Flexible level-1 consensus ensuring stable social choice: analysis and algorithms
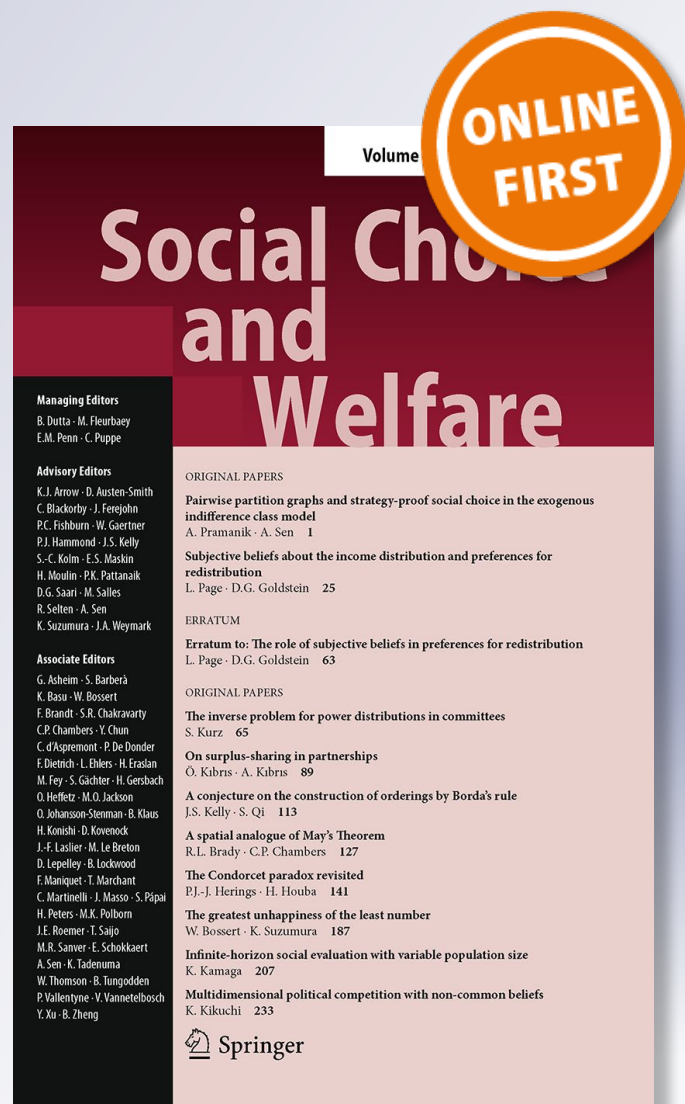
Mor Nitzan, Shmuel Nitzan & Erel Segal-Halevi

🦄 Springer

Springer

CrossMark

ORIGINAL PAPER

Flexible level-1 consensus ensuring stable social choice: analysis and algorithms

Mor Nitzan[1,2,3] · Shmuel Nitzan[4] ·
Erel Segal-Halevi[5,6]

Abstract *Level-1 consensus* is a recently-introduced property of a preference-profile. Intuitively, it means that there exists a preference relation which induces an ordering of all other preferences such that frequent preferences are those that are more similar to it. This is a desirable property, since it enhances the stability of social choice by guaranteeing that there exists a Condorcet winner and it is elected by all scoring rules. In this paper, we present an algorithm for checking whether a given preference profile exhibits level-1 consensus. We apply this algorithm to a large number of preference profiles, both real and randomly-generated, and find that level-1 consensus is very improbable. We support these empirical findings theoretically, by showing that, under the impartial culture assumption, the probability of level-1 consensus approaches zero when the number of individuals approaches infinity. Motivated by these observations, we show that the level-1 consensus property can be weakened while retaining its stability implications. We call this weaker property *Flexible Consensus*. We show, both empirically and theoretically, that it is considerably more probable than the original level-1 consensus. In particular, under the impartial culture assumption, the probability

✉ Shmuel Nitzan
nitzans@mail.biu.ac.il

1   Present Address: School of Computer Science, The Hebrew University, 91904 Jerusalem, Israel

2   Racah Institute of Physics, The Hebrew University, 91904 Jerusalem, Israel

3   Department of Microbiology and Molecular Genetics, Faculty of Medicine,
     The Hebrew University, 91120 Jerusalem, Israel

4   Department of Economics, Bar-Ilan University, 5290002 Ramat Gan, Israel

5   Department of Computer Science, Ariel University, 40700 Ariel, Israel

6   Department of Computer Science, Bar-Ilan University, 5290002 Ramat Gan, Israel

🙋 Springer

for Flexible Consensus converges to a positive number when the number of individuals approaches infinity.

1 Introduction

Recently, Mahajne et al. (2015) have proposed the concept of *level-1 consensus* of a preference profile, showing that it considerably enhances the stability of social choice. In particular, if a profile exhibits level-1 consensus around a given preference-relation $\succ_0$ with respect to the inversion metric, then:[1]

- There exists a Condorcet winner;
- The Condorcet winner is chosen by all scoring rules;
- With an odd number of individuals, the majority relation is transitive and coincides with $\succ_0$.

The current study starts by considering two questions:

(1) How can a profile be tested for level-1 consensus?
(2) How likely is it that level-1 consensus exists?

Questions of the former type have been recently studied with respect to various domain restrictions. For example, Escoffier et al. (2008) provide an efficient way to check whether a profile is *single-peaked*, Bredereck et al. (2013) provide an efficient way to check whether a profile is *single-crossing*, and Barberà and Moreno (2011) ask whether the satisfaction of their proposed *top monotonicity* condition (a sufficient condition for an extension of the median-voter theorem to hold) is easy to check. See Elkind et al. (2016) for a survey of recent developments in algorithmic checking of domain restrictions.

Questions of the latter type have been studied in the social choice literature with respect to various domain restrictions that guarantee social stability, e.g., the existence of Condorcet winners under the majority rule (Gehrlein 1981; Tsetlin et al. 2003) and the likelihood of single-peaked preferences (Lackner and Lackner 2017).

Our answer to the first question is an efficient algorithm for determining whether a preference-profile exhibits level-1 consensus. In case such a consensus exists, the algorithm identifies a preference relation around which it occurs.

Our answer to the second question is that level-1 consensus is highly improbable. We applied our algorithm on a recently-released dataset of 315 real-world profiles from various sources (Mattei and Walsh 2013) and found that none of them exhibits level-1-consensus. Moreover, experiments performed on thousands of profiles generated randomly according to Mallows' phi model (Mallows 1957) revealed that, for a wide range of parameter settings, profiles exhibiting level-1 consensus were extremely rare. As a partial explanation to these findings, we prove that under the standard probabilistic

---

[1] In fact, Mahajne et al. (2015) define a family of conditions called level-r consensus, where r is an integer between 1 and $K!/2$ and K is the number of alternatives . But for the sake of simplicity, in the present paper we focus on level-1 consensus which is the strongest condition in this family.

Note that recently Poliakov (2016) proved that level-r consensus is equivalent to level-1 consensus whenever $r \leq (K-1)!$.

setting of equally-probable preference relations, the probability of level-1 consensus goes to zero when the number of individuals goes to infinity.

Motivated by these results, we found a way to weaken the level-1 consensus property while keeping its stability properties. We call the weakened property *Flexible Consensus*. In the above mentioned dataset, we found that 39 out of 315 profiles exhibit Flexible Consensus. Flexible Consensus is also much more probable in the settings of the randomly-generated profiles we tested. In particular, under the impartial-culture assumption, the probability of Flexible Consensus is lower-bounded by a positive constant even when the number of individuals goes to infinity.

2 Definitions

Let $A = \{a_1, \ldots, a_K\}$ be a set of $K \geq 3$ alternatives and $N = \{1, \ldots, n\}$ a set of voters. We assume that each voter has a *strict total order* on the alternatives, i.e, for each two different alternatives a and b, either the voter strictly prefers a to b or the voter strictly prefers b to a. Let $\mathcal{P}$ be the subset of strict total orders on A. We will refer to the elements of $\mathcal{P}$ as *preference relations* or simply as *preferences*.

A *preference profile* or simply a *profile* is a list $\pi = (\succ_1, \ldots, \succ_n)$ of preference relations on A. For each $i \in N$, $\succ_i$ is the preference relation of voter i. We denote by $\mathcal{P}^n$ the set of all possible profiles.

