Limited Commitment: Mechanism Design meets Information Design

Data-Driven Decision Processes Bootcamp

Laura Doval
Columbia Business School and CEPR

Based on joint work with Vasiliki Skreta

(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: Google ad auction design

- $A_0 \subseteq (\{0,1\} \times \mathbb{R})^N$ 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 advertiser i's value for the slot; $\psi(\cdot|\theta) = \delta_\theta$.
- π is the rule that assigns the good to the advertiser w/ highest θ_i.

(Static) Mechanism Design: (in more standard textbook notation)

- Agents have private information: $\Theta = \times_{i=1}^{N} \Theta_i$ and agent i knows θ_i. That is,

$$\psi : \Theta \mapsto \Delta(\Theta_1 \times \cdots \times \Theta_N)$$

  is such that $\psi(\cdot|\theta) = \delta_\theta$.
- 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 $M_1, \ldots, M_N$ and an outcome function $f : \times_{i=1}^{N} M_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

Example: Google ad auction design

- $A_0 \subseteq (\{0,1\} \times \mathbb{R})^N$ 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 advertiser i's value for the slot; $\psi(\cdot|\theta) = \delta_\theta$.
- π is the rule that assigns the good to the advertiser w/ highest θ_i.

Given $\pi : \Theta \mapsto \Delta(A_0)$,

MD question: Can we design actions for each player $M_1, \dots, M_N$ and an outcome function $f : \times_{i=1}^{N} M_i \mapsto \Delta(A_0)$ such that π is the equilibrium outcome of the game defined by $\langle G, \psi, f \rangle$?

Given $\pi : \Theta \mapsto \Delta(A_0)$,

MD question: Can we design actions for each player $M_1, \ldots, M_N$ and an outcome function $f : \times_{i=1}^{N} M_i \mapsto \Delta(A_0)$ 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

Given $\pi : \Theta \mapsto \Delta(A_0)$,

MD question: Can we design actions for each player $M_1, \ldots, M_N$ and an outcome function $f : \times_{i=1}^{N} M_i \mapsto \Delta(A_0)$ 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 first order concern is that the information is in the hands of selfish players

Given $\pi : \Theta \mapsto \Delta(A_0)$,

MD question: Can we design actions for each player $M_1, \ldots, M_N$ and an outcome function $f : \times_{i=1}^{N} M_i \mapsto \Delta(A_0)$ 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 first order concern is that the information is in the hands of selfish players
- and the information is needed to know what is the "correct" outcome.

Given $\pi : \Theta \mapsto \Delta(A_0)$,

MD question: Can we design actions for each player $M_1, \ldots, M_N$ and an outcome function $f : \times_{i=1}^{N} M_i \mapsto \Delta(A_0)$ 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 first order concern is that the information is in the hands of selfish players
- and the information is needed to know what is the "correct" outcome.

- We need to be able to consider *all* possible games

Given $\pi : \Theta \mapsto \Delta(A_0)$,

MD question: Can we design actions for each player $M_1, \ldots, M_N$ and an outcome function $f : \times_{i=1}^{N} M_i \mapsto \Delta(A_0)$ 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 first order concern is that the information is in the hands of selfish players
- and the information is needed to know what is the "correct" outcome.

- We need to be able to consider *all* possible games
- Mechanism design provides us with a language to do this via the **revelation principle**.

Revelation principle: first-pass

Theorem (Gibbard, 1973; Myerson, 1979; Dasgupta et al, 1979)

There exists a game that has π as an equilibrium outcome if and only if the following game implements π:

1. *Actions $M_i = \Theta_i$*
2. *When players take actions $\theta' = (\theta'_1, \ldots, \theta'_N)$, the outcome is $f(\overline{\theta}) = \pi(\cdot|\theta')$.*

Furthermore, it is without loss of generality to assume that the players find it optimal to truthfully report their types.

Revelation principle: first-pass

Theorem (Gibbard, 1973; Myerson, 1979; Dasgupta et al, 1979)

There exists a game that has π as an equilibrium outcome if and only if the following game implements π:

1. *Actions* $M_i = \Theta_i$
2. *When players take actions* $\theta' = (\theta'_1, \ldots, \theta'_N)$, *the outcome is* $f(\overline{\theta}) = \pi(\cdot|\theta')$.

Furthermore, it is without loss of generality to assume that the players find it optimal to truthfully report their types.

M A

Θ

indirect mechanism

Revelation principle: first-pass

Theorem (Gibbard, 1973; Myerson, 1979; Dasgupta et al, 1979)

There exists a game that has π as an equilibrium outcome if and only if the following game implements π:

1. *Actions $M_i = \Theta_i$*
2. *When players take actions $\theta' = (\theta'_1, \ldots, \theta'_N)$, the outcome is $f(\overline{\theta}) = \pi(\cdot|\theta')$.*

Furthermore, it is without loss of generality to assume that the players find it optimal to truthfully report their types.

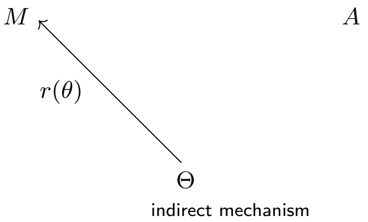

Revelation principle: first-pass

Theorem (Gibbard, 1973; Myerson, 1979; Dasgupta et al, 1979)

There exists a game that has π as an equilibrium outcome if and only if the following game implements π:

1. *Actions* $M_i = \Theta_i$
2. *When players take actions* $\theta' = (\theta'_1, \ldots, \theta'_N)$, *the outcome is* $f(\overline{\theta}) = \pi(\cdot|\theta')$.

Furthermore, it is without loss of generality to assume that the players find it optimal to truthfully report their types.

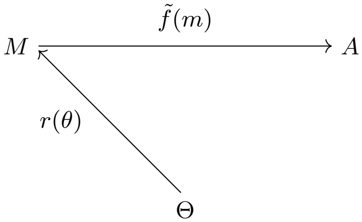

indirect mechanism

Revelation principle: first-pass

Theorem (Gibbard, 1973; Myerson, 1979; Dasgupta et al, 1979)

There exists a game that has π as an equilibrium outcome if and only if the following game implements π:

1. *Actions* $M_i = \Theta_i$
2. *When players take actions* $\theta' = (\theta'_1, \ldots, \theta'_N)$, *the outcome is* $f(\overline{\theta}) = \pi(\cdot|\theta')$.

Furthermore, it is without loss of generality to assume that the players find it optimal to truthfully report their types.

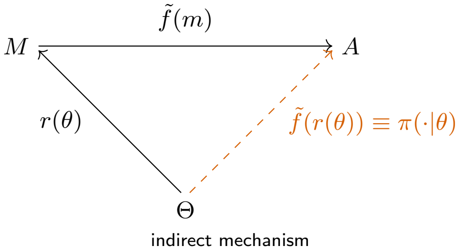

Revelation principle: first-pass

Theorem (Gibbard, 1973; Myerson, 1979; Dasgupta et al, 1979)

There exists a game that has π as an equilibrium outcome if and only if the following game implements π:

1. *Actions* $M_i = \Theta_i$
2. *When players take actions* $\theta' = (\theta'_1, \ldots, \theta'_N)$*, the outcome is* $f(\bar{\theta}) = \pi(\cdot|\theta')$*.*

Furthermore, it is without loss of generality to assume that the players find it optimal to truthfully report their types.

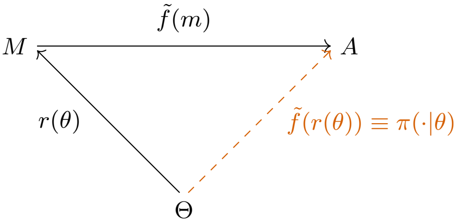

indirect mechanism

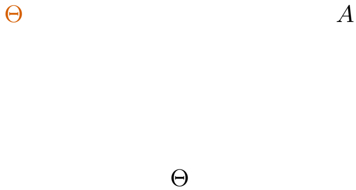

direct mechanism

Revelation principle: first-pass

Theorem (Gibbard, 1973; Myerson, 1979; Dasgupta et al, 1979)

There exists a game that has π as an equilibrium outcome if and only if the following game implements π:

1. *Actions $M_i = \Theta_i$*
2. *When players take actions $\theta' = (\theta'_1, \ldots, \theta'_N)$, the outcome is $f(\overline{\theta}) = \pi(\cdot|\theta')$.*

Furthermore, it is without loss of generality to assume that the players find it optimal to truthfully report their types.

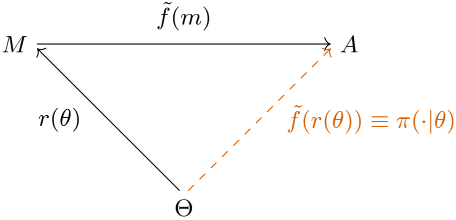

indirect mechanism

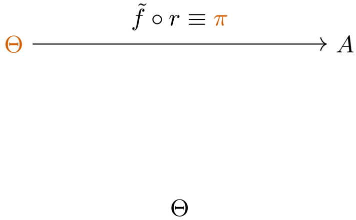

direct mechanism

Revelation principle: first-pass

Theorem (Gibbard, 1973; Myerson, 1979; Dasgupta et al, 1979)

There exists a game that has π as an equilibrium outcome if and only if the following game implements π:

1. *Actions* $M_i = \Theta_i$
2. *When players take actions* $\theta' = (\theta'_1, \ldots, \theta'_N)$, *the outcome is* $f(\bar{\theta}) = \pi(\cdot|\theta')$.

Furthermore, it is without loss of generality to assume that the players find it optimal to truthfully report their types.

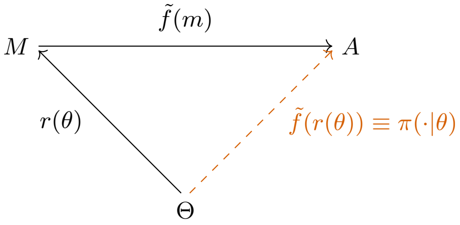

indirect mechanism

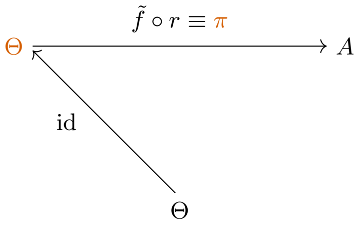

direct mechanism

The learning problem becomes *trivial*:

The learning problem becomes *trivial*:

- Abstract from the (maybe decentralized) learning process that ends up with the information that is ultimately an input in the designer's objective/rule

The learning problem becomes *trivial*:

- Abstract from the (maybe decentralized) learning process that ends up with the information that is ultimately an input in the designer's objective/rule
- Without loss, the owners of the information reveal it
  - Because it is in their best interest to do so!
  - they are compensated via rents and/or the allocation is distorted

The learning problem becomes *trivial*:

- Abstract from the (maybe decentralized) learning process that ends up with the information that is ultimately an input in the designer's objective/rule
- Without loss, the owners of the information reveal it
  - Because it is in their best interest to do so!
  - they are compensated via rents and/or the allocation is distorted

The RP does not mean that *all* mechanisms are truthful

Revelation principle: first-pass

The learning problem becomes *trivial*:

- Abstract from the (maybe decentralized) learning process that ends up with the information that is ultimately an input in the designer's objective/rule
- Without loss, the owners of the information reveal it
  - Because it is in their best interest to do so!
  - they are compensated via rents and/or the allocation is distorted

The RP does not mean that *all* mechanisms are truthful

- Many real world mechanisms are not truthful (e.g., first price auctions)
- not clear that truthful mechanisms are *better* (e.g., second price auctions) (c.f., Li, 2017, Akbarpour & Li, 2020)

Revelation principle: first pass

So why the obsession with the revelation principle?

- Truthful mechanisms are a good first cut abstraction,
- It is a recipe for constructing *algorithms* that implement allocations,
- It transforms an equilibrium problem into a *constrained* optimization problem.
- From the design perspective, if I cannot find a truthful mechanism that implements my desired rule then *no* mechanism does.

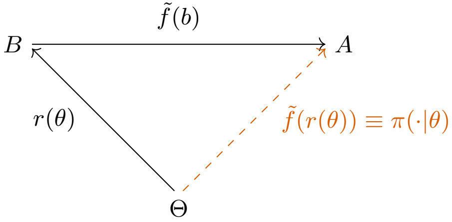

indirect mechanism

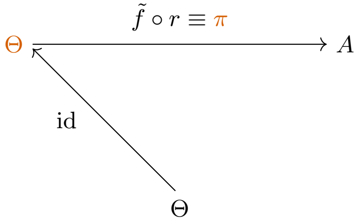

direct mechanism

Mechanism design in the wild

- Sponsored search auctions
- display advertising
- FCC spectrum auctions
- Kidney exchange
- Healthcare systems
- Recommendation systems
- Routing on the Internet
- Resource allocation in the cloud
- Platform design for a sharing economy
- Energy and electricity markets
- Bitcoin
- Participatory democracy
- Crowdsourcing

- Even closer to a data-driven decision process:
  - Repeated interactions
  - Persistent and/or evolving types

- Even closer to a data-driven decision process:
  - Repeated interactions
  - Persistent and/or evolving types
- (new) burgeoning area of dynamic mechanism design in Econ, CS, and OR
  - *internet auctions, government procurement, durable goods, regulation*

- Even closer to a data-driven decision process:
  - Repeated interactions
  - Persistent and/or evolving types
- (new) burgeoning area of dynamic mechanism design in Econ, CS, and OR
  - *internet auctions, government procurement, durable goods, regulation*
- The designer learns information that is relevant for today but also subsequent periods
  - *e.g., the optimal reserve price for today may not be optimal tomorrow*

- Even closer to a data-driven decision process:
  - Repeated interactions
  - Persistent and/or evolving types
- (new) burgeoning area of dynamic mechanism design in Econ, CS, and OR
  - *internet auctions, government procurement, durable goods, regulation*
- The designer learns information that is relevant for today but also subsequent periods
  - *e.g., the optimal reserve price for today may not be optimal tomorrow*
- what the designer learns today, they can use tomorrow: **ratchet effect**
  - *e.g., forward-looking bidders understand that bids today determine reserve prices tomorrow* $\Rightarrow$ *additional incentive to shave bids above and beyond the strategic and dynamic interaction*

Dynamic Mechanism Design

- If the designer has **full commitment** to the mechanism/algorithm, this is not an issue

Dynamic Mechanism Design

- If the designer has **full commitment** to the mechanism/algorithm, this is not an issue
- The revelation principle holds in dynamic environments (Myerson 1986):

Dynamic Mechanism Design

- If the designer has **full commitment** to the mechanism/algorithm, this is not an issue
- The revelation principle holds in dynamic environments (Myerson 1986):
  - Wlog, the designer asks the players to report their (new) private information,

Dynamic Mechanism Design

- If the designer has **full commitment** to the mechanism/algorithm, this is not an issue
- The revelation principle holds in dynamic environments (Myerson 1986):
  - Wlog, the designer asks the players to report their (new) private information,
  - Wlog, we focus on truthful equilibria* of this game.

