Information without Rents: Mechanism Design without Expected Utility

Ernesto Rivera Mora [1] Philipp Strack [2]

[1]University of Colorado, Boulder

[2]Yale University

Nov 2026



Joint work with **Philipp Strack**.

What role do EU preferences play in mechanism design?

- Expected Utility preferences are typically assumed in mech design

What role do EU preferences play in mechanism design?

- Expected Utility preferences are typically assumed in mech design
  - Revelation Principle: use static direct mechanisms without loss

What role do EU preferences play in mechanism design?

- Expected Utility preferences are typically assumed in mech design
  - Revelation Principle: use static direct mechanisms without loss
- Non-EU preferences play an important role in behavioral economics and applications

What role do EU preferences play in mechanism design?

- Expected Utility preferences are typically assumed in mech design
  - Revelation Principle: use static direct mechanisms without loss
- Non-EU preferences play an important role in behavioral economics and applications
  - Non-EU preferences - time-inconsistency

What role do EU preferences play in mechanism design?

- Expected Utility preferences are typically assumed in mech design
  - Revelation Principle: use static direct mechanisms without loss
- Non-EU preferences play an important role in behavioral economics and applications
  - Non-EU preferences - time-inconsistency
  - Dynamic mechanisms could be used to exploit time-inconsistency

Main Questions

Q1: When does the Revelation Principle hold?

Main Questions

Q1: When does the Revelation Principle hold?

Q2: What is the set of implementable allocations?

Main Questions

Q1: When does the Revelation Principle hold?

Q2: What is the set of implementable allocations?

Q3: What are the information rents for non-EU?

Main Questions

Q1: When does the Revelation Principle hold?

Result: Revelation Principle holds if and only if all types have EU preferences

Q2: What is the set of implementable allocations?

Q3: What are the information rents for non-EU?

Main Questions

Q1: When does the Revelation Principle hold?

Result: Revelation Principle holds if and only if all types have EU preferences

Q2: What is the set of implementable allocations?

Result: Fully characterize implementable allocations for mixture-averse preferences

Q3: What are the information rents for non-EU?

Main Questions

Q1: When does the Revelation Principle hold?

Result: Revelation Principle holds if and only if all types have EU preferences

Q2: What is the set of implementable allocations?

Result: Fully characterize implementable allocations for mixture-averse preferences

Q3: What are the information rents for non-EU?

Result: Dynamic mechanisms fully extract private information from agents with non-EU mixture-averse preferences for free

Literature — Mech Design for non-EU preferences

Static Mechanisms Gershkov et al. (2022); Eisenhuth (2019); Gershkov et al. (2023); Dütting et al. (2024); Dillenberger and Segal (2024); Gershkov et al. (2025)

Different Auction Formats Lange and Ratan (2010); Rosato and Tymula (2019); von Wangenheim (2021); Balzer et al. (2022); Balzer and Rosato (2025)

Time-inconsistency Green (1987); Karni and Schmeidler (1991); Amador et al. (2006); Heidhues and Kőszegi (2008); Galperti (2014); Gottlieb and Zhang (2021); Bose and Daripa (2009)

Closest work: Ahn, Zhang, and Zhang (2026): unidimensional continuum types with biseparable preferences (Including CRA with multiplicative types); link recursive utility

Model

Setup:

- Finite set of outcomes Ω.

Setup:

- Finite set of outcomes Ω.
  - Leading Example $\Omega := \underbrace{X}_{\text{allocations}} \times \underbrace{T}_{\text{transfers}}$

Setup:

- Finite set of outcomes Ω.
  - Leading Example $\Omega := \underbrace{X}_{\text{allocations}} \times \underbrace{T}_{\text{transfers}}$
- This talk: One agent

Setup:

- Finite set of outcomes Ω.
  - Leading Example $\Omega := \underbrace{X}_{\text{allocations}} \times \underbrace{T}_{\text{transfers}}$
- This talk: One agent
- Finite set of types Θ

Setup:

- Finite set of outcomes Ω.
  - Leading Example $\Omega := \underbrace{X}_{\text{allocations}} \times \underbrace{T}_{\text{transfers}}$
- This talk: One agent
- Finite set of types Θ
- Each type $\theta \in \Theta$ has preference $\succsim_\theta$ over lotteries $L := \Delta(\Omega)$

Assumptions:

- Preferences are complete, continuous, and transitive

Assumptions:

- Preferences are complete, continuous, and transitive

Assumptions:

- Preferences are complete, continuous, and transitive
- Reduction of compound lotteries

Assumptions:

- Preferences are complete, continuous, and transitive
- Reduction of compound lotteries

Assumptions:

- Preferences are complete, continuous, and transitive
- Reduction of compound lotteries
- Sophistication: agents correctly anticipate future choices

Assumptions:

- Preferences are complete, continuous, and transitive
- Reduction of compound lotteries
- Sophistication: agents correctly anticipate future choices

Assumptions:

- Preferences are complete, continuous, and transitive
- Reduction of compound lotteries
- Sophistication: agents correctly anticipate future choices
- Some outcome $\underline{\omega} \in \Omega$ is the worst for all types:

Assumptions:

- Preferences are complete, continuous, and transitive
- Reduction of compound lotteries
- Sophistication: agents correctly anticipate future choices
- Some outcome $\underline{\omega} \in \Omega$ is the worst for all types:
  - Degenerate worst lottery $\delta_{\underline{\omega}}$

Assumptions:

- Preferences are complete, continuous, and transitive
- Reduction of compound lotteries
- Sophistication: agents correctly anticipate future choices
- Some outcome $\underline{\omega} \in \Omega$ is the worst for all types:
  - Degenerate worst lottery $\delta_{\underline{\omega}}$
  - If p has full support, then $p \succ_\theta \lambda p + (1-\lambda)\delta_{\underline{\omega}} \succ_\theta \delta_{\underline{\omega}}, \quad \lambda \in (0,1)$

Definition (Mechanism)

A (dynamic) mechanism is a tuple (T, A, m) consisting of

- the number of periods T,
- the possible actions A, and
- a lottery mapping $m : A^T \times [0,1]^T \to L$.

Each period t the mechanism

1. reveals random number $r_t \sim U([0,1])$,
2. agent observes r_t and chooses action $a_t \in A$.

