The Persuasion Duality

Piotr Dworczak & Anton Kolotilin
(Northwestern University) (UNSW Business School)

University of Oxford, October, 2025

Motivation

- An explosion of interest in **Bayesian persuasion** following Kamenica and Gentzkow (2011)
- Concavification not always tractable
- A number of papers propose **duality** as a tool to solve information design problems:
    - Kolotilin (2018);
    - Dworczak and Martini (2019);
    - Dizdar and Kováč (2020);
    - Kolotilin, Corrao, and Wolitzky (2025);
    - (Galperti, Levkun, and Perego (2023));
    - ...

Introduction

- We propose a **unified duality approach** to persuasion.

Introduction

- We propose a **unified duality approach** to persuasion.
- The optimal dual variable, interpreted as a price function, is a **supergradient of the concave closure of the objective function at the prior** belief.

Introduction

- We propose a **unified duality approach** to persuasion.
- The optimal dual variable, interpreted as a price function, is a **supergradient of the concave closure of the objective function at the prior** belief.
- Our results unify and generalize existing duality results in persuasion.

Introduction

- We propose a **unified duality approach** to persuasion.
- The optimal dual variable, interpreted as a price function, is a **supergradient of the concave closure of the objective function at the prior** belief.
- Our results unify and generalize existing duality results in persuasion.
- This minicourse introduces methodology and illustrates it in applications.

Recap of Kamenica and Gentzkow (2011)

Sender chooses a signal/experiment to reveal information about a state to Receiver, who then takes an action.

Recap of Kamenica and Gentzkow (2011)

Sender chooses a signal/experiment to reveal information about a state to Receiver, who then takes an action.

- There is a state $\omega \in \Omega$ (compact or finite) with distribution μ_0.

Recap of Kamenica and Gentzkow (2011)

Sender chooses a signal/experiment to reveal information about a state to Receiver, who then takes an action.

- There is a state $\omega \in \Omega$ (compact or finite) with distribution μ_0.
- Sender designs a signal:
  - Signal realization space S;
  - Conditional probability $\pi(s|\omega)$ for each $s \in S$ and $\omega \in \Omega$;

Recap of Kamenica and Gentzkow (2011)

Sender chooses a signal/experiment to reveal information about a state to Receiver, who then takes an action.

- There is a state $\omega \in \Omega$ (compact or finite) with distribution μ_0.
- Sender designs a signal:
  - Signal realization space S;
  - Conditional probability $\pi(s|\omega)$ for each $s \in S$ and $\omega \in \Omega$;
- Receiver chooses action $a \in A$ (compact).

Recap of Kamenica and Gentzkow (2011)

Sender chooses a signal/experiment to reveal information about a state to Receiver, who then takes an action.

- There is a state $\omega \in \Omega$ (compact or finite) with distribution μ_0.
- Sender designs a signal:
  - Signal realization space S;
  - Conditional probability $\pi(s|\omega)$ for each $s \in S$ and $\omega \in \Omega$;
- Receiver chooses action $a \in A$ (compact).
- Receiver's utility is $u(a, \omega)$ and Sender's utility is $v(a, \omega)$; both are continuous

Belief-based approach

- After observing *s*, Receiver uses Bayes' rule to update his belief from the *prior* $\mu_0 \in \Delta(\Omega)$ to the *posterior* $\mu \in \Delta(\Omega)$.

Belief-based approach

- After observing s, Receiver uses Bayes' rule to update his belief from the *prior* $\mu_0 \in \Delta(\Omega)$ to the *posterior* $\mu \in \Delta(\Omega)$.

- Given posterior μ, Receiver takes a *best-response* action

$$a^*(\mu) \in \arg\max_{a \in A} \mathbb{E}_\mu[u(a, \omega)],$$

breaking possible ties in favor of Sender.

Belief-based approach

- After observing s, Receiver uses Bayes' rule to update his belief from the *prior* $\mu_0 \in \Delta(\Omega)$ to the *posterior* $\mu \in \Delta(\Omega)$.

- Given posterior μ, Receiver takes a *best-response* action

$$a^*(\mu) \in \arg\max_{a \in A} \mathbb{E}_\mu[u(a, \omega)],$$

  breaking possible ties in favor of Sender.

- Sender's *indirect utility* from posterior μ is

$$V(\mu) = \mathbb{E}_\mu[v(a^*(\mu), \omega)].$$

Belief-based approach

- After observing s, Receiver uses Bayes' rule to update his belief from the *prior* $\mu_0 \in \Delta(\Omega)$ to the *posterior* $\mu \in \Delta(\Omega)$.

- Given posterior μ, Receiver takes a *best-response* action

$$a^*(\mu) \in \arg\max_{a \in A} \mathbb{E}_\mu[u(a, \omega)],$$

  breaking possible ties in favor of Sender.

- Sender's *indirect utility* from posterior μ is

$$V(\mu) = \mathbb{E}_\mu[v(a^*(\mu), \omega)].$$

- A signal π induces a distribution τ over posteriors μ, so Sender's expected utility is $\mathbb{E}_\tau[V(\mu)]$.

Splitting lemma

Lemma. There exists a signal π that induces a distribution of posteriors $\tau \in \Delta(\Delta(\Omega))$ iff $\mathbb{E}_\tau[\mu] = \mu_0$.

Splitting lemma

Lemma. There exists a signal π that induces a distribution of posteriors $\tau \in \Delta(\Delta(\Omega))$ iff $\mathbb{E}_\tau[\mu] = \mu_0$.

Proof: The only if part follows from the law of iterated expectations.

Splitting lemma

Lemma. There exists a signal π that induces a distribution of posteriors $\tau \in \Delta(\Delta(\Omega))$ iff $\mathbb{E}_\tau[\mu] = \mu_0$.

Proof: The only if part follows from the law of iterated expectations. The if part is shown by construction. Indeed, in the finite case, define, for all $\omega \in \Omega$ and all μ in $\mathsf{supp}(\tau)$,

$$\pi(\mu|\omega) = \frac{\mu(\omega)\tau(\mu)}{\mu_0(\omega)}$$

$$\implies \Pr(\omega|\mu) = \frac{\pi(\mu|\omega)\mu_0(\omega)}{\tau(\mu)} = \mu(\omega).$$

Concavification

Sender's problem is to find a distribution of posteriors τ to

$$\text{maximize } \mathbb{E}_\tau[V(\mu)]$$

$$\text{subject to } \mathbb{E}_\tau[\mu] = \mu_0.$$

Smallest concave function that is everywhere greater than V is called *concave closure* of V and is denoted by $\widehat{V}$.

Concavification. The value of Sender's problem is $\widehat{V}(\mu_0)$.

Recap of linear programming

Fix an $m \times n$ matrix A, an n-vector b, and an m-vector c.

If the *primal problem* is to find an m-vector $x \geq 0$ to

$$\text{maximize } xc$$
$$\text{subject to } xA = b,$$

then the *dual problem* is to find an n-vector y to

$$\text{minimize } by$$
$$\text{subject to } Ay \geq c.$$

Recap of linear-programming duality

Primal: find $x \geq 0$ to	Dual: find y to
maximize xc	minimize by
subject to $xA = b$,	subject to $Ay \geq c$.

Recap of linear-programming duality

Primal: find $x \geq 0$ to	Dual: find y to
maximize xc	minimize by
subject to $xA = b$,	subject to $Ay \geq c$.

Weak duality. If x and y are feasible solutions, then

$$xc \leq xAy = by.$$

Recap of linear-programming duality

Primal: find $x \geq 0$ to	Dual: find y to
maximize xc	minimize by
subject to $xA = b$,	subject to $Ay \geq c$.

Weak duality. If x and y are feasible solutions, then

$$xc \leq xAy = by.$$

Optimality criterion. If feasible solutions x and y satisfy $xc = by$, then they are optimal, as, for any feasible solutions $\tilde{x}$ and $\tilde{y}$,

$$\tilde{x}c \leq by = xc \quad \text{and} \quad b\tilde{y} \geq xc = by.$$

Recap of linear-programming duality

Primal: find $x \geq 0$ to	Dual: find y to
maximize xc	minimize by
subject to $xA = b$,	subject to $Ay \geq c$.

Weak duality. If x and y are feasible solutions, then

$$xc \leq xAy = by.$$

Optimality criterion. If feasible solutions x and y satisfy $xc = by$, then they are optimal, as, for any feasible solutions $\tilde{x}$ and $\tilde{y}$,

$$\tilde{x}c \leq by = xc \quad \text{and} \quad b\tilde{y} \geq xc = by.$$

Strong duality. If the primal and dual admit feasible solutions, then both have optimal solutions x and y, and they satisfy $xc = by$.

Primal Problem in Dworczak and Kolotilin (2024)

Find a distribution of posteriors $\tau \in \Delta(\Delta(\Omega))$ to

$$\begin{aligned} \text{maximize} & \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu) \\ \text{subject to} & \int_{\Delta(\Omega)} \mu \mathrm{d}\tau(\mu) = \mu_0 \end{aligned} \tag{P}$$

Primal Problem in Dworczak and Kolotilin (2024)

Find a distribution of posteriors $\tau \in \Delta(\Delta(\Omega))$ to

$$\begin{aligned} \text{maximize} \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu) \\ \text{subject to} \int_{\Delta(\Omega)} \mu \mathrm{d}\tau(\mu) = \mu_0 \end{aligned} \tag{P}$$

where

- (Ω, ρ) is a compact metric space
- $\mu_0 \in \Delta(\Omega)$ is a prior belief
- $V : \Delta(\Omega) \to \mathbb{R}$ is bounded and u.s.c. in the weak* topology

Technical Background

- Let $M(\Omega)$ be the set of finite signed Borel measures on Ω

Technical Background

- Let $M(\Omega)$ be the set of finite signed Borel measures on Ω
- Let $\|\cdot\|_{KR}$ denote the Kantorovich-Rubinstein norm on $M(\Omega)$

Technical Background

- Let $M(\Omega)$ be the set of finite signed Borel measures on Ω
- Let $\|\cdot\|_{KR}$ denote the Kantorovich-Rubinstein norm on $M(\Omega)$
- **Fact 1**. $\|\cdot\|_{KR}$ metrizes the weak* topology on $\Delta(\Omega)$

Technical Background

- Let $M(\Omega)$ be the set of finite signed Borel measures on Ω
- Let $\|\cdot\|_{KR}$ denote the Kantorovich-Rubinstein norm on $M(\Omega)$
- **Fact 1.** $\|\cdot\|_{KR}$ metrizes the weak* topology on $\Delta(\Omega)$
- **Fact 2.** $(\Delta(\Omega), \|\cdot\|_{KR})$ is a compact metric space

Technical Background

- Let $M(\Omega)$ be the set of finite signed Borel measures on Ω
- Let $\|\cdot\|_{KR}$ denote the Kantorovich-Rubinstein norm on $M(\Omega)$
- **Fact 1.** $\|\cdot\|_{KR}$ metrizes the weak* topology on $\Delta(\Omega)$
- **Fact 2.** $(\Delta(\Omega), \|\cdot\|_{KR})$ is a compact metric space
- Let $L(\Omega)$ denote the set of Lipschitz functions on Ω

Technical Background

- Let $M(\Omega)$ be the set of finite signed Borel measures on Ω
- Let $\|\cdot\|_{KR}$ denote the Kantorovich-Rubinstein norm on $M(\Omega)$
- **Fact 1.** $\|\cdot\|_{KR}$ metrizes the weak* topology on $\Delta(\Omega)$
- **Fact 2.** $(\Delta(\Omega), \|\cdot\|_{KR})$ is a compact metric space
- Let $L(\Omega)$ denote the set of Lipschitz functions on Ω
- **Fact 3.** The space dual to $(M(\Omega), \|\cdot\|_{KR})$ is $L(\Omega)$

Technical Background

- Let $M(\Omega)$ be the set of finite signed Borel measures on Ω
- Let $\|\cdot\|_{KR}$ denote the Kantorovich-Rubinstein norm on $M(\Omega)$
- **Fact 1.** $\|\cdot\|_{KR}$ metrizes the weak* topology on $\Delta(\Omega)$
- **Fact 2.** $(\Delta(\Omega), \|\cdot\|_{KR})$ is a compact metric space
- Let $L(\Omega)$ denote the set of Lipschitz functions on Ω
- **Fact 3.** The space dual to $(M(\Omega), \|\cdot\|_{KR})$ is $L(\Omega)$
- Note that $(\Delta(\Delta(\Omega)), \|\cdot\|_{KR})$ is also a compact metric space

Dual problem

The dual problem is to find a *price function* $p \in L(\Omega)$ to

$$\begin{aligned} \text{minimize} & \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega) \\ \text{subject to} & \int_\Omega p(\omega)\mathrm{d}\mu(\omega) \geq V(\mu) \text{ for all } \mu \in \Delta(\Omega). \end{aligned} \tag{D}$$

Value functions

Let $\mathcal{T}(\mu_0)$ and $\mathcal{P}(V)$ be the primal and dual sets of feasible solutions.

