Optimal Allocation with Noisy Inspection

Nawaaz Khalfan

January 20, 2023

Inspection

A core economic activity

- employers **interview** potential employees
- public funds **assess** grant applications
- venture capitalists **evaluate** investment opportunities

Why **inspect**?

1. discovery or *information acquisition*
2. verification or *screening*

A class of problems

A **principal** receives an unknown reward from allocating to an **agent**.

The agent has imperfect private information about this unknown reward; they receive a unit reward from being allocated to.

The principal may elicit a report from the agent, as well as inspect the reward at a cost.

The principal can commit to a mechanism, but must do so without transfers.

How should the principal design the inspection and allocation mechanism to maximize their ex ante expected return?

Applications

Mechanism design problems with noisy information, costly inspection, and limited transfers are widespread.

1. **Hiring**: a firm seeks to fill an open position in their operation with a potential employee.

2. **Grant assignment**: a public fund is tasked with assessing a grant application.

3. **Impact investment**: a venture capitalist sets the mechanism by which it reviews and invests in startups.

A simple solution

Let r be the principal's **reward**, and s be the agent's **type**, sorted and labelled by the expected value of the reward.

Symmetric information benchmark:

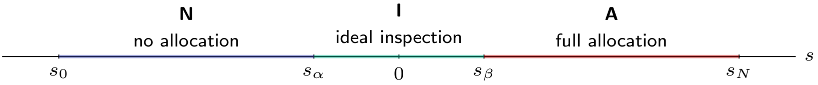

Optimal separating mechanism:

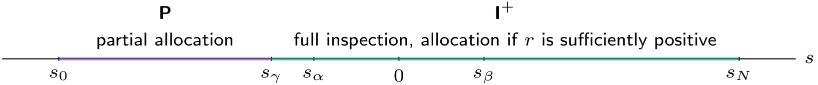

Losses

Three types of losses from private information:

1. over-allocation at the bottom,
2. over-inspection at the top and bottom, and
3. under-allocation post-inspection.

Symmetric information benchmark:

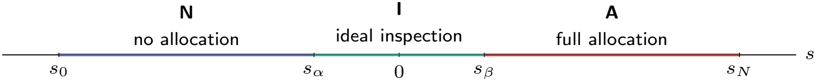

Optimal (separating) mechanism:

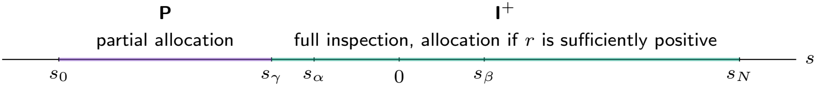

Literature

Perfect information: Green and Laffont (1986), Ben-Porath, Dekel and Lipman (2014), Mylovanov and Zapechelnyuk (2017), Epitropou and Vohra (2019).

Transfers: Townsend (1979), Border and Sobel (1987), Mookherjee and Png (1989), Alaei et al. (2020).

Limited transfers: Mylovanov and Zapechelnyuk (2017), Silva (2019*b*), Li (2021).

Efficient mechanisms: Ball and Kattwinkel (2019), Silva (2019*a*), Siegel and Strulovici (2021), Pereyra and Silva (2021), Erlanson and Kleiner (2020).

Scoring rules: McCarthy (1956), Savage (1971), Gneiting and Raftery (2007).

Environment

The agent is endowed with a signal, s, about the reward that defines the agents **type**.

The principal receives a reward, r, from **allocating** to the agent, and 0 otherwise.

The agent's payment is 1 if allocated to, and 0 otherwise.

The principal can **inspect** the agent to reveal the true reward, r, which costs a fixed $c > 0$ to their final payoff.

Direct transfers of value between the principal and agent are prohibited.

Signals

Suppose $s \in \{s_0, s_1, \ldots, s_N\}$, where $s = s_n$ with probability $p_n \in (0, 1)$, $\sum_n p_n = 1$, and P_n is the cmf.

If $s = s_n$, then the reward $r_n \sim \Pi_n$ where Π_n is absolutely continuous and admits a pdf π_n.

