Mathematical Social Sciences 96 (2018) 49–62

Contents lists available at ScienceDirect

Mathematical Social Sciences

journal homepage: www.elsevier.com/locate/mss

A probabilistic evaluation framework for preference aggregation reflecting group homogeneity

C. Jansen *, G. Schollmeyer, T. Augustin

Department of Statistics, LMU Munich, Ludwigsstr. 33, 80539 Munich, Germany

HIGHLIGHTS

- We propose a framework for evaluating preference aggregation functions.
- Our framework explicitly reflects group specific homogeneity.
- A new preference homogeneity measure is introduced.
- We approximate expected similarity by maximum entropy and credal set approaches.
- Different aggregation functions are compared in a simulation study.

ARTICLE INFO

Article history:
Received 5 August 2016
Received in revised form 22 March 2018
Accepted 4 September 2018
Available online 14 September 2018

ABSTRACT

Groups differ in the homogeneity of their members' preferences. Reflecting this, we propose a probabilistic criterion for evaluating and comparing the adequateness of preference aggregation procedures that takes into account information on the considered group's homogeneity structure. Further, we discuss two approaches for approximating our criterion if information is only imperfectly given and show how to estimate these approximations from data. As a preparation, we elaborate some general minimal requirements for measuring homogeneity and discuss a specific proposal for a homogeneity measure. Finally, we investigate our framework by comparing aggregation rules in a simulation study.

1. Introduction

One of the fundamental tasks in social choice theory is to define adequately justified rules for aggregating the preferences of a group of individuals into one global consensus order. Due to the generality of this problem, it is hardly surprising that many different rules have been proposed since the pioneering works by de Borda (1781), de Condorcet (1785) and Hare (1857) (see Brams and Fishburn, 2002 for a survey). More generally, the question of aggregating collections of binary relations in a meaningful way does not exclusively concern social choice theory, but also appears in classification problems in statistics (see, e.g., Maniqueta and Mongin, 2016), benchmarking of algorithms in the computer sciences (see, e.g., Mersmann et al., 2015) or problems of judgment aggregation in philosophy (see, e.g., Hartmann and Sprenger, 2012) to name only a few examples.

Given the diversity of aggregation rules, criteria for evaluating and comparing their quality need to be established. Many different criteria have been proposed, and comparisons of aggregation rules with respect to them have been studied intensively (see, e.g., Grofman and Feld, 2004). However, almost all these criteria are non-group-specific: They are intended to be valid independently of the group whose members' preferences are to be aggregated. But what is a perfectly adequate aggregation procedure for one group may not be as appropriate for another one. The adequateness of an aggregation procedure may, beyond compatibility with non-group-specific criteria, additionally depend on certain characteristics of the specific group under consideration. One such characteristic is the homogeneity of the group members (see Section 2.2 for a discussion of the literature on homogeneity). In this paper we propose a group-specific quality criterion for aggregation rules that takes into account information on the homogeneity of group members' preference structure. Moreover, we show different ways to approximate our criterion under partial probabilistic information and discuss how to estimate these approximations in the presence of data or expert knowledge.

More precisely, the paper is structured as follows: In Section 2, we discuss measures for quantifying the homogeneity of a group that is represented by a fixed profile $(R_1, \dots, R_n)$ of strict weak orders. Specifically, in Section 2.3, we elaborate a list of three minimal requirements that every reasonable measure should satisfy.

* Corresponding author.
E-mail address: christoph.jansen@stat.uni-muenchen.de (C. Jansen).

https://doi.org/10.1016/j.mathsocsci.2018.09.001

In Section 2.4, we then propose a concrete measure, the maximum consensus homogeneity, and discuss why it is reasonable beyond its mere compatibility with these minimal requirements. Section 3, after reviewing some basics on Bayesian theory in Section 3.1, introduces a framework for evaluating and comparing aggregation procedures in the presence of probabilistic information on the considered group. This involves three steps: In Section 3.2, we introduce an optimality criterion that requires perfect knowledge of the probabilities with respect to which the group constitutes different profiles $(R_1, \dots, R_n)$. Section 3.3 discusses approaches for approximating this criterion if the probabilistic information on the group is partial in the sense that only the probability distribution of some homogeneity measure is given. Finally, Section 3.4 discusses several statistical approaches for estimating this distribution in the presence of data, expert knowledge, or both. Section 4 starts by briefly reviewing some common aggregation procedures relevant to our context (Section 4.1). Afterwards, Section 4.2 summarizes an aggregation procedure recently proposed in Schollmeyer (2017).[1] In Section 5, we investigate the aggregation procedures reviewed, in respect to our criterion in a simulation study. Section 6 is reserved for concluding remarks as well as an outlook on future research questions.

2. Measuring homogeneity of preference profiles

We begin the section by introducing our notation and terminology (Section 2.1) and surveying some related work on the topic (Section 2.2). Subsequently, we establish and discuss a weak set of conditions (Section 2.3) as well as a concrete proposal (Section 2.4) for measuring the homogeneity of a fixed collection $(R_1, \dots, R_n)$ of strict weak orders each of which is representing the opinion of a member of a group of size n.

