(Sequential) Information Design

Data-Driven Decision Processes Bootcamp

Laura Doval

Columbia Business School

The mechanism design framework

Two lectures on what essentially are two special cases of **mechanism design**:

- Lecture #1: Information Design
- Lecture #2: Mechanism Design with Limited Commitment

Hopefully by the end of Lecture #2 that deep connections exist will be self-evident

For now, I will follow the (modern) Econ tradition of keeping them separate

Common Primitives:

- N agents, $i \in \{1, \ldots, N\}$,
- finite set of states of the world, Θ,
- common prior $\mu_0 \in \Delta(\Theta)$,
- Set $A \equiv A_0 \times A_1 \times \ldots A_N$ of alternatives.
- Payoffs $u_i : \Theta \times A \mapsto \mathbb{R}$.

I will call the tuple $G \equiv \{N, \Theta, \mu_0, A, (u_i)_{i \in N}\}$ the **base game**.

(Static) Mechanism Design:

- Agents have private information: T_i is the set of types of agent i and

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

describes the information player i has about θ and the types of other players.

(Static) Mechanism Design:

- Agents have private information: T_i is the set of types of agent i and

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

describes the information player i has about θ and the types of other players.

- Payoffs only depend on A_0

(Static) Mechanism Design:

- Agents have private information: T_i is the set of types of agent i and

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

describes the information player i has about θ and the types of other players.
- Payoffs only depend on A_0
- We are given a mapping $\pi : \Theta \mapsto \Delta(A_0)$.

(Static) Mechanism Design:

- Agents have private information: T_i is the set of types of agent i and

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

  describes the information player i has about θ and the types of other players.
- Payoffs only depend on A_0
- We are given a mapping $\pi : \Theta \mapsto \Delta(A_0)$.
- **Question:** Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

(Static) Mechanism Design:

- Agents have private information: T_i is the set of types of agent i and

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

  describes the information player i has about θ and the types of other players.
- Payoffs only depend on A_0
- We are given a mapping $\pi : \Theta \mapsto \Delta(A_0)$.
- **Question:** Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

Example: Ad auction design

(Static) Mechanism Design:

- Agents have private information: T_i is the set of types of agent i and

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

  describes the information player i has about θ and the types of other players.
- Payoffs only depend on A_0
- We are given a mapping $\pi : \Theta \mapsto \Delta(A_0)$.
- **Question:** Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

Example: Ad auction design

- One ad slot: $A_0 \subseteq (\{0,1\} \times \mathbb{R})^N$ and $(q,t) \in A_0$ if, and only if, $0 \leq \sum_{i=1}^{N} q_i \leq 1$.

(Static) Mechanism Design:

- Agents have private information: T_i is the set of types of agent i and

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

describes the information player i has about θ and the types of other players.

- Payoffs only depend on A_0
- We are given a mapping $\pi : \Theta \mapsto \Delta(A_0)$.
- **Question:** Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

Example: Ad auction design

- One ad slot: $A_0 \subseteq (\{0,1\} \times \mathbb{R})^N$ and $(q,t) \in A_0$ if, and only if, $0 \leq \sum_{i=1}^{N} q_i \leq 1$.
- $\Theta = \Theta_1 \times \ldots \Theta_N$; $T_i = \Theta_i$ denotes agent i's value for the ad; $\psi(\cdot|\theta) = \delta_\theta$.

(Static) Mechanism Design:

- Agents have private information: T_i is the set of types of agent i and

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

  describes the information player i has about θ and the types of other players.
- Payoffs only depend on A_0
- We are given a mapping $\pi : \Theta \mapsto \Delta(A_0)$.
- **Question:** Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

Example: Ad auction design

- One ad slot: $A_0 \subseteq (\{0,1\} \times \mathbb{R})^N$ and $(q,t) \in A_0$ if, and only if, $0 \leq \sum_{i=1}^{N} q_i \leq 1$.
- $\Theta = \Theta_1 \times \ldots \Theta_N$; $T_i = \Theta_i$ denotes agent i's value for the ad; $\psi(\cdot|\theta) = \delta_\theta$.
- π is the rule that assigns the slot to the advertiser w/highest θ_i.

Information Design:

- Only agents are endowed with actions, $|A_0| = 1$.

Information Design:

- Only agents are endowed with actions, $|A_0| = 1$.
- We are given a mapping $\pi : \Theta \mapsto \Delta(A)$.

Information Design:

- Only agents are endowed with actions, $|A_0| = 1$.
- We are given a mapping $\pi : \Theta \mapsto \Delta(A)$.
- **Question:** Can I find an information structure, that is, a set of types T_i for each player and a mapping

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

such that π is the equilibrium outcome of the game defined by $\langle G, \psi \rangle$?

Information Design:

- Only agents are endowed with actions, $|A_0| = 1$.
- We are given a mapping $\pi : \Theta \mapsto \Delta(A)$.
- **Question:** Can I find an information structure, that is, a set of types T_i for each player and a mapping

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N)$$

such that π is the equilibrium outcome of the game defined by $\langle G, \psi \rangle$?

Example: Say Google chooses the first-price auction

- $A_i = \mathbb{R}$ represents the bids of advertiser i
- Θ is the *common* value for the ad slot
- $u_i(a, \theta) = (\theta - a_i)\,\mathbb{1}[a_i = \max_j a_j]$
- π describes a possible distribution of bid profiles. (e.g., adversarial eqbm selection)

We are given a mapping $\pi : \Theta \mapsto \Delta(A)$.

MD: Given $\langle (T_i), \psi \rangle$, Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

ID: Given $A_1, \ldots, A_N$, Can we design $\langle (T_i), \psi \rangle$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi \rangle$?

We are given a mapping $\pi : \Theta \mapsto \Delta(A)$.

MD: Given $\langle (T_i), \psi \rangle$, Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

- The focus is on designing a game *given* an information structure

ID: Given $A_1, \ldots, A_N$, Can we design $\langle (T_i), \psi \rangle$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi \rangle$?

- The focus is on designing an information structure *given* a (base) game G.

We are given a mapping $\pi : \Theta \mapsto \Delta(A)$.

MD: Given $\langle (T_i), \psi \rangle$, Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

- The focus is on designing a game *given* an information structure
- The focus is on finding a game that rationalizes π as an outcome.

ID: Given $A_1, \ldots, A_N$, Can we design $\langle (T_i), \psi \rangle$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi \rangle$?