Dynamic Mechanism Design

- If the designer has **full commitment** to the mechanism/algorithm, this is not an issue
- The revelation principle holds in dynamic environments (Myerson 1986):
  - Wlog, the designer asks the players to report their (new) private information,
  - Wlog, we focus on truthful equilibria* of this game.
- What changes are the determinants of the agents' rents: they account for today's private information and the impact of today's information in future decisions

Dynamic Mechanism Design

- If the designer has **full commitment** to the mechanism/algorithm, this is not an issue
- The revelation principle holds in dynamic environments (Myerson 1986):
  - Wlog, the designer asks the players to report their (new) private information,
  - Wlog, we focus on truthful equilibria* of this game.
- What changes are the determinants of the agents' rents: they account for today's private information and the impact of today's information in future decisions
- Sometimes these rents are large enough that optimal mechanisms do not use the information learned

Dynamic Mechanism Design

- If the designer has **full commitment** to the mechanism/algorithm, this is not an issue
- The revelation principle holds in dynamic environments (Myerson 1986):
  - Wlog, the designer asks the players to report their (new) private information,
  - Wlog, we focus on truthful equilibria* of this game.
- What changes are the determinants of the agents' rents: they account for today's private information and the impact of today's information in future decisions
- Sometimes these rents are large enough that optimal mechanisms do not use the information learned
  - *e.g., sale of a durable good*

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\,(1 - F_1(\hat{p}))\,.$$

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\left(1 - F_1(\hat{p})\right).$$

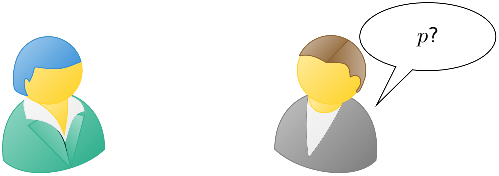

Period 1

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\,(1 - F_1(\hat{p}))\,.$$

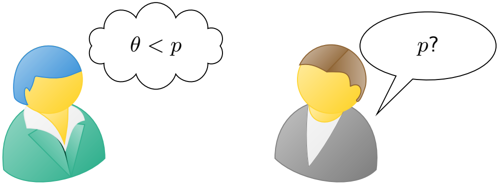

Period 1

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\,(1 - F_1(\hat{p}))\,.$$

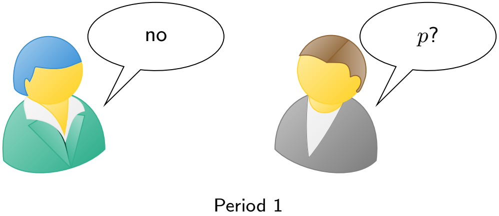

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\,(1 - F_1(\hat{p}))\,.$$

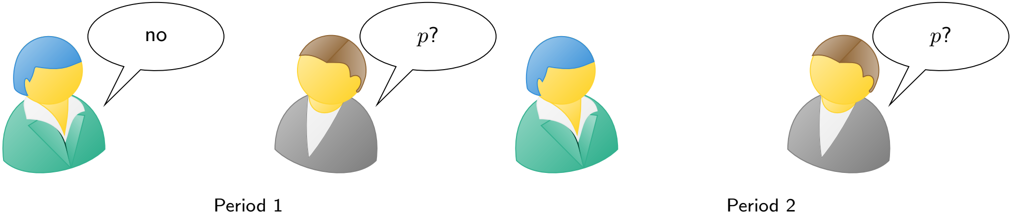

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\,(1 - F_1(\hat{p}))\,.$$

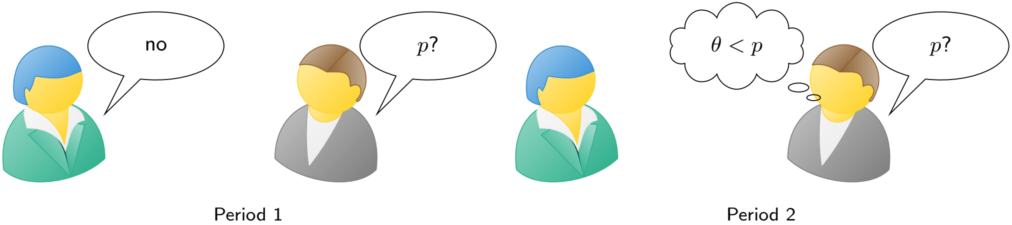

Dynamic Mechanism Design

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\,(1 - F_1(\hat{p}))\,.$$

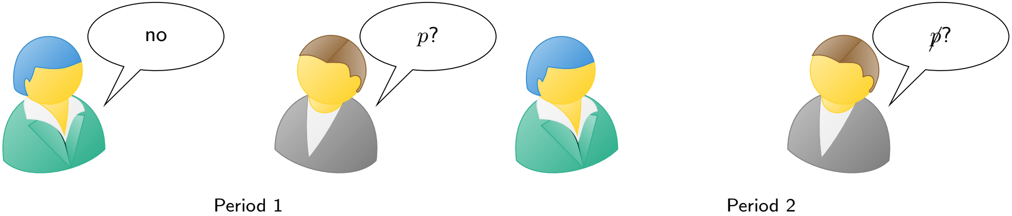

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\,(1 - F_1(\hat{p}))\,.$$

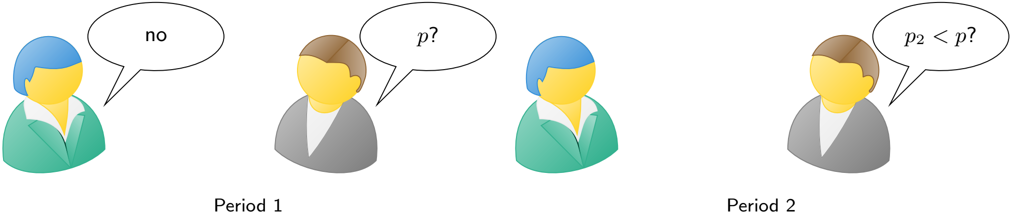

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\,(1 - F_1(\hat{p}))\,.$$

Period 1

Period 2

Consider the following example:

- Seller (designer) wants to sell **one** unit of a good to buyer over two periods
- Buyer's valuation for the good is private information $\theta \sim F_1$
- Revenue-maximizing mechanism is to set the same price p each period, where

$$p \in \arg\max_{\hat{p}} \hat{p}\,(1 - F_1(\hat{p}))\,.$$

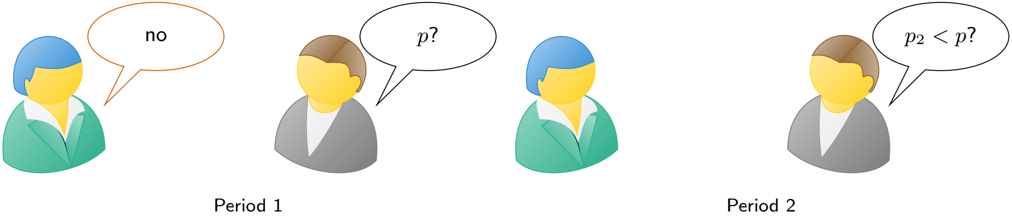

Oftentimes, optimal mechanisms are not *sequentially rational*:

- i.e., if we gave the designer the possibility to revise the decision rule given the new learned information, they would have an incentive to do so.

There are many examples with these features

- dynamic (ad) auctions (e.g., Google) (c.f., Kanoria & Nazerzadeh, 2014; Papadimitrou et al, 2014; Balseiro et al, 2022)
- repeated sales (e.g., Lobel & Paes Leme, 2017; Devanur et al, 2019; Immorlica et al, 2017)
- procurement (e.g., Gur et al, 2022)

Desiderata: a theory of mechanism design that does not rely so strongly on the assumption that the designer has full commitment

Commitment[C]

Instead, we would like to assume that the designer does not have full commitment

Commitment^C

Instead, we would like to assume that the designer does not have full commitment

The first issue is that the complement of commitment is way too big:

Instead, we would like to assume that the designer does not have full commitment

The first issue is that the complement of commitment is way too big:

- Short-term mechanisms that can be revised in each period (e.g., Freixas, Guesnerie, & Tirole, 1985, Laffont & Tirole, 1986, Hart & Tirole, 1988,. . .)

Instead, we would like to assume that the designer does not have full commitment

The first issue is that the complement of commitment is way too big:

- Short-term mechanisms that can be revised in each period (e.g., Freixas, Guesnerie, & Tirole, 1985, Laffont & Tirole, 1986, Hart & Tirole, 1988,. . .)
- Full-term mechanisms w/ 2-sided renegotiation (e.g., Hart & Tirole, 1988, Dewatripoint, 1989)

Commitment^C

Instead, we would like to assume that the designer does not have full commitment

The first issue is that the complement of commitment is way too big:

- Short-term mechanisms that can be revised in each period (e.g., Freixas, Guesnerie, & Tirole, 1985, Laffont & Tirole, 1986, Hart & Tirole, 1988,. . .)
- Full-term mechanisms w/ 2-sided renegotiation (e.g., Hart & Tirole, 1988, Dewatripoint, 1989)
- Full-term mechanisms w/ 1-sided renegotiation (e.g., Baron & Besanko, 1987)

Instead, we would like to assume that the designer does not have full commitment

The first issue is that the complement of commitment is way too big:

- Short-term mechanisms that can be revised in each period (e.g., Freixas, Guesnerie, & Tirole, 1985, Laffont & Tirole, 1986, Hart & Tirole, 1988,. . .)
- Full-term mechanisms w/ 2-sided renegotiation (e.g., Hart & Tirole, 1988, Dewatripoint, 1989)
- Full-term mechanisms w/ 1-sided renegotiation (e.g., Baron & Besanko, 1987)
- Long, but not full, term contract w/ renegotiation (e.g., Rey & Salanie, 1990)

Instead, we would like to assume that the designer does not have full commitment

The first issue is that the complement of commitment is way too big:

- Short-term mechanisms that can be revised in each period (e.g., Freixas, Guesnerie, & Tirole, 1985, Laffont & Tirole, 1986, Hart & Tirole, 1988,. . .)
- Full-term mechanisms w/ 2-sided renegotiation (e.g., Hart & Tirole, 1988, Dewatripoint, 1989)
- Full-term mechanisms w/ 1-sided renegotiation (e.g., Baron & Besanko, 1987)
- Long, but not full, term contract w/ renegotiation (e.g., Rey & Salanie, 1990)
- Cannot commit even to today's mechanism (e.g., Adams & Schwarz 2007, Vartianen 2013, Akbarpour & Li 2020)

Commitment[C]

Instead, we would like to assume that the designer does not have full commitment

The first issue is that the complement of commitment is way too big:

- Short-term mechanisms that can be revised in each period (e.g., Freixas, Guesnerie, & Tirole, 1985, Laffont & Tirole, 1986, Hart & Tirole, 1988,. . .)
- Full-term mechanisms w/ 2-sided renegotiation (e.g., Hart & Tirole, 1988, Dewatripoint, 1989)
- Full-term mechanisms w/ 1-sided renegotiation (e.g., Baron & Besanko, 1987)
- Long, but not full, term contract w/ renegotiation (e.g., Rey & Salanie, 1990)
- Cannot commit even to today's mechanism (e.g., Adams & Schwarz 2007, Vartianen 2013, Akbarpour & Li 2020)

Instead, we would like to assume that the designer does not have full commitment

The first issue is that the complement of commitment is way too big:

- Short-term mechanisms that can be revised in each period (e.g., Freixas, Guesnerie, & Tirole, 1985, Laffont & Tirole, 1986, Hart & Tirole, 1988,. . .)
- Full-term mechanisms w/ 2-sided renegotiation (e.g., Hart & Tirole, 1988, Dewatripoint, 1989)
- Full-term mechanisms w/ 1-sided renegotiation (e.g., Baron & Besanko, 1987)
- Long, but not full, term contract w/ renegotiation (e.g., Rey & Salanie, 1990)
- Cannot commit even to today's mechanism (e.g., Adams & Schwarz 2007, Vartianen 2013, Akbarpour & Li 2020)

Papers in CS & OR that study dynamic lack of commitment focus on this case as well: Papadimitrou et al, 2014; Lobel & Paes Leme, 2017; Devanur et al, 2019; Immorlica et al, 2017; Balseiro et al, 2022;Gur et al, 2022

Commitment to short-term mechanisms

Setting:

- Uninformed designer interacts with privately & persistently informed agent over time
- Designer can offer short-term mechanisms
- Designer can commit to today's mechanism, but not to the continuation ones.

Commitment to short-term mechanisms

Setting:

- Uninformed designer interacts with privately & persistently informed agent over time
- Designer can offer short-term mechanisms
- Designer can commit to today's mechanism, but not to the continuation ones.

Examples:

1. Regulation (c.f., Laffont & Tirole, 1988)
2. Procurement
3. Political Economy; e.g., taxation and social insurance,
4. Ad auctions, online shopping

Commitment to short-term mechanisms

Setting:

- Uninformed designer interacts with privately & persistently informed agent over time
- Designer can offer short-term mechanisms
- Designer can commit to today's mechanism, but not to the continuation ones.

