Strategic Attribute Learning

Jean-Michel Benkert [1] Ludmila Matysková [2] Egor Starkov [3]

[1]University of Bern

[2]University of Alicante

[3]University of Copenhagen

EEA-ESEM, Aug 2024

Delegated learning in complex environments

- Decision-maker and researcher
- One decision to be made regarding an uncertain project
- Project characterized by finitely many independent attributes
- Players disagree on the importance of the attributes
- Agent learns by allocating limited resources across attributes

- **Broad question:** how does misalignment affect learning?

Examples

1. Politician decides on a policy, but depends on advisor
   - the impact varies across social groups
   - politician prioritizes some social groups

Examples

1. Politician decides on a policy, but depends on advisor
   - the impact varies across social groups
   - politician prioritizes some social groups
2. Manager chooses firm strategy, but depends on analyst
   - they may disagree on the importance of different factors (e.g., regulatory vs competitive environment) for final strategy

Examples

1. Politician decides on a policy, but depends on advisor
   - the impact varies across social groups
   - politician prioritizes some social groups
2. Manager chooses firm strategy, but depends on analyst
   - they may disagree on the importance of different factors (e.g., regulatory vs competitive environment) for final strategy
3. Voter influenced by media
   - media may want the voter to pay more or less atention to certain issues

Preview of Results – Main results

- We characterize the equilibrium learning behavior.
- We show that it coincides with the solution of a modified single-player problem.
- We provide conditions under which the researcher abstains from (free) learning.
- We prove the equivalence to similar yet economically distinct frameworks.

Preview of Results – Applications 1

Organization

- Does a manger want to employ like-minded individuals?
- Preference for diversity: employer may prefer a more biased employee

Preview of Results – Applications 1

Organization

- Does a manger want to employ like-minded individuals?
- Preference for diversity: employer may prefer a more biased employee

Discrimination

- A discriminating policymaker can by tamed by detaching learning from decisions
- Impartial advisor maximizes welfare and mitigates inequality
- Eliminating inequality requires a counter-biased advisor

Preview of Results – Applications 2

Media polarization

- Competition between media outlets can lead to polarization...
- ...which is beneficial for the voter

Preview of Results – Applications 2

Media polarization

- Competition between media outlets can lead to polarization...
- ...which is beneficial for the voter

Dynamic preferences

- How does time-inconsistency affect learning and welfare?
- A sophisticated agent may engage in strategic ignorance
- Naivete can be beneficial

Outline

Framework

- Researcher learns, Decision-maker makes a decision

Framework

- Researcher learns, Decision-maker makes a decision
- K attributes $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_K)' \in \mathbb{R}^K$. This talk: $K = 2$.

Framework

- Researcher learns, Decision-maker makes a decision
- K attributes $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_K)' \in \mathbb{R}^K$. This talk: $K = 2$.
- Bliss points:

$$\begin{aligned} v^R(\boldsymbol{\theta}) &= \alpha_1^R \theta_1 + \ldots + \alpha_K^R \theta_K \\ v^{DM}(\boldsymbol{\theta}) &= \alpha_1^{DM} \theta_1 + \ldots + \alpha_K^{DM} \theta_K \end{aligned}$$

where weights $\alpha_k^i \geq 0$ are commonly known

Framework

- Researcher learns, Decision-maker makes a decision
- K attributes $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_K)' \in \mathbb{R}^K$. This talk: $K = 2$.
- Bliss points:

$$\begin{aligned} v^R(\boldsymbol{\theta}) &= \alpha_1^R \theta_1 + \ldots + \alpha_K^R \theta_K \\ v^{DM}(\boldsymbol{\theta}) &= \alpha_1^{DM} \theta_1 + \ldots + \alpha_K^{DM} \theta_K \end{aligned}$$

  where weights $\alpha_k^i \geq 0$ are commonly known

- DM's decision is $d \in \mathbb{R}$

Framework

- Researcher learns, Decision-maker makes a decision
- K attributes $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_K)' \in \mathbb{R}^K$. This talk: $K = 2$.
- Bliss points:

$$\begin{aligned} v^R(\boldsymbol{\theta}) &= \alpha_1^R \theta_1 + \ldots + \alpha_K^R \theta_K \\ v^{DM}(\boldsymbol{\theta}) &= \alpha_1^{DM} \theta_1 + \ldots + \alpha_K^{DM} \theta_K \end{aligned}$$

  where weights $\alpha_k^i \geq 0$ are commonly known
- DM's decision is $d \in \mathbb{R}$
- Utilities: $u^i(d, \boldsymbol{\theta}) = -(d - v^i(\boldsymbol{\theta}))^2$

Attributes and learning

Attributes are independent multivariate normal:

$$\tilde{\theta}_k \sim \mathcal{N}(\mu_k^0, \Sigma_k^0)$$

Attributes and learning

Attributes are independent multivariate normal:

$$\tilde{\theta}_k \sim \mathcal{N}(\mu_k^0, \Sigma_k^0)$$

Researcher has budget of tests $T > 0$:

Attributes and learning

Attributes are independent multivariate normal:

$$\tilde{\theta}_k \sim \mathcal{N}(\mu_k^0, \Sigma_k^0)$$

Researcher has budget of tests $T > 0$:

- chooses test allocation $\tau_1, \ldots, \tau_K \geq 0$ such that $\sum_k \tau_k \leq T$ (free disposal)
- Allocate τ_k tests to attribute $\theta_k \implies$ generate signal with precision τ_k:

$$\tilde{s}_k = \theta_k + \mathcal{N}\left(0, \frac{1}{\tau_k}\right)$$

Timing

1. R chooses test allocation $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_K)'$ s.t. $\sum_k \tau_k \leq T$ (observable)
2. Signal realizations $s_k = \theta_k + \varepsilon_k$ publicly observed
3. DM updates beliefs
4. DM chooses the decision
5. Payoffs are realized

Timing

1. R chooses test allocation $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_K)'$ s.t. $\sum_k \tau_k \leq T$ (observable)
2. Signal realizations $s_k = \theta_k + \varepsilon_k$ publicly observed
3. DM updates beliefs
4. DM chooses the decision
5. Payoffs are realized

Assumption: when R is indifferent, doesn't test.