- The focus is on designing an information structure *given* a (base) game G.

We are given a mapping $\pi : \Theta \mapsto \Delta(A)$.

MD: Given $\langle (T_i), \psi \rangle$, Can we design actions for each player $A_1, \dots, A_N$ and an outcome function $f : \times_{i=1}^N A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

- The focus is on designing a game *given* an information structure
- The focus is on finding a game that rationalizes π as an outcome.

ID: Given $A_1, \dots, A_N$, Can we design $\langle (T_i), \psi \rangle$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi \rangle$?

- The focus is on designing an information structure *given* a (base) game G.
- The focus is on finding an information structure that rationalizes π as an outcome.

We are given a mapping $\pi : \Theta \mapsto \Delta(A)$.

MD: Given $\langle (T_i), \psi \rangle$, Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

- Mechanism design gives us a language to represent all games in a concise manner

ID: Given $A_1, \ldots, A_N$, Can we design $\langle (T_i), \psi \rangle$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi \rangle$?

- Information design gives us a language to represent all information structures in a concise manner

We are given a mapping $\pi : \Theta \mapsto \Delta(A)$.

MD: Given $\langle (T_i), \psi \rangle$, Can we design actions for each player $A_1, \ldots, A_N$ and an outcome function $f : \times_{i=1}^{N} A_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

Each lecture will be about these representations:

- Information Structures (Lecture #1)
- Games (mechanisms) (Lecture #2)

ID: Given $A_1, \ldots, A_N$, Can we design $\langle (T_i), \psi \rangle$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi \rangle$?

- Information design gives us a language to represent all information structures in a concise manner

Information Design

Suppose we know the base game G, that is:

- Players: N players, $i \in \{1, \ldots, N\}$
- Actions: A_i: player i's actions; A: action profiles,
- Θ, finite set of states of the world,
- Payoffs: $u_i : A \times \Theta \mapsto \mathbb{R}$: player i's payoffs,
- (common) prior $\mu_0 \in \Delta_+(\Theta)$,

Suppose we know the base game G, that is:

- Players: N players, $i \in \{1, \ldots, N\}$
- Actions: A_i: player i's actions; A: action profiles,
- Θ, finite set of states of the world,
- Payoffs: $u_i : A \times \Theta \mapsto \mathbb{R}$: player i's payoffs,
- (common) prior $\mu_0 \in \Delta_+(\Theta)$,

but we do not know:

Suppose we know the base game G, that is:

- Players: N players, $i \in \{1, \ldots, N\}$
- Actions: A_i: player i's actions; A: action profiles,
- Θ, finite set of states of the world,
- Payoffs: $u_i : A \times \Theta \mapsto \mathbb{R}$: player i's payoffs,
- (common) prior $\mu_0 \in \Delta_+(\Theta)$,

but we do not know:

- Information structure: what the players know about their state of the world

Suppose we know the base game G, that is:

- Players: N players, $i \in \{1, \ldots, N\}$
- Actions: A_i: player i's actions; A: action profiles,
- Θ, finite set of states of the world,
- Payoffs: $u_i : A \times \Theta \mapsto \mathbb{R}$: player i's payoffs,
- (common) prior $\mu_0 \in \Delta_+(\Theta)$,

but we do not know:

- Information structure: what the players know about their state of the world
- Extensive form - information players' have about others' moves

Suppose we know the base game G, that is:

- Players: N players, $i \in \{1, \ldots, N\}$
- Actions: A_i: player i's actions; A: action profiles,
- Θ, finite set of states of the world,
- Payoffs: $u_i : A \times \Theta \mapsto \mathbb{R}$: player i's payoffs,
- (common) prior $\mu_0 \in \Delta_+(\Theta)$,

but we do not know:

- Information structure: what the players know about their state of the world
- Extensive form - information players' have about others' moves

What are all the possible *equilibrium* outcomes?

Selling a good to a buyer w/unknown demand

- A seller wants to sell one unit of a good to a buyer,

[Roesler & Szentes, 2017; Ravid, Roesler, & Szentes, 2020]

Selling a good to a buyer w/unknown demand

- A seller wants to sell one unit of a good to a buyer,
- It is common knowledge that the buyer's valuation for the good $v \in [\underline{v}, \overline{v}]$ is drawn from a distribution μ_0,

[Roesler & Szentes, 2017; Ravid, Roesler, & Szentes, 2020]

Selling a good to a buyer w/unknown demand

- A seller wants to sell one unit of a good to a buyer,
- It is common knowledge that the buyer's valuation for the good $v \in [\underline{v}, \overline{v}]$ is drawn from a distribution μ_0,
- Before observing the price, (an adversarial) nature can disclose information to the buyer about her value.

[Roesler & Szentes, 2017; Ravid, Roesler, & Szentes, 2020]

Selling a good to a buyer w/unknown demand

- A seller wants to sell one unit of a good to a buyer,
- It is common knowledge that the buyer's valuation for the good $v \in [\underline{v}, \overline{v}]$ is drawn from a distribution μ_0,
- Before observing the price, (an adversarial) nature can disclose information to the buyer about her value.
- Each choice of information structure defines a game between the buyer and the seller.

[Roesler & Szentes, 2017; Ravid, Roesler, & Szentes, 2020]

Selling a good to a buyer w/unknown demand

- A seller wants to sell one unit of a good to a buyer,
- It is common knowledge that the buyer's valuation for the good $v \in [\underline{v}, \overline{v}]$ is drawn from a distribution μ_0,
- Before observing the price, (an adversarial) nature can disclose information to the buyer about her value.
- Each choice of information structure defines a game between the buyer and the seller.
- If the seller knows the buyer is fully informed, then the (unique) equilibrium outcome is for the seller to choose

$$p \in \arg\max_{v} (1 - \mu_0(v))v$$

and the buyer accepts if and only if $p \leq v$.

[Roesler & Szentes, 2017; Ravid, Roesler, & Szentes, 2020]

Selling a good to a buyer w/unknown demand

- A seller wants to sell one unit of a good to a buyer,
- It is common knowledge that the buyer's valuation for the good $v \in [\underline{v}, \overline{v}]$ is drawn from a distribution μ_0,
- Before observing the price, (an adversarial) nature can disclose information to the buyer about her value.
- Each choice of information structure defines a game between the buyer and the seller.
- If the seller knows the buyer is fully informed, then the (unique) equilibrium outcome is for the seller to choose

$$p \in \arg\max_v (1 - \mu_0(v))v$$

  and the buyer accepts if and only if $p \leq v$.
- Alternatively, if the seller knows the buyer does not know v, then $p = \mathbb{E}[v]$ is the unique outcome.