Solution Concept

Subgame Perfect Equilibrium (SPE) as solution concept

▸ Formal definition

- Requires sequential rationality

Solution Concept

Subgame Perfect Equilibrium (SPE) as solution concept

▸ Formal definition

- Requires sequential rationality
  - Optimality after every history

Solution Concept

Subgame Perfect Equilibrium (SPE) as solution concept

Formal definition

- Requires sequential rationality
  - Optimality after every history
  - Assumes the agent follows the strategy in subsequent periods

Solution Concept

Subgame Perfect Equilibrium (SPE) as solution concept

▸ Formal definition

- Requires sequential rationality
  - Optimality after every history
  - Assumes the agent follows the strategy in subsequent periods
  - Dynamic inconsistencies may arise

Solution Concept

Subgame Perfect Equilibrium (SPE) as solution concept

▸ Formal definition

- Requires sequential rationality
  - Optimality after every history
  - Assumes the agent follows the strategy in subsequent periods
  - Dynamic inconsistencies may arise
- Agent cannot commit to future suboptimal behavior

Solution Concept

Subgame Perfect Equilibrium (SPE) as solution concept

▸ Formal definition

- Requires sequential rationality
  - Optimality after every history
  - Assumes the agent follows the strategy in subsequent periods
  - Dynamic inconsistencies may arise
- Agent cannot commit to future suboptimal behavior
- Assumes the agent is sophisticated

Solution Concept

Subgame Perfect Equilibrium (SPE) as solution concept

Formal definition

- Requires sequential rationality
  - Optimality after every history
  - Assumes the agent follows the strategy in subsequent periods
  - Dynamic inconsistencies may arise
- Agent cannot commit to future suboptimal behavior
- Assumes the agent is sophisticated
  - Each type θ correctly predicts her behavior in future stages

Solution Concept

Subgame Perfect Equilibrium (SPE) as solution concept

▸ Formal definition

- Requires sequential rationality
  - Optimality after every history
  - Assumes the agent follows the strategy in subsequent periods
  - Dynamic inconsistencies may arise
- Agent cannot commit to future suboptimal behavior
- Assumes the agent is sophisticated
  - Each type θ correctly predicts her behavior in future stages
  - Naive types: easy to create money pumps (see e.g. Ebert & Strack 2013)

Example: Monopolistic Screening

- Single buyer; type θ = value of the object

Example: Monopolistic Screening

- Single buyer; type θ = value of the object
- Mechanism sets price $\omega \in \mathbb{R}_+$; buyer buys iff $\theta \geq \omega$

Example: Monopolistic Screening

- Single buyer; type θ = value of the object
- Mechanism sets price $\omega \in \mathbb{R}_+$; buyer buys iff $\theta \geq \omega$
- Buyer payoff: $\max\{\theta - \omega, 0\}$;

Example: Monopolistic Screening

- Single buyer; type θ = value of the object
- Mechanism sets price $\omega \in \mathbb{R}_+$; buyer buys iff $\theta \geq \omega$
- Buyer payoff: $\max\{\theta - \omega, 0\}$;
- Seller payoff: $\omega \cdot 1[\theta \geq \omega]$

Example: Monopolistic Screening

- Single buyer; type θ = value of the object
- Mechanism sets price $\omega \in \mathbb{R}_+$; buyer buys iff $\theta \geq \omega$
- Buyer payoff: $\max\{\theta - \omega, 0\}$;
- Seller payoff: $\omega \cdot 1[\theta \geq \omega]$

Claims

- No risk-loving EU $\Rightarrow$ deterministic posted price; incomplete surplus extraction

Example: Monopolistic Screening

- Single buyer; type θ = value of the object
- Mechanism sets price $\omega \in \mathbb{R}_+$; buyer buys iff $\theta \geq \omega$
- Buyer payoff: $\max\{\theta - \omega, 0\}$;
- Seller payoff: $\omega \cdot 1[\theta \geq \omega]$

Claims

- No risk-loving EU $\Rightarrow$ deterministic posted price; incomplete surplus extraction
- Non-EU Mixture-Averse Preferences $\Rightarrow$ Extract full surplus; zero information rents

Example

- Available Prices: $\Omega = \{100, 50, 0\}$ (at 0 the object is free)

Example

- Available Prices: $\Omega = \{100, 50, 0\}$ (at 0 the object is free)
- Yaari Preferences with cumulative distortion $g(x) = x^2$
  - Two types (valuations) $\theta_1 = 50$ and $\theta_2 = 100$

Example

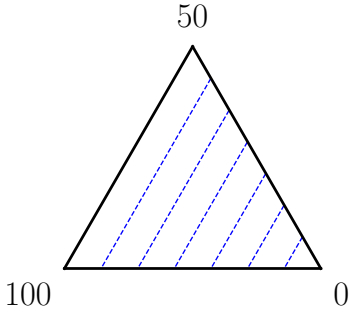
- Available Prices: $\Omega = \{100, 50, 0\}$ (at 0 the object is free)
- Yaari Preferences with cumulative distortion $g(x) = x^2$
  - Two types (valuations) $\theta_1 = 50$ and $\theta_2 = 100$
  - Utility for lottery $p \in \Delta(\{100, 50, 0\})$
    - Type $\theta_1 = 50$: $v_1(p) = 50\, p(0)^2$
    - Type $\theta_2 = 100$: $v_2(p) = 50\big(p(50) + p(0)\big)^2 + 50\big(p(0)\big)^2$

Claim: Dynamic mechanisms can extract all surplus

Simplex of Price Lotteries: Indifference Curves

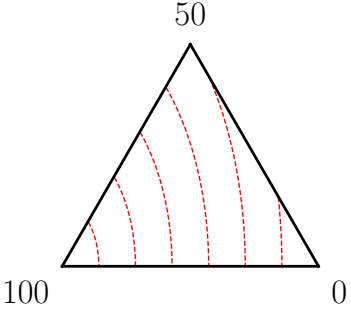
Low Type $\theta_1 = 50$

- Indifferent over $\omega \in \{100, 50\}$
- EU-equivalent

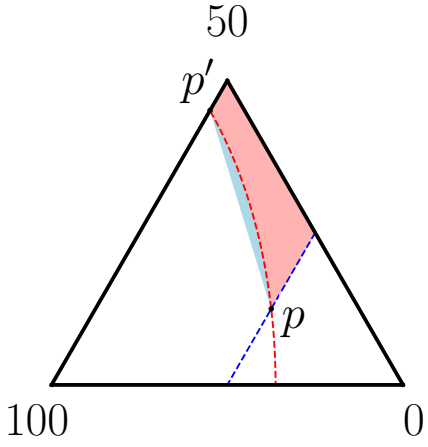


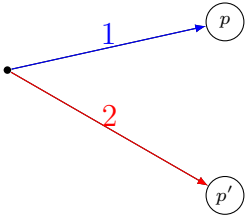
High Type $\theta_2 = 100$

- Strictly mixture-averse (non-EU)
- Curved indifference curves

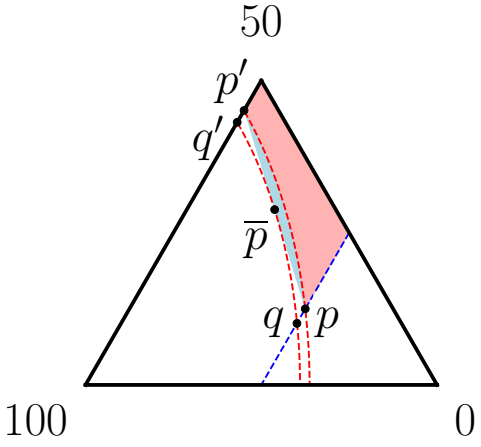


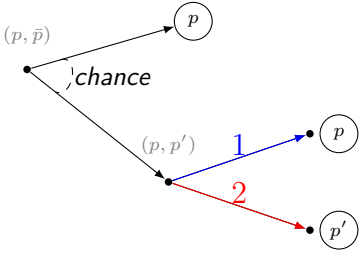
Implementing (p, p') to (θ_1, θ_2)



