Regret-Minimizing Project Choice

Yingni Guo, Eran Shmaya

BFI Economic Theory Conference
Aug 2023

Motivating example I

- A department has a hiring slot open

Motivating example I

- A department has a hiring slot open
- It learns that the available candidates are $\{\blacktriangle, \blacksquare, \blacklozenge\}$

Motivating example I

- A department has a hiring slot open
- It learns that the available candidates are $\{\blacktriangle, \blacksquare, \blacklozenge\}$
- The department proposes one candidate to the dean, who decides whether to make an offer or not

Motivating example II

- An employee has identified three approaches to cutting energy costs at a heater (Ross (1986)):

  $\{$use advanced combustion controls, replace heater, add heat exchangers$\}$

Motivating example II

- An employee has identified three approaches to cutting energy costs at a heater (Ross (1986)):

  $\{$use advanced combustion controls, replace heater, add heat exchangers$\}$
- These projects overlap; he can advocate at most one

Motivating example II

- An employee has identified three approaches to cutting energy costs at a heater (Ross (1986)):

  {use advanced combustion controls, replace heater, add heat exchangers}
- These projects overlap; he can advocate at most one
- The employee proposes a project to the headquarters, which decides whether to approve or not

Motivating example III

- A division has a chance to choose an office building

Motivating example III

- A division has a chance to choose an office building
- It learns that the available locations are $\{L, M, N, P\}$

Motivating example III

- A division has a chance to choose an office building
- It learns that the available locations are $\{L, M, N, P\}$
- The division proposes a short list for the headquarters to evaluate and choose

Project choice

- Project choice often involves two parties:

Project choice

- Project choice often involves two parties:
  - an agent learns which projects are available and proposes project(s)

Project choice

- Project choice often involves two parties:
  - an agent learns which projects are available and proposes project(s)
  - a principal evaluates the proposed project(s) and makes the choice

Project choice

- Project choice often involves two parties:
  - an agent learns which projects are available and proposes project(s)
  - a principal evaluates the proposed project(s) and makes the choice

- Proposing bias: the agent has a tendency to propose his favorite project and hide his less preferred ones

Questions

- How shall the principal encourage the agent to propose the principal's preferred projects?

Questions

- How shall the principal encourage the agent to propose the principal's preferred projects?

  If he wants a mechanism that works "fairly well" in all circumstances, what shall this mechanism look like?

Questions

- How shall the principal encourage the agent to propose the principal's preferred projects?

  If he wants a mechanism that works "fairly well" in all circumstances, what shall this mechanism look like?

- Even though the principal chooses at most one project, in some applications, the agent proposes multiple projects.

Questions

- How shall the principal encourage the agent to propose the principal's preferred projects?

  If he wants a mechanism that works "fairly well" in all circumstances, what shall this mechanism look like?

- Even though the principal chooses at most one project, in some applications, the agent proposes multiple projects.

  What is the strategic role of multiproject proposals?

Questions

- How shall the principal encourage the agent to propose the principal's preferred projects?

  If he wants a mechanism that works "fairly well" in all circumstances, what shall this mechanism look like?

- Even though the principal chooses at most one project, in some applications, the agent proposes multiple projects.

  What is the strategic role of multiproject proposals?

  Do several projects within the proposal have a chance of being chosen?

Roadmap

- Model
- Single-project environment
- Multiproject environment
- Discussion

Projects and payoffs

- $D = [\underline{u}, 1] \times [\underline{v}, 1]$: the set of all possible projects

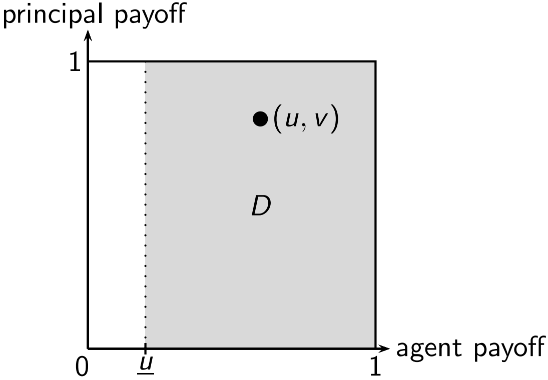

Projects and payoffs

- $D = [\underline{u}, 1] \times [\underline{v}, 1]$: the set of all possible projects

  For the talk $\underline{v} = 0$, $\underline{u} \in [0, 1]$

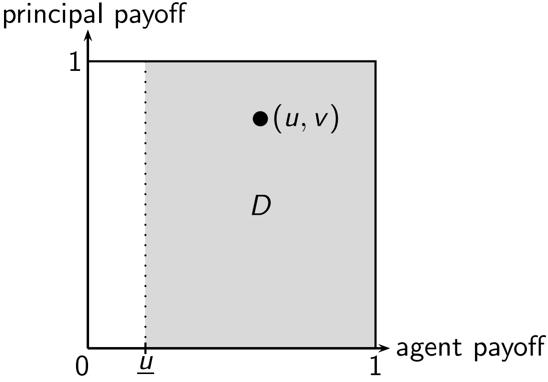

Projects and payoffs

- If project (u, v) is chosen: the agent gets u; the principal gets v

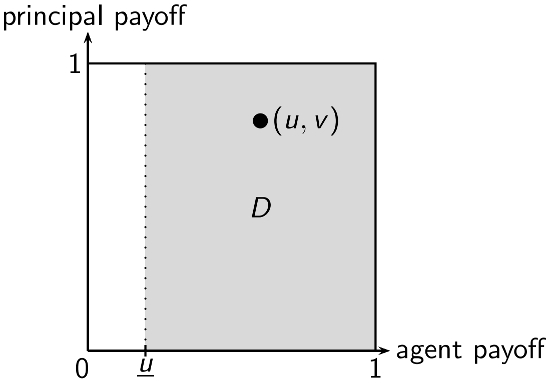

Projects and payoffs

- If project (u, v) is chosen: the agent gets u; the principal gets v
- If no project is chosen, both players get zero

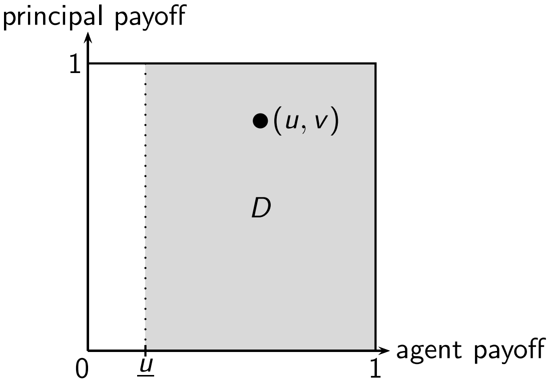

The agent's type

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

The agent's type

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

In our motivating examples:

- A department has a hiring slot open

The agent's type

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

In our motivating examples:

- A department has a hiring slot open

  It learns that the available candidates are $A = \{\blacktriangle, \blacksquare, \blacklozenge\}$

The agent's type

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

In our motivating examples:

- A department has a hiring slot open

  It learns that the available candidates are $A = \{\blacktriangle, \blacksquare, \blacklozenge\}$
- An employee has identified three approaches to cutting energy costs at a heater:

  $A = \{\text{use advanced combustion controls}, \text{ replace heater}, \text{ add heat exchangers}\}$

The agent's type

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

In our motivating examples:

- A department has a hiring slot open

  It learns that the available candidates are $A = \{\blacktriangle, \blacksquare, \blacklozenge\}$
- An employee has identified three approaches to cutting energy costs at a heater:

  $A = \{\text{use advanced combustion controls}, \ \text{replace heater}, \ \text{add heat exchangers}\}$
- A division has a chance to choose an office building

The agent's type

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

In our motivating examples:

- A department has a hiring slot open

  It learns that the available candidates are $A = \{\blacktriangle, \blacksquare, \blacklozenge\}$
- An employee has identified three approaches to cutting energy costs at a heater:

