Journal of Mathematical Psychology 66 (2015) 70–82

Contents lists available at ScienceDirect

Journal of Mathematical Psychology

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

Individual differences in the algebraic structure of preferences

Clintin P. Davis-Stober [a,*], Nicholas Brown [a], Daniel R. Cavagnaro [b]

[a] University of Missouri, USA

[b] Mihaylo College of Business and Economics, California State University Fullerton, P.O. Box 6848, Fullerton, CA 92834-6848, USA

HIGHLIGHTS

- We examine the algebraic structure of individual preference.
- We evaluate a large class of weak-order and lexicographic semiorder based theories.
- We present a new study as well as a re-analysis of an existing data set.
- We find that a majority of subjects' preferences are consistent with weak orders.
- We find that the remaining subjects are well-described by lexicographic semiorders.

ARTICLE INFO

Article history:
Received 14 December 2013
Received in revised form
29 October 2014
Available online 2 March 2015

Keywords:
Preferential choice
Choice modeling
Lexicographic semiorders
Weak orders
Convex polytopes
Random preference
Order-constrained inference

ABSTRACT

Two divergent theories regarding the algebraic structure of preferences are the strict weak-order (i.e., utility) representation, and the lexicographic semiorder representation. We carry out a novel comparison of these theories by formulating them as mixture models of ternary choice that are general yet parsimonious. We apply Bayesian model selection to see which representation (if any) best explains each decision maker's choices across multiple data sets. We report the results of a new experiment, which tests the robustness of each representation with respect to manipulations of stimuli, display format, and time pressure. We find that a majority of participants are best described by strict weak-ordered preferences with a substantial minority best described by lexicographically semi-ordered preferences.

1. Introduction

When buyers choose from among different bundles of goods and services, their choices are assumed to be based on underlying states-of-mind called preferences. Theorists commonly impose idealized conditions on these preferences in order to generate tractable models. One such condition is that preferences can be represented by a unidimensional, numerical utility function $U(\cdot)$ such that alternative A is preferred to alternative B if and only if $U(A) > U(B)$. The existence of such a utility representation is assumed by many of the most prominent models of decision making, including Expected Utility Theory (von Neumann & Morgenstern, 1947) and Cumulative Prospect Theory (Tversky & Kahneman, 1992). Yet, despite its centrality in modeling preferential choice, the decision-making literature is internally divided on the question of whether a numerical utility function can well describe the actual choices of individual decision makers. In particular, numerous studies, beginning with Tversky's (1969) 'Intransitivity of Preference,' have questioned whether the preferences of human decision makers satisfy transitivity, which is a necessary condition for the existence of a utility representation as described above — see Mellers and Biagini (1994), Fishburn (1991), and Regenwetter, Dana, and Davis-Stober (2011) for comprehensive reviews of the arguments both for and against transitivity.

Tversky's (1969) experiment, and many of those that followed (e.g. Birnbaum, 2010; Birnbaum & Gutierrez, 2007; Montgomery, 1977; Ranyard, 1977, 1982), were designed to elicit intransitive choice patterns arising from a particular ordered collection of semi-ordered preferences, called a lexicographic semiorder. The idea of representing preferences as semiorders was introduced by Luce (1956) (although see Armstrong, 1939; Georgescu-Roegen, 1936, for earlier, related work) and extended to lexicographic semiorders by Tversky (1969, 1972). The core feature of decision models based on semiorders is that "small" differences in attribute values are ignored by the decision maker. A canonical example is that of a decision maker choosing between two cups of coffee: one without sugar, the other with one microgram of sugar. Since

* Correspondence to: University of Missouri, Department of Psychological Sciences, Columbia, MO, 65211, USA.
E-mail address: stoberc@missouri.edu (C.P. Davis-Stober).

http://dx.doi.org/10.1016/j.jmp.2014.12.003
0022-2496/

this amount of sugar is below what a human tongue can detect, the decision maker would be indifferent between the two cups of coffee. Similarly, a decision maker would likely be indifferent between two similar goods whose difference in price is a single US penny. The idea of a *lexicographic* semiorder representation of preference is that, when comparing any two choice alternatives, attribute values are compared sequentially via semiorders until a set of attribute values are reached on which the choice alternatives differ by a sufficient margin, i.e., are "distinguishable." At that point the process stops and the alternative that is superior based on that attribute is preferred.

Not all lexicographic semiorders are compatible with a utility representation, and not all utility representations are compatible with a lexicographic semiorder. Hence, these two representations constitute divergent theories of the algebraic structure of preferences. Representing preferences as lexicographic semiorders is intuitively appealing for its apparent simplicity and realism, and provides a model of bounded rational choice that can be characterized by direct axioms on choice behavior (Manzini & Mariotti, 2012). However, as with the utility representation of preferences, the literature remains divided on whether lexical-based heuristics, such as lexicographic semiorders, can accurately describe real human choice data. Proponents of the lexicographic semiorder representation tout the ecological rationality of fast and frugal heuristics from which semi-ordered preferences can arise (Gigerenzer & Brighton, 2009), yet recent tests of lexicographic models such as Take the Best (Gigerenzer & Goldstein, 1996) and the Priority Heuristic[1] (Brandstätter, Gigerenzer, & Hertwig, 2006), have found mixed to little empirical support (e.g. Birnbaum, 2008; Birnbaum & Gutierrez, 2007; Glöckner & Betsch, 2008; Lee & Cummins, 2004).

