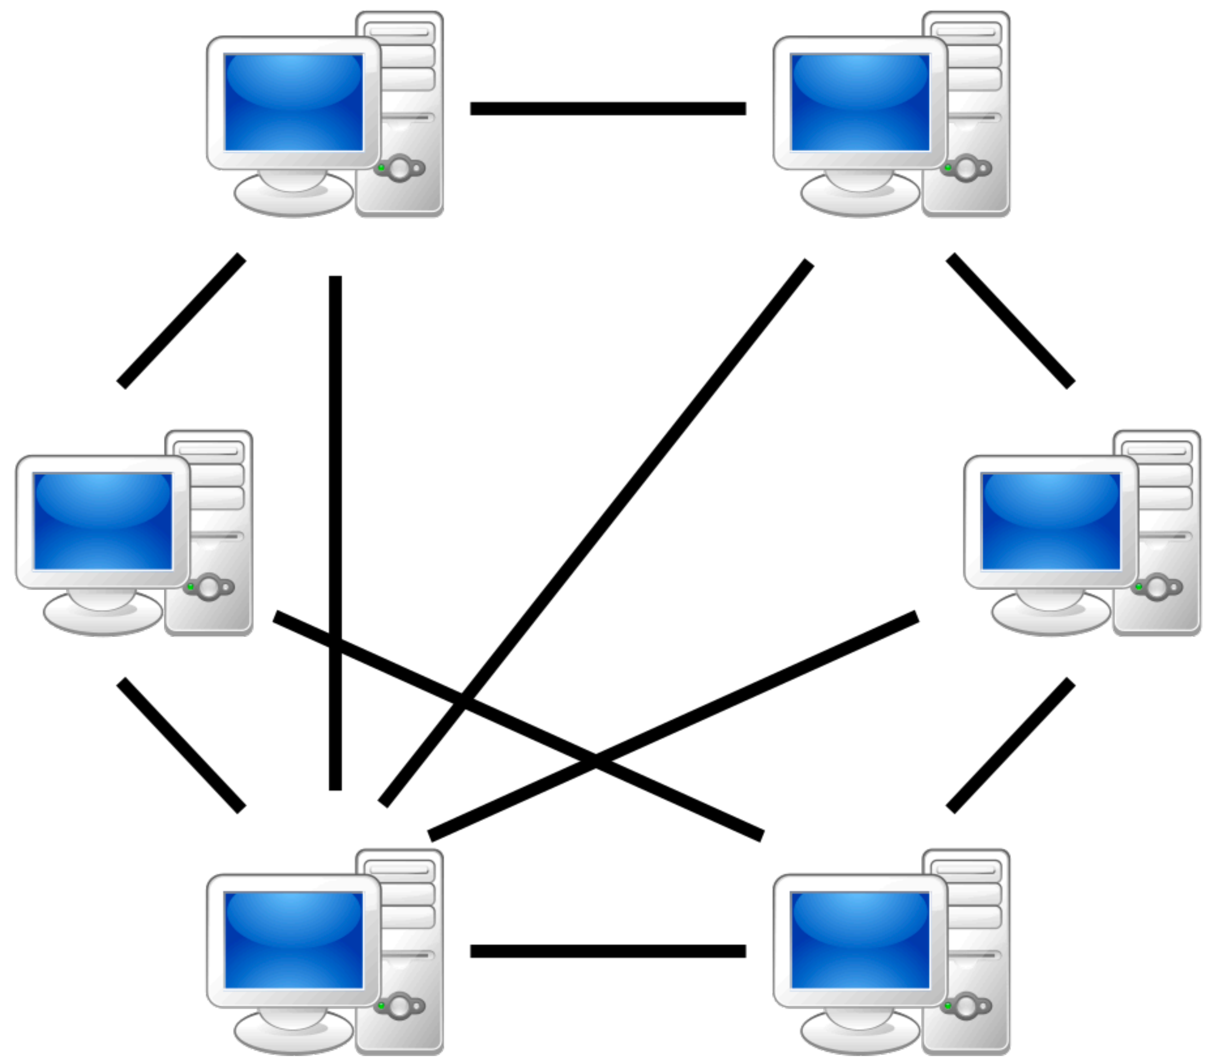

Decentralized: Gnutella

- First decentralized P2P network of its kind
- Design highlighted various incentive issues inherent in P2P networks
- Functionality of Gnutella rests on users conforming to the following behavior:
    - Upon receiving a file request, either upload the file to the requester (if the user has the file) or forward to other peers
- Problem with the design. Users were not given any incentive to actually behave in this way
- **Free-riding in Gnutella**:
    - A user is a free-rider who downloads but never uploads
  - A study by researchers showed that free-riding was the dominant behavior in Gnutella: 2/3rd of the users were free riders
  - In follow up study in 2005, free riding had climbed to 85% leading to the extinction of the system

File Sharing Game

- Consider two players: Aamir and Beth
- Aamir has a file that Beth wants and vice versa
- They simultaneously and independently decide whether or not to upload the requested file
- For each player, the benefit of receiving the file is 3 and the cost of uploading is 1 (bandwidth charges, opportunity costs, etc)

Aamir

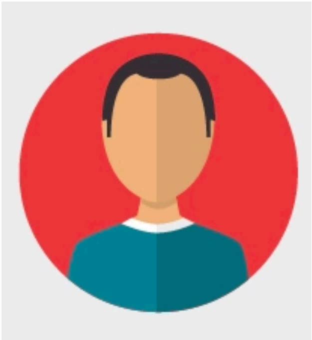

Beth

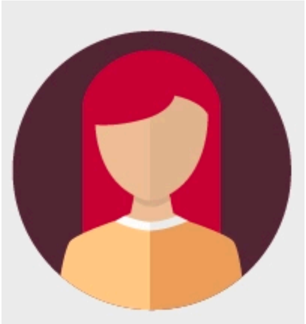

	Upload	Don't Upload
Upload		
Don't Upload		

Prisoner's Dilemma

- Our payoff matrix is just a variant of the prisoner's dilemma game from Lecture 2
- Each player has a strictly dominant strategy to defect
  - In this case, to not upload
- When Aamir and Beth play their dominant strategy neither uploads and each gets a payoff of zero
- Prisoner's dilemma summarizes the essential conflict between individual good and the collective good

	Upload	Don't Upload
Upload	$2, 2$	$-1, 3$
Don't Upload	$3, -1$	$0, 0$

Aamir

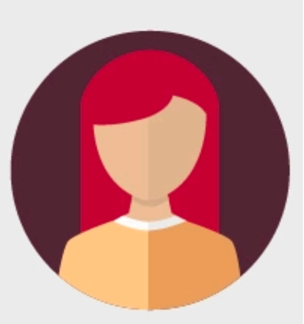

Beth

	C	D
C	a, a	b, c
D	c, b	d, d

$$c > a > d > b$$

Repeated Prisoner's Dilemma

- In real life examples of Prisoner's dilemma players do seem to cooperate: how can we explain this?
- **Intuition**: in real-life settings, the short term gain from defecting is outweighed by its long-term costs
- The model of repeated game try to capture this aspect:
  - How cooperation develops in long-term play
- **Idea**: Suppose we repeated Prisoner's dilemma n times
  - What behavior to be expect to see?

	Upload	Don't Upload
Upload	$2, 2$	$-1, 3$
Don't Upload	$3, -1$	$0, 0$

Aamir

Beth

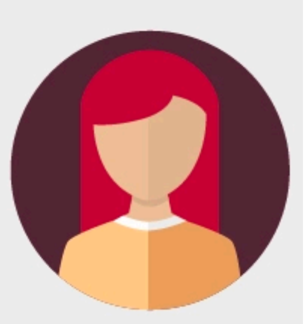

	C	D
C	a, a	b, c
D	c, b	d, d

$$c > a > d > b$$

Finitely Repeated Games

- Consider the following finitely-repeated Prisoner's dilemma
- **Finitely-repeated Prisoner's Dilemma**:
  - Aamir and Beth play the one-shot simultaneously move Prisoner's dilemma game n times (for some known $n \geq 1$)
  - The total utility of each is the sum of utilities across the n rounds
- Note that this is a sequential game in the sense that the action chosen by players in round i can depend on the history

Finitely Repeated Games

- Consider a normal-form game (**"stage game"**) $G = (N, \tilde{A}, \tilde{u})$ with action set $\tilde{A}_i$ for player i and utility $\tilde{u}_i(a)$ for player i on outcome profile $a = (a_1, \ldots, a_n)$
- **Finitely repeated game.** In a finitely-repeated game G^T, the stage game G is played by the same players for $T \geq 1$ rounds, such that each player has perfect information about the history of actions in all previous rounds
- To represent a repeated game in **extensive-form**, we allow simultaneous moves in each round
  - The history of the game is now a sequence of action profiles (instead of a sequence of individual actions)
  - The utility of each player is the sum of utilities along the history