  $A = \{\text{use advanced combustion controls}, \text{ replace heater}, \text{ add heat exchangers}\}$
- A division has a chance to choose an office building

  It learns that the available locations are $A = \{L, M, N, P\}$

The agent's proposal

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

- The agent proposes a set P of projects to the principal

  He can propose only available projects: $P \subseteq A$

The agent's proposal

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

- The agent proposes a set P of projects to the principal

  He can propose only available projects: $P \subseteq A$

- $|P| \leqslant K$ for some fixed number K

The agent's proposal

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

- The agent proposes a set P of projects to the principal

  He can propose only available projects: $P \subseteq A$

- $|P| \leqslant K$ for some fixed number K
  single-project environment: $K = 1$

The agent's proposal

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

- The agent proposes a set P of projects to the principal

  He can propose only available projects: $P \subseteq A$

- $|P| \leqslant K$ for some fixed number K

  single-project environment: $K = 1$

  multiproject environment: $K = \infty$

The agent's proposal

- $A \subseteq D$: a finite set of available projects

  A is the agent's private information (the agent's type)

- The agent proposes a set P of projects to the principal

  He can propose only available projects: $P \subseteq A$

- $|P| \leqslant K$ for some fixed number K

  single-project environment: $K = 1$

  multiproject environment: $K = \infty$

  intermediate environment: $2 \leqslant K < \infty$

Mechanism

- A mechanism ρ attaches to each proposal P a subprobability measure $\rho(\cdot|P)$ over P:

$$\rho((u,v)|P) \geqslant 0, \qquad \sum_{(u,v)\in P} \rho((u,v)|P) \leqslant 1$$

Mechanism

- A mechanism ρ attaches to each proposal P a subprobability measure $\rho(\cdot|P)$ over P:

$$\rho((u,v)|P) \geqslant 0, \quad \sum_{(u,v)\in P} \rho((u,v)|P) \leqslant 1$$

with probability $\rho((u,v)|P)$, project (u,v) is chosen;

Mechanism

- A mechanism ρ attaches to each proposal P a subprobability measure $\rho(\cdot|P)$ over P:

$$\rho((u,v)|P) \geqslant 0, \quad \sum_{(u,v)\in P} \rho((u,v)|P) \leqslant 1$$

with probability $\rho((u,v)|P)$, project (u,v) is chosen;
with probability $1 - \sum_{(u,v)\in P} \rho((u,v)|P)$, no project is chosen

Mechanism

- A mechanism ρ attaches to each proposal P a subprobability measure $\rho(\cdot|P)$ over P:

$$\rho((u,v)|P) \geqslant 0, \quad \sum_{(u,v)\in P} \rho((u,v)|P) \leqslant 1$$

  with probability $\rho((u,v)|P)$, project (u,v) is chosen;
  with probability $1 - \sum_{(u,v)\in P} \rho((u,v)|P)$, no project is chosen

- The principal commits to a mechanism ρ

The principal's problem

Given mechanism ρ and the agent's type A:

- the agent chooses a proposal P to maximize his expected payoff:

$$P \in \operatorname*{argmax}_{P \subseteq A, |P| \leqslant K} \sum_{(u,v) \in P} \rho((u,v)|P)u$$

The principal's problem

Given mechanism ρ and the agent's type A:

- the agent chooses a proposal P to maximize his expected payoff:

$$P \in \operatorname*{argmax}_{P \subseteq A, |P| \leqslant K} \sum_{(u,v) \in P} \rho((u,v)|P)u$$

- the principal's regret is

$$\mathrm{RGRT}(\rho, A) = \max_{(u,v) \in A} v - \sum_{(u,v) \in P} \rho((u,v)|P)v$$

The principal's problem

- The principal's worst-case regret under mechanism ρ is

$$\mathrm{WCR}(\rho) = \sup_{A \subseteq D, |A| < \infty} \mathrm{RGRT}(\rho, A),$$

where the supremum is over all possible types of the agent

The principal's problem

- The principal's worst-case regret under mechanism ρ is

$$\mathrm{WCR}(\rho) = \sup_{A \subseteq D, |A| < \infty} \mathrm{RGRT}(\rho, A),$$

  where the supremum is over all possible types of the agent

- The principal chooses ρ that minimizes his worst-case regret

Related literature

- Project choice: Armstrong and Vickers (2010), Nocke and Whinston (2013)
- Mechanism design with worst-case regret: Hurwicz and Shapiro (1978), Bergemann and Schlag (2008, 2011), Manski (2011), Renou and Schlag (2011), Beviá and Corchón (2019), Malladi (2022), Guo and Shmaya (2023)

Roadmap

- Model
- Single-project environment
- Multiproject environment
- Discussion

Single-project environment

- The agent can propose at most one project

Single-project environment

- The agent can propose at most one project
- We let $\alpha(u, v) \in [0, 1]$ denote the approval probability if the agent proposes project (u, v), instead of using $\rho((u, v)|\{(u, v)\})$

Intuition: deterministic mechanisms for now

- Suppose that only deterministic mechanisms are allowed

$$\alpha(u, v) \in \{0, 1\}$$

Intuition: deterministic mechanisms for now

- Suppose that only deterministic mechanisms are allowed

$$\alpha(u, v) \in \{0, 1\}$$

- Shall the principal approve ▲?

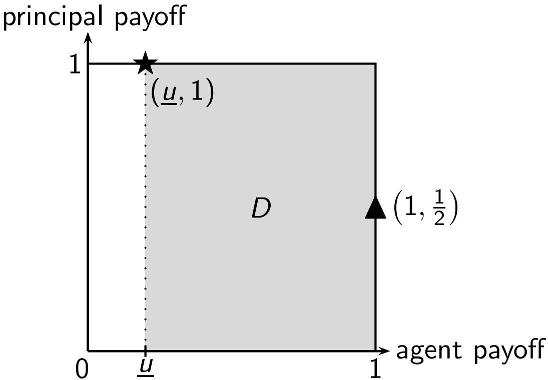

Intuition: deterministic mechanisms for now

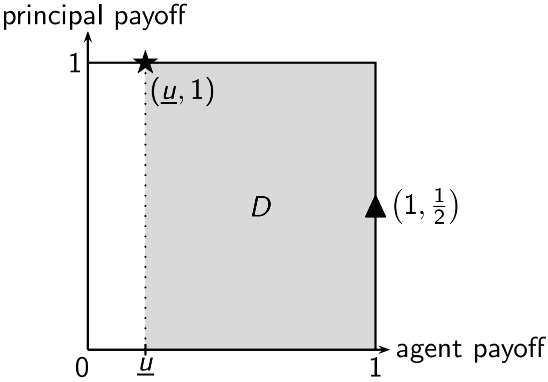

Intuition: deterministic mechanisms for now

- If the principal approves ▲ and $A = \{▲, ★\}$, the agent will propose ▲ and the principal suffers regret of $1/2$

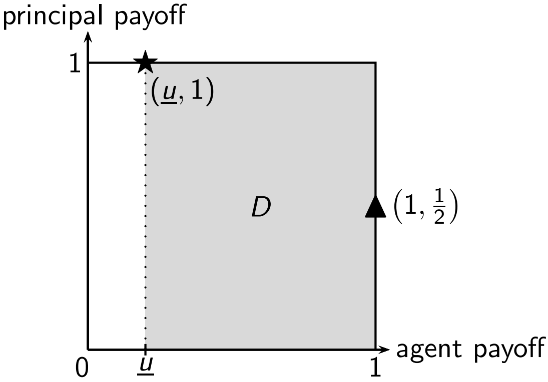

Intuition: deterministic mechanisms for now

- If the principal approves ▲ and $A = \{▲, ★\}$, the agent will propose ▲ and the principal suffers regret of $1/2$
- If the principal rejects ▲ and $A = \{▲\}$, the principal suffers regret of $1/2$

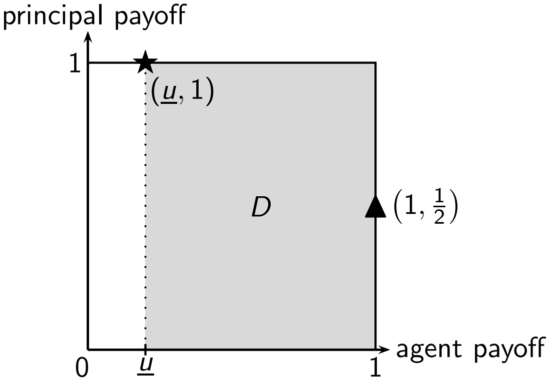

Deterministic mechanisms

Claim

In the single-project environment, the principal's worst-case regret under any deterministic mechanism is at least $1/2$.