Equilibrium concept: weak PBE

Related Literature

Strategic attribute selection

- Bardhi (2024), *Econometrica*
  - Different: independence, finite attributes, noisy signals

(Dynamic) non-strategic attribute learning

- Liang, Mu, and Syrgkanis (2022), *Econometrica*
  - Different: strategic framework, independence, specific utilities

Outline

Single player

Start by solving a single-player problem:

$$\alpha^{DM} = \alpha^{R} \equiv \alpha.$$

Single player

Start by solving a single-player problem:

$$\alpha^{DM} = \alpha^{R} \equiv \alpha.$$

- Decision is trivial: $d = \mathbb{E}[v|s,\tau]$
- At the learning stage, the goal is to minimize residual variance $\hat{\Sigma}_k := \mathbb{V}[v|s,\tau]$

Theorem 1: Optimal test allocation ($K = 2$)

WLOG let $\alpha_1 \Sigma_1^0 \geq \alpha_2 \Sigma_2^0$.

Theorem 1: Optimal test allocation ($K = 2$)

WLOG let $\alpha_1 \Sigma_1^0 \geq \alpha_2 \Sigma_2^0$. Define

$$\bar{T} = \frac{\alpha_1 \Sigma_1^0 - \alpha_2 \Sigma_2^0}{\alpha_2 \Sigma_1^0 \Sigma_2^0}.$$

Theorem 1: Optimal test allocation ($K = 2$)

WLOG let $\alpha_1 \Sigma_1^0 \geq \alpha_2 \Sigma_2^0$. Define

$$\bar{T} = \frac{\alpha_1 \Sigma_1^0 - \alpha_2 \Sigma_2^0}{\alpha_2 \Sigma_1^0 \Sigma_2^0}.$$

The single-player optimal test allocation strategy is as follows:

Theorem 1: Optimal test allocation ($K = 2$)

WLOG let $\alpha_1 \Sigma_1^0 \geq \alpha_2 \Sigma_2^0$. Define

$$\bar{T} = \frac{\alpha_1 \Sigma_1^0 - \alpha_2 \Sigma_2^0}{\alpha_2 \Sigma_1^0 \Sigma_2^0}.$$

The single-player optimal test allocation strategy is as follows:

1. Allocate $\min\{T, \bar{T}\}$ tests to attribute 1

Theorem 1: Optimal test allocation ($K = 2$)

WLOG let $\alpha_1 \Sigma_1^0 \geq \alpha_2 \Sigma_2^0$. Define

$$\bar{T} = \frac{\alpha_1 \Sigma_1^0 - \alpha_2 \Sigma_2^0}{\alpha_2 \Sigma_1^0 \Sigma_2^0}.$$

The single-player optimal test allocation strategy is as follows:

1. Allocate $\min\{T, \bar{T}\}$ tests to attribute 1
2. Allocate the remaining $T - \bar{T}$ tests (if any) in constant fraction

$$\left(\underbrace{\frac{\alpha_1}{\alpha_1 + \alpha_2}}_{\text{to attribute 1}}, \underbrace{\frac{\alpha_2}{\alpha_1 + \alpha_2}}_{\text{to attribute 2}} \right) \cdot (T - \bar{T})$$

Theorem 1: Optimal test allocation ($K = 2$)

WLOG let $\alpha_1 \Sigma_1^0 \geq \alpha_2 \Sigma_2^0$. Define

$$\bar{T} = \frac{\alpha_1 \Sigma_1^0 - \alpha_2 \Sigma_2^0}{\alpha_2 \Sigma_1^0 \Sigma_2^0}.$$

The single-player optimal test allocation strategy is as follows:

1. Allocate $\min\{T, \bar{T}\}$ tests to attribute 1
2. Allocate the remaining $T - \bar{T}$ tests (if any) in constant fraction

$$\Bigg(\underbrace{\frac{\alpha_1}{\alpha_1 + \alpha_2}}_{\text{to attribute 1}}, \underbrace{\frac{\alpha_2}{\alpha_1 + \alpha_2}}_{\text{to attribute 2}}\Bigg) \cdot (T - \bar{T})$$

- Step 1 equalizes weighted residual variances $\alpha_k \hat{\Sigma}_k$ example
- Optimal test allocation depends on weight ratio $\frac{\alpha_1}{\alpha_2}$

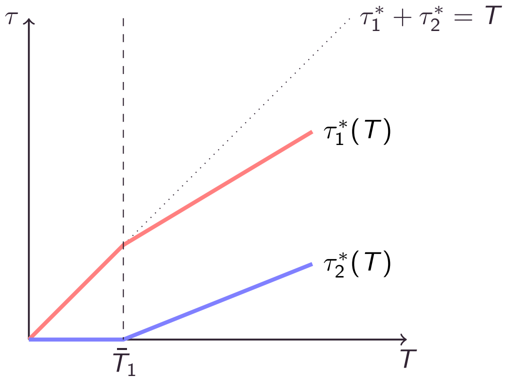
τ
$\tau_1^* + \tau_2^* = T$
$\tau_1^*(T)$
$\tau_2^*(T)$
$\bar{T}_1$
T

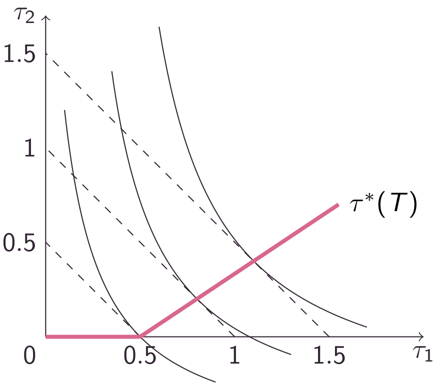
τ_2
1.5
1
0.5
0
0.5
1
1.5
τ_1
$\tau^*(T)$

