Journal of Mathematical Psychology 70 (2016) 1–11

Contents lists available at ScienceDirect

Journal of Mathematical Psychology

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

A statistical test of the equality of latent orders

Michael L. Kalish[a,*], John C. Dunn[b], Oleg P. Burdakov[c], Oleg Sysoev[c]

[a] *Syracuse University, USA*
[b] *University of Adelaide, Australia*
[c] *Linköping University, Sweden*

HIGHLIGHTS

- We present a measure of the difference in the latent orders of two variables.
- We present an algorithm for finding the minimum of this measure.
- We present a statistical test for the null hypothesis that the latent orders are the same.
- The test can be applied to any form of data, as long as an appropriate statistical model can be specified.
- The test allows hypothesis testing for designs analyzed with state trace analysis.

ARTICLE INFO

Article history:
Received 17 March 2015
Received in revised form
16 October 2015
Available online 19 November 2015

Keywords:
State-trace analysis
Monotonic regression
Hypothesis test

ABSTRACT

It is sometimes the case that a theory proposes that the population means on two variables should have the same rank order across a set of experimental conditions. This paper presents a test of this hypothesis. The test statistic is based on the coupled monotonic regression algorithm developed by the authors. The significance of the test statistic is determined by comparison to an empirical distribution specific to each case, obtained via non-parametric or semi-parametric bootstrap. We present an analysis of the power and Type I error control of the test based on numerical simulation. Partial order constraints placed on the variables may sometimes be theoretically justified. These constraints are easily incorporated into the computation of the test statistic and are shown to have substantial effects on power. The test can be applied to any form of data, as long as an appropriate statistical model can be specified.

1. Introduction

Consider an experiment in which data are obtained on two different variables across k different conditions. We would like to know if these data are drawn from populations whose means on the two variables have different orders. That is, we ask if the variables have unequal *latent orders*. This question arises in the theory of *state trace analysis* (STA) where inferences concerning the number of latent variables underlying changes in two or more dependent variables depend on the ordinal arrangements of their respective population means (Bamber, 1979; Prince, Brown, & Heathcote, 2012a). STA contrasts a *one-dimensional model*, in which changes in the dependent variables are mediated by one latent variable, and a *two-dimensional model*, in which changes are mediated by more than one latent variable (Loftus, Oberg, & Dillon, 2004; Newell & Dunn, 2008). Under the assumption of the one-dimensional model that each dependent variable is a (distinct) monotonic function of the single latent variable, this model predicts that the latent orders of the two variables are equal. It follows that if the variables have different latent orders across a set of experimental conditions then the effects must be mediated by more than one latent variable.

Implementation of STA requires a statistical procedure to test whether two sets of population means have the same order across a set of conditions. To our knowledge, at least three previous approaches to this problem have been proposed in the psychological literature. The first of these, described by Loftus et al. (2004), relies on reducing sampling error to near zero thereby using the observed sample means as a proxy for the population means. Clearly, this approach cannot be applied in situations with non-negligible sampling error and it lacks a means of quantifying when the sampling error is small enough to be ignored. The second

* Correspondence to: Department of Psychology, Syracuse University, Syracuse, NY, 13244, USA.
E-mail addresses: mlkalish@syr.edu (M.L. Kalish), john.c.dunn@adelaide.edu.au (J.C. Dunn), oleg.burdakov@liu.se (O.P. Burdakov), oleg.sysoev@liu.se (O. Sysoev).

http://dx.doi.org/10.1016/j.jmp.2015.10.004
0022-2496/

approach, described by Pratte and Rouder (2012), quantifies the effects of sampling error but is limited to particular theory-dependent variables and to a fixed two-by-two factorial design. The third approach, described by Prince et al. (2012a), uses Bayesian model selection to test whether two sets of population means have the same or different orders. While the approach is in principle quite general, the particular implementation described by Prince et al. (2012a) applies only to binomial data and to a relatively constrained factorial design. We discuss this approach in greater detail below and compare it to the test that we develop.

The test we present here is a null hypothesis statistical test (NHST), based on the computation of an empirical *p*-value of the data given the null hypothesis. Despite the well known problems with *p*-values (Wagenmakers, 2007), the evidence provided by them remains useful; e.g., it predicts future replicability (Open Science Collaboration, 2015).

The outline of the paper is as follows. First, we describe more fully the logic of our statistical test, based on an extension of monotonic regression (Burdakov, Dunn, & Kalish, 2012). In so doing, we introduce the concept of partial order constraints and foreshadow how they may be used to increase statistical power. Second, we describe a null hypothesis significance test of the equality of latent orders based on a bootstrap resampling procedure for estimating the empirical sampling distribution of the test statistic. Third, we examine the statistical power of our procedure for a fully randomized design with and without partial order constraints. Finally, we extend the procedure to binomial data and compare it to the Bayesian model selection approach developed by Prince et al. (2012a).

