

Representative Social Choice: From Learning Theory to AI Alignment

TIANYI QIU, Peking University, China

Social choice theory is the study of preference aggregation across a population, used both in mechanism design for human agents and in the democratic alignment of language models. In this study, we propose the *representative social choice* framework for the modeling of democratic representation in collective decisions, where the number of issues and individuals are too large for mechanisms to consider all preferences directly. These scenarios are widespread in real-world decision-making processes, such as jury trials, legislation, corporate governance, and, more recently, language model alignment. In representative social choice, the population is *represented* by a finite sample of individual-issue pairs based on which social choice decisions are made. We show that many of the deepest questions in representative social choice can be formulated as statistical learning problems, and prove the generalization properties of social choice mechanisms using the theory of machine learning. We further formulate axioms for representative social choice, and prove Arrow-like impossibility theorems with new combinatorial tools of analysis. Our framework introduces the representative approach to social choice, opening up research directions at the intersection of social choice, learning theory, and AI alignment.

JAIR Track: AI and Society

JAIR Associate Editor: Toby Walsh

JAIR Reference Format:

Tianyi Qiu. 2026. Representative Social Choice: From Learning Theory to AI Alignment. *Journal of Artificial Intelligence Research* 87, Article 13 (October 2026), 21 pages. DOI: [10.1613/jair.1.17524](https://doi.org/10.1613/jair.1.17524)

1 Introduction

Social choice theory is a field of study that deals with the aggregation of individual preferences to form a collective decision. It has been applied in domains such as economics (Feldman and Serrano 2006), political science (Coleman and Ferejohn 1986; Miller 1983), and computer science (Conitzer, Freedman, et al. 2024), to name a few. In these applications, the goal is to design mechanisms that aggregate individual preferences in a way that satisfies certain desirable properties, most especially fairness.

However, existing models in social choice theory tend to be simplistic and rely on relatively strong assumptions, including (1) *independent, single-issue choices*, and (2) *complete information on all preferences of all individuals*. In practice, these assumptions are often violated. In common large-scale elections, candidates have partisan policies that are correlated across a huge number of different issues, and it is also infeasible to collect preferences of all voters on all issues, due to the large number of issues and voters involved.

These problems are not merely practical details that can be ignored. They are fundamental to the theory of social choice itself, since these complexities are exactly what give rise to democratic *representation* — the idea that individuals can delegate their decision-making power to a small number of representatives who can make decisions on their behalf, when there are too many issues and too many individuals to consider all preferences directly. The introduction of representation leads to fundamental questions for social choice theory that are not well-understood, such as the problem of *generalization*: how can we ensure that the decisions made by the

Author's Contact Information: Tianyi Qiu, ORCID: 0009-0003-4554-3201, qutianyi.qty@gmail.com, Peking University, Beijing, China.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2026 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.17524](https://doi.org/10.1613/jair.1.17524)

representatives are representative of the population's preferences, when the representatives are only chosen based on a small number of individuals' opinions on a small number of issues?

In this paper, we propose a new framework for social choice theory that models these complexities, which we call *representative social choice*. As we will show in the following sections, many of the deepest questions in representative social choice can be formulated as statistical learning problems. This connection allows us to leverage the rich theory of statistical machine learning to formulate axioms and mechanisms for representative social choice, and to analyze their properties. While some work studies winner determination for a single election from voter samples (Bhattacharyya and Dey 2021), our framework addresses the aggregation of preferences over a profile of many issues, making generalization over both individuals and issues a central challenge.

Applications. The representative model can be applied in the modeling of all *collective decisions involving representation and delegation*, which include most real-world decision-making processes, such as:

- *Jury Trials*: Citizens delegate their decision-making power to a randomly selected jury.
- *Legislation Processes*: Legislature body members, which can be viewed as representative samples of the citizen population, delegate the population's legal decisions to a set of laws that they decide on. We have a multi-issue setting, since the laws cover a wide range of issues. The space of possible collective preference profiles is the space of profiles implementable by laws.
- *Corporate Governance*: Shareholders elect a board of directors to make decisions on their behalf. We have a multi-issue setting, since the board decides on a wide range of issues during its service.
- *AI Alignment*: Frontier AI systems, including large language models (LLMs), undergo the *alignment* process during training, where they are trained to make decisions that are aligned with human values (Bai, Jones, et al. 2022; Bai, Kadavath, et al. 2022), using preference datasets sampled from humans. Alignment can be viewed as a form of representative social choice, where the LLM is trained to make decisions that are representative of the human population's preferences, the latter represented by individual-issue pairs sampled from human evaluators (*i.e.*, the preference dataset). The space of possible collective preference profiles is the space of profiles that can be actualized as LLM policies, a *feature space* that has been the subject of much research in the field of statistical learning theory (Vapnik 1999).

2 Related Work

Social choice theory has had a long history (Nisan and Ronen 1999; Satterthwaite 1975; Young 1975), with more recent research studying its intersection with machine learning (Fish et al. 2023; Parkes and Procaccia 2013), and its applications in AI alignment (Conitzer, Freedman, et al. 2024; Ge, Halpern, et al. 2024; Huang et al. 2024; Klingefjord et al. 2024; Köpf et al. 2024; Mishra 2023; Prasad 2018). This paper contributes to all three fronts, by extending the study of social choice theory to include democratic representation, using tools from the theory of machine learning, and with important applications in both human society and AI alignment.

Social Choice Theory. Social choice theory studies the aggregation of individual preferences to form a collective decision. The founding work of the field (Arrow 1951) proved the famous *Arrow's impossibility theorem*, stating that no social choice mechanism can satisfy the axioms of *unrestricted domain*, *non-dictatorship*, *Pareto efficiency*, and *independence of irrelevant alternatives*. Many extensions of Arrow's theorem have been proposed, including the Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975) which introduces strategy-proofness to the analysis. Beyond impossibility results, social choice theory also studies concrete mechanisms for preference aggregation, such as *voting rules* (Taylor 2005), *scoring rules* (Young 1975), and *judgment aggregation* (List 2012), while featuring intersections with other fields such as mechanism design (Nisan and Ronen 1999) and computer science, the latter resulting in the study of computational social choice (Brandt et al. 2016). Here, we extend the study of social choice to the representative setting, where issues and individuals are too numerous for all

preferences to be considered directly. By doing so, we formalize the problem of democratic representation within a statistical learning framework, providing a learning-theoretic lens for social choice.

Machine Learning Methods in Social Choice Theory. Machine learning methods have been applied to social choice theory in a variety of ways. *Preference learning* studies the problem of learning preference models from observed data, with deep connections between maximum likelihood estimation (MLE) and random utility models used in social choice (Azari et al. 2012; Conitzer, Rognlie, et al. 2009). More recently, studies on the learnability of utility functions from pairwise comparison queries have emerged (Ge, Juba, et al. 2024), further bridging the gap between learning theory and preference aggregation. In the domain of open-ended outcomes, *generative social choice* applies machine learning to handle settings where the outcome space is not fixed in advance (Fish et al. 2023), and *dynamic social choice* uses reinforcement learning models to address evolving preferences (Parkes and Procaccia 2013). Additionally, social choice theory has developed its own approaches to settings with incomplete preference information, where not all voters express preferences over all alternatives (Halpern et al. 2026). Our work differs from these lines of research in that we focus on the *representative* setting, where a mechanism must generalize from a finite sample of individual-issue pairs to the full population, and we leverage VC theory and Rademacher complexity to provide formal generalization guarantees, a perspective that is complementary to preference learning and aggregation under incomplete information.

Applications of Social Choice in AI Alignment. Social choice theory has been applied to the study of AI alignment, where the goal is to design AI systems that make decisions that are aligned with human values. Current approaches to AI alignment involves the aggregation of human preferences, and social choice-based algorithms (Klingefjord et al. 2024; Köpf et al. 2024), experimental studies (Huang et al. 2024), conceptual frameworks (Conitzer, Freedman, et al. 2024; Mishra 2023; Prasad 2018; Zhi-Xuan et al. 2024), and axiomatic frameworks (Ge, Halpern, et al. 2024) have been proposed to address the problem. In this paper, we extend social choice theory to the representative setting, which can then be applied to the AI alignment setting, where human preferences are represented by a collection of individual-issue pairs sampled from human evaluators, and the AI system is trained to make decisions representative of the human population’s preferences.

3 Problem Settings of Representative Social Choice

In this section, we present the formal definitions of the representative social choice problem, generalizing the standard social choice formalism to include more realistic complexities.

Issues. We consider a discrete (but possibly infinite) set of N -ary issues $\mathcal{I}$, where each issue $i \in \mathcal{I}$ (e.g., in a given state or province, which construction project to launch this year?) comes with N outcomes $[N] = \{1, 2, \dots, N\}$. Each individual’s preference profile can therefore be represented as a mapping from $\mathcal{I}$ to $\text{LO}(N)$, where $\text{LO}(N)$ is the set of all linear orders over the N outcomes.

We define a *saliency distribution* $\mathcal{D}_{\mathcal{I}}$ with full support over $\mathcal{I}$, which represents the importance of different issues, and decides the probability of each issue being sampled in the representation process. If there are a finite number of equally important issues, then $\mathcal{D}_{\mathcal{I}}$ is the uniform distribution over $\mathcal{I}$. In practice, $\mathcal{D}_{\mathcal{I}}$ could be estimated from public opinion polls, or in AI alignment, from the empirical distribution of user prompts.

Population. We consider a possibly infinite population, represented by a distribution $\mathcal{D}_{\mathcal{P}} \in \Delta[\mathcal{P}]$ supported on $\mathcal{P} \subseteq \text{LO}(N)^{\mathcal{I}}$.¹ For any preference profile $C \in \mathcal{P}$, denote with $\mathcal{D}_{\mathcal{P}}(C)$ the probability (mass or density) that a random individual in the population has preference profile C over the outcomes of all issues.

¹We focus on linear orders as is standard in classical social choice (Arrow 1951). This framework could be extended to partial orders or cardinal utilities, but we leave this for future work. Here, $\Delta[\mathcal{P}]$ denotes the space of probability distributions over the support set $\mathcal{P}$, and A^B denotes the space of mappings from B to A .

Often, we only need to consider the marginal distribution of $\mathcal{D}_{\mathcal{P}}$, i.e., the mapping $\mathcal{M} : \mathcal{I} \rightarrow \Delta[\text{LO}(N)]$. For any issue $i \in \mathcal{I}$, $\mathcal{M}(i)$ is the distribution of preferences over the N outcomes of issue i in the population. For preference ordering $o \in \text{LO}(N)$, we denote with $\mathcal{M}(i)_o$ the probability that a random individual in the population has preference ordering o over outcomes of issue i .

