

Measuring Diversity of Preferences in a Group

Vahid Hashemi and Ulle Endriss¹

Abstract. We introduce a general framework for measuring the degree of diversity in the preferences held by the members of a group. We formalise and investigate three specific approaches within that framework: diversity as the range of distinct views held, diversity as aggregate distance between individual views, and diversity as distance of the group’s views to a single compromise view. While similarly attractive from an intuitive point of view, the three approaches display significant differences when analysed using both the axiomatic method and empirical studies.

1 INTRODUCTION

Preferences are ubiquitous in AI [5, 12]. Examples for application domains include recommender systems, planning, and configuration. Of particular interest is the case of preference handling in multiagent systems, where several agents each have their own individual preferences and we need to take decisions that are appropriate in view of such a profile of preferences. The normative, mathematical, and algorithmic aspects of this problem are studied in the field of (computational) social choice [3]. In social choice, preferences are taken to be linear orders over a finite set of alternatives. As is well known, many of the most interesting phenomena in social choice are in fact rare events. For example, while the notorious *Condorcet Paradox* manifests itself in around 25% of all theoretically possible preference profiles for 5 alternatives and a large number of voters, empirical studies suggest that it plays hardly any role in real-world elections of the same size [11, 17]. Another example is the fact that many of the computational hardness results for the strategic manipulation problem in voting rely on a very narrow basis of worst-case scenarios, while the vast majority of problem instances are in fact easy [18].

This divergence can be explained by the fact that the preference profiles we encounter in practice exhibit a certain amount of structure. The classical approach to modelling such structure are domain restrictions, the best known example of which is *single-peakedness* [9]. Yet, while an unconstrained model of social choice arguably is too broad, domain restrictions are often too narrow to accurately describe preference profiles that occur in practice. In this paper, we propose the exploration of a middle way. Our starting point is the basic idea that the less *diverse* the preferences in a profile are, the easier it should be to come to a mutually acceptable decision. For example, in the most extreme case where all agents share the exact same preference order, it will be trivial to make collective decisions. *Vice versa*, the more diversity we find in a profile, the more we should expect to encounter paradoxes, i.e., situations in which different social choice-theoretic principles would lead to opposing conclusions.

Our first contribution is to propose a formal model of preference diversity. At the centre of this model is the notion of *preference*

diversity index (PDI): a function mapping profiles to nonnegative numbers, with 0 denoting perfect agreement amongst all agents. The model does not commit to one specific interpretation of the term diversity. Rather, we use it to formalise three concrete interpretations: diversity as the range of distinct views held (*support-based* PDI), diversity as aggregate distance between individual views (*distance-based* PDI), and diversity as distance of the group’s views to a single compromise view (*compromise-based* PDI).

We formulate several intuitively appealing properties of PDI’s as *axioms* and classify our concrete indices in terms of which of these axioms they satisfy. We also provide an example for an *impossibility result*, showing that certain axioms are mutually incompatible, and a *characterisation result*, showing how one concrete PDI is fully determined by a certain combination of axioms. On the practical side, we have conducted a range of *experiments* that shed additional light on our PDI’s. In particular, we explore how the differences between synthetically generated preference data and data sampled from a real election profile manifest themselves in terms of the distribution of diversity over profiles. We also confirm that the likelihood for undesirable social choice-theoretic effects increases with diversity.

In Section 2 we introduce our model of preference diversity and define several specific PDI’s. Section 3 is devoted to the axiomatic analysis of diversity and Section 4 presents our experimental results. We conclude with a discussion of related work in Section 5 and a brief outlook on other approaches to defining PDI’s in Section 6.

2 MEASURING PREFERENCE DIVERSITY

In this section we introduce the concept of *preference diversity index* (PDI) and then define several concrete such indices.

2.1 Basic terminology and notation

Let \mathcal{X} be a finite set of m alternatives. We model preferences as (strict) *linear orders* over the set of alternatives (recall that a linear order R is a binary relation that is irreflexive, transitive, and complete). We write $\mathcal{L}(\mathcal{X})$ for the set of all preference/linear orders over \mathcal{X} . The *position* of $x \in \mathcal{X}$ in $R \in \mathcal{L}(\mathcal{X})$ is $\text{pos}_R(x) = |\{y \in \mathcal{X} \mid yRx\}|$.

Let $\mathcal{N} = \{1, \dots, n\}$ be a finite set of *voters* (or *agents*). A *profile* $\mathbf{R} = (R_1, \dots, R_n) \in \mathcal{L}(\mathcal{X})^n$ is a vector of preference orders, one for each voter. We write $N_{x \succ y}^{\mathbf{R}} = \{i \in \mathcal{N} \mid xR_i y\}$ for the set of voters who in profile \mathbf{R} say that they prefer x over y . The *support* of a profile $\mathbf{R} = (R_1, \dots, R_n)$ is the set of preference orders occurring in it: $\text{SUPP}(\mathbf{R}) = \{R_1\} \cup \dots \cup \{R_n\}$. We call a profile \mathbf{R} *unanimous* if $|\text{SUPP}(\mathbf{R})| = 1$, i.e., if it is of the form (R, \dots, R) .

2.2 Preference diversity orderings and indices

Given two profiles \mathbf{R} and \mathbf{R}' (both with n voters expressing preferences over m alternatives), we want to be able to make judgments

¹ Institute for Logic, Language and Computation, University of Amsterdam, email: vhashemi@gmail.com, ulle.endriss@uva.nl

about which of them we consider more diverse. Recall that a *weak order* is a binary relation that is reflexive, transitive, and complete.