Examples:

1. Regulation (c.f., Laffont & Tirole, 1988)
2. Procurement
3. Political Economy; e.g., taxation and social insurance,
4. Ad auctions, online shopping

Few papers analyze optimal mechanisms under limited commitment:

- Optimal mechanisms w/ finite horizon, e.g.,
  - Kumar (1985), Laffont & Tirole (1988), Bester and Strausz (2000,2001,2007), Hart & Tirole (1988), Skreta (2006,2015), Bisin & Rampini (2006), Deb & Said (2015), Fiocco & Strausz (2015), Beccutti & Möller (2018)
- Infinite Horizon under restrictions, e.g.,
  - Acharya & Ortner (2017), Gerardi & Maestri (2018)
  - **iid private information:** e.g., Sleet and Yeltekin (2006, 2008), Farhi, Sleet, Yeltekin, and Werning (2012), Golosov and Iovino (2021)

The second issue is that the revelation principle no longer holds under limited commitment:

The lack of commitment in repeated adverse-selection situations leads to substantial difficulties for contract theory.

Laffont & Tirole, 1993

- Substantial setback in terms of what we know about optimal policies under limited commitment.

The second issue is that the revelation principle no longer holds under limited commitment:

> *The lack of commitment in repeated adverse-selection situations leads to substantial difficulties for contract theory.*
>
> Laffont & Tirole, 1993

- Substantial setback in terms of what we know about optimal policies under limited commitment.

Revelation principle for mechanism design with limited commitment

We characterize a class of mechanisms and strategies that are enough to implement any outcome distribution that can be implemented under limited commitment.

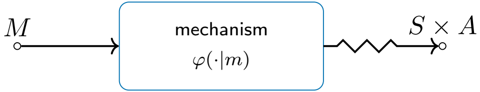

Mechanisms (Myerson ’82, Forges ’85)

- M is a set of input messages,
- S is a set of output messages,
- φ assigns to each input message a *joint* distribution over output messages and allocations

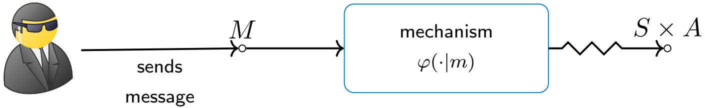

Mechanisms (Myerson '82, Forges '85)

- M is a set of input messages,
- S is a set of output messages,
- φ assigns to each input message a *joint* distribution over output messages and allocations

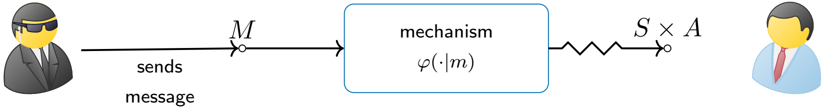

Mechanisms (Myerson '82, Forges '85)

- M is a set of input messages,
- S is a set of output messages,
- φ assigns to each input message a *joint* distribution over output messages and allocations

Mechanisms

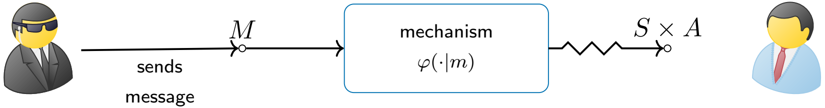

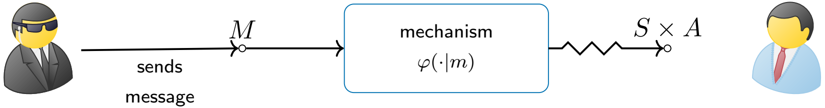

Revelation Principle under commitment

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Communication is **observable:** M and S have the same cardinality and φ is *invertible*.
- Equilibrium communication is **truthful**.

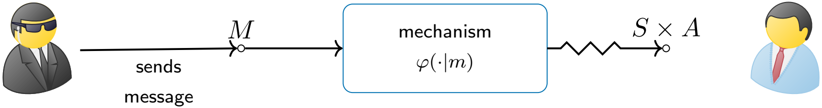

Limited Commitment 1: Bester & Strausz (ECMA, 2001)

Assume:

- Communication is **observable:** M and S have the same cardinality and φ is *invertible*,
- No randomization in the allocation, i.e., each output message is attached to one allocation.

What we know: Bester & Strausz (2001,2007)

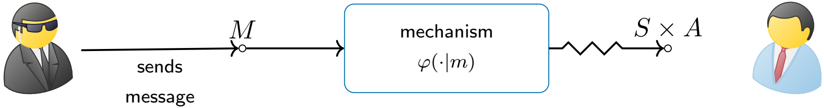

Limited Commitment 1: Bester & Strausz (ECMA, 2001)

Assume:

- Communication is **observable:** M and S have the same cardinality and φ is *invertible*,
- No randomization in the allocation, i.e., each output message is attached to one allocation.

Then, if the principal earns his highest payoff consistent with the agent's payoff, wlog

- Communication is **direct**, i.e., $M = \Theta$,

However, ~~Equilibrium communication is **truthful**~~. (c.f., Papadimitrou et al, 2014)

What we know: Bester & Strausz (2001,2007)

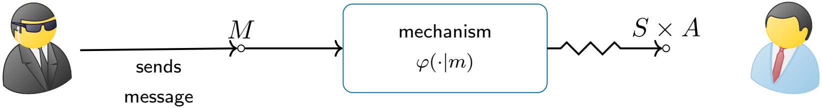

Limited Commitment 2: Bester & Strausz (JET, 2007)

Assume:

- ~~Communication is **observable**: M and S have the same cardinality and φ is invertible~~
- No randomization in the allocation, i.e., each output message is attached to one allocation.

What we know: Bester & Strausz (2001,2007)

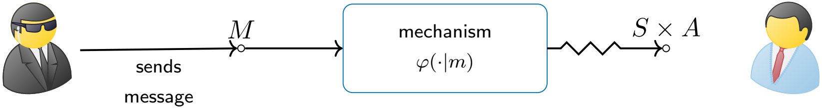

Limited Commitment 2: Bester & Strausz (JET, 2007)

Assume:

- ~~Communication is **observable:** M and S have the same cardinality and φ is invertible~~
- No randomization in the allocation, i.e., each output message is attached to one allocation.

Then, without loss of generality

- Communication is **direct**, i.e., $M = \Theta$,
- Equilibrium communication is **truthful**.

Revelation Principle for Limited Commitment

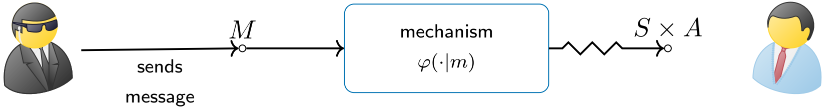

Revelation Principle for Limited Commitment

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
-
- Equilibrium communication is **truthful**

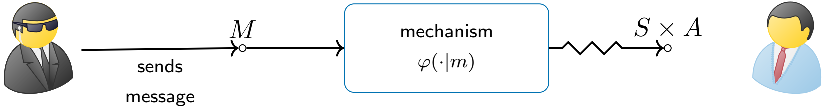

Revelation Principle for Limited Commitment

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**

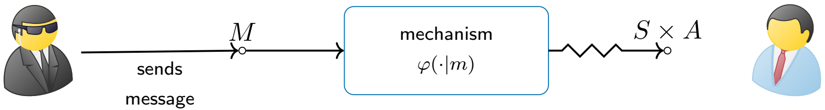

Revelation Principle for Limited Commitment

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**

Revelation Principle for Limited Commitment

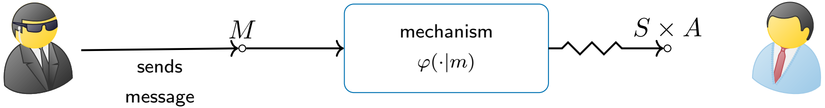

Revelation Principle for Limited Commitment

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A$$

Revelation Principle for Limited Commitment

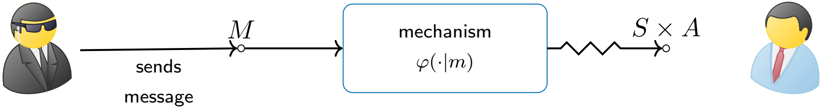

Revelation Principle for Limited Commitment

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms separate the design of the information from the design of the allocation

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A$$

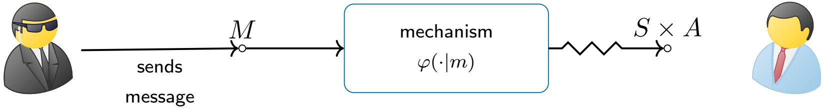

Revelation Principle for Limited Commitment

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms separate the design of the information from the design of the allocation

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\beta} \Delta(\Theta) \xrightarrow{\alpha} A$$

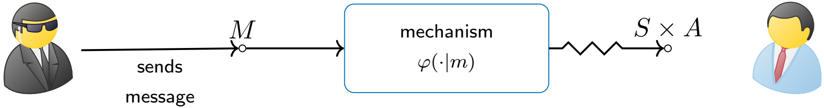

Revelation Principle for Limited Commitment

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms separate the design of the information from the design of the allocation

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\beta} \Delta(\Theta) \xrightarrow{\alpha} A$$

Direct-Blackwell mechanisms

Restoring the revelation principle: Allocations **and** information

In a dynamic setting, we need the mechanism to replicate

1. The type-by-type allocation distribution
2. **and** (at the very least) the type-by-type distribution of beliefs

Restoring the revelation principle: Allocations **and** information

In a dynamic setting, we need the mechanism to replicate

1. The type-by-type allocation distribution (mechanism design)
2. **and** (at the very least) the type-by-type distribution of beliefs (information design)

Restoring the revelation principle: Allocations and information

In a dynamic setting, we need the mechanism to replicate

1. The type-by-type allocation distribution (mechanism design)
2. **and** (at the very least) the type-by-type distribution of beliefs (information design)

- Instead of designing a mechanism for a *given* information structure

Restoring the revelation principle: Allocations and information

In a dynamic setting, we need the mechanism to replicate

1. The type-by-type allocation distribution (mechanism design)
2. **and** (at the very least) the type-by-type distribution of beliefs (information design)

- Instead of designing a mechanism for a *given* information structure
- **Direct Blackwell mechanisms** jointly design the mechanism and the information

Restoring the revelation principle: Allocations **and** information

In a dynamic setting, we need the mechanism to replicate

1. The type-by-type allocation distribution (mechanism design)
2. **and** (at the very least) the type-by-type distribution of beliefs (information design)

- Instead of designing a mechanism for a *given* information structure
- **Direct Blackwell mechanisms** jointly design the mechanism and the information

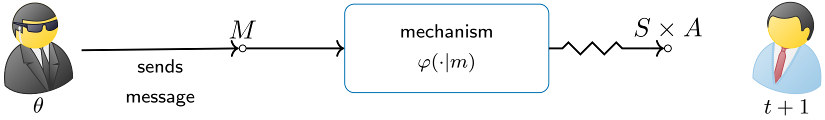

Restoring the revelation principle: Allocations **and** information

In a dynamic setting, we need the mechanism to replicate

1. The type-by-type allocation distribution (mechanism design)
2. **and** (at the very least) the type-by-type distribution of beliefs (information design)

- Instead of designing a mechanism for a *given* information structure
- **Direct Blackwell mechanisms** jointly design the mechanism and the information

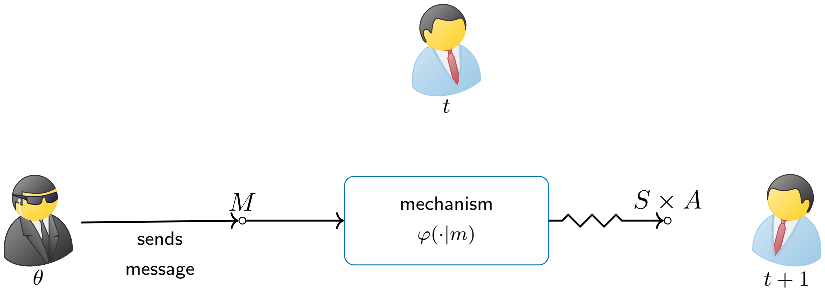

Restoring the revelation principle: Allocations **and** information

In a dynamic setting, we need the mechanism to replicate

1. The type-by-type allocation distribution (mechanism design)
2. **and** (at the very least) the type-by-type distribution of beliefs (information design)

- Instead of designing a mechanism for a *given* information structure
- **Direct Blackwell mechanisms** jointly design the mechanism and the information

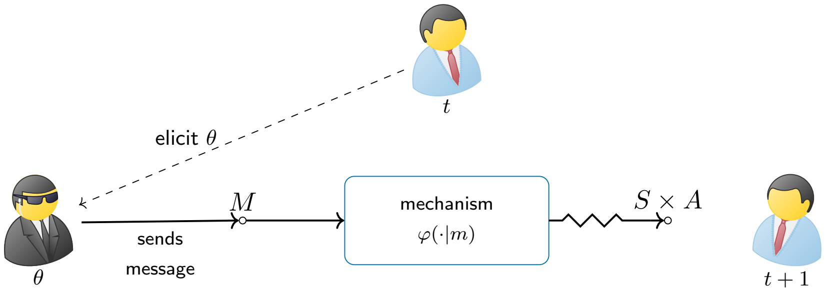

Restoring the revelation principle: Allocations **and** information

In a dynamic setting, we need the mechanism to replicate

1. The type-by-type allocation distribution (mechanism design)
2. **and** (at the very least) the type-by-type distribution of beliefs (information design)

- Instead of designing a mechanism for a *given* information structure
- **Direct Blackwell mechanisms** jointly design the mechanism and the information

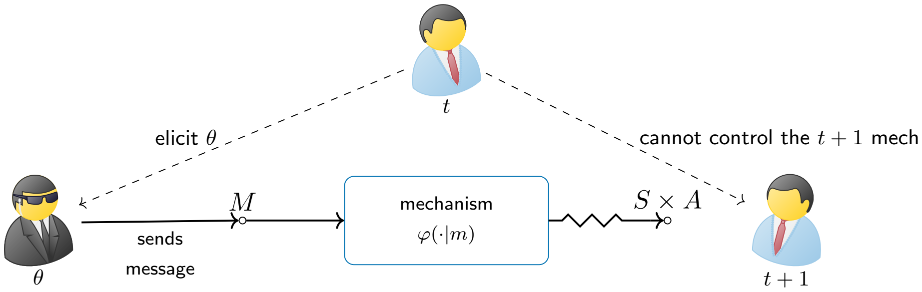

Restoring the revelation principle: Allocations **and** information

In a dynamic setting, we need the mechanism to replicate

1. The type-by-type allocation distribution (mechanism design)
2. **and** (at the very least) the type-by-type distribution of beliefs (information design)