Outcomes. The result of a decision-making process is a preference profile $C : \mathcal{I} \rightarrow \text{LO}(N)$, which represents the aggregated preference of the population generated by some mechanism.

The mechanism is *representational* if the decision is made based on a finite collection of individual-issue pairs $\mathcal{S} = \{(o_1, i_1), (o_2, i_2), \dots, (o_{|\mathcal{S}|}, i_{|\mathcal{S}|})\}$, with $i_k \in \mathcal{I}$ sampled from $\mathcal{D}_{\mathcal{I}}$, and $o_k \sim \mathcal{M}(i_k)$ is an individual's preference over the outcomes of issue i_k , sampled from the population distribution.

However, not all preference profiles are allowed. As a key feature of representative social choice, the mechanism is only allowed to output preference profiles from a limited *candidate space* $C \subseteq \text{LO}(N)^{\mathcal{I}}$ (e.g., mutually compatible combinations of per-state construction projects in the national policy case, or language model policies in the AI alignment case). A mechanism is thus a function $f : (\text{LO}(N) \times \mathcal{I})^* \rightarrow C$, where A^* denotes the set of all finite sequences of elements from A , that maps the sample collection $\mathcal{S}$ to a preference profile $C \in C$.

REMARK 3.1 (ON FIXED N). *For simplicity, we assume a fixed N for all issues. This can model varying N_i by setting $N = \max(N_i)$ and treating unused outcomes as ‘dummy’ alternatives that are always ranked last.*

REMARK 3.2 (AI ALIGNMENT AS REPRESENTATIVE SOCIAL CHOICE). *Representative social choice can model the problem of AI alignment, where the LLM is trained to make decisions representative of the human population's preferences. Here, the set of issues $\mathcal{I}$ is the space of all possible user prompts, and the N outcomes are the candidate responses. A candidate profile $C \in C$ is a specific AI policy (e.g., a set of model weights). The candidate space C is the set of all policies achievable via the alignment process (e.g., all models reachable by RLHF training). The sample $\mathcal{S}$ is the human preference dataset. Our framework thus analyzes how well a policy trained on $\mathcal{S}$ generalizes to the true human population's preferences over all possible prompts.*

Binary Setting as Special Case. In the binary ($N = 2$) case of representative social choice, each issue admits two outcomes. In language model alignment (Bai, Jones, et al. 2022), this corresponds to pairwise comparisons where a human annotator chooses between two candidate model responses (e.g., Response A vs. Response B) for a given prompt. This reduction in complexity will allow for the analysis of a well-defined majority vote mechanism.

4 Summary of Key Results

Here, we summarize the key results from our analysis of representative social choice.

Generalization Bounds (Section 5.1). Generalization is essential in representative social choice. It ensures that decisions made by the mechanism based on a finite sample of preferences are representative of the population. We show that for binary issues, the required sample size $|\mathcal{S}|$ to ensure decisions generalize to the whole population scales with the complexity of the candidate space C , a measure of its complexity (Theorem 1). Theorem 2 generalizes this bound to non-binary settings.

In essence, the bounds tell us that the larger the sample size, the better the mechanism's decision reflects the population's true preferences, with the sample size needing to grow in proportion to the complexity of the candidate space, as measured by the VC dimension (Vapnik 1999) or the Rademacher complexity (Mohri and Rostamizadeh 2008); as an analogy, when election candidates tailor their messaging in a fine-grained manner, the population needs to watch more debates to find broadly aligned candidates, or else candidates could easily *overfit* their message to a few flagship issues.

Majority Vote and Scoring Mechanisms (Section 5.2, 6.1). In the binary setting, this generalization analysis leads to our first result: the simple *majority vote* mechanism is approximately optimal for the whole population, as long as

enough samples are drawn (Corollary 5.6). We extend this to non-binary issues using *scoring mechanisms*, which assign scores to preference profiles. We show these mechanisms also generalize well, with sample requirements depending on the complexity of the issues and outcomes (Corollary 6.2).

Representative Impossibilities (Section 6.2, 6.3, 6.4). Finally, we show that, like its classical counterpart, representative social choice faces fundamental impossibilities. We introduce a weak impossibility theorem (Theorem 3) that directly extends Arrow’s theorem. We then strengthen it by introducing *privilege graphs* to analyze issue interdependence, giving a strong impossibility theorem (Theorem 4) that identifies the structural conditions, i.e., cycles in the privilege graph, that lead to impossibility. It both explains why Arrow’s result holds and extends it to the general, constrained-candidate-space setting.

These theorems show that, in representative social choice, we must make trade-offs between different desirable properties, such as fairness, utility maximization, and convergence. The two theorems differ in their range of applicability, where the latter allows interdependence between issues by lifting the constraint on C .

We consider our contribution in Section 5, 6.1, 6.2 to be primarily conceptual and stage-setting, establishing the representative framework with techniques well-known in other fields. Section 6.3 and 6.4, when establishing the conceptually important strong impossibility theorem, additionally introduce new combinatorial tools of analysis, *privileged orderings* and the *privilege graph*, which we believe to be of independent interest.

5 Binary Representative Social Choice

In this section, we consider the case of representative social choice where the issues are binary. A real-world example is the case of jury trials, where a randomly selected jury makes binary decisions (guilt or innocence) on behalf of the entire population. AI alignment, with its binary preference annotations, is another example.

5.1 Binary Generalization Bound

When aggregating preferences, the mechanism only sees a finite collection of individual-issue pairs, and therefore the decision is made based on that finite sample. This raises the question of *generalization*, i.e., how can we ensure that the decision made by the mechanism is representative of the population’s preferences, when the mechanism is only optimized for a small number of individual opinions on individual issues?

To answer this question, we can leverage the theory of statistical learning, which studies the generalization properties of learning/optimization algorithms based on finite samples. The reliability of generalization depends on the complexity of the candidate space C ; the more flexible the candidate space, the more likely a profile can be picked by the mechanism that specifically fits the sample (*overfitting*) rather than the broader population. To characterize complexity, we use the *Vapnik-Chervonenkis (VC) dimension*, which measures the capacity of a hypothesis space to fit arbitrary finite samples (Vapnik 1999). Example A.1 in Appendix A gives a simple candidate space with VC dimension 2 over an infinite issue space.

DEFINITION 5.1 (VAPNIK-CHERVONENKIS DIMENSION (VAPNIK 1999)). *Given any issue space I , we consider candidate profiles mapping I to $\text{LO}(2)$, and a candidate space $C \subseteq \text{LO}(2)^I$. The VC dimension of C is the cardinality of the largest finite set of issues $I' \subseteq I$ such that for any binary function $o \in \text{LO}(2)^{I'}$, there exists a preference profile $c \in C$ such that $c(i) = o(i)$ for all $i \in I'$. If I can be arbitrarily large, then the VC dimension is infinite. Here we assume that the VC dimension is nonzero.*

REMARK 5.2 (VC DIMENSION FOR GENERAL N -OUTCOME SETTINGS). *Definition 5.1 is stated for the binary case where candidate profiles map issues to $\text{LO}(2)$. For the general N -outcome setting, where candidates map issues to $\text{LO}(N)$, the VC dimension can be defined analogously by encoding each N -ary ranking as a binary vector over the $\binom{N}{2}$ pairwise comparisons between outcomes. The candidate space $C \subseteq \text{LO}(N)^I$ then induces a binary candidate space over issue-pair combinations, to which the binary analysis applies. In the general setting, we will instead work*

with the Rademacher complexity (Theorem 2), which extends the same capacity-based analysis to real-valued scoring rules.

We can now introduce the sample complexity theorem. For any preference profile C as the aggregation outcome, the theorem gives an upper bound on the difference between the *sample utility* $\frac{1}{|\mathcal{S}|} \sum_k \mathbf{1}_{C(i_k)=o_k}$ (the goodness of the aggregated profile, evaluated on the selected samples) and the *population utility* $\mathbb{E}_{i \sim \mathcal{D}_I} [\mathcal{M}(i)_{C(i)}]$ (the unknown utility of the aggregated profile for the entire population), as a function of the sample size $|\mathcal{S}|$ and the VC dimension of the candidate space C . Such a difference is called the *generalization error*.

THEOREM 1 (BINARY GENERALIZATION BOUND). *Let C be a candidate space with VC dimension $\text{VC}(C)$, and let $\epsilon > 0$ be a desired generalization error. Then, for any $\delta > 0$, with probability at least $1 - \delta$, the sample utility and population utility of any preference profile $C \in \mathcal{C}$ are ϵ -close, i.e.,*

$$\Pr \left[\left| \frac{1}{|\mathcal{S}|} \sum_{k=1}^{|\mathcal{S}|} \mathbf{1}_{C(i_k)=o_k} - \mathbb{E}_{i \sim \mathcal{D}_I} [\mathcal{M}(i)_{C(i)}] \right| \leq \epsilon, \quad \forall C \in \mathcal{C} \right] \geq 1 - \delta \quad (1)$$

as long as we have the following, for some constant $c > 0$:

$$|\mathcal{S}| \geq \frac{c}{\epsilon^2} \text{VC}(C) \left(\log \text{VC}(C) + \log \frac{1}{\epsilon} + \log \frac{1}{\delta} \right) \quad (2)$$

REMARK 5.3 (BOUNDED AND UNBOUNDED VC DIMENSION IN AI ALIGNMENT). *In the context of AI alignment (Remark 3.2), the VC dimension of the candidate space C depends on the expressiveness of the policy class under consideration. When C consists of a small, finite set of pre-trained candidate LLMs (e.g., selecting among k existing models to deploy), the VC dimension is at most $\log_2 k$, since the candidate space can shatter at most $\log_2 k$ prompts. More generally, if the alignment process is restricted to a parameterized family with limited capacity, such as prompt templates with a bounded number of tunable parameters, or low-rank adapter (LoRA) fine-tuning (Hu et al. 2022) with a fixed rank, the VC dimension remains bounded and grows with the number of effective parameters. Note also that preference annotations in alignment are typically pairwise comparisons between two candidate responses, so the binary analysis applies directly at the level of annotations, and the size of the underlying response space enters the bounds only through the expressiveness of the policy class C .*

However, when C encompasses all policies reachable by unconstrained RLHF training of a large neural network, the candidate space is highly expressive, and the VC dimension can grow very large or become effectively infinite. In such cases, the generalization bounds of Theorem 1 require correspondingly large preference datasets $\mathcal{S}$ to ensure that the aligned model truly represents population preferences rather than overfitting to the training data. This distinction is practically important: alignment pipelines that restrict the policy class (e.g., by limiting fine-tuning to specific layers or using a small model family) benefit from stronger generalization guarantees, while unconstrained training necessitates larger and more diverse preference datasets.