Value functions

Let $\mathcal{T}(\mu_0)$ and $\mathcal{P}(V)$ be the primal and dual sets of feasible solutions.

The *concave closure* of V at μ_0 is the value of the primal problem:

$$\widehat{V}(\mu_0) := \sup_{\tau \in \mathcal{T}(\mu_0)} \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu).$$

$\widehat{V}(\mu_0)$ is the supremum of z such that (z, μ_0) belongs to the convex hull of the graph of V on $\Delta(\Omega)$.

Value functions

Let $\mathcal{T}(\mu_0)$ and $\mathcal{P}(V)$ be the primal and dual sets of feasible solutions.

The *concave closure* of V at μ_0 is the value of the primal problem:

$$\widehat{V}(\mu_0) := \sup_{\tau \in \mathcal{T}(\mu_0)} \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu).$$

$\widehat{V}(\mu_0)$ is the supremum of z such that (z, μ_0) belongs to the convex hull of the graph of V on $\Delta(\Omega)$.

The *concave envelope* of V at μ_0 is the value of the dual problem:

$$\overline{V}(\mu_0) := \inf_{p \in \mathcal{P}(V)} \int_{\Omega} p(\omega) \mathrm{d}\mu_0(\omega).$$

Value functions

Let $\mathcal{T}(\mu_0)$ and $\mathcal{P}(V)$ be the primal and dual sets of feasible solutions.

The *concave closure* of V at μ_0 is the value of the primal problem:

$$\widehat{V}(\mu_0) := \sup_{\tau \in \mathcal{T}(\mu_0)} \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu).$$

$\widehat{V}(\mu_0)$ is the supremum of z such that (z, μ_0) belongs to the convex hull of the graph of V on $\Delta(\Omega)$.

The *concave envelope* of V at μ_0 is the value of the dual problem:

$$\overline{V}(\mu_0) := \inf_{p \in \mathcal{P}(V)} \int_{\Omega} p(\omega) \mathrm{d}\mu_0(\omega).$$

$\overline{V}(\mu_0)$ is the infimum at μ_0 of continuous linear functions on $M(\Omega)$ that bound V from above on $\Delta(\Omega)$, because the space $L(\Omega)$ is dual to $(M(\Omega), \|\cdot\|_{KR})$.