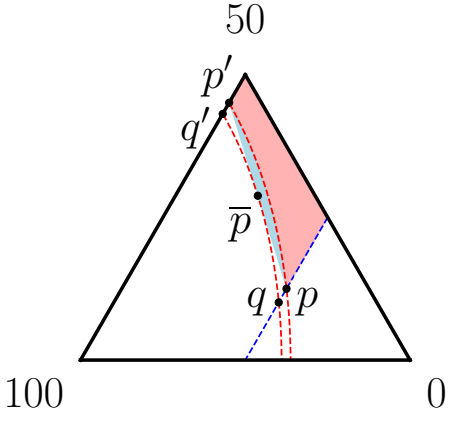


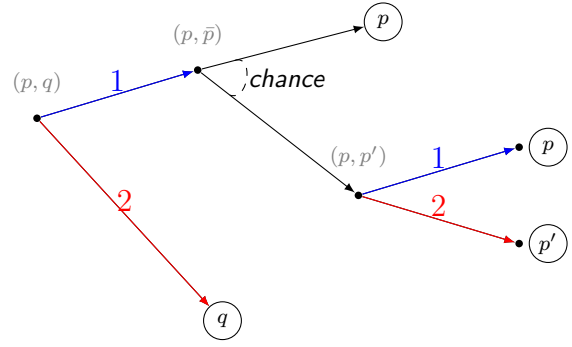
Implementing $(p, \bar{p})$ to (θ_1, θ_2)



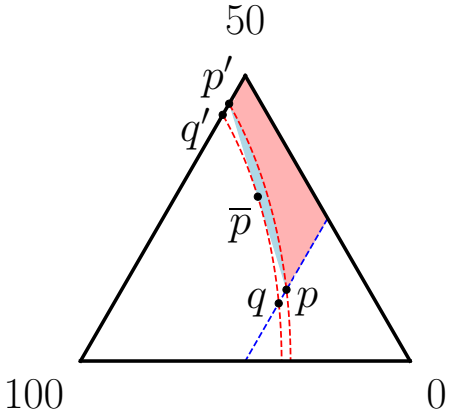


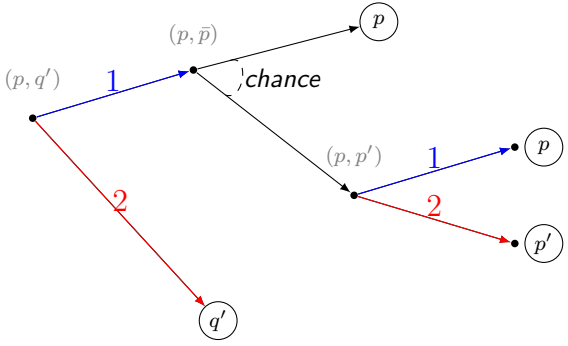
Implementing (p, q)



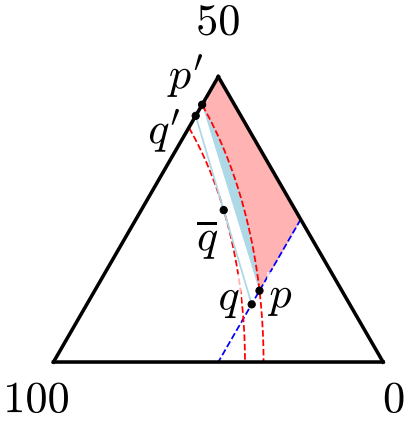


Implementing (p, q')





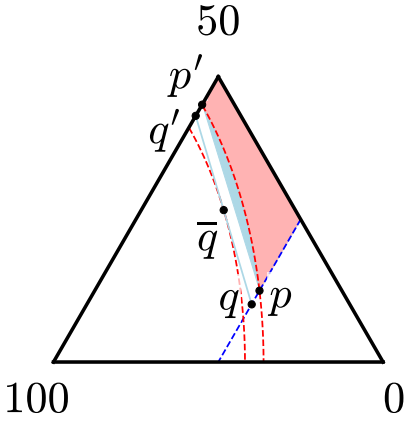
Implementing $(p, \overline{q})$



Randomize between

1. Mechanism implementing (p, q)
2. Mechanism implementing (p, q')

Implementing $(p, \overline{q})$



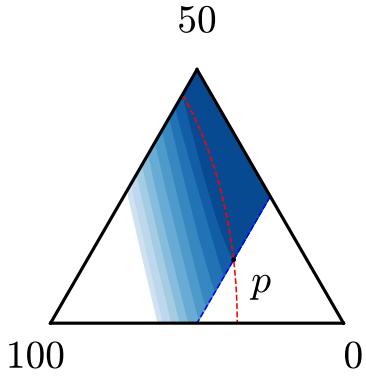
Randomize between

1. Mechanism implementing (p, q)
2. Mechanism implementing (p, q')

Can we achieve more allocations with more steps?

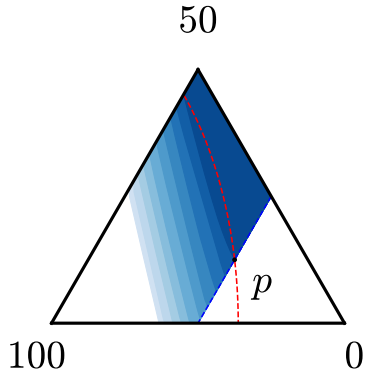
Allocations achievable with 10 steps

Lotteries implementable for high type $\theta_2 = 100$ with 10-period mechanisms conditional on low $\theta_1 = 50$ getting p:



Allocations achievable with 10 steps

Lotteries implementable for high type $\theta_2 = 100$ with 10-period mechanisms conditional on low $\theta_1 = 50$ getting p:



Similar for each interior lottery p $\Rightarrow$ Full surplus extraction in the limit

Beyond the example: Implementation

Implementable allocations. An *allocation* $(p_\theta)_{\theta\in\Theta} \in L^{|\Theta|}$ is implemented by (T, A, m) if it has an equilibrium (σ_θ) such that each type θ obtains p_θ after playing σ_θ.

