Does Representation Guarantee Welfare?

HARVARD UNIVERSITY · PURDUE UNIVERSITY · NEURAL INFORMATION PROCESSING SYSTEMS

Jakob de Raaij*, Ariel D. Procaccia*, Alexandros Psomas**

*Harvard University, **Purdue University

Key Take-Aways

- The current standard of selecting panels for citizens' assemblies to satisfy 1-*representation* guarantees social welfare maximization for utilitarian welfare but fails for all other p-mean welfare functions.
- Panels that satisfy our stronger notion of 2-*representation* have significantly better guarantees on approximately maximizing social welfare.
- 2-representative panels can be found in real-world datasets.

Citizens' Assemblies

- A *citizens' assembly* is a group of everyday people who come together to discuss a policy issue, with the goal of informing decision makers.
- A successful assembly gains legitimacy by accurately representing the underlying population. In practice, many assemblies strive to achieve this through both:

randomly selecting the panel of citizens — ensuring *descriptive representation* of the population

- Citizens' assemblies have been on the rise around the world, with prominent assemblies convened by Australian, European, and North American governments, while also gaining prominence in AI governance.

Descriptive Representation

A panel of citizens satisfies *descriptive representation* if it resembles the composition of the population along pre-defined features.

Definition. *A panel is k-representative if any intersection of up to k of the features appears with (approximately) the same frequency in the panel and the population.*

Example: The pre-defined features are shape (☆ or ♡) and color (🟠 or 🟡). The population $\mathcal{P}$ consists of 30% ★(orange), 20% ♡(orange), 30% ☆(yellow), 20% ♡(yellow). Consider two panels:

Panel 1 ($\mathcal{C}_1$)	Panel 2 ($\mathcal{C}_2$)
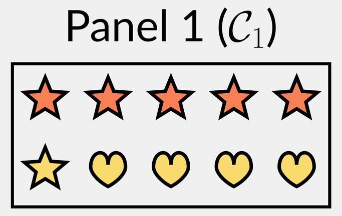	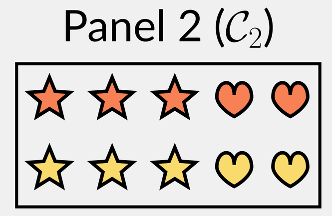
1-representation: ✓.	1-representation: ✓.
Shape: 60% ☆, 40% ♡	Shape: 60% ☆, 40% ♡
Color: 50% 🟠 and 50% 🟡.	Color: 50% 🟠 and 50% 🟡.
2-representation: ✗.	2-representation: ✓.
Shape×Color: 50% ☆(orange), 0% ♡(orange), 10% ☆(yellow) and 40% ♡(yellow).	Shape×Color: 30% ☆(orange), 20% ♡(orange), 30% ☆(yellow) and 20% ♡(yellow).

Problem Statement

- Prior work on algorithmically selecting a representative panels strives to find 1-representative panels.
- It is unclear if k-representative panels for $k \geq 2$ have better guarantees and are feasible in practice.

Question. For any integer k, does a k-representative committee making decisions to maximize its own welfare also optimize social welfare for the underlying population? If no, by how much is it off in the worst case?

Model: Assumptions on Utilities and Welfare

- For a given alternative a (i.e., a possible policy), the utility of an individual from a is a linear function of their feature values.
- The utilities of all individuals for any possible alternative a are in $[u_{\min}, 1]$.
- The welfare of a group $\mathcal{G}$ (the general population or the panel) for an alternative a is the *p-power mean* of the individual utilities, where $p \leq 1$:

$$u_{\mathcal{G}}(a) = \left(\frac{1}{|\mathcal{G}|}\sum_{i\in\mathcal{G}} u_i(a)^p\right)^{1/p} = \begin{cases} \frac{1}{|\mathcal{G}|}\sum_{i\in\mathcal{G}} u_i(a) & \text{if } p = 1 \quad \text{(utilitarian welfare)} \\ \sqrt[|\mathcal{G}|]{\prod_{i\in\mathcal{G}} u_i(a)} & \text{if } p \to 0 \quad \text{(Nash welfare)} \\ \min_{i\in\mathcal{G}} u_i(a) & \text{if } p \to -\infty \quad \text{(Rawlsian welfare)} \end{cases}$$

Example (continued): Consider a policy a favoring shape ☆ and color 🟡, giving utility 0.8 to ☆ and 0.1 to ♡, as well as utility 0 to 🟠 and 0.2 to 🟡. Thus, the utility of ★(orange) is $u_{★}(a) = u_{☆}(a) + u_{🟠}(a) = 0.8 + 0 = 0.8$, similarly $u_{♡(orange)}(a) = 0.1$, $u_{☆(yellow)}(a) = 1$, $u_{♡(yellow)}(a) = 0.3$.