Example: Repeated PD

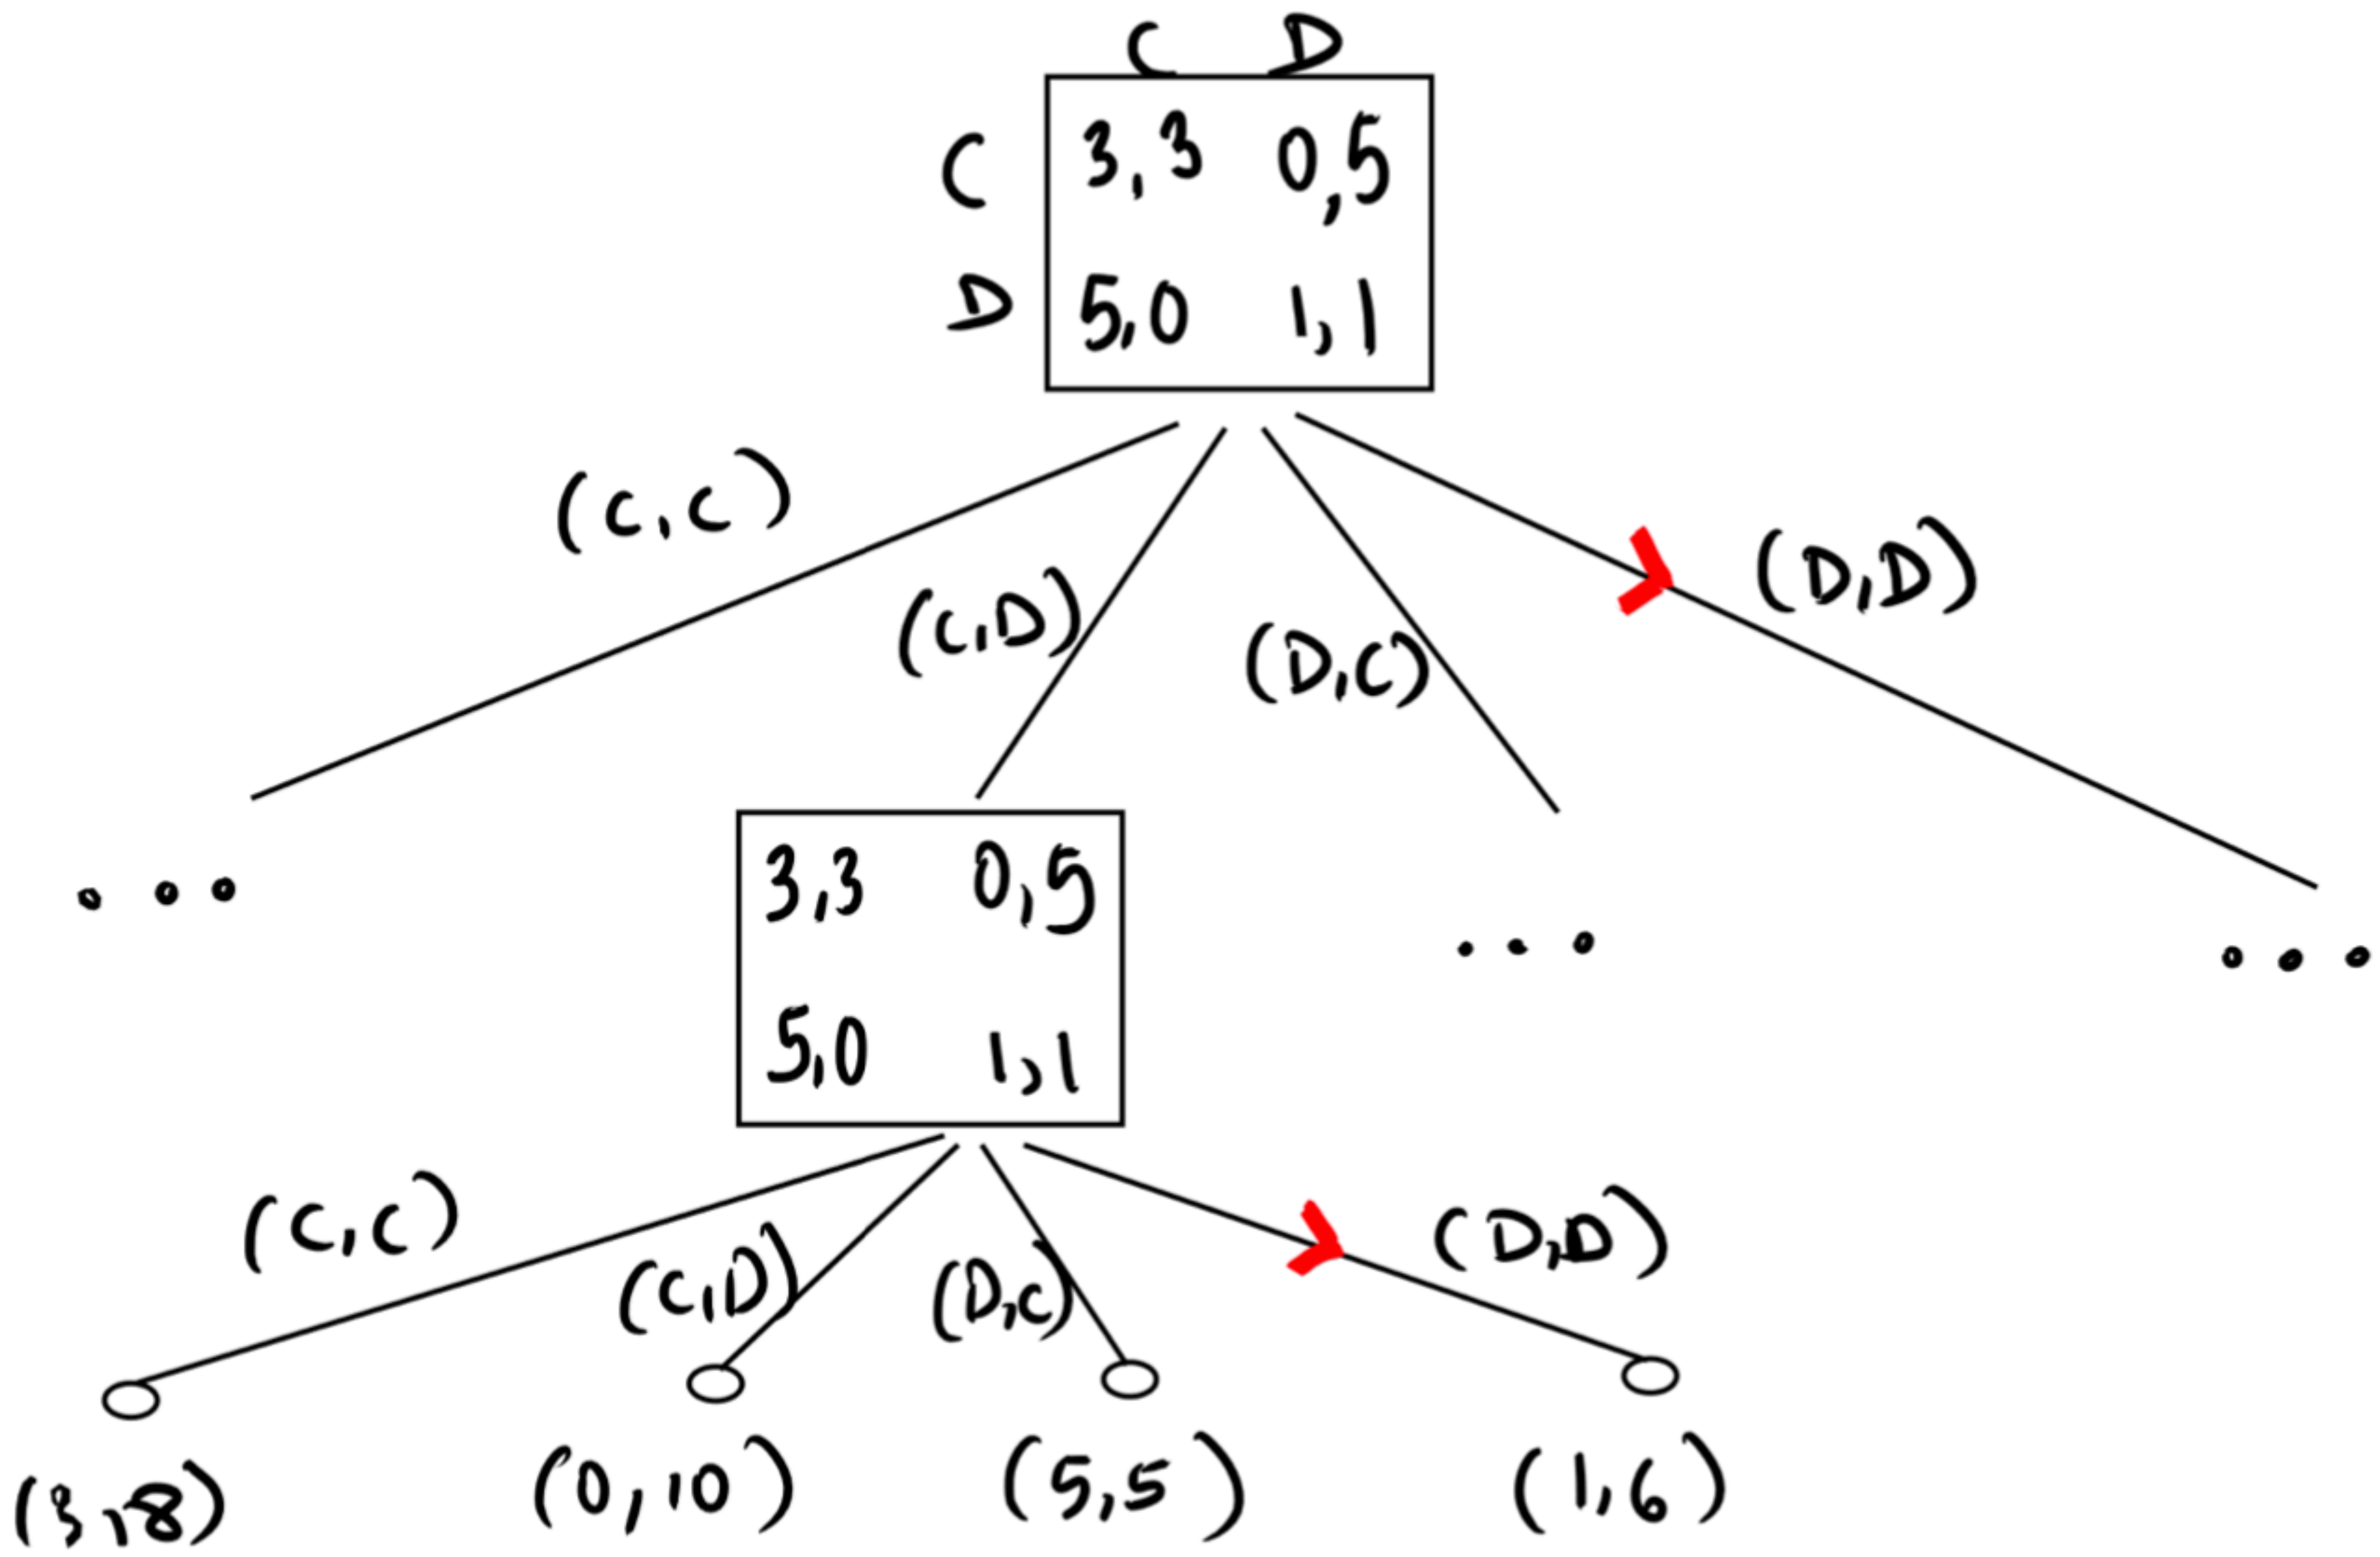

Credit: Textbook by Parkes and Seuken

Finitely Repeated Games

- Consider the following finitely-repeated Prisoner's dilemma
- **Finitely-repeated Prisoner's Dilemma**:
  - Aamir and Beth play the one-shot simultaneously move Prisoner's dilemma game n times (for some known $n \geq 1$)
  - The total utility of each is the sum of utilities across the n rounds
- Use backward induction: best response in round n?
- Given that (D, D) is best response in round n, what is the best response in round $n - 1$
  - Round $n - 1$ has no effect on future (round n): (D, D)
- Best response in any round?
- Unique SPE: **defect in each round**

Finitely Repeated Games

- We can generalize the trend we saw for repeated Prisoner's dilemma to any normal-form game
- **Theorem**. If the stage game has a unique Nash equilibrium, then the only subgame-perfect equilibrium strategy in a finitely-repeated game is to play the stage game Nash equilibrium strategy after each possible history.
- **Proof**. By backward induction, from the final period
  - Only best response in final round is to play the unique Nash equilibrium
  - In second-last round, the actions of the players does not effect the payoff in the last round, so best response is to play the unique Nash in this round and so on

Takeaways

- The model of finitely-repeated games fails to capture the emergence of cooperation in real-world settings
- What is possible missing from the model?
  - Element of **uncertainty about the future** (when does the game really come to an end?)
  - How players might value short term payment differently from payment that is long in the future
    - “**Discount**” future payments
- The model of **infinitely repeated games** captures this intuition
- In fact, there is a SPE of prisoner’s dilemma where (C, C) is sustained in every period: this shift happens when we move to infinitely repeated games

Infinitely Repeated Games

- **Infinitely-repeated Prisoner's dilemma**:
  - Aamir and Beth play the one-shot simultaneously move Prisoner's dilemma game
  - With probability δ (where $0 < \delta < 1$) the game continues for another round
  - With probability $1 - \delta$ the game ends at this round
- In the literature, this δ is often represented as players "**discounting the future payoffs**"
- Suppose players get utility a, a', a'' from three rounds, then their "discounted utility" is $a + \delta a' + \delta^2 a''$
- You can think of it as **expected utility** or **discounted utility**

Infinitely Repeated Games

- **Example**. Suppose the sequence of play in a three-round Prisoner's dilemma is $(C, C), (C, C)$ and (D, C)
- Suppose $\delta = 0.9$
- What is player 1's "discounted" utility
  - With probability 1 player 1 gets 2 in round 1
  - With probability 0.9 player 1 gets 2 in round 2
  - With probability 0.9^2 player 1 gets 3 in round 3
- Overall expected/discounted utility is thus $2 + 0.9(2) + 0.81(3) = 6.23$

	C	D
C	2, 2	−1, 3
D	3, −1	0, 0

Infinitely Repeated Games

- **Infinitely-repeated Prisoner's dilemma**:
  - Aamir and Beth play the one-shot simultaneously move Prisoner's dilemma game
  - With probability δ (where $0 < \delta < 1$) the game continues for another round
  - With probability $1 - \delta$ the game ends at this round
- Both players now want to maximize their total expected utility or total discounted utility defined as $u_i(h) = \sum_{k=0}^{\infty} \delta^k \tilde{u}_i(a^{(k)})$ where $a^{(k)}$ is the action profile chosen in the kth round
- Intuition if δ is sufficiently large, then cooperation should emerge
  - Note that $\delta = 0$, we are back to the one-shot game

When δ is sufficiently large: players are interpreted as "**patient**"

Infinitely Repeated Games

- Caution about infinitely-repeated games:
  - Can potentially have a **lot of equilibria**
  - Specific equilibria thus reduces predictive power
  - Known problem with the repeated-game framework and a topic of ongoing research
- Often the focus is on what type of strategies can be sustained by subgame-perfect and Nash equilibria
- First, we will discuss what type of **symmetric strategies** can be sustained in equilibrium
- Second, we will discuss what type of strategies seem to do well in an **asymmetric environment** based on **empirical study**

Trigger Strategies

- In repeated games, a **trigger strategy** essentially threatens the opponent with a “worse” punishment if they deviate from an implicit agreed upon action profile
- The most extreme (and unforgiving) trigger strategy is the **grim trigger strategy** that punishes forever after a single deviation
- Suppose Beth plays the following **grim trigger strategy**:
    - Start by cooperating, but if opponent has ever defected in the past, then defect
    - Otherwise, cooperate
- If Beth plays such a strategy, what is Aamir’s best response in a Nash or subgame perfect equilibrium?

Grim Trigger Strategy

- First let us reason about the Nash equilibria (before we think about subgame-perfect)
- If Aamir ever defects, Beth punishes him forevermore
- What is Aamir's best response?
- Suppose we are in stage i
  - If Aamir ever defected until now, then Beth's future behavior is fixed:
    - Aamir might as well defect forever
  - If Aamir has cooperated in the past $i-1$ stages for $i \geq 1$, should he continue to cooperate?

	C	D
C	$2, 2$	$-1, 3$
D	$3, -1$	$0, 0$

	stage i	stages $\geq i+1$
C	2	$\geq 2\delta$
D	3	≤ 0

Grim Trigger Strategy

- First let us reason about the Nash equilibria (before we think about subgame-perfect)
- If Aamir ever defects, Beth punishes him forevermore
- What is Aamir's best response?
- Suppose we are in stage i and Aamir has cooperated in each of the past $i-1$ stages for $i \geq 1$
  - Utility from playing C in stage i: $2 + 2\delta$
  - Utility from playing D in stage i: $3 + 0$
- For what values of δ is C a best response?
- Thus, the strategy pair **(Grim, Grim)** is a Nash equilibria if $\delta \geq 1/2$

	C	D
C	2, 2	−1, 3
D	3, − 1	0, 0

	stage i	stages $\geq i+1$
C	2	$\geq 2\delta$
D	3	≤ 0

Grim Trigger Strategy

- Thus, the strategy pair (Grim, Grim) is a Nash equilibria if $\delta \geq 1/2$
- Notice that this generates the outcome (C, C) is each stage of the game
- Is this also a subgame perfect equilibria?
- Intuitively, for a Nash equilibria to be subgame perfect, the **threats must be credible**
- Challenge in analyzing subgame perfect:
  - Need to account for behavior on histories that may never be reached under equilibrium
- Matters a lot how we exactly we define the strategies followed by the players

	C	D
C	$2, 2$	$-1, 3$
D	$3, -1$	$0, 0$

	stage i	stages $\geq i+1$
C	2	$\geq 2\delta$
D	3	≤ 0

Automaton Strategies

- Consider the following two automaton strategies for grim trigger
- Grim Trigger (left) is a subgame-perfect equilibrium for $\delta > 1/2$
- But the modified Grim Trigger (right) is **NOT** a subgame-perfect equilibrium for any $\delta \in (0,1)$