Suppose that the signals are ordered by the monotone likelihood ratio property, **MLRP**.

$$\pi_n(r_1)/\pi_m(r_1) \geq \pi_n(r_0)/\pi_m(r_0) \text{ for all } r_1 > r_0 \text{ and } n > m$$

Note that MLRP $\Rightarrow$ FOSD.

It's without loss to relabel the signals by their induced expected reward, so that $s_n = \mathbb{E}(r|s_n)$.

Timing

The timing of the game is as follows:

1. The principal commits to a mechanism, and nature assigns signals.
2. The agent observes their signal and submits a report to the principal.
3. The principal implements the mechanism conditional on the report and any reward realizations.
4. All remaining uncertainty is resolved, and rewards are distributed.

Mechanism

After the agent reports to the principal, what can the principal do?

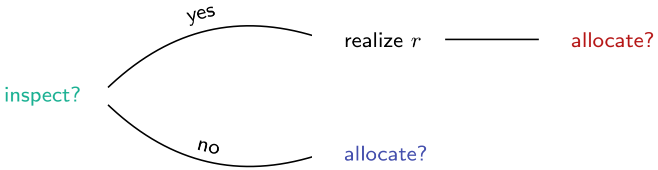

Then, a **mechanism** specifies for each type s,

- x_s: an inspection rule,
- y_s: a pre-inspection allocation, and
- $z_{s,r}$: a post-inspection allocation for each r.

These are potentially probabilistic choices, so are bounded between 0 and 1.

Rewards

Principal's objective:

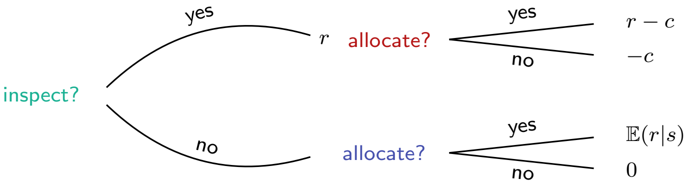

Agent's incentives: 1 if allocated to, 0 otherwise.

An **optimal allocation** is a mechanism that maximizes the ex ante expected objective subject to *incentive compatibility* (IC) for each type s:

$$u(s|s) \geq u(\hat{s}|s) \quad \forall \hat{s}$$

Optimal allocation

The principal's problem:

$$\max_{(x,y,z)} \quad \sum_n [(1-x_n)y_n\mathbb{E}(r|s_n) + x_n\psi_n(z_n)]p_n$$

$$\text{s.t.} \quad IC_{n,m} : (1-x_n)y_n + x_n\mathbb{E}(z_{n,r}|n) \geq (1-x_m)y_m + x_m\mathbb{E}(z_{m,r}|n) \quad \forall\ n,m$$

$$F : 0 \leq x_n, y_n, z_{n,r} \leq 1 \quad \forall\ r \quad \forall\ n$$

where:

- $\psi_n(z_n) \coloneqq \mathbb{E}(z_{s,r}.r|s) - c = \int r z_{n,r}\pi_{n,r}\, dr - c$, is the expected reward from inspecting n with post-inspection allocation rule z_n.