Definition 1. A *preference diversity order (PDO)* is a weak order \succsim declared on the space of preference profiles $\mathcal{L}(\mathcal{X})^n$ that respects $\mathbf{R} \succsim (R, \dots, R)$ for all $\mathbf{R} \in \mathcal{L}(\mathcal{X})^n$ and all $R \in \mathcal{L}(\mathcal{X})$.

That is, any PDO is required to classify unanimous profiles as being minimally diverse (and any two such profiles are equally diverse). We write \succ for the strict part of \succsim , and \sim for its indifference part.

Definition 2. A *preference diversity index (PDI)* is a function $\Delta : \mathcal{L}(\mathcal{X})^n \rightarrow \mathbb{R}^+ \cup \{0\}$, mapping profiles to the nonnegative reals, that respects $\Delta(R, \dots, R) = 0$ for any $R \in \mathcal{L}(\mathcal{X})$.

Let $\max(\Delta) = \max\{\Delta(\mathbf{R}) \mid \mathbf{R} \in \mathcal{L}(\mathcal{X})^n\}$. We say that a PDI Δ is *normalised* if it maps any given profile to the interval $[0, 1]$, and the maximum of 1 is reached for at least one profile, i.e., $\max(\Delta) = 1$. Every given PDI Δ gives rise to a normalised PDI Δ' by stipulating $\Delta'(\mathbf{R}) = \Delta(\mathbf{R}) / \max(\Delta)$ for every profile $\mathbf{R} \in \mathcal{L}(\mathcal{X})^n$.

A PDI Δ induces a PDO \succsim^Δ , by stipulating $\mathbf{R} \succsim^\Delta \mathbf{R}'$ if and only if $\Delta(\mathbf{R}) \geq \Delta(\mathbf{R}')$ for any two profiles $\mathbf{R}, \mathbf{R}' \in \mathcal{L}(\mathcal{X})^n$. Observe that, as Δ is required to map any unanimous profile to 0, any such profile will be correctly placed at the bottom of the corresponding PDO \succsim^Δ (i.e., the two definitions match). Every PDO \succsim (and thus every PDI) defines a partitioning of $\mathcal{L}(\mathcal{X})^n$ into equivalence classes w.r.t. \sim . We can think of these equivalence classes as the possible *levels* of diversity. Let the *dimension* of a PDI/PDO be the number of equivalence classes it defines (for fixed n and m).

2.3 Specific preference diversity indices

We now introduce three specific approaches to defining PDI's. The first approach is based on the idea that diversity may be measured in terms of the *number of distinct views* represented within a group. In its simplest form, this means that we count the number of distinct preference orders in a profile. This leads to the *simple support-based PDI* Δ_{supp} with $\Delta_{supp}(\mathbf{R}) = |\text{SUPP}(\mathbf{R})| - 1$ (we subtract 1 to ensure $\Delta_{supp}(R, \dots, R) = 0$). We can generalise this idea and count the number of distinct ordered k -tuples of alternatives appearing in a profile. The number of such tuples in *one* preference order is $\binom{m}{k}$. Let $\mathcal{L}_k(\mathcal{X})$ denote the set of ordered k -tuples of alternatives.

Definition 3. For a given $k \leq m$, the *support-based PDI* $\Delta_{supp}^{\ell=k}$ maps any given profile $\mathbf{R} \in \mathcal{L}(\mathcal{X})^n$ to the following value:

$$\Delta_{supp}^{\ell=k}(\mathbf{R}) = |\{T \in \mathcal{L}_k(\mathcal{X}) \mid T \subseteq R_i \text{ for some } i \in \mathcal{N}\}| - \binom{m}{k}$$

For example, $\Delta_{supp}^{\ell=2}$ counts the number of ordered pairs at least one agent accepts (above and beyond $\binom{m}{2}$). Note that $\Delta_{supp}^{\ell=2} \equiv \Delta_{supp}^{\ell=m}$.

Our second approach is based on the idea that diversity is related to the *distances between the individual views* held by the members of a group. We first require a notion of distance between two single preference orders R and R' , i.e., a function $\delta : \mathcal{L}(\mathcal{X}) \times \mathcal{L}(\mathcal{X}) \rightarrow \mathbb{R}$ meeting the familiar axioms for distances (nonnegativity, identity of indiscernibles, symmetry, and the triangle inequality). The following are all standard definitions that are widely used in the literature [6, 8]:

- Kendall's tau: $K(R, R') = \frac{1}{2} \cdot (|R \setminus R'| + |R' \setminus R|)$
- Spearman's footrule: $S(R, R') = \sum_{x \in \mathcal{X}} |\text{pos}_R(x) - \text{pos}_{R'}(x)|$
- Discrete distance: $D(R, R') = 0$ if $R = R'$, and $= 1$ otherwise

When R and R' are linear orders, our definition for K is equivalent to the more common $K(R, R') = \frac{1}{2} \cdot |\{(x, y) \mid xRy \text{ and } yR'x\}|$. Note

that we divide by 2 to ensure we count ordered pairs, not merely pairs of alternatives. To lift distances between pairs of voters to distances between the members of a group, we can use any aggregation operator $\Phi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ (that is nondecreasing, associative, commutative, and has identity element 0), such as \max or Σ (sum).