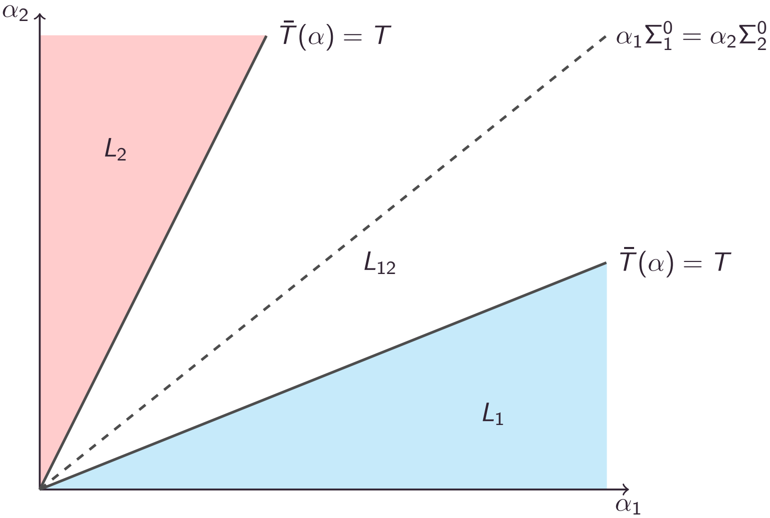
α_2
$\bar{T}(\alpha) = T$
$\alpha_1 \Sigma_1^0 = \alpha_2 \Sigma_2^0$
L_2
L_{12}
$\bar{T}(\alpha) = T$
L_1
α_1

Outline

Strategic problem

Researcher's expected utility

$$V^R(\tau) \equiv \mathbb{E}\left[-\left(\tilde{d}^{DM} - \tilde{v}^R\right)^2\right]$$

Strategic problem

Researcher's expected utility

$$V^R(\tau) \equiv \mathbb{E}\left[-\left(\tilde{d}^{DM} - \tilde{v}^R\right)^2\right] \quad = \quad -(v_0^{DM} - v_0^R)^2 - \mathbb{V}\left(\tilde{d}^{DM} - \tilde{v}^R\right)$$

Strategic problem

Researcher's expected utility

$$V^R(\tau) \equiv \mathbb{E}\left[-\left(\tilde{d}^{DM} - \tilde{v}^R\right)^2\right] = -(v_0^{DM} - v_0^R)^2 - \mathbb{V}\left(\tilde{d}^{DM} - \tilde{v}^R\right)$$

$$= V^A(\emptyset) + \underbrace{2cov[\tilde{v}^R, \tilde{d}^{DM}] - \mathbb{V}[\tilde{d}^{DM}]}_{\text{added value of learning}}$$

Strategic problem

Researcher's expected utility

$$V^R(\tau) \equiv \mathbb{E}\left[-\left(\tilde{d}^{DM} - \tilde{v}^R\right)^2\right] = -(v_0^{DM} - v_0^R)^2 - \mathbb{V}\left(\tilde{d}^{DM} - \tilde{v}^R\right)$$

$$= V^A(\emptyset) + \underbrace{2cov[\tilde{v}^R, \tilde{d}^{DM}] - \mathbb{V}[\tilde{d}^{DM}]}_{\text{added value of learning}}$$

Researcher considers two factors:

1. aligning DM's decision with R's objective: $\uparrow cov[\tilde{v}^R, \tilde{d}^{DM}]$
2. reducing excess variance of DM's decision: $\downarrow \left(\mathbb{V}[\tilde{d}^{DM}] - cov(\tilde{v}^R, \tilde{d}^{DM})\right)$

Main result

Theorem 2: Equilibrium Test Allocation

τ is an equilibrium test allocation iff τ solves the single-player problem with effective weights

$$\tilde{\alpha}_k := \sqrt{\max\left\{0, \left(\alpha_k^R\right)^2 - \left(\alpha_k^R - \alpha_k^{DM}\right)^2\right\}}$$

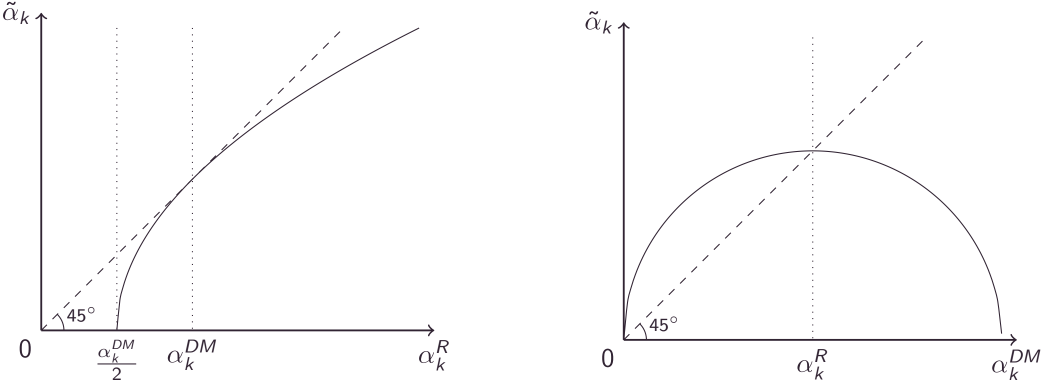

Effective weights $\tilde{\alpha}_k$ as functions of α_k^R (left) and α_k^{DM} (right)

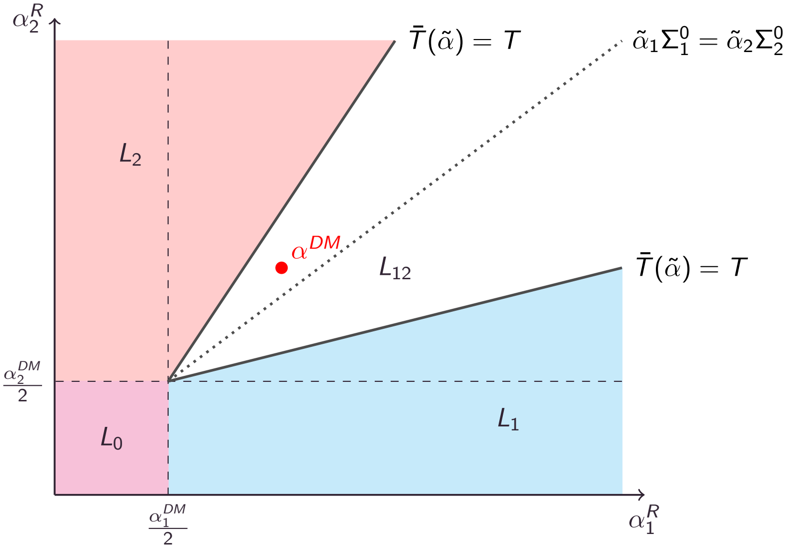

more on optimal misalignment

Takeaways

1. Tractability:
   equilibirum test allocation given by single-agent solution with distorted weights.

Takeaways

1. Tractability:
   equilibirum test allocation given by single-agent solution with distorted weights.
2. Monotonicity:
   higher $\alpha_k^R \Rightarrow$ higher effective weight $\tilde{\alpha}_k$.