First best policy, $*$

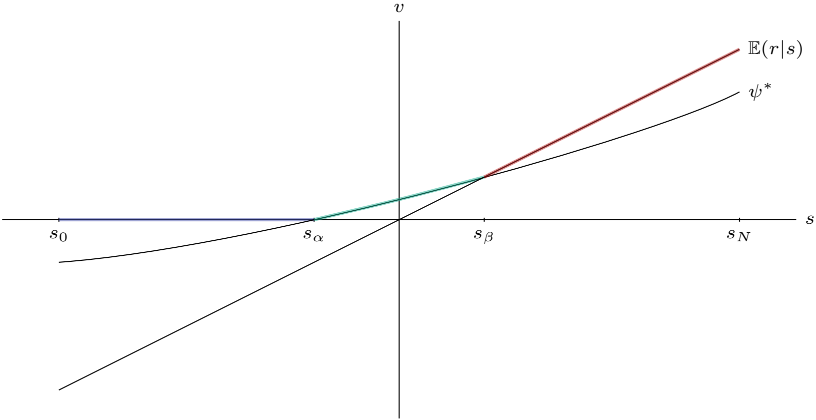

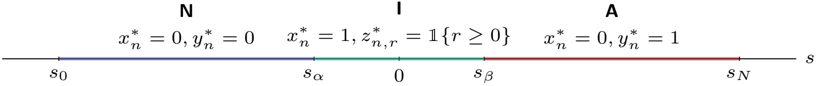

Second best policy, $\star$

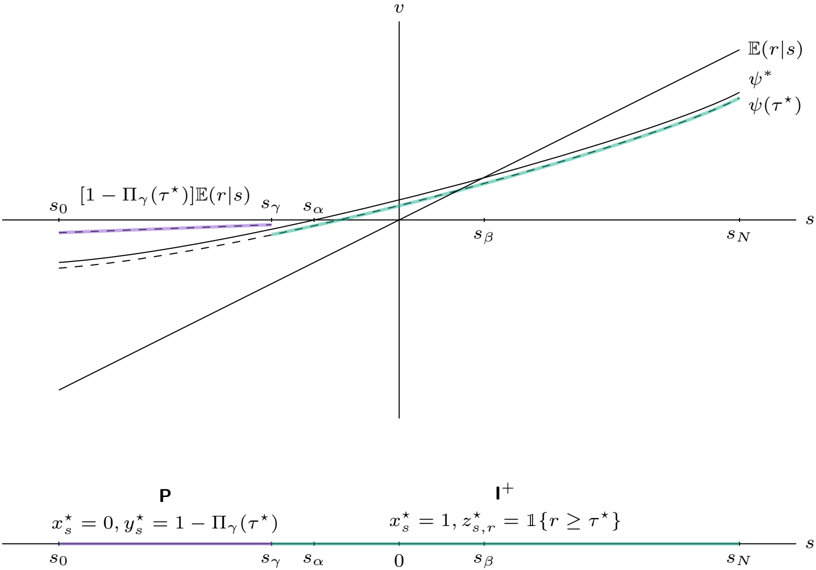

A solution recipe

Consider a **relaxation** of the principal's problem that only requires the upward local IC constraints to be satisfied. That is:

$$IC_{n,n+1} : (1-x_n)y_n + x_n \mathbb{E}(z_{n,r}|n) \geq (1-x_{n+1})y_{n+1} + x_{n+1}\mathbb{E}(z_{n+1,r}|n) \quad \forall\, n < N$$

Claim 1: Optimal post-inspection rules are threshold rules. That is, for each s_n there exists some τ_n such that allocation only occurs post-inspection if $r > \tau_n$.

Claim 2: Each upward local incentive compatibility constraint binds. That is, for each s_n, $u(s_n|s_n) = u(s_{n+1}|s_n)$.

Claim 3: Optimal inspection rules are themselves threshold rules. That is, there exists γ such that the agent is only inspected if $s_n > s_\gamma$.

$\Rightarrow$ Optimal post-inspection thresholds are constant: $\tau_n = \tau\ \forall n$.

1. Threshold post-inspection allocation

Claim 1: Optimal post-inspection rules are threshold mechanisms. That is, for each n there exists some τ_n such that:

$$z_{n,r} = \mathbb{1}\{r \geq \tau_n\}$$

Idea: For each n find the τ_n such that:

$$\int z_{n,r}\pi_{n,r}\,dr = \int \mathbb{1}\{r \geq \tau_n\}\pi_{n,r}\,dr$$

This transformation will always improve the objective, maintain the expected payoff for n, and weakly reduce the expected deviation payoff for $n-1$.