The orders of sample and population means

Consider two different dependent variables, x and y, observed across k different experimental conditions. Let $x_1, \ldots, x_k$, $y_1, \ldots, y_k$, be the k population means of each variable and let $X_1, \ldots, X_k, Y_1, \ldots, Y_k$, be the corresponding sample means. We define the (latent) order of x as a permutation, $O(x) = (i_1, i_2, \ldots, i_k)$, such that, $x_{i_1} \leq x_{i_2} \leq \ldots \leq x_{i_k}$. We wish to test the hypothesis that $O(x) = O(y)$, given the data. A desirable feature of such a test is that it should be sensitive to both the number and magnitude of differences in the two orders. Intuitively, given equal latent orders, numerically small violations of equality of the orders of the observed means are more likely than numerically large violations. This property is a feature of monotonic (or isotonic) regression (Robertson, Wright, & Dykstra, 1988). Our test is based on this method.

Monotonic regression

Monotonic regression addresses the problem of finding the best approximation, $\hat{X}$, to a set of observed values, X, under the constraint that $O(\hat{X})$ is known, either completely or partially. Let K be the set of integers, $\{1, 2, \ldots, k\}$. We represent a partial (or total) order on K by means of a subset of ordered pairs $(i, j) \in E \subseteq K \times K$.[1] An order, $O(\hat{X})$, is *consistent* with E if $\hat{X}_i \leq \hat{X}_j, \forall (i, j) \in E$. Formally, let X be a set of k values, let v be a set of corresponding weights, and let E be a partial order. Then monotonic regression finds a set of values, $\hat{X}$, consistent with E, that best approximates X in a weighted least-squares sense. That is, $\hat{X}$ solves the monotonic regression (MR) problem,

$$\min \sum_{i=1}^{k} v_i (X_i - \hat{X}_i)^2, \quad \text{subject to } \hat{X}_i \leq \hat{X}_j, \text{ for all } (i, j) \in E. \tag{1}$$

The choice of weights is critical for obtaining a meaningful 'best' $\hat{X}$. In this respect, we are guided by the property that the solution of Eq. (1) is the maximum likelihood estimate if the observations in each condition are independent and normally distributed with weights given by the precision of the data weighted by the number of observations in each condition (Robertson et al., 1988). That is,

$$v_i = \frac{n_{x_i}}{S_{X_i}^2} \qquad w_i = \frac{n_{y_i}}{S_{Y_i}^2} \tag{2}$$

where $S_{X_i}^2$ is the sample variance of variable x in condition i and $S_{Y_i}^2$ is the sample variance of variable y in condition i.

In many situations the observations in each condition are not independent, as when conditions are manipulated within participants rather than between. In this case the maximum likelihood estimate depends on the entire covariance matrix and the sets of weights, v_i and w_i, are replaced by appropriate matrices. For this reason, we generalize Eq. (2) in the following way. Suppose there are g groups of participants of size n_i, $i = 1, \ldots, g$, each measured under m different conditions on variable x. The total number of conditions is thus $k = gm$. Let $\mathbf{S}_i$ be the $m \times m$ covariance matrix for group i. Then the corresponding weight matrix is given by the following block-diagonal matrix,

$$V = \begin{bmatrix} n_1 \mathbf{S}_1^{-1} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & n_g \mathbf{S}_g^{-1} \end{bmatrix}. \tag{3}$$

The weight matrix, W, for variable y is similarly defined.[2] $\mathbf{S}_i^{-1}$ approximates the inverse of the population covariance matrix, $\mathbf{\Sigma}_i^{-1}$. A better estimate of $\mathbf{\Sigma}_i^{-1}$ can be obtained by first 'shrinking' $\mathbf{S}_i$, which reduces the unreliable off-diagonal elements but does not necessarily set all of them to zero (Ledoit & Wolf, 2004). We use Ledoit–Wolf method to adjust the weight matrices in our current approach.

Let X be a vector of k sample means and let $\hat{X}$ be a vector of values. Then, with the weight matrix V defined by Eq. (3), the MR problem is given by,

$$\min \left(X - \hat{X}\right)^T V \left(X - \hat{X}\right), \text{ subject to } \hat{X}_i \leq \hat{X}_j, \text{ for all } (i, j) \in E. \tag{4}$$

We write the problem corresponding to Eq. (4) as $\text{MR}(X, V, E)$ and the minimum value as $\omega(X, V, E)$, or, in shorthand form, as ω_X. Finding the solution to the MR problem is not trivial, but fast algorithms have been developed. If E is a total order then the MR problem can be solved using the *pool-adjacent-violators algorithm* (PAVA), a version of which was used in the original development of non-metric multidimensional scaling (Kruskal, 1964). Otherwise, the problem as posed in Eq. (4) can be solved using quadratic programming algorithms (de Leeuw, Hornik, & Mair, 2009). The functions *lsqlin* (equivalently, *quadprog*) and *lsei* implement this algorithm in MATLAB® and R (R Core Team, 2013) respectively. In addition, a rapid approximate solution may also be obtained using the *generalized pool-adjacent-violators* (GPAV) algorithm developed by Burdakov, Sysoev, Grimvall, and Hussian (2006).

---

[1] Unless otherwise stated, a partial order, E, is assumed to be transitively closed.

[2] We assume that observations on x and y are themselves independent.