Intuition: how randomization helps

- With randomized mechanisms, the principal approves ▲ with probability $\underline{u}$

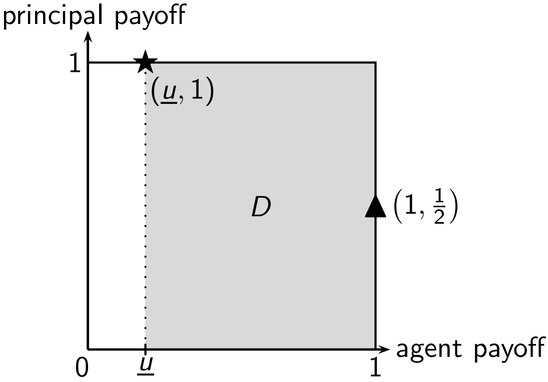

Intuition: how randomization helps

- With randomized mechanisms, the principal approves ▲ with probability $\underline{u}$
- If $A = \{▲, ★\}$, the agent is willing to propose ★

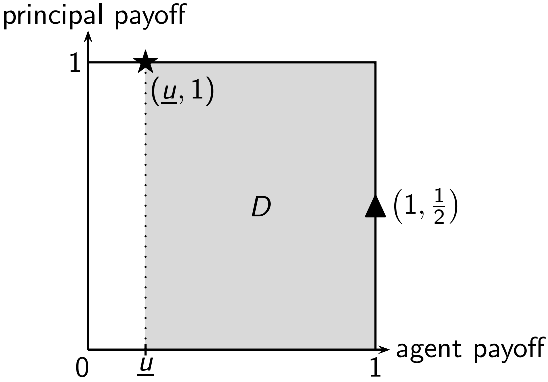

Intuition: how randomization helps

- With randomized mechanisms, the principal approves ▲ with probability $\underline{u}$
- If $A = \{▲, ★\}$, the agent is willing to propose ★
- If $A = \{▲\}$, the amount of inefficient rejection is reduced

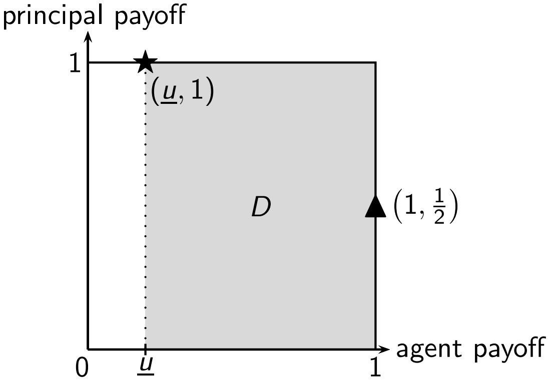

Lower bound on the worst-case regret

Theorem

(i) The worst-case regret under any mechanism is at least R^s:

$$R^s \equiv \max_{v \in [0,1]} \min \{1 - v, (1 - \underline{u})v\} = \frac{1 - \underline{u}}{2 - \underline{u}}.$$

(ii) The two-tier mechanism:

$$\alpha^s(u, v) = \begin{cases} 1, & \text{if } v \geqslant 1 - R^s \text{ or } u = 0 \\ \underline{u}/u, & \text{if } v < 1 - R^s \text{ and } u > 0 \end{cases}$$

has the worst-case regret of R^s.

Proof: worst-case regret under any α is at least R^s

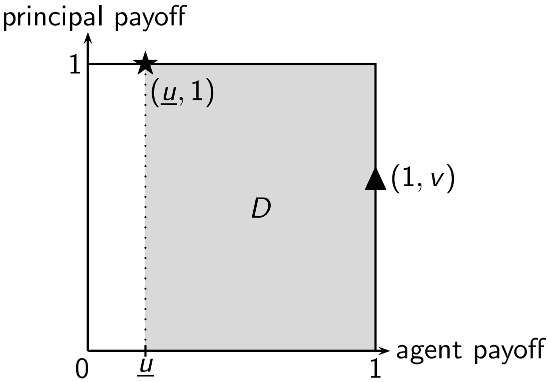

Proof: worst-case regret under any α is at least R^s

- If $\alpha(1, v) > \underline{u}$, then if $A = \{\blacktriangle, \bigstar\}$, the agent will propose $\blacktriangle$ so regret is at least $(1 - v)$

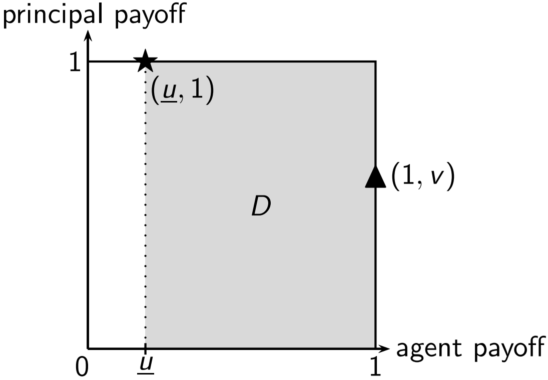

Proof: worst-case regret under any α is at least R^s

- If $\alpha(1, v) > \underline{u}$, then if $A = \{\blacktriangle, \bigstar\}$, the agent will propose $\blacktriangle$ so regret is at least $(1 - v)$
- If $\alpha(1, v) \leqslant \underline{u}$, then if $A = \{\blacktriangle\}$, regret is at least $(1 - \underline{u})v$

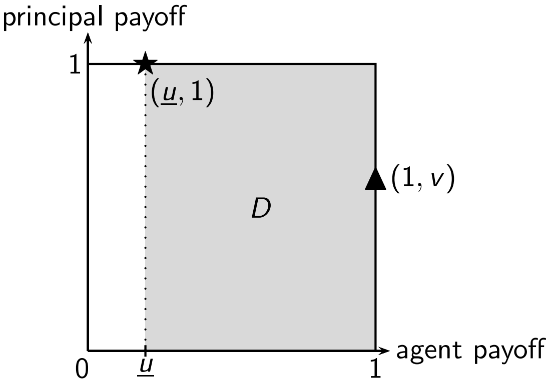

Proof: worst-case regret under any α is at least R^s

- If $\alpha(1, v) > \underline{u}$, then if $A = \{\blacktriangle, \bigstar\}$, the agent will propose $\blacktriangle$ so regret is at least $(1 - v)$
- If $\alpha(1, v) \leqslant \underline{u}$, then if $A = \{\blacktriangle\}$, regret is at least $(1 - \underline{u})v$

$\implies$ worst-case regret is at least $\min\{1 - v, (1 - \underline{u})v\}$ for any v

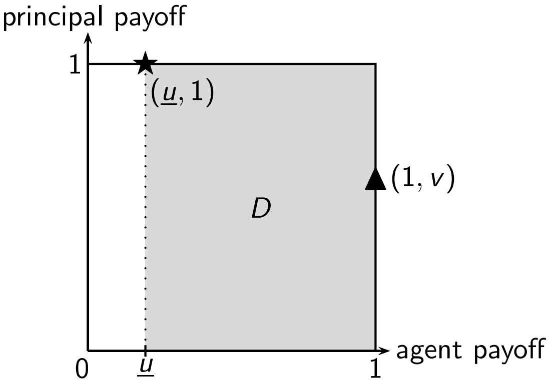

Optimal mechanism

Theorem

(i) The worst-case regret under any mechanism is at least R^s:

$$R^s \equiv \max_{v \in [0,1]} \min \{1 - v, (1 - \underline{u})v\} = \frac{1 - \underline{u}}{2 - \underline{u}}.$$

(ii) The two-tier mechanism:

$$\alpha^s(u, v) = \begin{cases} 1, & \text{if } v \geqslant 1 - R^s \text{ or } u = 0 \\ \underline{u}/u, & \text{if } v < 1 - R^s \text{ and } u > 0 \end{cases}$$

has the worst-case regret of R^s.