- Instead of designing a mechanism for a *given* information structure
- **Direct Blackwell mechanisms** jointly design the mechanism and the information

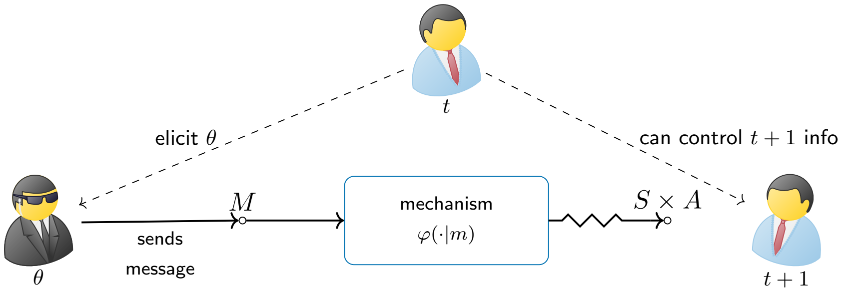

Revelation Principle for Limited Commitment

Revelation Principle for Limited Commitment (Doval & Skreta, 2021)

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms are **Direct Blackwell mechanisms**

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\beta} \Delta(\Theta) \xrightarrow{\alpha} A$$

Revelation Principle for Limited Commitment

Revelation Principle for Limited Commitment (Doval & Skreta, 2021)

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms are **Direct Blackwell mechanisms**

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\beta} \Delta(\Theta) \xrightarrow{\alpha} A$$

- Like the standard revelation principle, it reduces the agent's behavior and its impact on the principal's beliefs to a series of constraints the mechanism must satisfy:

Revelation Principle for Limited Commitment

Revelation Principle for Limited Commitment (Doval & Skreta, 2021)

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms are **Direct Blackwell mechanisms**

$$\Theta \xrightarrow{\quad\varphi\quad} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\quad\beta\quad} \Delta(\Theta) \xrightarrow{\quad\alpha\quad} A$$

- Like the standard revelation principle, it reduces the agent's behavior and its impact on the principal's beliefs to a series of constraints the mechanism must satisfy:
  - Truthtelling + participation + Bayes' plausibility constraint (designer's sequential rationality)

Revelation Principle for Limited Commitment

Revelation Principle for Limited Commitment (Doval & Skreta, 2021)

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms are **Direct Blackwell mechanisms**

$$\Theta \xrightarrow{\quad\varphi\quad} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\ \beta\ } \Delta(\Theta) \xrightarrow{\ \alpha\ } A$$

Revelation Principle for Limited Commitment

Revelation Principle for Limited Commitment (Doval & Skreta, 2021)

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms are **Direct Blackwell mechanisms**

$$\Theta \xrightarrow{\quad\varphi\quad} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\quad\beta\quad} \Delta(\Theta) \xrightarrow{\quad\alpha\quad} A$$

- New applications facilitated by the generality of the framework

Revelation Principle for Limited Commitment

Revelation Principle for Limited Commitment (Doval & Skreta, 2021)

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms are **Direct Blackwell mechanisms**

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\beta} \Delta(\Theta) \xrightarrow{\alpha} A$$

- New applications facilitated by the generality of the framework
  - no restrictions on the cardinality of Θ, on the length of the interaction + extension to Markov settings

Revelation Principle for Limited Commitment

Revelation Principle for Limited Commitment (Doval & Skreta, 2021)

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms are **Direct Blackwell mechanisms**

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\beta} \Delta(\Theta) \xrightarrow{\alpha} A$$

- New applications facilitated by the generality of the framework
  - no restrictions on the cardinality of Θ, on the length of the interaction + extension to Markov settings
  - Optimality of posted prices in infinite horizon-binary type durable goods model

Revelation Principle for Limited Commitment

Revelation Principle for Limited Commitment (Doval & Skreta, 2021)

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms are **Direct Blackwell mechanisms**

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\beta} \Delta(\Theta) \xrightarrow{\alpha} A$$

- New applications facilitated by the generality of the framework
  - no restrictions on the cardinality of Θ, on the length of the interaction + extension to Markov settings
  - Optimality of posted prices in infinite horizon-binary type durable goods model
  - Optimality of coarse product lines (menus) when purchase history leads to price discrimination

Revelation Principle for Limited Commitment

Revelation Principle for Limited Commitment (Doval & Skreta, 2021)

Without loss of generality,

- Communication is **direct**, i.e., $M = \Theta$.
- Output messages are **beliefs**, i.e., $S = \Delta(\Theta)$
- Equilibrium communication is **truthful**
- Equilibrium output messages coincide with **principal's equilibrium beliefs**
- Equilibrium mechanisms are **Direct Blackwell mechanisms**

$$\Theta \xrightarrow{\varphi} \Delta(\Theta) \times A \qquad = \qquad \Theta \xrightarrow{\beta} \Delta(\Theta) \xrightarrow{\alpha} A$$

- New applications facilitated by the generality of the framework
  - no restrictions on the cardinality of Θ, on the length of the interaction + extension to Markov settings
  - Optimality of posted prices in infinite horizon-binary type durable goods model
  - Optimality of coarse product lines (menus) when purchase history leads to price discrimination
- **Today:** Revisit the sale of a durable good w/ a continuum of types and finite horizon

Two final remarks

Two other reasons to care about MDLC in the context of DDDP and AGT:

1. Simplicity
2. Learning

Simple mechanisms

- Limiting the principal's commitment was also an attempt to justify simple mechanisms,
- ... the idea being that it would force the principal to condition his mechanism on less variables (e.g., non-clairvoyant mechanisms, Balseiro et al, 2022)
- It turns out that the optimal mechanism is not necessarily "simpler"
  - e.g., posted prices may no longer be optimal to sell durable goods in finite horizon settings,

Learning mechanisms

- Platforms use **learning algorithms** to optimize on prices/reserve prices based on historical data
  (c.f., Kanoria & Nazerzadeh, 2014; Haghtalab, Lykouris, Nietert,& Wei, 2022)

Learning mechanisms

- Platforms use **learning algorithms** to optimize on prices/reserve prices based on historical data

  (c.f., Kanoria & Nazerzadeh, 2014; Haghtalab, Lykouris, Nietert,& Wei, 2022)
- Not the same as having limited commitment

Learning mechanisms

- Platforms use **learning algorithms** to optimize on prices/reserve prices based on historical data

  (c.f., Kanoria & Nazerzadeh, 2014; Haghtalab, Lykouris, Nietert,& Wei, 2022)
- Not the same as having limited commitment
  - e.g., Amazon commits to its algorithm and how it outputs decisions as a function of the inputs

Learning mechanisms

- Platforms use **learning algorithms** to optimize on prices/reserve prices based on historical data (c.f., Kanoria & Nazerzadeh, 2014; Haghtalab, Lykouris, Nietert,& Wei, 2022)
- Not the same as having limited commitment
  - e.g., Amazon commits to its algorithm and how it outputs decisions as a function of the inputs
- And yet, these algorithms will do "the best" with the information collected so far according to some objective function

Learning mechanisms

- Platforms use **learning algorithms** to optimize on prices/reserve prices based on historical data (c.f., Kanoria & Nazerzadeh, 2014; Haghtalab, Lykouris, Nietert,& Wei, 2022)
- Not the same as having limited commitment
  - e.g., Amazon commits to its algorithm and how it outputs decisions as a function of the inputs
- And yet, these algorithms will do "the best" with the information collected so far according to some objective function
- This may lead to **strategic overfitting:** forward looking agents will have additional incentives to strategize

Learning mechanisms

- Platforms use **learning algorithms** to optimize on prices/reserve prices based on historical data (c.f., Kanoria & Nazerzadeh, 2014; Haghtalab, Lykouris, Nietert,& Wei, 2022)
- Not the same as having limited commitment
  - e.g., Amazon commits to its algorithm and how it outputs decisions as a function of the inputs
- And yet, these algorithms will do "the best" with the information collected so far according to some objective function
- This may lead to **strategic overfitting:** forward looking agents will have additional incentives to strategize
- The representation we obtain is very relevant

Learning mechanisms

- Platforms use **learning algorithms** to optimize on prices/reserve prices based on historical data (c.f., Kanoria & Nazerzadeh, 2014; Haghtalab, Lykouris, Nietert,& Wei, 2022)
- Not the same as having limited commitment
  - e.g., Amazon commits to its algorithm and how it outputs decisions as a function of the inputs
- And yet, these algorithms will do "the best" with the information collected so far according to some objective function
- This may lead to **strategic overfitting:** forward looking agents will have additional incentives to strategize
- The representation we obtain is very relevant
  - the algorithm takes the role of the "sequentially rational principal".

Learning mechanisms

- Platforms use **learning algorithms** to optimize on prices/reserve prices based on historical data (c.f., Kanoria & Nazerzadeh, 2014; Haghtalab, Lykouris, Nietert,& Wei, 2022)
- Not the same as having limited commitment
  - e.g., Amazon commits to its algorithm and how it outputs decisions as a function of the inputs
- And yet, these algorithms will do "the best" with the information collected so far according to some objective function
- This may lead to **strategic overfitting:** forward looking agents will have additional incentives to strategize
- The representation we obtain is very relevant
  - the algorithm takes the role of the "sequentially rational principal".
- Our result provides a way of representing these *Bayesian* algorithms and the outcomes that can arise from the strategic interaction with a forward looking agent.

Learning mechanisms

- Platforms use **learning algorithms** to optimize on prices/reserve prices based on historical data (c.f., Kanoria & Nazerzadeh, 2014; Haghtalab, Lykouris, Nietert,& Wei, 2022)
- Not the same as having limited commitment
  - e.g., Amazon commits to its algorithm and how it outputs decisions as a function of the inputs
- And yet, these algorithms will do "the best" with the information collected so far according to some objective function
- This may lead to **strategic overfitting:** forward looking agents will have additional incentives to strategize
- The representation we obtain is very relevant
  - the algorithm takes the role of the "sequentially rational principal".
- Our result provides a way of representing these *Bayesian* algorithms and the outcomes that can arise from the strategic interaction with a forward looking agent.
- The analyst is forced to jointly describe the way information is stored and how it is used to determine the allocation.

Sale of a durable good: binary types and two periods

Sale of a durable good

- A seller and a buyer interact over two periods.

Sale of a durable good

- A seller and a buyer interact over two periods.
- The seller owns one unit of a durable good and assigns value 0 to it.

Sale of a durable good

- A seller and a buyer interact over two periods.
- The seller owns one unit of a durable good and assigns value 0 to it.
- The buyer has private information indexed by $\theta \in \Theta \equiv \{\theta_L, \theta_H\}$ and $\mu_0 = \Pr(\theta = \theta_H)$

Sale of a durable good

- A seller and a buyer interact over two periods.
- The seller owns one unit of a durable good and assigns value 0 to it.
- The buyer has private information indexed by $\theta \in \Theta \equiv \{\theta_L, \theta_H\}$ and $\mu_0 = \Pr(\theta = \theta_H)$
- An allocation is a pair $(q, x) \in \{0, 1\} \times \mathbb{R}$,

Sale of a durable good

- A seller and a buyer interact over two periods.
- The seller owns one unit of a durable good and assigns value 0 to it.
- The buyer has private information indexed by $\theta \in \Theta \equiv \{\theta_L, \theta_H\}$ and $\mu_0 = \Pr(\theta = \theta_H)$
- An allocation is a pair $(q, x) \in \{0, 1\} \times \mathbb{R}$,
  - q indicates whether the good is sold $(q = 1)$ or not $(q = 0)$, and

Sale of a durable good

- A seller and a buyer interact over two periods.
- The seller owns one unit of a durable good and assigns value 0 to it.
- The buyer has private information indexed by $\theta \in \Theta \equiv \{\theta_L, \theta_H\}$ and $\mu_0 = \Pr(\theta = \theta_H)$
- An allocation is a pair $(q, x) \in \{0, 1\} \times \mathbb{R}$,
  - q indicates whether the good is sold $(q = 1)$ or not $(q = 0)$, and
  - x is a payment from the buyer to the seller.

Sale of a durable good

- A seller and a buyer interact over two periods.
- The seller owns one unit of a durable good and assigns value 0 to it.
- The buyer has private information indexed by $\theta \in \Theta \equiv \{\theta_L, \theta_H\}$ and $\mu_0 = \Pr(\theta = \theta_H)$
- An allocation is a pair $(q, x) \in \{0, 1\} \times \mathbb{R}$,
  - q indicates whether the good is sold $(q = 1)$ or not $(q = 0)$, and
  - x is a payment from the buyer to the seller.
- If the good is sold in the first period, the game ends.

Sale of a durable good

- A seller and a buyer interact over two periods.
- The seller owns one unit of a durable good and assigns value 0 to it.
- The buyer has private information indexed by $\theta \in \Theta \equiv \{\theta_L, \theta_H\}$ and $\mu_0 = \Pr(\theta = \theta_H)$
- An allocation is a pair $(q, x) \in \{0,1\} \times \mathbb{R}$,
  - q indicates whether the good is sold ($q = 1$) or not ($q = 0$), and
  - x is a payment from the buyer to the seller.
- If the good is sold in the first period, the game ends.
- If the final allocation is $\{(q_t, x_t)\}_{t \in \{1,2\}}$, buyer and seller's payoffs are

$$U(\cdot, \theta) = \sum_{t=1}^{2} \delta^{t-1} \left(q_t \theta - x_t\right) \text{ and } W(\cdot, \theta) = \sum_{t=1}^{2} \delta^{t-1} x_t$$

where $\delta \in (0,1)$ is a common discount factor.