Takeaways

1. Tractability:
   equilibirum test allocation given by single-agent solution with distorted weights.
2. Monotonicity:
   higher $\alpha_k^R \Rightarrow$ higher effective weight $\tilde{\alpha}_k$.
3. Ignorance:
   if α_k^{DM} too high compared to α_k^R, R thinks DM overreacts and refuses to learn.

Takeaways

1. Tractability:
   equilibirum test allocation given by single-agent solution with distorted weights.
2. Monotonicity:
   higher $\alpha_k^R \Rightarrow$ higher effective weight $\tilde{\alpha}_k$.
3. Ignorance:
   if α_k^{DM} too high compared to α_k^R, R thinks DM overreacts and refuses to learn.
4. Misalignment:
   if R is not sensitive enough (small α^R), DM may benefit from distorting R's weight ratio $\frac{\alpha_1^R}{\alpha_2^R}$ relative to her own.

Equivalent models

Framework A

DM takes a single decision, $d \in \mathbb{R}$; players' payoffs are

$$u_A^i(d, \theta) = -\sum_k \alpha_k^i \left(d - \theta_k\right)^2 ,$$

where $\sum_k \alpha_k^i = 1$ for both players.

Equivalent models

Framework A

DM takes a single decision, $d \in \mathbb{R}$; players' payoffs are

$$u_A^i(d,\theta) = -\sum_k \alpha_k^i \left(d - \theta_k\right)^2,$$

where $\sum_k \alpha_k^i = 1$ for both players.

Framework B

DM takes K different decisions $d_1, \ldots, d_K \in \mathbb{R}$; players' payoffs are

$$u_B^i(d,\theta) = -\sum_k (d_k - \alpha_k^i \theta_k)^2.$$

The equilibrium test allocation τ in these coincides with our baseline model.

Outline

Discrimination and inequality – Model 1/2

- Politician, advisor, two social groups
- Politician decides on policy $d \in \mathbb{R}$
- Two social groups $k \in \{1, 2\}$
  - Unknown optimal policy $\theta_k \in \mathbb{R}$
  - "Group" payoff: $u_k(d, \theta_k) = -(d - \theta_k)^2$

Discrimination and inequality – Model 2/2

- Politician decides on policy $d \in \mathbb{R}$, has payoff

$$u^{DM}(d, \theta_1, \theta_2; \delta) = \underbrace{\frac{1}{2}(1+\delta)}_{\alpha_1^{DM}(\delta)} u_1(d, \theta_1) + \underbrace{\frac{1}{2}(1-\delta)}_{\alpha_2^{DM}(\delta)} u_2(d, \theta_2)$$

where $\delta \in (0, 1)$ is the discrimination parameter.

Discrimination and inequality – Model 2/2

- Politician decides on policy $d \in \mathbb{R}$, has payoff

$$u^{DM}(d, \theta_1, \theta_2; \delta) = \underbrace{\frac{1}{2}(1+\delta)}_{\alpha_1^{DM}(\delta)} u_1(d, \theta_1) + \underbrace{\frac{1}{2}(1-\delta)}_{\alpha_2^{DM}(\delta)} u_2(d, \theta_2)$$

where $\delta \in (0, 1)$ is the discrimination parameter.

- Advisor chooses test allocation τ, has payoff

$$u^{R}(d, \theta_1, \theta_2; p) = \underbrace{\frac{1}{2}(1-p)}_{\alpha_1^{R}(p)} u_1(d, \theta_1) + \underbrace{\frac{1}{2}(1+p)}_{\alpha_2^{R}(p)} u_2(d, \theta_2)$$

where $p \in [0, 1)$ is the partiality parameter.

Discrimination and inequality – Impartial Advisor

Proposition (Impartial advisor)

When the advisor is impartial ($p = 0$), the equilibrium resource allocation is $\tau^* = \left(\frac{1}{2}, \frac{1}{2}\right)$ for any discrimination level δ.

Discrimination and inequality – Impartial Advisor

Proposition (Impartial advisor)

When the advisor is impartial ($p = 0$), the equilibrium resource allocation is $\tau^* = \left(\frac{1}{2}, \frac{1}{2}\right)$ for any discrimination level δ.

Intuition: $\uparrow$ discrimination δ has two opposing effects

Discrimination and inequality – Impartial Advisor

Proposition (Impartial advisor)

When the advisor is impartial ($p = 0$), the equilibrium resource allocation is $\tau^* = \left(\frac{1}{2}, \frac{1}{2}\right)$ for any discrimination level δ.

Intuition: $\uparrow$ discrimination δ has two opposing effects

- larger discrepancy between α^{DM} and α^{R}
  $\Rightarrow$ advisor wants to learn more about group 2

Discrimination and inequality – Impartial Advisor

Proposition (Impartial advisor)

When the advisor is impartial ($p = 0$), the equilibrium resource allocation is $\tau^* = \left(\frac{1}{2}, \frac{1}{2}\right)$ for any discrimination level δ.

Intuition: $\uparrow$ discrimination δ has two opposing effects

- larger discrepancy between α^{DM} and α^{R}
  $\Rightarrow$ advisor wants to learn more about group 2
- politician less sensitive to information about group 2
  $\Rightarrow$ advisor wants to learn less about group 2

Discrimination and inequality – Impartial Advisor

Proposition (Impartial advisor)

When the advisor is impartial ($p = 0$), the equilibrium resource allocation is $\tau^* = \left(\frac{1}{2}, \frac{1}{2}\right)$ for any discrimination level δ.

Intuition: $\uparrow$ discrimination δ has two opposing effects

- larger discrepancy between α^{DM} and α^{R}
  $\Rightarrow$ advisor wants to learn more about group 2
- politician less sensitive to information about group 2
  $\Rightarrow$ advisor wants to learn less about group 2
- impartial advisor: the two effects cancel out

Discrimination and inequality – Eliminating Inequality

Let $\mu^0 = (0, 0)$, $\Sigma_k^0 = 1$, $T = 1$.

Proposition (Equality)

For every $\delta \in (0, 1)$, there exists a unique $\hat{p}(\delta) > 0$ such that $u_1 = u_2$ in equilibrium. Further, $\hat{p}(\delta)$ is continuous and non-monotone in δ.