Let $\pi = (\succ_1, \ldots, \succ_n)$ be a profile. For each preference $\succ \in \mathcal{P}$, let $\mu_\pi(\succ) := |\{i \in N : \succ_i = \succ\}|$ = the number of voters whose preference is $\succ$, which in this paper is referred to as the *frequency* of $\succ$.

Definition 1 The *inversion distance* between two preferences $\succ, \succ'$, denoted $d(\succ, \succ')$, is the number of pairs of alternatives that are ranked differently by the two preferences, i.e, the number of sets $\{a, b\} \subseteq A$ such that $a \succ b$ and $b \succ' a$ or vice-versa.

The inversion-distance is a metric on $\mathcal{P}$ (Kemeny and Snell 1962). It can vary between 0 and $\binom{K}{2}$, the number of subsets of two alternatives.

For example, if there are three alternatives and $a_1 \succ a_3 \succ a_2$ and $a_2 \succ' a_3 \succ' a_1$, then $d(\succ, \succ') = 3$ since all three pairs of alternatives are ranked differently by $\succ$ and $\succ'$.

The following definition is due to Mahajne et al. (2015).

Definition 2 Let $\succ_0 \in \mathcal{P}$. A profile $\pi \in \mathcal{P}^n$ exhibits *consensus of level-1 around* $\succ_0$ if the following two conditions hold:

(1) For all pairs of preferences $\succ, \succ' \in \mathcal{P}$, $d(\succ, \succ_0) \leq d(\succ', \succ_0)$ implies $\mu_\pi(\succ) \geq \mu_\pi(\succ')$.
(2) There is at least one pair $\succ, \succ' \in \mathcal{P}$, such that $d(\succ, \succ_0) < d(\succ', \succ_0)$ and $\mu_\pi(\succ) > \mu_\pi(\succ')$.

3 Algorithm for detecting level-1 consensus

Algorithm 1 Detection of Level-1 Consensus

INPUT:

- A set of K alternatives, $A = \{a_1, \ldots, a_K\}$.
- A profile π containing n preference-relations on A (possibly with duplicates).

OUTPUT:

- If π exhibits level-1 consensus around some preference $\succ_0$, return $\succ_0$.
- Otherwise, return "no consensus".

ALGORITHM:

1. Calculate the frequency $\mu_\pi(\succ)$ of each preference $\succ \in \pi$.
2. Define $n' =$ number of distinct preferences in π.
3. If all frequencies are equal and $n' = K!$, return "no consensus; Condition 2 violated".
4. Order the preferences by descending frequency: $\mu_\pi(\succ_1) \geq \mu_\pi(\succ_2) \geq \cdots \geq \mu_\pi(\succ_{n'})$.
5. Set $M := \mu_\pi(\succ_1)$ to be the maximum frequency of a preference relation.
6. For $j = 1, 2, \ldots n'$ while $\mu_\pi(\succ_j) = M$:
   Check if Condition 1 is satisfied for π and $\succ_j$ (see Algorithm 2). If it is, then return $\succ_j$.
7. Return "no consensus; Condition 1 is violated for all candidates".

Algorithm 2 Check if Condition 1 is satisfied

INPUT:

- A profile π containing n preference-relations on A (possibly with duplicates) ordered by descending frequencies.
- A preference $\succ_0$.

OUTPUT: "True" if Condition 1 is satisfied for profile π with respect to $\succ_0$. "False" otherwise.
ALGORITHM:

1. Within every group of preferences with the same frequency, order the preferences by ascending inversion-distance from $\succ_0$.
2. For $i = 1, 2, \ldots, n' - 1$:
   If $\mu_\pi(\succ_i) > \mu_\pi(\succ_{i+1})$ and $d(\succ_i, \succ_0) \geq d(\succ_{i+1}, \succ_0)$, return False.
3. Set $\widehat{d} := d(\succ_{n'}, \succ_0)$.　　　　If $\sum_{j=0}^{\widehat{d}} |T(K, j)| = n$, return True.　　　　Else, return False.

Given a profile $\pi \in \mathcal{P}^n$, we would like to check whether there exists some preference-relation $\succ_0$ such that π exhibits level-1 consensus around it. Our solution is given by Algorithm 1 (using, as a sub-routine, Algorithm 2).

Theorem 1 *Algorithm 1 checks whether a profile π exhibits level-1 consensus in time:*

$$O(n'^2 K \sqrt{\log K} + n'^2 \log n')$$

where $n' \leq \min(K!, n)$ is the number of distinct preferences in π. In particular, the run-time is polynomial in the profile size.

This entire section is devoted to proving Theorem 1. We first explain why Algorithm 1 is correct. Then we analyze its run-time.

Algorithm 1 is based on the simple observation that the two conditions in Definition 2 are equivalent to the following:

(Condition 1)   For all $\succ, \succ' \in \mathcal{P}$, if $\mu_\pi(\succ') > \mu_\pi(\succ)$, then $d(\succ', \succ_0) < d(\succ, \succ_0)$.

(Condition 2)   There exists a pair $\succ, \succ' \in \mathcal{P}$ such that $\mu_\pi(\succ') > \mu_\pi(\succ)$.

The algorithm proceeds in several steps.

First, we calculate the frequency $\mu_\pi(\succ)$ of each of the preferences $\succ \in \pi$. Let n' be the number of distinct preferences in π. Note that $n' \leq n$ and also $n' \leq K!$, since with K alternatives there are at most $K!$ possible preferences. Now Condition 2 is easily checked: it is satisfied if-and-only-if (a) there exists a pair of preferences in π with different frequencies, or (b) $n' < K!$ (since this implies that there exists a preference not in π with frequency 0).

If Condition 2 is satisfied, it only remains to check whether Condition 1 is satisfied as well.

We order the preferences in descending order of $\mu_\pi(\succ)$, and rename them $\succ_1, \succ_2, \ldots, \succ_{n'}$, such that $\mu_\pi(\succ_1) \geq \mu_\pi(\succ_2) \geq \cdots \geq \mu_\pi(\succ_{n'})$. This enables us to identify the candidates for level-1-consensus. Since $d(\succ_0, \succ_0) = 0$, Condition 1 immediately implies that each candidate for level-1 consensus must be a preference with maximal frequency. So, the candidates are $\succ_1, \succ_2, \ldots, \succ_h$ such that $h \leq n'$ is the largest index for which $\mu_\pi(\succ_1) = \mu_\pi(\succ_2) = \cdots = \mu_\pi(\succ_h)$.

Now we can directly check Condition 1. This condition should be checked separately for each candidate preference $\succ_0$. Given a candidate-preference $\succ_0$, we can calculate its inversion-distance from each preference $\succ_i \in \pi$, $d(\succ_i, \succ_0)$. Now, we represent the profile π relative to $\succ_0$ in the form of a scatter-plot, which will lead to a straight-forward assessment of the profile's consensus status. Our scatter-plot is a plot whose x-axis denotes the distance $d(\succ_i, \succ_0)$ and whose y-axis denotes the frequency $\mu_\pi(\succ_i)$. Note that there may be several different preference relations $\succ_i$ with the same frequency, $\mu_\pi(\succ_i) = m$. Therefore, for each integer value on the y-axis of the scatter-plot, m, we may have several corresponding values on the x-axis, which can be represented by a horizontal segment whose maximum and minimum borders are given by $\max_{i:\mu_\pi(\succ_i)=m} d(\succ_i, \succ_0)$ and $\min_{i:\mu_\pi(\succ_i)=m} d(\succ_i, \succ_0)$, respectively.

Condition 1 above requires that, for every two frequencies $m_1 > m_2$, *all* preferences with frequency m_1 are closer to $\succ_0$ than *all* preferences with frequency m_2: $\max_{i:\mu_\pi(\succ_i)=m_1} d(\succ_i, \succ_0) < \min_{i:\mu_\pi(\succ_i)=m_2} d(\succ_i, \succ_0)$. Graphically (see Fig. 1), this means that when we scan the scatter plot from top to bottom, we must see non-overlapping intervals ordered strictly from left to right.

