Delegated Pandora's Box

Curtis Bechtel, Shaddin Dughmi, **Neel Patel**

University of Southern California

Problem Definition

- Principal looks to solve an optimization problem under uncertainty and delegates collecting the data (a.k.a. probing) to an expert Agent.
- **Example:** Firm (principal) delegates the candidate selection (constrained) to an outside recruitment agency (agent).
- Candidates (elements) are associated with stochastic reward for both $(X_e, Y_e) \sim \mu_e$ independently across elements.
- The recruitment agency (agent) adaptively interviews (probes) and learns realization of interviewed (probed) candidates (elements).
- Interviewing (probing) candidates (element) incurs costs.
- Non-delegated stochastic optimization problem in the above example is known as *Generalized Pandora's Box.*

Problem Statement

How can the principal (firm) incentivize the agent (requirement agency) to probe (interview) and select a set of elements (candidates) with high value for the principal (firm)?

Generalized Pandora's Box

- Given n independent random variables $X_1, \ldots, X_n \geq 0$, probing costs $c_1, \ldots, c_n$, packing constraints $\mathcal{I}$, (adaptively) probe elements T and select feasible $S \subseteq T$ that maximizes:

$$\mathbb{E}\left[X(S) - c(T)\right].$$

Who Pays the Probing Costs?

- **Fixed Cost (Standard Model [KK19]):** Principal and agent split equally, or a fixed percentage.
- **Free agent:** Principal pays the entire probing cost, agent provides expertise.
- **Custom Cost:** Principal stipulates percentage for each element.

Delegation Mechanisms

Mechanism design without money.

- Due to [KK19], we can focus on simple class of mechanisms: Single-Proposal Mechanism.
- Principal commits to menu $\mathcal{R}$ of acceptable valuated solutions.
  - Valuated soln: $S \in \mathcal{I}$ with each $e \in S$ tagged with acceptable x_e
- Agent adaptively probes subset of elements and proposes soln from $\mathcal{R}$ which is valid.

For example: Firm only will accept two candidates, s.t.

- At least one of them with PhD in CS/Math,
- At most one overseas, etc.

Technical Challenges

Pandora algorithms typically split each element e's distribution:

- Above and below the reservation/cap value τ_e (covers cost needs to be paid in expectation):

$$\mathbb{E}[(X_e - \tau_e)_+] = c_e$$

- When e is probed and $x_e \geq \tau_e$ then e is selected (selected set remains feasible).
- However, agent might have different preferences and ignore probed element with $x_e \geq \tau_e$

Example:
Principal needs to select one out of n candidates.

- For all e, $X_e = n$ w.p. $1/n$ and 0 otherwise, $Y_e = n$ w.p. $1/n$ and 0 otherwise (ind. of X_e).
- Agent probes e and observes $(x_e, y_e) = (n, 0)$ then agent will ignore e and move forward!

Delegation Gap

The delegation gap is the worst-case ratio of the principal's optimal delegated utility versus their optimal undelegated utility. **Question:** *Which packing constraints and cost division models exhibit a constant delegation gap? Can we characterize such a class of packing constraints?*

Standard Model

Proposition: No delegation mechanims can obtain constant delegation gap for rank one matroid constraints, even when costs of each items are discounted by a constant factor.

- Let for all e, $X_e = n$ w.p. $1/n$ and 0 otherwise, $Y_e = n$ w.p. $1/n$ and 0 otherwise (ind. of X_e), $c_e = 0.4$ for both.
- Let principal accepts $(n, n), (n, 0)$ realizations of e. If agent probes e then their expected gain is

$$n \cdot \Pr[X_e = n, Y_e = n] - 0.4 = \frac{1}{n} - 0.4 < 0$$

Free Agent Model

Theorem: For "nice enough" packing constraints, delegation gap is constant when costs are discounted by constant factor.

- α-strong Online Contention Resolution Schemes (OCRS) for $P_\mathcal{I}$ implies α-delegation gap when costs are discounted by $(1 - \alpha)$ factor.

Proof Sketch

- Let $Z_e = \min\{\tau_e, X_e\}$. We can show that OPT$\leq \mathbb{E}[\max_{S\in\mathcal{I}} z(S)]$
- Let $p_e = \Pr[e \in arg\max_{S\in\mathcal{I}} z(S)]$. Consider threshold t_e s.t. $\Pr[Z_e \geq t_e] = p_e$. Agent is allowed to probe e iff $\tau_e \geq t_e$.
- α-OCRS ensures that each element is selected by the agent with prob. $\geq \alpha \cdot p_e$

Custom Cost Model

Theorem: *α-OCRS for $P_\mathcal{I}$ implies $\alpha/2$-delegation gap.*

Proof Sketch

- The idea is similar to the free-agent case.
- Ask agent to pay the probing cost of elements for which principal has small expected value.