REMARK 5.4 (ON THE ROLE OF GENERALIZATION BOUNDS). *While Theorem 1 is an application of established Vapnik-Chervonenkis (VC) theory, its importance here is conceptual. It formally connects the ‘representativeness’ of a social choice mechanism to the complexity of the candidate space C . This quantifies a trade-off, where a more expressive set of possible candidates (e.g., more nuanced political platforms or AI policies) requires significantly more preference samples to ensure the chosen candidate truly represents the population.*

Intuitively, the sample complexity theorem states that the sample utility of any preference profile in the candidate space C approximates the population utility, as long as the sample size is sufficiently large and in proportion to the complexity of the candidate space. Note that here we cannot directly apply the tail inequalities (Hellman and Raviv 1970) to bound the generalization error, because when the profile C is picked to maximize sample utility, the sample utility ceases to be an unbiased estimator of the population utility.

5.2 Case Study: Majority Vote in the Binary Case

In this section, we consider the *majority vote* mechanism under the binary representative setting. It generalizes the well-known majority vote mechanism where the population directly votes a single issue, to the representative setting where the population is represented by a collection of individual-issue pairs.

DEFINITION 5.5 (MAJORITY VOTE MECHANISM). *The majority vote mechanism is a representational mechanism f_{maj} that outputs the preference profile C that maximizes the sample utility, i.e.,*

$$f_{\text{maj}}(\mathcal{S}) = \arg \max_{C \in \mathcal{C}} \frac{1}{|\mathcal{S}|} \sum_{k=1}^{|\mathcal{S}|} \mathbf{1}_{C(i_k)=o_k}. \quad (3)$$

The majority vote mechanism can be viewed as a voting process where each individual-issue pair in the sample collection $\mathcal{S}$ casts a vote for the outcome that the individual prefers, based on the individual's preference over the two outcomes of the issue. The candidate profile that receives the most votes is then selected as the aggregated preference profile. When there is only one issue and candidate space $C = \text{LO}(2)$, the majority vote mechanism reduces to the standard majority vote mechanism in social choice theory.

From Theorem 1, we know that the majority vote mechanism has good generalization properties when the VC dimension of the candidate space C is small, resulting in the following corollary.

COROLLARY 5.6 (MAJORITY VOTE APPROXIMATELY MAXIMIZES POPULATION UTILITY). *For any error requirement $\epsilon > 0$ and confidence requirement $\delta > 0$, the majority vote mechanism f_{maj} satisfies*

$$\Pr \left[U_{\mathcal{D}_I, \mathcal{M}}(f_{\text{maj}}(\mathcal{S})) \geq -2\epsilon + \max_{C \in \mathcal{C}} U_{\mathcal{D}_I, \mathcal{M}}(C) \right] \geq 1 - \delta \quad (4)$$

where the population utility

$$U_{\mathcal{D}_I, \mathcal{M}}(C) := \mathbb{E}_{i \sim \mathcal{D}_I} [\mathcal{M}(i)_{C(i)}], \quad (5)$$

as long as the sample size $|\mathcal{S}|$ satisfies

$$|\mathcal{S}| \geq \frac{c}{\epsilon^2} \text{VC}(C) \left(\log \text{VC}(C) + \log \frac{1}{\epsilon} + \log \frac{1}{\delta} \right) \quad (6)$$

Furthermore, we can verify that the majority vote mechanism satisfies the following formal axioms, analogous to the classical *Pareto efficiency* and *non-dictatorship* axioms in social choice theory (Fishburn 2015).

Axiom: Probabilistic Pareto Efficiency (PPE), Binary Case

Consider profiles $C, C' \in \mathcal{C}$ such that for the only issue $i \in \mathcal{I}$ that they disagree on, we have $1 \succ_C 0, 0 \succ_{C'} 1$ and the population is unanimous on $1 \succ 0$.^a For any such C, C' and a sufficiently large sample size $|\mathcal{S}|$, with probability at least $1 - e^{-\alpha|\mathcal{S}|}$,^b $f_{\text{maj}}(\mathcal{S}) \neq C'$. Likewise when $0 \succ_C 1, 1 \succ_{C'} 0$ and the population is unanimous on $0 \succ 1$.

^aIn other words, $\mathcal{M}(i)_{1 \succ 0} = 1$

^b... where $\alpha > 0$ is an arbitrary constant dependent only on C, C' (we use $1 - e^{-\Omega(|\mathcal{S}|)}$ to denote a probability that converges to 1 exponentially fast in $|\mathcal{S}|$, where $\Omega(|\mathcal{S}|)$ is any quantity that asymptotically grows no slower than $|\mathcal{S}|$), and $|\mathcal{S}|$ is the number of samples. From now on, we will abbreviate this expression to $1 - e^{-\Omega(|\mathcal{S}|)}$.

REMARK 5.7. *Introduction of the candidate space C leads to interdependence between issues, which is not present in the classical social choice setting. As a result, we could not simply require that the mechanism output $1 \succ 0$ when the population is unanimous on $1 \succ 0$ as in the classical setting: what if the population is unanimous on $1 \succ 0$ for one issue, but unanimous on $0 \succ 1$ for another issue, and the two issues are strongly correlated (e.g., every*

candidate profile agrees on the two issues)? Cases like these, while less extreme, are widespread in the real world. Our formulation of PPE avoids this problem by comparing two candidate profiles C and C' that keep the same preference on all issues except one, a way to ensure “all else being equal”.

$1 - e^{-\Omega(|S|)}$ is the convergence rate from Hoeffding’s inequality for independent samples. Intuitively, it’s the confidence with which you can tell a majority from a minority in the population based on a sample S .

Axiom: Probabilistic Non-Dictatorship (PND), Binary Case

For any issue $i \in \mathcal{I}$, for any subpopulation $\mathcal{D}'_{\mathcal{P}}$ that occupies a probability mass $|\mathcal{D}'_{\mathcal{P}}| < 0.5$ in the whole population $\mathcal{D}_{\mathcal{P}}$, at least one of below is true w.r.t. issue i :

- When $\mathcal{D}'_{\mathcal{P}}$ is unanimous on $0 > 1$, there exists a population distribution $\mathcal{D}_{\mathcal{P}}$ for which $f_{\text{maj}}(S)(i) = (1 > 0)$ with probability $1 - e^{-\Omega(|S|)}$.
- When $\mathcal{D}'_{\mathcal{P}}$ is unanimous on $1 > 0$, there exists a population distribution $\mathcal{D}_{\mathcal{P}}$ for which $f_{\text{maj}}(S)(i) = (0 > 1)$ with probability $1 - e^{-\Omega(|S|)}$.

REMARK 5.8. Here, the requirement that at least one (as opposed to both) condition is fulfilled is, again, a consequence of the candidate space C . In the extreme case, if all candidate profiles agree that $0 > 1$ for issue i , then at most one of the two conditions can be met.

The PND axiom above is stronger than the classical version, in the sense that not only individuals, but also coalitions, cannot dictate the aggregation result. In fact, our problem setting treats individuals as interchangeable, thereby implicitly forcing *anonymity* of the mechanism, a stronger property than PND. As a result, representational mechanisms always satisfy PND, a fact that is formalized in Lemma 6.5.

Note that the classical axiom of independence of irrelevant alternatives (IIA) (Fishburn 2015) is not applicable to the binary case, since irrelevant alternatives only exist when $N \geq 3$.

6 General Representative Social Choice

Having formulated and analyzed the binary case of representative social choice, we now extend our analysis to the general case where there can be an arbitrary number of outcomes for each issue.

6.1 Scoring Rules and Generalization Errors

Generalizing the majority vote mechanism to the general case meets challenges, as the majority vote is over candidate profiles as opposed to outcomes, and when $N > 2$, the way each individual-issue pair votes is no longer well-defined. Instead, we can consider a more general class of mechanisms called *scoring mechanisms* (Lepelley et al. 2000), which for every individual-issue pair, assign a score to each candidate profile, and then output the candidate profile with the highest average score.

As an example, the scoring rule could be an arbitrary distance measure that measures the degree of alignment between the sampled individual’s preference and the candidate profile’s preference. The mechanism then outputs the candidate profile that maximizes the average sample alignment score.

DEFINITION 6.1. A scoring mechanism is defined by a scoring rule $s : \text{LO}(N) \times \text{LO}(N) \rightarrow \mathbb{R}$, which assigns a score to each pair of preference orderings over the N outcomes. The mechanism then outputs the candidate profile that maximizes the average score over all individual-issue pairs in the sample collection S , i.e.,

$$f_s(S) = \arg \max_{C \in C} \frac{1}{|S|} \sum_{k=1}^{|S|} s(o_k, C(i_k)). \quad (7)$$

For scoring mechanisms, we can similarly define the sample score $\frac{1}{|\mathcal{S}|} \sum_{k=1}^{|\mathcal{S}|} s(o_k, C(i_k))$ and the population score $U_{\mathcal{D}_I, \mathcal{M}}^s := \mathbb{E}_{i \sim \mathcal{D}_I, o \sim \mathcal{M}(i)} [s(o, C(i))]$, and have the following guarantees on generalization.

THEOREM 2 (GENERALIZATION BOUND FOR SCORING MECHANISMS). *For any scoring rule s with value bounded by constants, the scoring mechanism f_s satisfies*

$$\Pr \left[U_{\mathcal{D}_I, \mathcal{M}}^s(f_s(\mathcal{S})) \geq -2\epsilon + \max_{C \in \mathcal{C}} U_{\mathcal{D}_I, \mathcal{M}}^s(C) \right] \geq 1 - \delta \quad (8)$$

as long as the sample size $|\mathcal{S}|$ satisfies

$$|\mathcal{S}| \geq \frac{c}{\epsilon^2} \log \frac{1}{\delta}, \quad \hat{R}_{|\mathcal{S}|}(\bar{\mathcal{C}}) \leq \frac{\epsilon}{c}, \quad \text{where } \bar{\mathcal{C}} := \{(o, i) \mapsto s(o, C(i)) \mid C \in \mathcal{C}\} \quad (9)$$

and $\hat{R}_{|\mathcal{S}|}(\bar{\mathcal{C}})$ is the empirical Rademacher complexity of $\bar{\mathcal{C}}$ with respect to the sample collection $\mathcal{S}$ (Mohri and Rostamizadeh 2008), a generalization of the VC dimension to real-valued (as opposed to binary) functions.