Three examples are shown in Fig. 1. The left example is positive: there are five non-overlapping intervals (two of which consist of a single point), and when they are scanned from top to bottom, they are ordered strictly from left to right. Therefore Condition 1 holds. The middle and right examples are negative: the second and third intervals from the top overlap. For instance, in the middle example the overlap is in a single point, $x = 2$. This point corresponds to two distinct preferences with different frequencies (5 and 4), both of which are at distance 2 from the candidate preference; these preferences violate Condition 1.
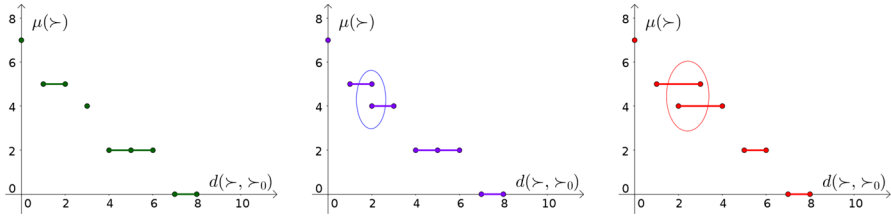
Fig. 1 Scatter-plot representation of profiles relative to a candidate preference enables to graphically check whether level-1 consensus around that candidate is satisfied. The left plot satisfies Condition 1 since the horizontal segments are decreasing and non-overlapping. The middle and right plots violate Condition 1 since there are overlapping segments (circled)

The process of 'scanning the scatter plot from top to bottom' can be formalized as follows (see Algorithm 2). Order the list of preferences lexicographically by two criteria: the primary criterion is descending frequency (as before), and the secondary criterion is ascending distance. So the preferences are partitioned to equivalence classes by their frequency: the classes are ordered by descending frequency, and within each equivalence-class, the preferences are ordered by ascending distance from $\succ_0$. Preferences with both the same frequency and the same distance are ordered arbitrarily. Under this ordering, the following lemma holds:

Lemma 1 *If Condition 1 is violated for* any *pair of preference relations in π, then it is violated for an adjacent pair $\succ_i, \succ_{i+1}$ for some i.*

The lemma is easy to understand based on the graphical criterion outlined above. A formal proof is provided in Appendix A.

Lemma 1 implies that in order to ensure that Condition 1 is satisfied for all preferences in π, it is sufficient to scan the ordered list of preferences from $\succ_1$ to $\succ_{n'}$, and check if there is some i such that $\mu_\pi(\succ_i) > \mu_\pi(\succ_{i+1})$ and $d(\succ_i, \succ_0) \geq d(\succ_{i+1}, \succ_0)$.

If Condition 1 holds for all preferences in π, it remains to check that it holds for preferences not in π, i.e, preferences with zero frequency. Let $\widehat{d}$ be the largest distance from a preference in π to $\succ_0$, i.e, $\widehat{d} = d(\succ_{n'}, \succ_0)$. Condition 1 implies that, if $\mu_\pi(\succ) = 0$, then $d(\succ, \succ_0) > \widehat{d}$. Therefore, we have to check that the distances between $\succ_0$ and preferences not in π are all larger than $\widehat{d}$. Equivalently, we can ensure that all preferences with distance at most $\widehat{d}$ are in π. This can be checked by calculating the *number* of possible preferences with distance at most $\widehat{d}$, and verifying that it is equal to the total number n' of distinct preferences in π. Since this number involves all possible preferences, it does not depend on the candidate $\succ_0$. Therefore, we can calculate this number assuming w.l.o.g that $\succ_0$ is the preference defined by $1 \succ_0 2 \succ_0 \ldots \succ_0 K$. So we have to calculate the number of permutations on K elements with at most $\widehat{d}$ *inversions* (out-of-order elements). This can be written as:

$$\sum_{j=0}^{\widehat{d}} |T(K, j)|$$

where $T(K, j)$ is the set of permutations on K elements with *exactly* j inversions.[2]
The number $|T(K, j)|$ can be calculated using the following recurrence relation

- $\forall K : |T(K, 0)| = 1$, since there is exactly one permutation with zero inversions—the identity permutation.
- $\forall j : |T(0, j)| = 0$, since there are no permutations with 0 elements.
- $\forall K, j : |T(K, j)| = \sum_{i=0}^{\min(K-1,j)} |T(K-1, j-i)|$: for any permutation of $1, \ldots, K$ with j inversions, let i be the number of elements that come after element K in that permutation. Since K is larger than all other elements, there are exactly i inversions involving K. Therefore, if we remove K, we get a permutation of $1, \ldots, K-1$ with exactly $j-i$ inversions. By summing the counts of these permutations for all possible values of i (namely, $i \geq 0, i \leq j, i \leq K-1$) we get $|T(K, j)|$.

The algorithm for detecting Level-1 Consensus is summarized in Algorithm 1.
We complete the proof of Theorem 1 by a run-time analysis.

Lemma 2 *The run-time of Algorithm 1 is:*

$$O(n'^2 K \sqrt{\log K} + n'^2 \log n')$$

where $n' \leq \min(K!, n)$ is the number of distinct preferences in π.

Proof We first analyze Algorithm 2. It has to calculate the distance between $\succ_0$ and each of the other $n' - 1$ preferences. Calculating the inversion distance between a given pair of preferences can be done by a recently-developed algorithm (Chan and Pătraşcu 2010) with a runtime of $O(K \sqrt{\log K})$. We then have to order the n' distinct preferences and then scan them from top to bottom. Ordering n' items can be done in time $O(n' \log n')$. The value of n' is at most the maximum of n (the number of voters) and $K!$ (the number of possible preferences). So the run-time of Algorithm 2 is

$$O(n' K \sqrt{\log K} + n' \log n')$$

As will be explained in the next section, the probability of having two preferences exhibiting exactly the same frequency is low, so in most cases we will have to apply Algorithm 2 only once. However, in the improbable case in which there are many preferences with the same frequency, we would have to apply it at most n' times. Therefore, the worst-case run-time of Algorithm 1 is n' times the run-time of Algorithm 2. □

---

[2] $|T(K, j)|$ is also known as the *Mahonian number*; see OEIS sequence A008302, https://oeis.org/A008302.

4 Probability of level-1 consensus

Equipped with a procedure for checking level-1 consensus, we set out to check how likely is this property in various settings. We conducted several simulation experiments.

In the first experiment we used the PrefLib database (Mattei and Walsh 2013), an online database of real-world preference-profiles collected from various sources. This database contains 315 full profiles, with different numbers of alternatives and voters; see Table 1 for statistics. For each profile, we used the algorithm described in the previous section to check whether there exists a 1-level consensus. The results were striking: none of the 315 profiles exhibited a level-1 consensus.

In the second experiment we used preference-profiles that were generated according to *Mallows' phi model* (Mallows 1957), which was claimed to favor level-1 consensus (Mahajne et al. 2015). Mallows' model assumes that there is a "correct" preference $\succ_*$, and the actual preferences of the voters are noisy variants of it. The probability of a preference $\succ$ depends on its inversion distance from the correct preference: $d(\succ, \succ_*)$. The strength of this dependence is determined by a parameter $\phi \in (0, 1]$, where lower ϕ means higher dependence; when $\phi \to 0$ the preferences of all voters are identical and equal to $\succ_*$, while when $\phi = 1$ the preference of each voter is selected uniformly at random from the $K!$ possible orderings on K items. In general, the probability of each preference-relation $\succ$ is given by (Lu and Boutilier 2014):

$$\text{Prob}\left[\ \succ\ \mid\ \phi, \succ_* \right] = \frac{1}{Z} \cdot \phi^{d(\succ, \succ_*)}$$

where Z is a normalization factor.

We considered all six combinations of $K \in \{3, 4, 5\}$ alternatives and $n \in \{100, 1000\}$ voters, where ϕ varied between 0 and 1. For each combination of K, n, ϕ we ran 1000 experiments and calculated (a) the fraction of profiles that exhibit level-1 consensus, (b) the fraction of profiles that are single-peaked,[3] and (c) the fraction of profiles that exhibit *Flexible Consensus*, which will be presented in the next section. The results are shown in Fig. 2.