[Roesler & Szentes, 2017; Ravid, Roesler, & Szentes, 2020]

Selling a good to a buyer w/unknown demand

- A seller wants to sell one unit of a good to a buyer,
- It is common knowledge that the buyer's valuation for the good $v \in [\underline{v}, \overline{v}]$ is drawn from a distribution μ_0,
- Before observing the price, (an adversarial) nature can disclose information to the buyer about her value.
- Each choice of information structure defines a game between the buyer and the seller.
- If the seller knows the buyer is fully informed, then the (unique) equilibrium outcome is for the seller to choose

$$p \in \arg\max_{v} (1 - \mu_0(v))v$$

and the buyer accepts if and only if $p \leq v$.

- Alternatively, if the seller knows the buyer does not know v, then $p = \mathbb{E}[v]$ is the unique outcome.
- **Question:** What prices are consistent with equilibrium under some information structure?

[Roesler & Szentes, 2017; Ravid, Roesler, & Szentes, 2020]

Revenue-maximizing auction w/adversarial equilibrium selection

- An auctioneer can sell to one of N buyers

[Bergemann, Brooks, & Morris (many), Du, 2018; Brooks and Du, 2020,2022]

Revenue-maximizing auction w/adversarial equilibrium selection

- An auctioneer can sell to one of N buyers
- The auctioneer is deciding between multiple auction formats

[Bergemann, Brooks, & Morris (many), Du, 2018; Brooks and Du, 2020,2022]

Revenue-maximizing auction w/adversarial equilibrium selection

- An auctioneer can sell to one of N buyers
- The auctioneer is deciding between multiple auction formats
- It is common knowledge that values are common and the bidders' value is $v \sim \mu_0$

[Bergemann, Brooks, & Morris (many), Du, 2018; Brooks and Du, 2020,2022]

Revenue-maximizing auction w/adversarial equilibrium selection

- An auctioneer can sell to one of N buyers
- The auctioneer is deciding between multiple auction formats
- It is common knowledge that values are common and the bidders' value is $v \sim \mu_0$
- The auctioneer, however, does not know if the bidders will receive further information before playing the auction.

[Bergemann, Brooks, & Morris (many), Du, 2018; Brooks and Du, 2020,2022]

Revenue-maximizing auction w/adversarial equilibrium selection

- An auctioneer can sell to one of N buyers
- The auctioneer is deciding between multiple auction formats
- It is common knowledge that values are common and the bidders' value is $v \sim \mu_0$
- The auctioneer, however, does not know if the bidders will receive further information before playing the auction.
- Each auction format and each *information structure* determine a game

[Bergemann, Brooks, & Morris (many), Du, 2018; Brooks and Du, 2020,2022]

Revenue-maximizing auction w/adversarial equilibrium selection

- An auctioneer can sell to one of N buyers
- The auctioneer is deciding between multiple auction formats
- It is common knowledge that values are common and the bidders' value is $v \sim \mu_0$
- The auctioneer, however, does not know if the bidders will receive further information before playing the auction.
- Each auction format and each *information structure* determine a game
- The auctioneer may want to choose an auction format that performs well across different information structures.

[Bergemann, Brooks, & Morris (many), Du, 2018; Brooks and Du, 2020,2022]

Revenue-maximizing auction w/adversarial equilibrium selection

- An auctioneer can sell to one of N buyers
- The auctioneer is deciding between multiple auction formats
- It is common knowledge that values are common and the bidders' value is $v \sim \mu_0$
- The auctioneer, however, does not know if the bidders will receive further information before playing the auction.
- Each auction format and each *information structure* determine a game
- The auctioneer may want to choose an auction format that performs well across different information structures.

Note that as the auctioneer moves the mechanism, the adversary can pick a different information structure.

[Bergemann, Brooks, & Morris (many), Du, 2018; Brooks and Du, 2020,2022]

Applications in Computer Science and Operations

Design perspective: e.g.,

1. Platforms and Crowdsourcing (Bimpikis et al., 2020; Gur et al., 2019; Papanastasiou et al, 2018; Yang et al, 2019)
2. Social and Economic Networks (Candogan and Drakopoulos, 2020; Candogan, 2019)
3. Revenue Management (Drakopoulos et al., 2018; Kücülgul et al., 2019; Lingebrink & Iyer, 2018)
4. Firm competition (Banerjee et. al, 2022)
5. Queues (Lingebrink & Iyer, 2019; Che & Tercieux, 2020)
6. Team formation (Banerjee & Hssaine, 2018)

Structure of the problem: e.g.,

1. Dughmi, 2017; Dughmi & Xu, 2016-7
2. Ariely & Babichenko, 2016; Arieli et al., 2020

Single-agent

Single-agent

- Suppose there is one player, i.e., $N = 1$
- Then, the question about the extensive form is mute.

Single-agent

- Suppose there is one player, i.e., $N = 1$
- Then, the question about the extensive form is mute.

To fix ideas, consider the following example:

Single-agent setting

- Suppose there is one player, i.e., $N = 1$
- Then, the question about the extensive form is mute.

To fix ideas, consider the following example:

- The player can contribute to a public good, $A = \{C, NC\}$

- Suppose there is one player, i.e., $N = 1$
- Then, the question about the extensive form is mute.

To fix ideas, consider the following example:

- The player can contribute to a public good, $A = \{C, NC\}$
- The value of the public good, $\theta \in \{\theta_L, \theta_H\}$

- Suppose there is one player, i.e., $N = 1$
- Then, the question about the extensive form is mute.

To fix ideas, consider the following example:

- The player can contribute to a public good, $A = \{C, NC\}$
- The value of the public good, $\theta \in \{\theta_L, \theta_H\}$
- The cost of the public good is $\theta_L < c < \theta_H$.

- Suppose there is one player, i.e., $N = 1$
- Then, the question about the extensive form is mute.

To fix ideas, consider the following example:

- The player can contribute to a public good, $A = \{C, NC\}$
- The value of the public good, $\theta \in \{\theta_L, \theta_H\}$
- The cost of the public good is $\theta_L < c < \theta_H$.
- Payoffs

$$u(a, \theta) = (\theta - c)\mathbb{1}[a = C].$$

- Suppose there is one player, i.e., $N = 1$
- Then, the question about the extensive form is mute.