Two-tier mechanism

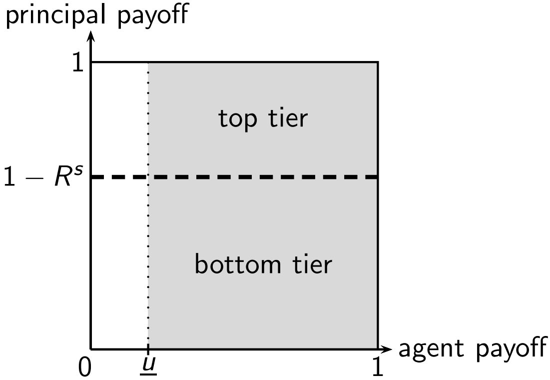

Two-tier mechanism

- If the agent proposes a top-tier project (u, v), his expected payoff is u

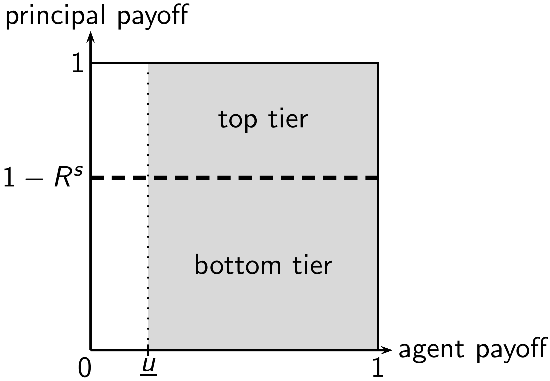

Two-tier mechanism

- If the agent proposes a top-tier project (u, v), his expected payoff is u
- If he proposes a bottom-tier project (u, v), his expected payoff is $\underline{u}$

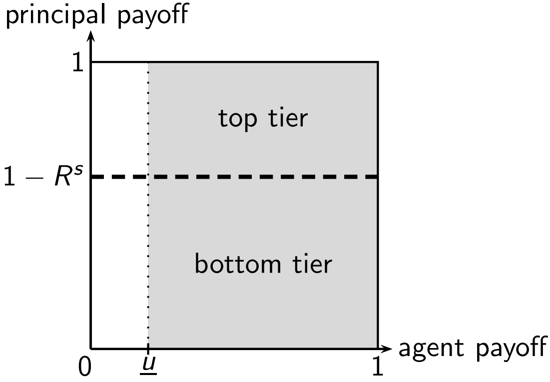

The agent's best response under α^s

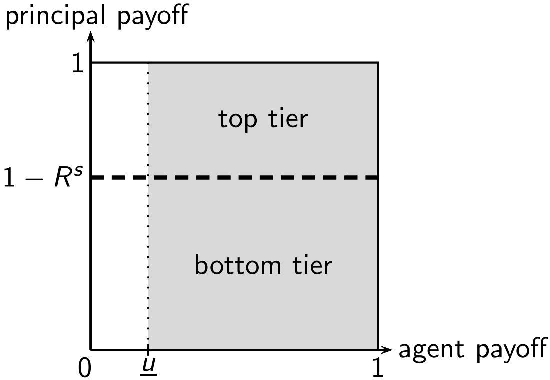

The agent's best response under α^s

- If the agent has projects with $v \geqslant 1 - R^s$, it is optimal to propose his favorite project among those with $v \geqslant 1 - R^s$

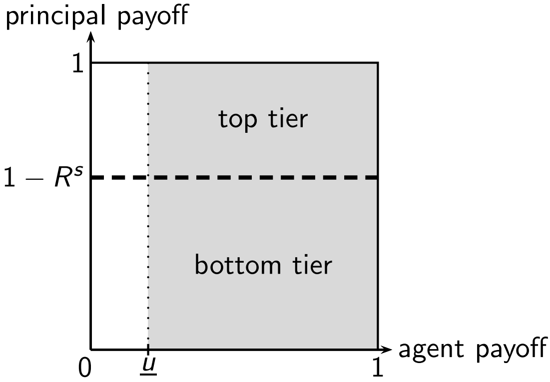

The agent's best response under α^s

- If the agent has projects with $v \geqslant 1 - R^s$, it is optimal to propose his favorite project among those with $v \geqslant 1 - R^s$
- Otherwise, it is optimal to propose a project that maximizes the principal's expected payoff $\alpha^s(u, v)v$

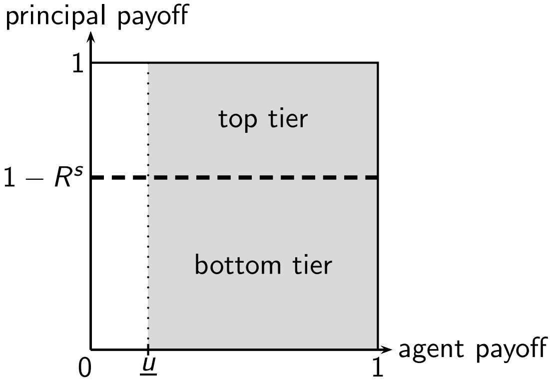

Proof: worst-case regret is R^s under α^s

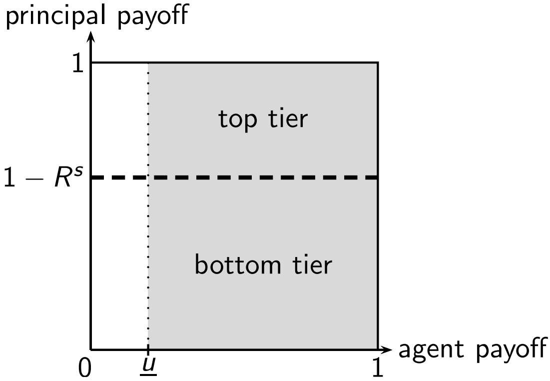

Proof: worst-case regret is R^s under α^s

- If the agent proposes a top-tier project (u, v), it is approved for sure. Regret is at most $(1 - v) \leqslant R^s$.