As can be expected, consensus always exists when $\phi = 0$, since in this case there is a deterministic consensus on the "true" preference. Additionally, when there are 100 voters and 3 alternatives and ϕ is small, a small positive percentage of the profiles exhibit a level-1 consensus (top left plot). In all other cases, the percentage of level-1 consensus profiles drops to 0 when $\phi \geq 0.05$.
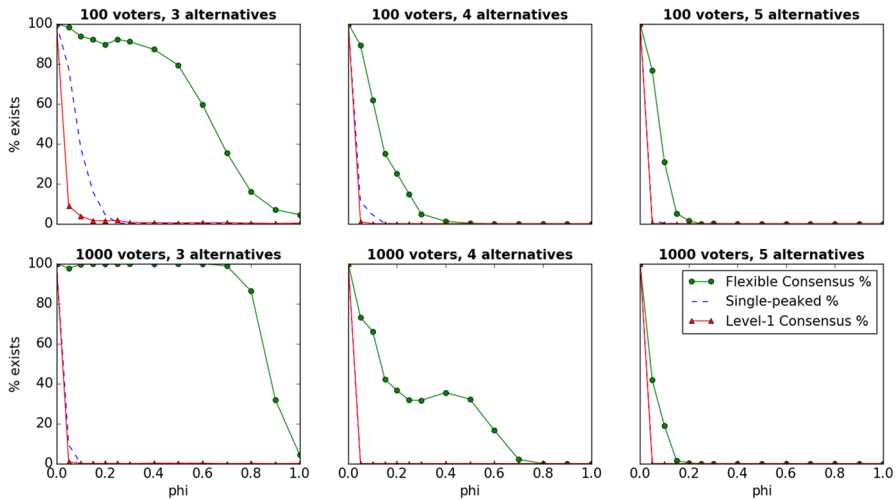
Why is level-1 consensus so rare?

Intuitively, the reason is that it requires groups of preferences to have exactly the same frequency in the population. Condition 1 implies that if $d(\succ_i, \succ_0) = d(\succ_j, \succ_0)$ then $\mu_\pi(\succ_i) = \mu_\pi(\succ_j)$. For every $K \geq 3$ and for every candidate $\succ_0$, there exist at least two preferences with the same distance from the candidate, $d(\succ_i, \succ_0) = d(\succ_j, \succ_0)$. Hence, a necessary condition for level-1 consensus is that there exist at

---

[3] This calculation was done for the sake of comparison. It was implemented using Nicholas Mattei's PrefLib tools, which are freely available at GitHub: https://github.com/nmattei/PrefLib-Tools.

Table 1 Summary of PrefLib (Mattei and Walsh 2013) data-sets used in our experiments

Code	Description	# Profiles	# Alternatives	# Voters
ED-00004 1-100: Netflix prize data (Bennett and Lanning 2007)	Rankings of movies by consumers	100	3	100–1000
ED-00004 101-200: Netflix prize data (Bennett and Lanning 2007)	Rankings of movies by consumers	100	4	100–1000
ED-00006: Skate data	Ranking of skaters by judges in competitions	20	10–25	8–10
ED-00009: AGH course selection	Ranking of courses by university students	2	7–9	≈ 150
ED-00011: Web search	Ranking of search-phrases by search-engines	3	100–250	5
ED-00012: T shirt	Ranking of T-shirt designs by researchers	1	10	30
ED-00014: Sushi data	Ranking of sushi kinds by consumers	1	10	5000
ED-00014: Sushi Data	Ranking of sushi kinds by consumers	1	10	5000
ED-00015: Clean web search	Ranking of search-phrases by search-engines	79	10–250	4
ED-00024: Mechanical turk dots (Mao et al. 2013)	Ranking of dot-sets by Amazon-Turk workers	4	4	≈ 800
ED-00025: Mechanical turk puzzle (Mao et al. 2013)	Ranking of puzzles by Amazon-Turk workers	4	4	≈ 800
ED-00032: Education surveys	Ranking of issues by informatics students	1	6	15

Fig. 2 Percentage of profiles, from a set of profiles selected at random according to Mallows' phi model, which exhibit level-1-consensus, single-peakedness, or Flexible Consensus (the latter is defined in Sect. 5). Note that the graphs of level-1 consensus and single-peakedness are almost overlapping in all plots except the top-left

least two preferences with exactly the same frequency. As the number of voters goes to ∞, the probability that any specific preference-relation appears with any specific frequency goes to 0. Therefore, the probability that two preference-relations have the same frequency goes to 0 too.

To formalize this intuition, we present an asymptotic upper bound on the probability of level-1 consensus for the case $\phi = 1$. This is the case of impartial culture, in which all $K!$ preferences are equally probable.

We select a preference-profile π according to the following random process, parametrized by K (the number of alternatives) and an integer-valued parameter m.

- Let $(\succ_i)_{i=1}^{K!}$ be an enumeration of the preferences in $\mathcal{P}$.
- For each i, draw a number B_i according to a binomial distribution with m trials and success-probability $1/K!$.
- Return a profile in which, for every i, there are B_i voters whose preference is $\succ_i$.

All the B_i's are i.i.d. random variables with mean value $\mu := \frac{m}{K!}$. The total number of preferences in the profile π is $n = \sum_{i=1}^{K!} B_i$; this is also a random variable, and its mean value is $E[n] = \sum_{i=1}^{K!} \frac{m}{K!} = m$. The process is symmetric with respect to the preferences in $\mathcal{P}$, so all preferences are equally likely, in accordance with the impartial-culture assumption.

Define $P^{consensus}(m, K)$ as the probability that the above random process yields a profile that exhibits Level-1 Consensus. The rest of the section is devoted to proving the following proposition:

Proposition 1 *For every K:*

$$\lim_{m \to \infty} P^{consensus}(m, K) = 0$$

Proof We first show an upper bound on the probability for level-1 consensus around a fixed preference.

Let $\succ_0$ be a fixed preference. Denote by $P^{\text{con-fixed}}(m, K)$ the probability that a profile π, selected according to the random process above with parameters m, K, exhibits Level-1 Consensus around $\succ_0$. Since all preferences are equally probable, $P^{\text{con-fixed}}(m, K)$ is the same for all $\succ_0$. We now present an approximate upper bound on $P^{\text{con-fixed}}(m, K)$.

Recall that $T(K, d)$ is the set of distinct preferences whose inversion-distance from $\succ_0$ is exactly d, where d can vary between 0 and $\binom{K}{2}$ (the number of pairs of K alternatives). The conditions for Level-1 Consensus imply that all preferences in $T(K, d)$ must have the same frequency in π, i.e:

$$\text{For all } i, j \text{ such that } \succ_i \in T(K, d) \text{ and } \succ_j \in T(K, d): \quad B_i = B_j \qquad (1)$$

Let $P^{\text{equal}}(m, p, t)$ be the probability that t i.i.d. random variables distributed like $Binom[m, p]$ are equal. Then, (1) implies that:

$$P^{\text{con-fixed}}(m, K) \leq \prod_{d=0}^{\binom{K}{2}} P^{\text{equal}}\left(m, \frac{1}{K!}, |T(K, d)|\right) \qquad (2)$$

In Appendix C, we show that $P^{\text{equal}}(m, p, t)$ can be approximated as:

$$P^{\text{equal}}(m, p, t) \approx \frac{1}{(\sqrt{2\pi pqm})^{t-1}}$$

where $q = 1 - p$ and the symbol $\approx$ means that the ratio of the expressions in its two sides goes to 1 as $m \to \infty$. Therefore, (2) can be approximated as:

$$P^{\text{con-fixed}}(m, K) \leq \approx 1 / \sqrt{\prod_{d=0}^{\binom{K}{2}} \left(2\pi \frac{1}{K!} m\right)^{[|T(k,d)|-1]}}$$

$$= 1 / \sqrt{\left(\frac{2\pi m}{K!}\right)^{\sum_{d=0}^{\binom{K}{2}}[|T(k,d)|-1]}} \qquad (3)$$

The sum in the exponent can be simplified as follows:

$$\sum_{d=0}^{\binom{K}{2}}[|T(k, d)| - 1] = \left[\sum_{d=0}^{\binom{K}{2}} |T(k, d)|\right] - \binom{K}{2} - 1$$

The sum of $|T(K, d)|$ over all possible values of d (i.e, d between 0 and $\binom{K}{2}$) equals the total number of different preferences over K items, which is $K!$:

$$\sum_{d=0}^{\binom{K}{2}} |T(k, d)| = K! \tag{4}$$

Substituting in (3) gives:

$$P^{\text{con-fixed}}(m, K) \leq \approx 1 / \sqrt{\left(\frac{2\pi m}{K!}\right)^{K! - \binom{K}{2} - 1}}$$

Recall that $P^{\text{con-fixed}}(m, K)$ is the probability of level-1 consensus around a *fixed* preference. The probability of level-1 consensus around *any* preference is, by the union bound, at most this probability times the number of possible preferences, i.e,

$$P^{\text{consensus}}(m, k) \leq K! \cdot P^{\text{con-fixed}}(m, K)$$

$$\leq \frac{K!}{\sqrt{\left(\frac{2\pi m}{K!}\right)^{K! - \binom{K}{2} - 1}}}$$

so for every fixed $K \geq 3$, $\lim_{m \to \infty} P^{\text{consensus}}(m, k) = 0$. $\qquad \square$

As an illustration, with $K = 3$ alternatives we get an upper bound of $6 / \sqrt{(2\pi m/6)^2} \in O(1/m)$; with $K = 4$ alternatives the upper bound is $24 / \sqrt{(2\pi m/24)^{17}} \in O(1/m^{8.5})$. The rate of convergence to zero is faster when K is larger.