The result follows directly from the Rademacher generalization bounds in statistical learning theory.

COROLLARY 6.2 (SCORING MECHANISMS APPROXIMATELY MAXIMIZE POPULATION SCORE). *When $|\mathcal{I}|$ is finite, for any scoring rule s with value bounded by constants, the scoring mechanism f_s satisfies*

$$\Pr \left[U_{\mathcal{D}_I, \mathcal{M}}^s(f_s(\mathcal{S})) \geq -2\epsilon + \max_{C \in \mathcal{C}} U_{\mathcal{D}_I, \mathcal{M}}^s(C) \right] \geq 1 - \delta \quad (10)$$

as long as we have the following, for some constant $c > 0$:

$$|\mathcal{S}| \geq \frac{c}{\epsilon^2} \left(|\mathcal{I}| N \log N + \log \frac{1}{\delta} \right) \quad (11)$$

Rademacher complexity generalizes the VC dimension to real-valued functions, and measures the capacity of a function class to fit arbitrary finite samples. Corollary 6.2 states that the sample score of any candidate profile in the candidate space $\mathcal{C}$ approximates the population score, as long as the sample size exceeds the ability of the candidate space to fit arbitrary finite samples — i.e., when the mechanism is forced to *generalize*.

Corollary 6.2 is agnostic towards the correlation structure among issues, and as a result, the sample complexity grows approximately linearly with the number of issues $|\mathcal{I}|$ and the number of outcomes N . When given an issue correlation structure, the sample complexity can potentially be reduced using Theorem 2.

6.2 Weak Representative Impossibility

In this section, we present axioms that ideal social choice mechanisms should satisfy, and present Arrow-like impossibility theorems, showing that no mechanism can satisfy all these axioms simultaneously. A weaker version of the impossibility theorem, which we present in this section, is a simple generalization of Arrow's impossibility theorem, while a stronger version shall be derived in Section 6.4. We introduce the necessary notations and then restate the axioms from the binary case, adapting them to our general setting.

DEFINITION 6.3 (OPERATION OF A PERMUTATION ON AN ORDERING). *For a set S , a linear order $o \in \text{LO}(S)$, and a permutation $\sigma \in \mathfrak{S}_S$, we define the operation $o \odot \sigma$ as the ordering obtained by applying the permutation σ to the elements of o . Specifically, for any $s_1, s_2 \in S$, we have $s_1 \succ_{o \odot \sigma} s_2$ if and only if $\sigma^{-1}(s_1) \succ_o \sigma^{-1}(s_2)$.*

For a subset $T \subseteq S$ and $\sigma \in \mathfrak{S}_T$, we similarly define $o \odot \sigma := o \odot \sigma|_S$, where $\sigma|_S$ is the extension of σ to the whole set S , mapping elements outside T to themselves. For a full profile $C \in \text{LO}(N)^{\mathcal{I}}$, we define $C \odot_i \sigma$ as the profile obtained by applying the permutation σ to $C(i)$, while keeping C unchanged for all other issues.

REMARK 6.4. We denote with $\sigma_{(c_1, c_2)}$ the 2-element permutation (transposition) that swaps c_1 and c_2 . Operation of a permutation on an ordering aims to capture “local changes” to a preference profile, where the permutation only affects a subset of the outcomes. They will be used in the definition of the axioms below.

Axiom: Probabilistic Pareto Efficiency (PPE)

Consider profiles $C, C' \in \mathcal{C}$ such that for the only issue $i \in \mathcal{I}$ that they disagree on, there exists $c, c' \in [N]$ such that $c \succ_C c'$, $C' = C \odot_i \sigma_{(c, c')}$ ^a and the population is unanimous on $c \succ c'$.^b For any such C, C' , with probability $1 - e^{\Omega(|S|)}$, $f(S) \neq C'$.

^aThe second condition means that the order between c, c' is the only disagreement they have on issue i .

^bIn other words, $M(i)_o = 0$ for all $o \in \text{LO}(N)$ that prefers c' over c .

Axiom: Probabilistic Non-Dictatorship (PND)

For any issue $i \in \mathcal{I}$ and $c, c' \in [N]$, for any subpopulation $\mathcal{D}'_{\mathcal{P}}$ that occupies a probability mass $|\mathcal{D}'_{\mathcal{P}}| < 0.5$ in the whole population $\mathcal{D}_{\mathcal{P}}$, at least one of below is true w.r.t. issue i :

- When $\mathcal{D}'_{\mathcal{P}}$ is unanimous on $c \succ c'$, there exists a population distribution $\mathcal{D}_{\mathcal{P}}$ for which $c' \succ_{f(S)} c$ with probability $1 - e^{\Omega(|S|)}$.
- When $\mathcal{D}'_{\mathcal{P}}$ is unanimous on $c' \succ c$, there exists a population distribution $\mathcal{D}_{\mathcal{P}}$ for which $c \succ_{f(S)} c'$ with probability $1 - e^{\Omega(|S|)}$.

In fact, *all* representational mechanisms satisfy PND, since our setting treats individuals as interchangeable and homogeneous, therefore implicitly forcing *anonymity* of the mechanism, which in turn implies PND.

LEMMA 6.5 (PROBABILISTIC NON-DICTATORSHIP FOR ALL REPRESENTATIONAL MECHANISMS). *For any $(\mathcal{I}, \mathcal{D}_{\mathcal{I}}, \mathcal{C})$ and any representational mechanism f , PND is satisfied.*

As a result, we shall remove PND from the explicit statements of the impossibility theorems, but it should be understood that PND is still implicitly satisfied. Finally, we define a weak version of the independence of irrelevant alternatives (IIA) axiom, as well as a new axiom specific to the representative setting.

Axiom: Weak Probabilistic Independence of Irrelevant Alternatives (W-PIIA)

When $\mathcal{C} = \text{LO}(N)^{\mathcal{I}}$, for two populations $\mathcal{D}_{\mathcal{P}}, \mathcal{D}'_{\mathcal{P}}$ that differ only in the preference over a single issue $i \in \mathcal{I}$ satisfying $M(i) \upharpoonright_{\text{LO}(\{c, c'\})} = M'(i) \upharpoonright_{\text{LO}(\{c, c'\})}$ ^a (where c, c' are any two elements of $[N]$),^b with probability $1 - e^{\Omega(|S|)}$, we have $f(S) \upharpoonright_{\text{LO}(\{c, c'\})} = f(S') \upharpoonright_{\text{LO}(\{c, c'\})}$.

^aHere, $M(i) \upharpoonright_{\text{LO}(\{c, c'\})}$ denotes the marginal distribution of preferences when restricted to the set $\{c, c'\}$.

^bMeaning that the population's distribution over the preference between c, c' are the same in the two populations.

REMARK 6.6. The assumption $\mathcal{C} = \text{LO}(N)^{\mathcal{I}}$ implies a structural independence in the candidate space: any preference over one issue can be combined with any preference over any other. As a result, we avoid “cross-issue” or “cross-outcome” structural dependencies; for example, we don't worry that a population's insistence on $c_1 \succ c_2$ might force the side effect of $c_3 \succ c_4$, simply because all available profiles in $\mathcal{C}$ happen to bundle $\{c_1, c_3\}$ together and $\{c_2, c_4\}$ together. This structural independence of $\mathcal{C}$ should be distinguished from statistical correlations in the population's preferences $\mathcal{D}_{\mathcal{P}}$ (which are not restricted by this assumption). This simplification makes $\mathcal{C} = \text{LO}(N)^{\mathcal{I}}$ the easiest candidate space to analyze. We will lift this assumption in Section 6.3 and 6.4.

Axiom: Weak Probabilistic Convergence (W-PC)

When $C = \text{LO}(N)^I$, for a population $\mathcal{D}_{\mathcal{P}}$ that's non-uniform on outcomes c, c' of issue $i \in I$,^a there exists an ordering $o \in \text{LO}(\{c, c'\})$ such that with probability $1 - e^{-\Omega(|S|)}$, $f(S)(i) \mid_{\text{LO}(\{c, c'\})} = o$.

^ai.e., $M(i) \mid_{\text{LO}(\{c, c'\})} (c > c') \neq 0.5$. We are only concerned with the marginal distribution.

REMARK 6.7. *W-PC requires that the mechanism not be torn between two outcomes when the population is not torn between them. This is to rule out “indecisive” mechanisms that are unable to make a decision with high probability (or requires too many samples to make that decision) when the population has clear preference. This need arises only in the representative setting where the mechanism is no longer deterministic.*

THEOREM 3 (WEAK REPRESENTATIVE IMPOSSIBILITY). *When $N \geq 3$, for any $(I, \mathcal{D}_I, C = \text{LO}(N)^I)$, no representational mechanism simultaneously satisfies PPE, W-PIIA, and W-PC for all $\mathcal{D}_{\mathcal{P}}$.*

The weak representative impossibility theorem is “weak” in the sense that it assumes $C = \text{LO}(N)^I$, i.e., no interdependence between issues or outcomes exist. This shall change in the next section.

6.3 Privileged Orderings and Privilege Graph

Before we present the strong representative impossibility theorem, we need tools to represent and analyze the structure of the candidate space C . To this end, we introduce the concept of *privileged orderings*, which are partial orderings that are preferred over all other alternative orderings in the candidate space.