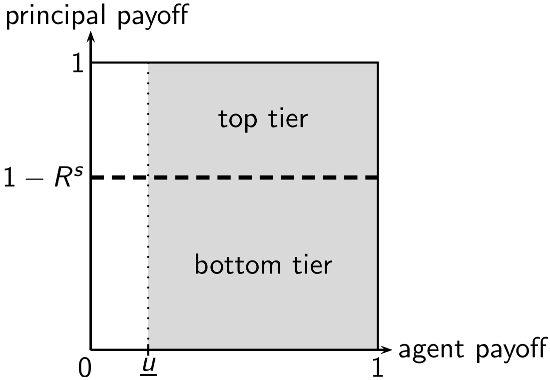

Proof: worst-case regret is R^s under α^s

- If the agent proposes a top-tier project (u, v), it is approved for sure. Regret is at most $(1 - v) \leqslant R^s$.
- Suppose that the agent proposes a bottom-tier (u, v). Let (u_p, v_p) be a principal's favorite project in A. Regret is:

$$v_p - \alpha^s(u, v)v \leqslant v_p - \alpha^s(u_p, v_p)v_p \leqslant (1 - \underline{u})v_p \leqslant R^s.$$

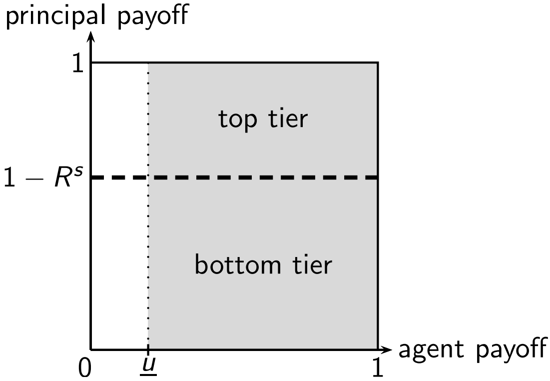

Worst-case regret R^s as a function of $\underline{u}$

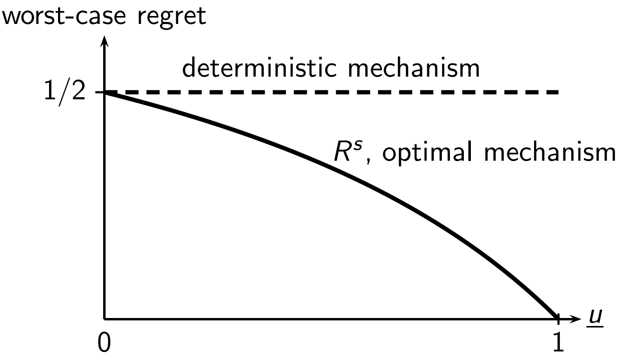

Comparison to Armstrong and Vickers (2010)

- The threshold for top tier depends only on v

Comparison to Armstrong and Vickers (2010)

- The threshold for top tier depends only on v
- We show how to reduce rejection in bottom tier without jeopardizing the agent's incentives to propose a top-tier project

Comparison to Armstrong and Vickers (2010)

- The threshold for top tier depends only on v
- We show how to reduce rejection in bottom tier without jeopardizing the agent's incentives to propose a top-tier project
- The approval probability $\alpha^s(u, v)$ is monotone decreasing in u

Roadmap

- Model
- Single-project environment
- Multiproject environment
- Discussion

Multiproject environment

- The agent can propose any subset of the available projects, $P \subset A$

Multiproject environment

- The agent can propose any subset of the available projects, $P \subset A$
- Given proposal P, $(u, v) \in P$ is chosen with probability $\rho((u, v)|P)$

Multiproject environment

- The agent can propose any subset of the available projects, $P \subset A$
- Given proposal P, $(u, v) \in P$ is chosen with probability $\rho((u, v)|P)$
- Revelation principle holds, so it is without loss to focus on mechanisms in which the agent optimally proposes all available projects, $P = A$

Intuition: single-project environment for now

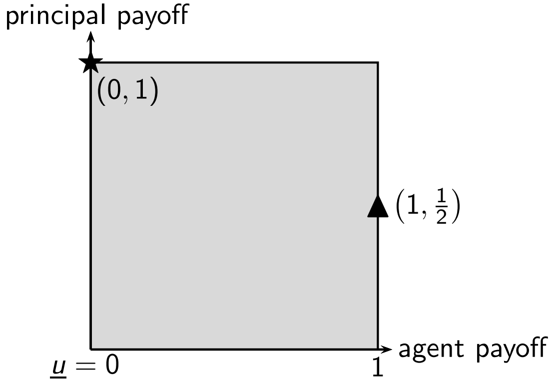

Intuition: single-project environment for now

- Let's revisit the single-project environment. Suppose that $\underline{u} = 0$

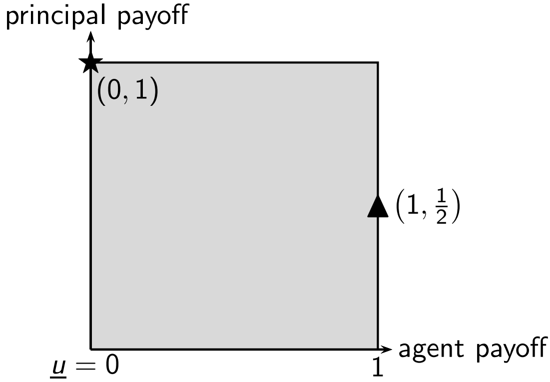

Intuition: single-project environment for now

- Let's revisit the single-project environment. Suppose that $\underline{u} = 0$
- If the principal want to incentivize the agent to propose ★ when $A = \{\bigstar, \blacktriangle\}$, he has to reject ▲

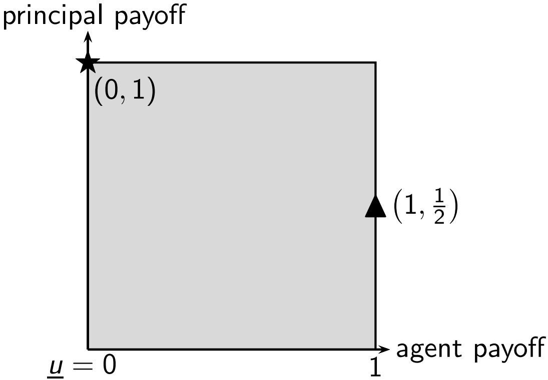

Intuition: single-project environment for now

- Let's revisit the single-project environment. Suppose that $\underline{u} = 0$
- If the principal want to incentivize the agent to propose ★ when $A = \{★, ▲\}$, he has to reject ▲
- Worst-case regret is $R^s = 1/2$

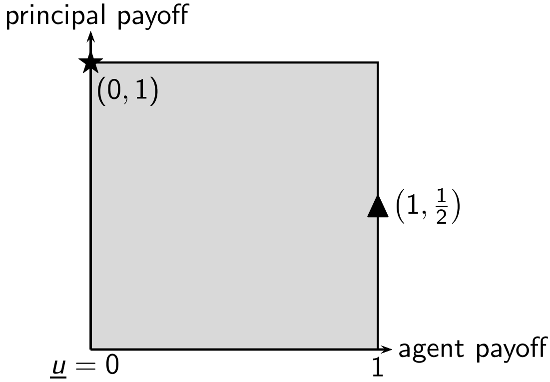

Intuition: how multiproject environment helps

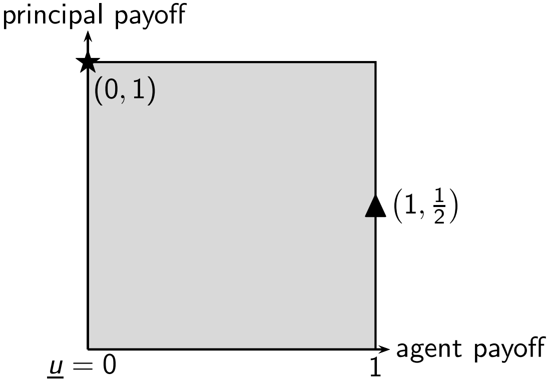

Intuition: how multiproject environment helps

- To incentive the agent to propose ★, his payoff from proposing {▲, ★} ⩾ his payoff from proposing ▲

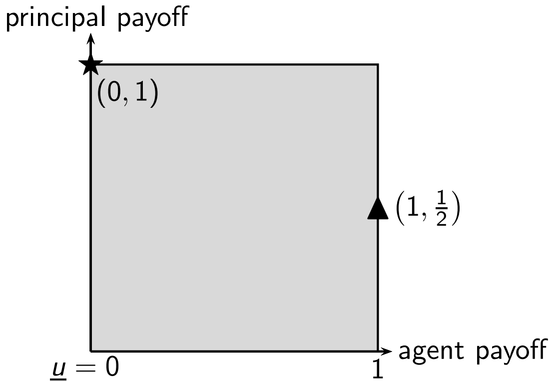

Intuition: how multiproject environment helps

- To incentive the agent to propose ★, his payoff from proposing $\{\blacktriangle, \bigstar\} \geqslant$ his payoff from proposing ▲
  - If the agent proposes ▲ alone, it is chosen with probability $1/2$