Beyond the example: Implementation

Implementable allocations. An *allocation* $(p_\theta)_{\theta\in\Theta} \in L^{|\Theta|}$ is implemented by (T, A, m) if it has an equilibrium (σ_θ) such that each type θ obtains p_θ after playing σ_θ.

Set of Implementable allocations. Write $\mathrm{IMP}^T \subseteq L^{|\Theta|}$ for the set of allocations implementable by some T-round mechanism

$$\mathrm{IMP} := \text{closure}\left(\bigcup_{T\in\mathbb{N}} \mathrm{IMP}^T\right).$$

Direct Mechanisms

Definition (Direct Mechanism)

A mechanism is ***direct*** if $T = 1$, $A = \Theta$, and $m(\theta, \cdot) \in \Delta(\Omega)$ is constant for each type $\theta \in \Theta$ (menu of fixed lotteries)

Definition (Direct Mechanism)

A mechanism is ***direct*** if $T = 1$, $A = \Theta$, and $m(\theta, \cdot) \in \Delta(\Omega)$ is constant for each type $\theta \in \Theta$ (menu of fixed lotteries)

Definition (Incentive Compatibility)

For each $\theta \in \Theta$, write

$$\mathrm{IC}_\theta := \left\{ (p_\theta)_{\theta \in \Theta} \in L^{|\Theta|} \ : \ p_\theta \succeq_\theta p_{\theta'} \text{ for all } \theta, \theta' \in \Theta \right\}.$$

for the allocations incentive compatible for θ (under direct mechanisms).

The Revelation Principle

An allocation $(p_\theta)_{\theta\in\Theta}$ is implementable by an IC direct mechanism iff

$$(p_\theta)_{\theta\in\Theta} \in \bigcap_{\theta\in\Theta} \mathrm{IC}_\theta.$$

An allocation $(p_\theta)_{\theta\in\Theta}$ is implementable by an IC direct mechanism iff

$$(p_\theta)_{\theta\in\Theta} \in \bigcap_{\theta\in\Theta} \mathrm{IC}_\theta.$$

Proposition (Static Revelation Principle)

If each type $\theta \in \Theta$ has expected utility preferences, then

$$\mathrm{IMP} = \bigcap_{\theta\in\Theta} \mathrm{IC}_\theta.$$

No loss in restricting to static direct mechanisms and truthful reports

Main Equivalence

Theorem 1

Fix an environment $\Omega, (\preceq_\theta)_\theta$. The following statements are equivalent:

(i) Each type $\theta \in \Theta$ has EU preferences.

(ii) $\bigcap_{\theta \in \Theta} \mathrm{IC}_\theta \;=\; \mathrm{IMP}$.

Main Equivalence

Theorem 1

Fix an environment $\Omega, (\preceq_\theta)_\theta$. The following statements are equivalent:

(i) Each type $\theta \in \Theta$ has EU preferences.

(ii) $\bigcap_{\theta \in \Theta} \mathrm{IC}_\theta \;=\; \mathrm{IMP}$.

- EU is both *necessary* and *sufficient* for the Revelation Principle.

Theorem 1

Fix an environment $\Omega, (\preceq_\theta)_\theta$. The following statements are equivalent:

(i) Each type $\theta \in \Theta$ has EU preferences.

(ii) $\bigcap_{\theta\in\Theta} \mathrm{IC}_\theta \;=\; \mathrm{IMP}$.

Proof that (ii) implies (i):

- Continuity + outcome $\underline{\omega}$
- $\bigcap_{\theta\in\Theta} \mathrm{IC}_\theta = \mathrm{IMP}$ is convex

$\Rightarrow$

Each type θ satisfies the Independence axiom

Discussion

Theorem 1 shows dynamic mechanisms implement more allocations when some types are non-EU.

Theorem 1 shows dynamic mechanisms implement more allocations when some types are non-EU.

- Which allocations are implementable?
- Which are implementable with T-period mechanisms?
- Is using 2-period mechanisms enough?
- Can we provide a substitute for the Revelation Principle?

Recursive Approach: Binary Types



Consider 2 types: For each set of allocation profiles $S \subseteq L^2$ define $\Phi(S) \subseteq L^2$ as

$$\text{conv}\Big\{(p_1, p_2) \colon \exists\ q_1, q_2 \in L \text{ s.t. } (p_1, q_2), (q_1, p_2) \in S \text{ s.t. } p_1 \succeq_1 q_1 \text{ and } p_2 \succeq_2 q_2\Big\}$$

Recursive Approach: Binary Types



Consider 2 types: For each set of allocation profiles $S \subseteq L^2$ define $\Phi(S) \subseteq L^2$ as

$$\text{conv}\Big\{ \underbrace{(p_1, p_2) \colon \exists\ q_1, q_2 \in L \text{ s.t. } (p_1, q_2), (q_1, p_2) \in S \text{ s.t. } p_1 \succeq_1 q_1 \text{ and } p_2 \succeq_2 q_2}_{\text{Achievable by allowing agent to self-select}} \Big\}$$

Recursive Approach: Binary Types



Consider 2 types: For each set of allocation profiles $S \subseteq L^2$ define $\Phi(S) \subseteq L^2$ as

$$\underbrace{\text{conv}}_{\text{Randomization}} \left\{ (p_1, p_2) \colon \exists\ q_1, q_2 \in L \text{ s.t. } (p_1, q_2), (q_1, p_2) \in S \text{ s.t. } p_1 \succeq_1 q_1 \text{ and } p_2 \succeq_2 q_2 \right\}$$

Recursive Approach: Binary Types



Consider 2 types: For each set of allocation profiles $S \subseteq L^2$ define $\Phi(S) \subseteq L^2$ as

$$\text{conv}\Big\{(p_1, p_2) \colon \exists\ q_1, q_2 \in L \text{ s.t. } (p_1, q_2), (q_1, p_2) \in S \text{ s.t. } p_1 \succeq_1 q_1 \text{ and } p_2 \succeq_2 q_2\Big\}$$

Write $\text{IMP}^0 := \{(p, p) \in L^2\}$

Lemma

The set of allocations implementable in T periods satisfies $\text{IMP}^T = \Phi(\text{IMP}^{T-1})$

Mixture-averse preferences

Type θ has mixture-averse preferences if for each $p \sim_\theta q$, $p \neq q$, and each $\alpha \in (0,1)$:

$$\alpha p + (1-\alpha)q \;\prec_\theta\; p \;\sim_\theta\; q.$$

Implementable Allocations

Mixture-averse preferences

Type θ has mixture-averse preferences if for each $p \sim_\theta q$, $p \neq q$, and each $\alpha \in (0,1)$:

$$\alpha p + (1-\alpha)q \ \prec_\theta \ p \ \sim_\theta \ q.$$

Theorem 2

Suppose each type θ is EU or mixture-averse. If $\mathrm{EU} \subseteq \Theta$ is the set of EU types, then