Primal problem

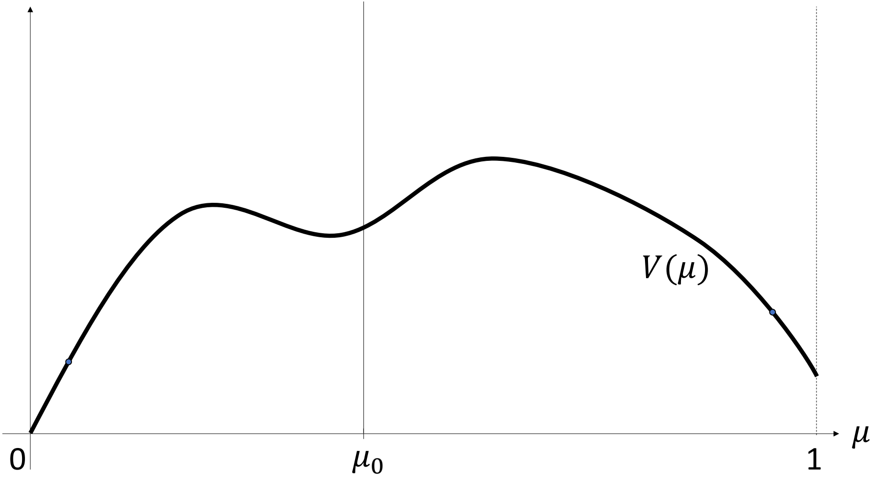

Primal problem

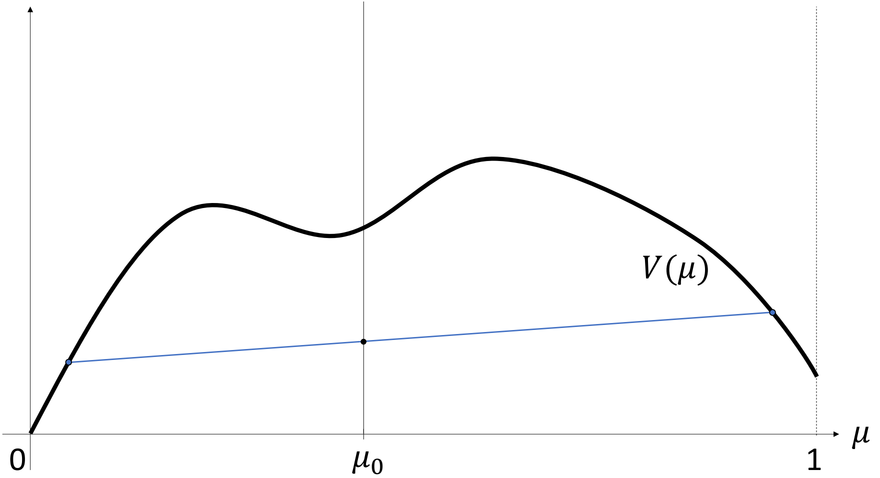

Primal problem

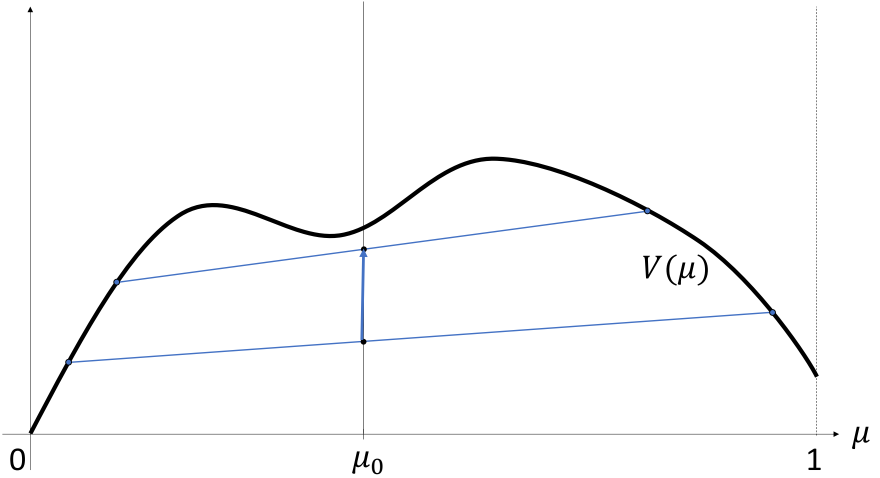

Primal problem

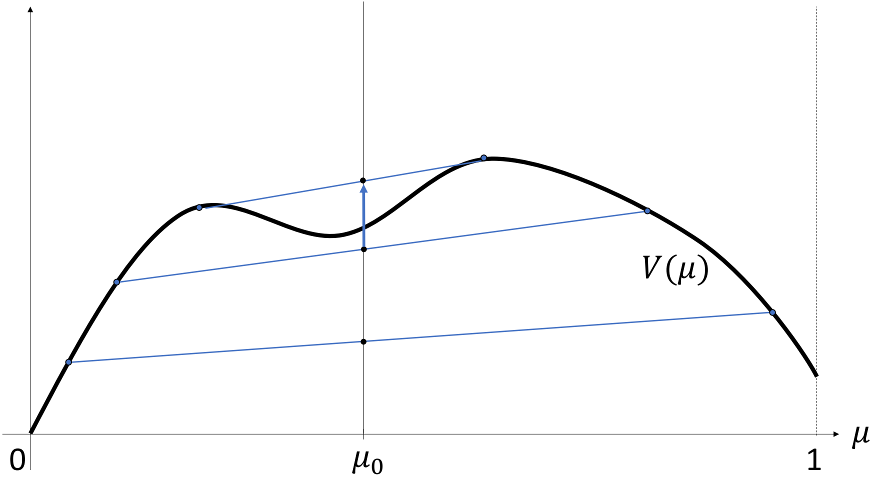

Primal problem

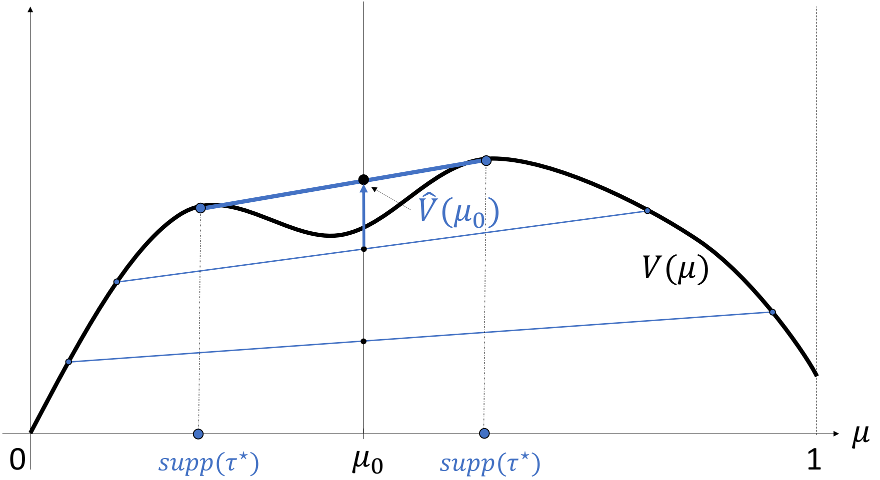

Dual problem

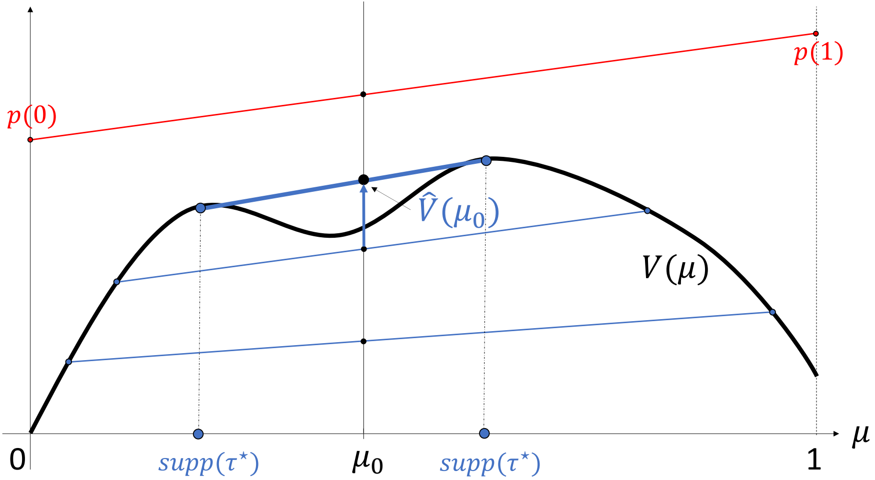

Dual problem

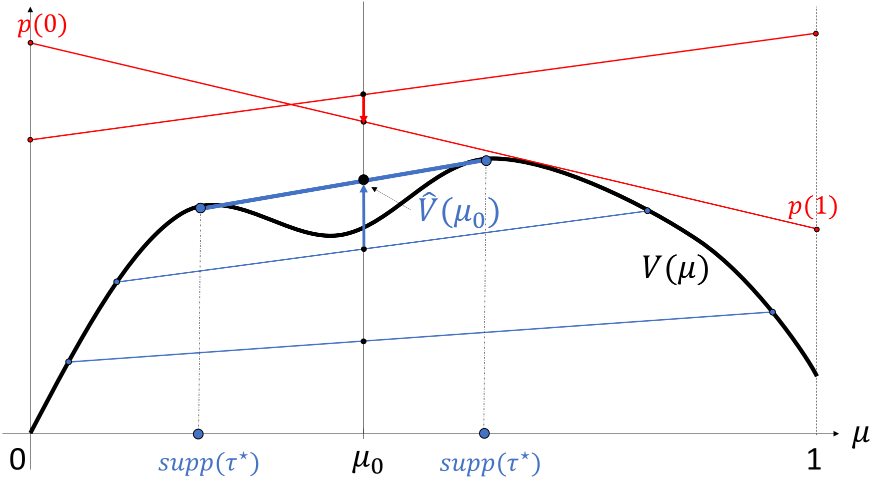

Dual problem

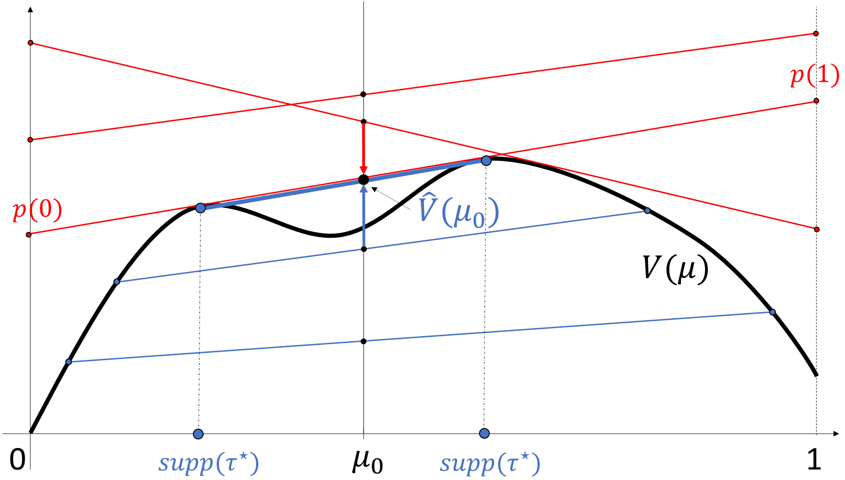

Dual problem

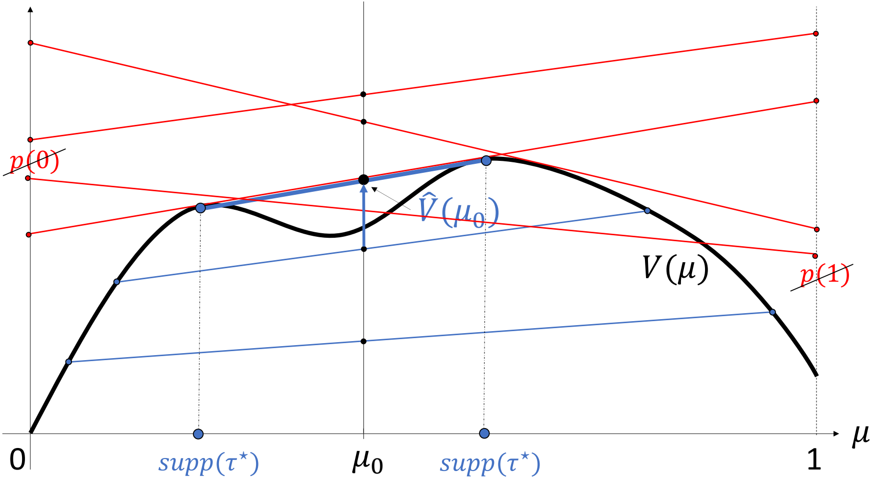

Dual problem

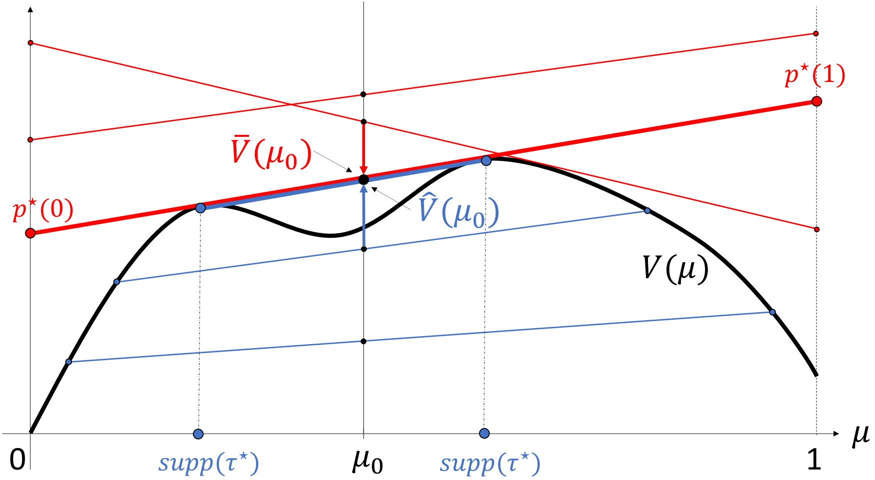

Interpretation as a production problem

We first interpret the primal problem.

Interpretation as a production problem

We first interpret the primal problem.

- There is Producer designing a production plan.

Interpretation as a production problem

We first interpret the primal problem.

- There is Producer designing a production plan.
- The set of resources is Ω.

Interpretation as a production problem

We first interpret the primal problem.

- There is Producer designing a production plan.
- The set of resources is Ω.
- Producer's endowment is $\mu_0 \in \Delta(\Omega)$.

Interpretation as a production problem

We first interpret the primal problem.

- There is Producer designing a production plan.
- The set of resources is Ω.
- Producer's endowment is $\mu_0 \in \Delta(\Omega)$.
- The set of linear production processes $\Delta(\Omega)$.

Interpretation as a production problem

We first interpret the primal problem.

- There is Producer designing a production plan.
- The set of resources is Ω.
- Producer's endowment is $\mu_0 \in \Delta(\Omega)$.
- The set of linear production processes $\Delta(\Omega)$.
- Process $\mu \in \Delta(\Omega)$ operated at unit level consumes measure μ of resources and generates income $V(\mu)$.

Interpretation as a production problem

We first interpret the primal problem.

- There is Producer designing a production plan.
- The set of resources is Ω.
- Producer's endowment is $\mu_0 \in \Delta(\Omega)$.
- The set of linear production processes $\Delta(\Omega)$.
- Process $\mu \in \Delta(\Omega)$ operated at unit level consumes measure μ of resources and generates income $V(\mu)$.
- Production plan $\tau \in \Delta(\Delta(\Omega))$ describes the level at which each process μ is operated.

Interpretation as a production problem

We first interpret the primal problem.

- There is Producer designing a production plan.
- The set of resources is Ω.
- Producer's endowment is $\mu_0 \in \Delta(\Omega)$.
- The set of linear production processes $\Delta(\Omega)$.
- Process $\mu \in \Delta(\Omega)$ operated at unit level consumes measure μ of resources and generates income $V(\mu)$.
- Production plan $\tau \in \Delta(\Delta(\Omega))$ describes the level at which each process μ is operated.
- The primal problem is to find a production plan $\tau \in \Delta(\Delta(\Omega))$ that exhausts endowment μ_0 and maximizes total income.

Interpretation as a production problem

We now interpret the dual problem.

Interpretation as a production problem

We now interpret the dual problem.

- Dealer wants to buy out the producer.

Interpretation as a production problem

We now interpret the dual problem.

- Dealer wants to buy out the producer.
- Dealer sets a unit price $p(\omega)$ for each resource $\omega \in \Omega$.

Interpretation as a production problem

We now interpret the dual problem.

- Dealer wants to buy out the producer.
- Dealer sets a unit price $p(\omega)$ for each resource $\omega \in \Omega$.
- Price function p is feasible for Dealer if Producer's income $V(\mu)$ from each process μ is below his cost $\int_\Omega p(\omega)\mathrm{d}\mu(\omega)$.

Interpretation as a production problem

We now interpret the dual problem.

- Dealer wants to buy out the producer.
- Dealer sets a unit price $p(\omega)$ for each resource $\omega \in \Omega$.
- Price function p is feasible for Dealer if Producer's income $V(\mu)$ from each process μ is below his cost $\int_\Omega p(\omega)\mathrm{d}\mu(\omega)$.
- The dual problem is to find feasible prices that minimize the total cost of buying up all the resources.

Interpretation as a production problem

We finally decentralize this production economy. A production plan τ and a price function *p* constitute a competitive equilibrium if:

Interpretation as a production problem

We finally decentralize this production economy. A production plan τ and a price function *p* constitute a competitive equilibrium if:

1. All resources are used in production:

$$\int \mu \mathrm{d}\tau(\mu) = \mu_0.$$

Interpretation as a production problem

We finally decentralize this production economy. A production plan τ and a price function *p* constitute a competitive equilibrium if:

1. All resources are used in production:

$$\int \mu \mathrm{d}\tau(\mu) = \mu_0.$$

2. Operating processes make zero profits:

$$\int p(\omega)\mathrm{d}\mu(\omega) = V(\mu), \quad \text{for all } \mu \in \mathrm{supp}(\tau).$$

Interpretation as a production problem

We finally decentralize this production economy. A production plan τ and a price function p constitute a competitive equilibrium if:

1. All resources are used in production:

$$\int \mu \mathrm{d}\tau(\mu) = \mu_0.$$

2. Operating processes make zero profits:

$$\int p(\omega)\mathrm{d}\mu(\omega) = V(\mu), \quad \text{for all } \mu \in \mathrm{supp}(\tau).$$

3. No entrant can make strictly positive profits:

$$\int p(\omega)\mathrm{d}\mu(\omega) \geq V(\mu), \quad \text{for all } \mu \in \Delta(\Omega).$$

Interpretation as a production problem

Weak duality is the first welfare theorem. If a pair (τ, p) is a competitive equilibrium, then τ solves Producer's problem, and p solves Dealer's problem.

Interpretation as a production problem

Weak duality is the first welfare theorem. If a pair (τ, p) is a competitive equilibrium, then τ solves Producer's problem, and p solves Dealer's problem.

This is because, for any feasible τ and p, we have

$$\int p \mathrm{d}\mu_0 - \int V \mathrm{d}\tau = \int \left(\int p \mathrm{d}\mu - V \right) \mathrm{d}\tau \geq 0$$

Interpretation as a production problem

Weak duality is the first welfare theorem. If a pair (τ, p) is a competitive equilibrium, then τ solves Producer's problem, and p solves Dealer's problem.

This is because, for any feasible τ and p, we have

$$\int p \mathrm{d}\mu_0 - \int V \mathrm{d}\tau = \int \left(\int p \mathrm{d}\mu - V \right) \mathrm{d}\tau \geq 0$$

Strong duality is the second welfare theorem. Producer's and Dealer's problems admit optimal solutions τ and p. Any such solutions constitute a competitive equilibrium.

Application of duality

- Assume that weak duality and strong duality hold.

Application of duality

- Assume that weak duality and strong duality hold.
- Assume that $\mathsf{supp}(\mu_0) = \Omega$.

Application of duality

- Assume that weak duality and strong duality hold.
- Assume that $\mathsf{supp}(\mu_0) = \Omega$.
- Weak duality yields sufficient optimality conditions.

Application of duality

- Assume that weak duality and strong duality hold.
- Assume that $\mathrm{supp}(\mu_0) = \Omega$.
- Weak duality yields sufficient optimality conditions.
- Strong duality yields necessary optimality conditions.

Application of duality

- Assume that weak duality and strong duality hold.
- Assume that $\text{supp}(\mu_0) = \Omega$.
- Weak duality yields sufficient optimality conditions.
- Strong duality yields necessary optimality conditions.
- Specifically, $\tau \in \mathcal{T}(\mu_0)$ is optimal iff there exists $p \in L(\Omega)$:

$$\int_\Omega p(\omega)d\mu(\omega) \geq V(\mu) \text{ for all } \mu \in \Delta(\Omega),$$
$$\int_\Omega p(\omega)d\mu(\omega) = V(\mu) \text{ for all } \mu \in \text{supp}(\tau).$$

Full disclosure

- Full disclosure τ_F satisfies $\mathsf{supp}(\tau_F) = \{\delta_\omega : \omega \in \Omega\}$.

Full disclosure

- Full disclosure τ_F satisfies $\mathsf{supp}(\tau_F) = \{\delta_\omega : \omega \in \Omega\}$.
- By weak duality, full disclosure is optimal if

$$\int_\Omega V(\delta_\omega)\mathrm{d}\mu(\omega) \geq V(\mu) \text{ for all } \mu \in \Delta(\Omega).$$

Full disclosure

- Full disclosure τ_F satisfies $\text{supp}(\tau_F) = \{\delta_\omega : \omega \in \Omega\}$.
- By weak duality, full disclosure is optimal if

$$\int_\Omega V(\delta_\omega)\mathrm{d}\mu(\omega) \geq V(\mu) \text{ for all } \mu \in \Delta(\Omega). \tag{F}$$

- Indeed, suppose (F) holds. Given a price function $p(\omega) = V(\delta_\omega)$, all operating operating processes make zero profits, and no entrant can make strictly positive profits, by (F).

Full disclosure

- Full disclosure τ_F satisfies $\text{supp}(\tau_F) = \{\delta_\omega : \omega \in \Omega\}$.
- By weak duality, full disclosure is optimal if

$$\int_\Omega V(\delta_\omega)\mathrm{d}\mu(\omega) \geq V(\mu) \text{ for all } \mu \in \Delta(\Omega). \tag{F}$$

- Indeed, suppose (F) holds. Given a price function $p(\omega) = V(\delta_\omega)$, all operating operating processes make zero profits, and no entrant can make strictly positive profits, by (F).
- By strong duality, if full disclosure is optimal, then (F) holds.