Discrimination and inequality – Eliminating Inequality

Let $\mu^0 = (0,0)$, $\Sigma_k^0 = 1$, $T = 1$.

Proposition (Equality)

For every $\delta \in (0,1)$, there exists a unique $\hat{p}(\delta) > 0$ such that $u_1 = u_2$ in equilibrium. Further, $\hat{p}(\delta)$ is continuous and non-monotone in δ.

- As $\uparrow \delta$, an advisor partial towards group 2 compensates by learning more about group 2.
- When $\delta \to 1$, even an impartial advisor chooses $\tau_1 \approx 0$ and "punishes" the favored group by effectively providing no information.
- But welfare (with equal group weights) is maximized with an impartial advisor, $p = 0$.

Media Polarization – Model

- Two media outlets $m = A, B$ aim to influence a decision of (median) voter V.
- The voter's decision d affects two policy issues $k = 1, 2$; each player wants the decision to reflect their preferred platform.
- The utility of player $i = A, B, V$ is (with $\alpha_1^i + \alpha_2^i = 1$)

$$u^i(d, \theta) = -\sum_{k=1}^{2} \alpha_k^i \left(d - \theta_k\right)^2 . \tag{1}$$

Media Polarization – Model

- Two media outlets $m = A, B$ aim to influence a decision of (median) voter V.
- The voter's decision d affects two policy issues $k = 1, 2$; each player wants the decision to reflect their preferred platform.
- The utility of player $i = A, B, V$ is (with $\alpha_1^i + \alpha_2^i = 1$)

$$u^i(d, \theta) = -\sum_{k=1}^{2} \alpha_k^i \left(d - \theta_k\right)^2 . \tag{1}$$

- Each media outlet chooses coverage $q^m = (q_1^m, q_2^m)' \in \mathbb{R}_+^2$ such that $q_1^m + q_2^m = 1$.
- The voter has a total budget of attention $T > 0$ and chooses how to allocate it between the two outlets: $t = (t^A, t^B)' \in \mathbb{R}_+^2$ such that $t^A + t^B \leq T$.
- Given q^m and t, the voter observes signals about each issue with precisions $t_k q_k^M$.

Media Polarization – Timing

1. The voter chooses her attention allocation t.
2. Media outlets observe t and simultaneously choose coverage q^m.
3. Nature draws state realizations θ and signal realizations s
4. The voter observes q^m and s, updates her beliefs and takes decision d.
5. Payoffs are realized.

Media Polarization – Results

Proposition

If $\alpha_1^A > \alpha_1^V > \alpha_1^B$ and T is large enough, then in the unique equilibrium, media is polarized, $q^A = (1,0)'$ and $q^B = (0,1)'$, but the voter achieves her optimal aggregate attention allocation τ_V^*.

Media Polarization – Results

Proposition

If $\alpha_1^A > \alpha_1^V > \alpha_1^B$ and T is large enough, then in the unique equilibrium, media is polarized, $q^A = (1,0)'$ and $q^B = (0,1)'$, but the voter achieves her optimal aggregate attention allocation τ_V^*.

Corollary

If $\alpha_1^A > \alpha_1^V > \alpha_1^B$ and T is large enough, then the voter strictly prefers a polarized media duopoly to a (moderate) monopoly.

Dynamic Preferences – Model 1/2

- A single agent who can acquire information before taking a decision.
- The agent's preferences can change in between.
- The agent can be naive or sophisticated about the potential change.
- Assume $\alpha_1^{DM} = \alpha_1^R$ and

$$\alpha_2^{DM} = \begin{cases} \alpha_2^R & \text{with probability } p, \\ c\alpha_2^R & \text{with probability } 1-p, \end{cases}$$

for some $c \geq 0$ with $c \neq 1$.

Dynamic Preferences – Model 2/2

- Thus, learning utility reads

$$u^R(d,\theta) = -\left(d - \alpha_1^R\theta_1 - \alpha_2^R\theta_2\right)^2,$$

and decision utility reads

$$u^{DM}(d,\theta) = \begin{cases} -\left(d - \alpha_1^R\theta_1 - \alpha_2^R\theta_2\right)^2 & \text{with probability } p, \\ -\left(d - \alpha_1^R\theta_1 - c\alpha_2^R\theta_2\right)^2 & \text{with probability } 1-p. \end{cases}$$

Dynamic Preferences – Results

Proposition

The sophisticate ignores attribute 2 for all $T > 0$ if and only if

$$c - 1 \geq \frac{1}{\sqrt{1 - p}}.$$

Dynamic Preferences – Results

Proposition

The sophisticate ignores attribute 2 for all $T > 0$ if and only if

$$c - 1 \geq \frac{1}{\sqrt{1-p}}.$$

Proposition (informal)

Depending on parameters and the choice of the welfare criterion, either the naif or the sophisticate may be better off.

Conclusion

- We contribute to the literature on strategic multi-attribute learning problems (Bardhi, 2024) by providing a tractable model to analyze preference misalignment.
- More broadly, we contribute to the literature on delegated expertise, which is typically considering only single-dimensional problems.
- In the context of strategic communication, our approach corresponds to a form of "constrained Bayesian persuasion".
- Our model builds on Liang et al. (2022), who show that a "greedy" learning strategy is optimal with correlated attributes while we show it is optimal with strategic motives.

Several applications

1. Diversity in Organizations
2. Discrimination and Inequality
3. Media Polarization
4. Dynamic Preferences

Example 1

$$\begin{pmatrix}\tilde{\theta}_1\\ \tilde{\theta}_2\end{pmatrix} \sim N\left(\begin{pmatrix}\mu_1^0\\ \mu_2^0\end{pmatrix}, \begin{pmatrix}8 & 0\\ 0 & 1\end{pmatrix}\right)$$

Bliss action is $v = \theta_1 + \theta_2$ (i.e., $\alpha_1 = \alpha_2 = 1$)

Test budget $T = 2$

Example 1

$$\begin{pmatrix}\tilde{\theta}_1\\ \tilde{\theta}_2\end{pmatrix} \sim N\left(\begin{pmatrix}\mu_1^0\\ \mu_2^0\end{pmatrix}, \begin{pmatrix}8 & 0\\ 0 & 1\end{pmatrix}\right)$$

Bliss action is $v = \theta_1 + \theta_2$ (i.e., $\alpha_1 = \alpha_2 = 1$)

Test budget $T = 2$

Optimal test allocation:

- allocate $\bar{T} = \frac{7}{8}$ tests to attribute 1 $\Rightarrow$ the posterior covariance matrix is $\begin{pmatrix}1 & 0\\ 0 & 1\end{pmatrix}$

Example 1

$$\begin{pmatrix} \tilde{\theta}_1 \\ \tilde{\theta}_2 \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_1^0 \\ \mu_2^0 \end{pmatrix}, \begin{pmatrix} 8 & 0 \\ 0 & 1 \end{pmatrix} \right)$$

Bliss action is $v = \theta_1 + \theta_2$ (i.e., $\alpha_1 = \alpha_2 = 1$)

Test budget $T = 2$

Optimal test allocation:

- allocate $\bar{T} = \frac{7}{8}$ tests to attribute 1 $\Rightarrow$ the posterior covariance matrix is $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$
- allocate the remaining $T - \bar{T} = \frac{9}{8}$ tests equally
- optimal test allocation: $\tau_1^* = \frac{25}{16}$, $\tau_2^* = \frac{9}{16}$

Example 2

$$\begin{pmatrix} \tilde{\theta}_1 \\ \tilde{\theta}_2 \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_1^0 \\ \mu_2^0 \end{pmatrix}, \begin{pmatrix} 8 & 0 \\ 0 & 1 \end{pmatrix} \right)$$

Bliss action is $v = \theta_1 + 2\theta_2$ (i.e., $\alpha_1 = 1$ and $\alpha_2 = 2$)

Test budget $T = 2$

Example 2

$$\begin{pmatrix} \tilde{\theta}_1 \\ \tilde{\theta}_2 \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_1^0 \\ \mu_2^0 \end{pmatrix}, \begin{pmatrix} 8 & 0 \\ 0 & 1 \end{pmatrix} \right)$$

Bliss action is $v = \theta_1 + 2\theta_2$ (i.e., $\alpha_1 = 1$ and $\alpha_2 = 2$)

Test budget $T = 2$

Optimal test allocation:

- allocate $\bar{T}_1 = \frac{3}{8}$ tests to attribute 1 $\Rightarrow$ the posterior covariance matrix is $\begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}$

Example 2

$$\begin{pmatrix}\tilde{\theta}_1\\ \tilde{\theta}_2\end{pmatrix} \sim N\left(\begin{pmatrix}\mu_1^0\\ \mu_2^0\end{pmatrix}, \begin{pmatrix}8 & 0\\ 0 & 1\end{pmatrix}\right)$$

Bliss action is $v = \theta_1 + 2\theta_2$ (i.e., $\alpha_1 = 1$ and $\alpha_2 = 2$)

Test budget $T = 2$

Optimal test allocation:

- allocate $\bar{T}_1 = \frac{3}{8}$ tests to attribute 1 $\Rightarrow$ the posterior covariance matrix is $\begin{pmatrix}2 & 0\\ 0 & 1\end{pmatrix}$
- allocate the remaining $T - \bar{T}_1 = \frac{13}{8}$ tests in fractions $\left(\frac{1}{3}, \frac{2}{3}\right)$
- optimal test allocation: $\tau_1^* = \frac{22}{24}$, $\tau_2^* = \frac{26}{24}$

back

Decomposition, $K = 2$

Decompose the agent's weights as

$$\alpha^R(\beta, \gamma) \equiv \beta\alpha^{DM} + \gamma\hat{\alpha}^P$$

where $\hat{\alpha}^P \equiv (-\alpha_2^{DM}, \alpha_1^{DM})$ is an orthogonal vector to α^{DM}.

Engagement β:

- $\uparrow \beta \iff$ agent more engaged

Distortion γ:

- $\gamma = 0$: agent is undistorted
- $\uparrow |\gamma|$: agent becomes more distorted

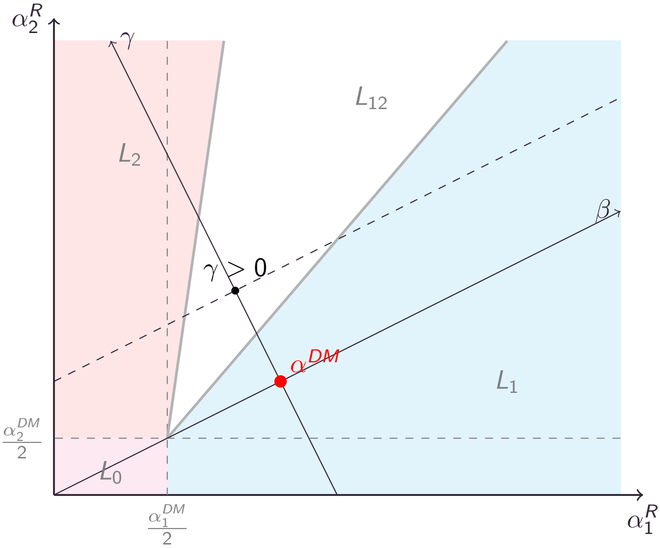

α_2^R
γ
L_{12}
L_2
β
$\gamma > 0$
α^{DM}
L_1
$\frac{\alpha_2^{DM}}{2}$
L_0
$\frac{\alpha_1^{DM}}{2}$
α_1^R

Proposition: Undistorted agent

Suppose the agent is undistorted ($\gamma = 0$). Then

1. If $\beta \leq 1/2$ (agent is disengaged), then the agent does no testing in equilibrium.
2. If $\beta > 1/2$ (agent is sufficiently engaged), then the agent chooses the principal's optimum in equilibrium.

Intuition:

- if $\beta \leq 1/2$: the principal's decision is too sensitive to new information, so the agent prefers to stick with status quo

Proposition: Undistorted agent

Suppose the agent is undistorted ($\gamma = 0$). Then

1. If $\beta \leq 1/2$ (agent is disengaged), then the agent does no testing in equilibrium.
2. If $\beta > 1/2$ (agent is sufficiently engaged), then the agent chooses the principal's optimum in equilibrium.