To fix ideas, consider the following example:

- The player can contribute to a public good, $A = \{C, NC\}$
- The value of the public good, $\theta \in \{\theta_L, \theta_H\}$
- The cost of the public good is $\theta_L < c < \theta_H$.
- Payoffs

$$u(a, \theta) = (\theta - c)\mathbb{1}[a = C].$$

- $\mu_0 \equiv \mu_0(\theta = \theta_H)$

- Suppose there is one player, i.e., $N = 1$
- Then, the question about the extensive form is mute.

To fix ideas, consider the following example:

- The player can contribute to a public good, $A = \{C, NC\}$
- The value of the public good, $\theta \in \{\theta_L, \theta_H\}$
- The cost of the public good is $\theta_L < c < \theta_H$.
- Payoffs

$$u(a, \theta) = (\theta - c)\mathbb{1}[a = C].$$

- $\mu_0 \equiv \mu_0(\theta = \theta_H)$
- Assume $\mu_0\theta_H + (1 - \mu_0)\theta_L < c$, i.e., at the prior it is not optimal to contribute.

Question: What distributions $\pi \in \Delta(\Theta \times \{C, NC\})$ are consistent with the player optimizing given an information structure?

Question: What distributions $\pi \in \Delta(\Theta \times \{C, NC\})$ are consistent with the player optimizing given an information structure?

What is an **information structure**?

- A set of signals $s \in S$
- A mapping $\hat{\pi} : \Theta \mapsto \Delta(S)$.

Question: What distributions $\pi \in \Delta(\Theta \times \{C, NC\})$ are consistent with the player optimizing given an information structure?

What is an **information structure**?

- A set of signals $s \in S$
- A mapping $\hat{\pi} : \Theta \mapsto \Delta(S)$.

Timing:

- Nature draws $\theta \sim \mu_0$ and $s \sim \hat{\pi}(\cdot|\theta)$
- The player observes s (but not θ) (knows μ_0 and $\hat{\pi}$)
- The player decides whether to contribute

Single-agent

- Observing $s \in S$, the player updates her belief

$$\mu_s(\theta) = \frac{\mu_0(\theta)\hat{\pi}(s|\theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta')\hat{\pi}(s|\theta')}$$

- Observing $s \in S$, the player updates her belief

$$\mu_s(\theta) = \frac{\mu_0(\theta)\hat{\pi}(s|\theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta')\hat{\pi}(s|\theta')}$$

- and then chooses

$$a^*_{\hat{\pi}}(\mu_s) \in \arg\max_{\tilde{a} \in A} \sum_{\theta \in \Theta} \mu_s(\theta) u(\tilde{a}, \theta),$$

where

$$a^*_{\hat{\pi}}(\mu_s) \begin{cases} = C & \text{if } \mu > \frac{c-\theta_L}{\theta_H - \theta_L} \\ = NC & \text{if } \mu < \frac{c-\theta_L}{\theta_H - \theta_L} \\ \in \{C, NC\} & \text{otherwise} \end{cases}$$

- Observing $s \in S$, the player updates her belief

$$\mu_s(\theta) = \frac{\mu_0(\theta)\hat{\pi}(s|\theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta')\hat{\pi}(s|\theta')}$$

- and then chooses

$$a^*_{\hat{\pi}}(\mu_s) \in \arg\max_{\tilde{a} \in A} \sum_{\theta \in \Theta} \mu_s(\theta) u(\tilde{a}, \theta),$$

where

$$a^*_{\hat{\pi}}(\mu_s) \begin{cases} = C & \text{if } \mu > \frac{c-\theta_L}{\theta_H - \theta_L} \\ = NC & \text{if } \mu < \frac{c-\theta_L}{\theta_H - \theta_L} \\ \in \{C, NC\} & \text{otherwise} \end{cases}$$

- Then, we can define the probability that action a is taken at θ under $\hat{\pi}$ as

$$\Pr_{\hat{\pi}}(a|\theta) = \sum_{s \in S} \hat{\pi}(s|\theta) \mathbb{1}[a = a^*_{\hat{\pi}}(\mu_s)]$$

- Observing $s \in S$, the player updates her belief

$$\mu_s(\theta) = \frac{\mu_0(\theta)\hat{\pi}(s|\theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta')\hat{\pi}(s|\theta')}$$

- and then chooses

$$a^*_{\hat{\pi}}(\mu_s) \in \arg\max_{\tilde{a} \in A} \sum_{\theta \in \Theta} \mu_s(\theta) u(\tilde{a}, \theta),$$

where

$$a^*_{\hat{\pi}}(\mu_s) \begin{cases} = C & \text{if } \mu > \frac{c-\theta_L}{\theta_H - \theta_L} \\ = NC & \text{if } \mu < \frac{c-\theta_L}{\theta_H - \theta_L} \\ \in \{C, NC\} & \text{otherwise} \end{cases}$$

- Then, we can define the probability that action a is taken at θ under $\hat{\pi}$ as

$$\Pr_{\hat{\pi}}(a|\theta) = \sum_{s \in S} \hat{\pi}(s|\theta) \mathbb{1}[a = a^*_{\hat{\pi}}(\mu_s)]$$

- Using the prior μ_0 we can construct:

$$\tilde{\pi}(a, \theta|\hat{\pi}) = \mu_0(\theta) \Pr_{\hat{\pi}}(a|\theta)$$

Definition

$\pi \in \Delta(\Theta \times A)$ is consistent with some information structure if there exists $\langle S, \hat{\pi} \rangle$ such that for all $\theta \in \Theta$ and $a \in A$,

$$\pi(a, \theta) = \tilde{\pi}(a, \theta | \hat{\pi})$$

Definition

$\pi \in \Delta(\Theta \times A)$ is consistent with some information structure if there exists $\langle S, \hat{\pi} \rangle$ such that for all $\theta \in \Theta$ and $a \in A$,

$$\pi(a, \theta) = \tilde{\pi}(a, \theta | \hat{\pi})$$

Goal

Characterize the set

$$\Pi(\mu_0) = \{\pi \in \Delta(\Theta \times A) : (\exists \langle S, \hat{\pi} \rangle) \pi \equiv \tilde{\pi}(\cdot | \hat{\pi})\}.$$