Definition 4. For a given distance $\delta : \mathcal{L}(\mathcal{X}) \times \mathcal{L}(\mathcal{X}) \rightarrow \mathbb{R}$ and aggregation operator $\Phi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$, the *distance-based PDI* $\Delta_{dist}^{\Phi, \delta}$ maps any given profile $\mathbf{R} \in \mathcal{L}(\mathcal{X})^n$ to the following value:

$$\Delta_{dist}^{\Phi, \delta}(\mathbf{R}) = \Phi(\delta(R_i, R_{i'}) \mid i, i' \in \mathcal{N} \text{ with } i < i')$$

That is, we first compute the $n(n-1)/2$ -vector of pairwise distances $\delta(R_i, R_{i'})$ and then apply Φ to that vector. The PDI $\Delta_{dist}^{\Sigma, K}$, for instance, measures diversity as the sum of the Kendall tau distances between all pairs of preferences in a profile. In this paper, we will largely focus on $\Phi = \Sigma$ and the effect of varying δ .

The idea underlying our third approach is to measure diversity as a group's accumulated *distance to a compromise view*. For instance, for a given profile \mathbf{R} , we may compute its *majority graph* $\text{MG}(\mathbf{R}) = \{(x, y) \mid |N_{x \succ y}^{\mathbf{R}}| > \frac{n}{2}\}$ and then measure the distance of the individual preferences to the compromise view represented by $\text{MG}(\mathbf{R})$. To measure the distance between a preference order and a compromise view, we will use the Kendall tau distance, although in principle also other distances could be used. Observe that K , as defined above, is a meaningful notion of distance between any two binary relations on \mathcal{X} , not just linear orders. We refer to functions mapping profiles to binary orders (representing compromise views) as *social welfare functions* (SWF), which is a slight generalisation of the common use of the term in social choice theory [10].

Definition 5. For a given SWF $F : \mathcal{L}(\mathcal{X})^n \rightarrow 2^{\mathcal{X} \times \mathcal{X}}$ and aggregation operator $\Phi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$, the *compromise-based PDI* $\Delta_{com}^{\Phi, F}$ maps any given profile $\mathbf{R} \in \mathcal{L}(\mathcal{X})^n$ to the following value:

$$\Delta_{com}^{\Phi, F}(\mathbf{R}) = \Phi(K(R_i, F(\mathbf{R})) \mid i \in \mathcal{N})$$

Thus, $\Delta_{com}^{\Sigma, \text{MG}}$, for instance, computes the sum of the distances of the individual preferences to the majority graph. Besides $F = \text{MG}$, we can use voting rules, e.g., the *Borda rule* [10], to define a compromise. Under Borda, each voter i gives as many points to x as there are other alternatives below x in i 's ranking; the Borda score of x is the sum of those points. This induces a SWF that for any given profile \mathbf{R} returns the weak order $\text{Bor}(\mathbf{R}) = \{(x, y) \mid \text{BordaScore}(x) \geq \text{BordaScore}(y)\}$. Thus, the PDI $\Delta_{com}^{\max, \text{Bor}}$ computes the maximal distance of any individual preference order to the ranking we obtain when we order alternatives in terms of their Borda score.²

3 AXIOMATIC ANALYSIS

In this section we motivate and formalise desirable properties that a specific manner of measuring preference diversity may or may not satisfy. That is, in the parlance of social choice theory [10], we introduce a number of *axioms* for preference diversity. We formulate these axioms in terms of PDO's rather than PDI's, i.e., we axiomatise the ordinal notion of "being more diverse than", rather than the cardinal notion of having a particular degree of diversity. The reason for this choice is that, while some details of the numerical representation of degrees of diversity is bound to be arbitrary, relative diversity

² Beware that this approach does not result in a well-defined PDI for every possible voting rule: e.g., if we rank alternatives in terms of their *plurality* score [10], for a unanimous profile (R, \dots, R) , we obtain a weak order of depth 2 rather than R , meaning that (R, \dots, R) will not be mapped to 0.

judgments should not and need not be. Our axioms will nevertheless apply to PDI's indirectly, given that every PDI induces a PDO.

We then use our axioms to organise the space of concrete ways of measuring preference diversity introduced earlier. We will also see that not all combinations of axioms can be satisfied together.

3.1 Axioms

Our first axiom is a basic symmetry requirement w.r.t. voters.

Axiom 1. A PDO \succcurlyeq is *anonymous* if, for every permutation $\sigma : \mathcal{N} \rightarrow \mathcal{N}$, we have $(R_1, \dots, R_n) \sim (R_{\sigma(1)}, \dots, R_{\sigma(n)})$.

The statement above is understood to apply to *all* preference profiles (R_1, \dots, R_n) . For the sake of readability, we shall keep such universal quantification over profiles implicit also in later axioms. Our next axiom, neutrality, postulates symmetry w.r.t. alternatives. For any permutation $\tau : \mathcal{X} \rightarrow \mathcal{X}$ on alternatives and any preference order $R \in \mathcal{L}(\mathcal{X})$, define $\tau(R) = \{(x, y) \mid \tau(x)R\tau(y)\}$.

Axiom 2. A PDO \succcurlyeq is *neutral* if, for every permutation $\tau : \mathcal{X} \rightarrow \mathcal{X}$, we have $(R_1, \dots, R_n) \sim (\tau(R_1), \dots, \tau(R_n))$.

Our next axiom says that no two profiles should be judged as being of equal diversity, unless anonymity and neutrality force us to do so.

