Near-Optimal Dropout-Robust Sortition

Maya Pal Gambhir
University of Pennsylvania

Bailey Flanigan
Massachusetts Institute of Technology

Aaron Roth
University of Pennsylvania

AISTATS 2026

What is a citzen’s assembly?

Citizens assemblies are a method for democratic decision-making where a **representative** group of citizens is **randomly** selected for *policy deliberation*.

The well studied problem: Sortition is the process of generating an assembly from a pool of volunteers, maintaining key properties like demographic representation.

The well studied problem: Sortition is the process of generating an assembly from a pool of volunteers, maintaining key properties like demographic representation.

Pool of Volunteers

Volunteers with attributes, skewed from population

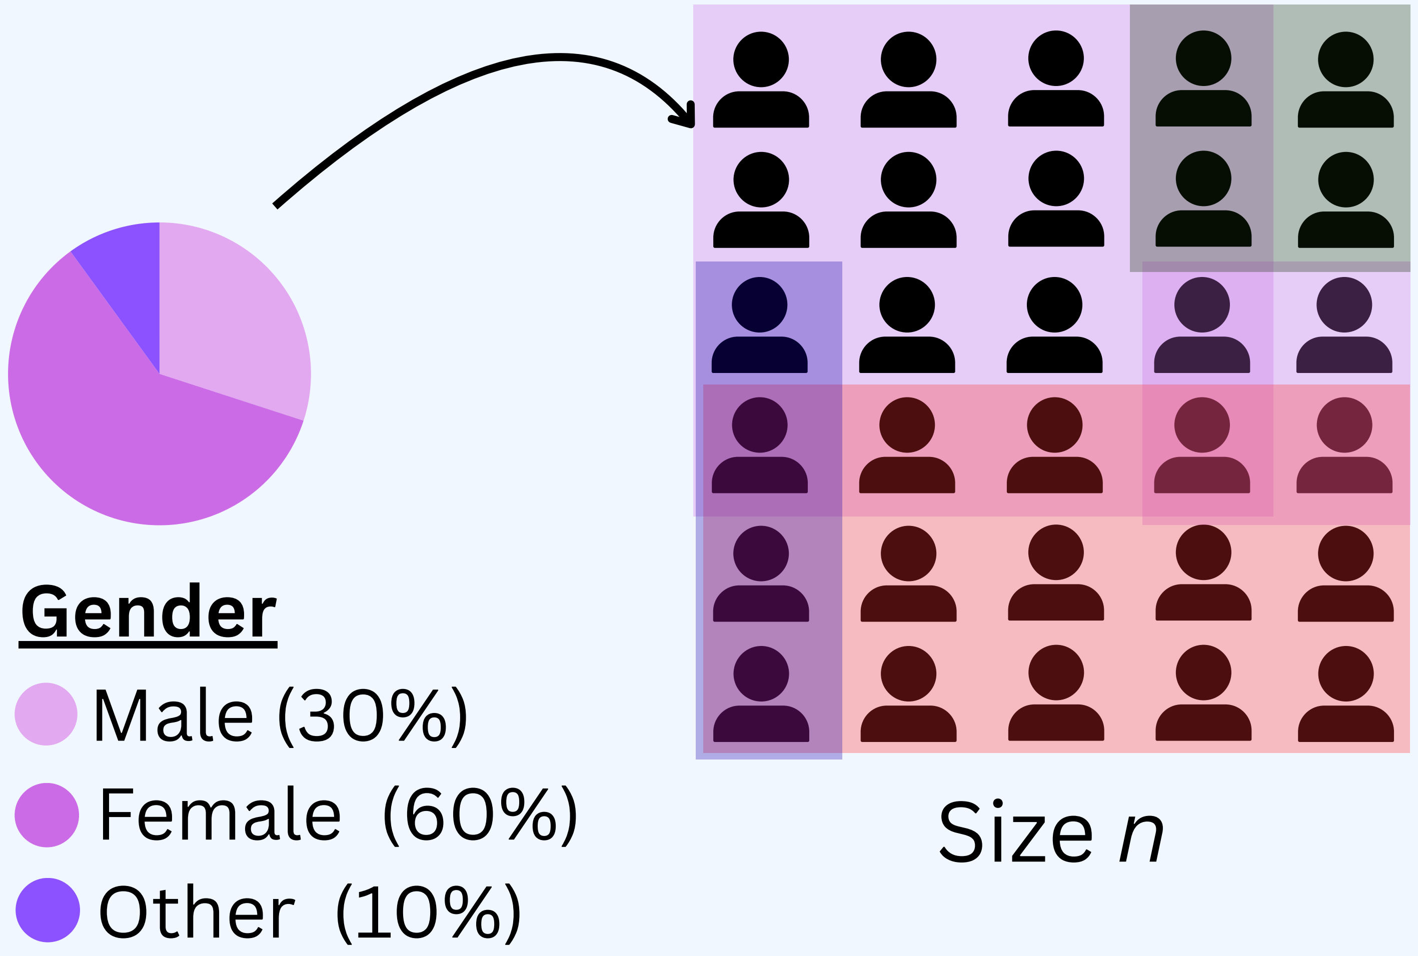

The well studied problem: Sortition is the process of generating an assembly from a pool of volunteers, maintaining key properties like demographic representation.