Suppose we have $\pi = \tilde{\pi}(\cdot|\hat{\pi})$:

$$\Pr_{\hat{\pi}}(a|\theta) = \sum_{s \in S} \hat{\pi}(s|\theta) \mathbb{1}[a = a^*_{\hat{\pi}}(\mu_s)]$$

$$\tilde{\pi}(a, \theta|\hat{\pi}) = \mu_0(\theta) \Pr_{\hat{\pi}}(a|\theta)$$

Suppose we have $\pi = \tilde{\pi}(\cdot|\hat{\pi})$:

$$\Pr_{\hat{\pi}}(a|\theta) = \sum_{s \in S} \hat{\pi}(s|\theta)\mathbb{1}[a = a^*_{\hat{\pi}}(\mu_s)]$$

$$\tilde{\pi}(a, \theta|\hat{\pi}) = \mu_0(\theta)\Pr_{\hat{\pi}}(a|\theta)$$

Two implications of this construction:

Suppose we have $\pi = \tilde{\pi}(\cdot|\hat{\pi})$:

$$\mathrm{Pr}_{\hat{\pi}}(a|\theta) = \sum_{s \in S} \hat{\pi}(s|\theta)\mathbb{1}[a = a^*_{\hat{\pi}}(\mu_s)]$$

$$\tilde{\pi}(a, \theta|\hat{\pi}) = \mu_0(\theta)\mathrm{Pr}_{\hat{\pi}}(a|\theta)$$

Two implications of this construction:

- "Martingale" property:

$$\sum_{a \in A} \pi(a, \theta) = \mu_0(\theta)$$

Suppose we have $\pi = \tilde{\pi}(\cdot|\hat{\pi})$:

$$\Pr_{\hat{\pi}}(a|\theta) = \sum_{s \in S} \hat{\pi}(s|\theta)\mathbb{1}[a = a^*_{\hat{\pi}}(\mu_s)]$$

$$\tilde{\pi}(a, \theta|\hat{\pi}) = \mu_0(\theta)\Pr_{\hat{\pi}}(a|\theta)$$

Two implications of this construction:

- "Martingale" property:

$$\sum_{a \in A} \pi(a, \theta) = \mu_0(\theta)$$

- Obedience: for all $a \in A$ such that $\sum_{\theta \in \Theta} \pi(a, \theta) > 0$

$$\sum_{\theta \in \Theta} \pi(a, \theta) \left[u(a, \theta) - u(a', \theta)\right] \geq 0 (\forall a' \in A)$$

Suppose we have $\pi = \tilde{\pi}(\cdot|\hat{\pi})$:

$$\Pr_{\hat{\pi}}(a|\theta) = \sum_{s\in S} \hat{\pi}(s|\theta)\mathbb{1}[a = a^*_{\hat{\pi}}(\mu_s)]$$

$$\tilde{\pi}(a,\theta|\hat{\pi}) = \mu_0(\theta)\Pr_{\hat{\pi}}(a|\theta)$$

Obedience:

$$\sum_{\theta\in\Theta} \pi(a,\theta)u(a,\theta) = \sum_{\theta\in\Theta} \mu_0(\theta) \sum_{s\in S} \hat{\pi}(s|\theta)\mathbb{1}[a = a^*(\mu_s)]$$

$$= \sum_{s\in S} \frac{\sum_{\theta'\in\Theta} \mu_0(\theta')\hat{\pi}(s|\theta')}{\sum_{\theta'\in\Theta} \mu_0(\theta')\hat{\pi}(s|\theta')} \sum_{\theta\in\Theta} \mu_0(\theta)\hat{\pi}(s|\theta)\mathbb{1}[a = a^*(\mu_s)]u(a,\theta)$$

$$= \sum_{s\in S} \Pr_{\hat{\pi}}(s) \sum_{\theta\in\Theta} \frac{\mu_0(\theta)\hat{\pi}(s|\theta)}{\Pr_{\hat{\pi}}(s)} \mathbb{1}[a = a^*(\mu_s)]u(a,\theta) = \sum_{s\in S} \Pr_{\hat{\pi}}(s) \sum_{\theta\in\Theta} \mu_s(\theta)u(a,\theta)\mathbb{1}[a = a^*(\mu_s)]$$

$$\geq \sum_{s\in S} \Pr_{\hat{\pi}}(s) \sum_{\theta\in\Theta} \mu_s(\theta)\mathbb{1}[a = a^*(\mu_s)]u(a',\theta)$$

Single-agent: Recommendation principle

Theorem (Myerson, 1982; Kamenica and Gentzkow, 2011)

$\pi \in \Pi(\mu_0)$ *if and only if*

$$\sum_{a \in A} \pi(a, \theta) = \mu_0(\theta) \tag{M}$$

$$(\forall a' \in A) \sum_{\theta \in \Theta} \pi(a, \theta) \left[u(a, \theta) - u(a', \theta)\right] \geq 0 \tag{O}$$

Theorem (Myerson, 1982; Kamenica and Gentzkow, 2011)

$\pi \in \Pi(\mu_0)$ *if and only if*

$$\sum_{a \in A} \pi(a, \theta) = \mu_0(\theta) \qquad \text{(M)}$$

$$(\forall a' \in A) \sum_{\theta \in \Theta} \pi(a, \theta) \left[u(a, \theta) - u(a', \theta)\right] \geq 0 \qquad \text{(O)}$$

What is the information structure that rationalizes such π?