Full disclosure

- Full disclosure τ_F satisfies $\text{supp}(\tau_F) = \{\delta_\omega : \omega \in \Omega\}$.
- By weak duality, full disclosure is optimal if

$$\int_\Omega V(\delta_\omega)\mathrm{d}\mu(\omega) \geq V(\mu) \text{ for all } \mu \in \Delta(\Omega). \tag{F}$$

- Indeed, suppose (F) holds. Given a price function $p(\omega) = V(\delta_\omega)$, all operating operating processes make zero profits, and no entrant can make strictly positive profits, by (F).
- By strong duality, if full disclosure is optimal, then (F) holds.
- Indeed, suppose τ_F is optimal. Since all operating processes make zero profit, the price function is given by $p(\omega) = V(\delta_\omega)$. Since no entrant can make strictly positive profits, (F) holds.

No disclosure

- No disclosure τ_N satisfies $\mathsf{supp}(\tau_N) = \{\mu_0\}$.

No disclosure

- No disclosure τ_N satisfies $\text{supp}(\tau_N) = \{\mu_0\}$.
- By weak duality, no disclosure is optimal if there exists a price function $p \in L(\Omega)$ such that

$$\int_\Omega p(\omega)\mathrm{d}\mu_0(\omega) = V(\mu_0),$$
$$\int_\Omega p(\omega)\mathrm{d}\mu(\omega) \geq V(\mu), \text{ for all } \mu \in \Delta(\Omega).$$

No disclosure

- No disclosure τ_N satisfies $\mathrm{supp}(\tau_N) = \{\mu_0\}$.

- By weak duality, no disclosure is optimal if there exists a price function $p \in L(\Omega)$ such that

$$\begin{aligned} \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega) &= V(\mu_0), \\ \int_\Omega p(\omega)\mathrm{d}\mu(\omega) &\geq V(\mu), \text{ for all } \mu \in \Delta(\Omega). \end{aligned} \tag{N}$$

- By strong duality, if no disclosure is optimal, then (N) holds.

Duality

- **Weak duality** states that $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$, that is, the concave closure is bounded above by the concave envelope.

Duality

- **Weak duality** states that $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$, that is, the concave closure is bounded above by the concave envelope.
- **No duality gap** requires the equality $\widehat{V}(\mu_0) = \overline{V}(\mu_0)$, that is, that the concave closure and the concave envelope coincide.

Duality

- **Weak duality** states that $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$, that is, the concave closure is bounded above by the concave envelope.

- **No duality gap** requires the equality $\widehat{V}(\mu_0) = \overline{V}(\mu_0)$, that is, that the concave closure and the concave envelope coincide.

- **Primal and dual attainment** additionally require existence of solutions to the primal and dual problems, respectively.

Duality

- **Weak duality** states that $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$, that is, the concave closure is bounded above by the concave envelope.

- **No duality gap** requires the equality $\widehat{V}(\mu_0) = \overline{V}(\mu_0)$, that is, that the concave closure and the concave envelope coincide.

- **Primal and dual attainment** additionally require existence of solutions to the primal and dual problems, respectively.

- We use the term **strong duality** when there is no duality gap and both primal and dual attainment hold.

Weak Duality

Weak Duality. $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$.

Weak Duality

Weak Duality. $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$.

Proof: For all $\tau \in \mathcal{T}(\mu_0)$ and $p \in \mathcal{P}(V)$, we have

$$\int_{\Delta(\Omega)} V(\mu)\mathrm{d}\tau(\mu) \leq \int_{\Delta(\Omega)} \int_{\Omega} p(\omega)\mathrm{d}\mu(\omega)\mathrm{d}\tau(\mu) = \int_{\Omega} p(\omega)\mathrm{d}\mu_0(\omega).$$

Finally, take the supremum over $\mathcal{T}(\mu_0)$ and the infimum over $\mathcal{P}(V)$.

Weak Duality

Weak Duality. $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$.

Proof: For all $\tau \in \mathcal{T}(\mu_0)$ and $p \in \mathcal{P}(V)$, we have

$$\int_{\Delta(\Omega)} V(\mu)\mathrm{d}\tau(\mu) \leq \int_{\Delta(\Omega)} \int_{\Omega} p(\omega)\mathrm{d}\mu(\omega)\mathrm{d}\tau(\mu) = \int_{\Omega} p(\omega)\mathrm{d}\mu_0(\omega).$$

Finally, take the supremum over $\mathcal{T}(\mu_0)$ and the infimum over $\mathcal{P}(V)$.

Remark. Weak duality holds even if V is not u.s.c., as long as it is measurable and bounded.

Weak Duality

Weak Duality. $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$.

Proof: For all $\tau \in \mathcal{T}(\mu_0)$ and $p \in \mathcal{P}(V)$, we have

$$\int_{\Delta(\Omega)} V(\mu)\mathrm{d}\tau(\mu) \leq \int_{\Delta(\Omega)} \int_{\Omega} p(\omega)\mathrm{d}\mu(\omega)\mathrm{d}\tau(\mu) = \int_{\Omega} p(\omega)\mathrm{d}\mu_0(\omega).$$

Finally, take the supremum over $\mathcal{T}(\mu_0)$ and the infimum over $\mathcal{P}(V)$.

Remark. Weak duality holds even if V is not u.s.c., as long as it is measurable and bounded.

Optimality criterion. If $\tau \in \mathcal{T}(\mu_0)$ and $p \in \mathcal{P}(V)$ satisfy $\int V\mathrm{d}\tau = \int p\mathrm{d}\mu_0$, then τ attains $\widehat{V}(\mu_0)$ and p attains $\overline{V}(\mu_0)$,

Weak Duality

Weak Duality. $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$.

Proof: For all $\tau \in \mathcal{T}(\mu_0)$ and $p \in \mathcal{P}(V)$, we have

$$\int_{\Delta(\Omega)} V(\mu)\mathrm{d}\tau(\mu) \leq \int_{\Delta(\Omega)} \int_{\Omega} p(\omega)\mathrm{d}\mu(\omega)\mathrm{d}\tau(\mu) = \int_{\Omega} p(\omega)\mathrm{d}\mu_0(\omega).$$

Finally, take the supremum over $\mathcal{T}(\mu_0)$ and the infimum over $\mathcal{P}(V)$.

Remark. Weak duality holds even if V is not u.s.c., as long as it is measurable and bounded.

Optimality criterion. If $\tau \in \mathcal{T}(\mu_0)$ and $p \in \mathcal{P}(V)$ satisfy $\int V\mathrm{d}\tau = \int p\mathrm{d}\mu_0$, then τ attains $\widehat{V}(\mu_0)$ and p attains $\overline{V}(\mu_0)$, as, for any $\tilde{\tau} \in \mathcal{T}(\mu)$ and $\tilde{p} \in \mathcal{P}(V)$,

$$\int V\mathrm{d}\tilde{\tau} \leq \int p\mathrm{d}\mu_0 = \int V\mathrm{d}\tau \quad \text{and} \quad \int \tilde{p}\mathrm{d}\mu_0 \geq \int V\mathrm{d}\tau = \int p\mathrm{d}\mu_0.$$

Primal attainment

Primal attainment. The primal problem has an optimal solution.

Proof: Follows from the Weierstrass theorem, because we maximize bounded upper semi-continuous function $\tau \to \int V \mathrm{d}\tau$ on compact set $\mathcal{T}(\mu_0)$. $\mathcal{T}(\mu_0)$ is compact, because it is a closed subset (by continuity of $\tau \to \int \mu \mathrm{d}\tau$) of a compact set $\Delta(\Delta(\Omega))$.

No duality gap

No duality gap. $\widehat{V}(\mu_0) = \overline{V}(\mu_0)$.

No duality gap

No duality gap. $\widehat{V}(\mu_0) = \overline{V}(\mu_0)$.

Proof outline: Uses the Fenchel-Moreau theorem.

No duality gap

No duality gap. $\widehat{V}(\mu_0) = \overline{V}(\mu_0)$.

Proof outline: Uses the Fenchel-Moreau theorem.

- Let E be a normed vector space and $E^\star$ its topological dual.

No duality gap

No duality gap. $\widehat{V}(\mu_0) = \overline{V}(\mu_0)$.

Proof outline: Uses the Fenchel-Moreau theorem.

- Let E be a normed vector space and $E^\star$ its topological dual.
- Let $\varphi : E \to \mathbb{R} \cup \{+\infty\}$, and let $\varphi^\star : E^\star \to \mathbb{R} \cup \{+\infty\}$ be its Legendre transform.

No duality gap

No duality gap. $\widehat{V}(\mu_0) = \overline{V}(\mu_0)$.

Proof outline: Uses the Fenchel-Moreau theorem.

- Let E be a normed vector space and $E^\star$ its topological dual.
- Let $\varphi : E \to \mathbb{R} \cup \{+\infty\}$, and let $\varphi^\star : E^\star \to \mathbb{R} \cup \{+\infty\}$ be its Legendre transform.
- Fenchel-Moreau Theorem: If $\varphi : E \to \mathbb{R} \cup \{+\infty\}$ is convex and lower semi-continuous, then $\varphi^{\star\star} = \varphi$.

No duality gap

Proof outline continued:

- Define the function φ on $M(\Omega)$ as

$$\varphi(\eta) = \begin{cases} -\sup_{\tau \in \mathcal{T}(\eta)} \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu), & \eta \in \Delta(\Omega), \\ +\infty, & \eta \notin \Delta(\Omega). \end{cases}$$

No duality gap

Proof outline continued:

- Define the function φ on $M(\Omega)$ as

$$\varphi(\eta) = \begin{cases} -\sup_{\tau \in \mathcal{T}(\eta)} \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu), & \eta \in \Delta(\Omega), \\ +\infty, & \eta \notin \Delta(\Omega). \end{cases}$$

- By direct computation, we have for $\mu_0 \in \Delta(\Omega)$,

$$\varphi^{\star\star}(\mu_0) = -\inf_{p \in \mathcal{P}(V)} \left\{ \int_\Omega p(\omega) \mathrm{d}\mu_0(\omega) \right\}.$$

No duality gap

Proof outline continued:

- Define the function φ on $M(\Omega)$ as

$$\varphi(\eta) = \begin{cases} -\sup_{\tau \in \mathcal{T}(\eta)} \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu), & \eta \in \Delta(\Omega), \\ +\infty, & \eta \notin \Delta(\Omega). \end{cases}$$

- By direct computation, we have for $\mu_0 \in \Delta(\Omega)$,

$$\varphi^{\star\star}(\mu_0) = -\inf_{p \in \mathcal{P}(V)} \left\{ \int_\Omega p(\omega) \mathrm{d}\mu_0(\omega) \right\}.$$

- By the F-M theorem, we get $\varphi = \varphi^{\star\star}$ on $\Delta(\Omega)$.

Dual attainment

$\widehat{V}$ is *superdifferentiable* at μ_0 if there is a continuous linear function H on $M(\Omega)$ (*supporting hyperplane* of $\widehat{V}$ at μ_0), represented as $H(\mu) = \int p \mathrm{d}\mu$ with $p \in L(\Omega)$ (*supergradient* of $\widehat{V}$ at μ_0), such that

$$\widehat{V}(\mu_0) = H(\mu_0),$$
$$\widehat{V}(\mu) \leq H(\mu), \quad \text{for all } \mu \in \Delta(\Omega).$$

Dual attainment

$\widehat{V}$ is *superdifferentiable* at μ_0 if there is a continuous linear function H on $M(\Omega)$ (*supporting hyperplane* of $\widehat{V}$ at μ_0), represented as $H(\mu) = \int p \mathrm{d}\mu$ with $p \in L(\Omega)$ (*supergradient* of $\widehat{V}$ at μ_0), such that

$$\begin{aligned} \widehat{V}(\mu_0) &= H(\mu_0), \\ \widehat{V}(\mu) &\leq H(\mu), \quad \text{for all } \mu \in \Delta(\Omega). \end{aligned}$$

When Ω is finite, the concave function $\widehat{V}$ on $\Delta(\Omega)$ is continuous on the interior of the domain, and superdifferentiable at all interior points (full-support priors).

Dual attainment

$\widehat{V}$ is *superdifferentiable* at μ_0 if there is a continuous linear function H on $M(\Omega)$ (*supporting hyperplane* of $\widehat{V}$ at μ_0), represented as $H(\mu) = \int p \mathrm{d}\mu$ with $p \in L(\Omega)$ (*supergradient* of $\widehat{V}$ at μ_0), such that

$$\begin{aligned} \widehat{V}(\mu_0) &= H(\mu_0), \\ \widehat{V}(\mu) &\leq H(\mu), \quad \text{for all } \mu \in \Delta(\Omega). \end{aligned}$$

When Ω is finite, the concave function $\widehat{V}$ on $\Delta(\Omega)$ is continuous on the interior of the domain, and superdifferentiable at all interior points (full-support priors).