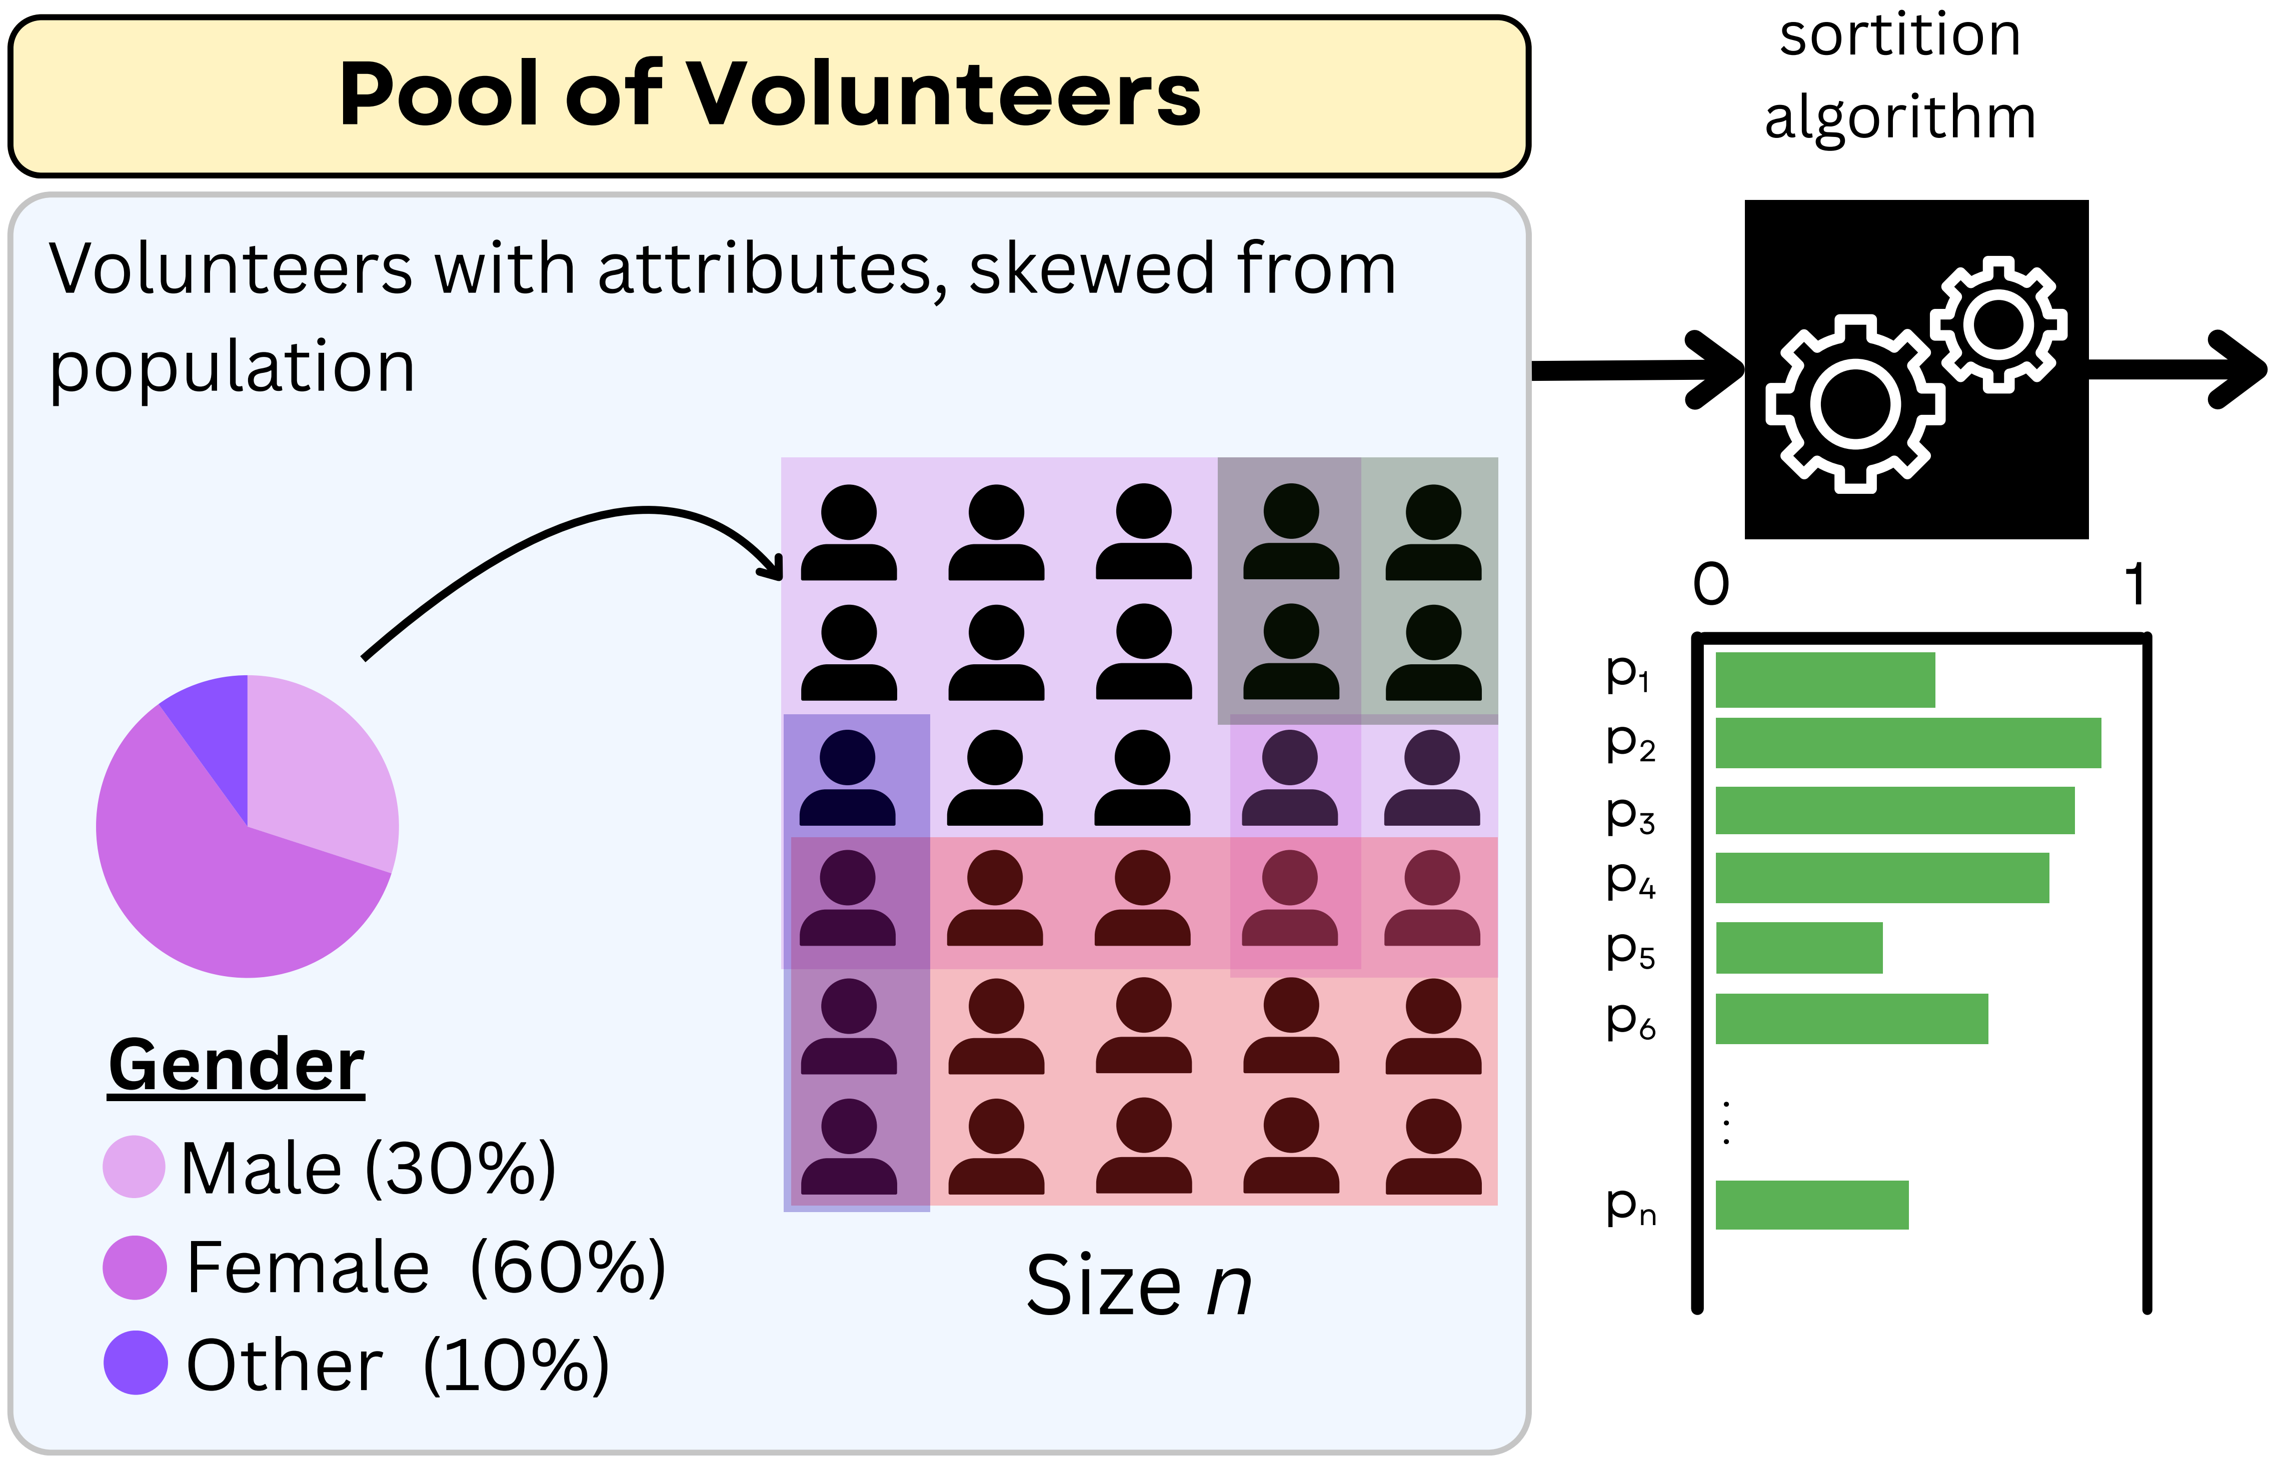

Sortition Overview

The well studied problem: Sortition is the process of generating an assembly from a pool of volunteers, maintaining key properties like demographic representation.