Assume $p = 0$: $u_{\mathcal{P}}(a) = u_{\mathcal{C}_2}(a) \approx 0.46$ but $u_{\mathcal{C}_1}(a) \approx 0.55$. We find that panel $\mathcal{C}_1$ overestimates the welfare of a, which may lead to it making suboptimal decisions.

Results

We show that a panel optimizing its own welfare makes decisions that are always (approximately) optimal for social welfare if and only if the panel's welfare always is an accurate estimate of the social welfare.

Theorem 1. *For $p = 1$, 1-representative panels perfectly estimate social welfare.*

Theorem 2 (informal). *For $p < 1$, 2-representative panels estimate the social welfare of an alternative more accurately than 1-representative panels, with the relative improvement increasing when $u_{\min}$ increases.*

Please scan the QR Code for a link to the full paper including the formal definitions, proofs, full experiment, extended discussion of results and limitations, and list of references.

Comparison of 1- and 2-representative Panels

- We say k-representative panels are ε-accurate if their welfare never differs by more than ε from the population welfare, for a given p and $u_{\min}$.
- Figure 1a shows the tight worst-case ε-accuracy of 1- (dashed) and 2-representative (solid) panels as a function of p, for different values of $u_{\min}$. Figure 1b shows the ratio of ε for 1- and 2-representative panels as a function of $u_{\min}$ for different p.
- 2-representative panels are ε-accurate for much smaller ε. If utilities are generally high, the relative improvement is the highest: For example, for $u_{\min} = 4/5$ and Nash welfare, accuracy increases by *2700%* for 2-representative panels.