Intuition:

- the principal's optimum depends on the ratio of his weights
- $\gamma = 0$: the weight ratio is the same for the principal and the agent

Undistorted agent

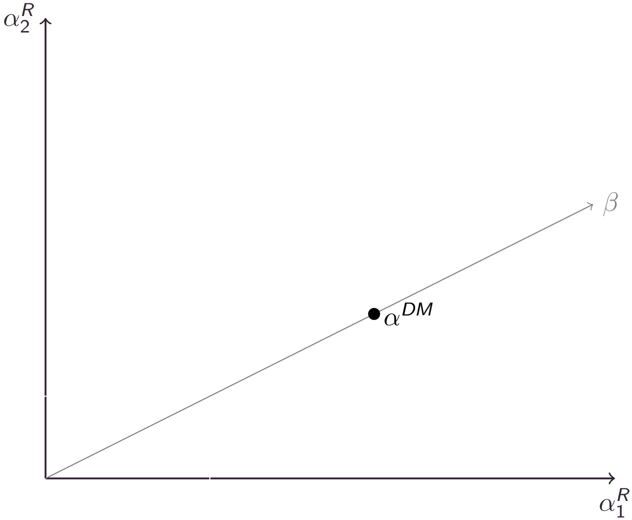

Undistorted agent

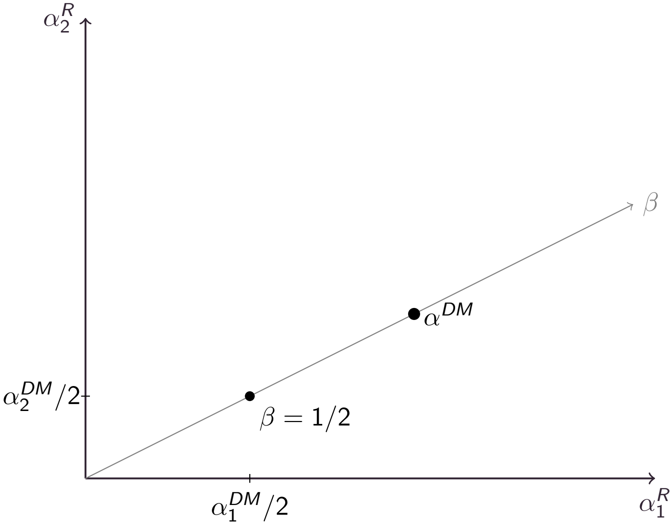

Undistorted agent

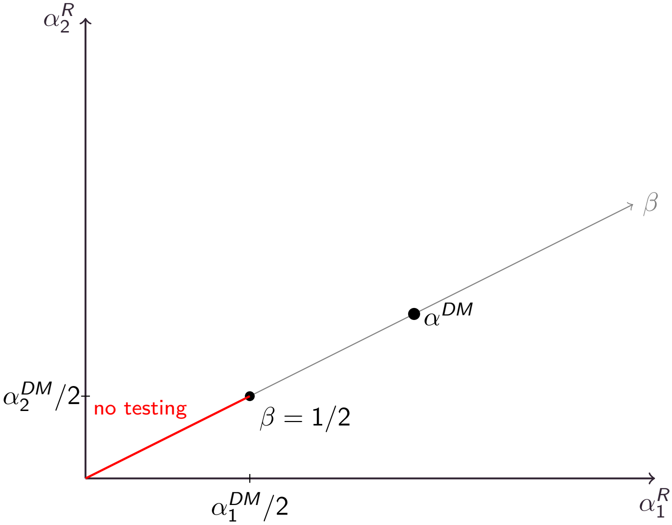

Undistorted agent

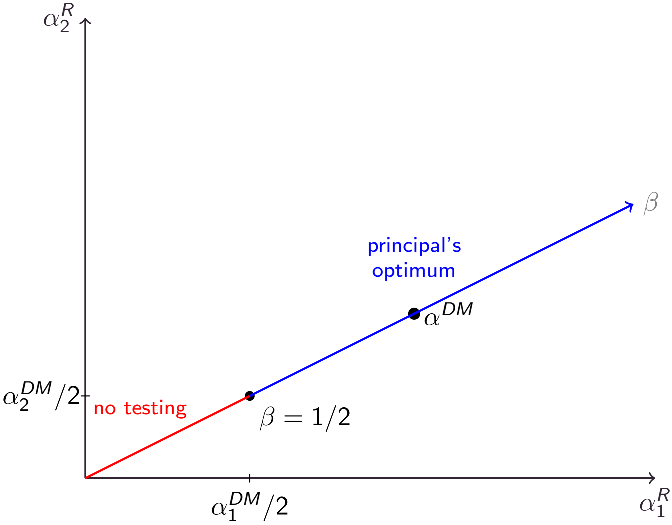

Preference for diversity

Proposition: preference over distortion

1. Fix sensitivity $\beta \leq 1/2$ (agent is disengaged). Then the principal's equilibrium payoff weakly increases in distortion $|\gamma|$.
2. Fix sensitivity $\beta > 1/2$ (agent is sufficiently engaged), then the principal's equilibrium payoff weakly decreases in distortion $|\gamma|$.

Intuition:

- When $\beta \leq 1/2$, an undistorted agent does no testing
  - distorted agent learns about one attribute
  - any info is better for the principal than no info