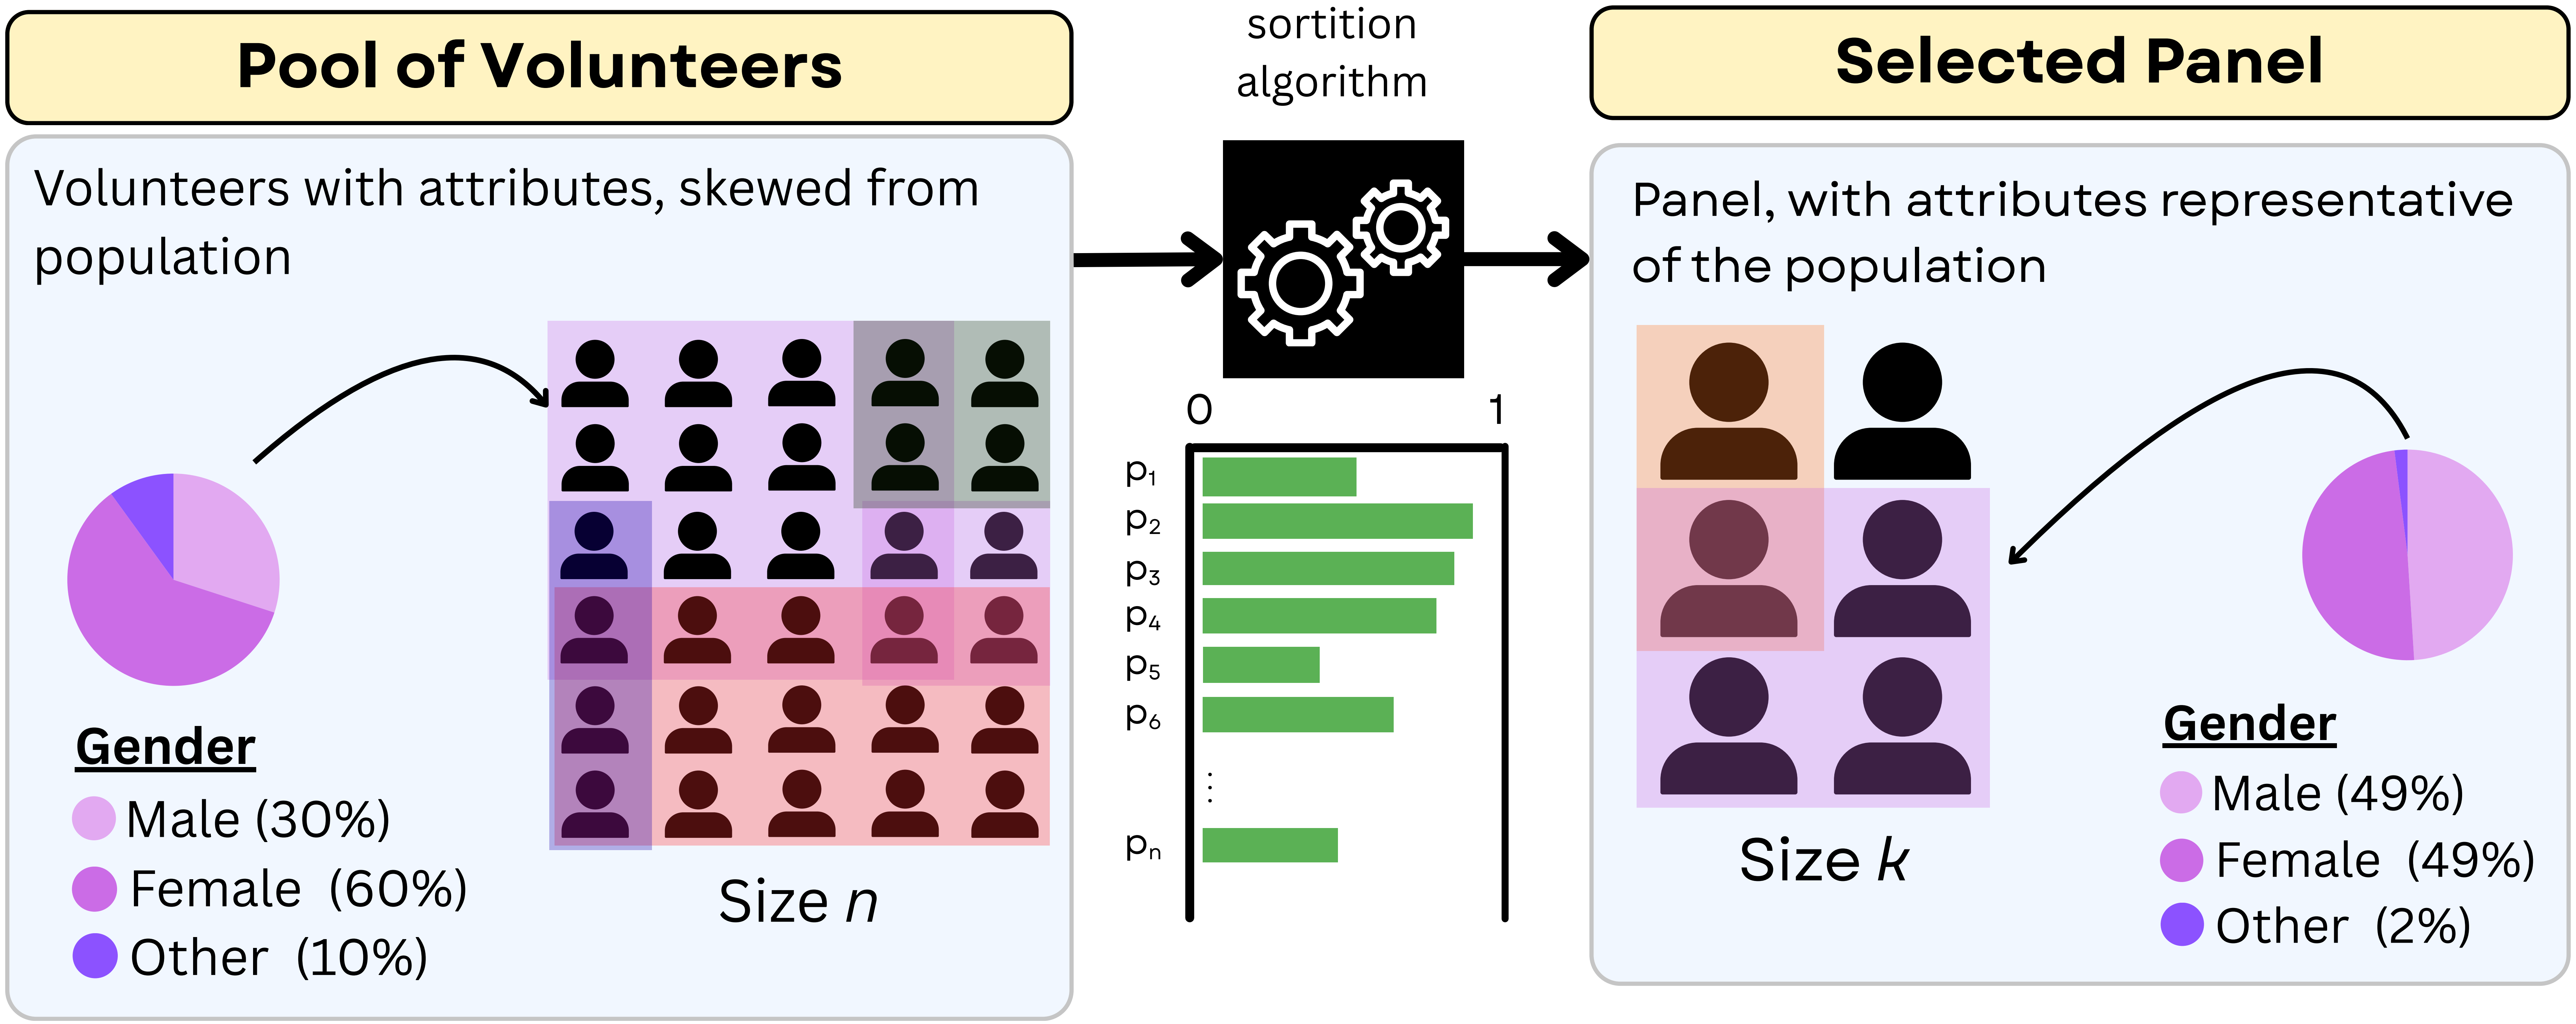

Desired Attributes for a Panel

Objective: Demographic representation *(offsets self-selection bias)*

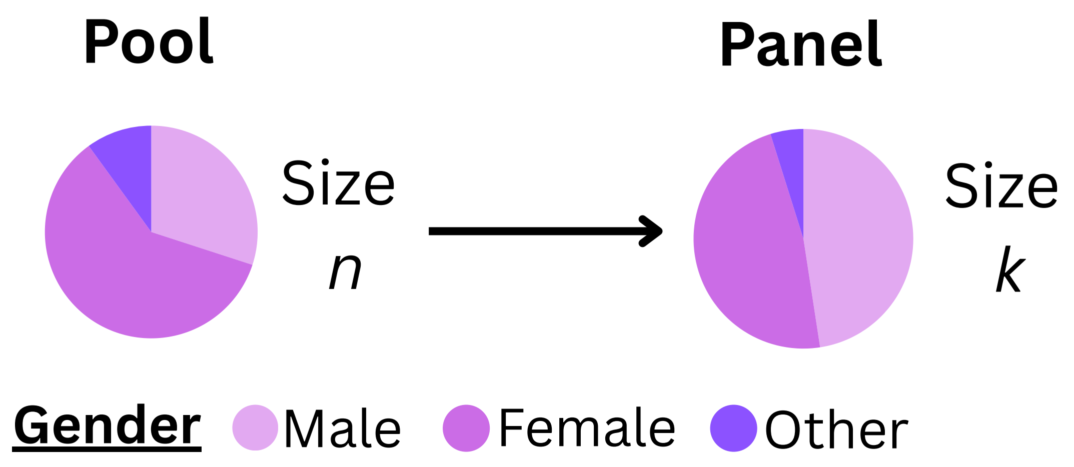

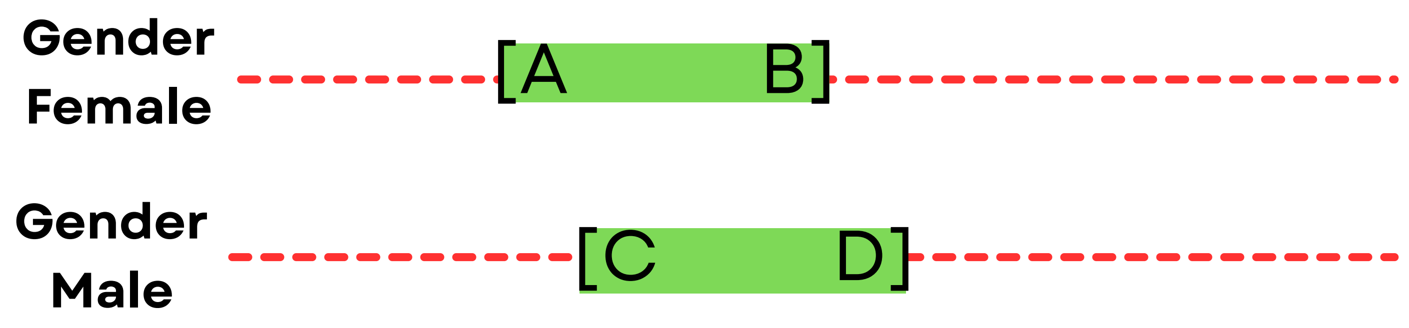

Representation is formalized with upper and lower **quotas** on each feature and value pair

Objective: Demographic representation *(offsets self-selection bias)*

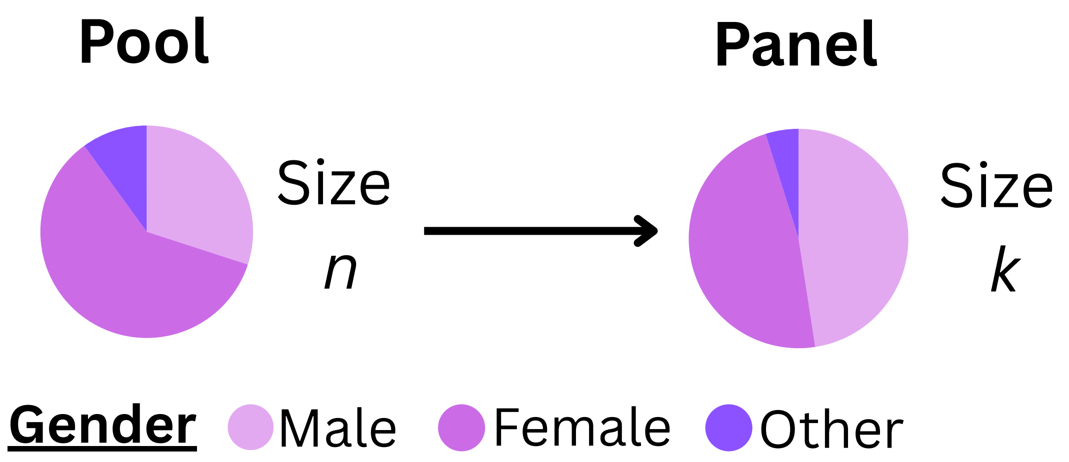

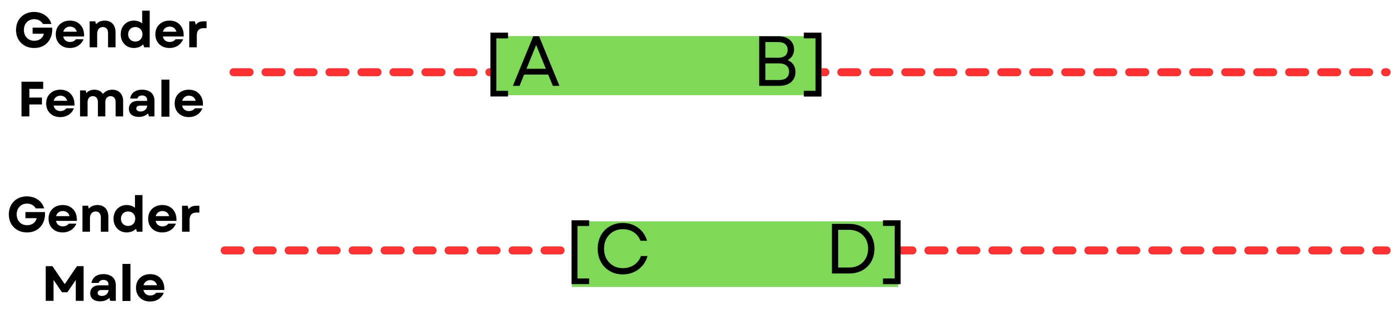

Representation is formalized with upper and lower **quotas** on each feature and value pair

Loss = maximum (relative) violation of any quota

Objective: Demographic representation *(offsets self-selection bias)*

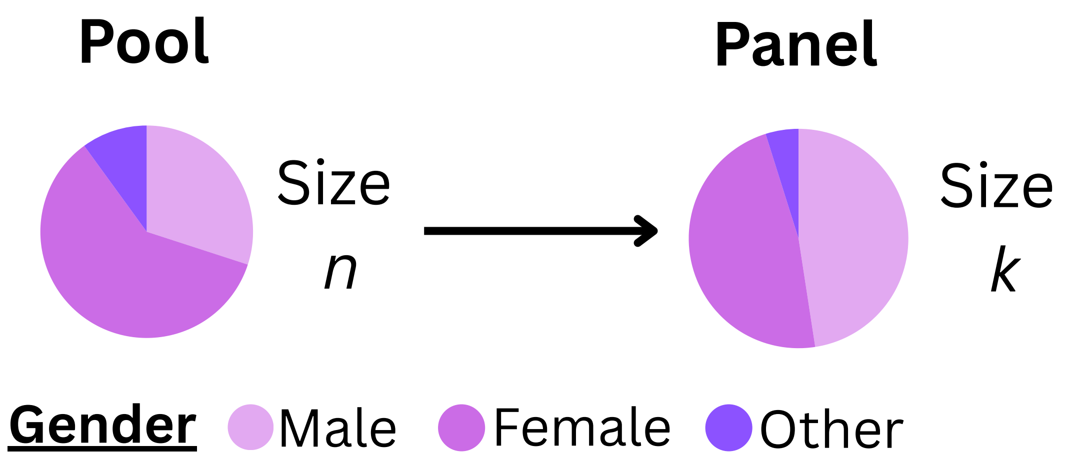

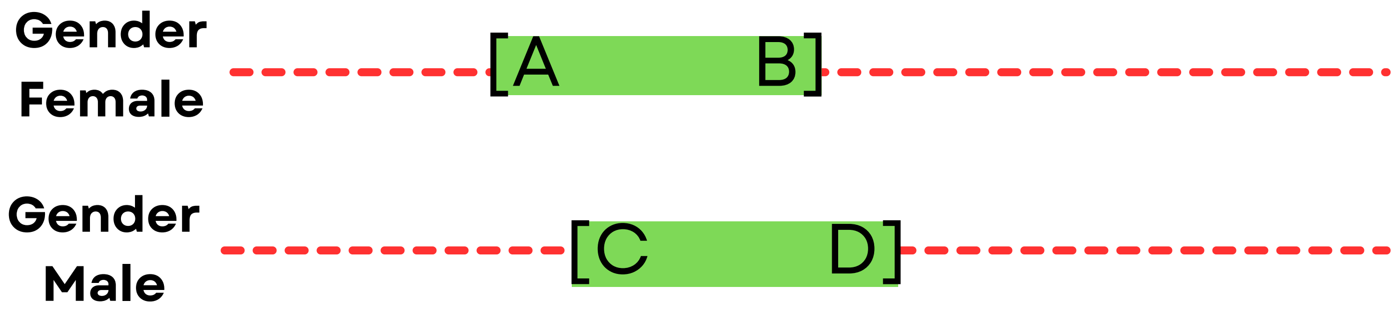

Representation is formalized with upper and lower **quotas** on each feature and value pair

Loss = maximum (relative) violation of any quota

Constraint 1: Predictable panel size *(for practicality and cost concerns)*

Desired Attributes for a Panel

Objective: Demographic representation *(offsets self-selection bias)*

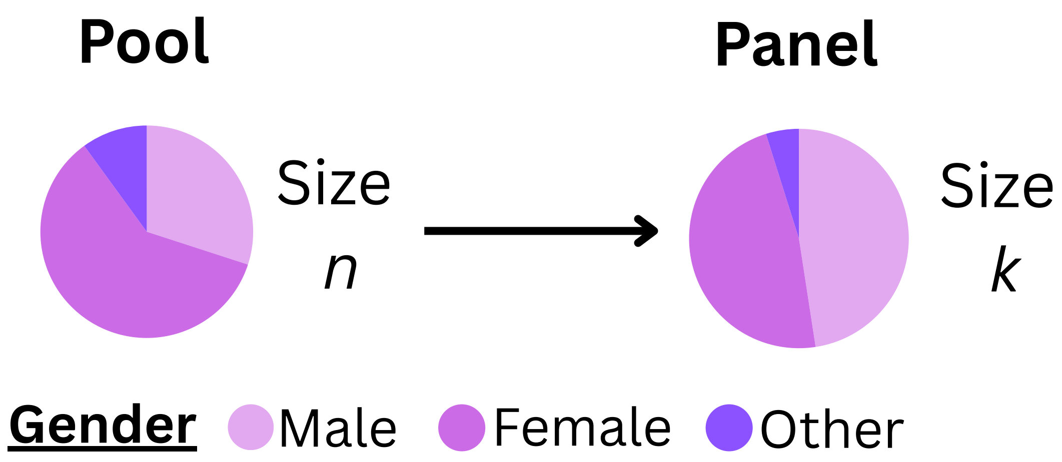

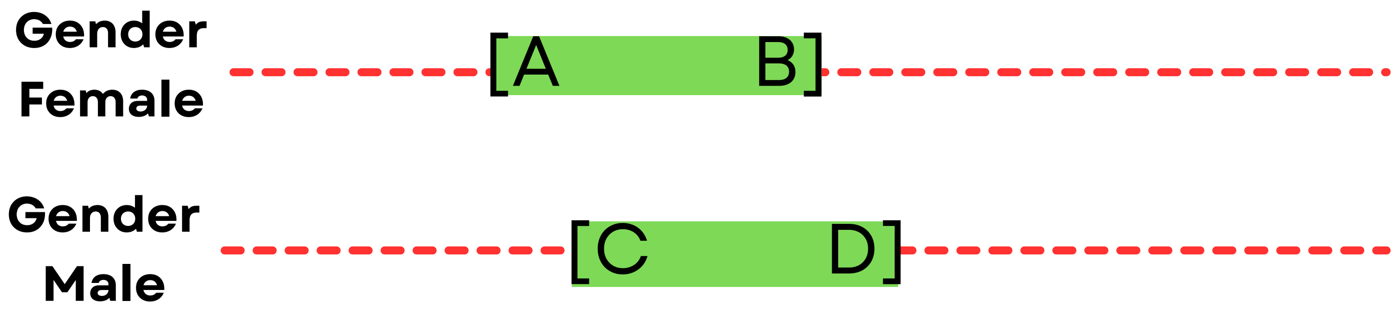

Representation is formalized with upper and lower **quotas** on each feature and value pair

<u>Loss = maximum (relative) violation of any quota</u>

Constraint 1: Predictable panel size *(for practicality and cost concerns)*

Constraint 2: limiting near 0/1 *selection probabilities prevents manipulation,* each probability must but be in some range $[\alpha, \beta]$

Problem: After constructing these assemblies, panelists drop out at the last-minute.

Problem: After constructing these assemblies, panelists drop out at the last-minute.