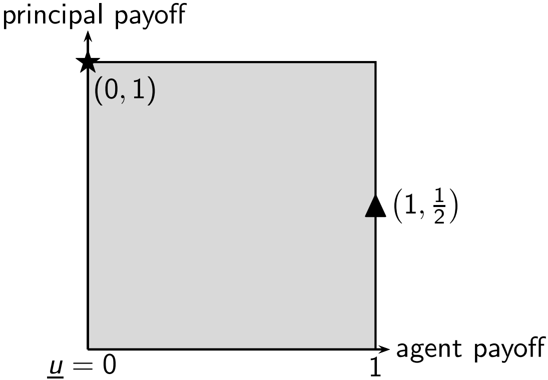

Intuition: how multiproject environment helps

- To incentive the agent to propose ★, his payoff from proposing {▲, ★} $\geqslant$ his payoff from proposing ▲
  - If the agent proposes ▲ alone, it is chosen with probability $1/2$
  - If the agent proposes {▲, ★}, each is chosen with probability $1/2$

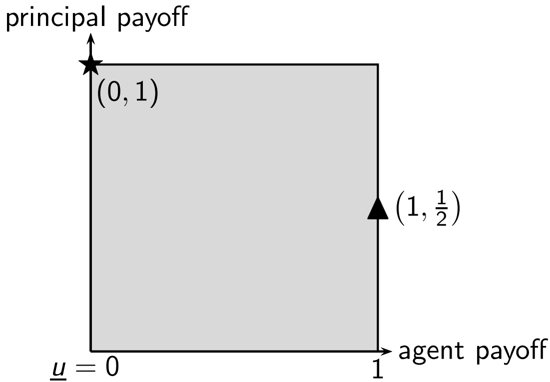

Intuition: how multiproject environment helps

- To incentive the agent to propose ★, his payoff from proposing {▲, ★} $\geqslant$ his payoff from proposing ▲
  - If the agent proposes ▲ alone, it is chosen with probability $1/2$
  - If the agent proposes {▲, ★}, each is chosen with probability $1/2$
- Regret is $1/4$ both when $A = \{▲, ★\}$ and when $A = \{▲\}$

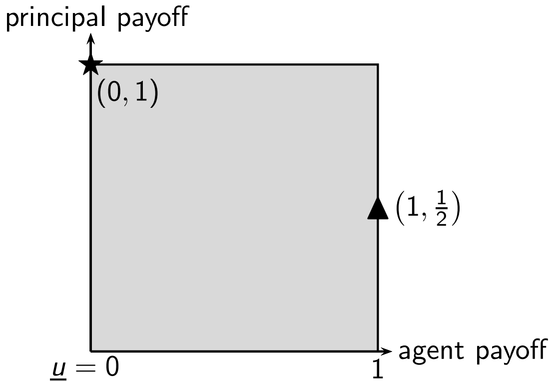

Definition of PMP mechanism

Let $\alpha : [\underline{u}, 1] \times [0, 1] \to [0, 1]$ be a function.

The **p**roposal-wide **m**aximal-**p**ayoff mechanism (PMP mechanism) induced by α is as follows:

- If the agent proposes one project (u, v), it is approved with probability $\alpha(u, v)$.
- If the agent proposes multiple projects, he is promised the maximal payoff from proposing each project alone (i.e., $\max_{(u,v)\in P} \alpha(u, v)u$).

Illustration of PMP mechanism

- If the agent proposes ▲ alone, it is chosen with probability $1/2$ and the principal's payoff is y

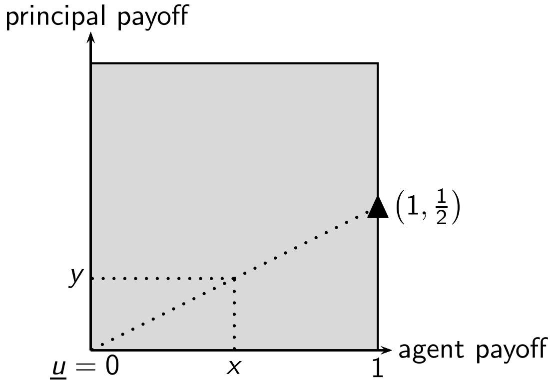

Illustration of PMP mechanism

- If the agent proposes ▲ alone, it is chosen with probability $1/2$ and the principal's payoff is y
- The multiproject environment allows the agent to also propose ★, a better fallback option than rejection. The principal's payoff goes up to z

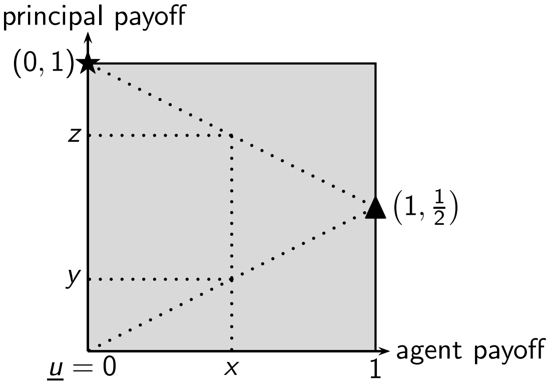

Lower bound on the worst-case regret

Theorem

For every $u \in [\underline{u}, 1]$ and $p \in [0,1]$, let $\gamma(u,p)$ be

$$\gamma(u,p) = \min\{q \in [0,1] : \underline{u} + q(u - \underline{u}) \geqslant pu\}.$$

(i) The worst-case regret under any mechanism is at least R^m:

$$R^m = \max_{(u,v)\in D} \min_{p\in[0,1]} \max\left\{(1-p)v, \gamma(u,p)(1-v)\right\}.$$

(ii) Let ρ^m be the PMP mechanism induced by

$$\alpha^m(u,v) = \max\{p \in [0,1] : \gamma(u,p)(1-v) \leqslant R^m\}.$$

Then ρ^m has the WCR of R^m.

Proof: worst-case regret under any ρ is at least R^m

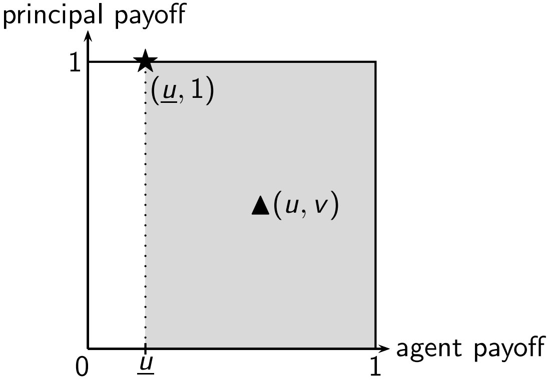

Proof: worst-case regret under any ρ is at least R^m

- Let p be the probability of approving (u, v) if it is proposed alone

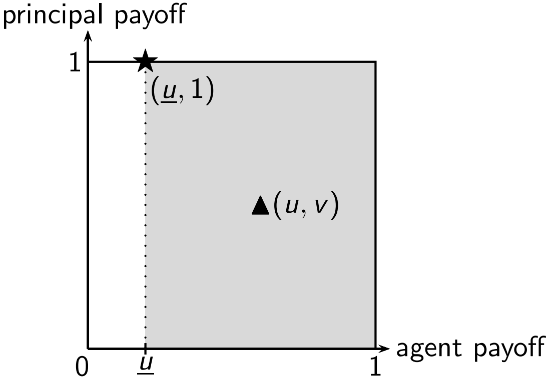

Proof: worst-case regret under any ρ is at least R^m

- Let p be the probability of approving (u, v) if it is proposed alone
- If the agent proposes $\{(u, v), (\underline{u}, 1)\}$, the probability of choosing (u, v) is at least