IC constraints of mixture-averse types can be ignored: $\quad \mathrm{IMP} = \bigcap_{\theta \in \mathrm{EU}} \mathrm{IC}_\theta.$

Implementable Allocations

Mixture-averse preferences

Type θ has mixture-averse preferences if for each $p \sim_\theta q$, $p \neq q$, and each $\alpha \in (0,1)$:

$$\alpha p + (1-\alpha)q \ \prec_\theta \ p \ \sim_\theta \ q.$$

Theorem 2

Suppose each type θ is EU or mixture-averse. If $\mathrm{EU} \subseteq \Theta$ is the set of EU types, then

IC constraints of mixture-averse types can be ignored: $\quad \mathrm{IMP} = \bigcap_{\theta \in \mathrm{EU}} \mathrm{IC}_\theta.$

*If EU is empty, then $\mathrm{IMP} = \cap_{\theta \in \mathrm{EU}} \mathrm{IC}_\theta := L^{|\Theta|}$.

Implementable Allocations

Mixture-averse preferences

Type θ has mixture-averse preferences if for each $p \sim_\theta q$, $p \neq q$, and each $\alpha \in (0,1)$:

$$\alpha p + (1-\alpha)q \ \prec_\theta \ p \ \sim_\theta \ q.$$

Theorem 2

Suppose each type θ is EU or mixture-averse. If $\mathrm{EU} \subseteq \Theta$ is the set of EU types, then

IC constraints of mixture-averse types can be ignored: $\quad \mathrm{IMP} = \bigcap_{\theta \in \mathrm{EU}} \mathrm{IC}_\theta.$

Intuition: Use time inconsistency to construct targeted punishments

Broad Applications

Theorem 2 imposes no assumptions on dimension of outcomes Ω, dimension of type space Θ, transfers, (linear) separability, supermodularity, etc.

Broad Applications

Theorem 2 imposes no assumptions on dimension of outcomes Ω, dimension of type space Θ, transfers, (linear) separability, supermodularity, etc.

Private information $\theta \in \Theta$ may include:

1. Preferences over outcomes (values of an object)
2. Degree of risk aversion (concavity in utility function)
3. Degree of dynamic inconsistency (quadratic vs. cubic distortion, etc)

Broad Applications

Theorem 2 imposes no assumptions on dimension of outcomes Ω, dimension of type space Θ, transfers, (linear) separability, supermodularity, etc.

Example: Kőszegi & Rabin, 2006 (Reference-dependent + Personal Equilibrium)

Expectations as reference points: Each type $\theta = (\pi, u, \gamma)$ with payoff $\pi : \Omega \to \mathbb{R}$, Bernoulli utility $u : \mathbb{R} \to \mathbb{R}$, loss-aversion parameter $\gamma > 1$,

$$v_\theta(p) = \sum_\omega p(\omega) \left[u(\pi(\omega)) + \sum_{\tilde{\omega}} p(\tilde{\omega})\, m_\gamma\big(u(\pi(\omega)) - u(\pi(\tilde{\omega}))\big) \right], \quad m_\gamma(x) := \min\{x, \gamma x\}$$

is mixture averse.

Details

Broad Applications

Theorem 2 imposes no assumptions on dimension of outcomes Ω, dimension of type space Θ, transfers, (linear) separability, supermodularity, etc.

Example: Kőszegi & Rabin, 2006 (Reference-dependent + Personal Equilibrium)

Expectations as reference points: Each type $\theta = (\pi, u, \gamma)$ with payoff $\pi : \Omega \to \mathbb{R}$, Bernoulli utility $u : \mathbb{R} \to \mathbb{R}$, loss-aversion parameter $\gamma > 1$,

$$v_\theta(p) = \sum_\omega p(\omega) \left[u(\pi(\omega)) + \sum_{\tilde{\omega}} p(\tilde{\omega})\, m_\gamma\big(u(\pi(\omega)) - u(\pi(\tilde{\omega}))\big) \right], \quad m_\gamma(x) := \min\{x, \gamma x\}$$

is mixture averse.

Details

Rank-dependent utility (Quiggin, 1982) under convex distortions

Application: Quasi-linear Transfers & Full-Surplus Extraction

- Outcome: $\Omega \subset X \times T$, with transfer $T \subset \mathbb{R}$.
- Utility of outcome $(x, t) \in \Omega$ is $v(\theta, x) - t$.

Application: Quasi-linear Transfers & Full-Surplus Extraction

- Outcome: $\Omega \subset X \times T$, with transfer $T \subset \mathbb{R}$.
- Utility of outcome $(x, t) \in \Omega$ is $v(\theta, x) - t$.

Corollary (Full-Surplus Extraction) ▸ Discussion

If all types have strictly mixture-averse preferences over $\Delta(\Omega)$, then the designer can fully extract value $v(\theta, x_\theta)$ received by each type θ.

Contemporaneous with Ahn, Zhang & Zhang (2026):

Look beyond the reduction compound axiom (e.g. recursive utility)

Proof of Theorem 2

Targeted Punishments (two types)

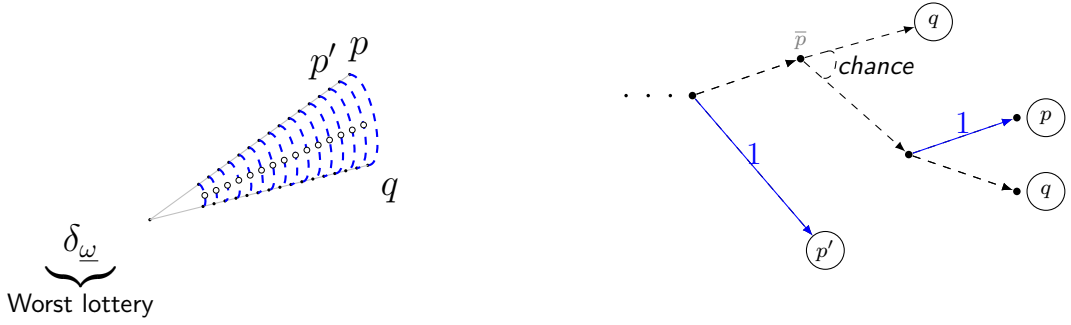
Lemma: Arbitrary Punishments to mixture-averse types

If type 1 has mixture-averse preferences and p, q are two distinct lotteries such that $q \succsim_2 p$, and $p \sim_1 q$, then $(\delta_{\underline{\omega}}, q) \in \mathrm{IMP}$.