Theorem (Myerson, 1982; Kamenica and Gentzkow, 2011)

$\pi \in \Pi(\mu_0)$ *if and only if*

$$\sum_{a \in A} \pi(a, \theta) = \mu_0(\theta) \qquad \text{(M)}$$

$$(\forall a' \in A) \sum_{\theta \in \Theta} \pi(a, \theta) \left[u(a, \theta) - u(a', \theta)\right] \geq 0 \qquad \text{(O)}$$

What is the information structure that rationalizes such π? $S = A$ and

Theorem (Myerson, 1982; Kamenica and Gentzkow, 2011)

$\pi \in \Pi(\mu_0)$ *if and only if*

$$\sum_{a \in A} \pi(a, \theta) = \mu_0(\theta) \tag{M}$$

$$(\forall a' \in A) \sum_{\theta \in \Theta} \pi(a, \theta) \left[u(a, \theta) - u(a', \theta)\right] \geq 0 \tag{O}$$

What is the information structure that rationalizes such π? $S = A$ and

$$\hat{\pi}(a|\theta) = \frac{\pi(a, \theta)}{\sum_{a' \in A} \pi(a', \theta)} = \frac{\pi(a, \theta)}{\mu_0(\theta)}.$$

Theorem (Myerson, 1982; Kamenica and Gentzkow, 2011)

$\pi \in \Pi(\mu_0)$ *if and only if*

$$\sum_{a \in A} \pi(a, \theta) = \mu_0(\theta) \tag{M}$$

$$(\forall a' \in A) \sum_{\theta \in \Theta} \pi(a, \theta) \left[u(a, \theta) - u(a', \theta)\right] \geq 0 \tag{O}$$

What is the information structure that rationalizes such π? $S = A$ and

$$\hat{\pi}(a|\theta) = \frac{\pi(a, \theta)}{\sum_{a' \in A} \pi(a', \theta)} = \frac{\pi(a, \theta)}{\mu_0(\theta)}.$$

The martingale property implies $\hat{\pi}(\cdot|\theta) \in \Delta(A)$.

- Suppose we have a third party whose payoff depends on the state of the world and the agent's action

- Suppose we have a third party whose payoff depends on the state of the world and the agent's action
- Given π that satisfies (M) and (O), the payoff of the third party is

$$\tilde{v}(\pi) \equiv \sum_{a,\theta} \pi(a,\theta) v(a,\theta)$$

- Suppose we have a third party whose payoff depends on the state of the world and the agent's action
- Given π that satisfies (M) and (O), the payoff of the third party is

$$\tilde{v}(\pi) \equiv \sum_{a,\theta} \pi(a,\theta) v(a,\theta)$$

- If the third party is a designer, then it is natural to think of

$$\max\{\tilde{v}(\pi) : \pi \in \Pi(\mu_0)\}$$

- Suppose we have a third party whose payoff depends on the state of the world and the agent's action
- Given π that satisfies (M) and (O), the payoff of the third party is

$$\tilde{v}(\pi) \equiv \sum_{a,\theta} \pi(a,\theta) v(a,\theta)$$

- If the third party is a designer, then it is natural to think of

$$\max\{\tilde{v}(\pi) : \pi \in \Pi(\mu_0)\}$$

- Alternatively, the third party may be concerned with adversarial selection:

$$\min\{\tilde{v}(\pi) : \pi \in \Pi(\mu_0)\}$$

- Suppose we have a third party whose payoff depends on the state of the world and the agent's action
- Given π that satisfies (M) and (O), the payoff of the third party is

$$\tilde{v}(\pi) \equiv \sum_{a,\theta} \pi(a,\theta) v(a,\theta)$$

- If the third party is a designer, then it is natural to think of

$$\max\{\tilde{v}(\pi) : \pi \in \Pi(\mu_0)\}$$

- Alternatively, the third party may be concerned with adversarial selection:

$$\min\{\tilde{v}(\pi) : \pi \in \Pi(\mu_0)\}$$

- Both are linear programs.

To characterize the set $\Pi(\mu_0)$ in the contribution example,

- the martingale property implies that it is enough to characterize the pair $\{\pi(C|\theta_H), \pi(C|\theta_L)\}$
- Obedience implies that

$$\mu_0 \pi(C|\theta_H)\theta_H + (1-\mu_0)\pi(C|\theta_L)\theta_L \geq c$$

Contribution game

To characterize the set $\Pi(\mu_0)$ in the contribution example,

- the martingale property implies that it is enough to characterize the pair $\{\pi(C|\theta_H), \pi(C|\theta_L)\}$
- Obedience implies that

$$\mu_0 \pi(C|\theta_H)\theta_H + (1-\mu_0)\pi(C|\theta_L)\theta_L \geq c$$

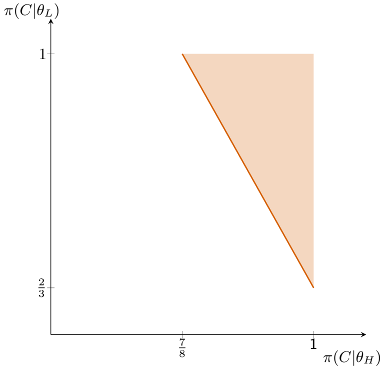

[Syrgkanis, Tamer, and Ziani, 2021]

An alternative approach to the single-agent question is the **belief approach**:

- Recall that $s \to \mu_s \to a^*(\mu_s)$:

$$a^*_{\hat{\pi}}(\mu_s) \in \arg\max_{\tilde{a} \in A} \sum_{\theta \in \Theta} \mu_s(\theta) u(\tilde{a}, \theta),$$

  where

$$a^*_{\hat{\pi}}(\mu_s) \begin{cases} = C & \text{if } \mu > \frac{c-\theta_L}{\theta_H - \theta_L} \\ = NC & \text{if } \mu < \frac{c-\theta_L}{\theta_H - \theta_L} \\ \in \{C, NC\} & \text{otherwise} \end{cases}$$

- The only part that can depend on $\hat{\pi}$ is what happens at the threshold belief, $\frac{c-\theta_L}{\theta_H - \theta_L}$.
- Except for that, we can replace the signals $s \in S$ for the beliefs they induce μ_s

Theorem (Kamenica & Gentzkow, 2011)

Fix a selection $a^*(\mu_s)$ *of the player's best-response correspondence. The following are equivalent:*

1. *There is a literal signal structure* $\langle \Delta(\Theta), \hat{\pi} \rangle$ *that induces* $\pi \in \Delta(\Theta \times A)$ *and satisfies:*

$$\mu(\theta) = \frac{\mu_0(\theta)\hat{\pi}(\mu|\theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta')\hat{\pi}(\mu|\theta')},$$

2. *There is an obedient signal structure* $\langle A, \hat{\pi} \rangle$ *that induces* $\pi \in \Delta(\Theta \times A)$.

Theorem (Kamenica & Gentzkow, 2011)

Fix a selection $a^*(\mu_s)$ of the player's best-response correspondence. The following are equivalent:

1. *There is a literal signal structure $\langle \Delta(\Theta), \hat{\pi} \rangle$ that induces $\pi \in \Delta(\Theta \times A)$ and satisfies:*

$$\mu(\theta) = \frac{\mu_0(\theta)\hat{\pi}(\mu|\theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta')\hat{\pi}(\mu|\theta')},$$

2. *There is an obedient signal structure $\langle A, \hat{\pi} \rangle$ that induces $\pi \in \Delta(\Theta \times A)$.*

In the single agent case, we can either

- recommend the agent what action to take,
- tell the agent what belief they should have.