When Ω is infinite,

- concavity does not imply continuity on the interior of the domain,
- the set $\Delta(\Omega)$ has an empty (relative) interior.

Dual attainment

Following Gale (1967), we say that $\widehat{V}$ has *bounded steepness* at μ_0 if there exists a constant L such that

$$\widehat{V}(\mu) - \widehat{V}(\mu_0) \leq L\|\mu - \mu_0\|_{KR}, \quad \text{for all } \mu \in \Delta(\Omega).$$

Dual attainment

Following Gale (1967), we say that $\widehat{V}$ has *bounded steepness* at μ_0 if there exists a constant L such that

$$\widehat{V}(\mu) - \widehat{V}(\mu_0) \leq L\|\mu - \mu_0\|_{KR}, \quad \text{for all } \mu \in \Delta(\Omega).$$

Intuitively, bounded steepness says that the marginal increase in the value of the persuasion problem is bounded above for a small perturbation of the prior.

Dual attainment

Dual Attainment. The following statements are equivalent:

1. The dual problem has an optimal solution.
2. $\widehat{V}$ is superdifferentiable at μ_0.
3. $\widehat{V}$ has bounded steepness at μ_0.

Dual attainment

Dual Attainment. The following statements are equivalent:

1. The dual problem has an optimal solution.
2. $\widehat{V}$ is superdifferentiable at μ_0.
3. $\widehat{V}$ has bounded steepness at μ_0.

Takeaway: Duality holds without any extra assumptions in finite state spaces, but additional regularity conditions are needed otherwise.

Dual attainment

Proof of dual attainment:

- Because we represent the weak* topology using the norm $\|\cdot\|_{KR}$, we can use the result of Gale (1967).

Dual attainment

Proof of dual attainment:

- Because we represent the weak* topology using the norm $\|\cdot\|_{KR}$, we can use the result of Gale (1967).
- By Duality Theorem in Gale (1967), $\widehat{V}$ is superdifferentiable at μ_0 if and only if $\widehat{V}$ has bounded steepness at μ_0.

Dual attainment

Proof of dual attainment:

- Because we represent the weak* topology using the norm $\|\cdot\|_{KR}$, we can use the result of Gale (1967).

- By Duality Theorem in Gale (1967), $\widehat{V}$ is superdifferentiable at μ_0 if and only if $\widehat{V}$ has bounded steepness at μ_0.

- Remains to prove that existence of solution to the dual problem is equivalent to superdifferentiability of $\widehat{V}$ at the prior μ_0.

Dual attainment

$\widehat{V}$ is superdifferentiable at $\mu_0 \implies$ dual attainment:

- Since $\widehat{V}$ is superdifferentiable at μ_0, there exists a continuous linear function H on $M(\Omega)$ such that

$$\widehat{V}(\mu_0) = H(\mu_0),$$
$$\widehat{V}(\mu) \leq H(\mu), \text{ for all } \mu \in \Delta(\Omega).$$

Dual attainment

$\widehat{V}$ is superdifferentiable at $\mu_0 \implies$ dual attainment:

- Since $\widehat{V}$ is superdifferentiable at μ_0, there exists a continuous linear function H on $M(\Omega)$ such that

$$\widehat{V}(\mu_0) = H(\mu_0),$$
$$\widehat{V}(\mu) \leq H(\mu), \text{ for all } \mu \in \Delta(\Omega).$$

- Since $(M(\Omega), \|\cdot\|_{KR})^\star = L(\Omega)$, there exists $p \in L(\Omega)$ such that...

Dual attainment

$\widehat{V}$ is superdifferentiable at $\mu_0 \implies$ dual attainment:

- Since $\widehat{V}$ is superdifferentiable at μ_0, there exists a continuous linear function H on $M(\Omega)$ such that

$$\widehat{V}(\mu_0) = \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega),$$
$$\widehat{V}(\mu) \leq \int_\Omega p(\omega)\mathrm{d}\mu(\omega), \text{ for all } \mu \in \Delta(\Omega).$$

- Since $(M(\Omega), \|\cdot\|_{KR})^\star = L(\Omega)$, there exists $p \in L(\Omega)$ such that...

Dual attainment

$\widehat{V}$ *is superdifferentiable at* $\mu_0 \implies$ *dual attainment:*

- Since $\widehat{V}$ is superdifferentiable at μ_0, there exists a continuous linear function H on $M(\Omega)$ such that

$$\widehat{V}(\mu_0) = \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega),$$
$$\widehat{V}(\mu) \leq \int_\Omega p(\omega)\mathrm{d}\mu(\omega), \text{ for all } \mu \in \Delta(\Omega).$$

- Since $(M(\Omega), \|\cdot\|_{KR})^\star = L(\Omega)$, there exists $p \in L(\Omega)$ such that...
- Since $V(\mu) \leq \widehat{V}(\mu)$, the second line says that p is feasible; the first line says that it achieves the lower bound, so p is optimal.

Dual attainment

$\widehat{V}$ is superdifferentiable at $\mu_0 \implies$ dual attainment:

- Since $\widehat{V}$ is superdifferentiable at μ_0, there exists a continuous linear function H on $M(\Omega)$ such that

$$\widehat{V}(\mu_0) = \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega),$$
$$\widehat{V}(\mu) \leq \int_\Omega p(\omega)\mathrm{d}\mu(\omega), \text{ for all } \mu \in \Delta(\Omega).$$

- Since $(M(\Omega), \|\cdot\|_{KR})^\star = L(\Omega)$, there exists $p \in L(\Omega)$ such that...
- Since $V(\mu) \leq \widehat{V}(\mu)$, the second line says that p is feasible; the first line says that it achieves the lower bound, so p is optimal.
- Main takeaway: The optimal price function p is a supergradient of the concave closure $\widehat{V}$ at the prior μ_0.

Dual attainment

Dual attainment $\implies \widehat{V}$ *is superdifferentiable at* μ_0*:*

- Since $p \in \mathcal{P}(V)$ is optimal, we have $p \in L(\Omega)$ and

$$\overline{V}(\mu_0) = \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega),$$
$$V(\mu) \leq \int_\Omega p(\omega)\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega).$$

Dual attainment

Dual attainment $\implies$ $\widehat{V}$ *is superdifferentiable at* μ_0*:*

- Since $p \in \mathcal{P}(V)$ is optimal, we have $p \in L(\Omega)$ and

$$\begin{aligned} \overline{V}(\mu_0) &= \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega), \\ V(\mu) &\leq \int_\Omega p(\omega)\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega). \end{aligned}$$

- By the absence of duality gap, $\overline{V}(\mu_0) = \widehat{V}(\mu_0)$.

Dual attainment

Dual attainment $\implies$ $\widehat{V}$ *is superdifferentiable at* μ_0*:*

- Since $p \in \mathcal{P}(V)$ is optimal, we have $p \in L(\Omega)$ and

$$\begin{aligned} \overline{V}(\mu_0) &= \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega), \\ V(\mu) &\leq \int_\Omega p(\omega)\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega). \end{aligned}$$

- By the absence of duality gap, $\overline{V}(\mu_0) = \widehat{V}(\mu_0)$.

- By weak duality, $\widehat{V}(\mu) \leq \int_\Omega p(\omega)\mathrm{d}\mu(\omega)$ for all $\mu \in \Delta(\Omega)$, so...

Dual attainment

Dual attainment $\Longrightarrow$ $\widehat{V}$ *is superdifferentiable at* μ_0*:*

- Since $p \in \mathcal{P}(V)$ is optimal, we have $p \in L(\Omega)$ and

$$\widehat{V}(\mu_0) = \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega),$$
$$\widehat{V}(\mu) \leq \int_\Omega p(\omega)\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega).$$

- By the absence of duality gap, $\overline{V}(\mu_0) = \widehat{V}(\mu_0)$.

- By weak duality, $\widehat{V}(\mu) \leq \int_\Omega p(\omega)\mathrm{d}\mu(\omega)$ for all $\mu \in \Delta(\Omega)$, so...

Dual attainment

Dual attainment $\implies$ $\widehat{V}$ *is superdifferentiable at* μ_0*:*

- Since $p \in \mathcal{P}(V)$ is optimal, we have $p \in L(\Omega)$ and

$$\widehat{V}(\mu_0) = \int_\Omega p(\omega) \mathrm{d}\mu_0(\omega),$$
$$\widehat{V}(\mu) \leq \int_\Omega p(\omega) \mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega).$$

- By the absence of duality gap, $\overline{V}(\mu_0) = \widehat{V}(\mu_0)$.

- By weak duality, $\widehat{V}(\mu) \leq \int_\Omega p(\omega) \mathrm{d}\mu(\omega)$ for all $\mu \in \Delta(\Omega)$, so...

- Thus, p is a supergradient of $\widehat{V}$ at μ_0.