Targeted Punishments (two types)

Lemma: Arbitrary Punishments to mixture-averse types

If type 1 has mixture-averse preferences and p, q are two distinct lotteries such that $q \succsim_2 p$, and $p \sim_1 q$, then $(\delta_{\underline{\omega}}, q) \in \text{IMP}$.

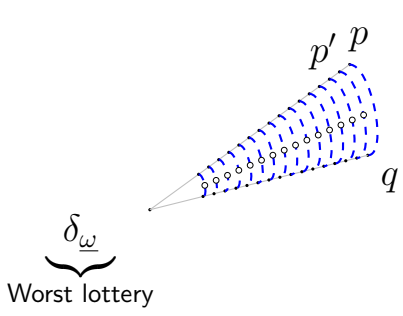


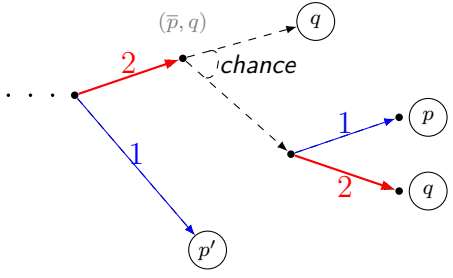
Mechanism leaving the mixture-averse type 1 indifferent at each point

Targeted Punishments (two types)

Lemma: Arbitrary Punishments to mixture-averse types

If type 1 has mixture-averse preferences and p, q are two distinct lotteries such that $q \succsim_2 p$, and $p \sim_1 q$, then $(\delta_{\underline{\omega}}, q) \in \text{IMP}$.





Mechanism leaving the mixture-averse type 1 indifferent at each point

Targeted Punishments (two types)

Lemma: Arbitrary Punishments to mixture-averse types

If type 1 has mixture-averse preferences and p, q are two distinct lotteries such that $q \succsim_2 p$, and $p \sim_1 q$, then $(\delta_{\underline{\omega}}, q) \in \mathrm{IMP}$.

Lemma (Theorem 2) for binary types

1. If type 1 has mixture-averse preferences, then $\mathrm{IC}_2 \subseteq \mathrm{IMP}$.
2. If both types have mixture-averse preferences, then $L^2 = \mathrm{IMP}$

Extensions

Extensions

- **Full (unique) implementation:** Characterization extends when all SPE must implement the target (requires preferences that are "sufficiently different") ▸ Details

Extensions

- **Full (unique) implementation:** Characterization extends when all SPE must implement the target (requires preferences that are "sufficiently different") Details
- **Multiple agents:** Revelation Principle holds iff all types of all agents are EU; Only mixture-averse types $\Rightarrow$ any jointly feasible allocation implementable Details

Extensions

- **Full (unique) implementation:** Characterization extends when all SPE must implement the target (requires preferences that are "sufficiently different") Details
- **Multiple agents:** Revelation Principle holds iff all types of all agents are EU; Only mixture-averse types $\Rightarrow$ any jointly feasible allocation implementable Details
- **Mixture-loving (concave) preferences:** Generally not all allocations are implementable Details

Extensions

- **Full (unique) implementation:** Characterization extends when all SPE must implement the target (requires preferences that are "sufficiently different") Details
- **Multiple agents:** Revelation Principle holds iff all types of all agents are EU; Only mixture-averse types $\Rightarrow$ any jointly feasible allocation implementable Details
- **Mixture-loving (concave) preferences:** Generally not all allocations are implementable Details
- **Ambiguity:** Results extend to Anscombe–Aumann preferences: mechanisms reveal objective randomization (not ambiguous states) Details

Discussion and Conclusion

Discussion & Conclusion

- Revelation Principle $\iff$ all agents are EU maximizers

Discussion & Conclusion

- Revelation Principle $\iff$ all agents are EU maximizers
- A non-EU type $\Rightarrow$ dynamic & stochastic mechanisms expand implementability

Discussion & Conclusion

- Revelation Principle $\iff$ all agents are EU maximizers
- A non-EU type $\Rightarrow$ dynamic & stochastic mechanisms expand implementability
  - Strictly mixture-averse $\Rightarrow$ No IC constraints $\Rightarrow$ Full surplus extraction

Discussion & Conclusion

- Revelation Principle $\iff$ all agents are EU maximizers
- A non-EU type $\Rightarrow$ dynamic & stochastic mechanisms expand implementability
  - Strictly mixture-averse $\Rightarrow$ No IC constraints $\Rightarrow$ Full surplus extraction
  - Mixture-loving: Characterize non-trivial limits on implementation

Discussion & Conclusion

- Revelation Principle $\iff$ all agents are EU maximizers
- A non-EU type $\Rightarrow$ dynamic & stochastic mechanisms expand implementability
  - Strictly mixture-averse $\Rightarrow$ No IC constraints $\Rightarrow$ Full surplus extraction
  - Mixture-loving: Characterize non-trivial limits on implementation
  - Betweenness preferences: Open question!

Discussion & Conclusion

- Revelation Principle $\iff$ all agents are EU maximizers
- A non-EU type $\Rightarrow$ dynamic & stochastic mechanisms expand implementability
  - Strictly mixture-averse $\Rightarrow$ No IC constraints $\Rightarrow$ Full surplus extraction
  - Mixture-loving: Characterize non-trivial limits on implementation
  - Betweenness preferences: Open question!

Implication: Restricting mechanisms is needed to avoid extreme conclusions:

- *Deterministic only*: contracts with full specification
- *Static only*: limits extraction but excludes impractical formats
- *Existing formats only*: empirically relevant but sacrifices optimal design

Thank You!

Questions?

Auxiliary Slides

Subgame Perfect Equilibrium

Back

- Fix a mechanism (T, A, m)
  - Non-terminal histories $H_t = A^{t-1} \times [0,1]^t$, $H = \bigcup_{1 \leq t \leq T} H_t$
- A *strategy* for type θ is a mapping $\sigma_\theta : H \to A$.
  - Write $P(a_t, \sigma_\theta \mid h_{t-1}) \in L$ for the outcome lottery conditional on choosing a_t at h_{t-1} and then σ_θ afterwards

Definition (Subgame Perfect Equilibrium)

A strategy profile $\sigma = (\sigma_\theta)_{\theta \in \Theta}$ is a subgame perfect equilibrium of (T, A, m) if for each $\theta \in \Theta$, $0 < t \leq T$, and $h_{t-1} \in H_t$,

$$\sigma_\theta(h_{t-1}) \in \arg\max_{a_t \in A} \left(P(a_t, \sigma_\theta \mid h_{t-1}) : \succsim_\theta \right).$$

Corollary (Full-Surplus Extraction)

Assume strictly mixture-averse preferences. For any deterministic $x : \Theta \to X$ and any $\epsilon > 0$, there exists a mechanism implementing an allocation $p \in \Delta(X \times T)^{|\Theta|}$ s.t.

$$\mathbb{E}_p[t] \geq \mathbb{E}_p[v(\theta, x)] - \epsilon, \quad \text{and} \quad v(\theta, x) - t \geq 0$$

- No *information rents*: agent's utility from participating is always 0
- Intuition: principal leverages time-inconsistency to extract rents
- Ahn & Zhang & Zhang (2026): similar result with one-dimensional continuum of types [0,1] and biseparable preferences

Full Implementation

Back

- *Partial implementation*: some SPE implements the target allocation
- *Full implementation* (IMP_F): every SPE implements the target (robust to equilibrium selection)

Definition

Preferences are *strongly different* if for each full-support lottery p there exist full-support lotteries q, q' such that $q' \succ_1 p \succ_1 q$ and $q \succ_2 p \succ_2 q'$.