2.1. Notation and terminology

Throughout the paper, C denotes a finite set of at least two consequences. The elements of C have to be ranked by the members of a specific group G_n of fixed size $n \geq 2$, where certain requirements of rationality regarding the individual orders involved are imposed. Specifically, we work with the following spaces of binary relations on C:

$$\mathcal{R} := \{R \subset C^2 : R \text{ asymmetric, negatively transitive}\} \tag{1}$$

$$\mathcal{Q} := \{Q \subset C^2 : Q \text{ asymmetric}\} \tag{2}$$

In the sequel, every $R \in \mathcal{R}$ is termed a strict weak order on C. For every $R \in \mathcal{R}$, define the usual equivalence relation $\sim_R$ on C by setting $a \sim_R b$ if and only if $(a, b) \notin R \wedge (b, a) \notin R$. Given this, interpret $(a, b) \in R$ as a is strictly preferred to b and $(a, b) \in \sim_R$ as indifference between a and b. The elements of $\mathcal{R}$ are associated with the individual orders of the group members. Hence, the group members are assumed to have asymmetric and negatively transitive preferences. Importantly, note that our model of the individual preferences excludes incomparability of consequences: For alternatives $a, b \in C$ chosen arbitrarily, every group member is thus assumed to be able to decide if she strictly prefers a to b, or b to a, or if she ranks them equally desirable. Thus, we explicitly assume that incomparability with respect to $R \in \mathcal{R}$ is interpreted as indifference (see, e.g., Kreps, 1988, Chapter 2 for a discussion of this convention).[2] For $n \geq 2$, an element $\underline{R} := (R_1, \dots, R_n) \in \mathcal{R}^n$ is called a preference profile on C and each component of $\underline{R}$ is interpreted as the opinion of a member of G_n about how the consequences in C should be ranked.

Contrarily, every element $Q \in \mathcal{Q}$ is called a consensus order (or group preference). Except for asymmetry, we do not impose any further restrictions on the consensus order. This allows for also investigating aggregation procedures for which the group preference is not always as well-behaved as the individual orders (this includes, e.g., Condorcet's method, see Section 4.1, which might yield intransitive consensus orders). In this context, every mapping $S : \mathcal{R}^n \to \mathcal{Q}$ is called a preference aggregation function. Particularly, for every preference profile $\underline{R} \in \mathcal{R}^n$, the image $S(\underline{R}) \in \mathcal{Q}$ is the consensus order of the group represented by $\underline{R}$ with respect to the aggregation procedure described by S.

2.2. Preference homogeneity in related work

In literature on social choice theory at least two different lines of how to establish a notion of homogeneity of groups can be identified. One line (see, e.g., Niemi, 1969; Jamison and Luce, 1972; Berg, 1985; Gehrlein and Lepelley, 2010; Lepelley and Valognes, 2003), which could be called "model-based", builds up stochastic models that govern the constitution of profiles and have specific parameters implicitly regulating the group's homogeneity. One prominent example is the multivariate Pólya–Eggenberger urn model (see, e.g., Johnson and Kotz, 1977), which has been used for instance in Berg (1985), Gehrlein and Lepelley (2010) and Lepelley and Valognes (2003) in order to analyze the relationship between group homogeneity and the probability of the voting paradox or the manipulability of different aggregation functions. The Pólya–Eggenberger model contains two other well-established models as special cases: impartial culture and impartial anonymous culture, which are also often presumed in studies of the voting paradox and the manipulability of aggregation procedures (see, e.g., Aleskerov et al., 2012; Diss et al., 2012; Pritchard and Slinko, 2006). Other model-based approaches, in which the orders in the profile are assumed to be randomly drawn with replacement, measure the homogeneity of the generating process by the probabilities p_i ($i = 1, \dots, |C|!$) with respect to which the order R_i is drawn: Natural measures of homogeneity are then the variance of the p_is used for instance in Abrams (1976) or the Herfindahl index $\sum_{i=1}^{|C|!} p_i^2$ used, for instance, in Gehrlein (1981). Measures that only rely on the values of the p_i's and not on the concrete associated orders R_i are called non-profile specific measures (see Gehrlein, 1981). Since they are related to the probabilities p_i, they are also called population specific homogeneity measures in Gehrlein and Lepelley (2010, p. 191).

A second line of establishing a notion of homogeneity, which can be called "data-based", relates homogeneity not to a probabilistic model but to the actually observed data in a profile. For example, in the above approaches, one can replace the probability p_i of observing the order R_i in a profile with the relative frequency of the associated order in the actually observed profile. Then one arrives at a notion of homogeneity that is not related to a generating process, but instead related to the observed profile. Such measures are called situation specific homogeneity measures in Gehrlein and Lepelley (2010, p. 192). A further type of such data-based measures are distance-based measures, which additionally utilize the information in the orders of the profile. These measures, arising not only in social choice theory but also in statistics and computer sciences (see, e.g., Fligner and Verducci, 1986; Dwork et al., 2001), rather rely on a geometric understanding and first introduce a distance between pairs of orders. Based on this distance, one defines a measure of heterogeneity by computing the average distance of all pairs of orders in the profile. Homogeneity of the profile is then measured by comparing the maximal distance to

[1] For an explanation of the procedure and a discussion see Section 4.2.

[2] An alternative approach would be to directly model the individual preferences by weak orders, i.e. complete and transitive binary relations $P \subset C^2$. To every such relation we then can associate its strict part $R_P \subset C^2$ by setting $(a, b) \in R_P$ if and only if $(a, b) \in P \wedge (b, a) \notin P$ for all $a, b \in C$. The relation R_P is then asymmetric and negatively transitive. Our model thus explicitly assumes that the individual orders $R \in \mathcal{R}$ arise as strict parts of a weak order.