Timing: At the beginning of each period $t \in \{1, 2\}$

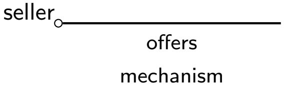

Timing: At the beginning of each period $t \in \{1, 2\}$

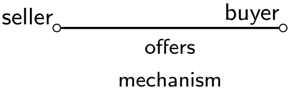

Timing: At the beginning of each period $t \in \{1, 2\}$

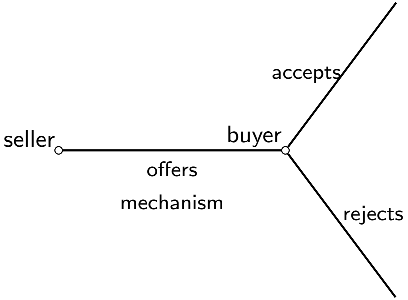

Timing: At the beginning of each period $t \in \{1, 2\}$

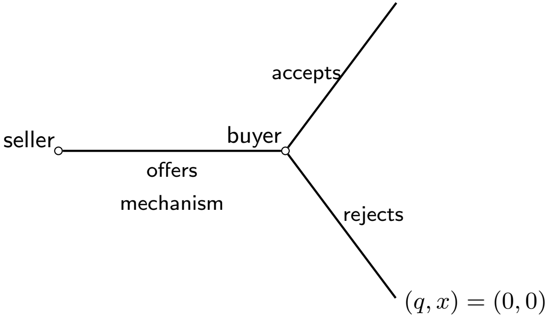

Timing: At the beginning of each period $t \in \{1, 2\}$

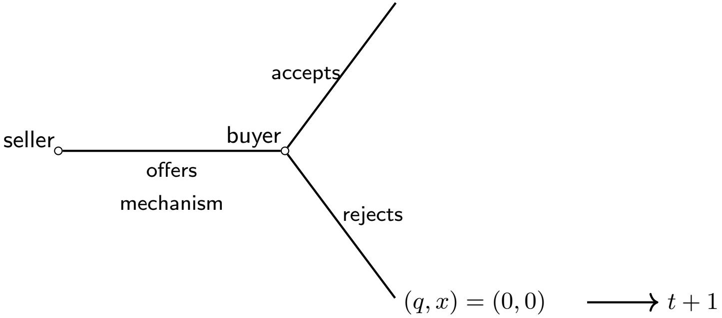

Timing: At the beginning of each period $t \in \{1, 2\}$

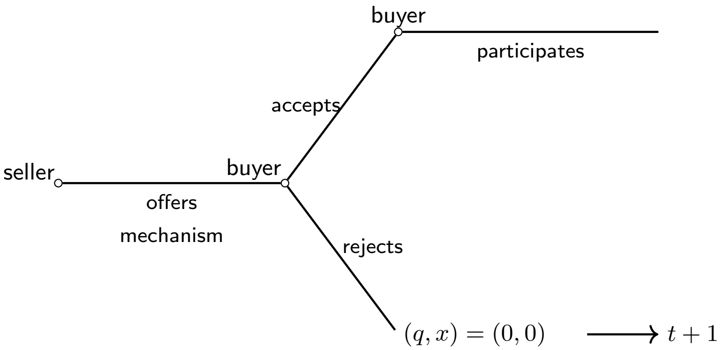

Timing: At the beginning of each period $t \in \{1, 2\}$

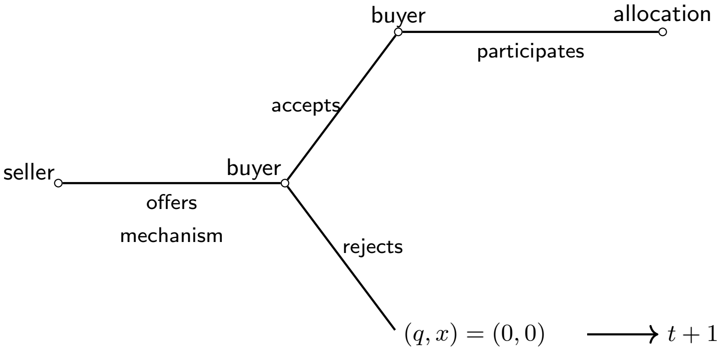

Timing: At the beginning of each period $t \in \{1, 2\}$

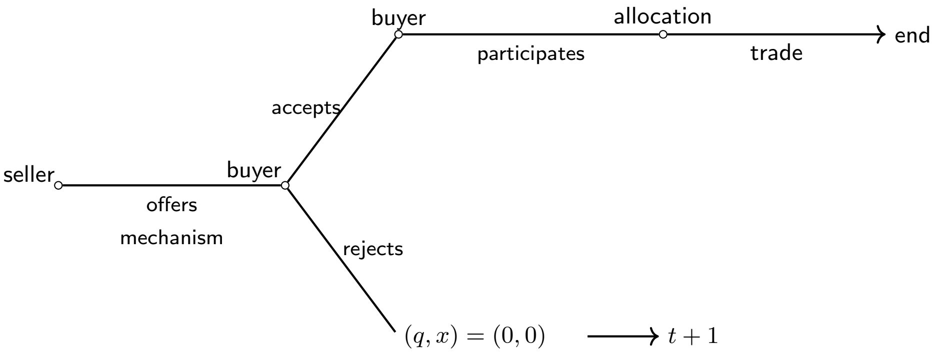

Timing: At the beginning of each period $t \in \{1, 2\}$

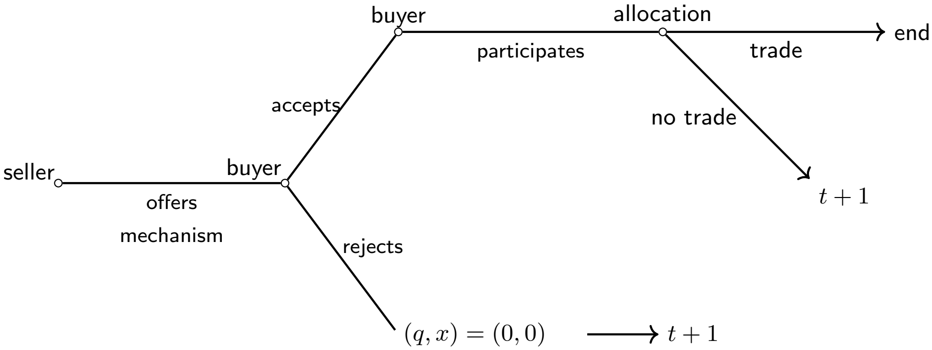

Sale of a durable good

- Final period: seller has full commitment. (Standard) Revelation principle applies.
- Let μ_2 denote the seller's belief that $\theta = \theta_H$.

Sale of a durable good

- Final period: seller has full commitment. (Standard) Revelation principle applies.
- Let μ_2 denote the seller's belief that $\theta = \theta_H$.
- The optimal mechanism is as follows:

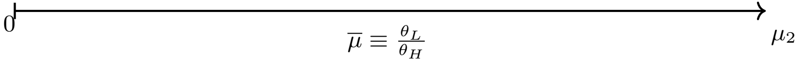

Sale of a durable good

- Final period: seller has full commitment. (Standard) Revelation principle applies.
- Let μ_2 denote the seller's belief that $\theta = \theta_H$.
- The optimal mechanism is as follows:

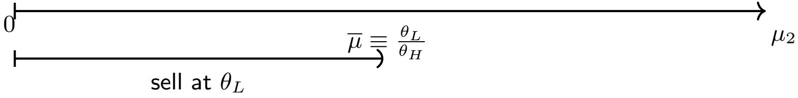

Sale of a durable good

- Final period: seller has full commitment. (Standard) Revelation principle applies.
- Let μ_2 denote the seller's belief that $\theta = \theta_H$.
- The optimal mechanism is as follows:

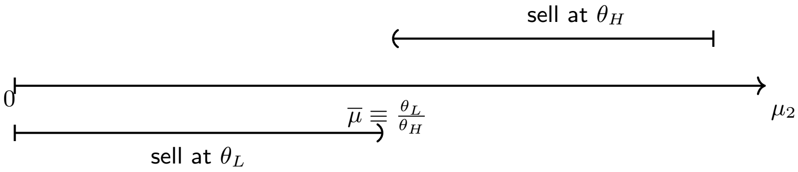

Sale of a durable good

- Final period: seller has full commitment. (Standard) Revelation principle applies.
- Let μ_2 denote the seller's belief that $\theta = \theta_H$.
- The optimal mechanism is as follows:

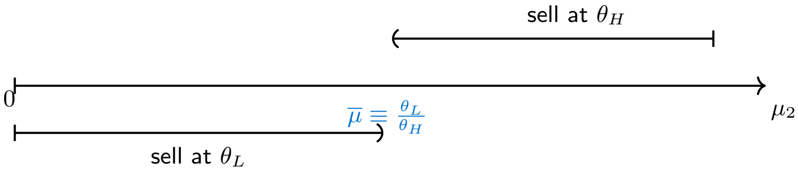

Sale of a durable good

- Final period: seller has full commitment. (Standard) Revelation principle applies.
- Let μ_2 denote the seller's belief that $\theta = \theta_H$.
- The optimal mechanism is as follows:

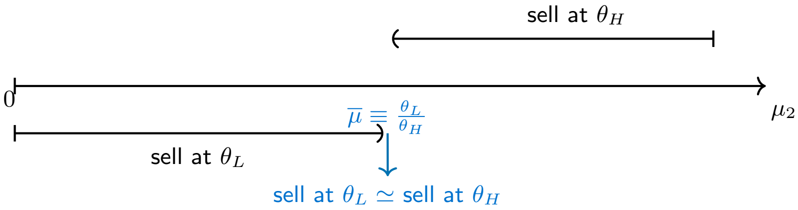

- Final period: seller has full commitment. (Standard) Revelation principle applies.
- Let μ_2 denote the seller's belief that $\theta = \theta_H$.
- The optimal mechanism is as follows:

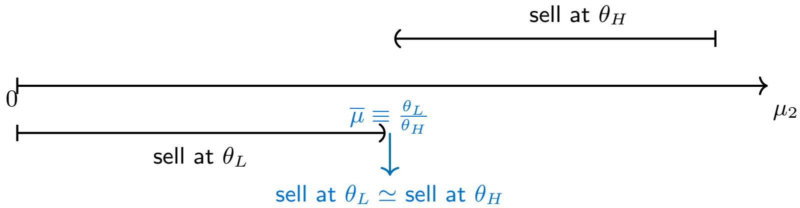

- Why $\overline{\mu}$? Whenever the seller sells to both types, he leaves rents $\mu_2 \Delta\theta$ to θ_H.

$$\theta_L = \mu_2(\theta_H - \Delta\theta) + (1-\mu_2)\theta_L = \mu_2\theta_H + (1-\mu_2)(\theta_L - \frac{\mu_2}{1-\mu_2}\Delta\theta)$$
$$= \mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_2)$$

Sale of a durable good

- Final period: seller has full commitment. (Standard) Revelation principle applies.
- Let μ_2 denote the seller's belief that $\theta = \theta_H$.
- The optimal mechanism is as follows:

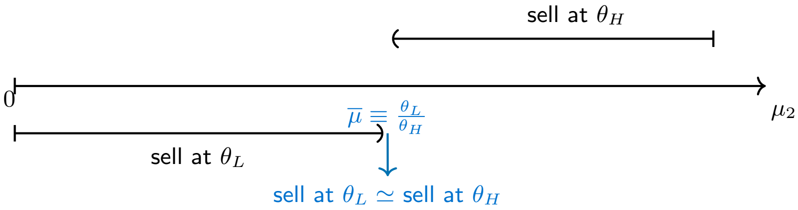

- Why $\overline{\mu}$? Whenever the seller sells to both types, he leaves rents $\mu_2 \Delta\theta$ to θ_H.

$$\theta_L = \mu_2(\theta_H - \Delta\theta) + (1-\mu_2)\theta_L = \mu_2\theta_H + (1-\mu_2)(\theta_L - \frac{\mu_2}{1-\mu_2}\Delta\theta)$$
$$= \mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_2)$$

When $\mu_2 = \overline{\mu}$, then $\hat{\theta}_L(\mu_2) = 0$.

Wrapping up:

$$R_2(\mu_2) = \begin{cases} \theta_L & \text{if } \mu_2 \leq \bar{\mu} \\ \mu_2\theta_H & \text{if } \mu_2 > \bar{\mu} \end{cases} = \begin{cases} \mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_2) & \text{if } \mu_2 \leq \bar{\mu} \\ \mu_2\theta_H & \text{if } \mu_2 > \bar{\mu} \end{cases}$$

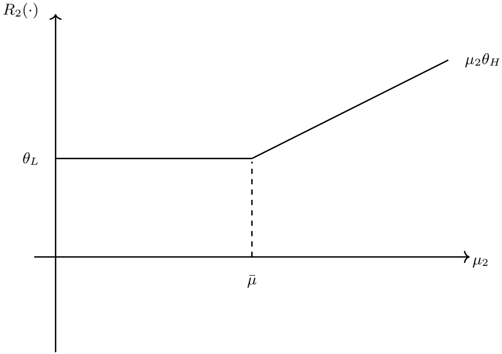

Seller's payoff in period 2

Sale of a durable good

- Recall μ_1 is the prior probability that $\theta = \theta_H$.