Duality

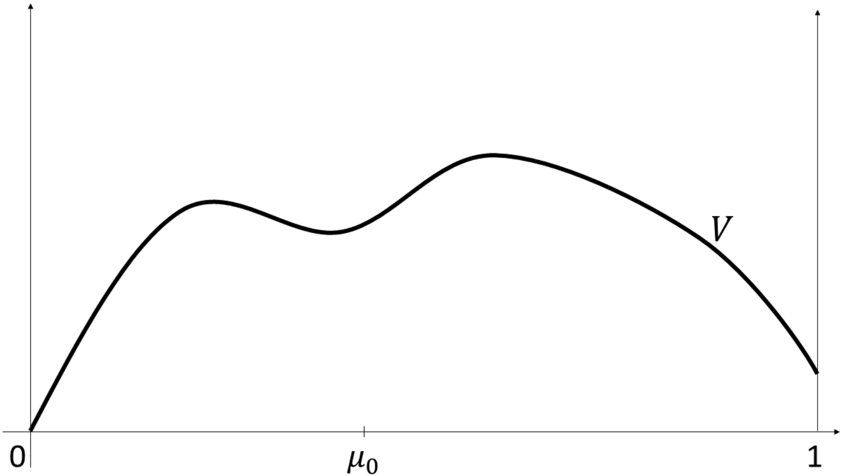

Duality

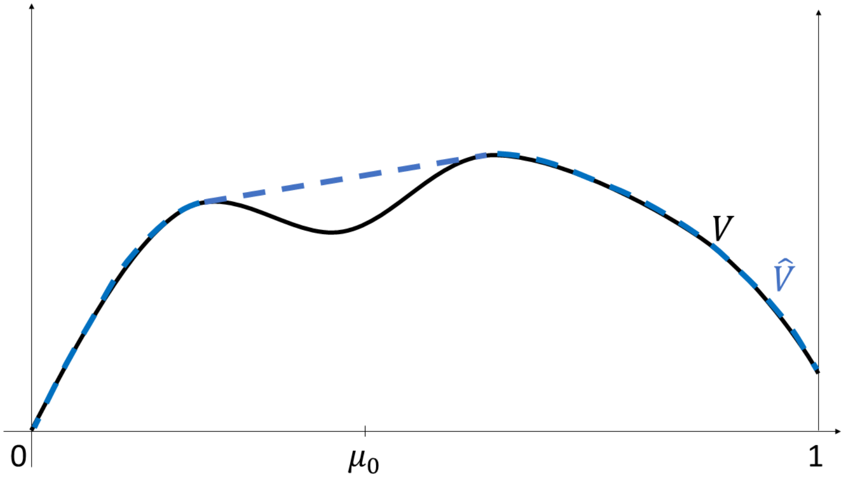

Duality

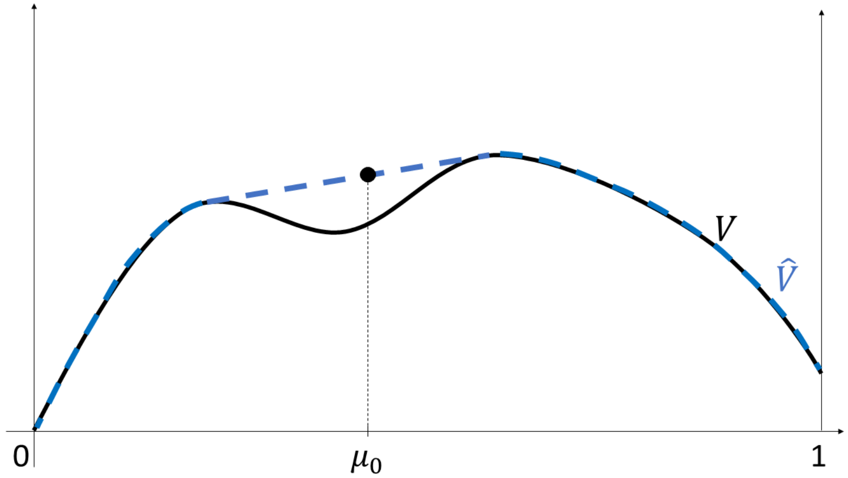

Duality

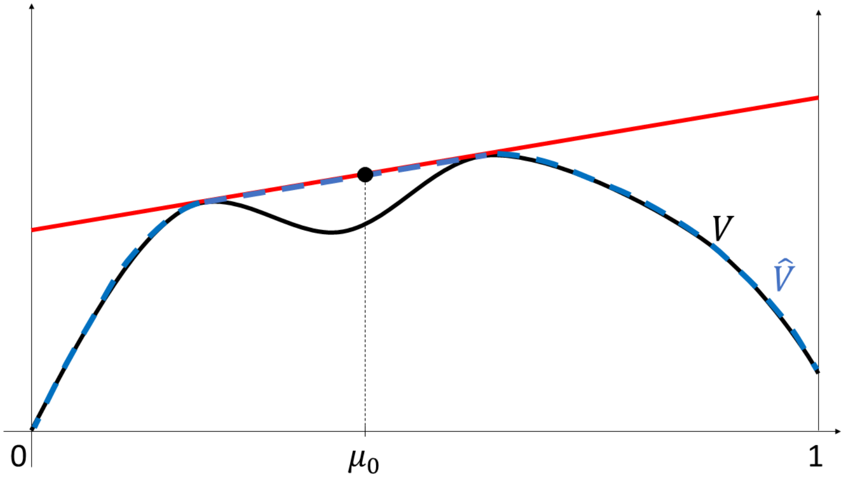

Duality

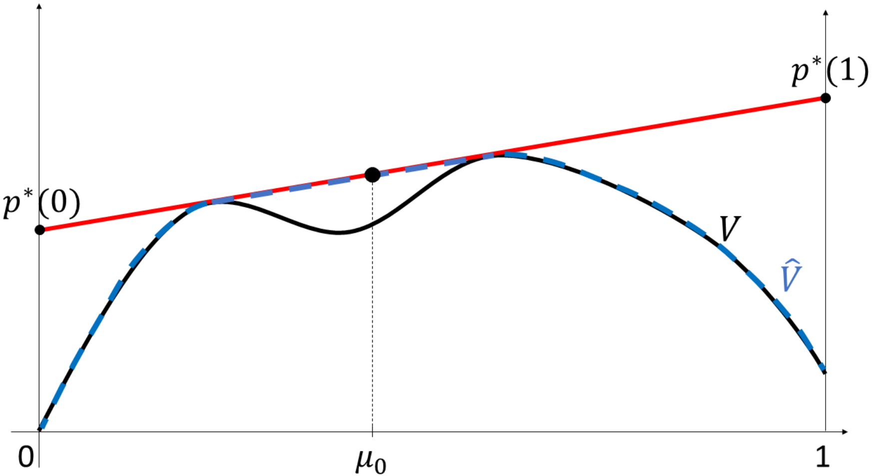

Duality

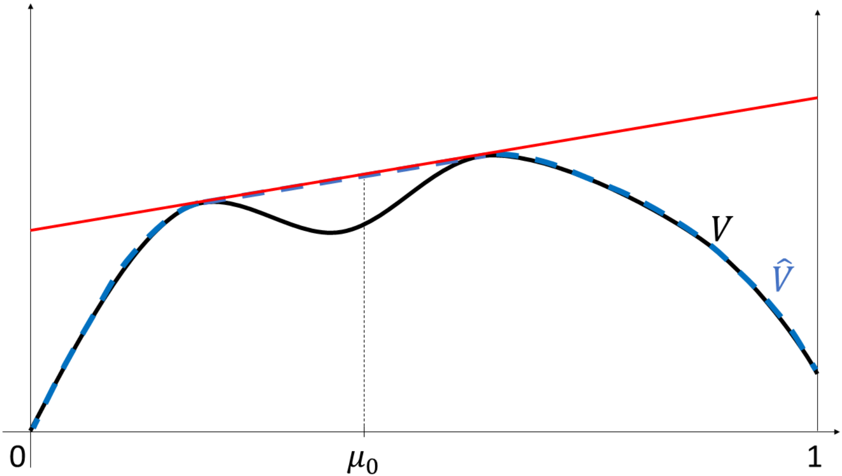

Duality

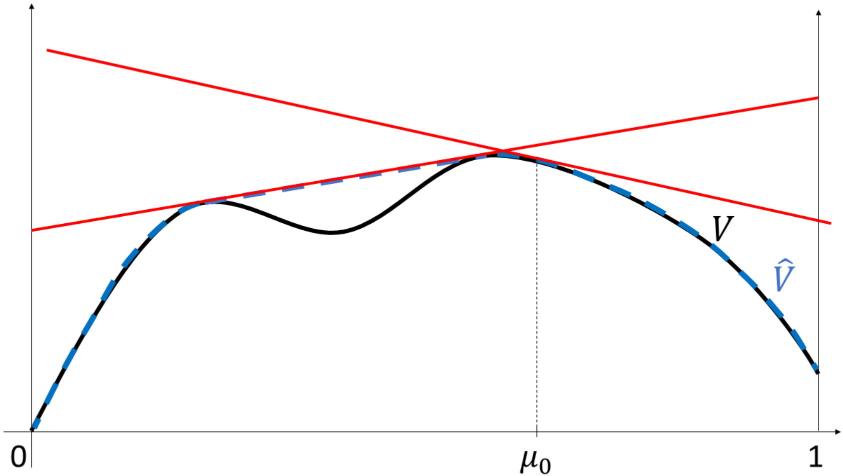

Duality

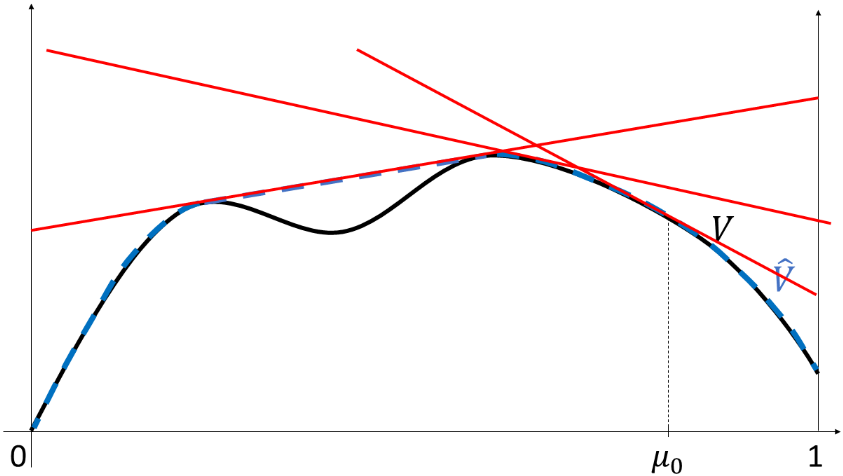

Duality

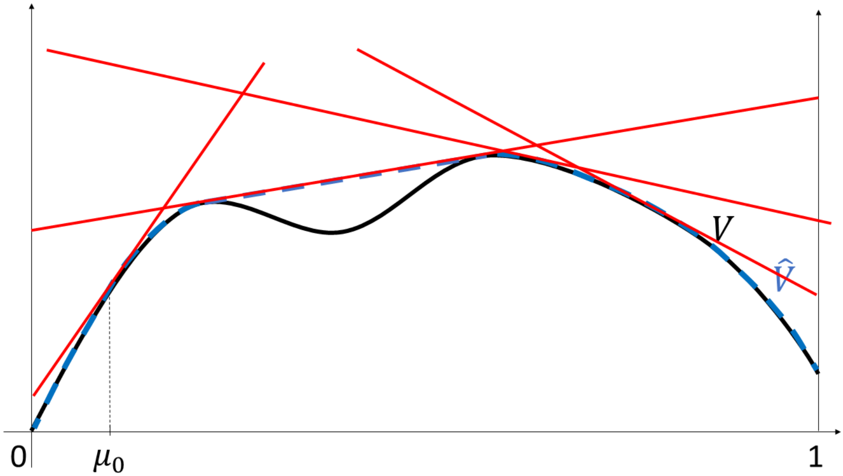

Duality

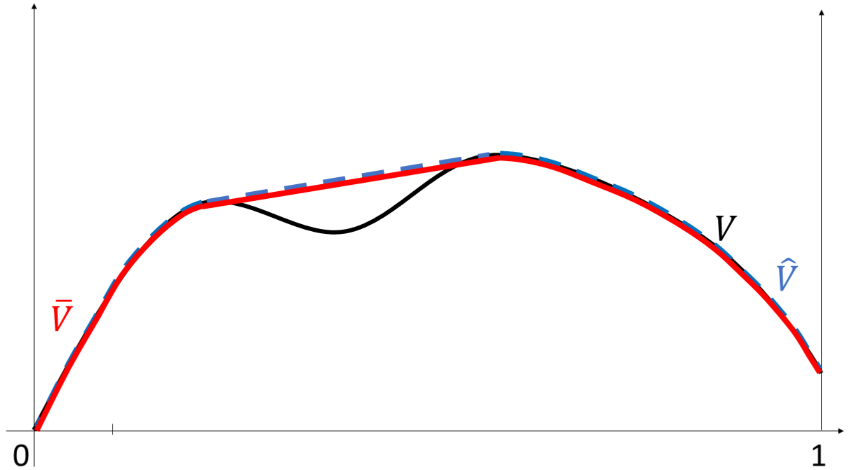

Lipschitz Preservation

Checking bounded steepness of $\widehat{V}$ at μ_0 can be quite difficult.

Lipschitz Preservation

Checking bounded steepness of $\widehat{V}$ at μ_0 can be quite difficult.

Lipschitz Preservation. If V is Lipschitz on $\Delta(\Omega)$, then so is $\widehat{V}$. Thus, $\widehat{V}$ has bounded steepness at each $\mu_0 \in \Delta(\Omega)$, ensuring that the dual problem has an optimal solution.

Lipschitz preservation

Lemma Let $\mu_0, \eta_0 \in \Delta(\Omega)$ and $\tau \in \mathcal{T}(\mu_0)$.

Lipschitz preservation

Lemma Let $\mu_0, \eta_0 \in \Delta(\Omega)$ and $\tau \in \mathcal{T}(\mu_0)$. There exists measurable $\eta : \Delta(\Omega) \to \Delta(\Omega)$ such that

$$\int_{\Delta(\Omega)} \eta(\mu) \mathrm{d}\tau(\mu) = \eta_0 \quad \text{and} \quad \int_{\Delta(\Omega)} \|\mu - \eta(\mu)\|_{KR} \, \mathrm{d}\tau(\mu) = \|\mu_0 - \eta_0\|_{KR} .$$

Lipschitz preservation

Lemma Let $\mu_0, \eta_0 \in \Delta(\Omega)$ and $\tau \in \mathcal{T}(\mu_0)$. There exists measurable $\eta : \Delta(\Omega) \to \Delta(\Omega)$ such that

$$\int_{\Delta(\Omega)} \eta(\mu) \mathrm{d}\tau(\mu) = \eta_0 \quad \text{and} \quad \int_{\Delta(\Omega)} \|\mu - \eta(\mu)\|_{KR} \, \mathrm{d}\tau(\mu) = \|\mu_0 - \eta_0\|_{KR} \, .$$

This gives us the conclusion because, for τ attaining $\widehat{V}(\mu_0)$, we have

Lipschitz preservation

Lemma Let $\mu_0, \eta_0 \in \Delta(\Omega)$ and $\tau \in \mathcal{T}(\mu_0)$. There exists measurable $\eta : \Delta(\Omega) \to \Delta(\Omega)$ such that

$$\int_{\Delta(\Omega)} \eta(\mu) \mathrm{d}\tau(\mu) = \eta_0 \quad \text{and} \quad \int_{\Delta(\Omega)} \|\mu - \eta(\mu)\|_{KR} \, \mathrm{d}\tau(\mu) = \|\mu_0 - \eta_0\|_{KR} \, .$$

This gives us the conclusion because, for τ attaining $\widehat{V}(\mu_0)$, we have

$$\widehat{V}(\mu_0) - \widehat{V}(\eta_0)$$

Lipschitz preservation

Lemma Let $\mu_0, \eta_0 \in \Delta(\Omega)$ and $\tau \in \mathcal{T}(\mu_0)$. There exists measurable $\eta : \Delta(\Omega) \to \Delta(\Omega)$ such that

$$\int_{\Delta(\Omega)} \eta(\mu) \mathrm{d}\tau(\mu) = \eta_0 \quad \text{and} \quad \int_{\Delta(\Omega)} \|\mu - \eta(\mu)\|_{KR} \, \mathrm{d}\tau(\mu) = \|\mu_0 - \eta_0\|_{KR} .$$