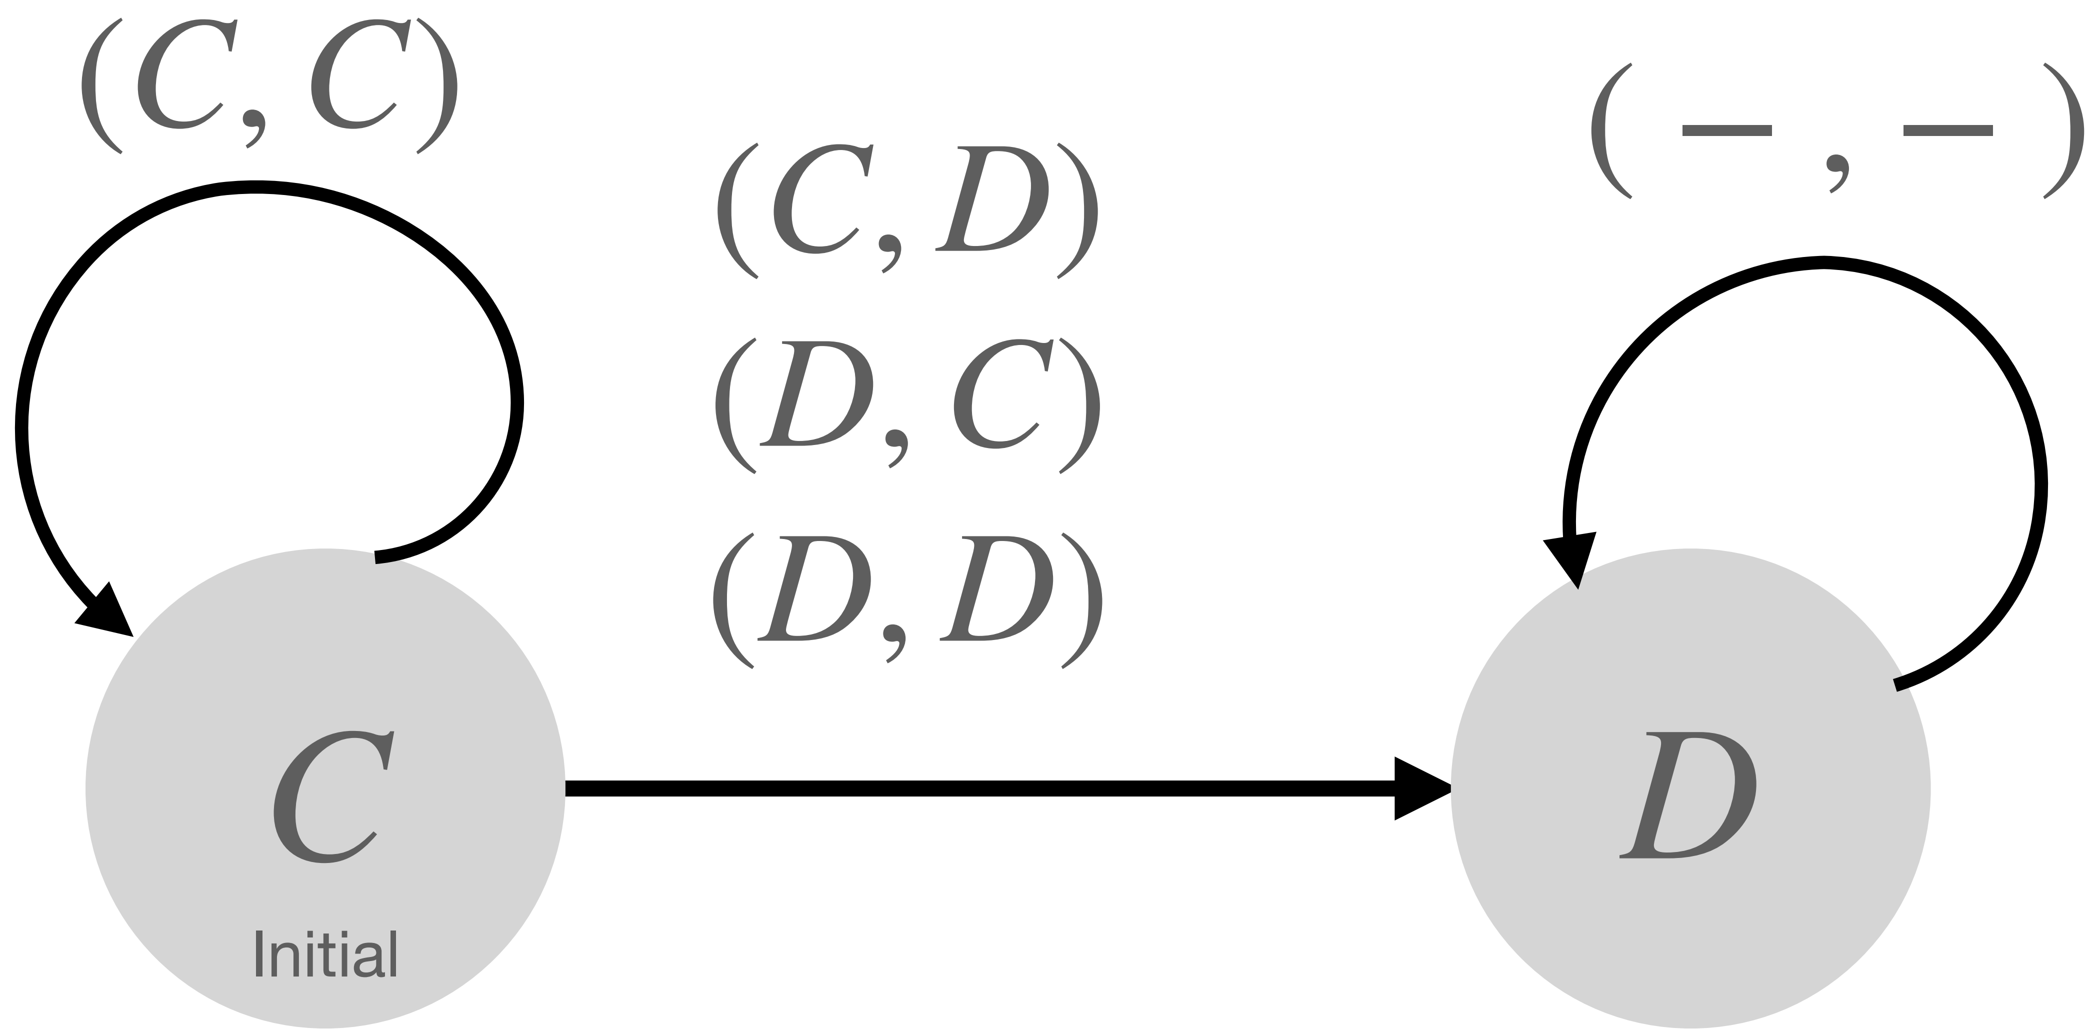

Grim Trigger

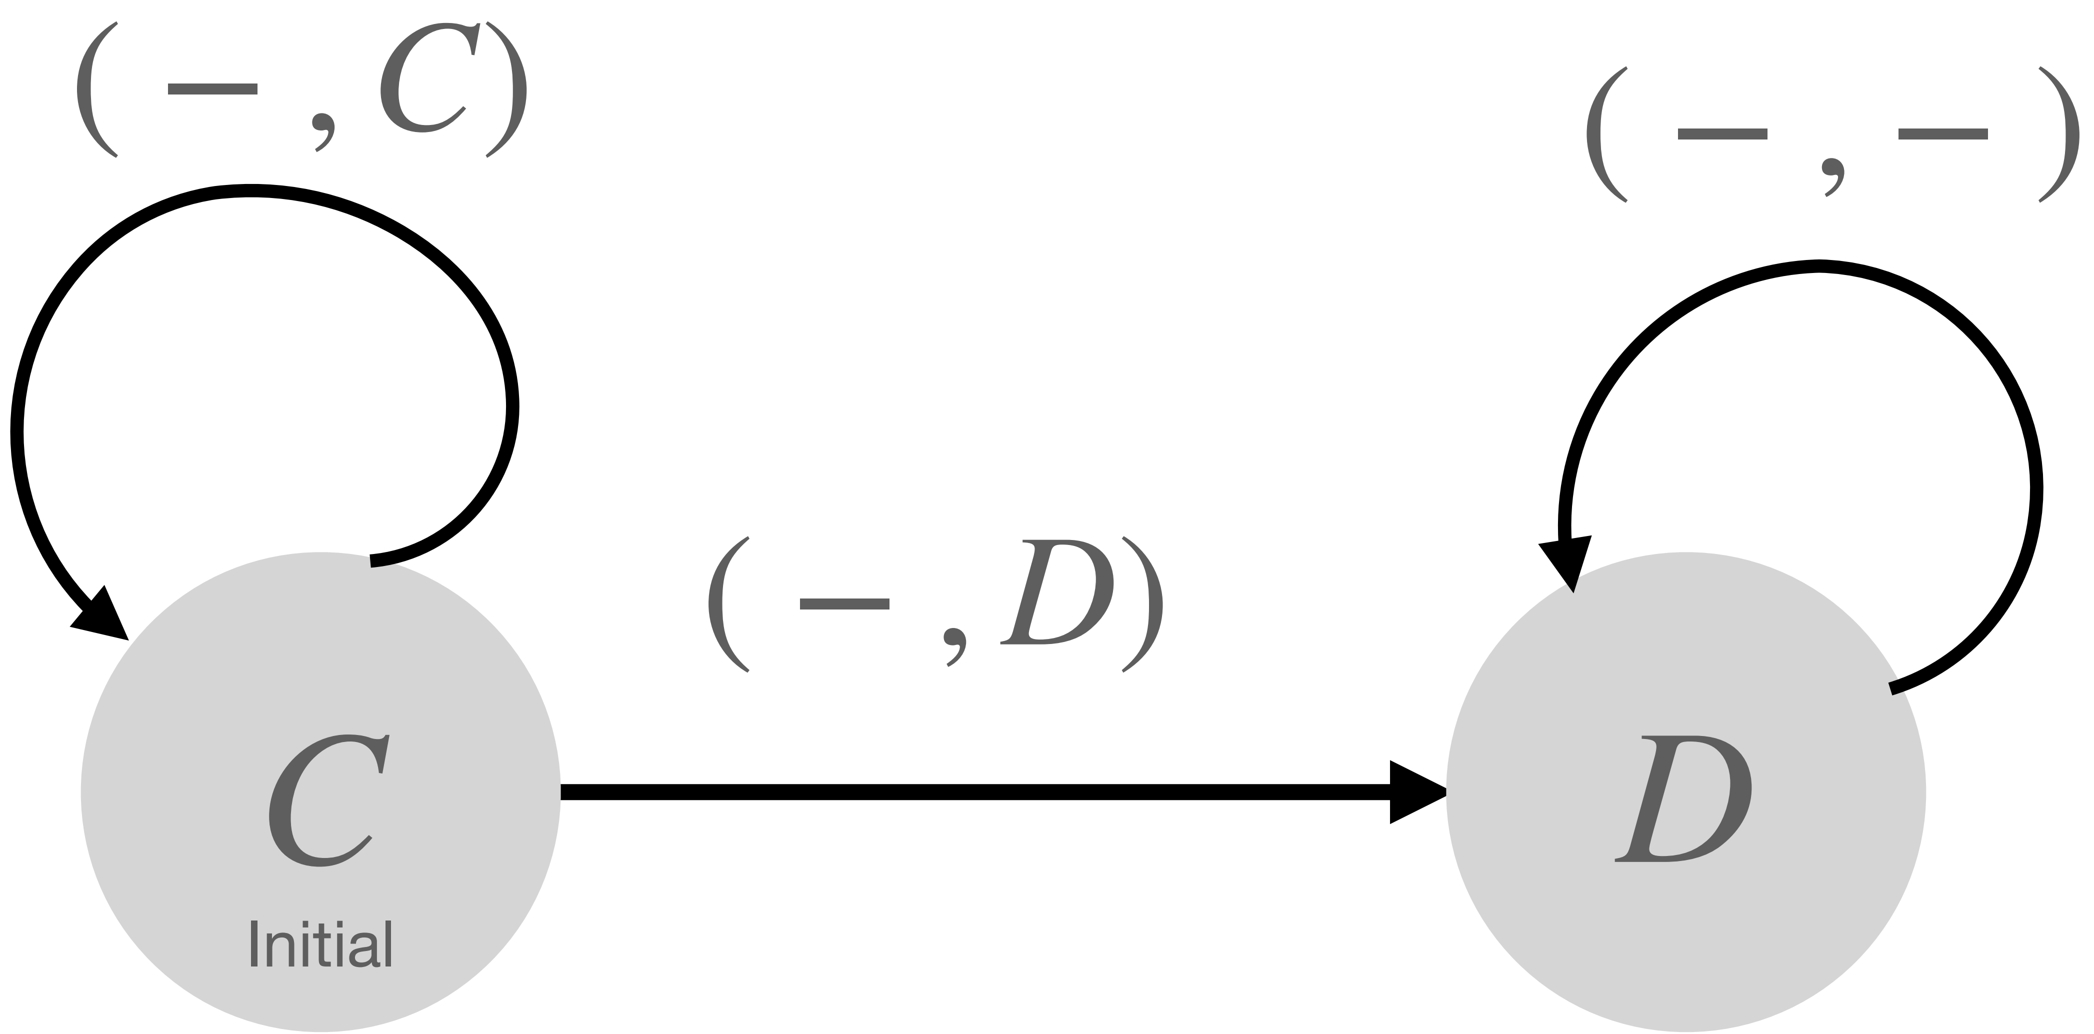

Modified Grim Trigger

Single Deviation Principle

- Single-deviation principle holds in an infinitely repeated game with discounting
- **Theorem**. A strategy profile is a subgame-perfect equilibrium in an infinitely-repeated game with discounting if, and only if, there is no useful single deviation.
- We will use this result without proof
- To reason about single deviation, its is useful to draw out the "reduced" decision tree for each strategy
  - (Game tree on board for Grim trigger)
  - The actions $(C, D), (D, C), (D, D)$ all have the same outcome and can be "reduced" for analysis

	C	D
C	$2, 2$	$-1, 3$
D	$3, -1$	$0, 0$

Grim in SPE

- **Lemma**. Grim Trigger strategy is a symmetric subgame-perfect Nash for all $\delta > 1/2$.
- **Proof**. We only need to consider two types of histories at stage $i \geq 1$:
- **Case 1: Cooperation history** at stage i. Histories in which D has never been played by any player in stages $i-1$
- **Case 2: Defection history** at stage i. Histories where at least one player has played D in stages $i-1$.
- For case 1, suppose both players continue with (C, C) prescribed by Grim trigger, then their payoff is
  - $\geq 2 + 2\delta$

	C	D
C	2, 2	−1, 3
D	3, − 1	0, 0

Grim in SPE

- **Lemma**. Grim Trigger strategy is a symmetric subgame-perfect Nash for all $\delta > 1/2$.
- **Proof**. We only need to consider two types of histories at stage $i \geq 1$:
- **Case 1: Cooperation history** at stage i. Histories in which D has never been played by any player in stages $i - 1$
- For case 1, suppose both players continue with (C, C) prescribed by Grim trigger, then their payoff is
  - $\geq 2 + 2\delta$
- If either player deviates to D in stage i, their payoff is
  - At most $3 + 0$ (for all future stages)
- Thus for $\delta > 1/2$, this deviation is not beneficial for Case 1

	C	D
C	$2,\ 2$	$-1,\ 3$
D	$3,\ -1$	$0,\ 0$

Grim in SPE

- **Lemma**. Grim Trigger strategy is a symmetric subgame-perfect Nash for all $\delta > 1/2$.
- **Proof**. We only need to consider two types of histories at stage $i \geq 1$:
- **Case 2: Defection history** at stage i. Histories where at least one player has played D in stages $i - 1$.
- If both players follow grim trigger strategy (D, D) in stage i
  - Their payoff is 0
- If either player deviates to C, their payoff becomes
  - $-1 + 0$
- Thus single-deviation is not useful for this case as well