Theorem

If $\Theta = \{1, 2\}$, $\Theta = \mathrm{EU} \cup \mathrm{MA}$, and preferences are strongly different, then

$$\mathrm{IMP}_F = \bigcap_{\theta \in \mathrm{EU}} \mathrm{IC}_\theta.$$

Multiple Agents — Setup & Feasibility

Back

- Agents: $i = 1, \ldots, n$; types: Θ^i
- Outcome: $\omega = (\omega^1, \ldots, \omega^n) \in \Omega^1 \times \cdots \times \Omega^n$
- Preferences: agent i depends only on distribution of ω^i
- Independent lotteries: $(p^1, \ldots, p^n) \in \Delta(\Omega^1)^{|\Theta^1|} \times \cdots \times \Delta(\Omega^n)^{|\Theta^n|}$
- Exogenous joint feasible set: $F \subseteq \times_{i=1}^n \Delta(\Omega^i)^{|\Theta^i|}$

Example: Border condition

- Single object + transfers: $\Omega^i = \{0, 1\} \times T$, $T \subset \mathbb{R}$ finite
- Prior of types μ; Allocation probability $q^i(\theta^i) = \mathbb{P}[\omega_1^i = 1 \mid \theta^i]$:

$$\sum_{i=1}^{n} \sum_{\theta^i \in A^i} q^i(\theta^i) \mu^i(\theta^i) \;\leq\; 1 - \prod_{i=1}^{n} \Big(1 - \sum_{\theta^i \in A^i} \mu^i(\theta^i)\Big) \quad \forall (A^i)$$

Multiple Agents — No Marginal Constraints & Main Result

Back

Theorem

1. The revelation principle holds if and only if $\Theta^i \subseteq \mathrm{EU}$ for all i.
2. Suppose $\Theta^i \subseteq \mathrm{MA}$ (mixture-averse) for all i, then $\mathrm{IMP} = F$.

Concave Preferences

Back

- Types have concave utility v_θ: prefer mixtures over extreme lotteries
- Examples: maxmin and cautious utility preferences
- Degenerate-lottery value: $b_\theta(p) := \sum_{\omega\in\Omega} p(\omega)\, v_\theta(\delta_\omega)$
- $B_\theta := \left\{p : v_\theta(p_\theta) \geq \max_{\theta'\in\Theta} b_\theta(p_{\theta'})\right\}$

Proposition

If each v_θ is concave, then $\mathrm{IMP} \subseteq \cap_{\theta\in\Theta} B_\theta$.

- b_θ is the maximum punishment assignable to type θ in any mechanism

Ambiguity

Back

- Anscombe–Aumann framework: states S, consequences X, acts $\omega : S \to X$,
  - State-independent preferences over lotteries of acts $\Delta(\Omega)$
- Our mechanisms reveal objective randomization (flipping a coin)
- Do not reveal information about the ambiguous state (Bose and Daripa (2009))

Proposition

1. $\mathrm{IMP} = \cap_{\theta \in \Theta} \mathrm{IC}_\theta$ iff all types have EU preferences
2. If $\Theta \subseteq \mathrm{EU} \cup \mathrm{MA}$, then $\mathrm{IMP} = \cap_{\theta \in \mathrm{EU}} \mathrm{IC}_\theta$
3. If all types have concave preferences, then $\mathrm{IMP} \subseteq \cap_{\theta \in \Theta} B_\theta$

Relation to Dubra, Maccheroni, and Ok (2004)

Back

Theorem 1 is related to **Dubra, Maccheroni, and Ok (2004)**: They show any preference relation that is continuous and satisfies independence (but not necessarily completeness):

(i) satisfies the strong version of independence, and

(ii) is represented by the intersection of some expected utility preferences.

Our Thm 1 is stronger:

- If in addition $\succeq^*$ is defined by the set $\mathrm{IC}_1 \cap \mathrm{IC}_2$, then $\succeq_1$ and $\succeq_2$ *each* has an expected utility representation.
- Uses the fact that there exist worst outcome $\underline{\omega}$.

Theorem 1: Sketch of Proof

Focus on showing $RP \Rightarrow EU$ for the case of two types

Key: Relate RP to the Independence Axiom

Definition

A set $C \subseteq L^2$ satisfies *strong independence* (SI) if for all lotteries $p, q, r \in L$ and all $\alpha \in (0,1)$:

$$(p,q) \in C \iff \big(\alpha p + (1-\alpha)r, \alpha q + (1-\alpha)r\big) \in C.$$

1. $\mathrm{IC}_1 \cap \mathrm{IC}_2 = \mathrm{IMP}$ is convex $\Rightarrow$ $\mathrm{IC}_1 \cap \mathrm{IC}_2$ satisfies SI
2. $\Rightarrow$ $\mathrm{IC}_1^c \cap \mathrm{IC}_2^c$, $\mathrm{IC}_1 \cap \mathrm{IC}_2^c$, $\mathrm{IC}_1^c \cap \mathrm{IC}_2$ satisfy SI
3. $\Rightarrow$ IC_1 and IC_2 satisfy SI
4. $\Rightarrow$ $\succeq_1$ and $\succeq_2$ satisfy the independence axiom

Key: Relate RP to the Independent Axiom

Definition

A set $C \subseteq L^2$ satisfies *strong independence* (SI) if for all lotteries $p, q, r \in L$ and all $\alpha \in (0,1)$:

$$(p,q) \in C \iff \big(\alpha p + (1-\alpha)r, \alpha q + (1-\alpha)r\big) \in C.$$

1. $\mathrm{IC}_1 \cap \mathrm{IC}_2 = \mathrm{IMP}$ is convex $\Rightarrow$ $\mathrm{IC}_1 \cap \mathrm{IC}_2$ satisfies SI ▸ Details
2. $\Rightarrow$ $\mathrm{IC}_1^c \cap \mathrm{IC}_2^c$, $\mathrm{IC}_1 \cap \mathrm{IC}_2^c$, $\mathrm{IC}_1^c \cap \mathrm{IC}_2$ satisfy SI
3. $\Rightarrow$ IC_1 and IC_2 satisfy SI
4. $\Rightarrow$ $\succeq_1$ and $\succeq_2$ satisfy the independence axiom

Theorem 2: Sketch of Proof

Generalization to more than 2 types

Extending to more than two types

Induction: Assume all allocations implementable for k mixture-averse types;

Goal: implement $(p_1, ..., p_{k+1})$ for $k+1$ mixture-averse types.