Axiom 3. A PDO \succcurlyeq is *strongly discernible* if $\mathbf{R} \sim \mathbf{R}'$ implies $\mathbf{R} = (\tau(R'_{\sigma(1)}), \dots, \tau(R'_{\sigma(n)}))$ for some $\sigma : \mathcal{N} \rightarrow \mathcal{N}$ and $\tau : \mathcal{X} \rightarrow \mathcal{X}$.

Strong discernability is a demanding requirement. Intuitively speaking, it excludes PDO's with a low dimension. The next axiom is much weaker (and implied by strong discernability). It only requires the bottom level to be distinct from the others.

Axiom 4. A PDO \succcurlyeq is *weakly discernible* if \mathbf{R} being unanimous and \mathbf{R}' not being unanimous together imply $\mathbf{R}' \succ \mathbf{R}$.

One possible position to take would be to say that diversity should be a function of the variety of views taken by members of a society, but that it should not depend on the frequency with which any particular such view is taken. That is, one might argue, the level of diversity of a profile should only depend on its support.

Axiom 5. A PDO \succcurlyeq is *support-invariant* if $\text{SUPP}(\mathbf{R}) = \text{SUPP}(\mathbf{R}')$ implies $\mathbf{R} \sim \mathbf{R}'$.

Observe that support-invariance implies anonymity. A different position to take would be to say that every single preference order matters. That is, it should not be possible to determine the level of diversity of a profile by only inspecting a proper subset of its elements.

Axiom 6. A PDO \succcurlyeq is *nonlocal* if for every profile $\mathbf{R} = (R_1, \dots, R_n) \in \mathcal{L}(\mathcal{X})^n$ and every voter $i \in \mathcal{N}$ there exists an order $R' \in \mathcal{L}(\mathcal{X})$ such that $\mathbf{R} \not\sim (R_1, \dots, R_{i-1}, R', R_{i+1}, \dots, R_n)$.

Our next axiom is adopted from the literature on ranking opportunity sets for measuring freedom of choice [16]. For any profile \mathbf{R} for n voters and individual preference R , let $\mathbf{R} \oplus R = (R_1, \dots, R_n, R)$ be the profile for $n+1$ voters we obtain by adding R to the first profile.³

Axiom 7. A PDO \succcurlyeq is *independent* if it is the case that $\mathbf{R} \succcurlyeq \mathbf{R}'$ if and only if $\mathbf{R} \oplus R \succcurlyeq \mathbf{R}' \oplus R$ for every two profiles $\mathbf{R}, \mathbf{R}' \in \mathcal{L}(\mathcal{X})^n$ and every preference $R \notin \text{SUPP}(\mathbf{R}) \cup \text{SUPP}(\mathbf{R}')$.

³ Note that, strictly speaking, Axiom 7 speaks about a *family* of PDO's (one for each n), even if it does not directly compare profiles of different size.

Finally, we consider two possible definitions of *monotonicity*. What they have in common is that they identify situations in which one or more voters change their preferences by moving closer to the views of the rest of the group, which intuitively should reduce diversity. First, suppose a single voter abandons her own preference order and instead adopts the preferences of one of the other voters.

Axiom 8. A PDO \succcurlyeq is *monotonic* if $\mathbf{R} \succcurlyeq \mathbf{R}'$ whenever there exist $j, k \in \mathcal{N}$ such that $R'_j = R_k$ and $R'_i = R_i$ for all $i \neq j$.

Observe that our monotonicity axiom implies support-invariance: if $\text{SUPP}(\mathbf{R}) = \text{SUPP}(\mathbf{R}')$, then we can move from \mathbf{R} to \mathbf{R}' (and *vice versa*) via a sequence of monotonicity-moves.

Now suppose one or several voters each swap two adjacent alternatives x and y in their preference orders. Under what circumstances should we consider such a move as having reduced diversity?

Axiom 9. A PDI \succcurlyeq is *swap-monotonic* if $\mathbf{R} \succcurlyeq \mathbf{R}'$ holds whenever there exist alternatives $x, y \in \mathcal{X}$ such that $|N_{x \succ y}^{\mathbf{R}}| \geq |N_{y \succ x}^{\mathbf{R}'}|$, $N_{x \succ y}^{\mathbf{R}'} = \mathcal{N}$, and $N_{w \succ z}^{\mathbf{R}} = N_{w \succ z}^{\mathbf{R}'}$ for all $\{w, z\} \neq \{x, y\}$.

That is, (a) before the move from \mathbf{R} to \mathbf{R}' there is a (possibly weak) majority for $x \succ y$, (b) after the move *all* voters agree on $x \succ y$, and (c) no other relative rankings change in the process. The axiom says that such a move decreases (or at most maintains) diversity. This axiom is relatively weak: it only applies if every voter either already ranks x above y , or if she ranks y *directly* above x and thus has the opportunity to swap them without affecting other rankings.

3.2 Results

Which PDO's satisfy which axioms? First, there is a group of three very weak axioms that will be satisfied by any reasonable PDO. In particular, as is easy to check, they are satisfied by the three specific families of PDO's defined in Section 2.3.

Fact 1. Every PDO induced by a PDI of the form $\Delta_{supp}^{\ell=k}$, $\Delta_{dist}^{\Phi, \delta}$, or $\Delta_{com}^{\Phi, F}$ with $k \in \{1, \dots, m\}$, $\Phi \in \{\Sigma, \max\}$, $\delta \in \{K, S, D\}$, and F being an anonymous and neutral SWF is anonymous, neutral, and weakly discernible.