	C	D
C	2, 2	−1, 3
D	3, −1	0, 0

Modified Grim is not SPE

	C	D
C	2, 2	–1, 3
D	3, – 1	0, 0

- **Lemma**. Modified Grim Trigger strategy is not a symmetric subgame-perfect equilibrium for any $\delta \in (0,1)$
- **Proof**. The difference now is that the action profiles (C, D) and (D, C) diverge in the game tree
- To show this is not a subgame-perfect equilibrium we look at the root of the tree (stage 0)
- Consider the subgame following outcome (C, D)
- Suppose player 1 adheres to modified grim and plays D in next stage
- **Claim**. It is not optimal for player 2 to play C according to modified grim

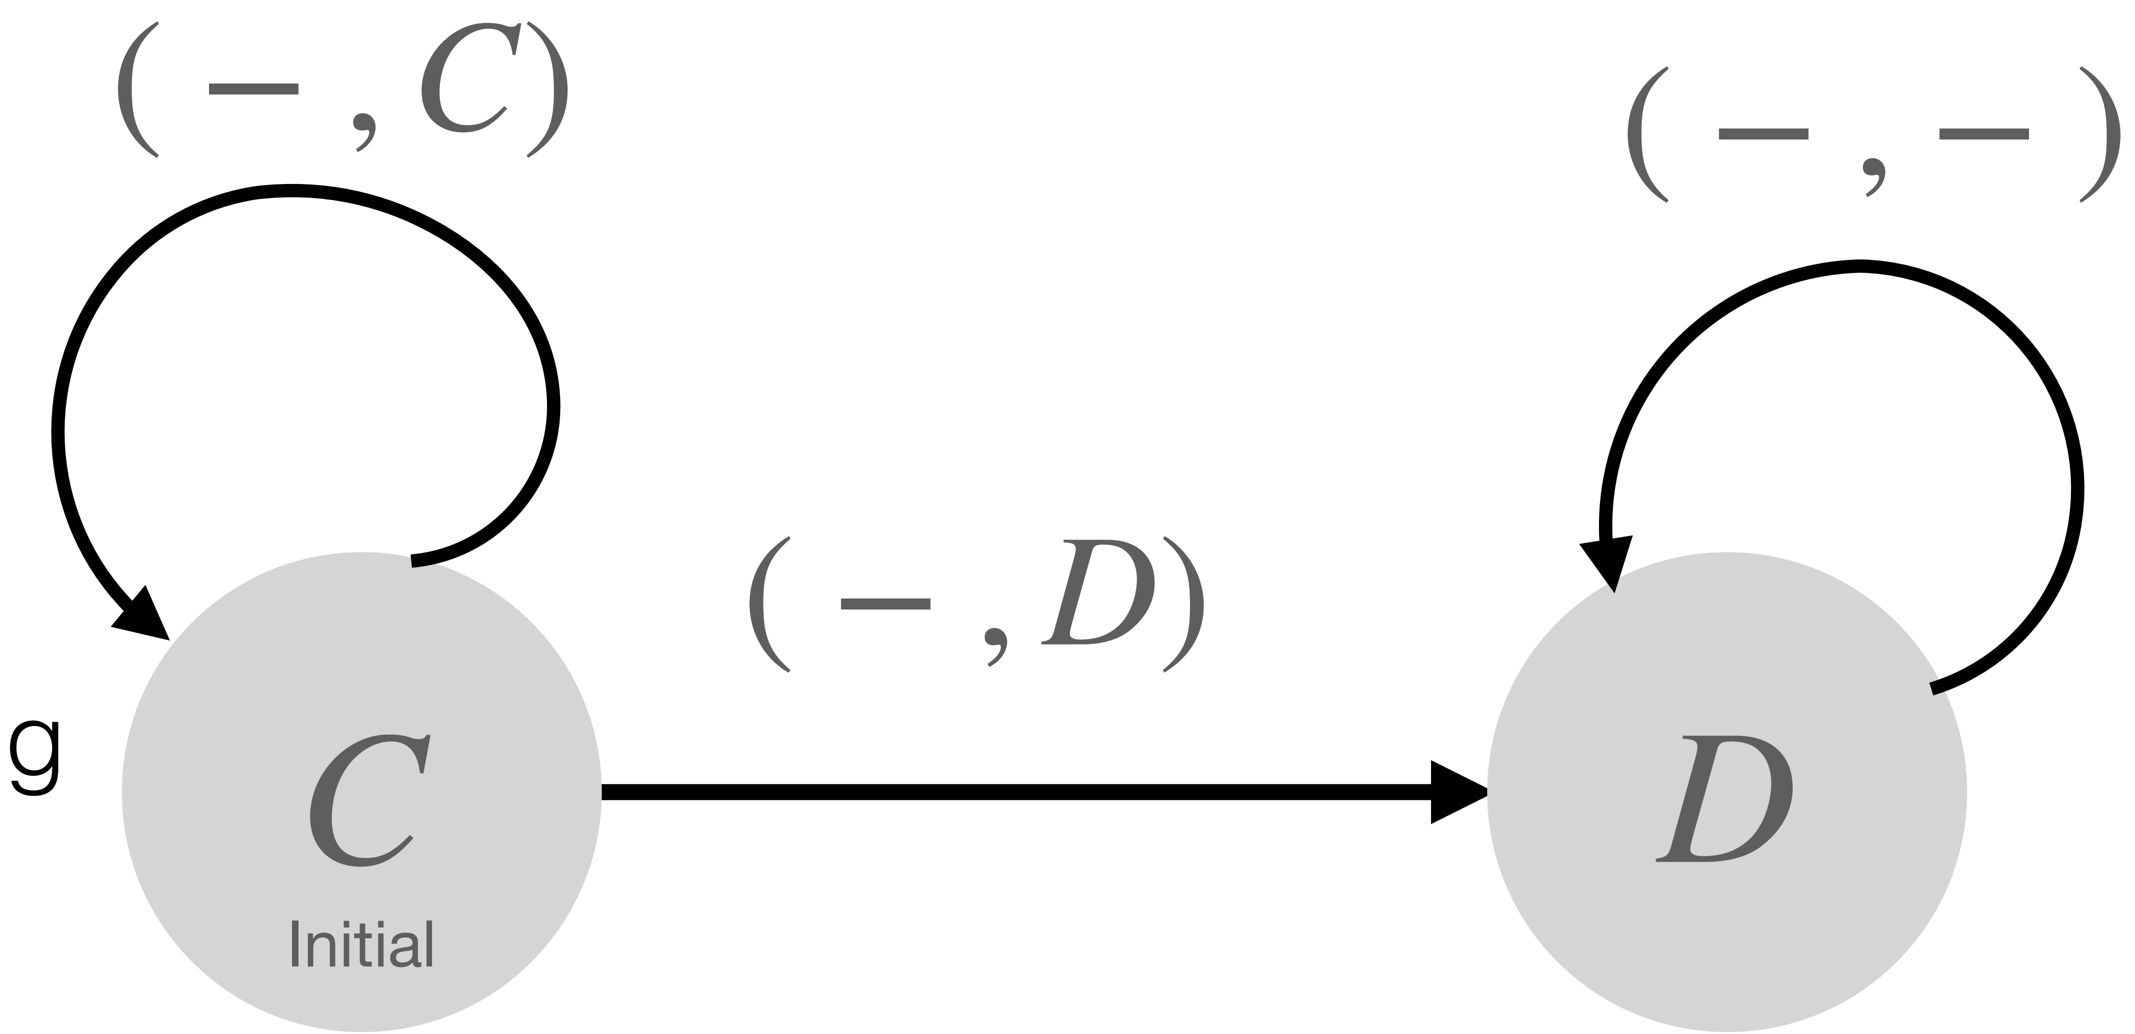

Modified Grim Trigger

Modified Grim is not SPE

	C	D
C	2, 2	–1, 3
D	3, – 1	0, 0

- **Lemma**. Modified Grim Trigger strategy is not a symmetric subgame-perfect equilibrium for any $\delta \in (0,1)$
- **Proof**. Consider the subgame following outcome (C, D)
- Suppose player 1 adheres to modified grim and plays D in next stage
- **Claim**. It is not optimal for player 2 to play C according to modified grim
- If player 2 adheres to modified grim, then outcome in next stage is (D, C) followed by (D, D) forever
  - Payoff -1
- If player 2 deviates to (D, D) gets at least 0

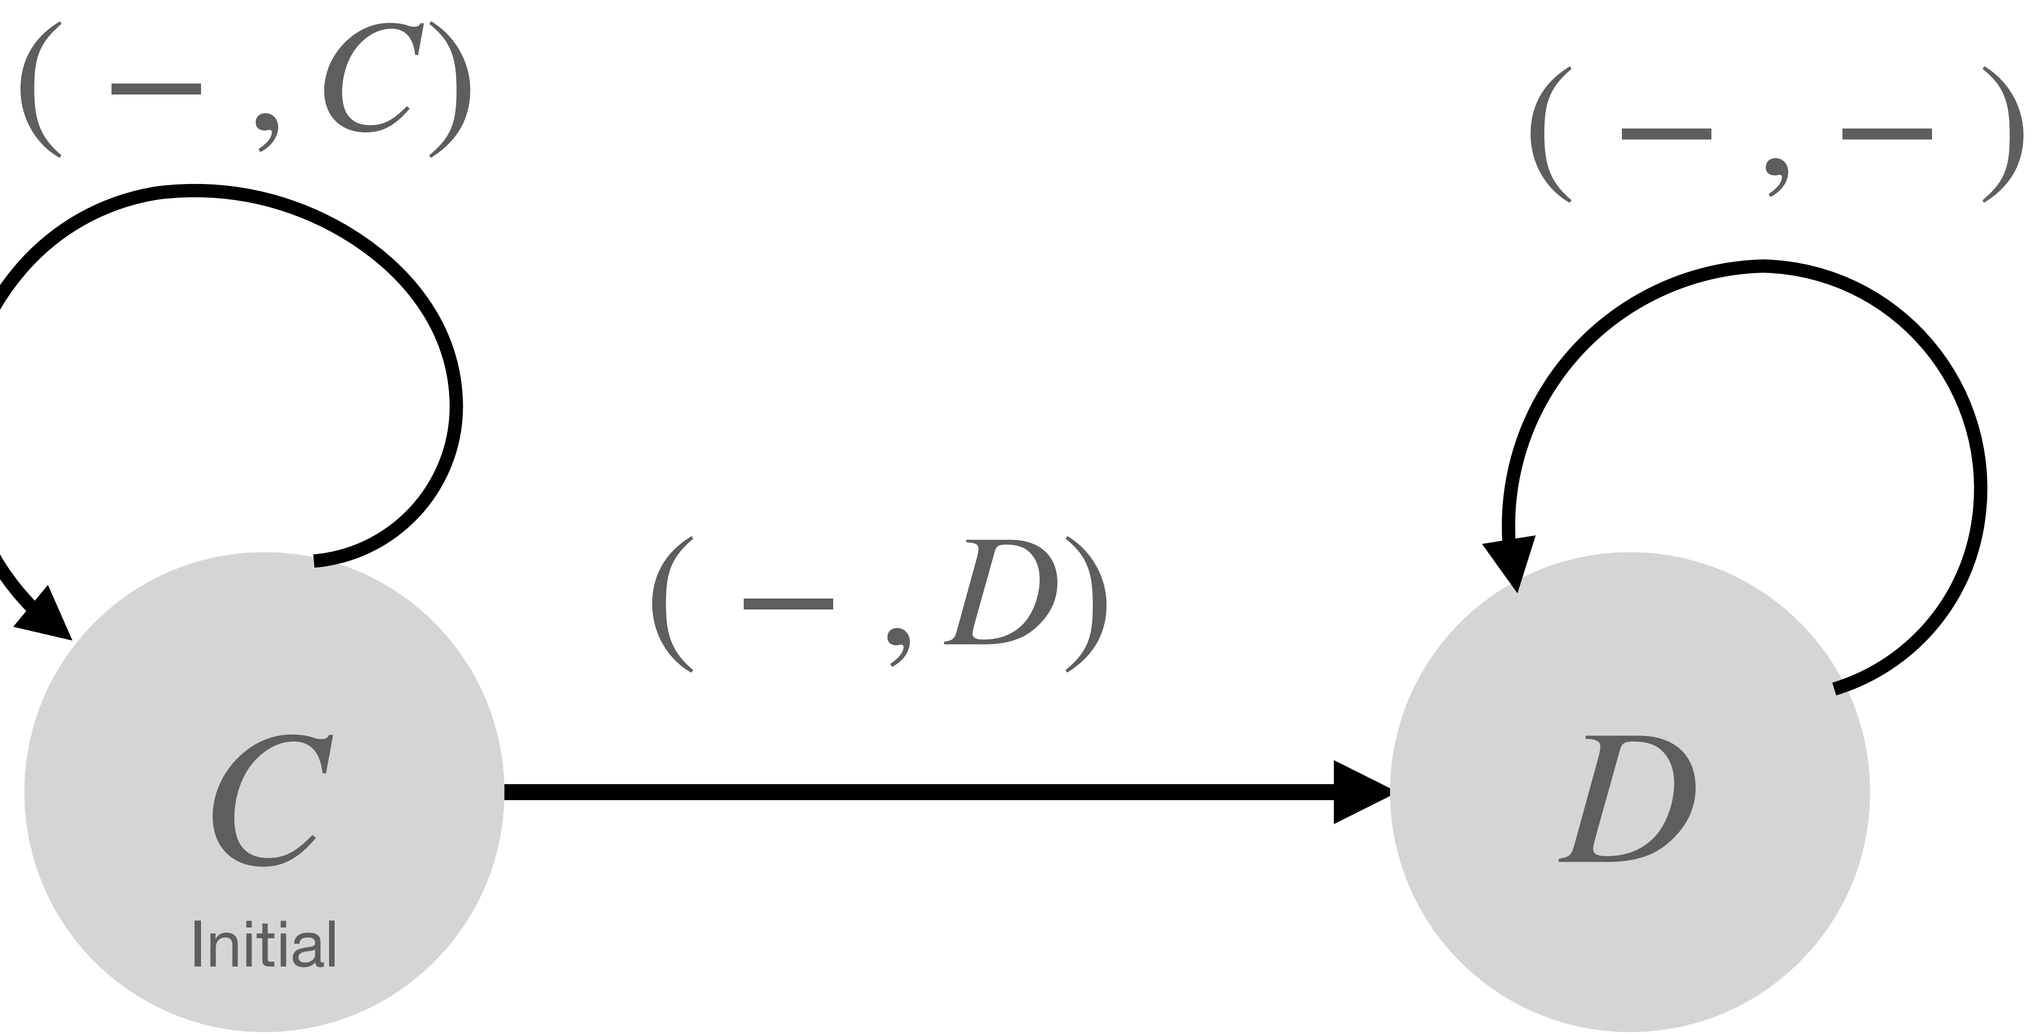

Modified Grim Trigger

Takeaway: Nash vs SPE

- Even though "on the equilibrium path" no player should deviate to D is opponent has not deviated to D in the past
- SPE requires that a threat be credible even on histories that may never be played in equilibrium
- SPE is fragile wrt slight changes in how strategy is defined in repeated Prisoner's dilemma