The transformation

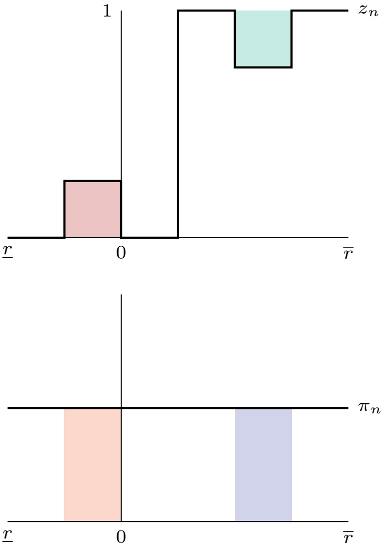

MLRP

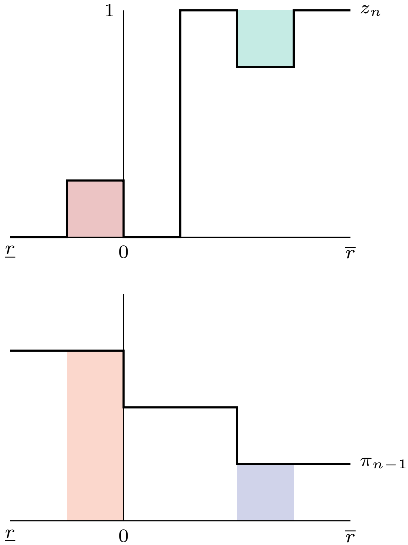

FOSD

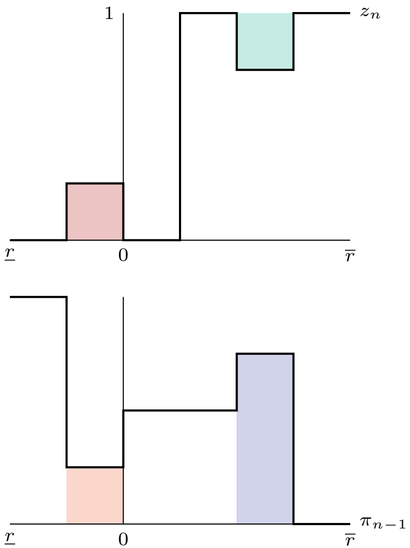

Threshold tests

Post-inspection allocations are then determined by a simple threshold test.

For the agent:

$$\mathbb{E}(z_{n,r}|n) = \int \mathbb{1}\{r \geq \tau_n\}\pi_{n,r}\, dr = 1 - \Pi_n(\tau_n)$$

For convenience, let's denote:

$$\overline{\Pi}_n(\tau) \coloneqq 1 - \Pi_n(\tau)$$

2. Binding ULIC

Claim 2: Each upward local incentive compatibility constraint binds. That is, for each $n < N$:

$$(1 - x_n)y_n + x_n\overline{\Pi}_n(\tau_n) = (1 - x_{n+1})y_{n+1} + x_{n+1}\overline{\Pi}_n(\tau_{n+1})$$

Idea: Consider the following partition:

1. $S_0 := \{n \mid 0 \geq \mathbb{E}(r|s_n),\ 0 \geq \psi_n(\tau_n)\}$
2. $S_\alpha := \{n \mid 0 \geq \mathbb{E}(r|s_n),\ \psi_n(\tau_n) > 0\}$
3. $S_\beta := \{n \mid \mathbb{E}(r|s_n) > 0,\ \psi_n(\tau_n) > \mathbb{E}(r|s_n)\}$
4. $S_1 := \{n \mid \mathbb{E}(r|s_n) > 0,\ \mathbb{E}(r|s_n) \geq \psi_n(\tau_n)\}$

Note, that if $\tau_n = 0$ for each n, this corresponds with our first best policy:

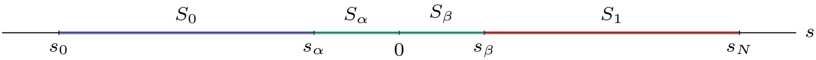

3. Threshold inspection rules

Claim 3: Optimal inspection rules are threshold mechanisms. That is, there exists n_0 such that $x_n = \mathbb{1}\{n \geq n_0\}$.

Idea: We can now rewrite $(1 - x_n)y_n$ recursively:

$$
\begin{aligned}
(1 - x_n)y_n &= (1 - x_{n+1})y_{n+1} \quad + x_{n+1}\overline{\Pi}_n(\tau_{n+1}) - x_n\overline{\Pi}_n(\tau_n) \\
&= (1 - x_{n+2})y_{n+2} \quad + x_{n+2}\overline{\Pi}_{n+1}(\tau_{n+2}) - x_{n+1}\overline{\Pi}_{n+1}(\tau_{n+1}) \\
&\qquad\qquad\qquad\qquad + x_{n+1}\overline{\Pi}_n(\tau_{n+1}) - x_n\overline{\Pi}_n(\tau_n) \\
&= \cdots \\
&= (1 - x_N)y_N \qquad + \sum_{m=n}^{N-1} [x_{m+1}\overline{\Pi}_m(\tau_{m+1}) - x_m\overline{\Pi}_m(\tau_m)]
\end{aligned}
$$