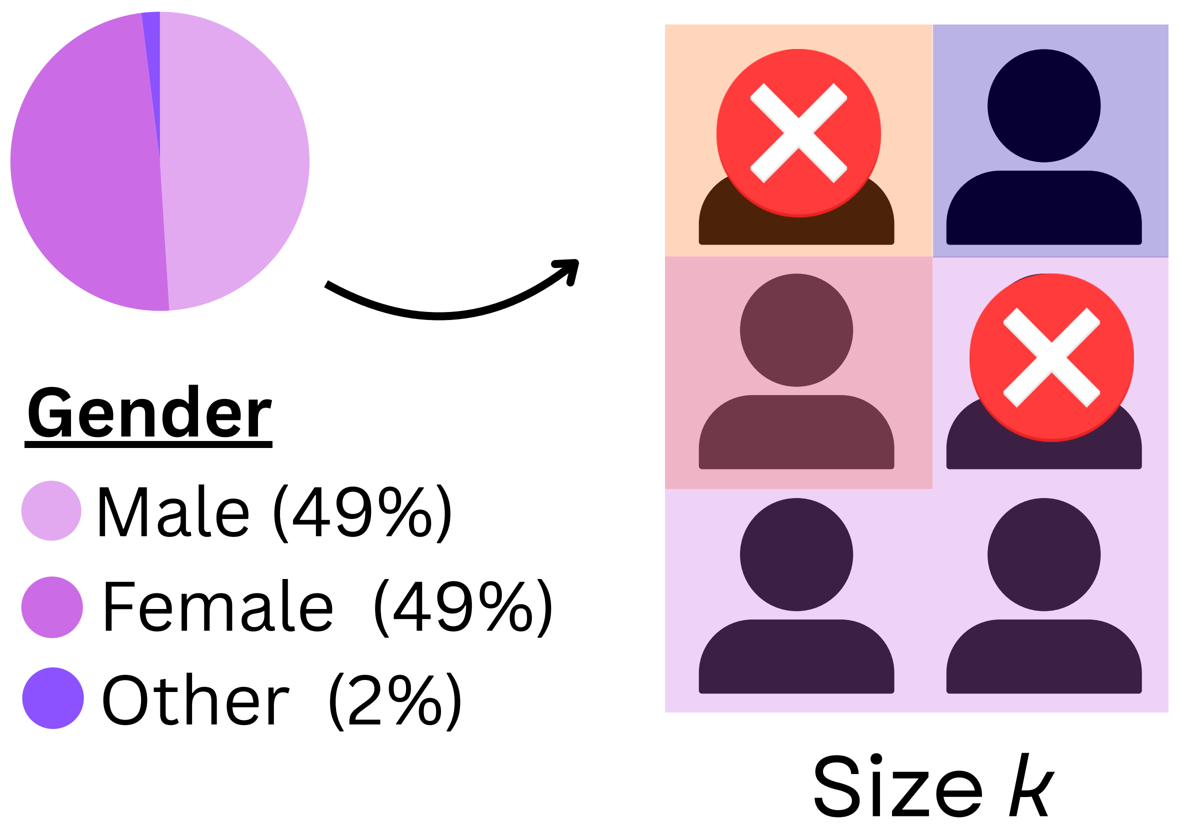

Problem: After constructing these assemblies, panelists drop out at the last-minute.

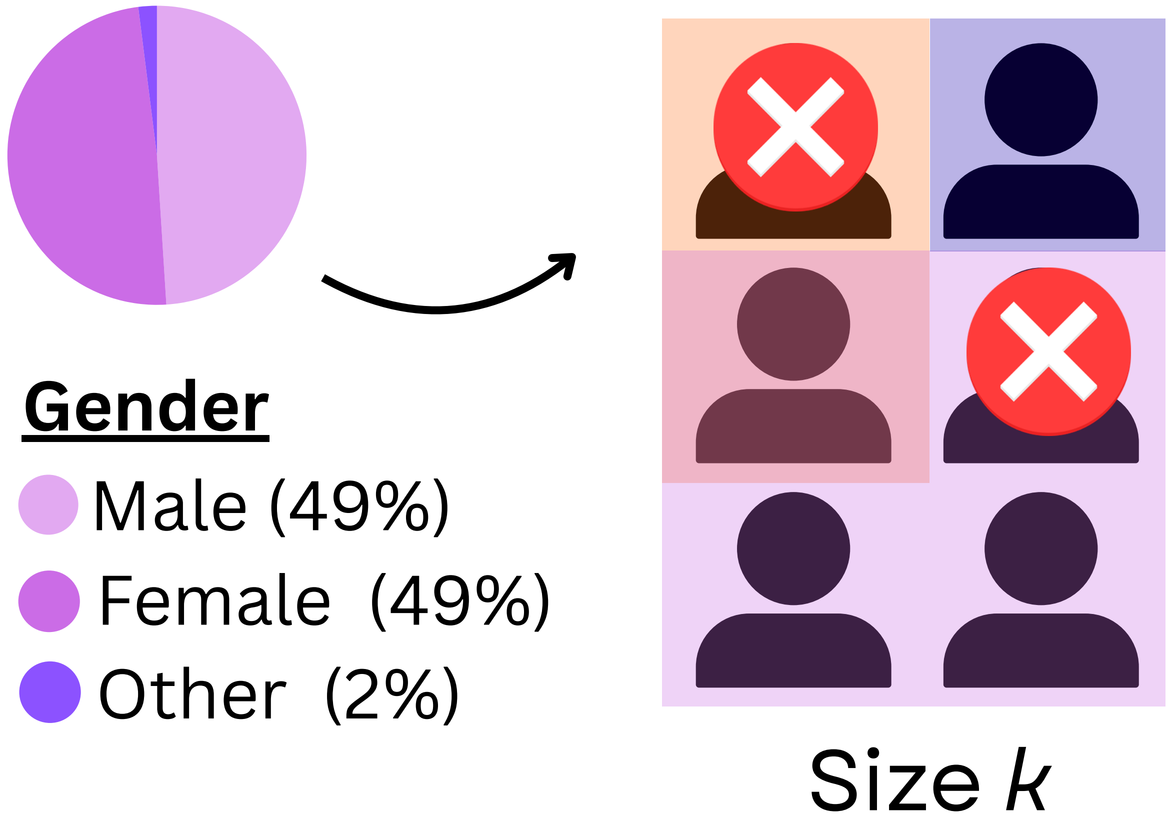

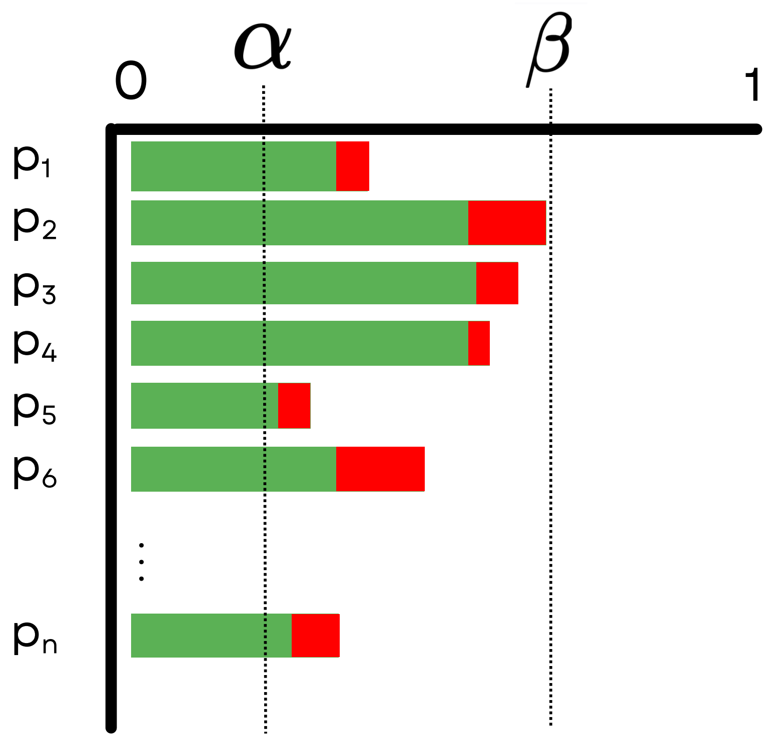

$$p_i = p_i' - d_i$$

Problem: After constructing these assemblies, panelists drop out at the last-minute.

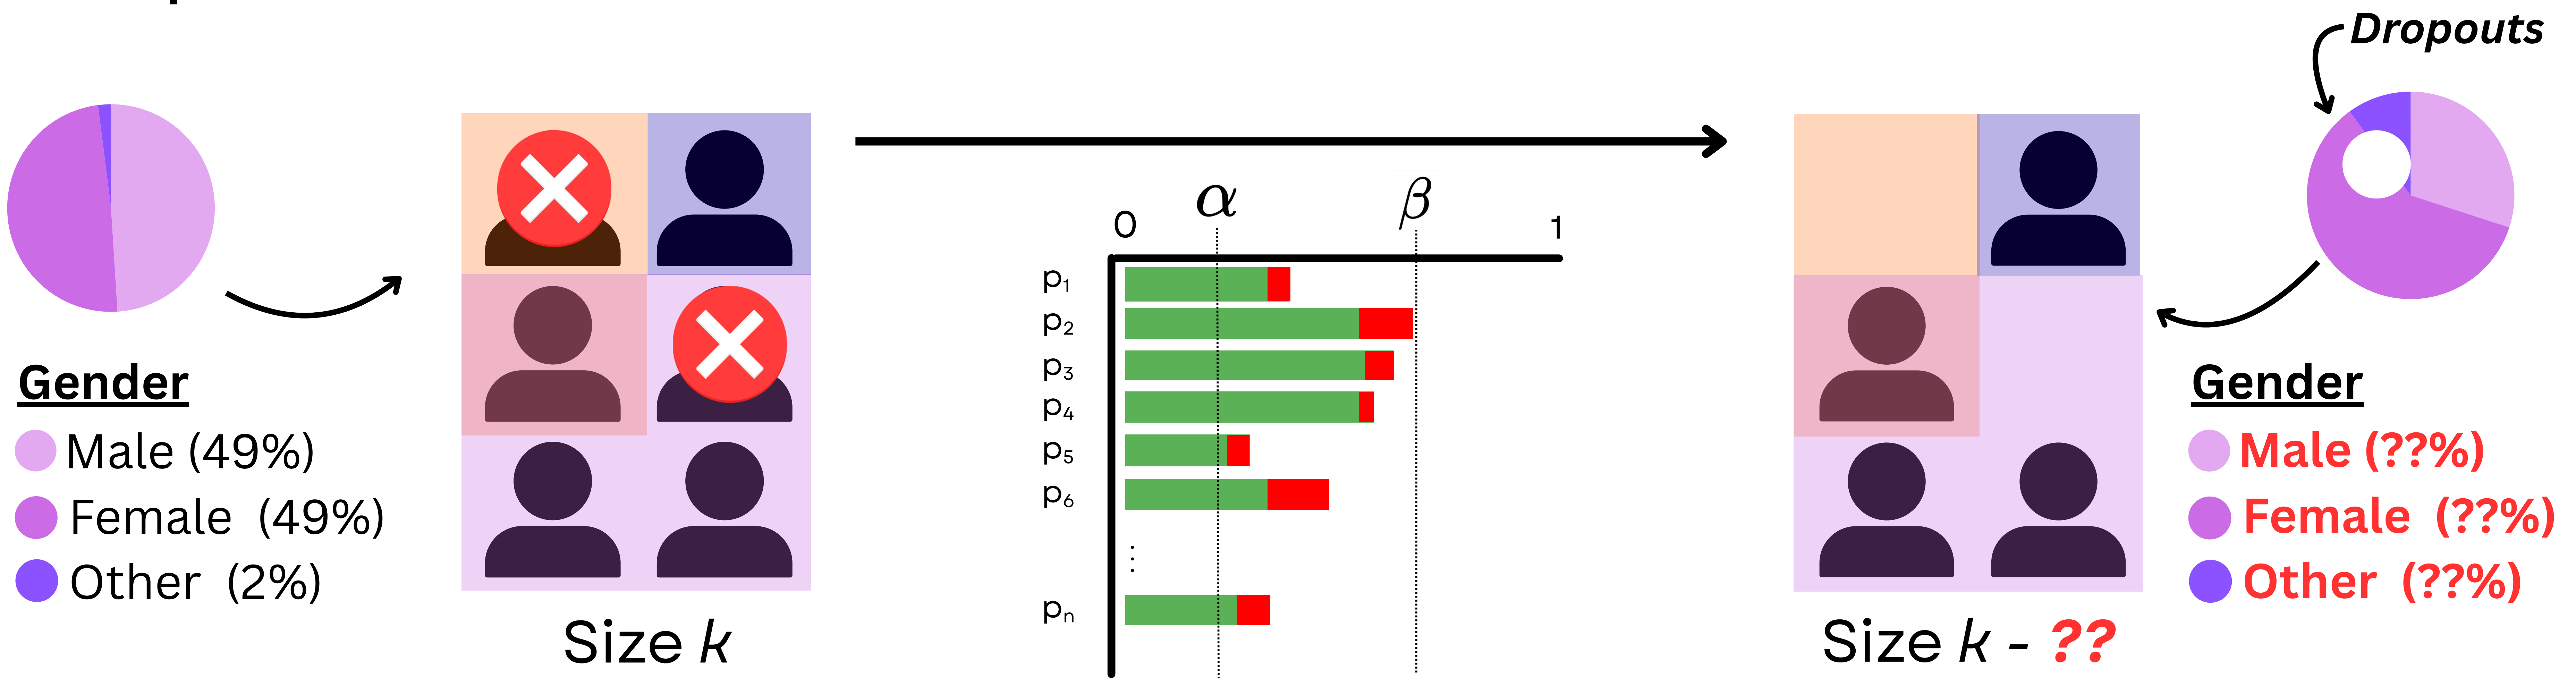

$$p_i = p'_i - d_i$$

Problem: After constructing these assemblies, panelists drop out at the last-minute.