This gives us the conclusion because, for τ attaining $\widehat{V}(\mu_0)$, we have

$$\widehat{V}(\mu_0) - \widehat{V}(\eta_0) \leq \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu)$$

Lipschitz preservation

Lemma Let $\mu_0, \eta_0 \in \Delta(\Omega)$ and $\tau \in \mathcal{T}(\mu_0)$. There exists measurable $\eta : \Delta(\Omega) \to \Delta(\Omega)$ such that

$$\int_{\Delta(\Omega)} \eta(\mu) \mathrm{d}\tau(\mu) = \eta_0 \quad \text{and} \quad \int_{\Delta(\Omega)} \|\mu - \eta(\mu)\|_{KR} \, \mathrm{d}\tau(\mu) = \|\mu_0 - \eta_0\|_{KR} \, .$$

This gives us the conclusion because, for τ attaining $\widehat{V}(\mu_0)$, we have

$$\widehat{V}(\mu_0) - \widehat{V}(\eta_0) \leq \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu) - \int_{\Delta(\Omega)} V(\eta(\mu)) \mathrm{d}\tau(\mu)$$

Lipschitz preservation

Lemma Let $\mu_0, \eta_0 \in \Delta(\Omega)$ and $\tau \in \mathcal{T}(\mu_0)$. There exists measurable $\eta : \Delta(\Omega) \to \Delta(\Omega)$ such that

$$\int_{\Delta(\Omega)} \eta(\mu) \mathrm{d}\tau(\mu) = \eta_0 \quad \text{and} \quad \int_{\Delta(\Omega)} \|\mu - \eta(\mu)\|_{KR} \, \mathrm{d}\tau(\mu) = \|\mu_0 - \eta_0\|_{KR} .$$

This gives us the conclusion because, for τ attaining $\widehat{V}(\mu_0)$, we have

$$\begin{aligned} \widehat{V}(\mu_0) - \widehat{V}(\eta_0) &\leq \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu) - \int_{\Delta(\Omega)} V(\eta(\mu)) \mathrm{d}\tau(\mu) \\ &\leq \int_{\Delta(\Omega)} L \, \|\mu - \eta(\mu)\|_{KR} \, \mathrm{d}\tau(\mu) \end{aligned}$$

Lipschitz preservation

Lemma Let $\mu_0, \eta_0 \in \Delta(\Omega)$ and $\tau \in \mathcal{T}(\mu_0)$. There exists measurable $\eta : \Delta(\Omega) \to \Delta(\Omega)$ such that

$$\int_{\Delta(\Omega)} \eta(\mu) \mathrm{d}\tau(\mu) = \eta_0 \quad \text{and} \quad \int_{\Delta(\Omega)} \|\mu - \eta(\mu)\|_{KR} \, \mathrm{d}\tau(\mu) = \|\mu_0 - \eta_0\|_{KR} \, .$$

This gives us the conclusion because, for τ attaining $\widehat{V}(\mu_0)$, we have

$$\begin{aligned} \widehat{V}(\mu_0) - \widehat{V}(\eta_0) &\leq \int_{\Delta(\Omega)} V(\mu) \mathrm{d}\tau(\mu) - \int_{\Delta(\Omega)} V(\eta(\mu)) \mathrm{d}\tau(\mu) \\ &\leq \int_{\Delta(\Omega)} L \, \|\mu - \eta(\mu)\|_{KR} \, \mathrm{d}\tau(\mu) = L \, \|\mu_0 - \eta_0\|_{KR} \, . \end{aligned}$$

Complementary slackness

Complementary slackness. Let V be Lipschitz. Distribution $\tau \in \mathcal{T}(\mu_0)$ is an optimal solution iff there exists $p \in \mathcal{P}(V)$ such that

$$V(\mu) = \int_\Omega p(\omega)\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \mathrm{supp}(\tau). \tag{C}$$

Complementary slackness

Complementary slackness. Let V be Lipschitz. Distribution $\tau \in \mathcal{T}(\mu_0)$ is an optimal solution iff there exists $p \in \mathcal{P}(V)$ such that

$$V(\mu) = \int_\Omega p(\omega)\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \mathrm{supp}(\tau). \tag{C}$$

Proof: The if part follows from weak duality: If (C) holds, then

$$\int_{\Delta(\Omega)} V(\mu)\mathrm{d}\tau(\mu) = \int_{\Delta(\Omega)} \int_\Omega p(\omega)\mathrm{d}\mu(\omega)\mathrm{d}\tau(\mu) = \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega).$$

The only if part follows from no duality gap and dual attainment: If $\tau \in \mathcal{T}(\mu_0)$ is optimal, then there exists an optimal $p \in \mathcal{P}(V)$ such that

$$\int_{\Delta(\Omega)} \left(\int_\Omega p(\omega)\mathrm{d}\mu(\omega) - V(\mu) \right) \mathrm{d}\tau(\mu) = 0.$$

Persuasion and Matching

Anton Kolotilin & Roberto Corrao & Alexander Wolitzky

Non-linear persuasion

In *non-linear persuasion* of Kolotilin, Corrao, and Wolitzky (2023), states and actions are one-dimensional but the utilities are not linear in the state.

Non-linear persuasion

In *non-linear persuasion* of Kolotilin, Corrao, and Wolitzky (2023), states and actions are one-dimensional but the utilities are not linear in the state.

Duality theorem characterizes optimal signals in terms of the value of matching different states and actions together.

Non-linear persuasion

In *non-linear persuasion* of Kolotilin, Corrao, and Wolitzky (2023), states and actions are one-dimensional but the utilities are not linear in the state.

Duality theorem characterizes optimal signals in terms of the value of matching different states and actions together.

- Similar to classical matching or optimal transport, except supply on one side of the market (marginal over actions) is endogenous and determined by receiver's obedience condition.

Non-linear persuasion

In *non-linear persuasion* of Kolotilin, Corrao, and Wolitzky (2023), states and actions are one-dimensional but the utilities are not linear in the state.

Duality theorem characterizes optimal signals in terms of the value of matching different states and actions together.

- Similar to classical matching or optimal transport, except supply on one side of the market (marginal over actions) is endogenous and determined by receiver's obedience condition.

- Useful for characterizing key properties of optimal signals, or even a unique optimal signal.

Non-linear persuasion

In *non-linear persuasion* of Kolotilin, Corrao, and Wolitzky (2023), states and actions are one-dimensional but the utilities are not linear in the state.

Duality theorem characterizes optimal signals in terms of the value of matching different states and actions together.

- Similar to classical matching or optimal transport, except supply on one side of the market (marginal over actions) is endogenous and determined by receiver's obedience condition.
- Useful for characterizing key properties of optimal signals, or even a unique optimal signal.
- For example, each optimal signal is pairwise under a non-singularity condition on utilities.

Non-linear persuasion

- One-dimensional state space $\Omega = [0, 1]$ and action space $A = [0, 1]$

Non-linear persuasion

- One-dimensional state space $\Omega = [0, 1]$ and action space $A = [0, 1]$

- Twice continuously differentiable Sender's utility $v(a, \omega)$ and Receiver's utility $u(a, \omega)$

Non-linear persuasion

- One-dimensional state space $\Omega = [0, 1]$ and action space $A = [0, 1]$
- Twice continuously differentiable Sender's utility $v(a, \omega)$ and Receiver's utility $u(a, \omega)$
- Strictly concave Receiver's utility ($u_{aa}(a, \omega) < 0$) with interior optimal action $a^*(\mu)$ given by $\int_\Omega u_a(a^*(\mu), \omega)\mathrm{d}\mu(\omega) = 0$ for all μ.

Non-linear persuasion

- One-dimensional state space $\Omega = [0, 1]$ and action space $A = [0, 1]$
- Twice continuously differentiable Sender's utility $v(a, \omega)$ and Receiver's utility $u(a, \omega)$
- Strictly concave Receiver's utility ($u_{aa}(a, \omega) < 0$) with interior optimal action $a^*(\mu)$ given by $\int_\Omega u_a(a^*(\mu), \omega)\mathrm{d}\mu(\omega) = 0$ for all μ.

Lipschitz property. Indirect utility $V(\mu) = \int_\Omega v(a^*(\mu), \omega)\mathrm{d}\mu(\omega)$ is Lipschitz in μ, so general duality applies.

General duality

The primal problem is to find $\tau \in \Delta(\Delta(\Omega))$ to

$$\begin{aligned} \text{maximize} & \int_{\Delta(\Omega)} V(\mu)\mathrm{d}\tau(\mu) \\ \text{subject to} & \int_{\Delta(\Omega)} \mu \mathrm{d}\tau(\mu) = \mu_0, \end{aligned} \tag{P}$$

General duality

The primal problem is to find $\tau \in \Delta(\Delta(\Omega))$ to

$$\begin{aligned} &\text{maximize} \int_{\Delta(\Omega)} V(\mu)\mathrm{d}\tau(\mu) \\ &\text{subject to} \int_{\Delta(\Omega)} \mu \mathrm{d}\tau(\mu) = \mu_0, \end{aligned} \tag{P}$$

The dual problem is to find $p \in L(\Omega)$ to

$$\begin{aligned} &\text{minimize} \int_{\Omega} p(\omega)\mathrm{d}\mu_0(\omega) \\ &\text{subject to } V(\mu) \leq \int_{\Omega} p(\omega)\mathrm{d}\mu(\omega) \text{ for all } \mu \in \Delta(\Omega). \end{aligned} \tag{D}$$

General duality

The primal problem is to find $\tau \in \Delta(\Delta(\Omega))$ to

$$\begin{aligned} &\text{maximize} \int_{\Delta(\Omega)} V(\mu)\mathrm{d}\tau(\mu) \\ &\text{subject to} \int_{\Delta(\Omega)} \mu\mathrm{d}\tau(\mu) = \mu_0, \end{aligned} \tag{P}$$

The dual problem is to find $p \in L(\Omega)$ to

$$\begin{aligned} &\text{minimize} \int_{\Omega} p(\omega)\mathrm{d}\mu_0(\omega) \\ &\text{subject to } V(\mu) \leq \int_{\Omega} p(\omega)\mathrm{d}\mu(\omega) \text{ for all } \mu \in \Delta(\Omega). \end{aligned} \tag{D}$$

General duality theorem. If V is Lipschitz in the Kantorovich Rubinstein metric, then *weak duality* and *strong duality* hold.

Action-based primal problem

Action-based primal is to find a *joint distribution* $\pi \in \Delta(A \times \Omega)$ to

$$\begin{aligned} \text{maximize } & \int_{A\times\Omega} v(a,\omega)\mathrm{d}\pi(a,\omega) \\ \text{subject to } & \int_{A\times\widetilde{\Omega}} \mathrm{d}\pi(a,\omega) = \int_{\widetilde{\Omega}} \mathrm{d}\mu_0(\omega) \text{ for all } \widetilde{\Omega} \subset \Omega, \\ & \int_{\widetilde{A}\times\Omega} u_a(a,\omega)\mathrm{d}\pi(a,\omega) = 0 \text{ for all } \widetilde{A} \subset A, \end{aligned} \tag{P'}$$

Action-based primal problem

Action-based primal is to find a *joint distribution* $\pi \in \Delta(A \times \Omega)$ to

$$\begin{aligned} \text{maximize} & \int_{A\times\Omega} v(a,\omega)\mathrm{d}\pi(a,\omega) \\ \text{subject to} & \int_{A\times\widetilde{\Omega}} \mathrm{d}\pi(a,\omega) = \int_{\widetilde{\Omega}} \mathrm{d}\mu_0(\omega) \text{ for all } \widetilde{\Omega} \subset \Omega, \\ & \int_{\widetilde{A}\times\Omega} u_a(a,\omega)\mathrm{d}\pi(a,\omega) = 0 \text{ for all } \widetilde{A} \subset A, \end{aligned} \tag{P'}$$

- The first constraint is the action-based version of Bayes plausibility, which says that the marginal of π on Ω equals μ_0.

Action-based primal problem

Action-based primal is to find a *joint distribution* $\pi \in \Delta(A \times \Omega)$ to

$$\begin{aligned} \text{maximize} \int_{A\times\Omega} v(a,\omega)\mathrm{d}\pi(a,\omega) \\ \text{subject to} \int_{A\times\widetilde{\Omega}} \mathrm{d}\pi(a,\omega) = \int_{\widetilde{\Omega}} \mathrm{d}\mu_0(\omega) \text{ for all } \widetilde{\Omega} \subset \Omega, \\ \int_{\widetilde{A}\times\Omega} u_a(a,\omega)\mathrm{d}\pi(a,\omega) = 0 \text{ for all } \widetilde{A} \subset A, \end{aligned} \tag{P'}$$

- The first constraint is the action-based version of Bayes plausibility, which says that the marginal of π on Ω equals μ_0.
- The second constraint is the obedience constraint, which says that the expected marginal utility equals zero at the recommended action given the belief it induces.