A linear objective

Our value function becomes:

$$\begin{aligned} v = \quad & (1 - x_N) y_N \mathbb{E}(r) \\ + \quad & x_N [\overline{\Pi}_{N-1}(\tau_N) \mathbb{E}(r | s \leq s_{N-1}) P_{N-1} + \psi_N(\tau_N) p_N] \\ + \sum_{n=1}^{N-1} \quad & x_n [\overline{\Pi}_{n-1}(\tau_n) \mathbb{E}(r | s \leq s_{n-1}) P_{n-1} - \overline{\Pi}_n(\tau_n) \mathbb{E}(r | s \leq s_n) P_n + \psi_n(\tau_n) p_n] \\ + \quad & x_0 [-\overline{\Pi}_0(\tau_0) \mathbb{E}(r | s_0) p_0 + \psi_0(\tau_0) p_0] \end{aligned}$$

This is a linear function in x_n. Similar to the proof of claim 2, we can then use variation arguments to prove that, $x_n \in \{0, 1\}$.

For example, our constraint directly implies that for any consecutive signals such that $x_n = x_{n+1} = 1$, then $\tau_n = \tau_{n+1}$, as $\overline{\Pi}_n(\tau_{n+1}) = \overline{\Pi}_n(\tau_n)$.

Optimal separating policy

Given **Claims 1-3**, we are only left to optimize by selecting:

- γ: the first type to inspect, and
- τ: the threshold for passing those who are inspected.

This is given by:

- the value of those signals below γ, that we partially allocate to, and
- the value of those signals above γ, that we inspect with threshold τ.

$$\max_{\gamma,\tau} \; Pr(r > \tau | s_\gamma)\mathbb{E}(r | s \leq s_\gamma) \cdot Pr(s \leq s_\gamma) + v(\mathbf{I}(\tau) | s > s_\gamma) \cdot Pr(s > s_\gamma)$$

This satisfies the **global** incentive compatibility constraints for all γ and τ, and thus must be a solution to the original problem.

Second best solution

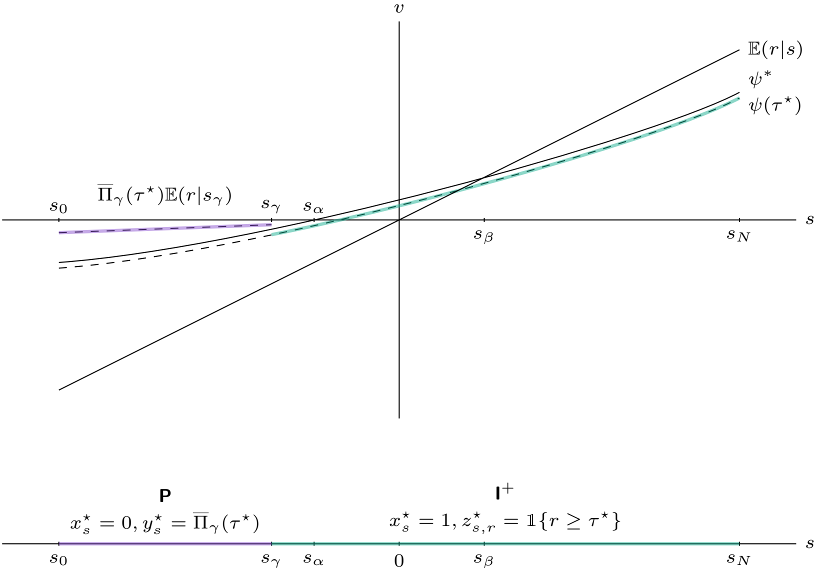

Losses

Public information benchmark:

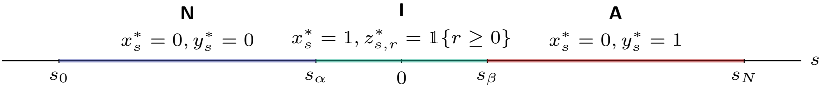

Optimal (separating) mechanism:

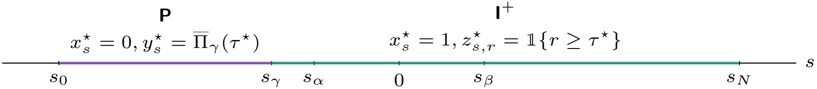

Three types of losses from private information:

1. over-allocation: for $s \in [s_0, s_\gamma]$, $y_s^\star = \overline{\Pi}_\gamma(\tau^\star) > 0$,
2. over-inspection: for $s \in [s_\gamma, s_\alpha] \cup [s_\beta, s_N]$, $x_s^\star = 1$, and
3. under-allocation post-inspection: for $s \in [s_\gamma, s_N]$ and $r \in [0, \tau^\star]$, $z_{s,r}^\star = 0$.

Noisy inspection

Optimal inspection balances *discovery* and *verification*.

When agents have **noisy private information**, the principal:

- over-inspects high and low types,
- under-allocates to agents who are inspected, and
- over-allocates to agents who are not inspected.

Weakening commitment magnifies the losses from over-allocating to agents who aren't inspected.

For separating to be optimal, signals need to be sufficiently accurate, costs sufficiently small and information sufficiently valuable.

Outstanding questions?

References

Alaei, Saeed, Alexandre Belloni, Ali Makhdoumi and Azarakhsh Malekian. 2020. "Optimal Auction Design with Deferred Inspection and Reward." *SSRN* .

Ball, Ian and Deniz Kattwinkel. 2019. "Probabilistic Verification in Mechanism Design." *Working Paper* .

Ben-Porath, Elchanan, Eddie Dekel and Barton L. Lipman. 2014. "Optimal Allocation with Costly Verification." *American Economic Review* 104(12):3779–3813.

Border, Kim C. and Joel Sobel. 1987. "Samurai Accountant: A Theory of Auditing and Plunder." *Review of Economic Studies* 54(4):525–540.

Epitropou, Markos and Rakesh Vohra. 2019. "Dynamic Mechanisms with Verification." *PIER Working Papers* 19:002.

Erlanson, Albin and Andreas Kleiner. 2020. "Costly verification in collective decisions." *Theoretical Economics* 15(3):923–954.

Gneiting, Tilmann and Adrian E. Raftery. 2007. "Strictly Proper Scoring Rules, Prediction, and Estimation." *Journal of the American Statistical Association* 102(477):359–378.

Green, Jerry R. and Jean-Jacques Laffont. 1986. "Partially Verifiable Information and Mechanism Design." *The Review of Economic Studies* 53(3):447–456.

Li, Yunan. 2021. "Mechanism design with financially constrained agents and costly verification." *Theoretical Economics* 16(3):1139–1194.

McCarthy, John. 1956. "Measures Of The Value Of Information." *Proceedings of the National Academy of Sciences* 42(9):654–655.

Mookherjee, Dilip and Ivan Png. 1989. "Optimal Auditing, Insurance, and Redistribution." *The Quarterly Journal of Economics* 104(2):399–415.

Mylovanov, Tymofiy and Andriy Zapechelnyuk. 2017. "Optimal Allocation with Ex Post Verification and Limited Penalties." *American Economic Review* 107(9):2666–94.

Pereyra, Juan and Francisco Silva. 2021. "Optimal object assignment mechanisms with imperfect type verification." *Working Paper* .

Savage, Leonard J. 1971. "Elicitation of Personal Probabilities and Expectations." *Journal of the American Statistical Association* 66(336):783–801.

Siegel, Ron and Bruno Strulovici. 2021. "Judicial Mechanism Design." *Working Paper* .

Silva, Francisco. 2019*a*. "If We Confess Our Sins." *International Economic Review* 60(3):1389–1412.

Silva, Francisco. 2019*b*. "Renegotiation-proof mechanism design with imperfect type verification." *Theoretical Economics* 14(3):971–1014.

Townsend, Robert M. 1979. "Optimal contracts and competitive markets with costly state verification." *Journal of Economic Theory* 21(2):265–293.