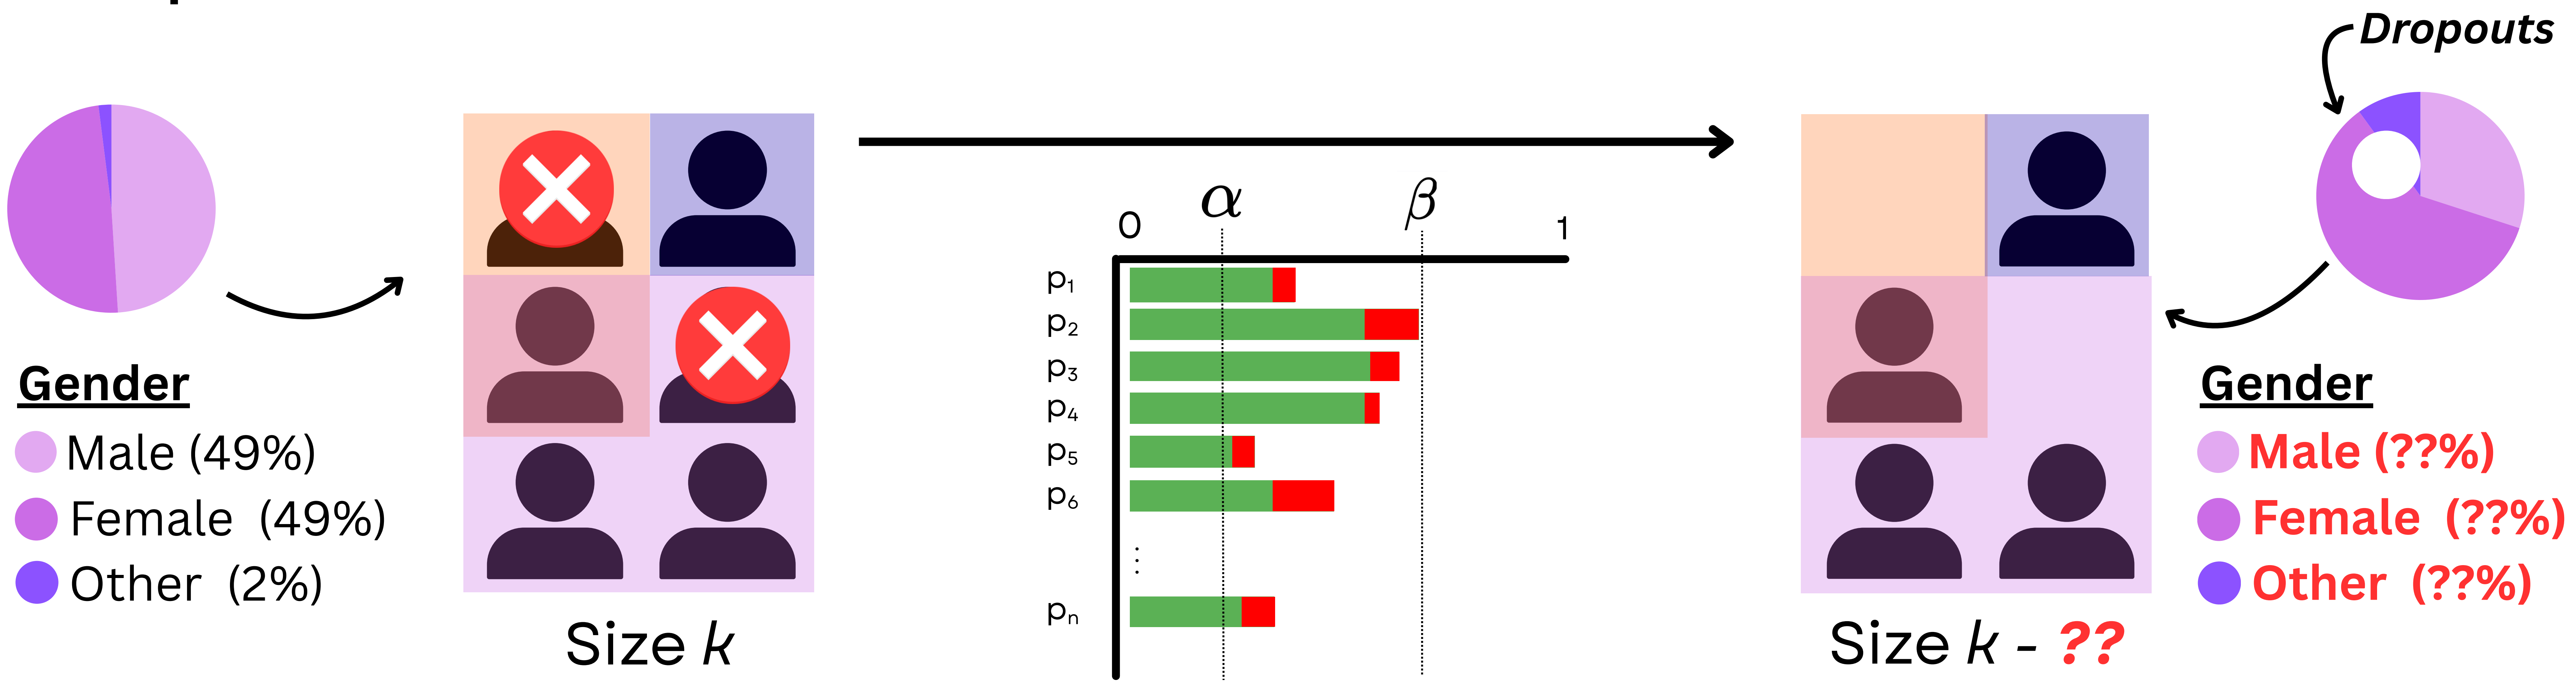

$$p_i = p_i' - d_i$$

How can we ensure our post-dropouts panel has our desired attributes?

We can construct good panels when we know dropout probabilities **d** *(Assos et. al, 2025),* but what if **d** is ***unknown*** or ***adversarial*****?**

We can construct good panels when we know dropout probabilities **d** *(Assos et. al, 2025),* but what if **d** is ***unknown*** or ***adversarial*?**

Say we have **noisy, *estimated* marginal dropout probabilities** $d_{estimated}$, and a **target robustness** γ

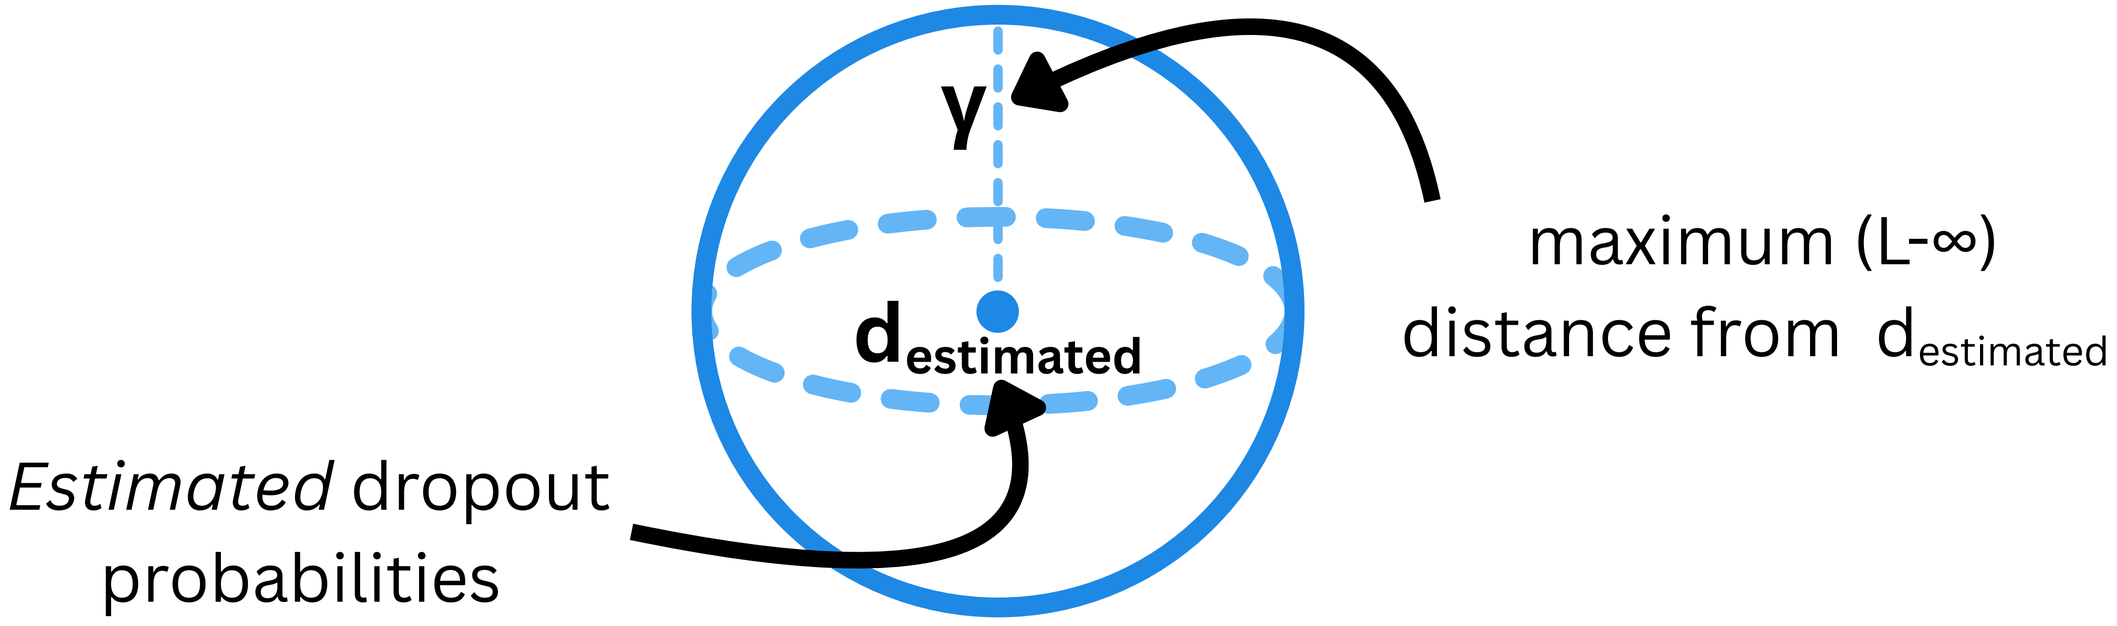

We can construct good panels when we know dropout probabilities **d** *(Assos et. al, 2025),* but what if **d** is ***unknown*** or ***adversarial*?**

Say we have **noisy, *estimated* marginal dropout probabilities** $d_{estimated}$, and a **target robustness** γ