DEFINITION 6.8 (PRIVILEGED ORDERING). *For an issue $i \in I$ and a subset of outcomes $T = \{c_1, c_2, \dots, c_k\} \subseteq [N]$ ($k \geq 2$), we call $o \in \text{LO}(T) : c_1 > c_2 > \dots > c_k$ a privileged ordering if for any extension $C \in \text{LO}(N)^I$ of o ² and permutation $\sigma \in \mathfrak{S}_T$ of T , we have $C \odot_i \sigma \in C \implies C \in C$.*

REMARK 6.9. *Intuitively, a privileged ordering o is a partial ordering that is preferred in the candidate space C over all other alternative partial orderings. Any full preference profile in C that disagrees with o must have a counterpart in C that agrees with o while keeping the rest of the profile unchanged. These privileged orderings are in fact surprisingly common in practice, as will be showcased in Example 6.11 and 6.12.*

DEFINITION 6.10 (PRIVILEGE GRAPH). *For an issue $i \in I$, we define the privilege graph G_i as a directed graph with vertices as the outcomes in $[N]$, and an edge from u to v iff there exists a privileged ordering $u > v$. We call i cyclically privileged if its privilege graph contains a simple directed cycle of length at least 3.³*

EXAMPLE 6.11 (PRIVILEGED ORDERINGS IN THE REAL WORLD). *Consider three persons $I = \{A, B, C\}$ on trial before the jury. A is charged with burglary (\$5k, first trial), B also with burglary (\$50k, second trial), and C with fraud. The $N = 3$ possible outcomes for each defendant include acquittal (**a**), community service (**c**), and imprisonment (**i**). The jury ranks the outcomes for each defendant in order of recommendation. In this hypothetical case, the following factors may lead to privileged orderings:*

- **Monotonicity and fairness constraints.** *The jury may consider it unfair to punish A harder than B given the difference in the amount of burglary. Thus, locally changing the recommended outcome of A from **c** to **i** may violate this monotonicity (namely when the outcome of B is **c**), but changing from **i** to **c** will not. This implies $\mathbf{c} > \mathbf{i}$ and (analogously) $\mathbf{a} > \mathbf{c}$ being privileged for A , and $\mathbf{i} > \mathbf{c}, \mathbf{c} > \mathbf{a}$ for B .*

²i.e., C must agree with o on issue i that $c_1 > \dots > c_k$, but can differ on the remaining outcomes of i and on other issues.

³Note that for our purposes, we don't consider graphs containing only 2-cycles as cyclic.

- **Local independence.** While the crimes of A and B are similar and therefore correlated in the jury's judgment, C is unrelated to the other two. As a result, C may be the Cartesian product of the candidate space $C_{(A,B)}$ for (A, B) and C_C for C . When $C_C = \text{LO}(\{a, c, i\})$ and $C = C_{(A,B)} \times \text{LO}(\{a, c, i\})$, it can be verified that C 's privilege graph G_C is a complete digraph, and every ordering is privileged for C .
- **Default outcomes.** Assume that B is sentenced to c in the first trial, and the jury is inclined to keep the sentence since no new substantial evidence is presented. This may lead to the privileged orderings $c > i$ and $c > a$ for B , since c , being the default outcome, is a plausible substitute for any other outcome.

These considerations are general, going far beyond the example's legal setting. Locally independent decisions, and decisions with a default outcome, are common in the real world; and many decisions (e.g., those involving numerical balancing of costs and benefits), according to common sense, should be monotonic as well. There are other reasons for privileged orderings too, and the three factors listed above are only for illustration.

EXAMPLE 6.12 (PRIVILEGED ORDERINGS IN AI AND AI ALIGNMENT). Outside of the human case, privileged orderings (Definition 6.10) are a natural model for the inductive biases of AI systems (Baxter 2000), including deep neural networks. For instance, an LLM pretrained on web text might have a structural bias (a 'privilege') for certain types of responses (e.g., coherent ones) over others (Perez et al. 2022; Santurkar et al. 2023), and alignment training may be unable to make the model output a response that violates a deep-seated bias; this creates a non-trivial candidate space $C \subseteq \text{LO}(N)^I$ where impossibility results can arise. More examples include:

- **Language models.** Language models are known to display a wide range of human behavioral tendencies (Lampinen et al. 2024; Perez et al. 2022; Santurkar et al. 2023), including ones discussed in Example 6.11. Tendencies including consistent treatment of similar decisions, default outcomes, monotonicity, and many more, can all lead to privileged orderings. In the alignment process, these pretrained tendencies serve as inductive biases that set limits on the possible outcomes of alignment training.
- **Reinforcement learning agents.** In reinforcement learning, the reward function is a key component that guides the learning process. It has been demonstrated to induce agents to learn deeply ingrained hierarchies or equivalences between outcomes (Di Langosco et al. 2022; H. Wang et al. 2024). These prior tendencies set limits on the possible outcomes of later learning, leading to privileged orderings.
- **Vision models.** Vision models have been shown to exhibit spatial locality (K. Li et al. 2022), translation invariance (Kauderer-Abrams 2017), simplicity bias (Shah et al. 2020), and other inductive biases. These biases, which are often necessary for the models to learn effectively, operate by setting limits on the possible outcomes of learning, leading to privileged orderings.

While these tendencies in AI systems are often probabilistic and less clear-cut than in the human case, privileged orderings can still be a useful abstraction for understanding and approximating their behavior.

Next, we characterize important properties of privileged orderings and privilege graphs.

LEMMA 6.13 (TRANSITIVITY OF PRIVILEGED GRAPH). For an issue $i \in I$ and $u, v, w \in [N]$, if $u > v$ and $v > w$ are both privileged, then $u > w$ is privileged.

LEMMA 6.14 (CLOSURE OF PRIVILEGED ORDERINGS UNDER CONCATENATION). Let $c_1, c_2, \dots, c_{k+m-1} \in [N]$ ($k, m \geq 2$) be distinct outcomes in an issue i . If $o : c_1 > c_2 > \dots > c_k$ and $o' : c_k > c_{k+1} > c_{k+2} > \dots > c_{k+m-1}$ are both privileged orderings, then $o \oplus o' : c_1 > c_2 > \dots > c_{k+m-1}$ is also a privileged ordering.

COROLLARY 6.15. For issue $i \in I$ and $o : c_1 > c_2 > \dots > c_k$ ($k \geq 2$) containing distinct outcomes in i , o is a privileged ordering if a (possibly non-simple) path exists in G_i that passes through $c_1, c_2, \dots, c_k$ in that order.

The privilege graph G_i captures all binary privileged orderings in issue i , and will be the primary way in which we represent the structure of the candidate space C . Properties stronger than Lemma 6.13 and 6.14 often do not

hold, e.g., a subsequence of a privileged ordering is not necessarily privileged. Also, the condition in Corollary 6.15 is a sufficient but not necessary condition for a privileged ordering.

6.4 Strong Representative Impossibility

Before stating the strong impossibility, we first present the strong version of the PIIA and PC axioms.

Axiom: Strong Probabilistic Independence of Irrelevant Alternatives (S-PIIA)

For arbitrary C and $c, c' \in [N]$ such that $c > c'$ and $c' > c$ are both privileged orderings in some issue $i \in \mathcal{I}$, for two populations $\mathcal{D}_P, \mathcal{D}'_P$ satisfying $\mathcal{M}(i) \upharpoonright_{\text{LO}(\{c, c'\})} = \mathcal{M}'(i) \upharpoonright_{\text{LO}(\{c, c'\})}$,^a with probability $1 - e^{-\Omega(|S|)}$, we have $f(S) \upharpoonright_{\text{LO}(\{c, c'\})} = f(S') \upharpoonright_{\text{LO}(\{c, c'\})}$.

^aMeaning that the population's distribution over the preference between c, c' are the same in the two populations.

Axiom: Strong Probabilistic Convergence (S-PC)

For arbitrary C and $c, c' \in [N]$ such that $c > c'$ and $c' > c$ are both privileged orderings in some issue $i \in \mathcal{I}$, for a population $\mathcal{D}_P$ that's non-uniform on $\{c, c'\}$,^a there exists an ordering $o \in \text{LO}(\{c, c'\})$ such that with probability $1 - e^{-\Omega(|S|)}$, $f(S)(i) \upharpoonright_{\text{LO}(\{c, c'\})} = o$.

^ai.e., $\mathcal{M}(i) \upharpoonright_{\text{LO}(\{c, c'\})} (c > c') \neq 0.5$. We are only concerned with the marginal distribution.

REMARK 6.16. *S-PIIA and S-PC are generalizations of W-PIIA and W-PC, respectively, to arbitrary candidate spaces C . This is achieved by limiting comparisons to pairs that are privileged orderings in both directions, ensuring that they are locally independent from other outcomes and issues.*

This finally leads us to the strong representative impossibility theorem.

THEOREM 4 (STRONG REPRESENTATIVE IMPOSSIBILITY). *For any $(\mathcal{I}, \mathcal{D}_I, C)$, when there is at least one cyclically privileged issue,⁴ no representational mechanism simultaneously satisfies PPE, S-PIIA, and S-PC for all $\mathcal{D}_P$.*

Meanwhile, given any mapping ϕ from each issue i to a privilege graph $\phi(i)$ without simple cycles of length at least 3, there exist a candidate space C whose privilege graph G_i (for each i) is $\phi(i)$ or its supergraph, and a representational mechanism f over $(\mathcal{I}, \mathcal{D}_I, C)$ satisfying PPE, S-PIIA, and S-PC for all $\mathcal{D}_P$.

Theorem 4 is a generalization of Arrow's theorem and Theorem 3 to arbitrary candidate spaces C . It shows that the cyclicity is both necessary and sufficient for the impossibility of a representative mechanism satisfying the axioms. However, the "necessary" part is not as strong as one would hope, since it constructs counterexample for at least one C associated with each non-cyclic privilege graph, instead of constructing counterexamples for all C associated with the privilege graph. It is an open question whether such a stronger necessity holds, and shall be the subject of future research.

The $N \geq 3$ condition ("at least 3 vertices in the N -clique")⁵ in Theorem 3 and the original Arrow's theorem is replaced by the cyclicity condition here ("at least 3 vertices in a cycle"). Intuitively, the latter is a more precise condition that identifies a necessary-and-sufficient "minimal structure" in the candidate space C leading to impossibility. It can be checked in linear time (in the number of vertices and edges) for any privilege graph, since it's equivalent to the existence of strongly connected components of size at least 3.

⁴i.e., issues whose privilege graphs contain simple cycles of length at least 3 (Definition 6.10).

⁵When $C = \text{LO}(N)^I$, as is assumed by Arrow's theorem and Thm. 3, the N outcomes form a clique in the privilege graph.

7 Conclusion

In this paper, we have formulated the problem of representative social choice, where a mechanism aggregates the preferences of a population based on a finite sample of individual-issue pairs. We have derived results that reflect both optimistic and pessimistic aspects of representative social choice.

Implications for AI Alignment. Representative social choice can model the alignment of AI systems to diverse human preferences. Our generalization bounds (Theorems 1 and 2) directly apply to alignment. They quantify how many human preference annotations are needed for a trained model to faithfully represent population preferences, as a function of the expressiveness of the model class. For instance, alignment pipelines that restrict the policy class, such as selecting among a finite set of pre-trained models or fine-tuning with low-rank adapters, enjoy tighter generalization guarantees than unconstrained RLHF training (Remark 5.3).