5 Flexible consensus

Motivated by the low probability of a level-1 consensus, we suggest below a weakened variant of this property termed *Flexible Consensus*. It is equivalent to a weakened version of Condition 1, without Condition 2.

Definition 3 Let $\succ_0 \in \mathcal{P}$. A profile $\pi \in \mathcal{P}^n$ exhibits *Flexible Consensus around* $\succ_0$ if the following condition holds:

(Flexible Condition 1)
For all $\succ', \succ$, if $\mu_\pi(\succ') > \mu_\pi(\succ)$, then $d(\succ', \succ_0) \leq d(\succ, \succ_0)$.

The only difference between Condition 1 and Flexible Condition 1 is that $d(\succ', \succ_0) < d(\succ, \succ_0)$ is replaced by $d(\succ', \succ_0) \leq d(\succ, \succ_0)$. It will be shown below that this apparently minor change significantly increases the probability that the condition is satisfied, while keeping the desirable stability properties

of the original condition. Moreover, these stability properties hold even without Condition 2.[4]

5.1 First stability property: existence of a Condorcet winner

Definition 4 Given a profile π, the *majority relation M_π* is defined as follows: $a M_\pi b$ iff, in a vote between alternatives a and b, a beats b by a weak majority. I.e, the number of preferences in π by which $a \succ b$ is at least as large as the number of preferences in π by which $b \succ a$.

Lemma 3 *Let $\pi \in \mathcal{P}^n$ be a profile of n voters and $\succ_0 \in \mathcal{P}$ a preference-relation such that Flexible Condition 1 is satisfied. Then, for any two alternatives a, b:*

(1) *If $a \succ_0 b$ (a is preferred to b by $\succ_0$) then $a M_\pi b$.*
(2) *If n is odd then the opposite is also true: if $a M_\pi b$ then $a \succ_0 b$.*
(3) *If n is even then the opposite is "almost" true: if $a M_\pi b$ then either $a \succ_0 b$ or $a \succ_1 b$, where $\succ_1$ is another preference around which there is Flexible Consensus.*[5]

Proof (1) Suppose that $a \succ_0 b$. Partition $\mathcal{P}$, the set of $K!$ possible preferences, to two subsets:

- The subset $C(a > b)$ containing the $K!/2$ preferences for which $a \succ b$;
- The subset $C(b > a)$ containing the $K!/2$ preferences for which $b \succ a$.

Let $w^{ab} : C(b > a) \rightarrow C(a > b)$ be the bijection that takes a preference in $C(b > a)$ and switches a with b in the ranking. Since $a \succ_0 b$, this switch brings every preference in $C(b > a)$ at least one step closer to $\succ_0$. I.e, for every preference $\succ \in C(b > a)$ it holds that $d(w^{ab}(\succ), \succ_0) < d(\succ, \succ_0)$ (see proof in Appendix B).

By Flexible Condition 1, this implies that $\mu_\pi(w^{ab}(\succ)) \geq \mu_\pi(\succ)$. So for every preference $\succ$ by which b is preferred to a corresponds a unique preference $w^{ab}(\succ)$ by which a is preferred to b, which is at least as frequent. Therefore, a beats b by a weak majority: $a M_\pi b$.

(2) When n is odd and $a M_\pi b$, the majority must be strict, so it is not true that $b M_\pi a$. Hence, by (1), it is not true that $b \succ_0 a$. By our assumption, $\succ_0$ is a strict total order. Therefore, $a \succ_0 b$.

(3) When n is even and $a M_\pi b$, there are two cases: if $a \succ_0 b$ then we are done. If $b \succ_0 a$, then $a \succ_1 b$, where $\succ_1 = w^{ab}(\succ_0)$. It remains to prove that there is Flexible Consensus around $w^{ab}(\succ_0)$.

Step I Since $b \succ_0 a$, the argument in (1) shows that, for every preference $\succ \in$ $C(a > b)$, the frequency $\mu_\pi(w^{ab}(\succ)) \geq \mu_\pi(\succ)$. But since $a M_\pi b$, all these inequalities must be equalities, i.e, for every preference $\succ \in C(a > b)$, we must have $\mu_\pi(w^{ab}(\succ)) = \mu_\pi(\succ)$.

Step II For every two preferences $\succ, \succ'$: $d(\succ, \succ') = d(w^{ab}(\succ), w^{ab}(\succ'))$. To see this, consider the pairs of alternatives that are inversed between $\succ, \succ'$. If, in each such pair, we replace a by b and b by a, then we get exactly the pairs of alternatives that are inversed between $w^{ab}(\succ)$ and $w^{ab}(\succ')$.

Step III Let $\succ', \succ$ be two preferences for which $d(\succ', w^{ab}(\succ_0)) < d(\succ, w^{ab}(\succ_0))$. Then, by Step II, $d(w^{ab}(\succ'), \succ_0) < d(w^{ab}(\succ), \succ_0)$. Since there is Flexible Consensus around $\succ_0$, this implies: $\mu_\pi(w^{ab}(\succ')) \geq \mu_\pi(w^{ab}(\succ))$. By Step I, this implies: $\mu_\pi(\succ') \geq \mu_\pi(\succ)$. Hence, there is Flexible Consensus around $w^{ab}(\succ_0)$. $\qquad\square$

Definition 5 For every preference $\succ_0 \in \mathcal{P}$, the alternative ranked first according to $\succ_0$ is denoted by $Best(\succ_0)$

Definition 6 Given a profile $\pi \in \mathcal{P}^n$, a *weak Condorcet winner* of π is an alternative a that beats all other alternatives by a weak majority, i.e, for any other alternative b, $a M_\pi b$.

Theorem 2 *Let $\pi \in \mathcal{P}^n$ be a profile and $\succ_0 \in \mathcal{P}$ a preference-relation around which there is Flexible Consensus. Then $Best(\succ_0)$ is a weak Condorcet winner of π.*

Moreover, if n is odd then $\succ_0$ coincides with the majority relation M_π, and $\succ_0$ is the unique preference in $\mathcal{P}$ for which Flexible Condition 1 is satisfied.

Proof Let $a := Best(\succ_0)$. So for every $b \neq a$, $a \succ_0 b$. By Lemma 3, this implies that $a M_\pi b$. Hence, a is a weak Condorcet winner of π.

When n is odd, Lemma 3 implies that $a \succ_0 b$ iff $a M_\pi b$, so $\succ_0$ coincides with the ordering induced by M_π. This is true for any preference in $\mathcal{P}$ for which Flexible Condition 1 holds, so any such preference coincides with $\succ_0$. $\qquad\square$

5.2 Second stability property: agreement of scoring rules

A *scoring rule* is a rule for selecting an alternative based on a profile.

Definition 7 A *scoring rule* is a rule characterized by a vector S of length K, $S_1 \geq \cdots \geq S_K$. Given a profile π, for each preference $\succ \in \pi$, the rule assigns score S_1 to the alternative ranked first by $\succ$, S_2 to the alternative ranked second by $\succ$, and so on. The rule then sums the scores assigned to each alternative by all preferences in π, and selects the alternative/s that received the highest total score.

In general, every scoring-rule S might select a different alternative. But below we show that, if a profile exhibits Flexible Consensus, then there is an alternative which is selected all scoring rules.