Extending to more than two types

Induction: Assume all allocations implementable for k mixture-averse types;

Goal: implement $(p_1, ..., p_{k+1})$ for $k+1$ mixture-averse types. By induction:

- Mechanism M_{-2}: implements $(p_1, \textbf{?}, p_3, \ldots, p_{k+1})$
- Mechanism M_{-1}: implements $(\textbf{?}, p_2, p_3, \ldots, p_{k+1})$

Extending to more than two types

Induction: Assume all allocations implementable for k mixture-averse types;

Goal: implement $(p_1, ..., p_{k+1})$ for $k+1$ mixture-averse types. By induction:

- Mechanism M_{-2}: implements $(p_1, \textbf{?}, p_3, \ldots, p_{k+1})$
- Mechanism M_{-1}: implements $(\textbf{?}, p_2, p_3, \ldots, p_{k+1})$

Auxiliary problem: Mechanism with 3 "outcomes":

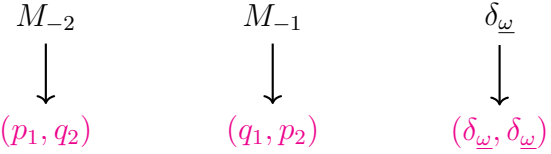
M_{-2} M_{-1} $\delta_{\underline{\omega}}$ (worst lottery)

- Each type $\theta = 3, \ldots, k+1$ chooses M_{-1} (or M_{-2}) to get p_θ
- Reduces problem to two-type case for types 1 and 2.

▸ Details

Auxiliary problem

Back



Goal: Each type $\theta \in \{1, 2\}$ gets p_θ

- If $\{p_\theta, q_\theta, \delta_{\underline{\omega}}\}$ are *linearly independent*, then θ has mixture-averse preferences over $\Delta(\{p_\theta, q_\theta, \delta_{\underline{\omega}}\}) \Rightarrow$ Theorem 2 applies
- If $\{p_\theta, q_\theta, \delta_{\underline{\omega}}\}$ are *linearly dependent*, construct p_θ more attractive than q_θ

Kőszegi–Rabin is mixture-averse

Back

Write $x_\omega := u(\pi(\omega))$. If $\gamma > 1$ and $u \circ \pi$ is one-to-one, then

$$v_\theta(p) = \sum_\omega p(\omega)\, x_\omega + (1-\gamma) \sum_{\omega<\omega'} p(\omega)p(\omega')\, |x_\omega - x_{\omega'}|.$$

The first term is linear in p. The second is a strictly convex quadratic form on the simplex when $\gamma > 1$: for $p \neq q$ and $\alpha \in (0,1)$,

$$v_\theta\big(\alpha p + (1-\alpha)q\big) < \alpha v_\theta(p) + (1-\alpha) v_\theta(q).$$

Hence if $p \sim_\theta q$ and $p \neq q$, the mixture is strictly worse: v_θ is mixture-averse.

References

David Ahn, Mengxi Zhang, and Mu Zhang. Mechanism design beyond expected utility: Dynamic (in) consistency and full surplus extraction. *Working Paper*, 2026.

Manuel Amador, Iván Werning, and George-Marios Angeletos. Commitment vs. flexibility. *Econometrica*, 74(2):365–396, 2006.

Benjamin Balzer and Antonio Rosato. Never say never: Optimal exclusion and reserve prices with expectations-based loss-averse buyers. *Journal of Economic Theory*, 228:106045, 2025.

Benjamin Balzer, Antonio Rosato, and Jonas von Wangenheim. Dutch vs. first-price auctions with expectations-based loss-averse bidders. *Journal of Economic Theory*, 205:105545, 2022.

Subir Bose and Arup Daripa. A dynamic mechanism and surplus extraction under ambiguity. *Journal of Economic theory*, 144(5):2084–2114, 2009.

David Dillenberger and Uzi Segal. *Allocation Mechanisms with Mixture-Averse Preferences*. Penn Institute for Economic Research, Department of Economics, University of , 2024.

Paul Dütting, Michal Feldman, Daniel Peretz, and Larry Samuelson. Ambiguous contracts. *Econometrica*, 92(6):1967–1992, 2024.

Roland Eisenhuth. Reference-dependent mechanism design. *Economic Theory Bulletin*, 7(1):77–103, 2019.

Simone Galperti. Optimal screening of time inconsistency. *Econometrica*, 2014.

Alex Gershkov, Benny Moldovanu, Philipp Strack, and Mengxi Zhang. Optimal auctions: Non-expected utility and constant risk aversion. *The review of economic studies*, 89(5): 2630–2662, 2022.

Alex Gershkov, Benny Moldovanu, Philipp Strack, and Mengxi Zhang. Optimal insurance: dual utility, random losses, and adverse selection. *American Economic Review*, 113(10):2581–2614, 2023.

Alex Gershkov, Benny Moldovanu, Philipp Strack, and Mengxi Zhang. Optimal security design for risk-averse investors. *American Economic Review*, 115(6):2050–2092, 2025.

Daniel Gottlieb and Xingtan Zhang. Long-term contracting with time-inconsistent agents. *Econometrica*, 89(2):793–824, 2021.

Jerry Green. making book against oneself, the independence axiom, and nonlinear utility theory. *The Quarterly Journal of Economics*, 102(4):785–796, 1987.

Paul Heidhues and Botond Kőszegi. Competition and price variation when consumers are loss averse. *American Economic Review*, 98(4):1245–1268, 2008.

Edi Karni and David Schmeidler. Atemporal dynamic consistency and expected utility theory. *Journal of Economic Theory*, 54(2):401–408, 1991.

Andreas Lange and Anmol Ratan. Multi-dimensional reference-dependent preferences in sealed-bid auctions–how (most) laboratory experiments differ from the field. *Games and Economic Behavior*, 68(2):634–645, 2010.

Antonio Rosato and Agnieszka A Tymula. Loss aversion and competition in vickrey auctions: Money ain't no good. *Games and Economic Behavior*, 115:188–208, 2019.

Jonas von Wangenheim. English versus vickrey auctions with loss-averse bidders. *Journal of Economic Theory*, 197:105328, 2021.