Theorem (Kamenica & Gentzkow, 2011)

Fix a selection $a^*(\mu_s)$ of the player's best-response correspondence. The following are equivalent:

1. *There is a literal signal structure $\langle \Delta(\Theta), \hat{\pi} \rangle$ that induces $\pi \in \Delta(\Theta \times A)$ and satisfies:*

$$\mu(\theta) = \frac{\mu_0(\theta)\hat{\pi}(\mu|\theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta')\hat{\pi}(\mu|\theta')},$$

2. *There is an obedient signal structure $\langle A, \hat{\pi} \rangle$ that induces $\pi \in \Delta(\Theta \times A)$.*

In the single agent case, we can either

- recommend the agent what action to take,
- tell the agent what belief they should have.

Each approach has its downsides:

- Belief-approach requires knowing how agent breaks ties
- Action approach can be complicated if the action space is complicated (Lecture #2)

A literal signal structure is a Blackwell-experiment and it induces a distribution over beliefs

$$\tau(\mu) = \sum_{\theta \in \Theta} \mu_0(\theta) \hat{\pi}(\mu|\theta).$$

So we can alternatively work with $\tau \in \Delta(\Delta(\Theta))$ if we know which ones are feasible:

A literal signal structure is a Blackwell-experiment and it induces a distribution over beliefs

$$\tau(\mu) = \sum_{\theta \in \Theta} \mu_0(\theta)\hat{\pi}(\mu|\theta).$$

So we can alternatively work with $\tau \in \Delta(\Delta(\Theta))$ if we know which ones are feasible:

Theorem (Blackwell, 1951; Aumann & Maschler, 1965; Kamenica & Gentzkow, 2011)

$\tau \in \Delta(\Delta(\Theta))$ *is consistent with a signal structure and prior* μ_0 *if and only if*

$$(\forall \theta \in \Theta) \sum_{\mu \in \Delta(\Theta)} \tau(\mu)\mu(\theta) = \mu_0(\theta).$$

Back to many players

Ingredients:

- N players, $i \in \{1, \dots, N\}$
- A_i: player i's actions; A: action profiles,
- Θ, finite set of states of the world,
- $u_i : A \times \Theta \mapsto \mathbb{R}$: player i's payoffs,
- (common) prior $\mu_0 \in \Delta_+(\Theta)$

Questions

1. Suppose players take their actions simultaneously. What is the set of distributions over action profiles

$$\pi \in \Delta(\Theta \times A)$$

   that is consistent with equilibrium under *some* information structure?
2. Same question, but we know neither the information structure nor the extensive form.

Ingredients:

- N players, $i \in \{1, \ldots, N\}$
- A_i: player i's actions; A: action profiles,
- Θ, finite set of states of the world,
- $u_i : A \times \Theta \mapsto \mathbb{R}$: player i's payoffs,
- (common) prior $\mu_0 \in \Delta_+(\Theta)$

Questions

1. Suppose players take their actions simultaneously. What is the set of distributions over action profiles

$$\pi \in \Delta(\Theta \times A)$$

that is consistent with equilibrium under *some* information structure?

2. Same question, but we know neither the information structure nor the extensive form.

An **information structure** is a tuple $\{T_1, \ldots, T_N\}$ of type spaces and a mapping

$$\psi : \Theta \mapsto \Delta(T_1 \times \cdots \times T_N).$$

- Each player knows $\langle T_1, \ldots, T_N, \psi \rangle$
- Each player observes t_i (but not t_{-i} or θ) before taking their action.
- After observing t_i, player i also needs a *conjecture* of how players choose their actions on the basis of information.
- Assume players play Bayes' Nash equilibrium.

Definition

A strategy profile $(\sigma_i)_{i=1}^N$, $\sigma_i : T_i \mapsto \Delta(A_i)$ is a Bayes' Nash equilibrium of the base game G under $\langle T, \psi \rangle$, if for all $i \in \{1, \ldots, N\}$, $t_i \in T_i$, $a_i \in A_i$, and $a_i' \in A_i$, the following holds:

$$\sum_{\theta \in \Theta} \mu_0(\theta) \sum_{t_{-i} \in T_{-i}} \psi(t_i, t_{-i}|\theta) \sum_{a_{-i} \in A_{-i}} \prod_{j \neq i} \sigma_j(a_j|t_j) \left[u_i(a_i, a_{-i}, \theta) - u_i(a_i', a_{-i}, \theta)\right] \geq 0$$

Note that from here we can again construct a joint probability $\pi \in \Delta(\Theta \times A)$. Namely,

$$\Pr_\psi(a|\theta) = \sum_{t \in T} \psi(t|\theta) \prod_{i=1}^N \sigma_i(a_i|t_i).$$

$$\pi(a, \theta) = \mu_0(\theta) \Pr_\psi(a|\theta)$$

Question: Which $\pi \in \Delta(\Theta \times A)$ are consistent with BNE under some information structure in base game G? Call the set of such π, $\Pi^*(G, \mu_0)$.

Obedience

Definition

$\pi \in \Delta(\Theta \times A)$ is **obedient** if for all $i \in \{1, \dots, N\}$, all $a_i \in A_i$ and all $a_i' \in A_i$,

$$\sum_{\theta \in \Theta} \sum_{a_{-i} \in A_{-i}} \pi(a_i, a_{-i}, \theta) \left[u(a_i, a_{-i}, \theta) - u(a_i', a_{-i}, \theta) \right] \geq 0$$

Definition

$\pi \in \Delta(\Theta \times A)$ is **obedient** if for all $i \in \{1, \ldots, N\}$, all $a_i \in A_i$ and all $a_i' \in A_i$,

$$\sum_{\theta \in \Theta} \sum_{a_{-i} \in A_{-i}} \pi(a_i, a_{-i}, \theta) \left[u(a_i, a_{-i}, \theta) - u(a_i', a_{-i}, \theta) \right] \geq 0$$

Definition (Bergemann and Morris, 2016)

$\pi \in \Delta(\Theta \times A)$ is a **Bayes' correlated equilibrium** if

1. π is obedient,
2. π satisfies the martingale property at μ_0.