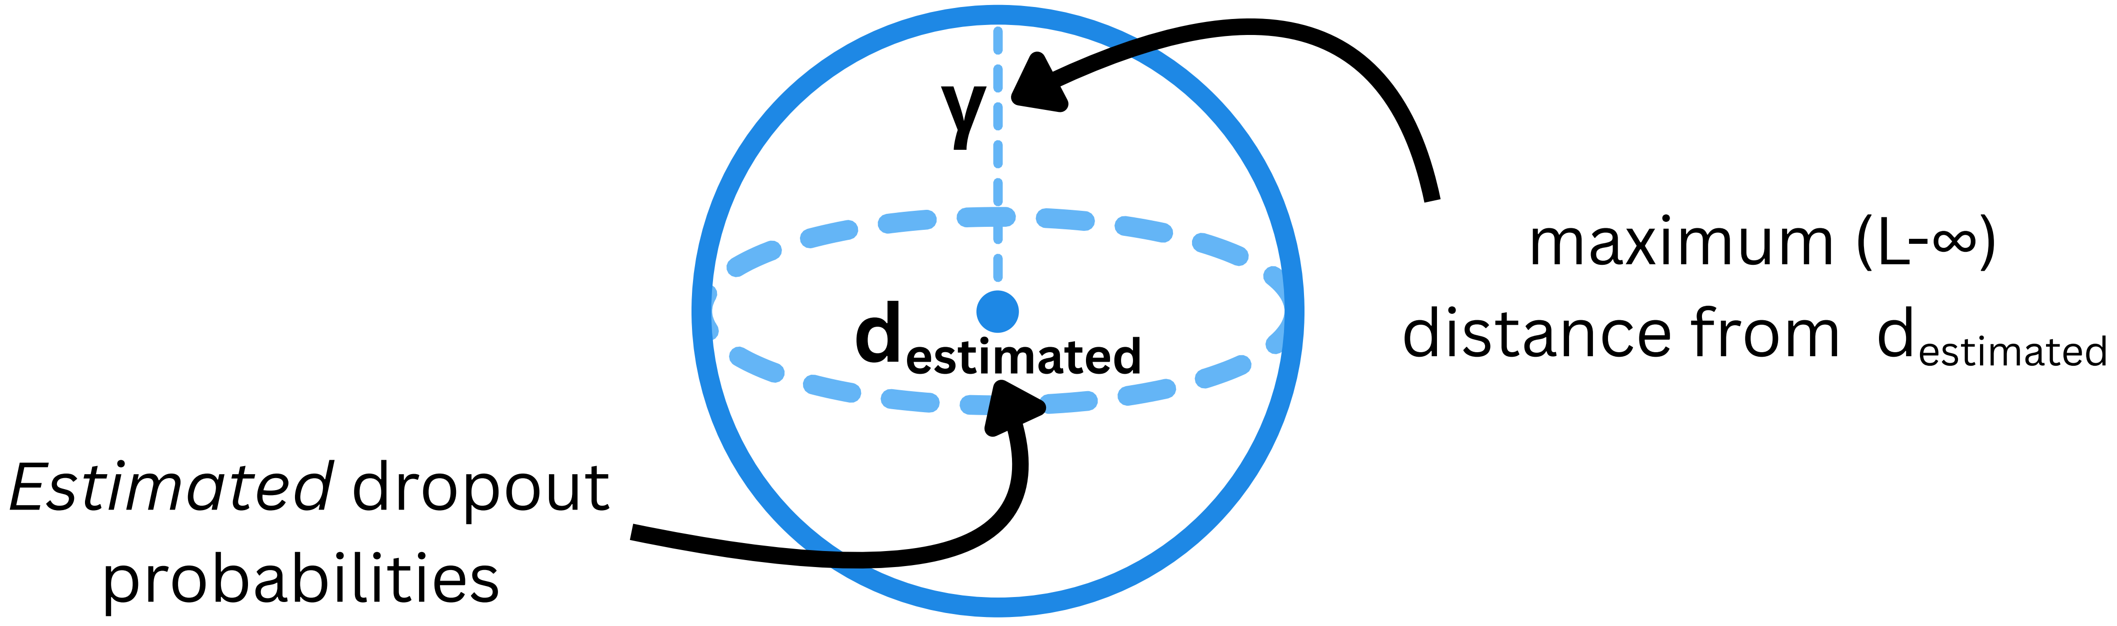

Research Question: Can we construct a panel robust to dropouts chosen by a constrained adversary?

Minimax Optimization Problem

$$\min_{\text{probabilities } \mathbf{p}} \; \max_{\text{distributions } \delta} \mathcal{L}(\mathbf{p}, \delta)$$

MIN chooses selection probabilities, MAX chooses a distribution over sets of dropouts.

Minimax Optimization Problem

$$\min_{\text{probabilities } \mathbf{p}} \max_{\text{distributions } \delta} \mathcal{L}(\mathbf{p}, \delta)$$

MIN chooses selection probabilities, MAX chooses a distribution over sets of dropouts.

Algorithm Sketch

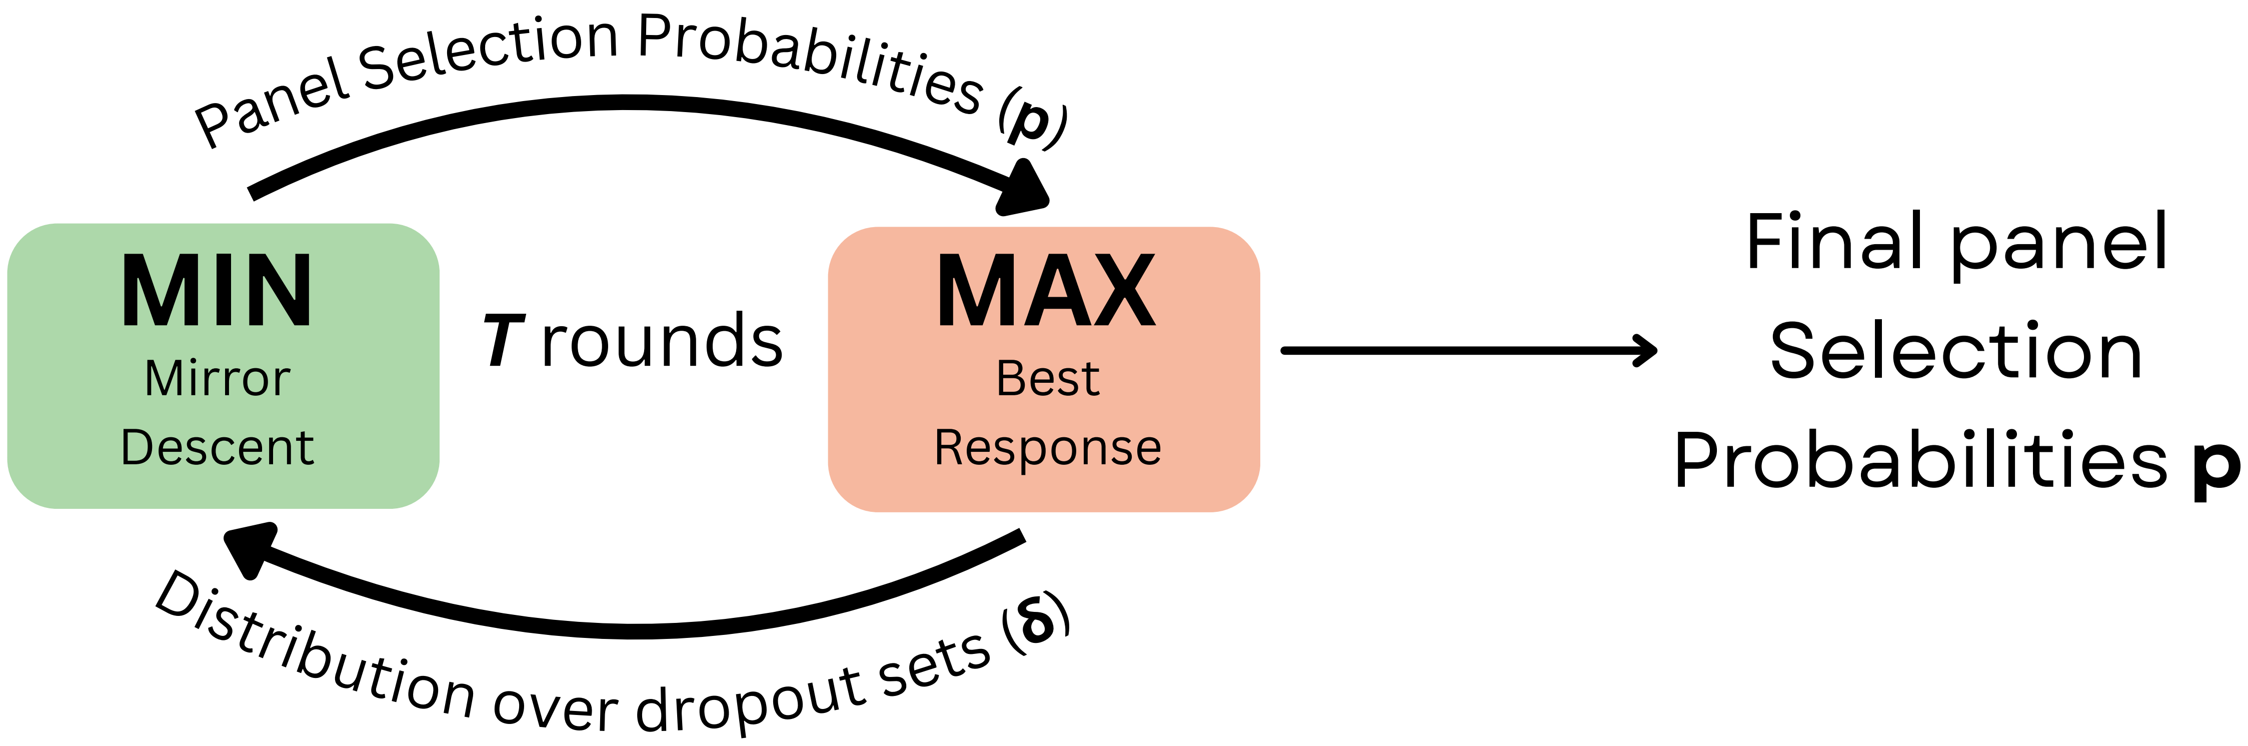

Minimax Optimization Problem

$$\min_{\text{probabilities } \mathbf{p}} \; \max_{\text{distributions } \delta} \; \mathcal{L}(\mathbf{p}, \delta)$$

MIN chooses selection probabilities, MAX chooses a distribution over sets of dropouts.

Algorithm Sketch

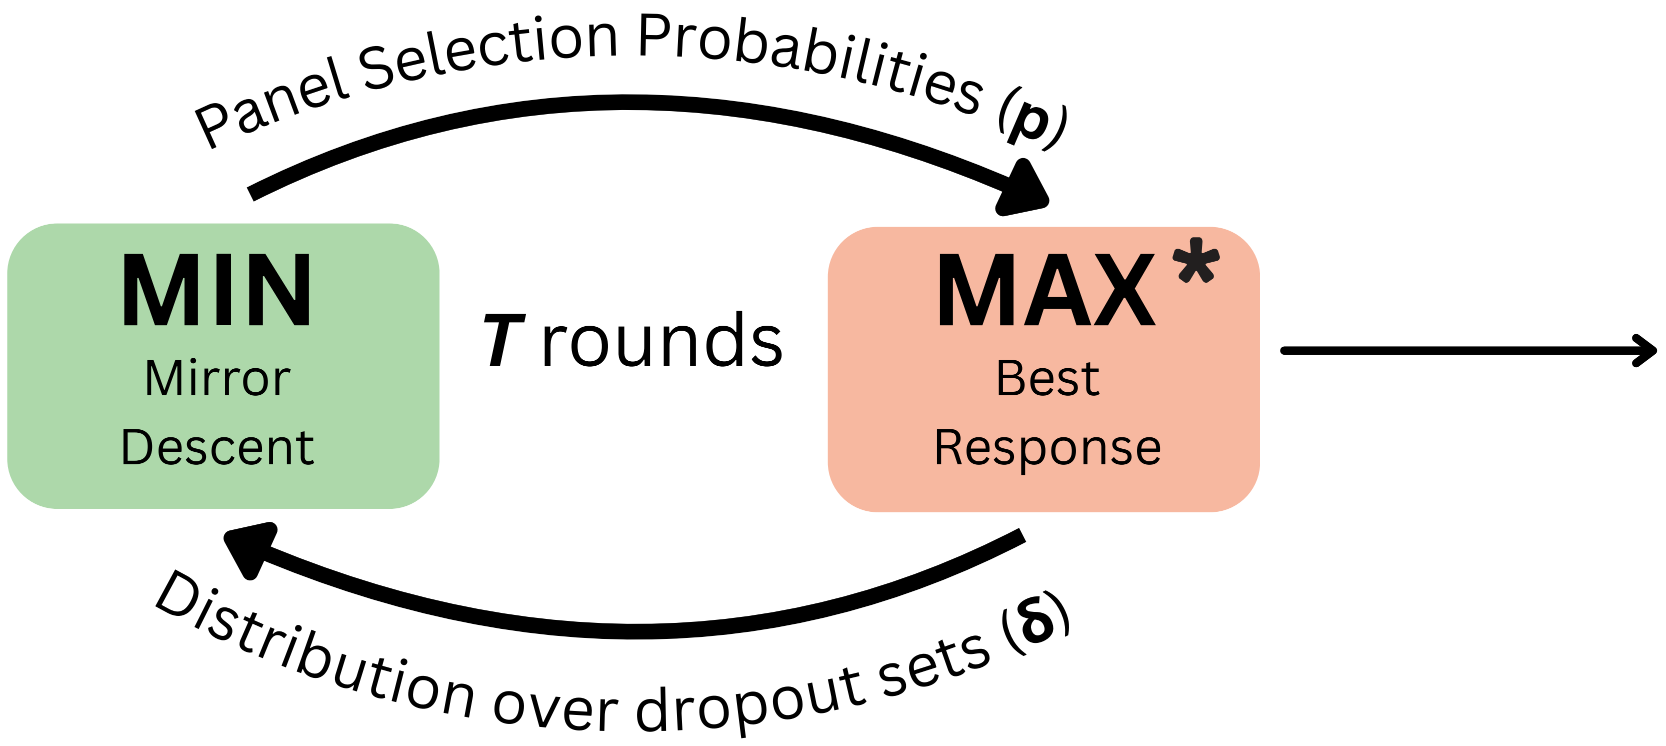

Final panel Selection Probabilities **p**

* δ is defined over distributions with exponentially large support, but we compute best response effciently

Minimax Optimization Problem

$$\min_{\text{probabilities } \mathbf{p}} \max_{\text{distributions } \delta} \mathcal{L}(\mathbf{p}, \delta)$$

MIN chooses selection probabilities, MAX chooses a distribution over sets of dropouts.