At the other extreme, the axiom of strong discernability is not satisfied by any of our specific PDO's. The following impossibility result illustrates the overly demanding character of this axiom.

Proposition 2. For $m > 2$ and $n > m!$, no PDO can be both support-invariant and strongly discernible.

Proof. We first derive an *upper bound* on the dimension of any PDO that is support-invariant. The number of possible preference orders is $m!$. A support-invariant PDO has to determine the level of a given profile \mathbf{R} based on $\text{SUPP}(\mathbf{R})$ alone. There are $2^{m!} - 1$ nonempty subsets of the set of all possible preferences, i.e., there are at most $2^{m!} - 1$ possible sets of support. Hence, the maximal dimension of any support-invariant PDO is $2^{m!} - 1$.

Next, we derive a *lower bound* on the dimension of any PDO that is strongly discernible. There are $(m!)^n$ distinct profiles. Let us first partition this space into clusters of profiles such that any two profiles that are reachable from one another via a permutation on agents are placed into the same cluster. There are $\binom{m!}{n} = \binom{m!+n-1}{n}$ such clusters:⁴ for each of the $m!$ possible preferences we have to decide

⁴ Recall from basic combinatorics that $\binom{n}{k} = \binom{n+k-1}{k}$ is the number of solutions to the equation $x_1 + \dots + x_n = k$ in nonnegative integers.

how many agents should hold that preference, with the total number of agents adding up to n . Let us now again partition this space of clusters into larger clusters, such that any two profiles reachable from each other via a permutation of alternatives are also in the same cluster. The number of these clusters is the lowest possible dimension of any PDO that is strongly discernible. Computing this number is a demanding combinatorial problem that has been studied, amongst others, by Egecioğlu [7]. Closed formulas are known only for certain special cases. However, for our purposes a lower bound is sufficient. There are $m!$ possible permutations of the alternatives. Hence, each of the large clusters can contain at most $m!$ of the small clusters. Thus, $\binom{m!}{n} / m!$ is a lower bound on the number of clusters.

$$\begin{aligned} \frac{\binom{m!}{n}}{m!} &= \frac{m! + n - 1}{n} \times \frac{m! + n - 2}{n - 1} \times \dots \times \frac{m! + 1}{2} \\ &\geq \frac{m! + m!}{m! + 1} \times \frac{m! + m! - 1}{m!} \times \dots \times \frac{m! + 1}{2} \end{aligned}$$

The denominator of the leftmost factor is equal to the numerator of the rightmost one. So we can rewrite as $\prod_{i=2}^{m!} \frac{m!+i}{i}$. Now, all of the $m! - 1$ factors of this product are at least equal to 2. The first one (with $i = 2$) furthermore is at least equal to 4 (for $m > 2$). Hence, $\binom{m!}{n} / m! \geq 2^{m!}$. This concludes the proof, as it shows that the upper bound is strictly smaller than the lower bound derived. \square

Our next result is a characterisation of the *simple support-based* PDO, i.e., the PDO induced by $\Delta_{supp}^{\ell=m}$. That is, this is the PDO \succsim defined as $\mathbf{R} \succsim \mathbf{R}'$ if and only if $|\text{SUPP}(\mathbf{R})| \geq |\text{SUPP}(\mathbf{R}')|$. As we shall see in Section 5, this result is closely related to a classical theorem on ranking opportunity sets due to Pattanaik and Xu [16].

Proposition 3. *A PDO is support-invariant, independent, and weakly discernible if and only if it is the simple support-based PDO.*

Proof (sketch). First, observe that the simple support-based PDO clearly satisfies all three axioms. For the other direction, let \succsim be any PDI that is support-invariant, independent, and weakly discernible. We need to show that $\mathbf{R} \succsim \mathbf{R}'$ if and only if $|\text{SUPP}(\mathbf{R})| \geq |\text{SUPP}(\mathbf{R}')|$. This is equivalent to proving the following two claims:

- (1) $|\text{SUPP}(\mathbf{R})| = |\text{SUPP}(\mathbf{R}')|$ implies $\mathbf{R} \sim \mathbf{R}'$.
- (2) $|\text{SUPP}(\mathbf{R})| = |\text{SUPP}(\mathbf{R}')| + 1$ implies $\mathbf{R} \succ \mathbf{R}'$.

We shall make repeated use of the following fact: By support invariance, for every profile \mathbf{R} and every preference $R \in \text{SUPP}(\mathbf{R})$, there exists a profile \mathbf{R}' of the same size that has the same support and in which R occurs exactly once.

We first prove claim (1) by induction on $k = |\text{SUPP}(\mathbf{R})|$. If $k = 1$, then both profiles are unanimous and we are done. Now assume the claim holds for k and consider two profiles \mathbf{R} and \mathbf{R}' with support of size $k+1$. First, suppose \mathbf{R} and \mathbf{R}' share at least one preference R . W.l.o.g., assume R occurs exactly once in each of them. Let $\hat{\mathbf{R}}$ be the rest of \mathbf{R} and let $\hat{\mathbf{R}}'$ be the rest of \mathbf{R}' , i.e., $\mathbf{R} = \hat{\mathbf{R}} \oplus R$, $\mathbf{R}' = \hat{\mathbf{R}}' \oplus R$, and $R \notin \text{SUPP}(\hat{\mathbf{R}}) \cup \text{SUPP}(\hat{\mathbf{R}}')$. As $\hat{\mathbf{R}} \sim \hat{\mathbf{R}}'$ by the induction hypothesis, we thus obtain $\mathbf{R} \sim \mathbf{R}'$ from the left-to-right direction of the independence axiom. In case \mathbf{R} and \mathbf{R}' do not share any single preference R , this construction is not applicable. In this case, let R be a preference with single occurrence in \mathbf{R} and let $\hat{\mathbf{R}}$ be such that $\mathbf{R} = \hat{\mathbf{R}} \oplus R$. Now consider \mathbf{R}' and $\hat{\mathbf{R}} \oplus \mathbf{R}'$ for some $\mathbf{R}' \in \text{SUPP}(\mathbf{R}')$. These two profiles do share a preference, so we have $\mathbf{R}' \sim \hat{\mathbf{R}} \oplus \mathbf{R}'$. We can then repeat the same argument for \mathbf{R} and $\hat{\mathbf{R}} \oplus \mathbf{R}'$, which also share a preference, and obtain $\mathbf{R} \sim \hat{\mathbf{R}} \oplus \mathbf{R}'$. Thus, $\mathbf{R} \sim \mathbf{R}'$ follows in all cases.