Lemma 4 *Let $\pi \in \mathcal{P}^n$ be a profile and $\succ_0 \in \mathcal{P}$ a preference-relation such that Flexible Condition 1 is satisfied. Then, for any two alternatives a, b and any scoring-rule S, if $a \succ_0 b$ then the score of a is at least as large as the score of b.*

Proof Similarly to the proof of Lemma 3, we partition $\mathcal{P}$ into two subsets, $C(a > b)$ and $C(b > a)$, and define the bijection w^{ab} between them.

For every scoring-rule S and preference $\succ \in \mathcal{P}$, define $\Delta_{S,a,b}(\succ)$ as the difference between the score of a in $\succ$ and the score of b in $\succ$. By definition of a scoring rule:

- For every preference $\succ \in C(a > b)$, $\Delta_{S,a,b}(\succ)$ is weakly-positive.
- For every preference $\succ \in C(b > a)$, $\Delta_{S,a,b}(\succ)$ is weakly-negative.
- For every preference $\succ \in \mathcal{P}$, $\Delta_{S,a,b}(\succ) = -\Delta_{S,a,b}(w^{ab}(\succ))$.

Given the scoring rule S and the profile π, define $\Delta_{S,a,b}(\pi)$ as the difference between the total score of a in π and the total score of b in π. Then, by definition:

$$
\begin{aligned}
\Delta_{S,a,b}(\pi) &= \sum_{\succ \in \mathcal{P}} \mu_\pi(\succ) \cdot \Delta_{S,a,b}(\succ) \\
&= \sum_{\succ \in C(a>b)} \mu_\pi(\succ) \cdot \Delta_{S,a,b}(\succ) + \sum_{\succ \in C(b>a)} \mu_\pi(\succ) \cdot \Delta_{S,a,b}(\succ) \\
&= \sum_{\succ \in C(a>b)} \left[\mu_\pi(\succ) \cdot \Delta_{S,a,b}(\succ) + \mu_\pi(w^{ab}(\succ)) \cdot \Delta_{S,a,b}(w^{ab}(\succ)) \right] \\
&= \sum_{\succ \in C(a>b)} \left[\mu_\pi(\succ) \cdot \Delta_{S,a,b}(\succ) - \mu_\pi(w^{ab}(\succ)) \cdot \Delta_{S,a,b}(\succ) \right] \\
&= \sum_{\succ \in C(a>b)} \Delta_{S,a,b}(\succ) \cdot \left[\mu_\pi(\succ) - \mu_\pi(w^{ab}(\succ)) \right]
\end{aligned}
$$

Since $a \succ_0 b$, for every preference $\succ \in C(a > b)$, the lemma in Appendix B implies that $d(w^{ab}(\succ), \succ_0) > d(\succ, \succ_0)$. Hence, by Flexible Condition 1, $\mu_\pi(\succ) \geq \mu_\pi(w^{ab}(\succ))$. Hence, all terms in the last sum are weakly-positive. Hence, $\Delta_S(\pi) \geq 0$ and the lemma is proved. $\qquad \square$

Theorem 3 *Let $\pi \in \mathcal{P}^n$ be a profile and $\succ_0 \in \mathcal{P}$ a preference-relation such that Flexible Condition 1 is satisfied. Then the score assigned to $Best(\succ_0)$ by every scoring rule is at least as high as the score assigned to any other alternative by the same rule.*

Proof Follows immediately from Lemma 4. $\qquad \square$

6 Algorithm for detecting flexible consensus

Algorithm 3 Detection of Flexible Consensus

INPUT:

- A set of K alternatives, $A = \{a_1, \ldots, a_K\}$.
- A profile π containing n preference-relations on A (possibly with duplicates).

OUTPUT:

- If π exhibits Flexible Consensus around some preference $\succ_0$, return $\succ_0$.
- Otherwise, return "no consensus".

ALGORITHM:

1. Calculate the frequency $\mu_\pi(\succ)$ of each preference $\succ \in \pi$.
2. Define $n' =$ number of distinct preferences in π.
3. Order the preferences by descending frequency: $\mu_\pi(\succ_1) \geq \mu_\pi(\succ_2) \geq \cdots \geq \mu_\pi(\succ_{n'})$.
4. Set $M := \mu_\pi(\succ_1)$ to be the maximum frequency of a preference relation.
5. For $j = 1, 2, \ldots n'$ while $\mu_\pi(\succ_j) = M$: Check if Flexible Condition 1 is satisfied for π and $\succ_j$ (see Algorithm 4). If it is, then return $\succ_j$.
6. Return "no consensus; Flexible Condition 1 is violated for all candidates".

Checking the existence of Flexible Consensus is very similar to checking Level-1 Consensus. The check is presented in Algorithm 3. It is very similar to Algorithm 1; the only differences are that we do not need to check Condition 2 (since it does not exist in Flexible-Level-1-Consensus), and instead of checking Condition 1 for each candidate, we check Flexible-Condition-1.

The procedure for checking Flexible Condition 1 is presented as Algorithm 4. It is very similar to the one for checking Condition 1 in Algorithm 2. There are two differences: the inequality that leads to the failure of the procedure is $d(\succ_i, \succ_0) > d(\succ_{i+1}, \succ_0)$ (instead of $d(\succ_i, \succ_0) \geq d(\succ_{i+1}, \succ_0)$), and in the last step we have to check that there is no preference outside of π whose distance to $\succ_0$ is less than $\widehat{d}$ (instead of less-than-or-equal-to $\widehat{d}$). Hence, by following the same proof as in Sect. 3, it is easy to prove:

Theorem 4 *Algorithm* 3 *checks whether a profile π exhibits Flexible Consensus in time:*

$$O(n'^2 K \sqrt{\log K} + n'^2 \log n')$$

where $n' \leq \min(K!, n)$ is the number of distinct preferences in π.

Algorithm 4 Check if Flexible Condition 1 is satisfied

INPUT:

- A profile π containing n preference-relations on A (possibly with duplicates) ordered by descending frequencies.
- A preference $\succ_0$.

OUTPUT: "True" if Flexible Condition 1 is satisfied for π and $\succ_0$. "False" otherwise.

ALGORITHM:

1. Order preferences with the same frequency by ascending inversion-distance from $\succ_0$.
2. For $i = 1, 2, \ldots, n' - 1$:
   If $\mu_{\underline{\pi}}(\succ_i) > \mu_\pi(\succ_{i+1})$ and $d(\succ_i, \succ_0) > d(\succ_{i+1}, \succ_0)$, return False.
3. Set $\widehat{d} := d(\succ_{n'}, \succ_0)$. Define n^* as the number of profiles in π whose distance to $\succ_0$ is at most $\widehat{d} - 1$.
   If $\sum_{j=0}^{\widehat{d}-1} T(K, j) = n^*$, return True.   Else, return False.

---

7 Probability of flexible consensus

We applied Algorithm 3 to the same experimental settings described in Sect. 4 and estimated the probability of having Flexible Consensus. Out of the 315 PrefLib profiles, **39** exhibit Flexible Consensus. All 39 profiles are from the dataset labeled "ED-00004 1-100", where all profiles have 3 alternatives. This means that 39% of all these profiles with 3 alternatives exhibited the Flexible Consensus (in contrast to **0** which exhibited level-1 consensus).

The results of the experiments on random profiles are shown in Fig. 2; it is evident that in all settings, including the most challenging setting of impartial culture ($\phi = 1$), Flexible Consensus is substantially more probable than level-1 consensus.

We complement the empirical findings with a theoretical analysis of the asymptotic probability of Flexible Consensus under the impartial culture assumption. We consider a profile generated by the random process described before Proposition 1. We denote by $P^{\text{flexible-consensus}}(m, K)$ the probability that such a random profile exhibits Flexible Consensus. In stark contrast to Proposition 1, we prove that this probability is always larger than a positive constant, even when the number of voters goes to infinity:

Proposition 2 *For every K there exists a constant $C_K > 0$ such that:*

$$\forall m : P^{\text{flexible-consensus}}(m, K) \geq C_K$$

Proof As explained in the proof of Proposition 1, the probability that two or more preferences have exactly the same frequency goes to 0 when $m \to \infty$, so for simplicity we neglect this possibility and assume that each preference relation has a different frequency. Note that this assumption can only decrease the probability of Flexible Consensus, since it implies that there is a unique preference with maximum frequency, so there is only one candidate that can possibly satisfy Flexible Condition 1. We denote this candidate by $\succ_0$. Below we calculate the probability that Flexible Condition 1 holds for this preference.