Algorithm Sketch

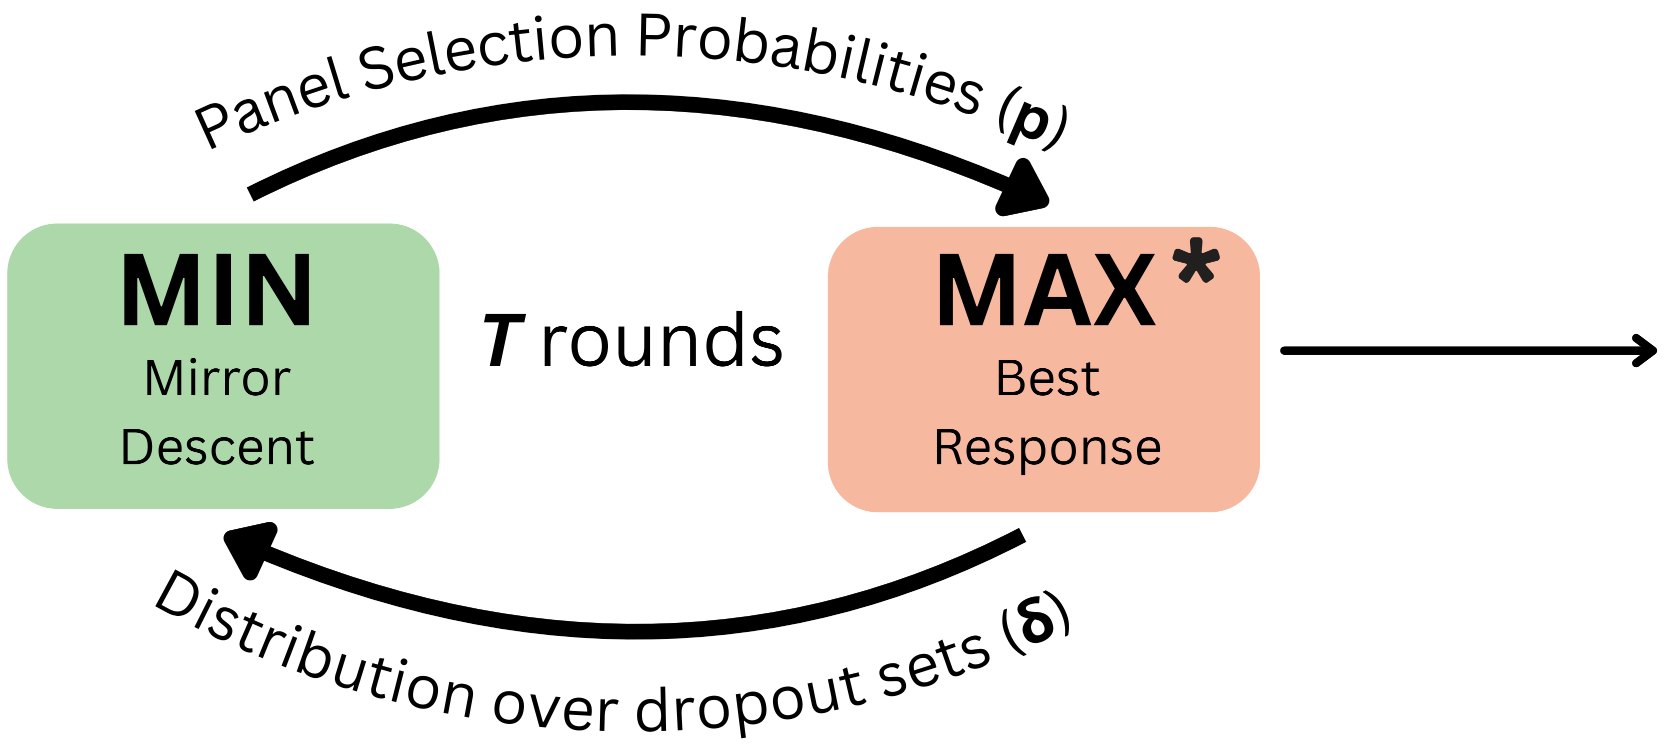

Final panel Selection Probabilities **p**

* δ is defined over distributions with exponentially large support, but we compute best response effciently

We can efficiently construct panels that minimize loss (demographic misrepresenation) after adversarial dropouts occur

Is it better to be overestimate or underestimate the adversary? *It is better to underestimate the strength of the adversary.*

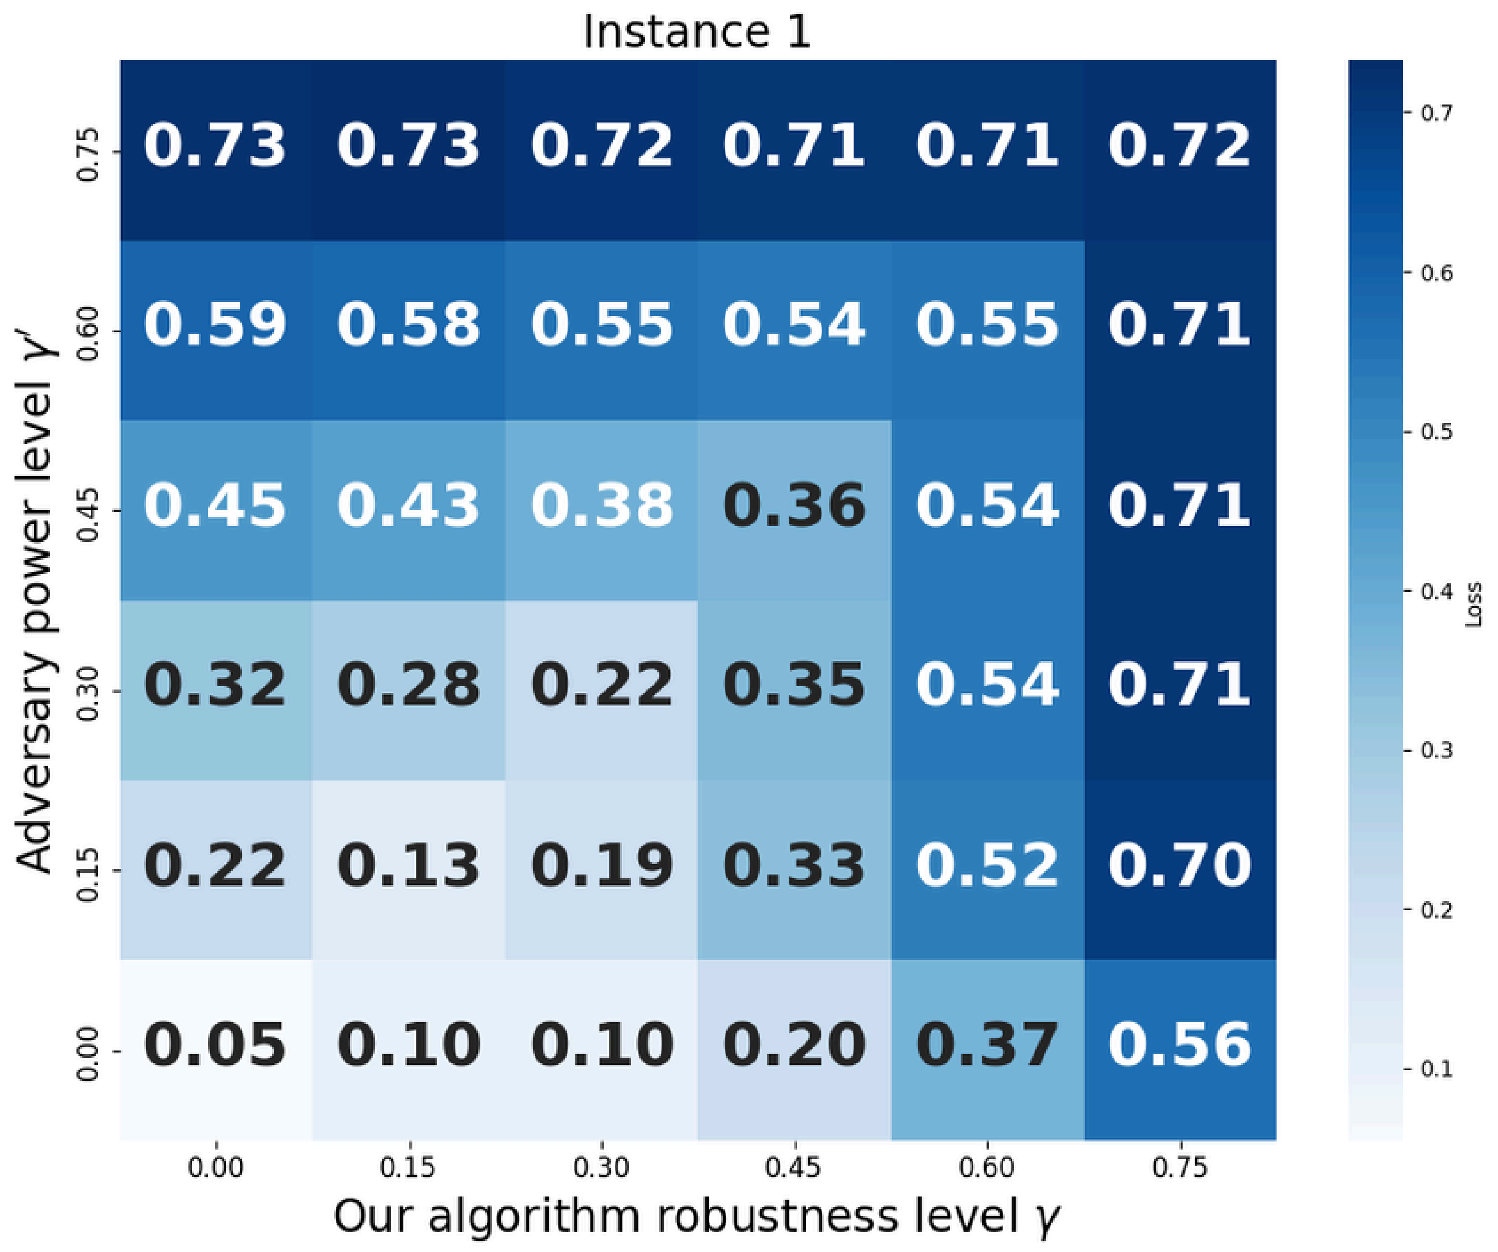

Experiments

Is it better to be overestimate or underestimate the adversary? *It is better to underestimate the strength of the adversary.*