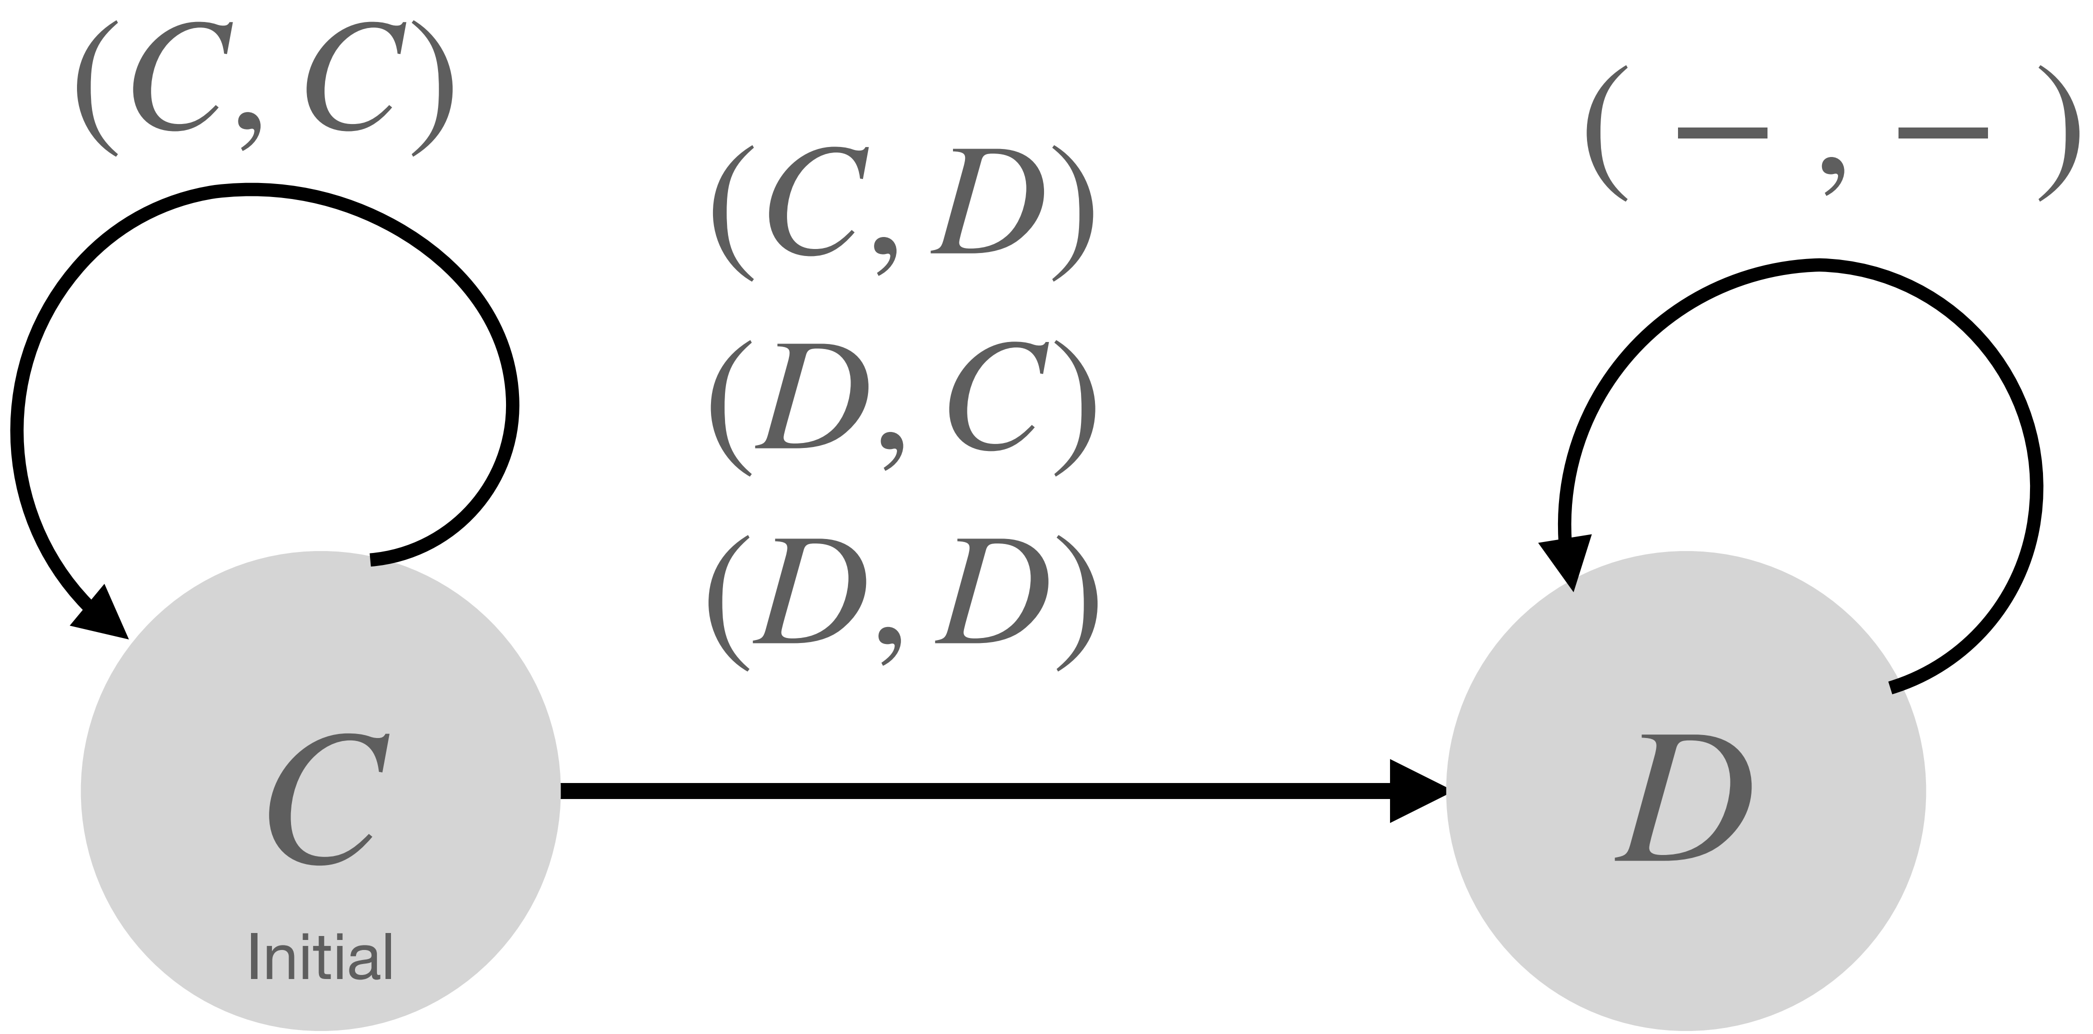

Grim Trigger

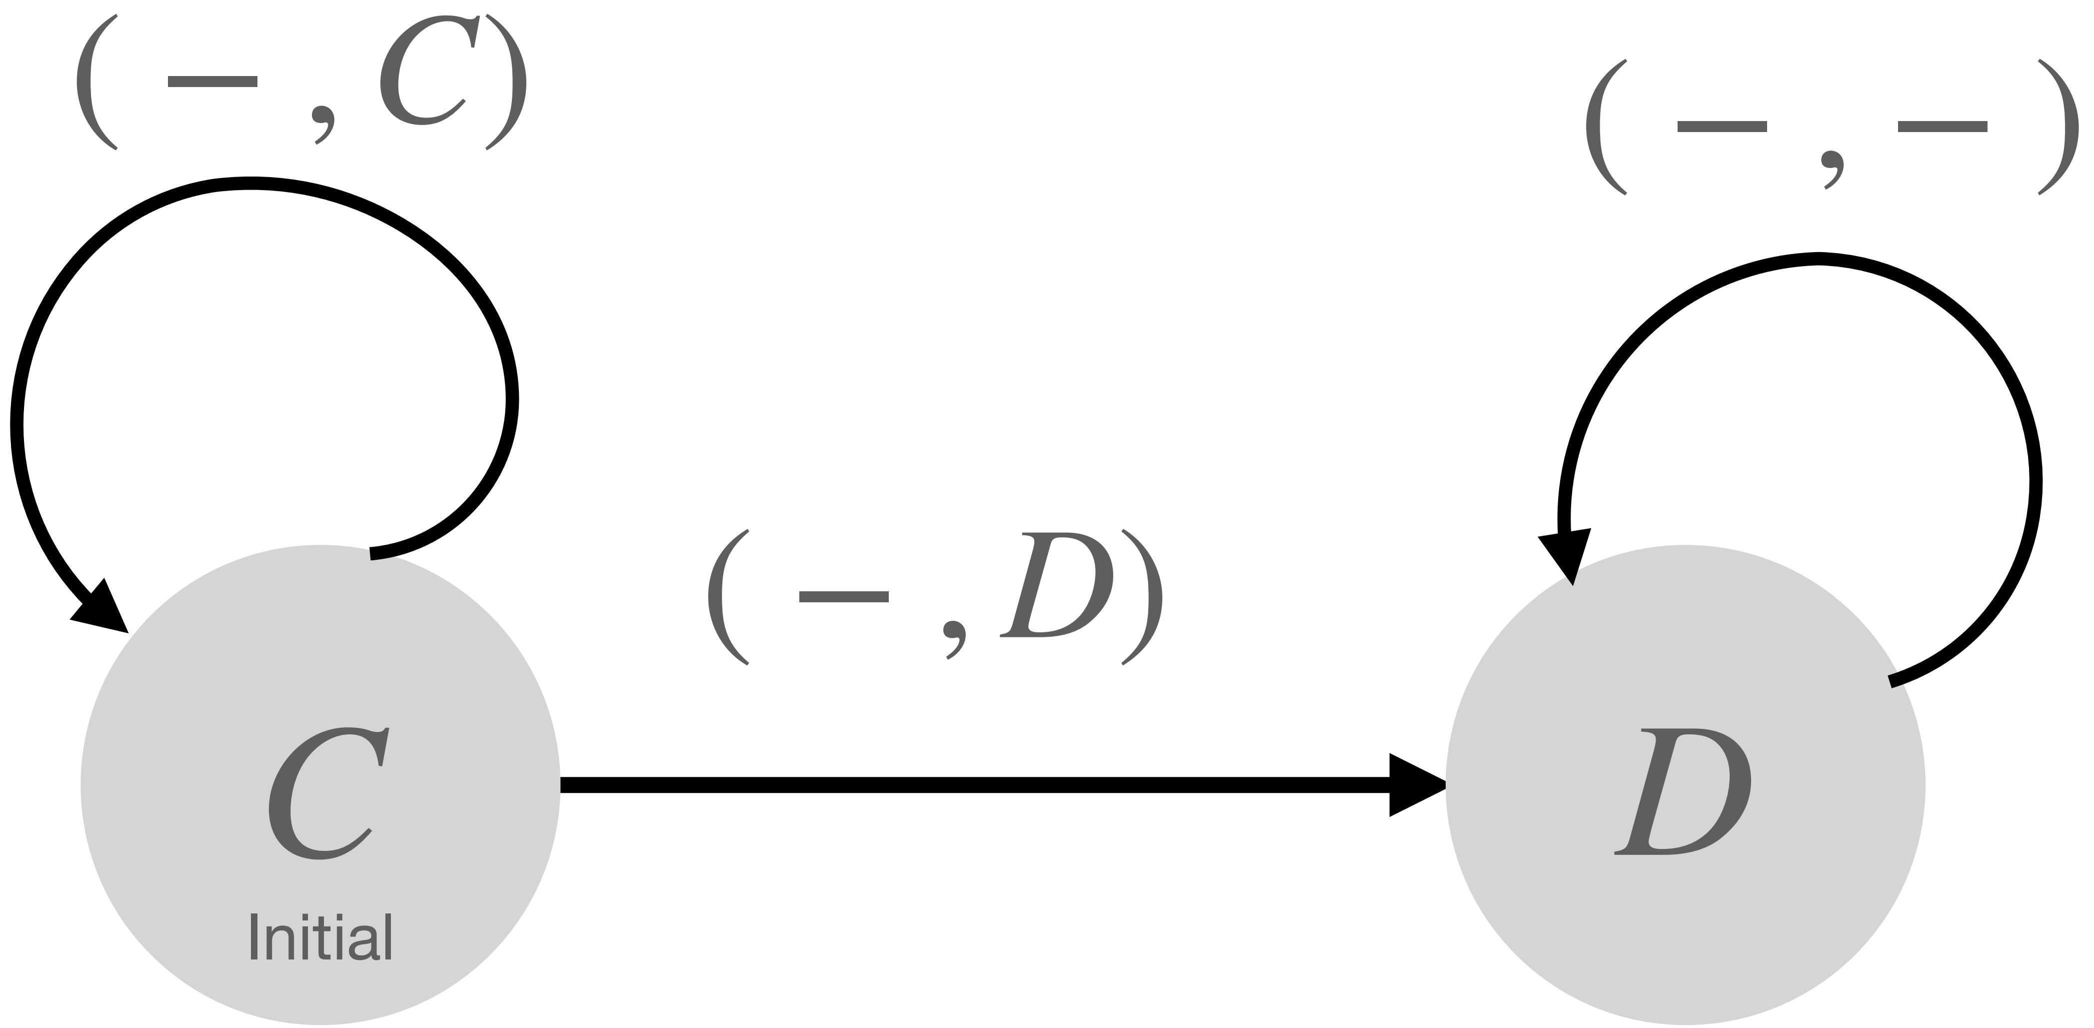

Modified Grim Trigger

Tit-for-Tat

- Grim trigger strategy is pretty extreme (holding a grudge in perpetuity for a single defection)
- Even though Grim is at equilibrium with itself, a more robust strategy (for asymmetric environments) should involve some forgiveness
- **Tit-for-tat strategy**:
  - Start by cooperating
  - Do in stage i whatever the opponent does in stage $i-1$
- Thus, tit-for-tat starts optimistically, punishes immediately and forgives quickly
- Turns out to be a good strategy in repeated prisoner's dilemma