For claim (2), we can use the same technique. For the base case of the induction we now use weak discernability to show that $\mathbf{R} \succ \mathbf{R}'$ when $|\text{SUPP}(\mathbf{R})| = 2$ and $|\text{SUPP}(\mathbf{R}')| = 1$. For the induction step we now use the right-to-left direction of independence (which is equivalent to the left-to-right direction for \succ rather than \succsim). \square

We stress that the crucial axiom in this last result is independence. In particular, support-invariance only says that a profile's diversity must be computable from its support, but it does not say that the support's cardinality needs to play any role in this process.

So far we have discussed our weakest and our most restrictive axioms. The remaining axioms tend to be satisfied by some reasonable PDO's and not by others, which means that they are helpful in structuring the space of all reasonable PDO's. For the main specific PDI's considered in this paper (and, more precisely, for the PDO's they induce), Table 1 summarises which of them satisfy which axioms. In the interest of space, we do not include proofs for the claims made in the table, but in most cases these claims are relatively easy to verify. Note that in some cases we state a sufficient (not always necessary) condition for a particular PDI to satisfy a particular axiom. For example, swap-monotonicity is satisfied by a compromise-based PDI if the SWF it is based on satisfies the *Pareto principle* and Arrow's *independence of irrelevant alternatives* [10]. Strong discernability is omitted from the table, as it is not satisfied by any of our PDI's.

	$\Delta_{supp}^{\ell=k}$	$\Delta_{dist}^{\Sigma, \delta}$	$\Delta_{dist}^{\max, \delta}$	$\Delta_{com}^{\Sigma, F}$	$\Delta_{com}^{\max, F}$
Support-invar.	✓	×	✓	×	×
Nonlocality	$n \leq k!$	✓	×	✓	×
Independence	$k = m$	×	×	×	×
Monotonicity	✓	×	✓	×	×
Swap-monoton.	✓	$\delta = K$	$\delta = K$	F is Arrovian	

Table 1. Classification of PDI's in terms of axioms.

4 EXPERIMENTAL ANALYSIS

In this section, we report on an experimental analysis of our PDI's. The experiments conducted fall into two classes. In the first kind of experiment we draw profiles from a given distribution and plot diversity values against the frequency of drawing profiles with these values. In the second kind, we investigate to what extent increasing diversity correlates with an increase of unwanted social choice-theoretic effects, such as the existence of Condorcet cycles.

Our findings are relative to the distribution over preference profiles from which we sample. We use two distributions. The first is the synthetic distribution generated by the *impartial culture assumption* (IC). This is the assumption that every possible profile is equally likely to occur. Despite its well-known limitations [17], this is the most widely used assumption in experimental work on social choice theory and serves as a useful base line. To generate the second distribution we have sampled from the (second) *AGH Course Selection* dataset available from PREFLIB, an online library of datasets concerning preferences [13]. This is a dataset with the complete preferences of 153 students regarding 7 course modules, collected at AGH University of Science and Technology, Kraków, in 2004. We have generated profiles by choosing, uniformly at random, 50 individual preferences regarding the first 5 courses.

All experimental results presented here concern scenarios with 5 alternatives and 50 voters; the effects are similar for scenarios of similar but different size. For each experiment we have drawn 1 million




Figure 1. Preference diversity (x -axis) against frequency (y -axis) in impartial cultures and amongst AGH students. [$n = 50, m = 5$]




Figure 2. Diversity for $\Delta_{dist}^{\Sigma,K}$ / IC data (x -axis). Condorcet winners/cycles; agreement between voting rules; voter satisfaction (y -axis). [$n = 50, m = 5$]

profiles from the relevant distribution. However, for most PDI's, profiles with very low or very high diversity have extremely low probability of occurring. For example, only $m!$ in $(m!)^n$ profiles are unanimous and thus have diversity 0 under every PDI. To be able to present our data in an illustrative manner, we therefore apply the following *pseudo-normalisation*. For a given PDI Δ and a given sample of profiles, let α_{\min} be the largest real number such that at most 1% of the profiles have a diversity value below α_{\min} . Analogously, let α_{\max} be the smallest real number such that at most 1% of the profiles have a diversity value above α_{\max} . We then plot the pseudo-normalised PDI Δ' with $\Delta'(\mathbf{R}) = \frac{\Delta(\mathbf{R}) - \alpha_{\min}}{\alpha_{\max} - \alpha_{\min}}$. Note that, strictly speaking, Δ' is not a PDI itself, as it can return values below 0. Also, as we plot diversity values from 0 to 1 only, up to 2% of the data may not be shown. What we gain in return is that we do not need to plot very long tails that only represent insignificantly small amounts of data. For all our plots, the x -axis ranges from 0 to 1.