Let $\mathrm{BCE}(G, \mu_0)$ denote the set of Bayes' correlated equilibrium.

Obedience

Definition

$\pi \in \Delta(\Theta \times A)$ is **obedient** if for all $i \in \{1, \ldots, N\}$, all $a_i \in A_i$ and all $a_i' \in A_i$,

$$\sum_{\theta \in \Theta} \sum_{a_{-i} \in A_{-i}} \pi(a_i, a_{-i}, \theta) \left[u(a_i, a_{-i}, \theta) - u(a_i', a_{-i}, \theta)\right] \geq 0$$

Definition (Bergemann and Morris, 2016)

$\pi \in \Delta(\Theta \times A)$ is a **Bayes' correlated equilibrium** if

1. π is obedient,
2. π satisfies the martingale property at μ_0.

Let $\mathrm{BCE}(G, \mu_0)$ denote the set of Bayes' correlated equilibrium.

When $|\Theta| = 1$, this is correlated equilibrium.

Characterization: Recommendation Principle

Theorem (Bergemann and Morris, 2016)

An outcome distribution $\pi \in \Delta(\Theta \times A)$ *is consistent with equilibrium in* G *under some information structure* $\langle T, \psi \rangle$ *if and only if it is a Bayes' correlated equilibrium.*

That is,

$$\Pi^*(G, \mu_0) = \mathrm{BCE}(G, \mu_0).$$

- Again, the information structure is the one that *recommends* the player what action to do and nothing else.

Sequential Information Design

Ingredients:

- N players, $i \in \{1, \ldots, N\}$
- A_i: player i's actions; A: action profiles,
- Θ, finite set of states of the world,
- $u_i : A \times \Theta \mapsto \mathbb{R}$: player i's payoffs,
- (common) prior $\mu_0 \in \Delta_+(\Theta)$

Question: What is the set of distributions over action profiles

$$\pi \in \Delta(\Theta \times A)$$

that is consistent with equilibrium under some information structure *and* extensive form?

		P_3	
		C_3	D_3
P_2	C_2	(g, g)	$(0, \Delta)$
	D_2	$(\Delta, 0)$	(b, b)

Prisoner's dilemma: $\Delta > g > b > 0$

		P_3	
		C_3	D_3
P_2	C_2	(g, g)	$(0, \Delta)$
	D_2	$(\Delta, 0)$	(b, b)

Prisoner's dilemma: $\Delta > g > b > 0$

- If $g > \frac{1}{2}\Delta + \frac{1}{2}b$, there is an extensive form in which (C, C) is the equilibrium outcome.

Prisoner's dilemma

		P_3	
		C_3	D_3
P_2	C_2	(g, g)	$(0, \Delta)$
	D_2	$(\Delta, 0)$	(b, b)

Prisoner's dilemma: $\Delta > g > b > 0$

- If $g > \frac{1}{2}\Delta + \frac{1}{2}b$, there is an extensive form in which (C, C) is the equilibrium outcome.
- Flip a coin who moves first

Prisoner's dilemma

		P_3	
		C_3	D_3
P_2	C_2	(g, g)	$(0, \Delta)$
	D_2	$(\Delta, 0)$	(b, b)

Prisoner's dilemma: $\Delta > g > b > 0$

- If $g > \frac{1}{2}\Delta + \frac{1}{2}b$, there is an extensive form in which (C, C) is the equilibrium outcome.
- Flip a coin who moves first
- Approach the first mover and announce

			P_3	
			C_3	D_3
P_2		C_2	(g, g)	$(0, \Delta)$
		D_2	$(\Delta, 0)$	(b, b)

Prisoner's dilemma: $\Delta > g > b > 0$

- If $g > \frac{1}{2}\Delta + \frac{1}{2}b$, there is an extensive form in which (C, C) is the equilibrium outcome.
- Flip a coin who moves first
- Approach the first mover and announce
  - If he plays C, then the recommendation to the second player is C

		P_3	
		C_3	D_3
P_2	C_2	(g, g)	$(0, \Delta)$
	D_2	$(\Delta, 0)$	(b, b)

Prisoner's dilemma: $\Delta > g > b > 0$

- If $g > \frac{1}{2}\Delta + \frac{1}{2}b$, there is an extensive form in which (C, C) is the equilibrium outcome.
- Flip a coin who moves first
- Approach the first mover and announce
    - If he plays C, then the recommendation to the second player is C
    - If he plays D, then the recommendation to the second player is D

		P_3	
		C_3	D_3
P_2	C_2	(g, g)	$(0, \Delta)$
	D_2	$(\Delta, 0)$	(b, b)

Prisoner's dilemma: $\Delta > g > b > 0$

- If $g > \frac{1}{2}\Delta + \frac{1}{2}b$, there is an extensive form in which (C, C) is the equilibrium outcome.
- Flip a coin who moves first
- Approach the first mover and announce
  - If he plays C, then the recommendation to the second player is C
  - If he plays D, then the recommendation to the second player is D

		P_3	
		C_3	D_3
P_2	C_2	(g, g)	$(0, \Delta)$
	D_2	$(\Delta, 0)$	(b, b)

Prisoner's dilemma: $\Delta > g > b > 0$

- If $g > \frac{1}{2}\Delta + \frac{1}{2}b$, there is an extensive form in which (C, C) is the equilibrium outcome.
- Flip a coin who moves first
- Approach the first mover and announce
  - If he plays C, then the recommendation to the second player is C
  - If he plays D, then the recommendation to the second player is D
- Note that if a player hears a recommendation of D, it is dominant to play D – they know they are moving second and the other played D

			P_3	
			C_3	D_3
P_2		C_2	(g, g)	$(0, \Delta)$
		D_2	$(\Delta, 0)$	(b, b)

Prisoner's dilemma: $\Delta > g > b > 0$

- If $g > \frac{1}{2}\Delta + \frac{1}{2}b$, there is an extensive form in which (C, C) is the equilibrium outcome.
- Flip a coin who moves first
- Approach the first mover and announce
  - If he plays C, then the recommendation to the second player is C
  - If he plays D, then the recommendation to the second player is D
- Note that if a player hears a recommendation of D, it is dominant to play D – they know they are moving second and the other played D
- If they get told C,

$$g \geq \frac{1}{2}\Delta + \frac{1}{2}b$$