Tit-for-Tat

- We show that TfT is a symmetric Nash equilibrium
- Suppose Beth is playing TfT, and we consider if Aamir can benefit by deviating in some stage i
  - Notice that stage i's payoff only depends on stage $i-1$
  - Deviating to D will lead to 3 now but at most 0 in next round
  - Cooperating now will lead to 2 now and 2δ in the next round
- Thus for $\delta > 1/2$, Aamir has no incentive to deviate
- In **Homework 8**, you are asked to show that
  - TfT is not a SPE for any $\delta < 1$
  - Modifying TfT to become a conditional cooperator (always return to C after D) is a SPE for sufficiently large δ

Repeated Prisoner’s Dilemma (Empirically)

Alexrod IPD Tournaments

- In ~1980, Robert Axelrod invited colleagues to enter into a tournament for computer programs playing the repeated Prisoner's dilemma
- There were **15 contestants** and each program played the other 14 in a repeated PD game with 200 stages (so, a round-robin tournament)
- The payoff of a program was averaged over all 200 stages of all 14 matches
- The winning strategy (submitted by Anatol Rapoport) was **Tit-for-Tat**!
- Tit-for-Tat was the shortest entry and many other programs were (trying to be) sophisticated
- What makes this even more surprising
  - TfT cannot win a head-to-head match with any opponent! **(Homework 8)**

Alexrod IPD Tournaments

- Axelrod circulated the results of the first tournament and solicited entries for a second tournament with the same rules
- This time there were **62 entries**
- Rapoport resubmitted **Tit-for-Tat**, completely unchanged, and won again!
    - This happened even though other programs were explicitly tailored to exploit Tit-for-Tat
    - Turns out they imploded against each other!
- **Python Axelrod library** (https://axelrod.readthedocs.io/en/stable/) has implementations of all the entries in the tournament
- Their Github has extensive documentation and is useful for potential project on the topic of Iterated Prisoner's Dilemma