4.1 Diversity distribution across cultures

Figure 1 shows, for both the IC and the AGH data, the relative frequency of each diversity value for four of our PDI's. Recall that each plot is showing around 99.8% of the data, after pseudo-normalisation. We can make two observations. First, all four PDI's result in what we judge to be reasonable frequency distributions, for both IC and AGH: very high and very low diversity are very rare, and there is a clear peak. Second, the AGH data results in a distribution where the peak is further to the left than for the IC data. This is what we would expect, and what we would want a good PDI to show: real preference profiles have more internal structure than purely random data, so we would expect to see less diversity. The simple support-based PDI is least able to show this difference.

A feature of the data that, due to our pseudo-normalisation, is not shown in Figure 1 is the number of distinct levels that the 1 million profiles we sampled ended up in. This data is shown, for the four PDI's of Figure 1 and five additional ones, in Table 2. We can make two observations. First, the support-based PDI's and the distance-based PDI using the max-operator make use of very few levels. This arguably makes them less attractive than the other PDI's. Second, the range of levels used is generally (much) larger for the AGH data than

for the IC data (which explains the increased levels of noise for the AGH data in Figure 1). In particular, an IC profile is very unlikely to have very low diversity. Thus, the range of levels observed is another criterion we can use to tell apart synthetic data and data based on real preferences. Overall, the distance-based PDI's using the Σ -operator for aggregation emerge as the most useful PDI's.

PDI	IC	AGH	PDI	IC	AGH	PDI	IC	AGH
$\Delta_{supp}^{\ell=m}$	22	13	$\Delta_{dist}^{\Sigma,D}$	34	244	$\Delta_{com}^{\Sigma,Bor}$	84	85
$\Delta_{supp}^{\ell=2}$	1	2	$\Delta_{dist}^{\Sigma,S}$	462	1170	$\Delta_{com}^{\Sigma,MG}$	94	88
$\Delta_{supp}^{\ell=3}$	4	12	$\Delta_{dist}^{\Sigma,K}$	660	1561	$\Delta_{dist}^{\Sigma,max,K}$	2	3

Table 2. Observed number of levels ($n = 50, m = 5$).

4.2 Impact on social choice-theoretic effects

Intuitively speaking, the less diverse a profile, the better behaved it should be from the perspective of social choice theory. Next, we report on three experiments where we put this intuition to the test for the PDI $\Delta_{dist}^{\Sigma,K}$ and data generated using the IC assumption. The results are shown in Figure 2 (diversity values against percentages).

In the first experiment we have measured the frequency of observing an *Condorcet cycle* (a cycle in the majority graph) in a profile and the frequency of a profile having a *Condorcet winner* (an alternative that wins against any other alternative in a pairwise majority contest).⁵ Figure 2 shows that, as diversity increases, so does the probability of encountering a Condorcet cycle, while the probability of finding a Condorcet winner decreases. This is exactly the behaviour we would like a good PDI to display, as it helps us predict good and bad social choice-theoretic phenomena.

The second experiment concerns the extent to which different voting rules agree on the winner for a given profile. For two *irresolute* voting rules, which may sometimes return a set of tied winners, we require a suitable definition for their degree of agreement under a given profile. For voting rules F_1 and F_2 , let W_1 and W_2 be the sets of winners we obtain. We define their *degree of agreement* as

⁵ Note that you may observe a Condorcet cycle and still find a Condorcet winner (namely, when the cycle does not occur amongst the top alternatives).

$\frac{|W_1 \cap W_2|}{|W_1| \times |W_2|}$. This is the probability of picking the same unique winner if each voting rule were to be paired with a uniformly random tie-breaking rule. Figure 2 shows the average degree of agreement for profiles with a given PDI-value for three pairs of well-known voting rules [10]: Plurality/Borda, Plurality/Copeland, Borda/Copeland. The plurality rule is widely regarded as a low-quality rule and this shows also here, as it disagrees considerably with the other two rules. This effect increases drastically as diversity increases.

Finally, we have computed the average *voter satisfaction* under the Borda rule. To this end, we define the satisfaction of a voter as the number of alternatives she ranks below the Borda winner. When normalised to percent, a unanimous profile would result in a satisfaction of 100%, while a satisfaction below 50% is not possible for the Borda rule. Figure 2 again clearly shows how voter satisfaction decreases with increased diversity and how it gets close to the absolute minimum of 50% for very high (and rare) values of diversity.

5 RELATED WORK

Our model is related to the literature on *freedom of choice* concerned with the ranking of alternative opportunity sets [15, 16], dealing with questions such as whether a choice between a bike and a car provides more freedom than the choice between a red car and a blue car. Conceptual differences aside, an important mathematical difference between ranking opportunity sets and ranking preference profiles in terms of diversity is that we only compare profiles of *the same size*, while two opportunity sets to be compared may have different cardinalities. This means that no direct transfer of results is possible. Still, a seminal result in this field, due to Pattanaik and Xu [16], has inspired our Proposition 3. They show that the only method of ranking opportunity sets satisfying three basic axioms they propose is the method of simply counting the number of options in each set. Their axioms are *independence* (of which ours is a direct translation), *indifference between no-choice situations* requiring any two singletons to be ranked at the same level (this requirement is part of our definition of a PDI), and a *strict monotonicity* axiom comparing sets of cardinality 1 and 2. The latter is not meaningful, or even expressible, in our framework. However, our weak discernability axiom has similar consequences. Pattanaik and Xu interpret their result as an impossibility result, given that simply counting opportunities is an overly simplistic way of measuring freedom of choice. As our empirical results suggest that the simple support-based PDI is not very attractive, Proposition 3 may be also be considered an impossibility result.