On the pessimistic side, our impossibility results (Theorems 3 and 4) reveal concrete trade-offs for alignment. For example, an alignment mechanism that satisfies Probabilistic Pareto Efficiency (i.e., ensuring the model does not systematically ignore any user group’s preferences) cannot simultaneously satisfy independence axioms when aggregating preferences across diverse prompts with three or more candidate responses. In practice, this means that RLHF reward models trained to maximize aggregate human approval may inadvertently introduce systematic biases in how they trade off between different user populations’ preferences. The privilege graph analysis (Section 6.3) further identifies the structural conditions under which these trade-offs are unavoidable, providing alignment practitioners with a diagnostic tool for understanding when their preference aggregation procedure is subject to impossibility constraints.

Limitations and Future Directions. We focused on the generalization properties of representational mechanisms without studying other important properties, such as incentive compatibility and computational tractability. Future research could explore them and their interactions with generalization. Furthermore, our model does not account for noise or inconsistency in human preferences, a significant practical challenge. Other open directions include methods for learning the scoring function s itself (Theorem 2) from data, or extending the model to handle out-of-distribution generalization to unseen outcomes or even unseen issues, which is vital for generative AI models.

Broader Impact Statement. We aim to advance our understanding of democratic representation in collective decisions, with anticipated positive impacts. It helps develop more robust alignment strategies for AI systems, contributing to the equitable alignment of AI systems with diverse human preferences.

Acknowledgments

Many thanks to Anand Siththaranjan, Yifeng Ding, Yaowen Ye, Xiaoyuan Zhu, and Karim Abdel Sadek for helpful discussions. The work is supported by the National Natural Science Foundation of China (624B100026).

References

- K. J. Arrow. 1951. *Social choice and individual values*. Vol. 12. Yale University Press.
- H. Azari, D. Parkes, and L. Xia. 2012. “Random utility theory for social choice.” *Advances in Neural Information Processing Systems*, 25.
- Y. Bai, A. Jones, et al.. 2022. *Training a helpful and harmless assistant with reinforcement learning from human feedback*. (2022).
- Y. Bai, S. Kadavath, et al.. 2022. *Constitutional AI: Harmlessness from AI feedback*. (2022).
- J. Baxter. 2000. “A model of inductive bias learning.” *Journal of artificial intelligence research*, 12, 149–198.
- A. Bhattacharyya and P. Dey. 2021. “Predicting winner and estimating margin of victory in elections using sampling.” *Artificial Intelligence*, 296, 103476.
- O. Bousquet, S. Boucheron, and G. Lugosi. 2003. “Introduction to statistical learning theory.” In: *Advanced Topics in Machine Learning*. University College London, 169–207.
- F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.

- J. Coleman and J. Ferejohn. 1986. "Democracy and social choice." *Ethics*, 97, 1, 6–25.
- V. Conitzer, M. Rognlie, and L. Xia. 2009. "Preference Functions that Score Rankings and Maximum Likelihood Estimation." In: *IJCAI*. Vol. 9, 109–115.
- V. Conitzer, R. Freedman, et al.. 2024. "Position: Social Choice Should Guide AI Alignment in Dealing with Diverse Human Feedback." In: *Forty-first International Conference on Machine Learning*.
- L. L. Di Langosco, J. Koch, L. D. Sharkey, J. Pfau, and D. Krueger. 2022. "Goal misgeneralization in deep reinforcement learning." In: *International Conference on Machine Learning*. PMLR, 12004–12019.
- A. M. Feldman and R. Serrano. 2006. *Welfare economics and social choice theory*. Springer Science & Business Media.
- S. Fish, P. Gözl, D. C. Parkes, A. D. Procaccia, G. Rusak, I. Shapira, and M. Wüthrich. 2023. *Generative social choice*. (2023).
- P. C. Fishburn. 2015. *The theory of social choice*. Princeton University Press.
- L. Ge, B. Juba, and Y. Vorobeychik. 2024. "Learning linear utility functions from pairwise comparison queries." *arXiv preprint arXiv:2405.02612*.
- L. Ge, D. Halpern, E. Micha, A. D. Procaccia, I. Shapira, Y. Vorobeychik, and J. Wu. 2024. *Axioms for AI Alignment from Human Feedback*. (2024).
- A. Gibbard. 1973. *Manipulation of voting schemes: a general result*. (1973).
- D. Halpern, G. Kehne, A. D. Procaccia, J. Tucker-Foltz, and M. Wüthrich. 2026. "Representation with incomplete votes." *Theory and Decision*, 1–40.
- M. Hellman and J. Raviv. 1970. "Probability of error, equivocation, and the Chernoff bound." *IEEE Transactions on Information Theory*, 16, 4, 368–372.
- E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al.. 2022. "Lora: Low-rank adaptation of large language models." *ICLR*, 1, 2, 3.
- S. Huang, D. Siddarth, L. Lovitt, T. I. Liao, E. Durmus, A. Tamkin, and D. Ganguli. 2024. "Collective Constitutional AI: Aligning a Language Model with Public Input." In: *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1395–1417.
- E. Kauderer-Abrams. 2017. *Quantifying translation-invariance in convolutional neural networks*. (2017).
- O. Klingefjord, R. Lowe, and J. Edelman. 2024. *What are human values, and how do we align AI to them?* (2024).
- A. Köpf et al.. 2024. "Openassistant conversations-democratizing large language model alignment." *Advances in Neural Information Processing Systems*, 36.
- A. K. Lampinen, I. Dasgupta, S. C. Chan, H. R. Sheahan, A. Creswell, D. Kumaran, J. L. McClelland, and F. Hill. 2024. "Language models, like humans, show content effects on reasoning tasks." *PNAS nexus*, 3, 7.
- D. Lepelley, P. Pierron, and F. Valognes. 2000. "Scoring rules, Condorcet efficiency and social homogeneity." *Theory and Decision*, 49, 2, 175–196.
- K. Li, R. Yu, Z. Wang, L. Yuan, G. Song, and J. Chen. 2022. "Locality guidance for improving vision transformers on tiny datasets." In: *European Conference on Computer Vision*. Springer, 110–127.
- C. List. 2012. "The theory of judgment aggregation: an introductory review." *Synthese*, 187, 1, 179–207.
- N. R. Miller. 1983. "Pluralism and social choice." *American Political Science Review*, 77, 3, 734–747.
- A. Mishra. 2023. *AI alignment and social choice: Fundamental limitations and policy implications*. (2023).
- M. Mohri and A. Rostamizadeh. 2008. "Rademacher complexity bounds for non-iid processes." *Advances in neural information processing systems*, 21.
- N. Nisan and A. Ronen. 1999. "Algorithmic mechanism design." In: *Proceedings of the thirty-first annual ACM Symposium on Theory of Computing*, 129–140.
- D. Parkes and A. Procaccia. 2013. "Dynamic social choice with evolving preferences." In: *Proceedings of the AAAI conference on artificial intelligence*. Vol. 27, 767–773.
- E. Perez et al.. 2022. *Discovering language model behaviors with model-written evaluations*. (2022).
- M. Prasad. 2018. *Social choice and the value alignment problem*. (2018).
- S. Santurkar, E. Durmus, F. Ladhak, C. Lee, P. Liang, and T. Hashimoto. 2023. "Whose opinions do language models reflect?" In: *International Conference on Machine Learning*. PMLR, 29971–30004.
- M. A. Satterthwaite. 1975. "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions." *Journal of economic theory*, 10, 2, 187–217.
- H. Shah, K. Tamuly, A. Raghunathan, P. Jain, and P. Netrapalli. 2020. "The pitfalls of simplicity bias in neural networks." *Advances in Neural Information Processing Systems*, 33, 9573–9585.
- A. D. Taylor. 2005. *Social choice and the mathematics of manipulation*. Cambridge University Press.
- V. N. Vapnik. 1999. "An overview of statistical learning theory." *IEEE transactions on neural networks*, 10, 5, 988–999.
- H. Wang, E. Miahi, M. White, M. C. Machado, Z. Abbas, R. Kumaraswamy, V. Liu, and A. White. 2024. "Investigating the properties of neural network representations in reinforcement learning." *Artificial Intelligence*, 330, 104100.
- H. P. Young. 1975. "Social choice scoring functions." *SIAM Journal on Applied Mathematics*, 28, 4, 824–838.
- T. Zhi-Xuan, M. Carroll, M. Franklin, and H. Ashton. 2024. *Beyond Preferences in AI Alignment*. (2024).

A Additional Definitions

The formal definition of the Vapnik-Chervonenkis dimension is given in Definition 5.1 in Section 5.1. We provide an additional example here for illustration.

EXAMPLE A.1 (SINGLE-PEAKED PREFERENCES). Consider a case where issues $i \in \mathcal{I} = \mathbb{R}$ are points on a political spectrum, and outcomes are $N = 2$ ('approve', 'disapprove'). A candidate $C \in \mathcal{C}$ might be defined by an 'interval' $[p_C - w_C, p_C + w_C]$, such that they 'approve' of issue i if and only if i falls within their interval. The class of such 'interval' functions has a finite VC dimension of 2, even though $\mathcal{I}$ is infinite. This models candidates with coherent, single-peaked political platforms.

B Proofs in Binary Representative Social Choice

In this appendix, we provide the proofs of the results in Section 5.

B.1 Proof of Theorem 1

PROOF. We adapt the standard VC generalization bound (Vapnik 1999), which states that for a class $\mathcal{C}$ of binary functions with finite $\text{VC}(\mathcal{C})$, for any $\epsilon, \delta > 0$, the uniform convergence bound $\Pr[\sup_{C \in \mathcal{C}} |\hat{R}_S(C) - R(C)| \leq \epsilon] \geq 1 - \delta$ holds⁶ if the sample size $|S|$ is polynomial in $(1/\epsilon)$, $\log(1/\delta)$, and $\text{VC}(\mathcal{C})$ (as stated in the theorem). Our proof reduces the representative social choice problem to this standard learning setting.

First consider the case where population has size 1, and the resulting population $\mathcal{M}$ always maps an issue to a one-point distribution (i.e., a deterministic preference).

In this case, $\mathcal{M}(i)_{C(i)} \in \{0, 1\}$ represents whether the aggregated profile C is correct on issue i . The population utility $E_{i \sim \mathcal{D}_I} [\mathcal{M}(i)_{C(i)}]$ can thus be formulated as the *population error* in statistical learning settings, and the result in Theorem 1 follows directly from the VC generalization error bound in statistical learning theory (Vapnik 1999).