Relaxing commitment

There are three natural relaxations to the commitment assumption:

1. **pre-inspection commitment**: the principal can commit to pre-inspection allocations and an inspection rule but cannot commit to post-inspection allocations,
2. **pre-assessment commitment**: the principal cannot commit to either an inspection rule or post-inspection allocations, but can commit to pre-inspection allocations, and
3. **no commitment**: the principal cannot commit to allocations or an inspection rule.

For **no commitment**, the principal can only choose between the pooling mechanisms and reports convey no information. We know what this looks like, so let's turn to the first two relaxations.

Pre-assessment commitment

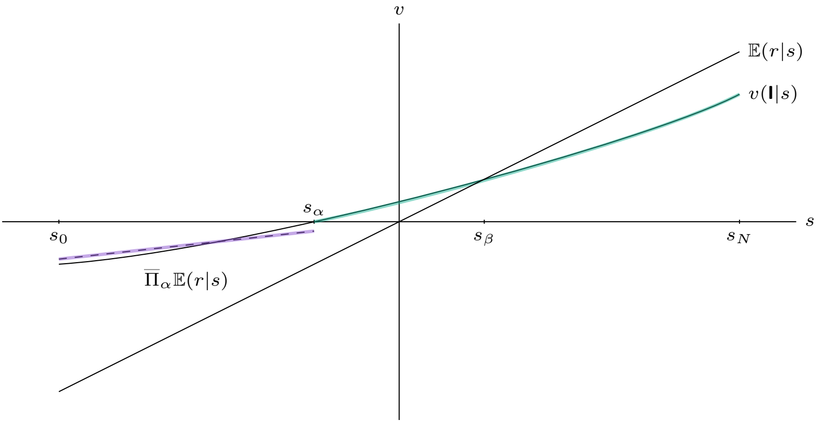

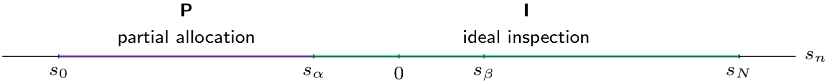

Pre-inspection commitment

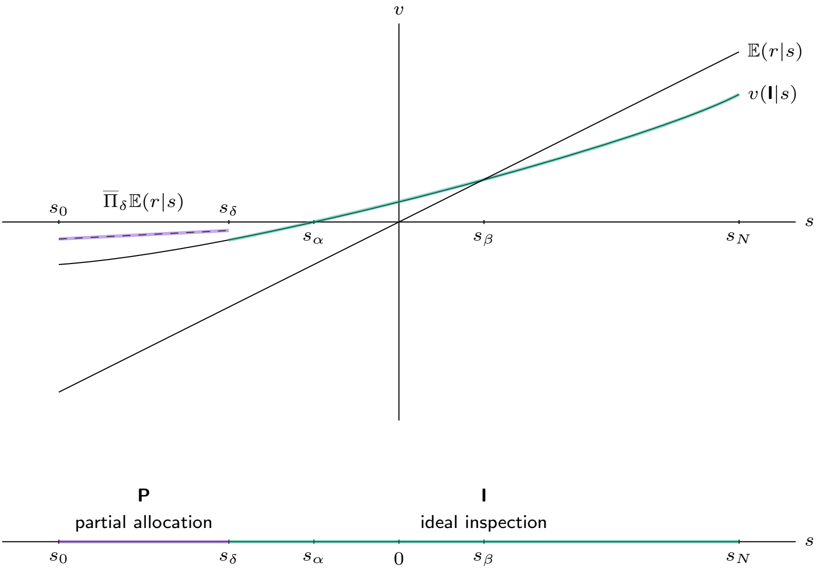

Full commitment

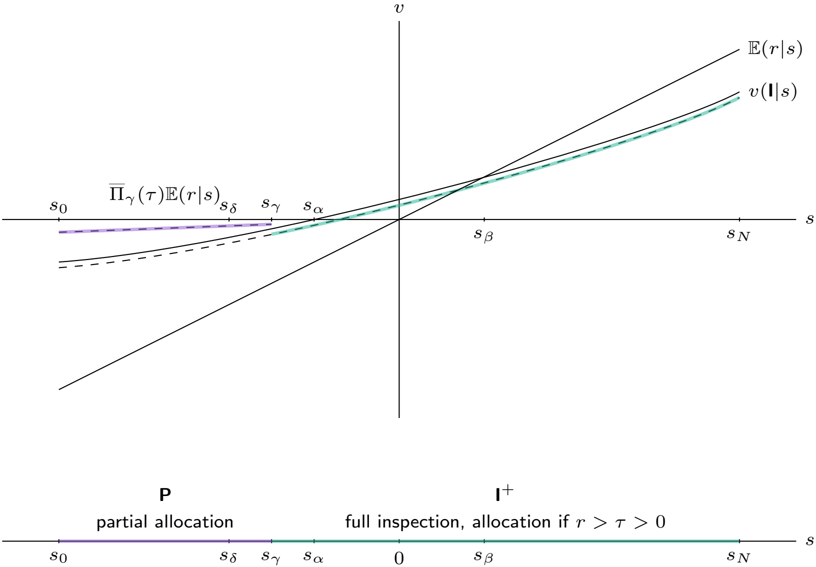

Gaussian environment