More expressive models of diversity, such as the multi-attribute approach of Nehring and Puppe [15] with its applications to the study of biodiversity, are not directly comparable to our setting.

Most closely related to our model is recent work on the *cohesiveness* (or the degree of *consensus*) of a profile [1, 2], which is the opposite of our notion of diversity. These studies focus on a generalisation of the Kendall tau distance, i.e., on measures based on averaging over pairwise distances between preferences (which can be seen as a special case of our distance-based measures) or the dual of this definition (averaging over the differences in the support of all possible pairs of alternatives). They also define several axioms (similar to some of ours) that characterise this class of measures. They do not, however, study the relationship between cohesiveness and social choice-theoretic phenomena.

Our compromise-based PDI's are related to *distance-based rationalisations* of voting rules [8, 14]. Such a rationalisation consists of a distance measure and a notion of consensus profile (e.g., a unanimous profile or one with a Condorcet winner): the winners are the alterna-

tives that win in the consensus profile that is closest (in terms of the distance measure) to the actual profile. What our compromise-based PDI's measure is precisely such a distance to a unanimous profile.

6 CONCLUSION

We have introduced the concept of *preference diversity*, together with a formal model facilitating the analysis of this concept. Besides being of interest in its own right, we also hope that PDI's may serve as a useful tool for parameterising data in research on preference handling and social choice, including applications in AI.

In the interest of space, we have focussed on three families of specific PDI's, but there is in fact a rich landscape of additional options that should be investigated in depth. For instance, we may count the maximal number preferences sharing a common subpreferences of a given length ℓ ; we may measure the maximal distance between all preferences in a given profile and all preferences not in the profile (to see how close a profile is to covering the full space of possibilities); or we may measure the distance to a single-peaked profile. In fact, the latter is a problem that already has received some attention in the literature [4]. Finally, we may use other distances and other aggregation operators (e.g., max-of-min) than those mentioned in Section 2.3.

REFERENCES

- [1] J. Alcalde-Unzu and M. Vorsatz, 'Measuring the cohesiveness of preferences: An axiomatic analysis', *Social Choice and Welfare*, **41**(4), 965–988, (2013).
- [2] R. Bosch, *Characterizations of Voting Rules and Consensus Measures*, Ph.D. dissertation, University of Tilburg, 2006.
- [3] F. Brandt, V. Conitzer, and U. Endriss, 'Computational social choice', in *Multiagent Systems*, ed., G. Weiss, 213–283, MIT Press, (2013).
- [4] R. Bredereck, J. Chen, and G. J. Woeginger, 'Are there any nicely structured preference profiles nearby?', in *Proc. 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, (2013).
- [5] C. Domshlak, E. Hüllermeier, S. Kaci, and H. Prade, 'Preferences in AI: An overview', *Artificial Intelligence*, **175**(7), 1037–1052, (2011).
- [6] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar, 'Rank aggregation methods for the web', in *Proc. 10th International World Wide Web Conference (WWW)*. ACM, (2001).
- [7] Ö. Eğecioğlu, 'Uniform generation of anonymous and neutral preference profiles for social choice rules', *Monte Carlo Methods and Applications*, **15**(3), 241–255, (2009).
- [8] E. Elkind, P. Faliszewski, and A. Slinko, 'Distance rationalization of voting rules', in *Proc. 3rd International Workshop on Computational Social Choice (COMSOC)*. University of Düsseldorf, (2010).
- [9] W. Gaertner, *Domain Conditions in Social Choice Theory*. Cambridge University Press, 2001.
- [10] W. Gaertner, *A Primer in Social Choice Theory*, LSE Perspectives in Economic Analysis. Oxford University Press, 2006.
- [11] W. V. Gehrlein, 'Condorcet's paradox', *Theory and Decision*, **15**(2), 161–197, (1983).
- [12] J. Goldsmith and U. Junker, 'Preference handling for artificial intelligence', *AI Magazine*, **29**(4), 9–12, (2008).
- [13] N. Mattei and T. Walsh, 'Preflib: A library of preference data', in *Proc. 3rd International Conference on Algorithmic Decision Theory (ADT)*. Springer-Verlag, (2013). <http://www.preflib.org>.
- [14] T. Meskanen and H. Nurmi, 'Closeness counts in social choice', in *Power, Freedom, and Voting*, Springer-Verlag, (2008).
- [15] K. Nehring and C. Puppe, 'A theory of diversity', *Econometrica*, **70**(3), 1155–1198, (2002).
- [16] P. K. Pattanaik and Y. Xu, 'On ranking opportunity sets in terms of freedom of choice', *Recherches Économiques de Louvain*, **56**(3/4), 383–390, (1990).
- [17] M. Regenwetter, B. Grofman, A. A. J. Marley, and I. Tsetlin, *Behavioral Social Choice: Probabilistic Models, Statistical Inference, and Applications*, Cambridge University Press, 2006.
- [18] T. Walsh, 'Where are the hard manipulation problems?', *Journal of Artificial Intelligence Research*, **42**, 1–29, (2011).