To generalize the result to the case where the population has arbitrary cardinality, we can construct the following reduction to the deterministic case. We make the following definitions:

$$\tilde{\mathcal{I}} := \mathcal{I} \times \text{LO}(2) \quad (12)$$

$$\tilde{\mathcal{D}}_{\tilde{\mathcal{I}}}(i, b) := \mathcal{D}_I(i) \cdot \mathcal{M}(i)_b \quad (13)$$

$$\tilde{\mathcal{M}}(i, b)_{b'} := \mathbf{1}_{b=b'} \quad (14)$$

$$\tilde{\mathcal{C}} := \{\tilde{C} : (i, b) \mapsto C(i) \mid C \in \mathcal{C}\} \quad (15)$$

$$\tilde{\mathcal{S}} := \{(o_1, (i_1, o_1)), (o_2, (i_2, o_2)), \dots, (o_{|S|}, (i_{|S|}, o_{|S|}))\} \quad (16)$$

where $b, b' \in \text{LO}(2)$.

In other words, we duplicate each issue i into two issues $(i, 0)$ and $(i, 1)$, and for each issue i , we construct a population $\tilde{\mathcal{M}}(i, b)$ that always maps the issue to a one-point distribution, where the point is b . We then construct a saliency distribution $\tilde{\mathcal{D}}_{\tilde{\mathcal{I}}}$ that accounts for the original population probabilities over preferences. It can be

⁶...where $\hat{R}_S(C)$ gives the empirical accuracy of function C on samples S and $R(C)$ gives the expected accuracy of C on a random sample.

verified that

$$\mathbb{E}_{(i,b) \sim \tilde{\mathcal{D}}_{\tilde{I}}} [\tilde{\mathcal{M}}(i, b)_{\tilde{C}(i,b)}] = \mathbb{E}_{i \sim \mathcal{D}_I} [\mathcal{M}(i)_{C(i)}] \quad (17)$$

$$\frac{1}{|\tilde{\mathcal{S}}|} \sum_{k=1}^{|\tilde{\mathcal{S}}|} \mathbf{1}_{\tilde{C}((i_k, o_k))=o_k} = \frac{1}{|\mathcal{S}|} \sum_{k=1}^{|\mathcal{S}|} \mathbf{1}_{C(i_k)=o_k} \quad (18)$$

$$\text{VC}(\tilde{\mathcal{C}}) = \text{VC}(\mathcal{C}) \quad (19)$$

Since Theorem 1 holds for the deterministic case $(\tilde{\mathcal{I}}, \tilde{\mathcal{D}}_{\tilde{I}}, \tilde{\mathcal{M}}, \tilde{\mathcal{C}})$, it also holds for the original case $(\mathcal{I}, \mathcal{D}_I, \mathcal{M}, \mathcal{C})$. $\square$

B.2 Proof of Corollary 5.6

PROOF. From Theorem 1, with probability $1 - \delta$, we have

$$U_{\mathcal{D}_I, \mathcal{M}}(f_{\text{maj}}(\mathcal{S})) \geq -\epsilon + \frac{1}{|\mathcal{S}|} \sum_{k=1}^{|\mathcal{S}|} \mathbf{1}_{f_{\text{maj}}(\mathcal{S})(i_k)=o_k} \quad (20)$$

$$= \max_{C \in \mathcal{C}} \left(-\epsilon + \frac{1}{|\mathcal{S}|} \sum_{k=1}^{|\mathcal{S}|} \mathbf{1}_{f_{\text{maj}}^C(i_k)=o_k} \right) \quad (21)$$

$$\geq -2\epsilon + \max_{C \in \mathcal{C}} U_{\mathcal{D}_I, \mathcal{M}}(C) \quad (22)$$

$\square$

C Proofs in General Representative Social Choice

In this appendix, we provide the proofs of the results in Section 6.

C.1 Proof of Corollary 6.2

PROOF. Proof of the corollary is done by bounding the empirical Rademacher complexity of the function class $\tilde{\mathcal{C}}$ using Massart's Lemma (Bousquet et al. 2003), which gives us

$$\hat{\mathbf{R}}_{|\mathcal{S}|}(\tilde{\mathcal{C}}) \leq c \sqrt{\frac{\log |\tilde{\mathcal{C}}|}{|\mathcal{S}|}}$$

for some constant $c > 0$.

Given that the cardinality of $\tilde{\mathcal{C}}$ is $(N!)^{|\mathcal{I}|}$, the Corollary is not hard to verify by plugging the bound into Theorem 2. $\square$

C.2 Proof of Lemma 6.5

PROOF. Assume otherwise, that under some mechanism f , a subpopulation $\mathcal{D}'_{\mathcal{P}}$ occupying a probability mass strictly less than 0.5 can dictate the aggregation result in both directions, by being unanimous on either $c' > c$ or $c > c'$.

Take another subpopulation $\mathcal{D}''_{\mathcal{P}}$ that's disjoint with $\mathcal{D}'_{\mathcal{P}}$ and has the same probability mass. Let $\mathcal{D}'_{\mathcal{P}}$ be unanimous on $c' > c$, and $\mathcal{D}''_{\mathcal{P}}$ be unanimous on $c > c'$. Since populations of the same probability mass are indistinguishable to the mechanism, $\mathcal{D}''_{\mathcal{P}}$ must also dictate the aggregation result in both directions, leading to a contradiction. $\square$

C.3 Proof of Theorem 3

PROOF. The proof proceeds by a simple reduction to Arrow's impossibility theorem. Given $N \geq 3$, any instance of the classical Arrow problem with an odd number n of voters can be simulated. We divide the population $\mathcal{D}_P$ into n disjoint subpopulations, each with mass $1/n$, and set each to be unanimous on one voter's preference profile.

Since n is odd, there are no ties. The $1 - e^{\Omega(|S|)}$ convergence from W-PC and W-PIIA implies that, as $|S| \rightarrow \infty$, the probability of our mechanism *disagreeing* with the deterministic outcome dictated by the population majority goes to 0. This allows us to apply the logic of the deterministic Arrow's theorem to the limiting behavior of our mechanism.

The independence between issues (from $C = \text{LO}(N)^I$) allows us to handle the multi-issue setting by examining each issue in isolation. The reduction is thus complete.

Since Arrow's theorem shows that for any problem instance with $N \geq 3$, no deterministic social choice mechanism can satisfy Pareto efficiency, independence of irrelevant alternatives, and non-dictatorship simultaneously, we are able to reduce at least one such hard instance to every instance $(I, \mathcal{D}_I, C = \text{LO}(N)^I)$ of the weak representative setting, and the theorem follows. $\square$

C.4 Proof of Lemma 6.13

PROOF. For any extension $C \in \text{LO}(N)^I$ of $u > w$, we need to show that $C' := C \odot_i \sigma_{(u,w)} \in C \implies C \in C$. Assuming $C' \in C$, we consider ordering between u, v, w in $C(i)$.

When $u >_C w >_C v$, we have $v >_{C'} w >_{C'} u$. Since $u > v$ is privileged, $C'' := C' \odot_i \sigma_{(u,v)} \in C$, and we have $u >_{C''} w >_{C''} v$. Since $v > w$ is privileged, $C''' := C'' \odot_i \sigma_{(v,w)} \in C$. It can be verified that $u >_{C'''} v >_{C'''} w$ and $C''' = C$. The case $v >_C u >_C w$ is analogous.

When $u >_C v >_C w$, we have $w >_{C'} v >_{C'} u$. Since $v > w$ is privileged, $C'' := C' \odot_i \sigma_{(v,w)} \in C$. Now $v >_{C''} w >_{C''} u$, reducing to the first case. $\square$

C.5 Proof of Lemma 6.14

PROOF. For any extension $C \in \text{LO}(N)^I$ of $o \oplus o'$ and permutation $\sigma \in \mathfrak{S}_{\{c_1, c_2, \dots, c_{k+m-1}\}}$, we need to show that $C' := C \odot_i \sigma \in C \implies C \in C$.

Similar to the proof of Lemma 6.13, there are two transformations that we are allowed to apply to C' to obtain C : first, we can apply some permutation $\sigma_o \in \mathfrak{S}_{\{c_1, c_2, \dots, c_k\}}$ on C to sort $\{c_1, c_2, \dots, c_k\}$ in this order if it isn't already sorted; second, we can apply some permutation $\sigma_{o'} \in \mathfrak{S}_{\{c_k, c_{k+1}, \dots, c_{k+m-1}\}}$ on C to sort $\{c_k, c_{k+1}, \dots, c_{k+m-1}\}$ in this order if it isn't already. Since o and o' are privileged, these two transformations preserve membership in C .

We repeatedly apply these two transformations in arbitrary order, until neither of them can be applied — i.e., until both $\{c_1, c_2, \dots, c_k\}$ and $\{c_k, c_{k+1}, \dots, c_{k+m-1}\}$ are sorted in the correct order. Since $(c_1 > c_2 > \dots > c_k) \wedge (c_k > c_{k+1} > c_{k+2} > \dots > c_{k+m-1}) \implies c_1 > c_2 > \dots > c_{k+m-1}$, when such a process terminates, we have $C \in C$.

We still need to show that the process indeed terminates. We define a *potential function* Φ such that for any $E \in \text{LO}(N)^I$, $\Phi(E) = \text{Inv}_{c_1 > \dots > c_{k+m-1}}(E(i)) = \sum_{1 \leq x < y \leq k+m-1} \mathbf{1}_{c_y >_{E(i)} c_x}$, the number of inversions in $E(i)$ with respect to the ordering $o \oplus o' : c_1 > \dots > c_{k+m-1}$.

We show that the transformation $\odot_i \sigma_o$ strictly decreases $\text{Inv}_{o \oplus o'}$. We first decompose $\odot_i \sigma_o$ into a series of transpositions $\odot_i \sigma_{(c_{j_1}, c_{j_2})}$ such that $j_1 < j_2$ and $c_{j_2} > c_{j_1}$ in the pre-transposition ordering. It can be verified that this transposition removes the inversion (c_{j_1}, c_{j_2}) without introducing new inversions. Therefore the transformation $\odot_i \sigma_o$ strictly decreases Φ , and likewise for $\odot_i \sigma_{o'}$. Since the initial $\Phi(C')$ is finite and Φ is always non-negative, the process must terminate. $\square$

C.6 The Field Expansion Lemma

The Field Expansion Lemma is a key instrument in proving the strong representative impossibility theorem. It shows that the decisiveness of a subpopulation over a specific preference can be expanded to decisiveness over all pairwise orderings among a set of outcomes. We state and prove the lemma in this section.