Evolution and IPD

The Evolution of Cooperation

Robert Axelrod and William D. Hamilton

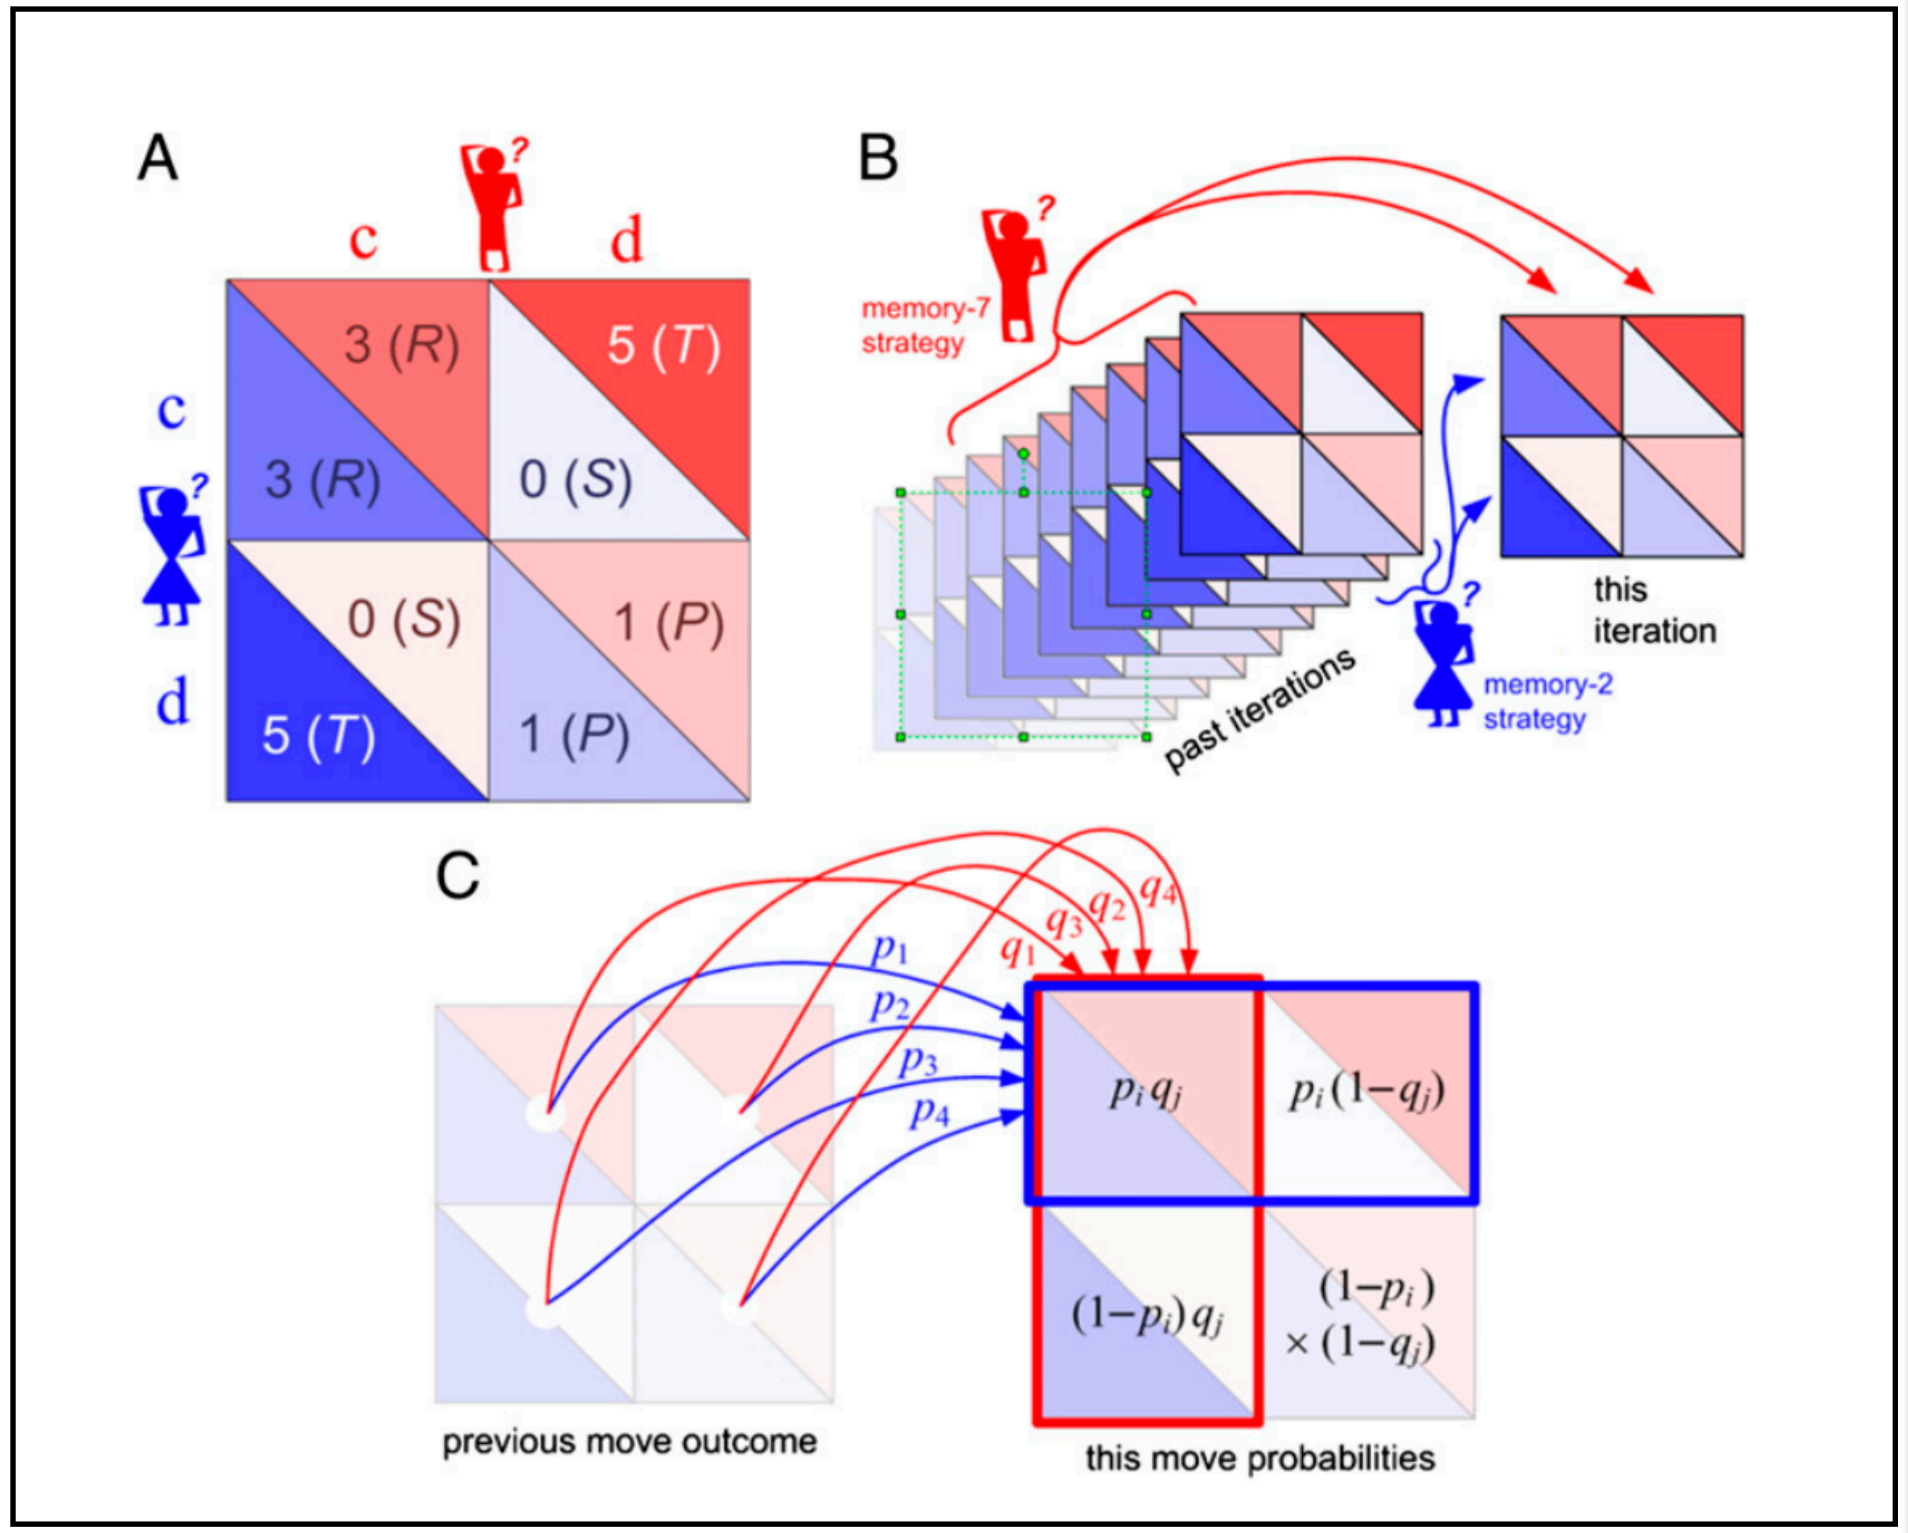

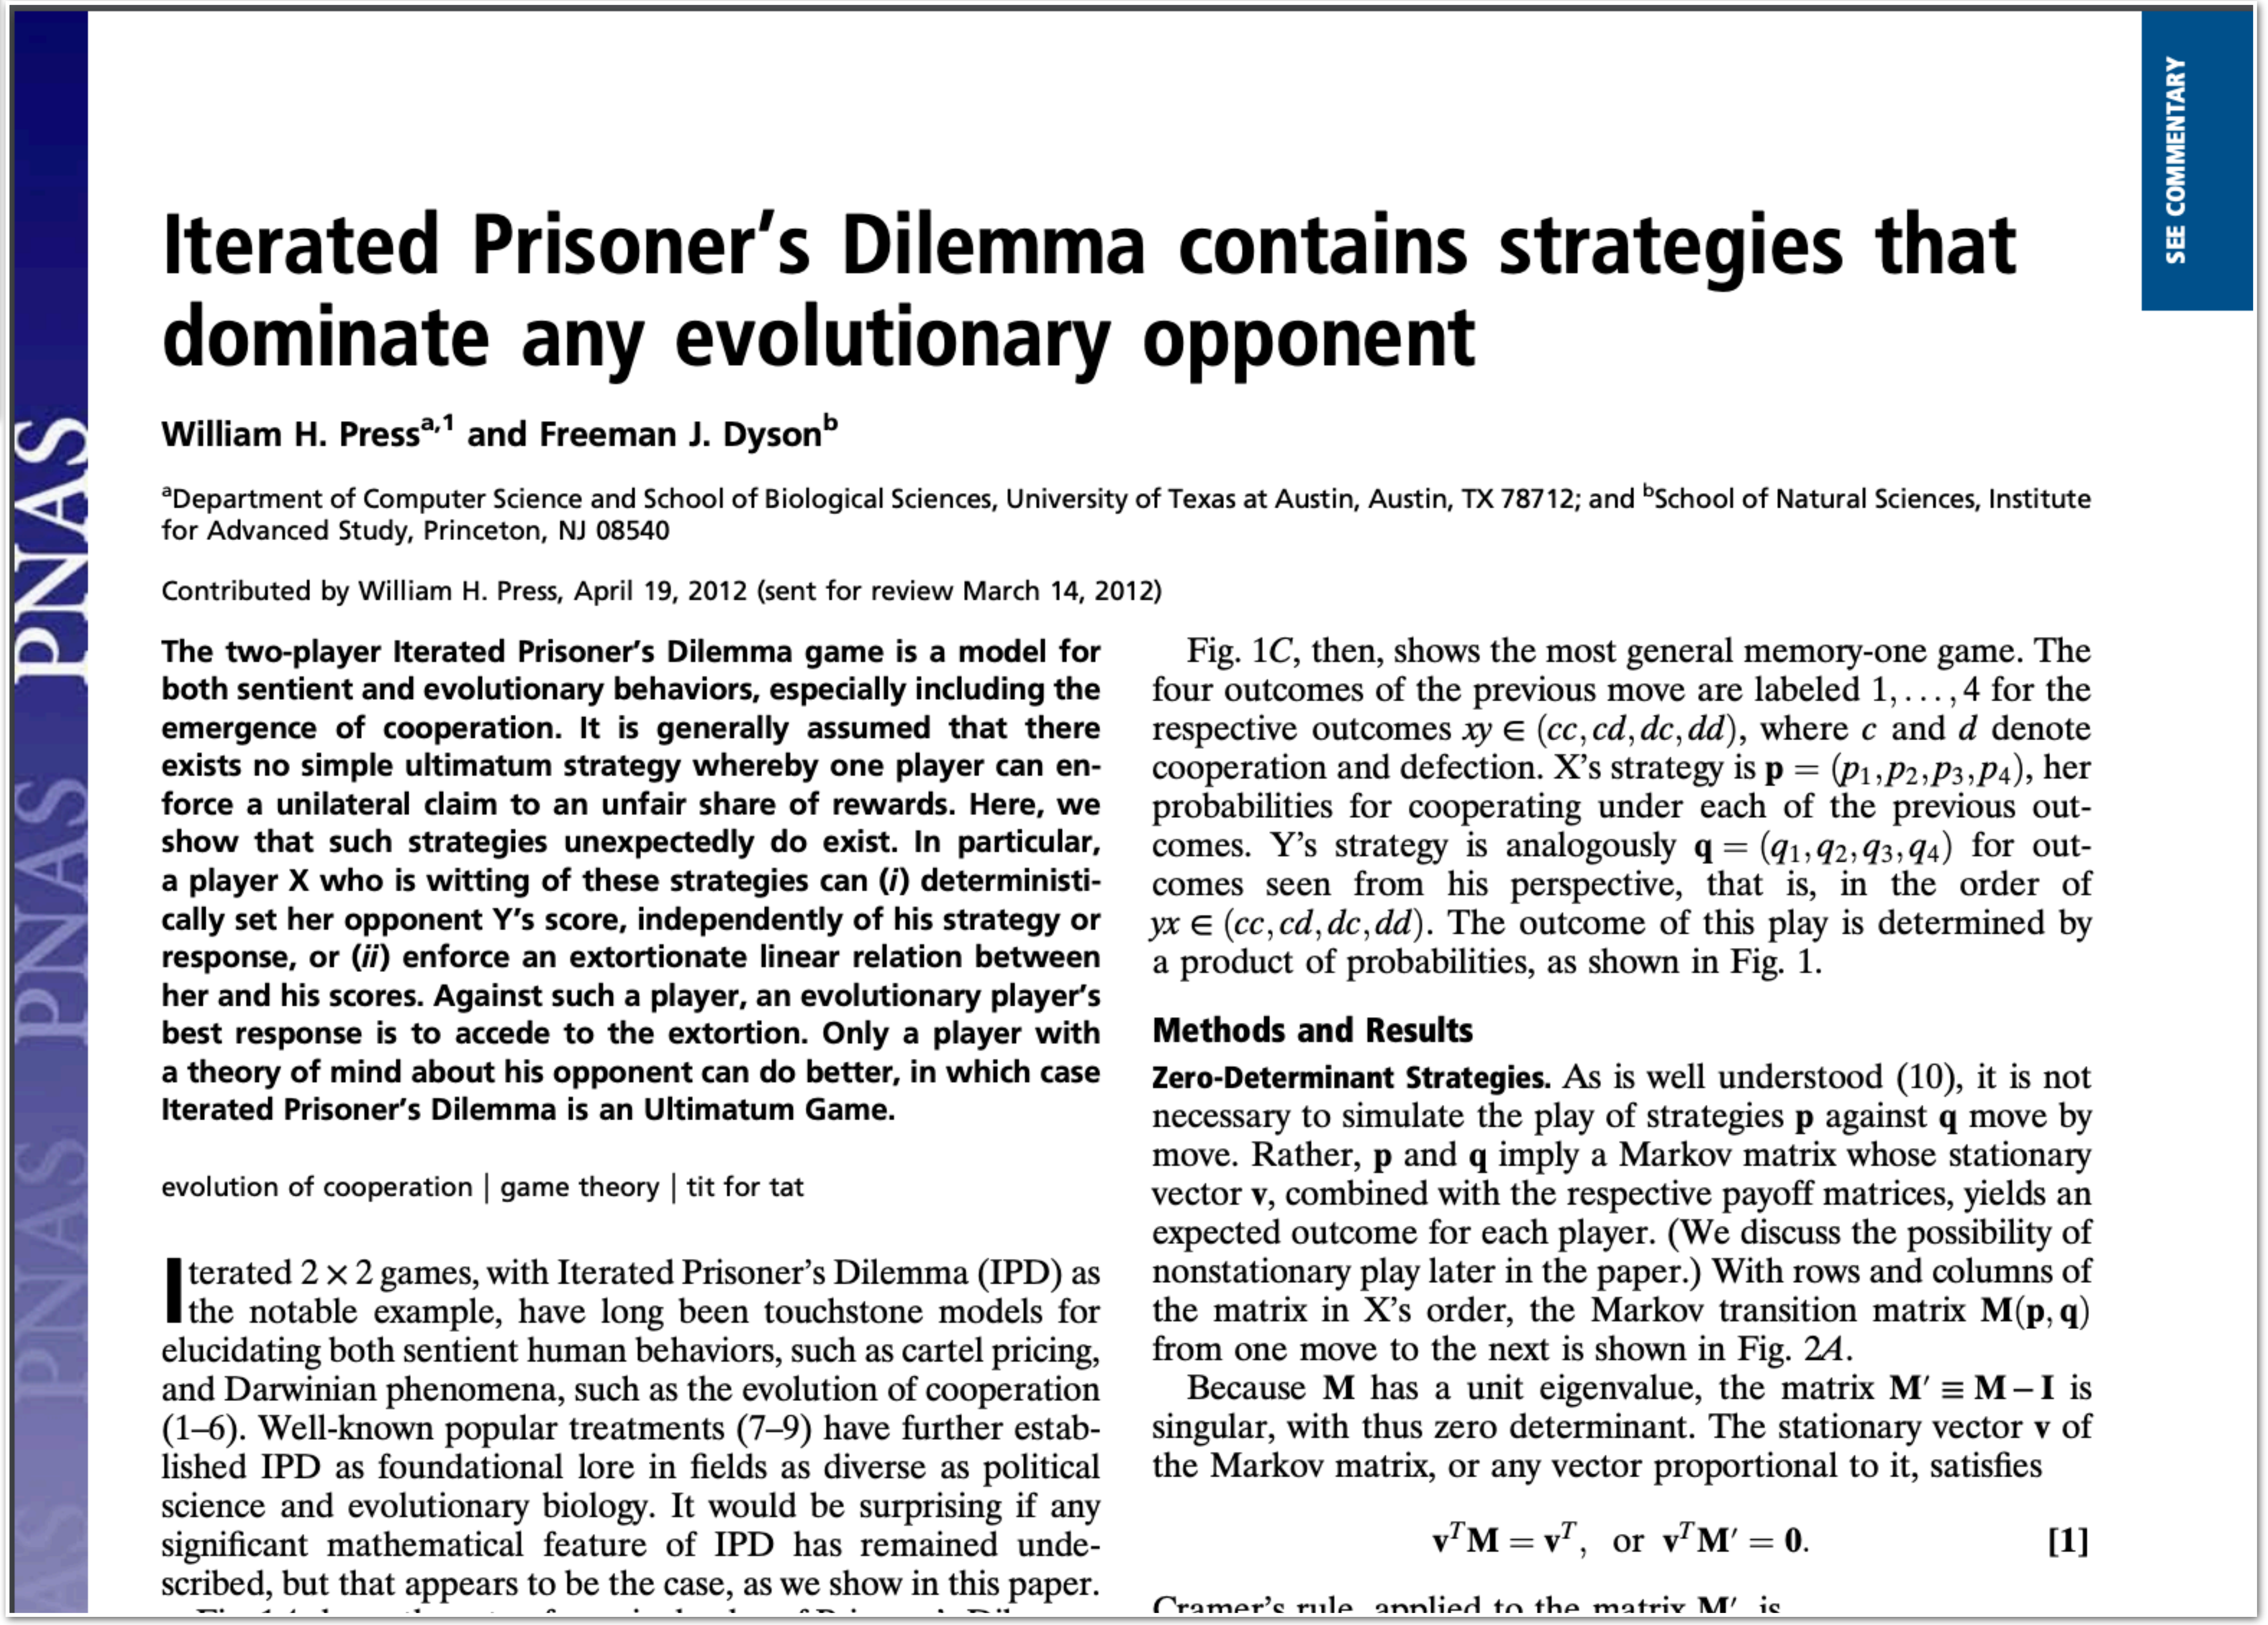

Iterated Prisoner's Dilemma contains strategies that dominate any evolutionary opponent

William H. Press[a,1] and Freeman J. Dyson[b]

[a]Department of Computer Science and School of Biological Sciences, University of Texas at Austin, Austin, TX 78712; and [b]School of Natural Sciences, Institute for Advanced Study, Princeton, NJ 08540

Contributed by William H. Press, April 19, 2012 (sent for review March 14, 2012)

The two-player Iterated Prisoner's Dilemma game is a model for both sentient and evolutionary behaviors, especially including the emergence of cooperation. It is generally assumed that there exists no simple ultimatum strategy whereby one player can enforce a unilateral claim to an unfair share of rewards. Here, we show that such strategies unexpectedly do exist. In particular, a player X who is witting of these strategies can (*i*) deterministically set her opponent Y's score, independently of his strategy or response, or (*ii*) enforce an extortionate linear relation between her and his scores. Against such a player, an evolutionary player's best response is to accede to the extortion. Only a player with a theory of mind about his opponent can do better, in which case Iterated Prisoner's Dilemma is an Ultimatum Game.

evolution of cooperation | game theory | tit for tat

Iterated 2×2 games, with Iterated Prisoner's Dilemma (IPD) as the notable example, have long been touchstone models for elucidating both sentient human behaviors, such as cartel pricing, and Darwinian phenomena, such as the evolution of cooperation (1–6). Well-known popular treatments (7–9) have further established IPD as foundational lore in fields as diverse as political science and evolutionary biology. It would be surprising if any significant mathematical feature of IPD has remained undescribed, but that appears to be the case, as we show in this paper.

Fig. 1*C*, then, shows the most general memory-one game. The four outcomes of the previous move are labeled $1, \ldots, 4$ for the respective outcomes $xy \in (cc, cd, dc, dd)$, where c and d denote cooperation and defection. X's strategy is $\mathbf{p} = (p_1, p_2, p_3, p_4)$, her probabilities for cooperating under each of the previous outcomes. Y's strategy is analogously $\mathbf{q} = (q_1, q_2, q_3, q_4)$ for outcomes seen from his perspective, that is, in the order of $yx \in (cc, cd, dc, dd)$. The outcome of this play is determined by a product of probabilities, as shown in Fig. 1.

Methods and Results

Zero-Determinant Strategies. As is well understood (10), it is not necessary to simulate the play of strategies $\mathbf{p}$ against $\mathbf{q}$ move by move. Rather, $\mathbf{p}$ and $\mathbf{q}$ imply a Markov matrix whose stationary vector $\mathbf{v}$, combined with the respective payoff matrices, yields an expected outcome for each player. (We discuss the possibility of nonstationary play later in the paper.) With rows and columns of the matrix in X's order, the Markov transition matrix $\mathbf{M}(\mathbf{p}, \mathbf{q})$ from one move to the next is shown in Fig. 2*A*.

Because $\mathbf{M}$ has a unit eigenvalue, the matrix $\mathbf{M}' \equiv \mathbf{M} - \mathbf{I}$ is singular, with thus zero determinant. The stationary vector $\mathbf{v}$ of the Markov matrix, or any vector proportional to it, satisfies

$$\mathbf{v}^T\mathbf{M} = \mathbf{v}^T, \quad \text{or} \quad \mathbf{v}^T\mathbf{M}' = \mathbf{0}. \qquad [1]$$

Evolution of Trust

- Simulation that compares different strategies in a visually appealing format: https://ncase.me/trust/
  - Adds the notion of error