Suppose the prior over rewards is given by: $r \sim N(\mu, 1)$, and the agent receives a signal of this reward, $\hat{s} = r + \varepsilon$, where $\varepsilon \sim N(0, \sigma^2)$.

Relabelling the signal by the expected reward given the signal, the posterior distribution of rewards, Π_s, is given by: $r \mid s \sim N(s, \hat{\sigma}^2)$ where:

$$s = \frac{\sigma^2}{\sigma^2 + 1}\left[\mu + \frac{\hat{s}}{\sigma^2}\right] \quad \text{and} \quad \hat{\sigma}^2 = \frac{\sigma^2}{\sigma^2 + 1}$$

The induced distribution of signals, P, is then given by: $s \sim N(\mu, \frac{1}{\sigma^2+1})$.

The environment is by a triple:

- μ, the ex-ante expected reward of allocating to an agent,
- $\alpha := 1/\sigma^2$, the precision of the agent's signal of the reward, and
- c, the inspection cost to the principal.

Pooling equilibria

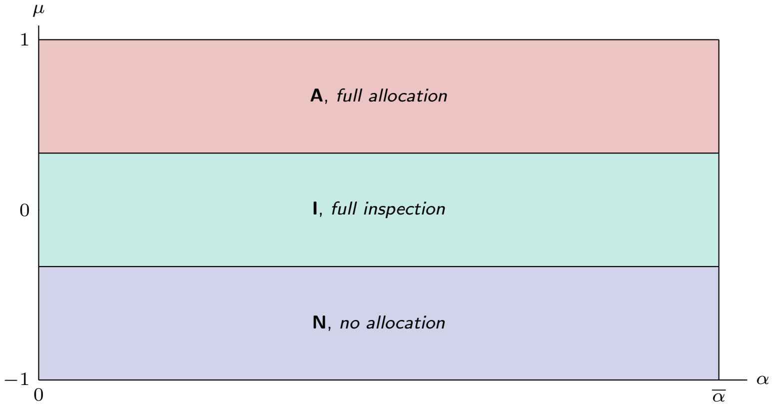

Figure: third-best policy as a function of precision, α, and prior mean, μ

Comparative statics

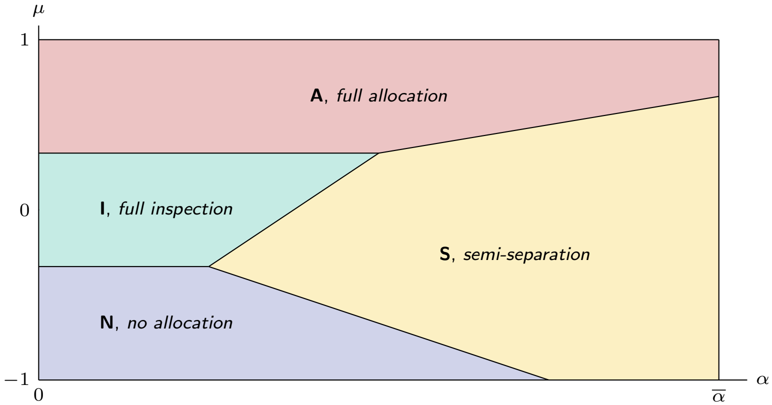

Figure: second-best policy as a function of precision, α, and prior mean, μ