Action-based dual problem

Action-based dual is to find $p \in L(\Omega)$ and $q \in B(A)$ to

$$\begin{aligned} &\text{minimize} \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega) \\ &\text{subject to } p(\omega) \geq v(a,\omega) + q(a)u_a(a,\omega) \text{ for all } (a,\omega) \in A \times \Omega \end{aligned} \tag{D'}$$

Action-based dual problem

Action-based dual is to find $p \in L(\Omega)$ and $q \in B(A)$ to

$$\begin{aligned} &\text{minimize } \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega) \\ &\text{subject to } p(\omega) \geq v(a,\omega) + q(a)u_a(a,\omega) \text{ for all } (a,\omega) \in A \times \Omega \end{aligned} \tag{D'}$$

Interpretation: Price of state ω is no less than Sender's value from inducing any action a at this state, where this value is the sum of

- Sender's utility, $v(a,\omega)$, and
- Value of relaxing obedience at a, $q(a)$, times amount by which obedience at a is relaxed when a is induced at ω, $u_a(a,\omega)$.

Action-based dual problem

Action-based dual is to find $p \in L(\Omega)$ and $q \in B(A)$ to

$$\begin{aligned} &\text{minimize } \int_\Omega p(\omega)\mathrm{d}\mu_0(\omega) \\ &\text{subject to } p(\omega) \geq v(a,\omega) + q(a)u_a(a,\omega) \text{ for all } (a,\omega) \in A \times \Omega \end{aligned} \tag{D'}$$

Interpretation: Price of state ω is no less than Sender's value from inducing any action a at this state, where this value is the sum of

- Sender's utility, $v(a,\omega)$, and
- Value of relaxing obedience at a, $q(a)$, times amount by which obedience at a is relaxed when a is induced at ω, $u_a(a,\omega)$.

First-order approach is key for dual constraint
taking such a simple form.

Action-based duality

Bridging belief-based and action-based duality. A price function $p \in L(\Omega)$ is feasible (optimal) for (D) iff there exists $q \in B(A)$ such that (p, q) is feasible (optimal) for (D').

Action-based duality

Bridging belief-based and action-based duality. A price function $p \in L(\Omega)$ is feasible (optimal) for (D) iff there exists $q \in B(A)$ such that (p, q) is feasible (optimal) for (D').

Proof outline: If (p, q) is feasible for (D'), then, for all $\mu \in \Delta(\Omega)$,

$$\begin{aligned}\int_\Omega p(\omega)\mathrm{d}\mu(\omega) &\geq \int_\Omega (v(a^*(\mu), \omega) + q(a^*(\mu))u_a(a^*(\mu), \omega))\mathrm{d}\mu(\omega) \\ &= \int_\Omega v(a^*(\mu), \omega)\mathrm{d}\mu(\omega) = V(\mu),\end{aligned}$$

so p is feasible for (D).

Action-based duality

Proof outline continued: If p is feasible for (D), then, for all $\omega_1, \omega_2 \in \Omega$ and $a \in A$ such that $u_a(a, \omega_1) < 0 < u_a(a, \omega_2)$, feasibility for

$$\mu = \frac{u_a(a, \omega_2)}{u_a(a, \omega_2) - u_a(a, \omega_1)} \delta_{\omega_1} + \frac{-u_a(a, \omega_1)}{u_a(a, \omega_2) - u_a(a, \omega_1)} \delta_{\omega_2}$$

Action-based duality

Proof outline continued: If p is feasible for (D), then, for all $\omega_1, \omega_2 \in \Omega$ and $a \in A$ such that $u_a(a, \omega_1) < 0 < u_a(a, \omega_2)$, feasibility for

$$\mu = \frac{u_a(a, \omega_2)}{u_a(a, \omega_2) - u_a(a, \omega_1)} \delta_{\omega_1} + \frac{-u_a(a, \omega_1)}{u_a(a, \omega_2) - u_a(a, \omega_1)} \delta_{\omega_2}$$

$$\implies \frac{p(\omega_2) - v(a, \omega_2)}{u_a(a, \omega_2)} \geq \frac{v(a, \omega_1) - p(\omega_1)}{-u_a(a, \omega_1)}$$

Action-based duality

Proof outline continued: If p is feasible for (D), then, for all $\omega_1, \omega_2 \in \Omega$ and $a \in A$ such that $u_a(a, \omega_1) < 0 < u_a(a, \omega_2)$, feasibility for

$$\mu = \frac{u_a(a, \omega_2)}{u_a(a, \omega_2) - u_a(a, \omega_1)} \delta_{\omega_1} + \frac{-u_a(a, \omega_1)}{u_a(a, \omega_2) - u_a(a, \omega_1)} \delta_{\omega_2}$$

$$\implies \frac{p(\omega_2) - v(a, \omega_2)}{u_a(a, \omega_2)} \geq \frac{v(a, \omega_1) - p(\omega_1)}{-u_a(a, \omega_1)}$$

$$\implies \inf_{\omega_2 : u_a(a, \omega_2) > 0} \frac{p(\omega_2) - v(a, \omega_2)}{u_a(a, \omega_2)} \geq \sup_{\omega_1 : u_a(a, \omega_1) < 0} \frac{v(a, \omega_1) - p(\omega_1)}{-u_a(a, \omega_1)}$$

Action-based duality

Proof outline continued: If p is feasible for (D), then, for all $\omega_1, \omega_2 \in \Omega$ and $a \in A$ such that $u_a(a, \omega_1) < 0 < u_a(a, \omega_2)$, feasibility for

$$\mu = \frac{u_a(a,\omega_2)}{u_a(a,\omega_2) - u_a(a,\omega_1)}\delta_{\omega_1} + \frac{-u_a(a,\omega_1)}{u_a(a,\omega_2) - u_a(a,\omega_1)}\delta_{\omega_2}$$

$$\implies \frac{p(\omega_2) - v(a,\omega_2)}{u_a(a,\omega_2)} \geq \frac{v(a,\omega_1) - p(\omega_1)}{-u_a(a,\omega_1)}$$

$$\implies \inf_{\omega_2 : u_a(a,\omega_2) > 0} \frac{p(\omega_2) - v(a,\omega_2)}{u_a(a,\omega_2)} \geq \sup_{\omega_1 : u_a(a,\omega_1) < 0} \frac{v(a,\omega_1) - p(\omega_1)}{-u_a(a,\omega_1)}$$

Thus, we can squeeze in $q(a)$ between the LHS and RHS, so that $p(\omega) \geq v(a,\omega) + q(a)u_a(a,\omega)$, yielding feasibility for (D').

First-order condition for the action-based dual

Given $q \in B(A)$, it is optimal for (D') to set

$$p(\omega) = \sup_{a \in A} \{v(a, \omega) + q(a) u_a(a, \omega)\}, \text{ for all } \omega \in \Omega.$$

First-order condition for the action-based dual

Given $q \in B(A)$, it is optimal for (D') to set

$$p(\omega) = \sup_{a \in A}\{v(a,\omega) + q(a)u_a(a,\omega)\}, \text{ for all } \omega \in \Omega.$$

Strong duality implies that for all μ in the support of an optimal τ and for all ω in the support of μ, the first-order condition holds:

$$v_a(a^*(\mu),\omega) + q(a^*(\mu))u_{aa}(a^*(\mu),\omega) + q'(a^*(\mu))u_a(a^*(\mu),\omega) = 0,$$

First-order condition for the action-based dual

Given $q \in B(A)$, it is optimal for (D') to set

$$p(\omega) = \sup_{a \in A}\{v(a,\omega) + q(a)u_a(a,\omega)\}, \text{ for all } \omega \in \Omega.$$

Strong duality implies that for all μ in the support of an optimal τ and for all ω in the support of μ, the first-order condition holds:

$$v_a(a^*(\mu),\omega) + q(a^*(\mu))u_{aa}(a^*(\mu),\omega) + q'(a^*(\mu))u_a(a^*(\mu),\omega) = 0,$$

and thus,

$$q(a^*(\mu)) = \frac{\mathbb{E}_\mu[v_a(a^*(\mu),\omega)]}{-\mathbb{E}_\mu[u_{aa}(a^*(\mu),\omega)]}.$$

Pairwise Signals

- A signal τ is pairwise if each induced posterior $\mu \in \text{supp}(\tau)$ has at most binary support ($|\text{supp}(\mu)| \leq 2$):

$$\mu = (1 - \lambda)\delta_\omega + \lambda\delta_{\omega'}, \quad \text{for } \omega \leq \omega' \text{ and } \lambda \in [0, 1].$$

- For example, for uniform μ_0 and any ω^*, the signal that reveals ω if $\omega < \omega^*$, and reveals that the state is either ω or $1 + \omega^* - \omega$ with equal probability if $\omega \in [\omega^*, \frac{1+\omega^*}{2}]$, is pairwise.

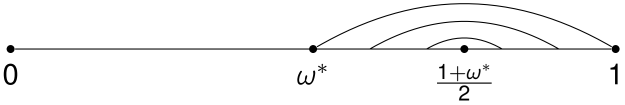

- Full disclosure ($\text{supp}(\tau_F) = \{\delta_\omega : \omega \in \Omega\}$) is pairwise.

- No disclosure ($\text{supp}(\tau_N) = \{\mu_0\}$) is not pairwise.

Optimality of pairwise signals

Every optimal signal is pairwise if for all a and $\omega_1 < \omega_2 < \omega_3$, we have

$$\det \begin{pmatrix} v_a(a,\omega_1) & v_a(a,\omega_2) & v_a(a,\omega_3) \\ u_a(a,\omega_1) & u_a(a,\omega_2) & u_a(a,\omega_3) \\ u_{aa}(a,\theta_1) & u_{aa}(a,\omega_2) & u_{aa}(a,\omega_3) \end{pmatrix} \neq 0. \qquad \text{(N)}$$

- Follows from the strong-duality FOC, as there do not exist an action a, three states $\omega_1 < \omega_2 < \omega_3$, and a vector $(q(a), q'(a))$:

$$v_a(a,\omega_i) + q(a)u_{aa}(a,\omega_i) + q'(a)u_a(a,\omega_i), \text{ for } i = 1,2,3.$$

- (N) holds, for example, if $u(a,\omega) = -(a-\omega)^2$ and $v(a,\omega) = aw(\omega)$ with strictly convex or concave w.
- (N) fails in the linear case where U_a and V_a are linear in θ (so pooling intervals of states can be optimal in the linear case).

Intuition

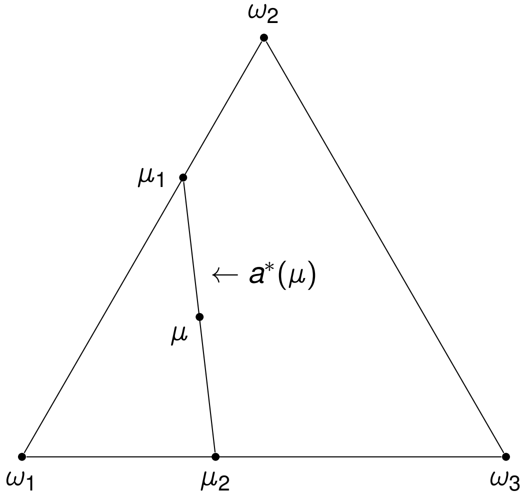

Linear persuasion

We now assume that $V(\mu) = v\left(\mathbb{E}_\mu[\omega]\right)$ for all $\mu \in \Delta([0,1])$.

This is a special case of non-linear persuasion with $u(a,\omega) = -(a-\omega)^2$ and state-independent v.

Only the distribution η of posterior means matters, where η is the A-marginal distribution of π: $\eta(\widetilde{A}) = \pi(\widetilde{A},\Omega)$ for all $\tilde{A} \subset A$.

Primal and Dual

The primal problem simplifies to finding $\eta \in \Delta([0, 1])$ to

$$\begin{aligned} &\text{maximize} \int v(a)\mathrm{d}\eta(a) \\ &\text{subject to } \eta \in MPC(\mu_0). \end{aligned} \tag{P''}$$

The dual problem simplifies to finding $p \in L(\Omega)$ to

$$\begin{aligned} &\text{minimize} \int p(\omega)\mathrm{d}\mu_0(\omega) \\ &\text{subject to } p \text{ is convex and } p \geq v. \end{aligned} \tag{D''}$$