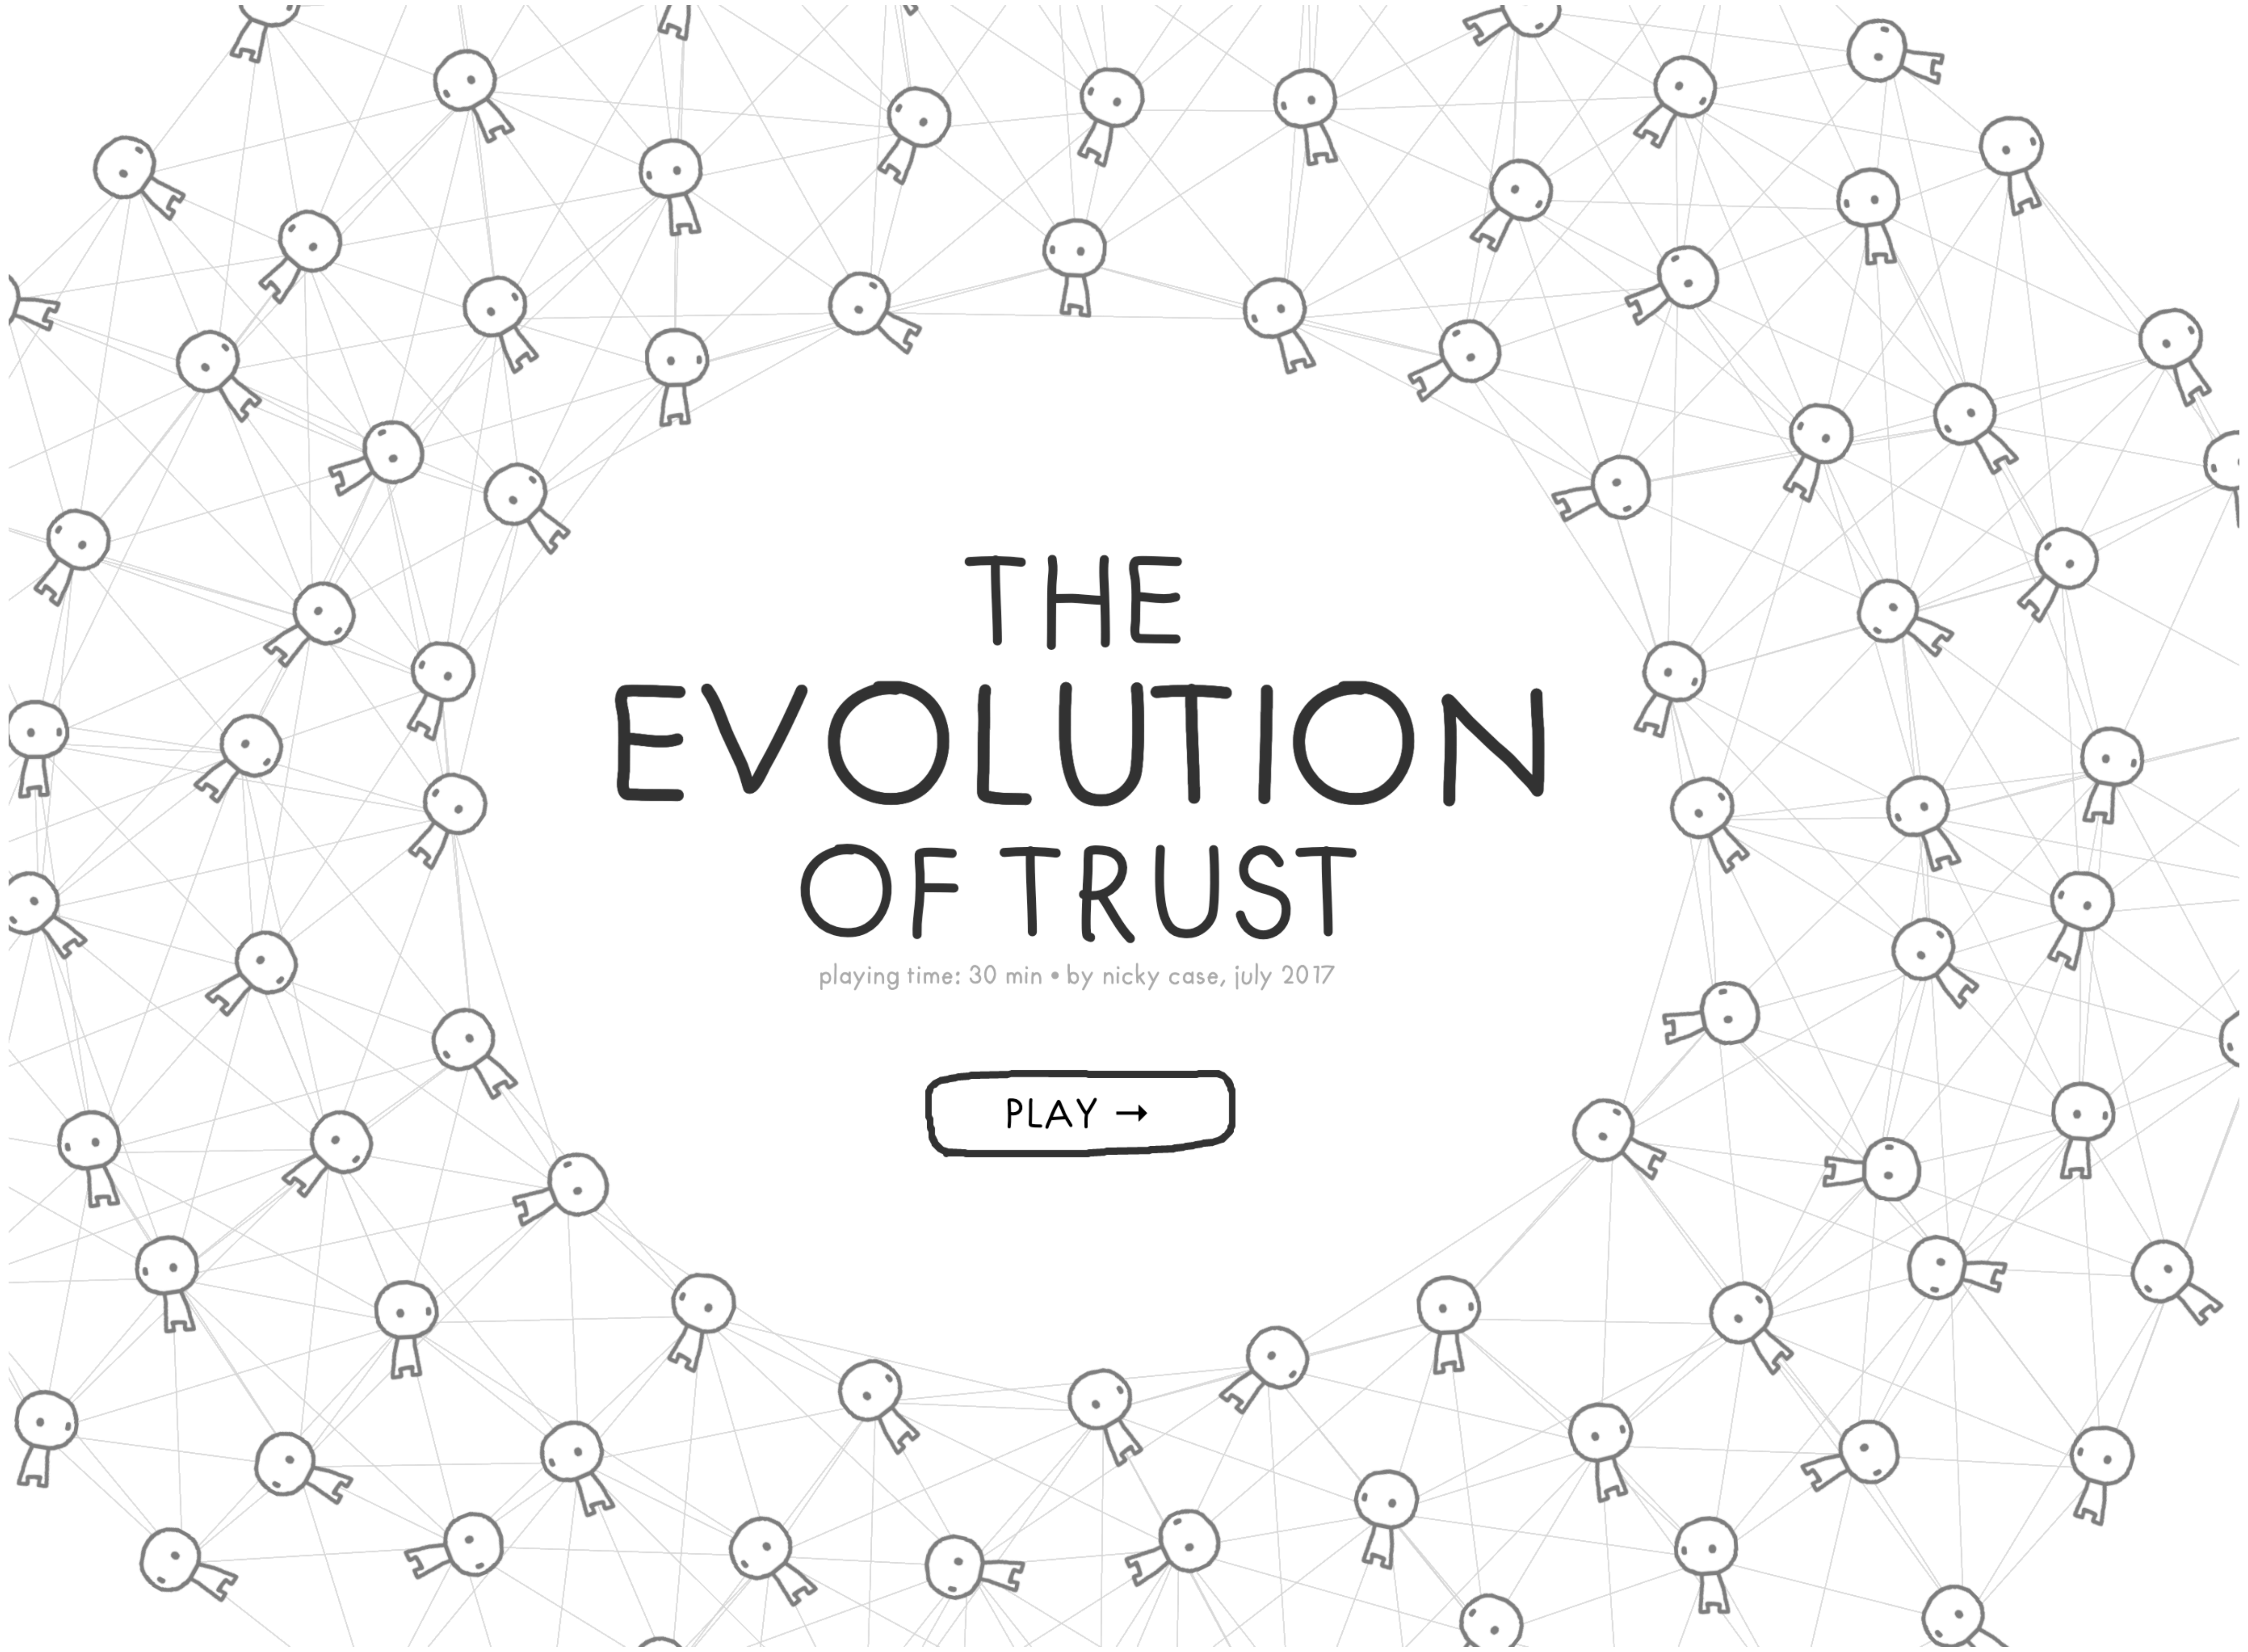

Project Overview

CS357 Project

- Goals of the final project:
  - Dig deeper into a topic in AGT
  - Analyze incentives inherent in a system or class of algorithms by applying the frameworks developed in class
  - Read and understand **research papers** in the field
- Ideal: groups of 2
- Scope and timeline
  - 3 week project: ~30 hours of effort
  - Project proposal April 22 (Fri)
  - Final project report due May 19 (Thurs)

Deliverables

- Proposal (google form, 150 word proposal):   due Fri April 22, 5 pm
  - Identify topic and partner
  - Describe high level goals and plan for next 3 weeks
  - Identify related research papers related to topic
- 1 page LaTeX report: due Fri April 29
- 2 pages LaTeX report including research background: due Thur May 5
- Project presentations: last week of classes: due May 12 ~5-10 slides
- Final project report (at least 7 pages) due May 19
- **Check ins:**   Meet with me at least twice for 20 mins

Suggested Project

- Each project must explore the **role of incentives and strategic behavior** in a market/ system or a class of algorithms
- Suggested project: Understanding Strategic Behavior in BitTorrent
  - BitTorrent simulator will be provided
  - Implement BitTorrent reference client
  - Read related research papers and implement strategic clients
    - BitTyrant: Do incentives build robustness in BitTorrent? by Piatek et al.
    - BitTorrent is an Auction: Analyzing and Improving BitTorrent's Incentives by Levin et al.
- Open ended: design your own client, perform advanced analysis to compare existing strategic behavior

Many Domains

- Each project must explore the **role of incentives and strategic behavior** in a market/ system or a class of algorithms
- Role of strategic behavior in markets without money
    - One-sided matching markets (Boston choice, kidney exchange)
    - Two-sided matching markets (Gale Shapley algorithm)
    - Voting systems and rank aggregation, fair division
- Role of strategic behavior in markets with money or transfers
    - Auctions, market equilibria, etc
- Game theory projects: evolutionary, repeated games, and sequential games etc
- Examples from F20: Will post two project reports from Fall 2020 on GLOW

Guide to Choosing a Topic

- Pick a topic you will enjoy working on (something fun!)
- Pick something that you will learn a lot from (something useful!)
- Pick something that is relevant to the course (something relevant!)
- Keep the scope in mind, make sure you have something that you can make progress on in ~3 weeks
  - Something not too ambitious but not too trivial
- Try to find CS/Econ papers related to the topic you have in mind
  - Topic must be technically interested
  - Related to concepts covered in class

Topic Ideas and Papers

- Will post a more detailed summary of suggested project topic
- Will post a bunch of **recent relevant research papers**
- Some ideas of topics you can explore
- If you have an idea and want to know if it is viable
    - Come talk to me!
    - I can suggest research papers related to topic
    - We can refine the idea together
- Projects can be mostly implementation or mostly theoretical or between
    - BitTorrent project is very implementation heavy

Theory vs Implementation

- Scope of implementation project similar to suggested project
- No matter theory or implementation
    - Project must involve reading and understanding **at least two research papers**
- Even when implementing, important to understand the theory and concepts
- Purely theory: plan on going above and beyond just surveying existing literature:
    - New insights/observations or conjectures
    - Filling in important details, examples or cases that may have been overlooked
    - Simulating strategic behavior can help analyze situations where theoretical analysis may be tricky
- Always try to identify what your contribution in the project is going to be

Paper Reading Advice

- Use the three-pass approach suggested by Keshav
- First pass:
  - Read the abstract, introduction, conclusion and boldface
  - What information you can glean from this:
    - What are the paper's main contributions?
    - Is the paper well written?
    - Main keywords that the paper is about
- At this point, you should be able to identify the "closest" related work to the current paper: which can often be a useful resource

Paper Reading Advice

- Second pass:
  - Read the paper more thoroughly (but ignore proofs)
  - Pay special attention to theorem/lemma statements or graphs
    - What assumptions are the authors making?
    - Do the definitions used make sense?
      - Create examples and identify the defined notions?
- Start writing down questions about what you don't understand
  - It may be the answers are hidden in the paper but not obvious
- Take lots of notes! Use your own examples to explain the work

Paper Reading Advice

- Third pass: Virtually reimplement the paper
    - For implementation project this is literal; for theory projects, this means understand the concepts as well as if you coded it up
- At this stage, you should have your own "perspective" on the work
    - How would you do things differently if you wrote the paper?
    - Is there a notion the authors do not explain well, that you can supplement through your project?
    - How would you teach this paper to the class?
    - What are the most interesting and challenging aspects?
    - What new directions can this work take if you had more time?

Literature Survey

- For your technical writeup, you'd need to summarize the closely related literature to the topic you picked
- How should you approach it?
  - First, pay close attention to what the authors identify as the most related work to theirs
  - Follow "the trail of citations"
  - Find what's the state of the art is on a topic
  - Use google scholar!

Incentives in BitTorrent