DEFINITION C.1 (DECISIVENESS). *Given any representational mechanism f , for a subpopulation $\mathcal{D}'_{\mathcal{P}}$ with probability mass $0 < |\mathcal{D}'_{\mathcal{P}}| \leq 1$ in the larger population $\mathcal{D}_{\mathcal{P}}$, we say that it is decisive over $c > c'$ in issue i if when $\mathcal{D}'_{\mathcal{P}}$ is unanimous on $c > c'$, whatever the preference specification $\mathcal{D}_{\mathcal{P}}$ of the whole population is, $c \succ_{f(S)} c'$ with probability $1 - e^{-\Omega(|S|)}$. We say that it is weakly decisive if when $\mathcal{D}'_{\mathcal{P}}$ is unanimous on $c > c'$ and the rest of the population is unanimous on $c' > c$, $c \succ_{f(S)} c'$ with probability $1 - e^{-\Omega(|S|)}$.*

REMARK C.2. *Given the guaranteed anonymity of our problem setting, all subpopulations of the same probability mass are entirely interchangeable, and thus the decisiveness of a subpopulation over a specific preference only depends on the subpopulation's mass.*

LEMMA C.3 (FIELD EXPANSION LEMMA). *Given any representational mechanism f satisfying PPE and S-PIIA, consider outcomes u, v, w of issue $i \in \mathcal{I}$ such that all 6 pairwise orderings among them are privileged, and a subpopulation $\mathcal{D}'_{\mathcal{P}}$ that is weakly decisive over $u > v$. We then have that $\mathcal{D}'_{\mathcal{P}}$ is decisive over $u > w$.*

Note that by analogy, it will also be decisive over $w > v$. With repeated application and transitivity, it follows that $\mathcal{D}'_{\mathcal{P}}$ is decisive over all 6 pairwise orderings among u, v, w .

PROOF. Assume that $\mathcal{D}'_{\mathcal{P}}$ is unanimous on $u > w$. By S-PIIA, we can assume that $\mathcal{D}'_{\mathcal{P}}$ is further unanimous on $u > v > w$, while the rest of the population is unanimous on $v > u$ and $v > w$. This does not affect $f(S) \upharpoonright_{\text{LO}(\{u, w\})}$.

By weak decisiveness, we have $u \succ_{f(S)} v$ with probability $1 - e^{-\Omega(|S|)}$. Due to the privilegedness of $v > w$, PPE can be applied to every extension of the opposite ordering $w > v$ (where the extension serve as C'), guaranteeing that $v \succ_{f(S)} w$ with probability $1 - e^{-\Omega(|S|)}$. Taken together, we have $u \succ_{f(S)} w$ with probability $1 - e^{-\Omega(|S|)}$, and thus $\mathcal{D}'_{\mathcal{P}}$ is decisive over $u > w$. The rest follows. $\square$

Finally, we review a few graph theoretical concepts that will prove useful in the proof of the strong representative impossibility theorem.

DEFINITION C.4 (STRONGLY CONNECTED COMPONENT). *For a directed graph G , a strongly connected component (SCC) is a maximal subgraph of G (including those of size 1) in which there is a directed path between every ordered pair of vertices. The SCCs of G form a partition of the vertex set of G .*

DEFINITION C.5 (CONDENSATION). *The condensation $\text{cond}(G)$ of a directed graph G is a directed acyclic graph (DAG) where each vertex represents an SCC of G , and there is an edge from SCC A to SCC B iff there is an edge from a vertex in A to a vertex in B in G .*

REMARK C.6. *Intuitively, SCC is the unit of bidirectional connectivity in a directed graph, and condensation is a way to simplify a directed graph into a DAG by contracting SCCs into single vertices. There is a directed path from u to v in G iff there is a directed path from the SCC containing u to the SCC containing v in $\text{cond}(G)$. Both SCC and condensation can be computed in linear time in the number of vertices and edges of the graph, using algorithms such as Tarjan's algorithm and Kosaraju's algorithm.*

Note that given the transitivity of the privilege graph, any SCC in the privilege graph is a complete digraph.

C.7 Proof of Theorem 4

PROOF. We first show that issues with cyclic privilege graphs imply the impossibility of a mechanism satisfying the axioms. We assume the existence of such an f and show that it leads to a contradiction.

For any cyclic privileged issue i , by Lemma 6.13, we know that any directed simple cycle in G_i induces a complete digraph on the outcomes in the cycle. We arbitrarily pick outcomes u, v, w from the cycle. Arbitrarily partition the population into equal portions $\mathcal{D}^a_{\mathcal{P}}$ and $\mathcal{D}^b_{\mathcal{P}}$ of probability mass $\frac{2}{3}$ and $\frac{1}{3}$ respectively. Let $\mathcal{D}^a_{\mathcal{P}}$ be unanimous on $u > v$ and $\mathcal{D}^b_{\mathcal{P}}$ be unanimous on $v > u$. By S-PC, the representational mechanism f must converge upon $u > v$ or $v > u$ with probability $1 - e^{-\Omega(|S|)}$ in the full population.

Assume that f converges upon $u > v$, in which case $\mathcal{D}^a_{\mathcal{P}}$ is weakly decisive over $u > v$. By the field expansion lemma, $\mathcal{D}^a_{\mathcal{P}}$ is decisive over $u > v$. We further partition $\mathcal{D}^a_{\mathcal{P}}$ into 1 : 2 portions $\mathcal{D}^{a,1}_{\mathcal{P}}$ and $\mathcal{D}^{a,2}_{\mathcal{P}}$. Then, we let:

- $\mathcal{D}^{a,1}_{\mathcal{P}}$ (probability mass $\frac{2}{9}$) be unanimous on $u > v > w$,
- $\mathcal{D}^{a,2}_{\mathcal{P}}$ (probability mass $\frac{4}{9}$) be unanimous on $w > u > v$, and
- $\mathcal{D}^b_{\mathcal{P}}$ (probability mass $\frac{1}{3}$) be unanimous on $v > w > u$.

Since $\mathcal{D}^a_{\mathcal{P}}$ is decisive, we have $u >_{f(S)} v$ with probability $1 - e^{-\Omega(|S|)}$. Since the probability masses of all subpopulations have denominators 3 or 9 (both odd numbers), there can be no ties, and f must converge due to S-PC. Given that $u >_{f(S)} v$, with probability $1 - e^{-\Omega(|S|)}$, we have either $u >_{f(S)} w$ or $w >_{f(S)} v$. In the former case, $\mathcal{D}^{a,1}_{\mathcal{P}}$ is weakly decisive over $u > w$, and in the latter case, $\mathcal{D}^{a,2}_{\mathcal{P}}$ is weakly decisive over $w > v$. In either case, the subpopulation is decisive by the field expansion lemma, but its probability mass ($\frac{2}{9}$ or $\frac{4}{9}$) is smaller than 0.5, contradicting PND (Lemma 6.5).

We can similarly show that f converging upon $v > u$ leads to a contradiction. Therefore, no such f can exist. Note that in the proof above, we have been misusing $|S|$ to denote the number of samples that falls into issue i , which is not a problem since i 's probability mass $\Pr_{\mathcal{D}_I}[i] > 0$ is a positive constant (we can pick the i with the largest probability mass, which makes the mass only dependent on $\mathcal{D}_I$ itself), and the asymptotic behavior of the convergence probability is the same.

Then, assuming that all issues (except possibly those with zero probability in $\mathcal{D}_I$) have specific non-cyclic privilege graphs, we construct a candidate space C consistent with the privilege graphs, and a mechanism f that satisfies the axioms for all $\mathcal{D}_{\mathcal{P}}$.

First consider the case where $|I| = 1$, $I = \{i\}$. Since $\phi(i)$ is transitive (Lemma 6.13) but doesn't have simple cycles of length at least 3, it follows that $\phi(i)$ doesn't have any SCC containing 3 or more vertices. Fix an arbitrary topological order $(g_1, \dots, g_m)$ of vertices in the DAG $\text{cond}(\phi(i))$, where $\{g_j\}_{j=1}^m$ constitute a partition of $[N]$, and $|\{g_j\}| \in \{1, 2\}$ for each j ; let $g_{k_1}, \dots, g_{k_l}$ be the SCCs of size 2. Now consider translating $(g_1, \dots, g_m)$ into an ordering in $\text{LO}(N)$, where the vertices in each SCC are ordered arbitrarily, and vertices in different SCCs are ordered according to the topological order $(g_1, \dots, g_m)$. There are 2^l ways to perform such translation, and we define $C \subset \text{LO}(N)$ to be the set of these 2^l orderings. It can be verified that the G_i resulting from C contains $\phi(i)$ as a subgraph.

We then define the mechanism f . For each $g_{k_j} = \{u_j, v_j\}$ ($1 \leq j \leq l$), it examines if a majority in the samples prefer u_j over v_j . If yes, it outputs the ordering that places u_j above v_j ; otherwise, it outputs the ordering that places v_j above u_j . Ties are broken arbitrarily. Let us now verify that f satisfies PPE, S-PIIA, and S-PC for all $\mathcal{D}_{\mathcal{P}}$.

- PPE: Relative orderings between outcomes in different SCCs are fixed by the topological order, and therefore the outcomes c, c' in the PPE axiom must be in the same SCC. f resolves this preference using a majority vote, which, by Hoeffding's inequality, has a convergence probability of $1 - e^{-\Omega(|S|)}$.
- S-PIIA: It can be verified that if $c > c'$ and $c' > c$ are both privileged orderings in i , then c, c' must be in the same SCC. f resolves this preference using a majority vote that considers only the preference between c and c' , and thus satisfies S-PIIA.
- S-PC: Again, c, c' must be in the same SCC. f resolves this preference using a majority vote, which, by Hoeffding's inequality, has a convergence probability of $1 - e^{-\Omega(|S|)}$ when the population is non-uniform on c, c' .

In the general case where $|\mathcal{I}| > 1$, we can apply the above construction to each issue independently, and define C to be the Cartesian product of the constructed candidate spaces. The mechanism f is then defined to apply the constructed mechanism for each issue independently. It can be verified that f satisfies PPE, S-PIIA, and S-PC for all $\mathcal{D}_{\mathcal{P}}$. $\square$

Received 02 November 2024; accepted 17 July 2026