▸ $t = 2$

- Recall μ_1 is the prior probability that $\theta = \theta_H$.
- A mechanism is a tuple

$$M \xrightarrow{\varphi(\cdot|m)} S \times A$$

  - M is the set of input messages
  - S is the set of output messages
  - $\varphi : M \mapsto \Delta(\mathsf{S} \times A)$

▸ $t = 2$

- Recall μ_1 is the prior probability that $\theta = \theta_H$.
- A mechanism is a tuple

$$M \xrightarrow{\varphi(\cdot|m)} S \times A$$

  - M is the set of input messages
  - S is the set of output messages
  - $\varphi : M \mapsto \Delta(\mathsf{S} \times A)$ (finite support)– without loss with finitely many types

▸ $t = 2$

- Recall μ_1 is the prior probability that $\theta = \theta_H$.
- A mechanism is a tuple

$$\text{Buyer} \xrightarrow{\text{sends a message}} M \xrightarrow{\varphi(\cdot|m)} S \times A$$

- M is the set of input messages
- S is the set of output messages
- $\varphi : M \mapsto \Delta(\mathsf{S} \times A)$ (finite support)– without loss with finitely many types

$t = 2$

- Recall μ_1 is the prior probability that $\theta = \theta_H$.
- A mechanism is a tuple

$$\text{Buyer} \xrightarrow{\text{sends a message}} M \xrightarrow{\varphi(\cdot|m)} \overbrace{S \times A}^{\text{Seller}}$$

  - M is the set of input messages
  - S is the set of output messages
  - $\varphi : M \mapsto \Delta(\mathsf{S} \times A)$ (finite support)– without loss with finitely many types

▸ $t = 2$

Input messages

Some simplifications:

Input messages

Some simplifications:

- $M_1 = ?$

Some simplifications:

- $M_1 =?$

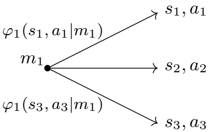

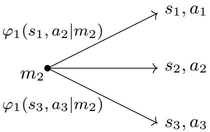

Some simplifications:

- $M_1 =?$

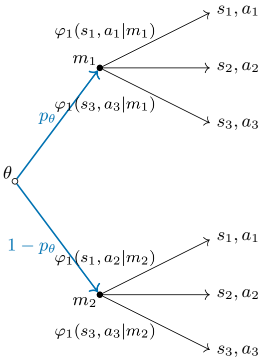

Input messages

Some simplifications:

- $M_1 = ?$

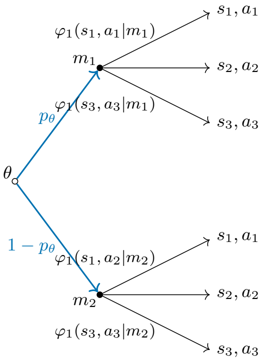

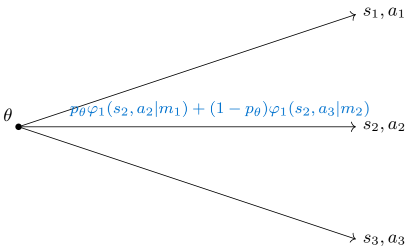

Some simplifications:

- $M_1 = ?$

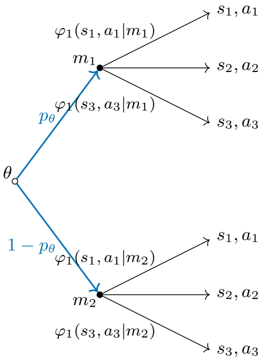

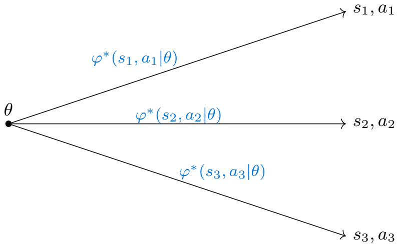

Some simplifications:

- $M_1 = \Theta$

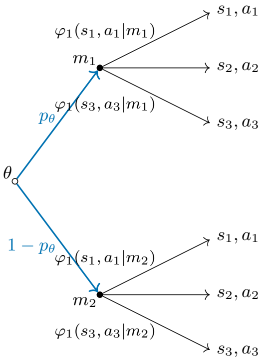

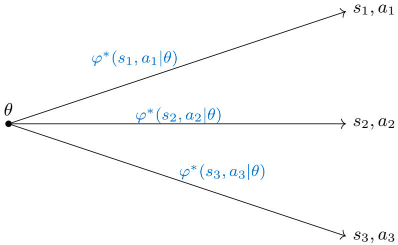

Input messages

Some simplifications:

- $M_1 = \Theta$

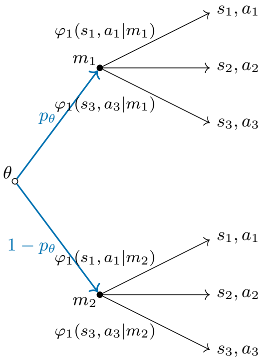

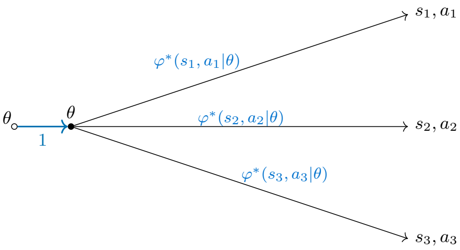

Input messages

Some simplifications:

- $M_1 = \Theta$

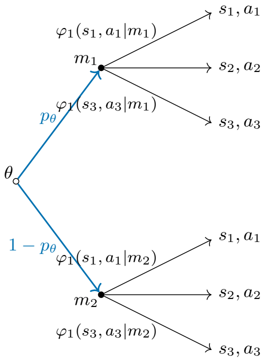

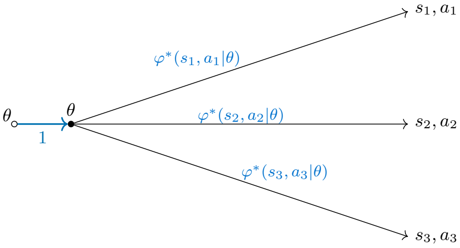

- Bester and Strausz (JET, 2007)

Output messages

$S_1 =?$

Output messages

$S_1 = ?$

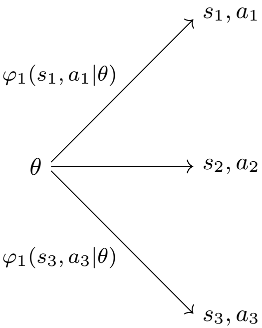

Output messages

period 1

$S_1 = ?$

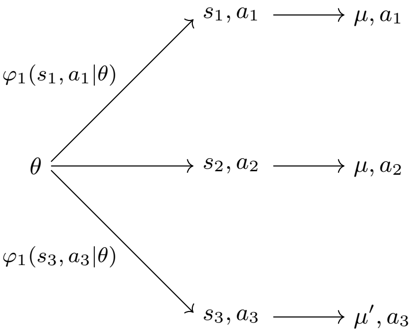

Output messages

$S_1 = ?$

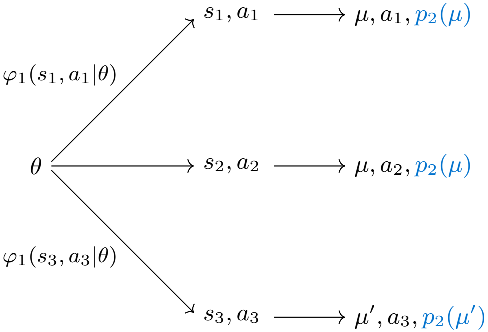

Output messages

$S_1 = ?$

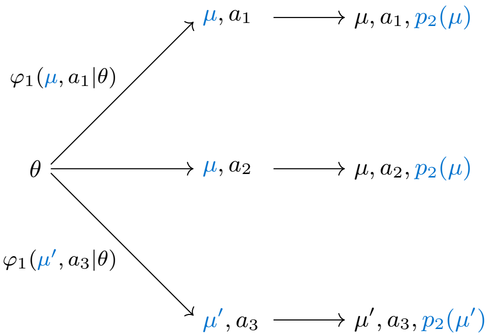

Output messages

$S_1 = \Delta(\Theta)$

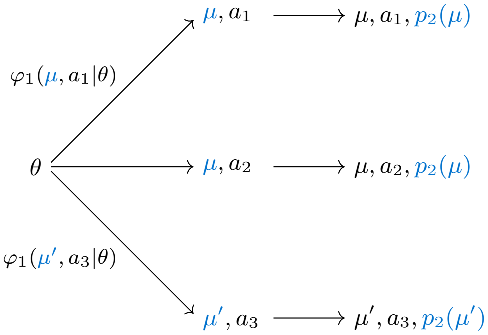

Separating information and allocation design

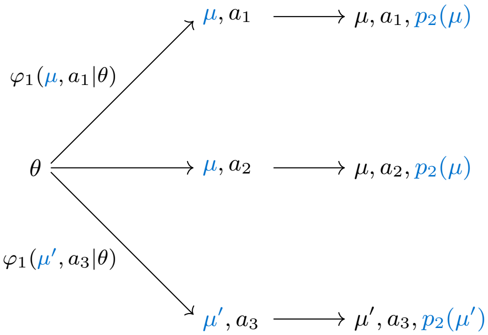

Separating information and allocation design

period 1

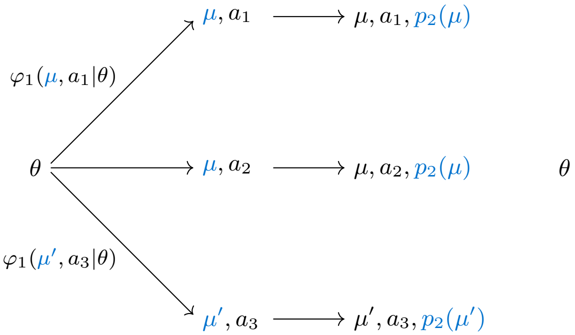

Separating information and allocation design

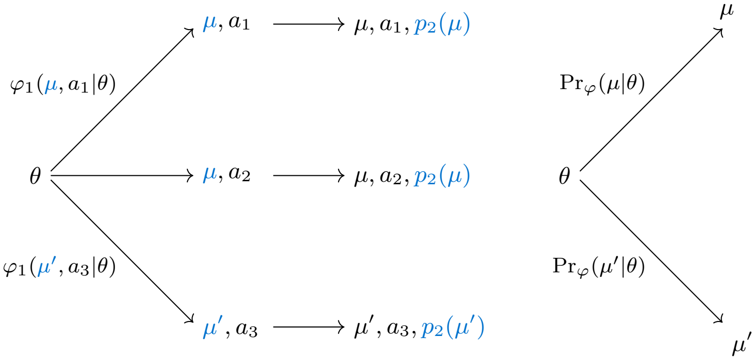

- $\mathrm{Pr}_\varphi(\mu|\theta) = \varphi_1(\mu, a_1|\theta) + \varphi_1(\mu, a_2|\theta)$

Separating information and allocation design

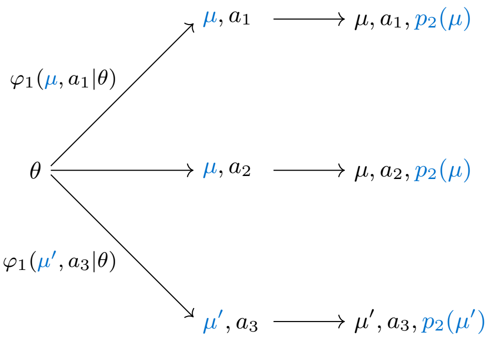

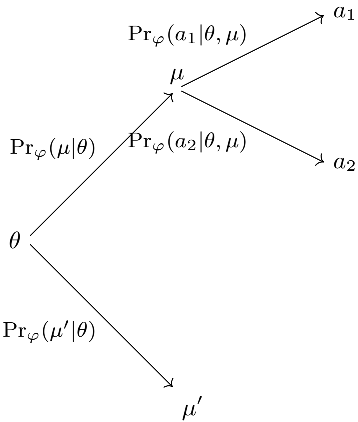

- $\Pr_\varphi(\mu|\theta) = \varphi_1(\mu, a_1|\theta) + \varphi_1(\mu, a_2|\theta)$

Separating information and allocation design

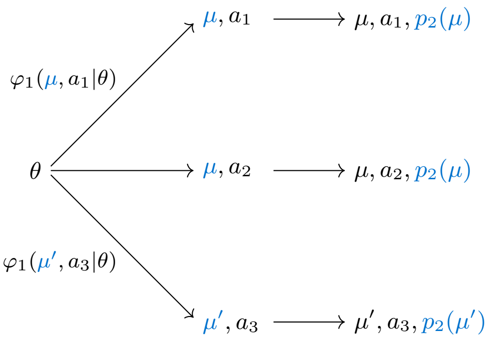

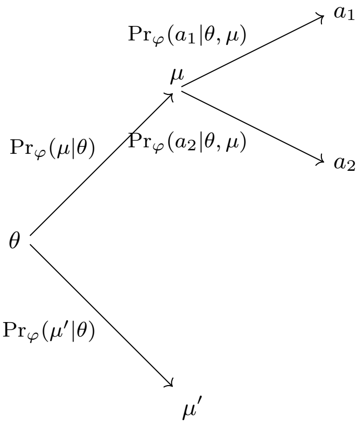

- $\Pr_\varphi(\mu|\theta) = \varphi_1(\mu, a_1|\theta) + \varphi_1(\mu, a_2|\theta)$
- $\Pr_\varphi(a_1|\theta, \mu) = \varphi_1(\mu, a_1|\theta)/Pr_\varphi(\mu|\theta)$

Separating information and allocation design

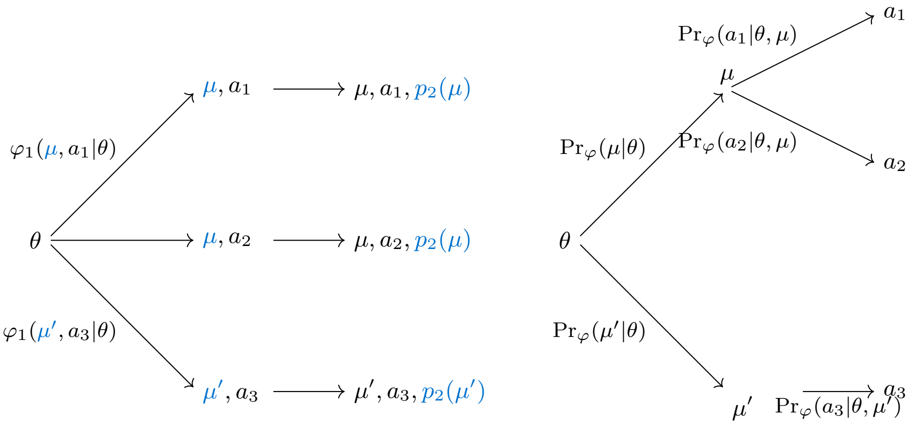

- $\mathrm{Pr}_\varphi(\mu|\theta) = \varphi_1(\mu, a_1|\theta) + \varphi_1(\mu, a_2|\theta)$
- $\mathrm{Pr}_\varphi(a_1|\theta, \mu) = \varphi_1(\mu, a_1|\theta)/Pr_\varphi(\mu|\theta)$