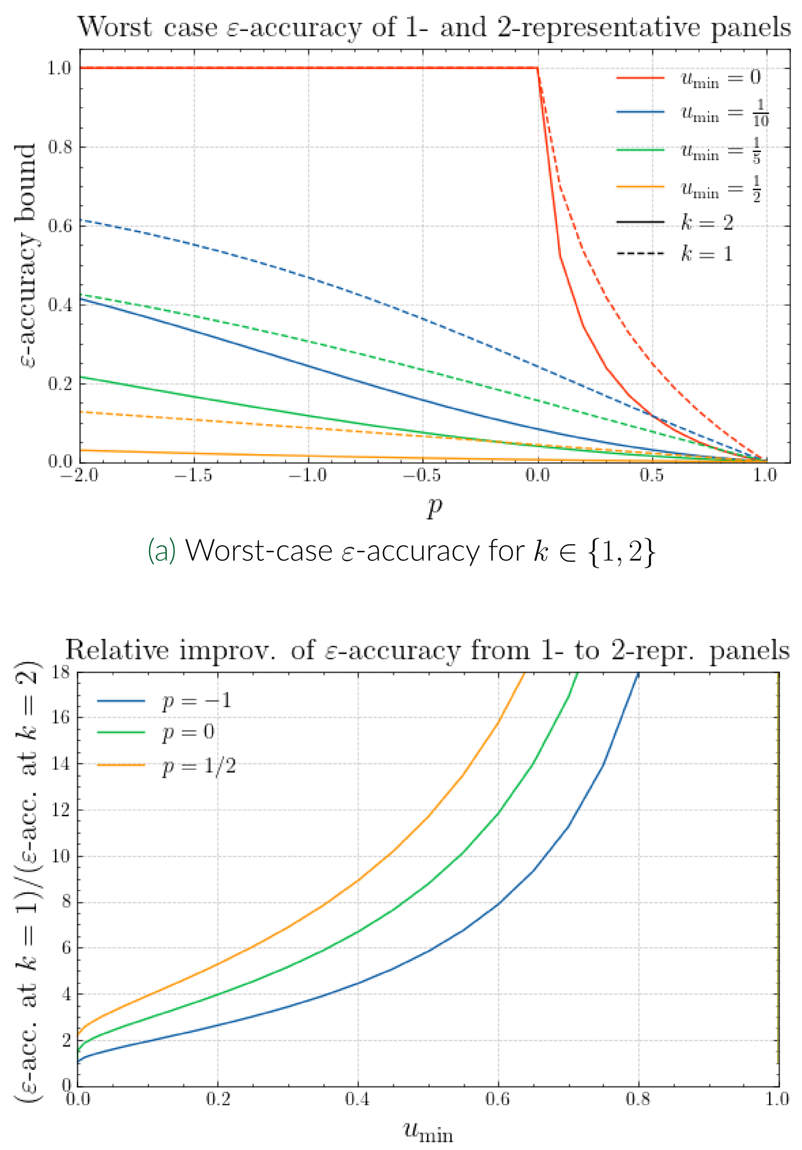

(a) Worst-case ε-accuracy for $k \in \{1, 2\}$

(b) Ratio of ε-accuracy bounds for $k \in \{1, 2\}$

Figure 1. Comparisons of ε-accuracy bounds.

Experiments: Existence of k-representative Panels

- The panels of citizens' assemblies are usually chosen from a *pool of volunteers*, which may not be representative of the population.
- **Question:** For volunteer pools of real citizens' assemblies, what is the largest panel we can select that is k-representative, up to rounding feature intersection quotas up or down to the next integer?
- We examined 4 datasets: 1 from a Western European nationwide assembly and 3 from Australian state-wide assemblies. In all cases, 2-representative panels of the desired size or close to it exist. The results for two assemblies are plotted in Figure 2.

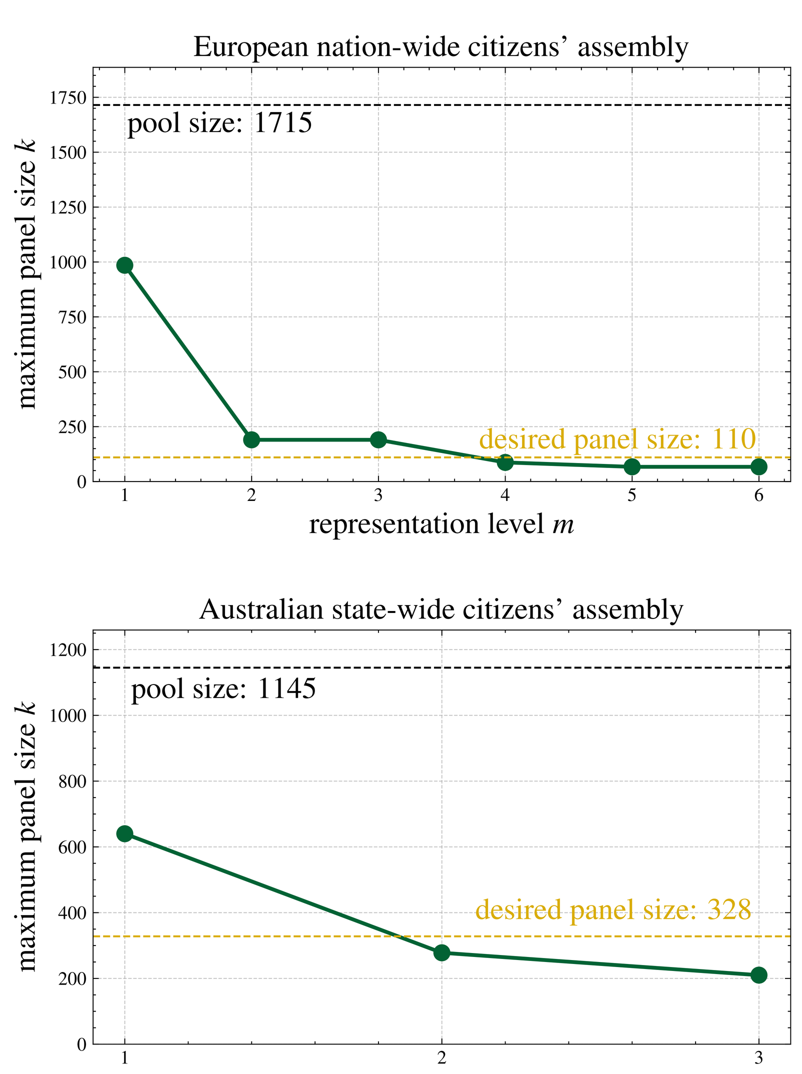

Figure 2. The maximum size of k-representative panels that can be found in the pool for different k is shown in green. The dashed yellow line is the desired panel size.