$$\gamma(u, p) = \min\{q \in [0, 1] : qu + (1 - q)\underline{u} \geqslant pu\}$$

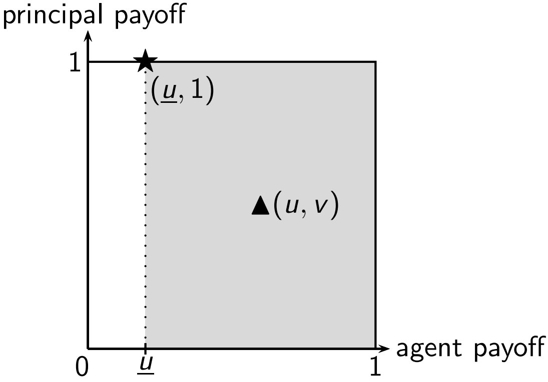

Proof: worst-case regret under any ρ is at least R^m

- Let p be the probability of approving (u, v) if it is proposed alone
- If the agent proposes $\{(u, v), (\underline{u}, 1)\}$, the probability of choosing (u, v) is at least

$$\gamma(u, p) = \min\{q \in [0, 1] : qu + (1 - q)\underline{u} \geqslant pu\}$$

$$\implies \text{worst-case regret is at least } \max\{(1 - p)v, \gamma(u, p)(1 - v)\}$$

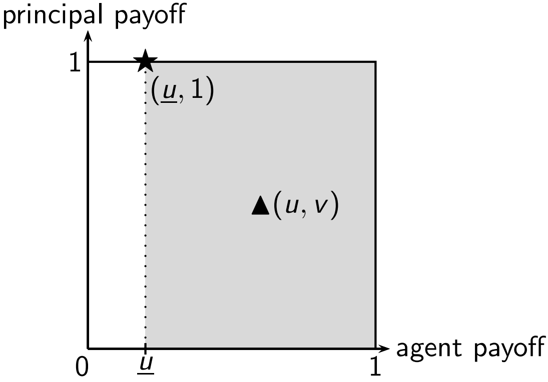

Proof: worst-case regret under any ρ is at least R^m

- Let p be the probability of approving (u, v) if it is proposed alone
- If the agent proposes $\{(u, v), (\underline{u}, 1)\}$, the probability of choosing (u, v) is at least

$$\gamma(u, p) = \min\{q \in [0, 1] : qu + (1 - q)\underline{u} \geqslant pu\}$$

$$\implies \text{worst-case regret is at least } \max\{(1 - p)v, \gamma(u, p)(1 - v)\}$$

Proof: worst-case regret under any ρ is at least R^m

- Let p be the probability of approving (u, v) if it is proposed alone
- If the agent proposes $\{(u, v), (\underline{u}, 1)\}$, the probability of choosing (u, v) is at least

$$\gamma(u, p) = \min\{q \in [0, 1] : qu + (1 - q)\underline{u} \geqslant pu\}$$

$$\implies \text{worst-case regret is at least } \max\{(1 - p)v, \gamma(u, p)(1 - v)\}$$

- There are two more operators in R^m:

$$R^m = \max_{(u,v)\in D} \min_{p\in[0,1]} \max\{(1 - p)v, \gamma(u, p)(1 - v)\}$$

Proof: worst-case regret under any ρ is at least R^m

- Let p be the probability of approving (u, v) if it is proposed alone
- If the agent proposes $\{(u, v), (\underline{u}, 1)\}$, the probability of choosing (u, v) is at least

$$\gamma(u,p) = \min\{q \in [0,1] : qu + (1-q)\underline{u} \geqslant pu\}$$
$$\implies \text{worst-case regret is at least } \max\{(1-p)v, \gamma(u,p)(1-v)\}$$

- There are two more operators in R^m:

$$R^m = \max_{(u,v)\in D} \min_{p\in[0,1]} \max\{(1-p)v, \gamma(u,p)(1-v)\}$$

  - $\min_{p\in[0,1]}$ reflects the fact that the principal can choose p

Proof: worst-case regret under any ρ is at least R^m

- Let p be the probability of approving (u, v) if it is proposed alone
- If the agent proposes $\{(u, v), (\underline{u}, 1)\}$, the probability of choosing (u, v) is at least

$$\gamma(u, p) = \min\{q \in [0, 1] : qu + (1 - q)\underline{u} \geqslant pu\}$$
$$\implies \text{worst-case regret is at least } \max\{(1 - p)v, \gamma(u, p)(1 - v)\}$$

- There are two more operators in R^m:

$$R^m = \max_{(u,v)\in D} \min_{p\in[0,1]} \max\{(1 - p)v, \gamma(u, p)(1 - v)\}$$

  - $\min_{p\in[0,1]}$ reflects the fact that the principal can choose p
  - $\max_{(u,v)\in D}$ reflects the fact that the argument holds for any (u, v)

Optimal mechanism

Theorem

For every $u \in [\underline{u}, 1]$ and $p \in [0,1]$, let $\gamma(u,p)$ be

$$\gamma(u,p) = \min\{q \in [0,1] : \underline{u} + q(u - \underline{u}) \geqslant pu\}.$$

(i) The worst-case regret under any mechanism is at least R^m:

$$R^m = \max_{(u,v)\in D} \min_{p\in[0,1]} \max\left\{(1-p)v, \gamma(u,p)(1-v)\right\}.$$

(ii) Let ρ^m be the PMP mechanism induced by

$$\alpha^m(u,v) = \max\{p \in [0,1] : \gamma(u,p)(1-v) \leqslant R^m\}.$$

Then ρ^m has the WCR of R^m.

Intuition for the definition of α^m

$$\alpha^m(u,v) = \max\{p \in [0,1] : \gamma(u,p)(1-v) \leqslant R^m\}$$

- $\alpha^m(u,v)$ is the highest probability p of approving (u,v) if it is proposed alone such that:

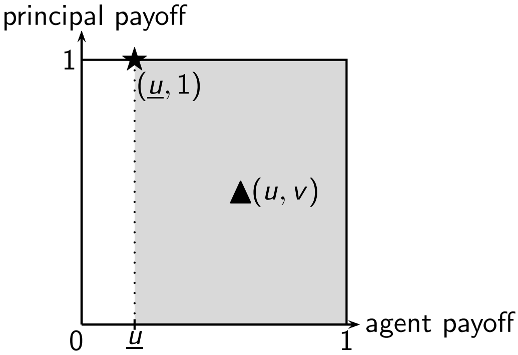

Intuition for the definition of α^m

$$\alpha^m(u, v) = \max\{p \in [0, 1] : \gamma(u, p)(1 - v) \leqslant R^m\}$$

- $\alpha^m(u, v)$ is the highest probability p of approving (u, v) if it is proposed alone such that:
  - the agent can get $\alpha^m(u, v)u$ when he proposes $\{(u, v), (\underline{u}, 1)\}$

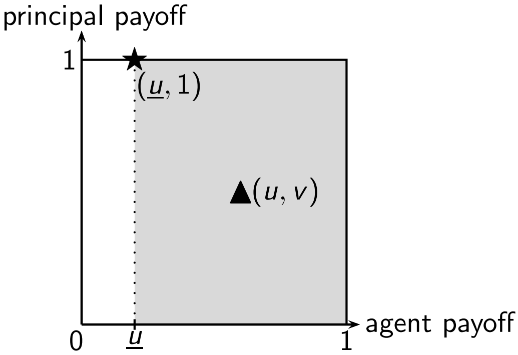

Intuition for the definition of α^m

$$\alpha^m(u,v) = \max\{p \in [0,1] : \gamma(u,p)(1-v) \leqslant R^m\}$$

- $\alpha^m(u,v)$ is the highest probability p of approving (u,v) if it is proposed alone such that:
  - the agent can get $\alpha^m(u,v)u$ when he proposes $\{(u,v),(\underline{u},1)\}$
  - and the regret from $\{(u,v),(\underline{u},1)\}$ is at most R^m