Preference for diversity

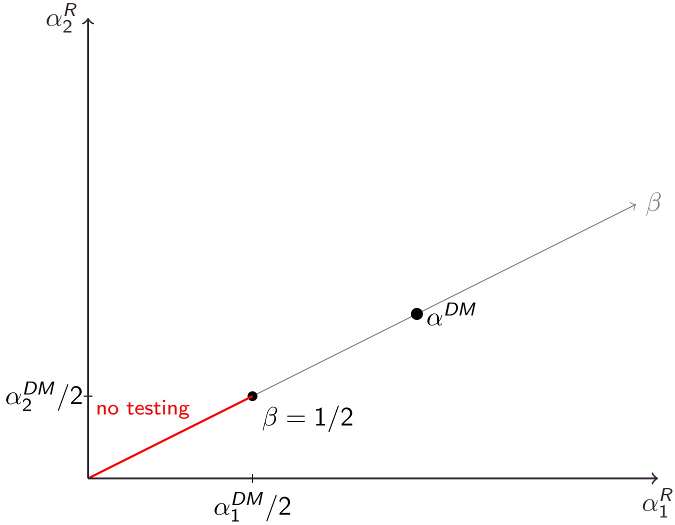

Preference for diversity

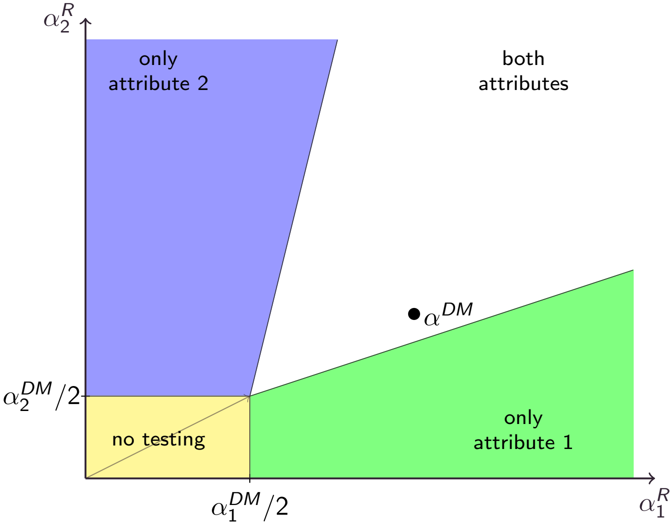

Preference for diversity

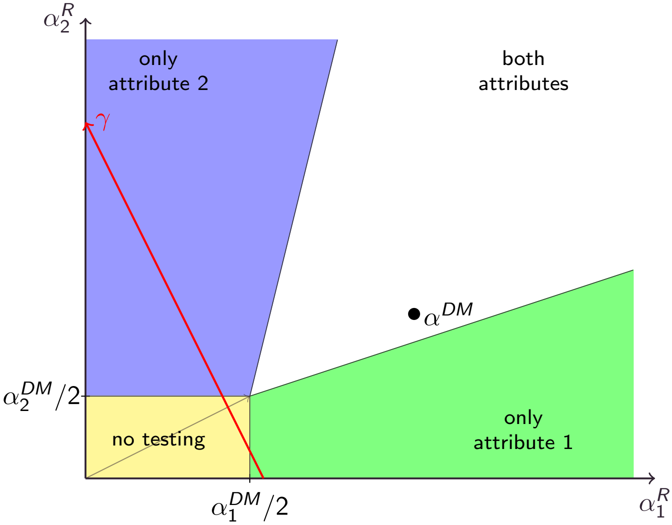

Preference for no diversity

Proposition: preference over distortion

1. Fix sensitivity $\beta \leq 1/2$ (agent is disengaged). Then the principal's equilibrium payoff weakly increases in distortion $|\gamma|$.
2. Fix sensitivity $\beta > 1/2$ (agent is sufficiently engaged), then the principal's equilibrium payoff weakly decreases in distortion $|\gamma|$.

Intuition:

- When $\beta > 1/2$, an undistorted agent chooses the principal's optimum
  - distorted agent distorts the test allocation

Preference for no diversity

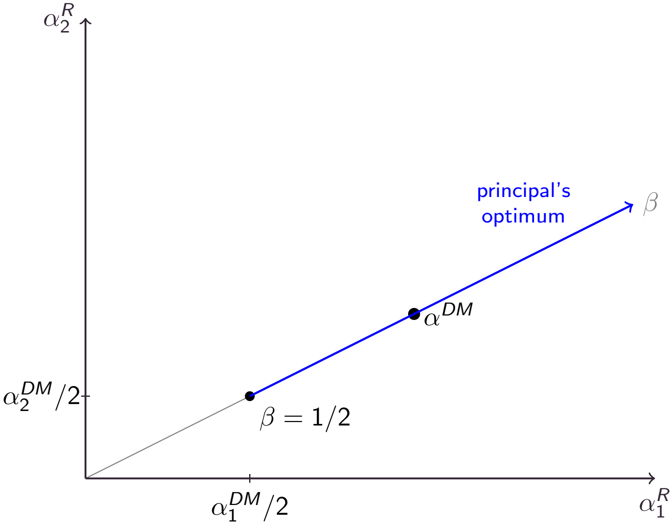

Preference for no diversity

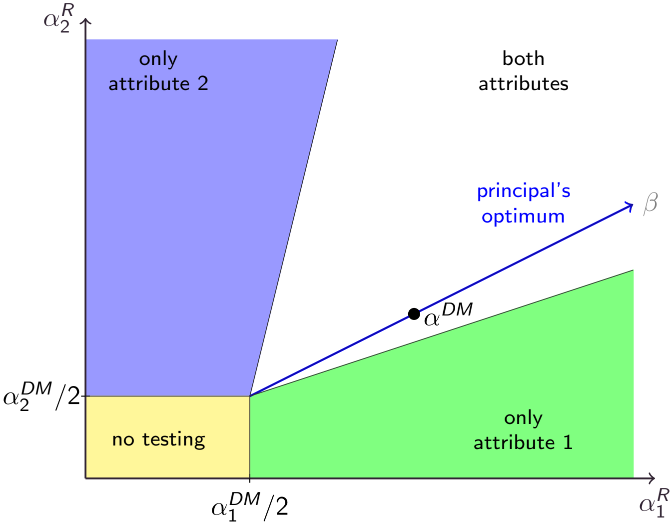

Preference for no diversity

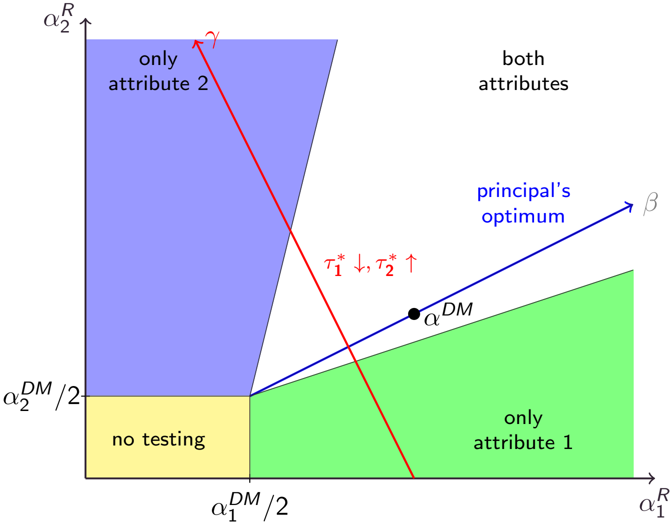

Preference for engagement

Proposition: preference for engagement

Fix distortion $\gamma \in \mathbb{R}$. Then:

- the principal's equilibrium payoff weakly increases in the agent's sensitivity β;
- as $\beta \to \infty$, the agent chooses the principal's optimum.

Preference for sensitivity