Separating information and allocation design

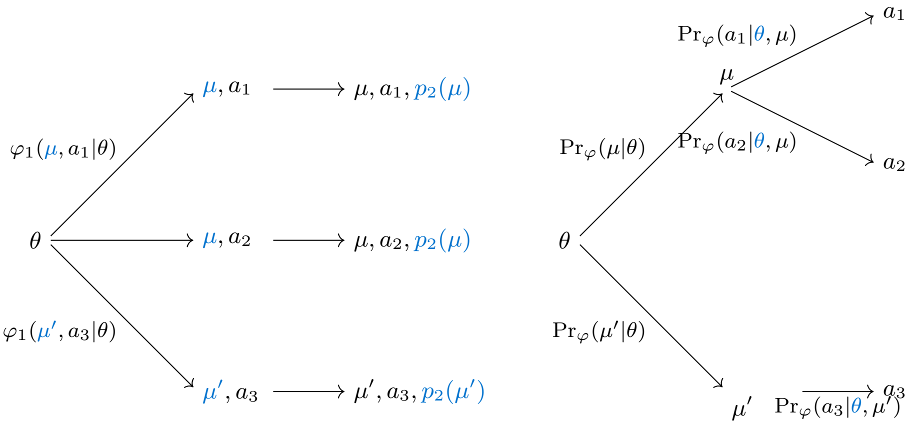

- $\mathrm{Pr}_\varphi(\mu|\theta) = \varphi_1(\mu, a_1|\theta) + \varphi_1(\mu, a_2|\theta)$
- $\mathrm{Pr}_\varphi(a_1|\theta, \mu) = \varphi_1(\mu, a_1|\theta)/Pr_\varphi(\mu|\theta)$

Separating information and allocation design

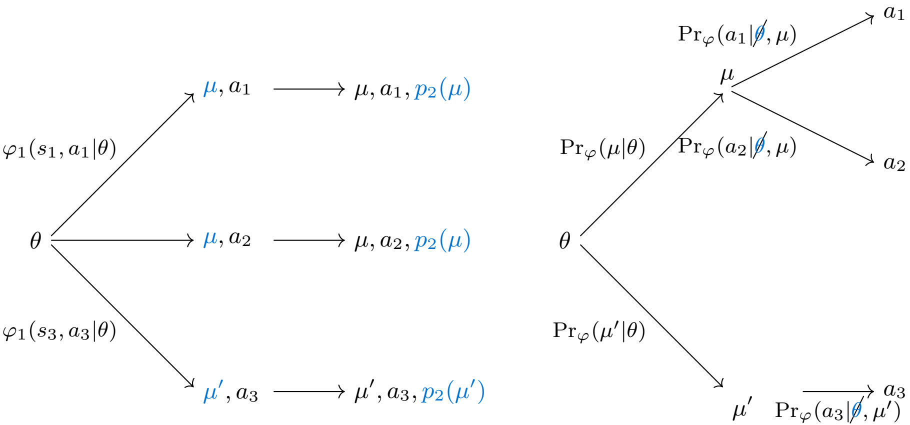

- $\mathrm{Pr}_\varphi(\mu|\theta) = \varphi_1(\mu, a_1|\theta) + \varphi_1(\mu, a_2|\theta)$
- $\mathrm{Pr}_\varphi(a_1|\theta, \mu) = \varphi_1(\mu, a_1|\theta)/Pr_\varphi(\mu|\theta)$

Separating information and allocation design

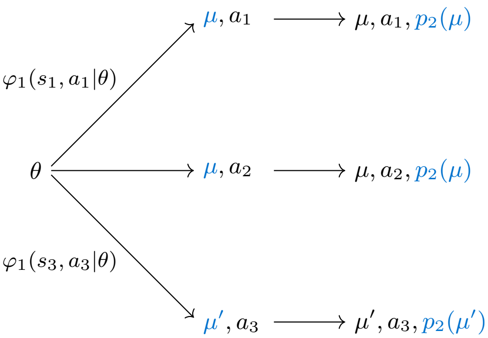

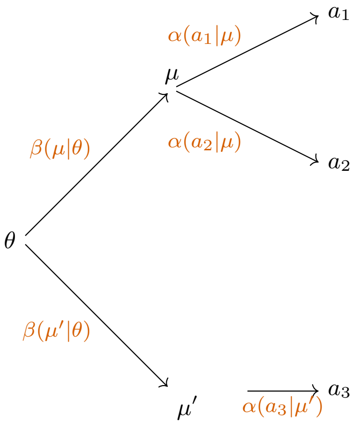

- $\Pr_\varphi(\mu|\theta) = \varphi_1(\mu, a_1|\theta) + \varphi_1(\mu, a_2|\theta)$
- $\Pr_\varphi(a_1|\theta, \mu) = \varphi_1(\mu, a_1|\theta)/Pr_\varphi(\mu|\theta)$

Separating information and allocation design

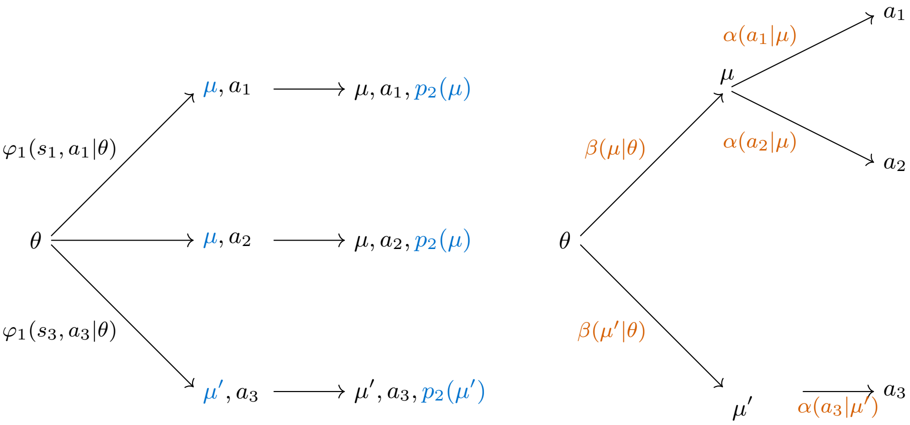

- Separate the design of the information from that of the allocation
- β is the mechanism's disclosure rule and α is the mechanism's allocation rule.

One last simplification

Quasilinearity + separation between allocation and information:

- No need to randomize on transfers: $x(\mu_2)$ is the (expected) payment when output message is μ_2
- $q(\mu_2)$ is the probability of selling the good when output message is μ_2

Thus, the seller's optimal outcome solves:

$$\max_{\text{mechanisms}} \text{Revenue}$$

where $M_1 = \Theta, S_1 = \Delta(\Theta)$, $\varphi = \beta \otimes \alpha$ subject to

- Participation
- Truthtelling
- Consistency between beliefs and output messages.

Seller optimal outcome

Thus, the seller's optimal outcome solves:

Thus, the seller's optimal outcome solves:

$$[x(\mu_2) + (1 - q(\mu_2))\delta R_2(\mu_2)]$$

Seller optimal outcome

Thus, the seller's optimal outcome solves:

$$\left(\sum_{\theta\in\Theta}\mu_1(\theta)\beta(\mu_2|\theta)\right)\left[x(\mu_2)+(1-q(\mu_2))\delta R_2(\mu_2)\right]$$

Thus, the seller's optimal outcome solves:

$$\max_{\beta,q,x} \sum_{\mu_2 \in \Delta(\Theta)} \left(\sum_{\theta \in \Theta} \mu_1(\theta)\beta(\mu_2|\theta) \right) [x(\mu_2) + (1 - q(\mu_2))\delta R_2(\mu_2)]$$

Thus, the seller's optimal outcome solves:

$$R_1(\mu_1) \equiv \max_{\beta,q,x} \sum_{\mu_2 \in \Delta(\Theta)} \left(\sum_{\theta \in \Theta} \mu_1(\theta)\beta(\mu_2|\theta) \right) [x(\mu_2) + (1 - q(\mu_2))\delta R_2(\mu_2)],$$

Thus, the seller's optimal outcome solves:

$$R_1(\mu_1) \equiv \max_{\beta,q,x} \sum_{\mu_2 \in \Delta(\Theta)} \left(\sum_{\theta \in \Theta} \mu_1(\theta)\beta(\mu_2|\theta) \right) [x(\mu_2) + (1 - q(\mu_2))\delta R_2(\mu_2)],$$

subject to for all $\theta \in \{\theta_L, \theta_H\}$:

Thus, the seller's optimal outcome solves:

$$R_1(\mu_1) \equiv \max_{\beta,q,x} \sum_{\mu_2 \in \Delta(\Theta)} \left(\sum_{\theta \in \Theta} \mu_1(\theta)\beta(\mu_2|\theta) \right) [x(\mu_2) + (1 - q(\mu_2))\delta R_2(\mu_2)],$$

subject to for all $\theta \in \{\theta_L, \theta_H\}$:

Participation$_\theta$: $$\sum_{\mu_2 \in \Delta(\Theta)} \beta(\mu_2|\theta)(\theta q(\mu_2) - x(\mu_2) + (1 - q(\mu_2))\delta u^*(\mu_2, \theta)) \geq 0$$

Thus, the seller's optimal outcome solves:

$$R_1(\mu_1) \equiv \max_{\beta,q,x} \sum_{\mu_2 \in \Delta(\Theta)} \left(\sum_{\theta \in \Theta} \mu_1(\theta)\beta(\mu_2|\theta) \right) [x(\mu_2) + (1 - q(\mu_2))\delta R_2(\mu_2)],$$

subject to for all $\theta \in \{\theta_L, \theta_H\}$:

$$\text{Participation}_{\theta}: \quad \sum_{\mu_2 \in \Delta(\Theta)} \beta(\mu_2|\theta)(\theta q(\mu_2) - x(\mu_2) + (1 - q(\mu_2))\delta u^*(\mu_2, \theta)) \geq 0$$

$$\text{Truthtelling}_{\theta,\theta'}: \quad \sum_{\mu_2 \in \Delta(\Theta)} (\beta(\mu_2|\theta) - \beta(\mu_2|\theta'))(\theta q(\mu_2) - x(\mu_2) + (1 - q(\mu_2))\delta u^*(\mu_2, \theta)) \geq 0$$

Thus, the seller's optimal outcome solves:

$$R_1(\mu_1) \equiv \max_{\beta,q,x} \sum_{\mu_2 \in \Delta(\Theta)} \left(\sum_{\theta \in \Theta} \mu_1(\theta)\beta(\mu_2|\theta) \right) [x(\mu_2) + (1 - q(\mu_2))\delta R_2(\mu_2)],$$

subject to for all $\theta \in \{\theta_L, \theta_H\}$:

Participation$_\theta$: $$\sum_{\mu_2 \in \Delta(\Theta)} \beta(\mu_2|\theta)(\theta q(\mu_2) - x(\mu_2) + (1 - q(\mu_2))\delta u^*(\mu_2, \theta)) \geq 0$$

Truthtelling$_{\theta,\theta'}$: $$\sum_{\mu_2 \in \Delta(\Theta)} (\beta(\mu_2|\theta) - \beta(\mu_2|\theta'))(\theta q(\mu_2) - x(\mu_2) + (1 - q(\mu_2))\delta u^*(\mu_2, \theta)) \geq 0$$

Consistency$_{\mu_2}$: $$\mu_2(\theta_H)[\sum_\theta \mu_1(\theta)\beta(\mu_2|\theta)] = \mu_1(\theta_H)\beta(\mu_2|\theta_H)$$

Seller optimal outcome

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

$$q(\mu_2)$$

Seller optimal outcome

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

$$q(\mu_2)(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1))$$

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

$$q(\mu_2)(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) + (1-q(\mu_2))$$

Seller optimal outcome

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

$$q(\mu_2)(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) + (1-q(\mu_2))\delta R(\mu_2;\mu_1)$$

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

$$\left(\sum_{\theta\in\Theta}\mu_1(\theta)\beta(\mu_2|\theta)\right)\left[q(\mu_2)(\mu_2\theta_H+(1-\mu_2)\hat{\theta}_L(\mu_1))+(1-q(\mu_2))\delta R(\mu_2;\mu_1)\right]$$

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

$$\tau(\mu_2)\left[q(\mu_2)(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) + (1-q(\mu_2))\delta R(\mu_2;\mu_1)\right]$$

Seller optimal outcome

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

$$R_1(\mu_1) = \max_{\tau,q} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \left[q(\mu_2)(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) + (1-q(\mu_2))\delta R(\mu_2;\mu_1) \right]$$

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

$$R_1(\mu_1) = \max_{\tau,q} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \left[q(\mu_2)(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) + (1-q(\mu_2))\delta R(\mu_2;\mu_1) \right]$$

subject to

$$\sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2)\mu_2(\theta_H) = \mu_1(\theta_H)$$

At the optimum, the following hold:

- The seller extracts all surplus from low-valuation buyer ($\Rightarrow$ participation binds for θ_L)
- High-valuation buyer is indifferent between reporting θ_H and θ_L ($\Rightarrow$ Truthtelling binds for θ_H)

Hence, we can rewrite the problem as:

$$R_1(\mu_1) = \max_{\tau,q} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \left[q(\mu_2)(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) + (1-q(\mu_2))\delta R(\mu_2;\mu_1) \right]$$

subject to

$$\sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2)\mu_2(\theta_H) = \mu_1(\theta_H)$$

Sale of a durable good: $t = 1$

$$\delta R_2(\mu_2; \mu_1) = \begin{cases} \delta(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) & \text{if } \mu_2 < \bar{\mu} \\ \delta\mu_2\theta_H & \text{if } \mu_2 > \bar{\mu} \end{cases}$$

Sale of a durable good: $t = 1$

$$\delta R_2(\mu_2; \mu_1) = \begin{cases} \delta(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) & \text{if } \mu_2 < \bar{\mu} \\ \delta\mu_2\theta_H & \text{if } \mu_2 > \bar{\mu} \end{cases}$$

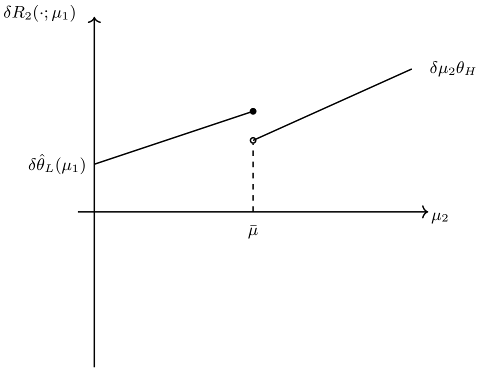

$\mu_1 \leq \bar{\mu}$

Sale of a durable good: $t = 1$

$$\delta R_2(\mu_2; \mu_1) = \begin{cases} \delta(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) & \text{if } \mu_2 < \bar{\mu} \\ \delta\mu_2\theta_H & \text{if } \mu_2 > \bar{\mu} \end{cases}$$