For every $i \geq 1$, define:

$$F_i := \{ \mu_\pi(\succ) \mid d(\succ, \succ_0) = i \}$$

so F_i contains the frequencies of all preferences whose distance from $\succ_0$ is exactly i. Note that F_i is non-empty only when $i \leq \binom{K}{2}$, since $\binom{K}{2}$ is the maximum possible inversion-distance between two preferences on K alternatives.

Flexible Condition 1 is equivalent to the requirement that each member of the set F_i is larger than each member of the set F_j, for every $i < j$.[6] However, it does not impose any restrictions on the frequencies *within* F_i.

Let $F := \cup_i F_i$ = the set of frequencies of the $K! - 1$ preferences different than $\succ_0$. The total number of different orders on F is $|F|!$. The total number of orders that satisfy Flexible Condition 1 is $|F_1|! \cdot |F_2|! \cdots |F_{\binom{K}{2}}|!$. Since all preferences are equally likely, all $|F|!$ orders are equally likely. Therefore, the probability that the order of frequencies satisfies Flexible Condition 1 is at least:

$$\frac{|F_1|!|F_2|! \cdots |F_{\binom{K}{2}}|!}{|F|!} \tag{5}$$

which is a positive constant that does not approach 0 even when $m \to \infty$. □

As an illustration, we calculate the lower bound for $K = 3$ alternatives. In this case we have $|F_1| = 2$ and $|F_2| = 2$ and $|F_3| = 1$ and $|F| = 2 + 2 + 1 = 5$. Therefore, the probability that Flexible Condition 1 is satisfied in a random impartial-culture profile is at least $2! \cdot 2! \cdot 1!/5! = 1/30 \approx 0.033$. In our experiments with $\phi = 1$, the fraction of profiles with Flexible Consensus (in 1000 experiments) was 0.045 for 100 voters and 0.043 for 1000 voters. This is slightly higher than the lower bound of 0.033, which can be explained by the fact that, when n is finite, there is a positive probability that two preferences have the same frequency.

When $K > 3$, the asymptotic probability of Flexible Condition 1 in an impartial culture remains positive, though much lower. For example, for $K = 4$ the lower bound is approximately 10^{-12}. This is consistent with the fact that we found no profiles that exhibit Flexible Consensus in our experiments with $\phi = 1$ and $K \geq 4$.

8 Conclusion

We presented a practical procedure for checking whether a preference-profile exhibits a level-1 consensus. We realized that this property is highly improbable, and found a weaker property, Flexible Consensus, that preserves the desirable stability properties of the social choice. Furthermore, Flexible Consensus is provably more likely than level-1 consensus. This was demonstrated both theoretically (for the impartial culture setting) and empirically and over a database of real-world profiles.

Our experiments can be reproduced by re-running the code, which is freely available through the following GitHub fork: https://github.com/erelsgl/PrefLib-Tools.

---

[6] In fact, it is equivalent to the requirement that each member of F_i is larger than *or equal to* each member of F_j, but as explained above, we neglect the possibility that two preferences have equal frequency.

Appendices

Appendix A: To verify Condition 1, it is sufficient to check adjacent preferences

This section provides a formal proof to the following lemma used in Sect. 3.

Lemma 1 *Suppose that the preference relations in a profile π are ordered by two criteria: first by frequency $\mu_\pi(\succ_i)$, then by distance $d(\succ_i, \succ_0)$, where $\succ_0$ is a fixed preference. In this ordering, if Condition 1 is violated for any pair of preference relations in π, then it is violated for an adjacent pair $\succ_i, \succ_{i+1}$ for some i.*

Proof Suppose that there exist indices $i < j$ such that Condition 1 is violated for the pair $\succ_i, \succ_j$, i.e, $\mu_\pi(\succ_i) > \mu_\pi(\succ_j)$, yet $d(\succ_i, \succ_0) \geq d(\succ_j, \succ_0)$. We now prove the lemma by induction on the difference of indices, $j - i$.

Base: If $j - i = 1$, then $\succ_i$ and $\succ_j$ are already adjacent, so we are done.

Step: Suppose that $j - i > 1$. We prove that there exists a pair with a smaller difference that violates Condition 1. We consider several cases.

Case 1 i is not the largest index in its equivalence class. i.e, there exists $i' > i$ such that $\mu_\pi(\succ_{i'}) = \mu_\pi(\succ_i)$. Then, by the secondary ordering criterion, $d(\succ_{i'}, \succ_0) \geq d(\succ_i, \succ_0)$, Condition 1 is violated for the pair $\succ_{i'}$ and $\succ_j$.

Case 2 j is not the smallest index in its equivalence class. i.e, there exists $j' < j$ such that $\mu_\pi(\succ_{j'}) = \mu_\pi(\succ_j)$. Then, by the secondary ordering criterion, $d(\succ_j, \succ_0) \geq d(\succ_{j'}, \succ_0)$, Condition 1 is violated for the pair $\succ_i$ and $\succ_{j'}$. Otherwise, i is the largest index in its equivalence class, j is the smallest index in its equivalence class, but still $i + 1 < j$. This means that the equivalence classes of i and j are not adjacent, i.e, $\mu_\pi(\succ_i) > \mu_\pi(\succ_{i+1}) > \mu_\pi(\succ_j)$. Now there are two remaining cases:

Case 3 $d(\succ_i, \succ_0) \geq d(\succ_{i+1}, \succ_0)$, in which case Condition 1 is violated for the adjacent pair $\succ_i$ and $\succ_{i+1}$ and we are done.

Case 4 $d(\succ_{i+1}, \succ_0) > d(\succ_i, \succ_0)$. This implies $d(\succ_{i+1}, \succ_0) > d(\succ_j, \succ_0)$, so Condition 1 is violated for the pair $\succ_{i+1}$ and $\succ_j$ and we are done. □

Appendix B: The effect of a switch on the inversion-distance

This section provides a formal proof to an intuitive claim made within the proof of Lemma 3. Let a, b be two fixed alternatives. Let $C(a > b)$ be the set of preferences by which $a \succ b$ and $C(b > a)$ the set of preferences by which $b \succ a$. Let $w^{ab} : C(b > a) \to C(a > b)$ be a function that takes a preference-relation and creates a new preference-relation by switching the position of a and b in the ranking.

Lemma 5 *If $a \succ_0 b$, then for every preference $\succ_1 \in C(b > a)$:*

$$d(\succ_1, \succ_0) \; > \; d(w^{ab}(\succ_1), \succ_0)$$

Proof For every preference $\succ$, define $D(\succ, \succ_0)$ as the set of pairs-of-alternatives $\{i, j\}$ that are ranked differently in $\succ$ and in $\succ_0$. By definition, the inversion distance is the cardinality of this set, $d(\succ, \succ_0) = |D(\succ, \succ_0)|$. Therefore, it is sufficient to show that there are more pairs in $D(\succ_1, \succ_0)$ than in $D(w^{ab}(\succ_1), \succ_0)$. To show this, we consider all possible pairs-of-alternatives; for each pair, we calculate its contribution to the difference in cardinalities $|D(\succ_1, \succ_0)| - |D(w^{ab}(\succ_1), \succ_0)|$, and show that the net difference is positive.