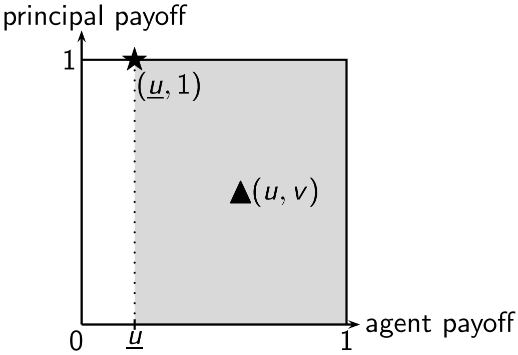

Two tier if the agent proposes one project

- The explicit expression for $\alpha^m(u, v)$ is:

$$\alpha^m(u,v) = \begin{cases} 1, & \text{if } v \geqslant 1 - R^m \text{ or } u = \underline{u}; \\ \left(1 - \frac{R^m}{1-v}\right)\underline{u}/u + \frac{R^m}{1-v}, & \text{if } v < 1 - R^m \text{ and } u > \underline{u}. \end{cases}$$

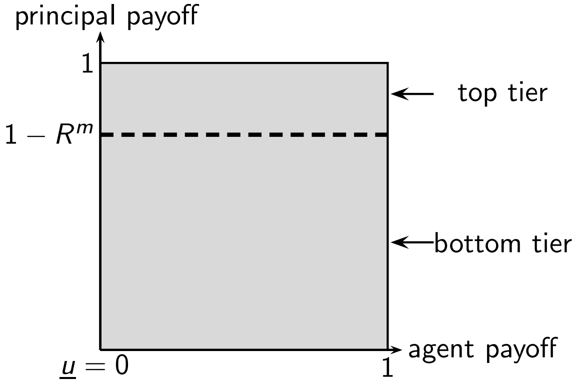

When do several projects have a chance of being chosen

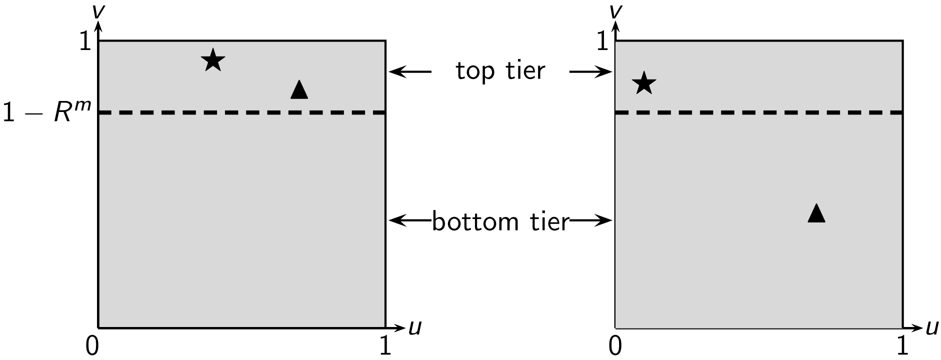

agent's favorite project
is in top tier

agent's favorite project
is in bottom tier

Worst-case regret: single-project vs multiproject

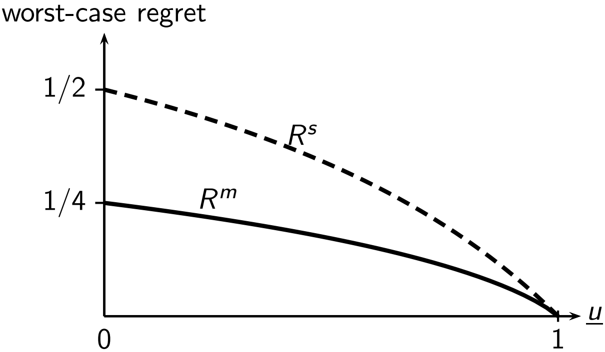

Roadmap

- Model
- Single-project environment
- Multiproject environment
- Discussion

Intermediate environments

- The agent can propose up to K projects

Proposition (Two is enough in minimizing worst-case regret)

For any $K \geqslant 2$,

Intermediate environments

- The agent can propose up to K projects

Proposition (Two is enough in minimizing worst-case regret)

For any $K \geqslant 2$,

1. the WCR under any mechanism is at least R^m;

Intermediate environments

- The agent can propose up to K projects

Proposition (Two is enough in minimizing worst-case regret)

For any $K \geqslant 2$,

1. the WCR under any mechanism is at least R^m;
2. under the PMP mechanism ρ^m induced by $\alpha^m(u, v)$, it is optimal for the agent to propose
   - the project that maximizes $\alpha^m(u, v)u$
   - and the principal's favorite project;

   the corresponding choice function has the WCR of R^m.

Conclusion

- We study prior-free mechanisms that counteract the agent's proposing bias

Conclusion

- We study prior-free mechanisms that counteract the agent's proposing bias
- If the agent can propose only a single project, the optimal mechanism has a two-tier structure

Conclusion

- We study prior-free mechanisms that counteract the agent's proposing bias
- If the agent can propose only a single project, the optimal mechanism has a two-tier structure
  - a top-tier project is approved with no questions asked

Conclusion

- We study prior-free mechanisms that counteract the agent's proposing bias
- If the agent can propose only a single project, the optimal mechanism has a two-tier structure
  - a top-tier project is approved with no questions asked
  - a bottom-tier project will be scrutinized

Conclusion

- We study prior-free mechanisms that counteract the agent's proposing bias
- If the agent can propose only a single project, the optimal mechanism has a two-tier structure
  - a top-tier project is approved with no questions asked
  - a bottom-tier project will be scrutinized
  - the more a bottom-tier project benefits the agent, the less likely it will be approved

Conclusion

- We study prior-free mechanisms that counteract the agent's proposing bias
- If the agent can propose only a single project, the optimal mechanism has a two-tier structure
  - a top-tier project is approved with no questions asked
  - a bottom-tier project will be scrutinized
  - the more a bottom-tier project benefits the agent, the less likely it will be approved
- If the agent can propose multiple projects,

Conclusion

- We study prior-free mechanisms that counteract the agent's proposing bias
- If the agent can propose only a single project, the optimal mechanism has a two-tier structure
  - a top-tier project is approved with no questions asked
  - a bottom-tier project will be scrutinized
  - the more a bottom-tier project benefits the agent, the less likely it will be approved
- If the agent can propose multiple projects,
  - he can provide the principal with better fallback options than rejection

Conclusion

- We study prior-free mechanisms that counteract the agent's proposing bias
- If the agent can propose only a single project, the optimal mechanism has a two-tier structure
  - a top-tier project is approved with no questions asked
  - a bottom-tier project will be scrutinized
  - the more a bottom-tier project benefits the agent, the less likely it will be approved
- If the agent can propose multiple projects,
  - he can provide the principal with better fallback options than rejection
  - the principal does not always choose his favorite project among proposed ones

Questions for future research

- The principal may not observe the agent's payoffs from proposed projects

Questions for future research

- The principal may not observe the agent's payoffs from proposed projects
- The project-choice process may involve a deeper hierarchy involving three or more parties

Questions for future research

- The principal may not observe the agent's payoffs from proposed projects
- The project-choice process may involve a deeper hierarchy involving three or more parties
- The principal may be able to acquire information about the agent's type at a cost

Questions for future research

- The principal may not observe the agent's payoffs from proposed projects
- The project-choice process may involve a deeper hierarchy involving three or more parties
- The principal may be able to acquire information about the agent's type at a cost
- The agent may need to exert effort to discover projects

Thank you!