Why has the literature been unable to reach a consensus on the algebraic structure of preferences despite decades of research? We argue that four, long-standing, theoretical and methodological conventions have hindered progress:

1. Due to the mathematical complexity of higher-order choice structures, studies have used a binary forced choice framework that does not explicitly include indifference, even though indifference is a defining aspect of the lexicographic semiorder structure (see Regenwetter & Davis-Stober, 2012, for an in-depth discussion of this limitation).
2. Due to a lack of consensus on how to appropriately specify algebraic structures as stochastic models, different studies have used different stochastic specifications to test the algebraic structures, e.g., trembling hand (Harless & Camerer, 1994), true-and-error (Birnbaum, 2011), and random utility (Regenwetter, Dana, Davis-Stober, & Guo, 2011). This confounds the test of the algebraic structure with a test of the stochastic framework. See Hey (2005) and Loomes, Moffatt, and Sugden (2002) for discussion.
3. Due to the limitations of statistical analyses based solely on classical goodness-of-fit tests and/or the complications of order-constrained inference, many previous analyses have not appropriately penalized models for complexity (e.g. Tversky, 1969); — see Davis-Stober & Brown, 2011, for discussion. This can bias the results in favor of models that are more flexible but not necessarily more generalizable.
4. Due to the limitations of statistical analyses that are not well-suited for non-nested model comparison, many previous studies have not directly compared competing theories to one another (e.g., Regenwetter, Dana, Davis-Stober, & Guo, 2011; Regenwetter & Davis-Stober, 2012). As a result, rather than offering an alternative explanation to "rationalize" violating data, statistically significant violations are attributed to either irrational behavior or Type I error.

[1] Depending on the choice stimuli being considered, the Priority Heuristic is either a lexicographic semiorder or a lexicographic interval order.

In this article, we aim to bring some clarity to the issue of preference representation by addressing all of the limitations listed above with a re-analysis of existing data as well as our own new experiment. Using a Bayesian model selection methodology to directly evaluate competing mixture specifications of 'ternary choice,' we find that a majority of decision makers (about 80%) are best described by a numerical utility representation, while a substantial minority (about 20%) are best described by a lexicographic semiorder representation. Very few participants seem to violate both representations, even though decision makers choosing randomly would do so more than 99.9% of the time due to the extremely strong restrictions placed on choice data by our parsimonious mixture framework. Further details on our approach, and how it addresses the limitations listed above, are given next.

1.1. Mixture models of ternary choice

To address the limitations of binary choice, we test a new class of lexicographic semiorder mixture models (LSMM) for ternary choice data (Davis-Stober, 2010, 2012). The ternary choice framework extends the standard binary choice framework by allowing participants to report indifference, instead of forcing them to always report a strict preference, thereby providing a richer mapping between true preference and choice behavior. Under an LSMM, at every experimental time point, decision makers (DMs) are required to make choices consistent with a lexicographic semiorder over the choice alternatives, with the particular lexicographic semiorder used by the DM allowed to shift over the course of an experiment. In other words, on each choice, the DM is assumed to draw his or her preference from a mixture of lexicographic semiorders. Moving to ternary choice is critical for this model to be testable, because a mixture model on binary choice data would be unable to distinguish between the case of a DM truly being indifferent between two alternatives and a DM having a mixture of opposite, strict preferences. This class of model provides a very general instantiation of the lexicographic semiorder representation, while also being testable and extremely parsimonious, as we will show.

To test whether a numerical utility representation or a lexicographic semiorder representation provides a better description of individual preferences, we will compare the fit of our class of LSMMs to that of a viable competitor model. However, to overcome the limitations of methodological convention (2) from the list above, the competitor model should use the same stochastic specification. Such a competitor is provided by the weak order mixture model (WOMM) of Regenwetter and Davis-Stober (2012). Like the LSMM, the WOMM is also defined for ternary choice, and also allows DMs to move from one preference state to another. However, the WOMM requires that DMs always make choices consistent with a strict weak order (i.e., a ranking with ties) over the choice alternatives, rather than a lexicographic semiorder. In the ternary choice framework, a strict weak order representation of preferences is equivalent to a numerical utility representation, so the WOMM provides a general instantiation of the numerical utility representation that includes other utility-based models like Expected Utility Theory and Cumulative Prospect Theory as special cases.

Both the WOMM and LSMM models can be described as a type of distribution-free "random preference model" (e.g. Heyer & Niederée, 1992; Loomes & Sugden, 1995). This perspective allows for a general test of the mathematical structures of interest. By operating at the level of preference relations, we are not assuming any particular functional form that gives rise to them, i.e., we are not engaging in model-fitting. Similarly, by allowing an arbitrary distribution over the preference relations of interest, we are not limiting our results or analyses to any particular choice of mixture distribution or estimation method. In other words, the models