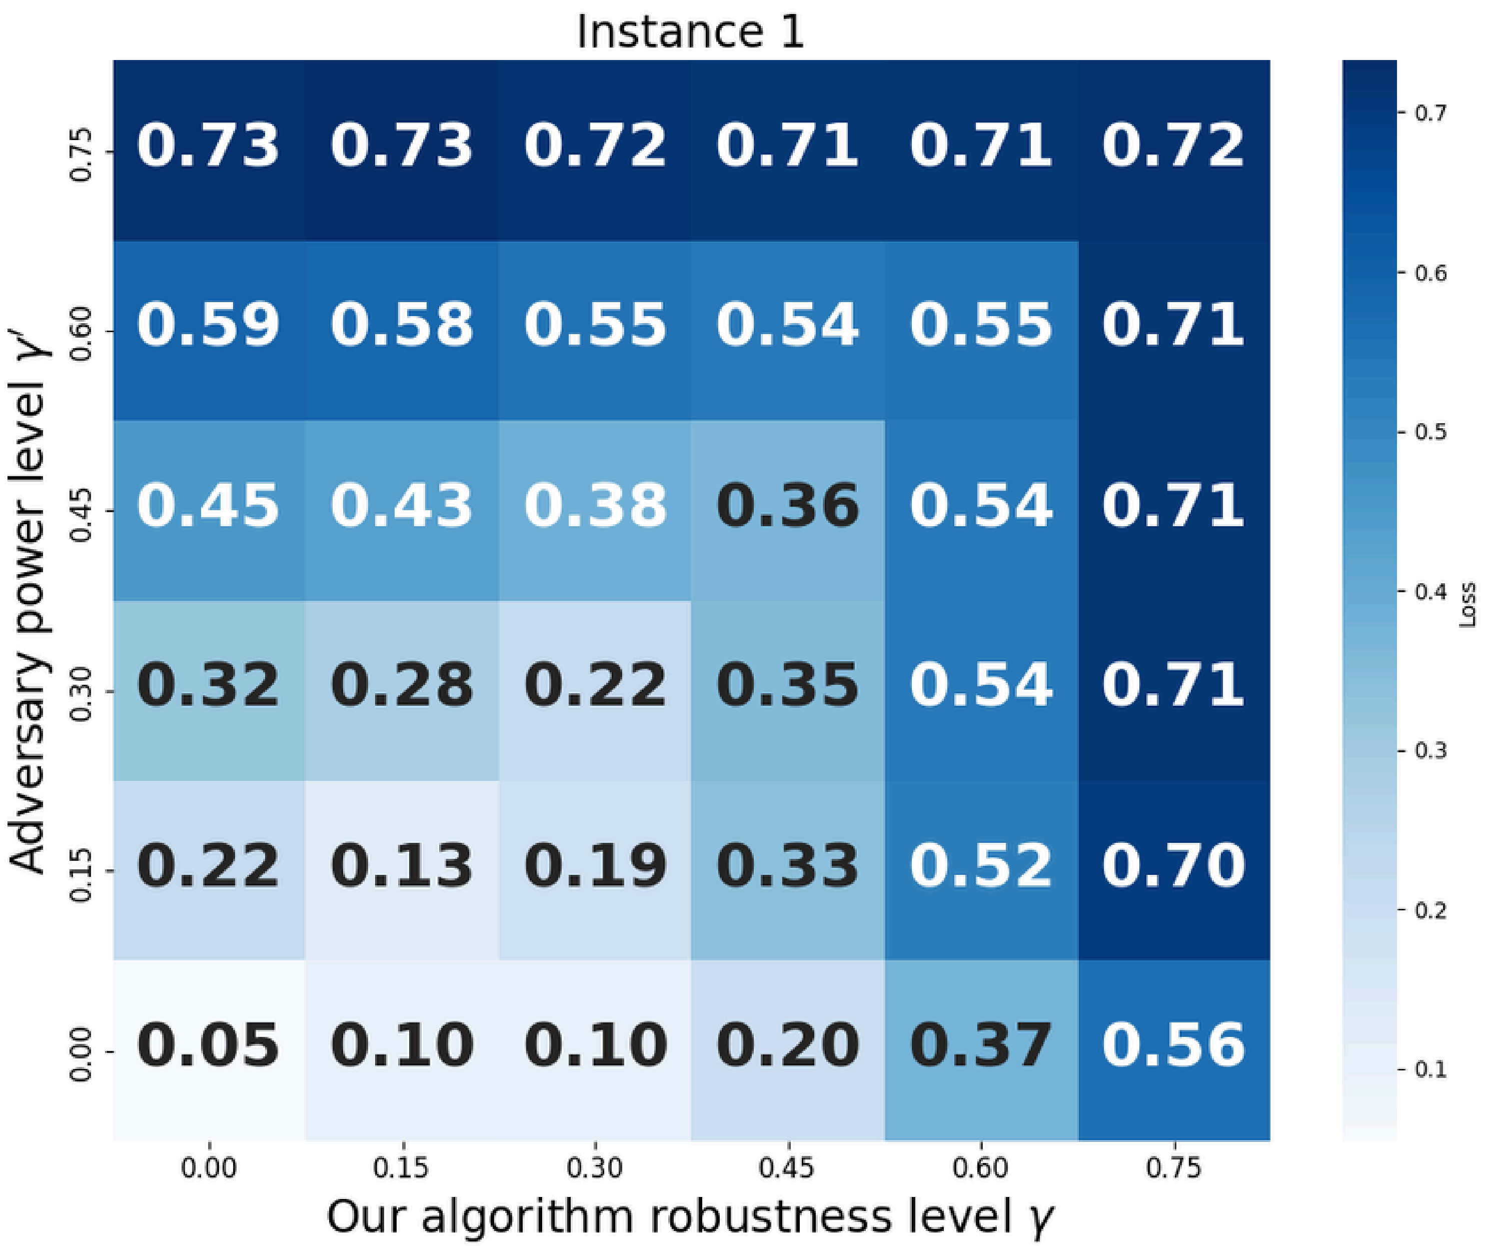

What is the loss trade off for constraining selection probabilties? *Constraint tightness varies (approximately) linearly with increases in loss.*

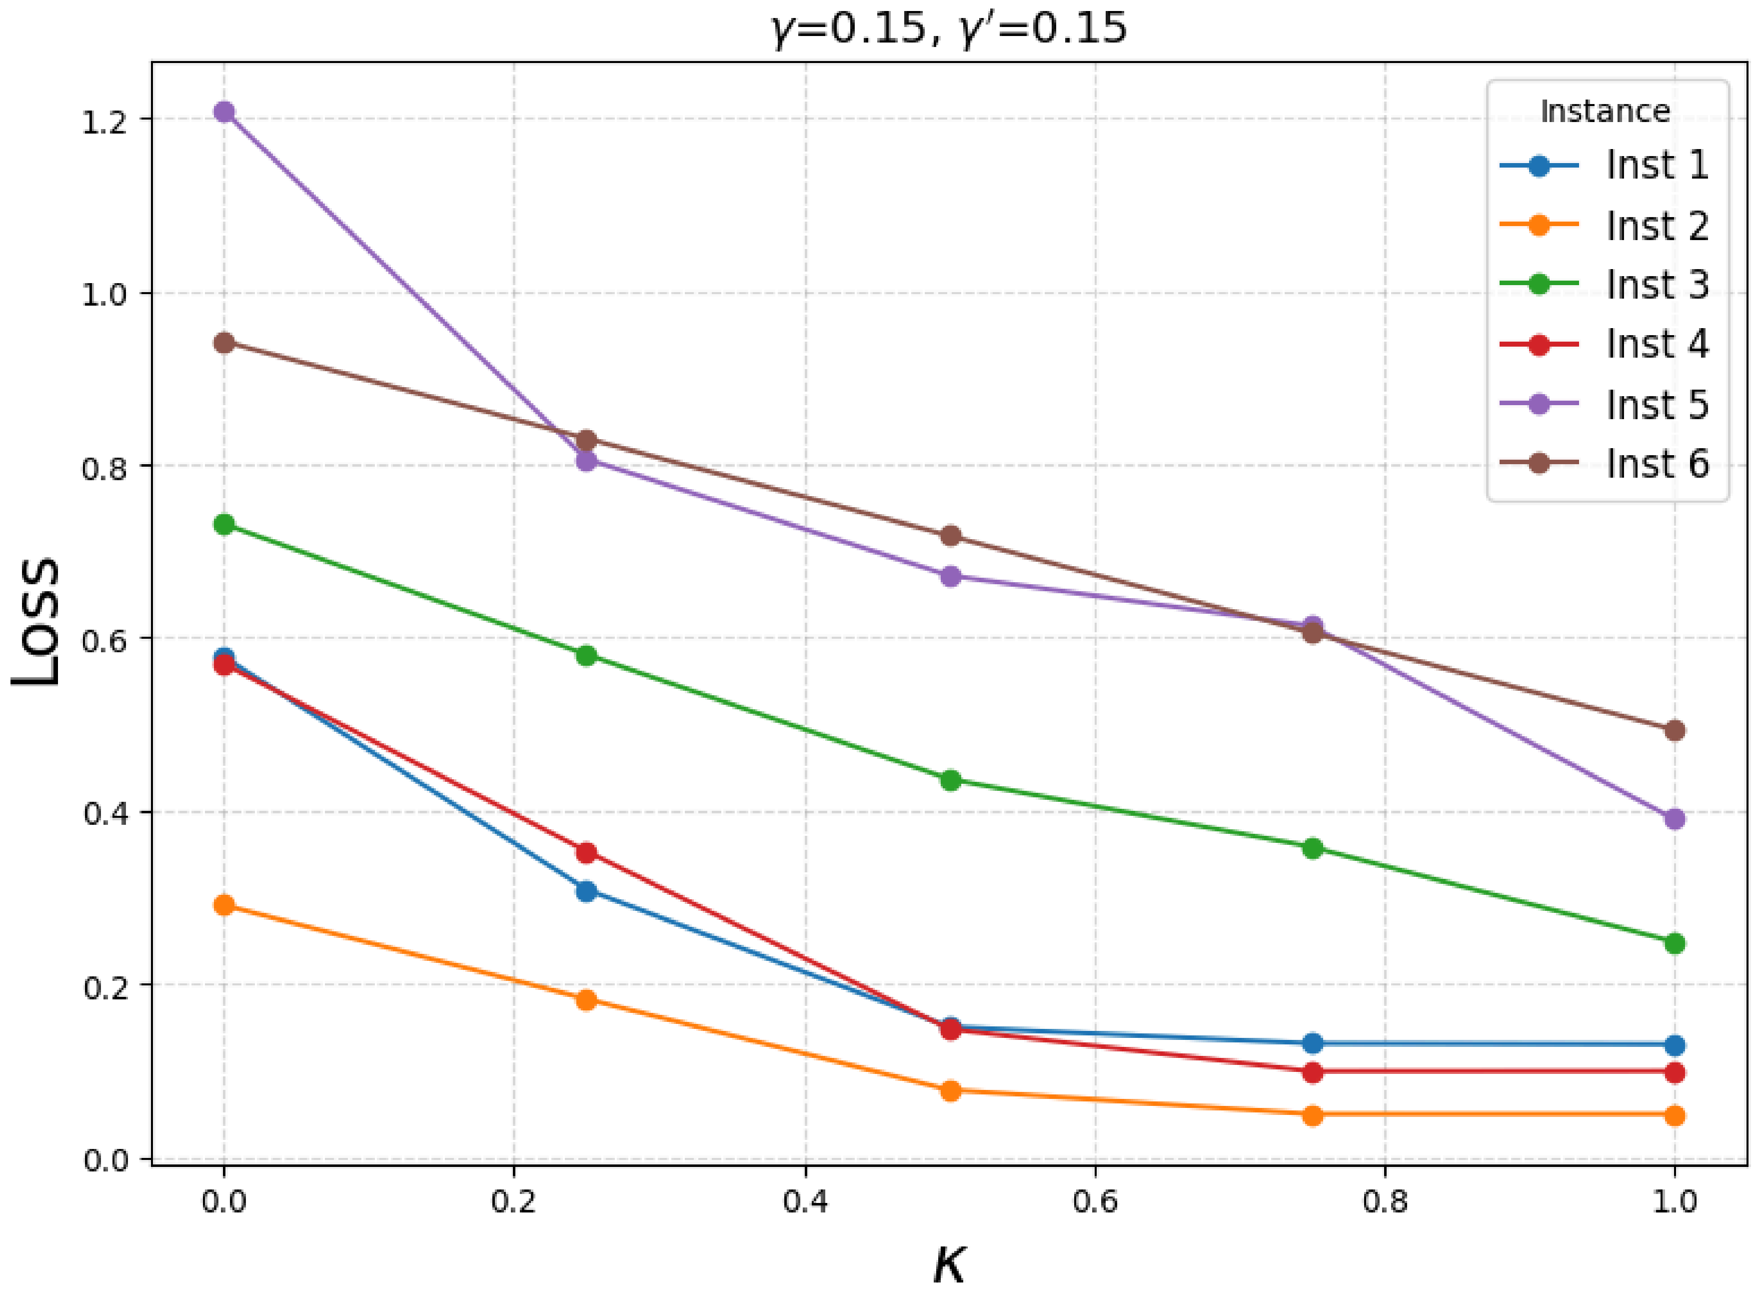

Thanks!

Poster #184

Read the full paper