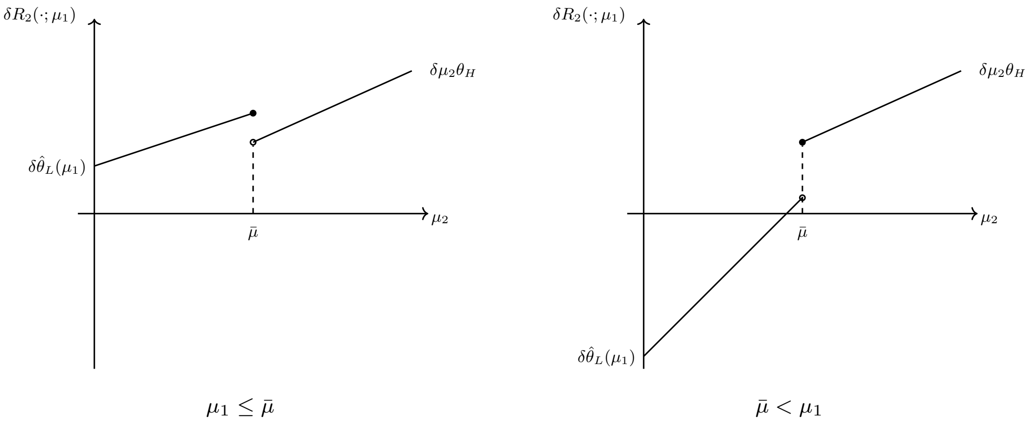

Sale of a durable good: $t = 1$

$$\delta R_2(\mu_2; \mu_1) = \begin{cases} \delta(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) & \text{if } \mu_2 < \overline{\mu} \\ \delta\mu_2\theta_H & \text{if } \mu_2 > \bar{\mu} \end{cases}$$

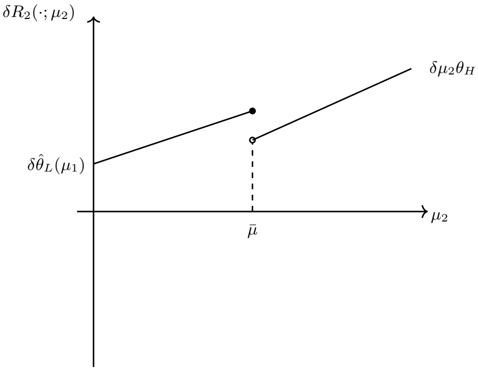

$\mu_1 \leq \bar{\mu}$

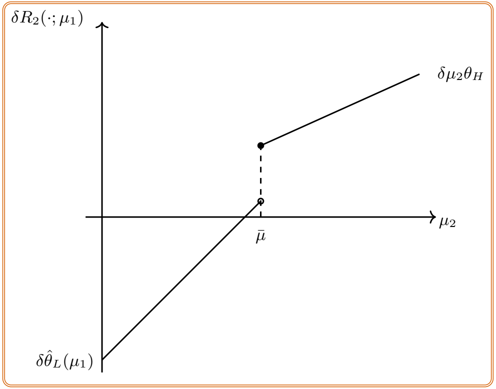

$\bar{\mu} < \mu_1$

Seller optimal outcome

$$\max_{\tau,q} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \left[q(\mu_2)(\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) + (1-q(\mu_2))\delta R_2(\mu_2;\mu_1) \right]$$

Seller optimal outcome

$$\max_{\tau,q} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \left[0 \times (\mu_2 \theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) + 1 \times \delta R_2(\mu_2;\mu_1) \right]$$

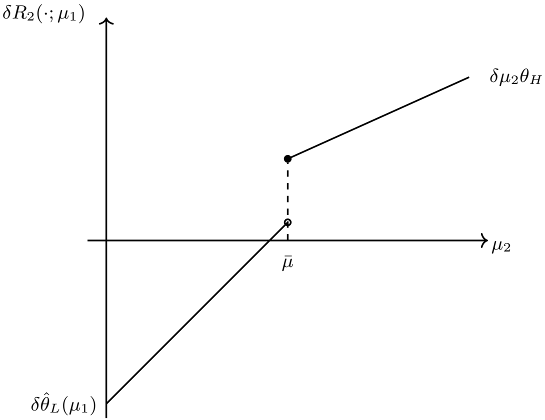

Seller optimal outcome

$$\max_{\tau,q} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \left[1 \times (\mu_2 \theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1)) + 0 \times \delta R_2(\mu_2; \mu_1) \right]$$

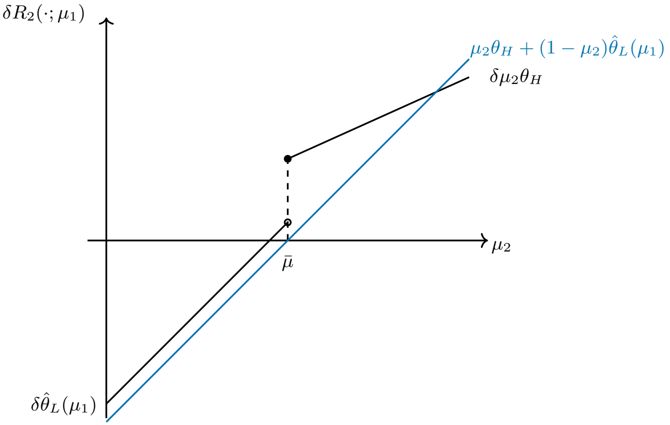

Seller optimal outcome

$$\max_{\tau} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \max\{\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1), \delta R(\mu_2;\mu_1)\}$$

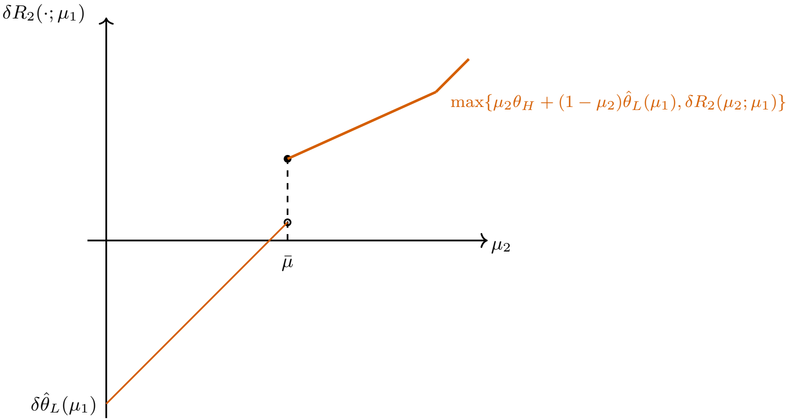

Seller optimal outcome

$$\max_{\tau} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \max\{\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1), \delta R(\mu_2;\mu_1)\}$$

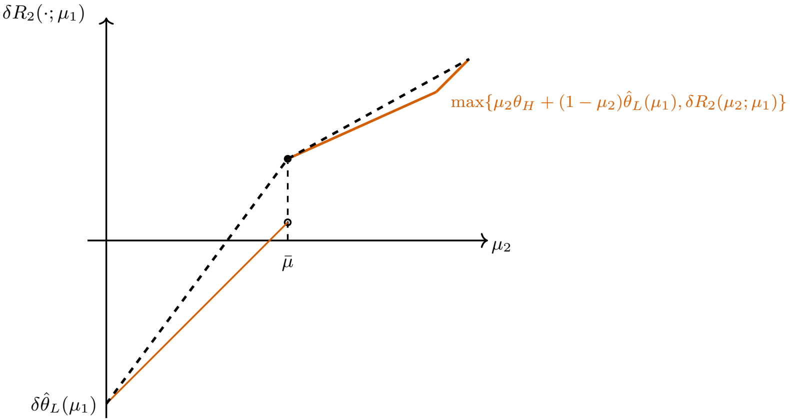

Seller optimal outcome

$$\max_{\tau} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \max\{\mu_2 \theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1), \delta R(\mu_2;\mu_1)\}$$

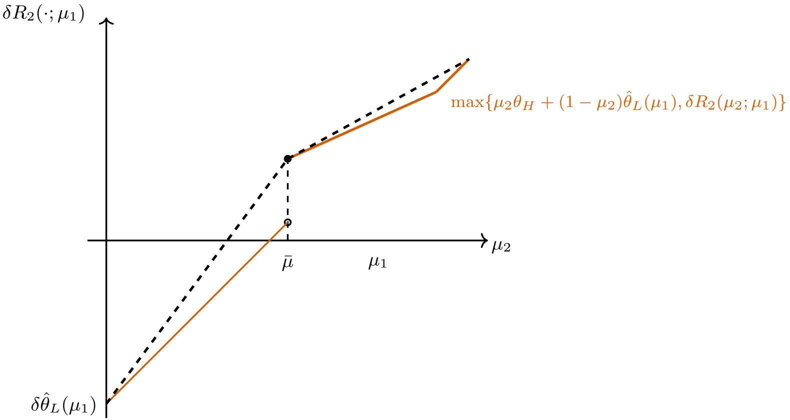

Seller optimal outcome

$$\max_{\tau} \sum_{\mu_2 \in \Delta(\Theta)} \tau(\mu_2) \max\{\mu_2\theta_H + (1-\mu_2)\hat{\theta}_L(\mu_1), \delta R(\mu_2;\mu_1)\}$$

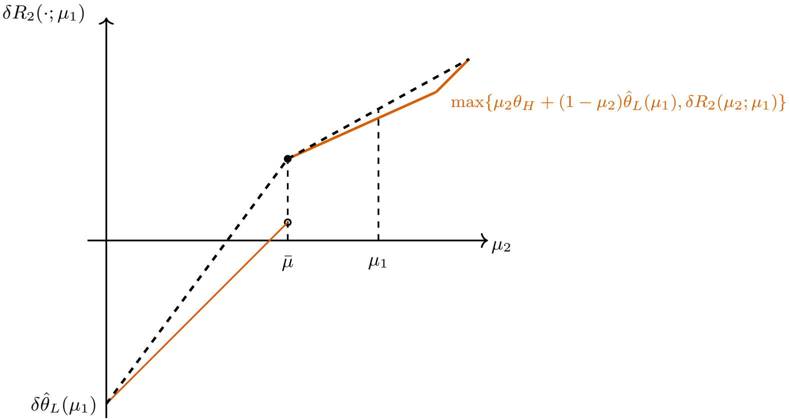

Seller optimal outcome

- Seller splits μ_1 between $\mu_2 = \bar{\mu}$ and $\mu_2 = 1$
- He sells when $\mu_2 = 1$ $(q(1) = 1)$ and delays when $\mu_2 = \bar{\mu}$ $(q(\bar{\mu}) = 0)$
- Posted price of θ_H in both periods.

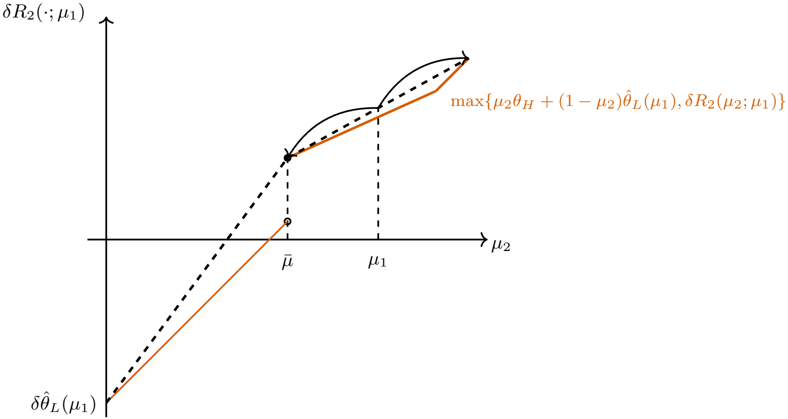

Seller optimal outcome

- Seller splits μ_1 between $\mu_2 = \bar{\mu}$ and $\mu_2 = 1$
- He sells when $\mu_2 = 1$ $(q(1) = 1)$ and delays when $\mu_2 = \bar{\mu}$ $(q(\bar{\mu}) = 0)$
- Posted price of θ_H in both periods.

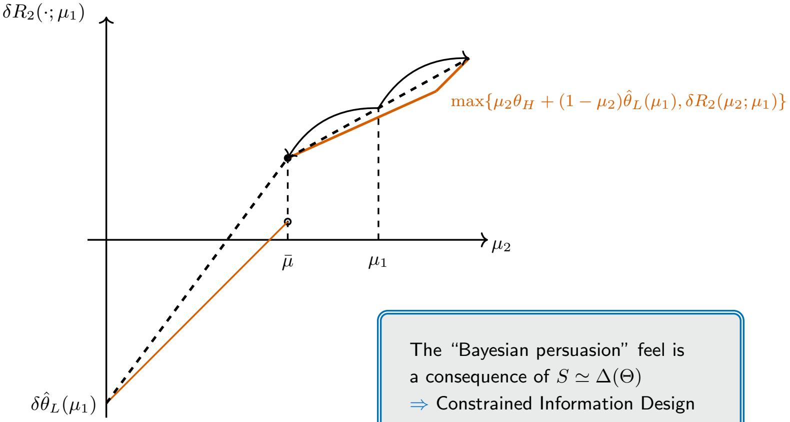

The "Bayesian persuasion" feel is a consequence of $S \simeq \Delta(\Theta)$
$\Rightarrow$ Constrained Information Design

Economic trade-off: tailor the allocation to the agent's report vs. learning about the agent's type.

- No such trade-off when there is commitment: acquired information can always be "forgotten."
- The seller slows down learning:
  - Similar to Kanoria & Nazerzadeh, 2014; Abernethy et al., 2019; Haghtalab, Lykouris, Nietert,& Wei, 2022

Open questions

Open questions

This is a problem that had been open in Economics for 30 years. There's much to do!

1. Most glaring: multiple agents (the existing counterexamples do not survive with our mechanisms)
   - How to aggregate the information from the multiple agents? (e.g., Halpern & Teague, 2006)
2. More practical: How to implement direct-Blackwell mechanisms?
   - Multiple (infinite?) rounds of indirect observable communication?
   - Cryptographic commitments? (e.g., Ferreira & Weinberg, 2020)