- The pair $\{a, b\}$ is in $D(\succ_1, \succ_0)$ but not in $D(w^{ab}(\succ_1), \succ_0)$, so this pair contributes $+1$ to the difference.
- Any pair that contains neither a nor b is not affected by the switch. I.e, each pair $\{c, e\}$ where $c, e \neq a, b$ is in $D(\succ_1, \succ_0)$ if-and-only-if it is in $D(w^{ab}(\succ_1), \succ_0)$, so it contributes 0 to the difference.
- Let c be an alternative that is ranked by $\succ_1$ above b or below a, i.e, either $c \succ_1 b \succ_1 a$ or $b \succ_1 a \succ_1 c$. Then, the order between c to a or b in $\succ_1$ is not affected by the switch, so $\{c, a\}$ is in $D(\succ_1, \succ_0)$ if-and-only-if it is in $D(w^{ab}(\succ_1), \succ_0)$, so it contributes 0 to the difference. The same is true for $\{c, b\}$.
- Finally, let c be an alternative that is ranked by $\succ_1$ between a and b, i.e, $b \succ_1 c \succ_1 a$. Then, the switch w^{ab} changes the direction of both the pair $\{c, a\}$ and the pair $\{c, b\}$:
  - If $c \succ_0 a \succ_0 b$, then the pair $\{c, a\}$ is in $D(\succ_0, w^{ab}(\succ_1))$ but not in $D(\succ_0, \succ_1)$, and the pair $\{c, b\}$ is in $D(\succ_0, \succ_1)$ but not in $D(\succ_0, w^{ab}(\succ_1))$, so these pairs contribute $+1 - 1 = 0$ to the difference.
  - If $a \succ_0 c \succ_0 b$, then both the pair $\{c, a\}$ and the pair $\{c, b\}$ are in $D(\succ_0, \succ_1)$ but not in $D(\succ_0, w^{ab}(\succ_1))$, so these pairs contribute $+1 + 1 = 2$ to the difference.
  - If $a \succ_0 b \succ_0 c$, then the pair $\{c, a\}$ is in $D(\succ_0, \succ_1)$ but not in $D(\succ_0, w^{ab}(\succ_1))$, and the pair $\{c, b\}$ is in $D(\succ_0, w^{ab}(\succ_1))$ but not in $D(\succ_0, \succ_1)$, so these pairs contribute $+1 - 1 = 0$ to the difference.

We proved that the contribution of each pair of alternatives is at least 0, and the contribution of the pair $\{a, b\}$ is $+1$. Therefore, the difference $|D(\succ_1, \succ_0)| - |D(w^{ab}(\succ_1), \succ_0)|$ is positive and the lemma is proved. $\qquad\square$

Appendix C: Probability that binomial random variables are equal

This section justifies the following approximation, used in Sect. 4.

Let $B_1, \ldots, B_t$ be i.i.d. random variables distributed binomially with m trials and success-probability $p < 1/2$. Let $q := 1 - p$. Define:

$$P^{\text{equal}}(m, p, t) := \Pr[B_1 = \cdots = B_t]$$

Then, for every $t \geq 1$, when m is sufficiently large,

$$P^{\text{equal}}(m, p, t) \approx \frac{1}{(\sqrt{2\pi pqm})^{t-1}}$$

Proof The value of each of the variables B_i can be any integer between 0 and m. Therefore we can present P^{equal} as a sum of probabilities of disjoint events:

$$P^{\text{equal}}(m, p, t) = \sum_{i=0}^{m} \Pr[B_1 = \cdots = B_t = i]$$

Since the B_i are i.i.d:

$$P^{\text{equal}}(m, p, t) = \sum_{i=0}^{m} (\Pr[B_1 = i])^t$$

$$= \sum_{i=0}^{m} \left(\binom{m}{i} p^i q^{m-i} \right)^t$$

where $q := 1 - p$. Using Stirling's approximation, when $m, i, m - i$ are sufficiently large:

$$\binom{m}{i} \approx \sqrt{\frac{m}{2\pi i(m-i)}} \cdot \frac{m^m}{i^i (m-i)^{m-i}}$$

Substitute this in P^{equal} and approximate the sum by an integral:

$$P^{\text{equal}}(m, p, t) \approx \int_{y=0}^{m} \left(\sqrt{\frac{m}{2\pi y(m-y)}} \cdot \frac{m^m}{y^y (m-y)^{m-y}} \cdot p^y q^{m-y} \right)^t dy$$

Substitute $y = mx$ and $dy = mdx$:

$$P^{\text{equal}}(m, p, t) \approx m \int_{x=0}^{1} \left(\sqrt{\frac{1}{2\pi mx(1-x)}} \cdot \frac{p^{mx} q^{m-mx}}{x^{mx} (1-x)^{m-mx}} \right)^t dx$$

$$= m^{1-t/2} \int_{x=0}^{1} \left(\sqrt{\frac{1}{2\pi x(1-x)}} \cdot \frac{p^{mx} q^{m-mx}}{x^{mx} (1-x)^{m-mx}} \right)^t dx$$

The integral can be approximated by Laplace's method. Define:

$$h(x) \quad := \left(\sqrt{\tfrac{1}{2\pi x(1-x)}} \right)^t$$

$$g(x) := t \cdot \left[x \ln(\tfrac{p}{x}) + (1-x) \ln(\tfrac{q}{1-x}) \right]$$

Then:

$$P^{\text{equal}}(m, p, t) \approx m^{1-t/2} \int_{x=0}^{1} h(x) e^{mg(x)} dx$$

The function $g(x)$ is twice continuously differentiable on $(0, 1)$ and has a unique maximum at $x_0 = \frac{p}{p+q} = p$; the maximum value is $g(x_0) = 0$. Moreover, $g''(x_0) = -\frac{1}{pq} < 0$. Therefore, by Laplace's method:

$$P^{\text{equal}}(m, p, t) \approx m^{1-t/2} \sqrt{\frac{2\pi}{-mg''(x_0)}} \cdot h(x_0) \cdot e^{mg(x_0)}$$

where the symbol $\approx$ means that the ratio between the expressions in its two sides goes to 1 as $m \to \infty$. Substituting the functions g and h gives:

$$P^{\text{equal}}(m, p, t) \approx m^{1-t/2} \sqrt{\frac{2\pi}{m/(pq)}} \cdot \left(\frac{1}{2\pi pq}\right)^{t/2} \cdot e^0$$
$$= m^{(1-t)/2} \cdot (2\pi pq)^{(1-t)/2}$$
$$= (2\pi pqm)^{(1-t)/2}$$
$$P^{\text{equal}}(m, p, t) \approx \frac{1}{(\sqrt{2\pi pqm})^{t-1}}$$

$\square$

Note that $P^{\text{equal}}(m, p, 1) = 1$, which is trivially true, since a single random variable always equals itself. When $t \geq 2$, $P^{\text{equal}}(m, p, 1) \to 0$ as $m \to \infty$.

References

Barberà S, Moreno B (2011) Top monotonicity: a common root for single peakedness, single crossing and the median voter result. Games Econ Behav 73(2):345–359

Bennett J , Lanning S (2007) The Netflix prize. In: Proceedings of the KDD Cup Workshop 2007, 3–6. ACM, New York

Bredereck R, Chen J, Woeginger GJ (2013) A characterization of the single-crossing domain. Soc Choice Welf 41(4):989–998

Chan TM , Pătrașcu M (2010) Counting inversions, offline orthogonal range counting, and related problems. In: Proceedings of the twenty-first annual ACM-SIAM symposium on discrete algorithms, SODA '10, 161–173, Philadelphia, PA, USA. Society for Industrial and Applied Mathematics

Elkind E, Lackner M, Peters D (2016) Preference restrictions in computational social choice: recent progress. IJCAI 16:4062–4065

Escoffier B, Lang J, Öztürk M (2008) Single-peaked consistency and its complexity. ECAI 8:366–370

Gehrlein WV (1981) The expected probability of Condorcet's paradox. Econ Lett 7(1):33–37

Kemeny JG, Snell JL (1962) Mathematical models in the social sciences. Blaisdell Publishing Company, Boston

Lackner M-L, Lackner M (2017) On the likelihood of single-peaked preferences. Soc Choice Welf 48(4):717–745

Lu T, Boutilier C (2014) Effective sampling and learning for mallows models with pairwise-preference data. J Mach Learn Res 15(1):3783–3829

Mahajne M, Nitzan S, Volij O (2015) Level r consensus and stable social choice. Soc Choice Welf 45(4):805–817

Mallows CL (1957) Non-null ranking models. i. Biometrika 44(1/2):114–130

Mao A, Procaccia AD, Chen Y (2013) Better human computation through principled voting. In: Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, AAAI'13, pp 1142–1148. AAAI Press

Mattei N, Walsh T (2013) Preflib: a library of preference data http://preflib.org. In: Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013), Lecture Notes in Artificial Intelligence. Springer

Poliakov NL (2016) Note on level r consensus. arXiv preprint arXiv:1606.04816

Tsetlin I, Regenwetter M, Grofman B (2003) The impartial culture maximizes the probability of majority cycles. Soc Choice Welf 21(3):387–398