ERASMUS UNIVERSITY ROTTERDAM
ERASMUS SCHOOL OF ECONOMICS

MASTER THESIS
ECONOMETRICS AND MANAGEMENT SCIENCE
BUSINESS ANALYTICS AND QUANTITATIVE MARKETING

Mediation and Moderation Analyses with the BANOVA R-package and Everything You Need to Know to Do it Right

---

Author:
Anna Kopyakova

Student number:
505237

Supervisor:
Prof. Dr. Michel Wedel

Second assessor:
Prof. Dr. Richard Paap

December 16, 2020

Abstract

In the social sciences, moderation and mediation analyses are widely used techniques for understanding how and why a certain effect occurs. This thesis implements two new functions in the framework of the BANOVA R-package, that are useful in these domains. The package allows its users to estimate a wide range of single-level and multi-level Bayesian models with Normal and non-Normal dependent variables, along with Bayesian ANOVA analysis. The new functionality in the package allows one to calculate Bayesian simple effects, relevant for answering if and for whom a given effect is observed. Furthermore, Bayesian mediation analysis with multiple correlated mediators was implemented to uncover why a certain effect manifests itself. To fit multiple mediators a seemingly unrelated regressions (SUR) model was added to the package. This method can also be used as a stand-alone technique for the analysis of Multivariate Normal dependent variables. The new functions are demonstrated via applications to two data sets from previously published studies. This paper additionally presents a detailed description of the BANOVA, along with a review of relevant concepts in Bayesian inference, which are useful for the users who are interested in an in-depth understanding of the analyses performed by the package.

Keywords: Moderation, mediation, Bayesian inference, Bayesian Analysis of variance, BANOVA R-package, simple effects, multiple mediators, seemingly unrelated regressions (SUR)

Contents

1 Introduction

Science is a journey in search of true beliefs about the world. It begins by identifying a certain phenomenon of interest and establishing associations between relevant concepts. Once the body of evidence for these effects is substantial, the research starts to search for an understanding of the underlying mechanisms behind the observed phenomenon. Understanding when and how a certain effect occurs is a necessary prerequisite for effective interventions that can nudge the world around us in a desired direction. Moderation and mediation analyses are the two most popular techniques in social sciences which investigate these underlying mechanisms (Hayes, 2017).

Moderation analysis helps to answer if or for whom a certain effect occurs. Its main goal is to model the relationship between two or more variables X and Y as a function of a third variable, referred to as a moderator W (Cohen et al., 2013; Hayes, 2017). The influence of W on the size and direction of the observed effects of X on Y is studied by computing conditional or simple effects of X at different values or levels of W. As an example, consider a study by Cohen and Wills (1985) who investigated the moderating influence of social support on the relationship between depression and stress. They found that the effect of stress on depression is stronger at lower levels of social support, which means that the association changes in strength depending on the value of the moderator.

Mediation analysis answers how a certain effect occurs by allowing for X to exhibit its influence on Y directly, as well as indirectly via another variable M (MacKinnon et al., 2007; Hayes, 2017). The latter variable is referred to as a mediator, and multiple mediators can operate simultaneously. The indirect effect of X, also referred to as mediated effect, can be found by combining outcomes of a regression of Y on X and M, with that of a regression of M on X. An example of a mediation model can be found in a study of Dixson et al. (2018), who showed that in adolescents the influence of socioeconomic status on academic achievements is mediated by hope. The authors found that socioeconomic status positively affects students' hope, which in turn positively affected their academic achievements.

The estimation of these techniques in commercially available statistical software is performed via regression-based techniques in the traditional frequentist framework. These methods, however, suffer from some disadvantages: inability to model individual heterogeneity and multilevel data, stringent assumption of Normality, poor performance on small samples, unintuitive interpretation of confidence intervals and p-values, as well as, common for frequentist statistical methods, a lack of information about the uncertainty in some of the estimates. Furthermore, violation of the assumptions made by these methods leads to invalid results of the analysis. Thus, a lack of knowledge about the pitfalls of the traditional techniques, and an absence of alternative methods brought us to where we are now: social sciences are descending into a reproducibility crisis.

However, an alternative does exist. The salvation comes from the flexibility of Bayesian methods, the application of which is helped by the wide availability of computational power. Bayesian Analysis of Variance or BANOVA is an open-source R-package (Dong and Wedel, 2017; Wedel and Dong, 2019), which offers a range of analyses with Bayesian hierarchical models. The package allows users to analyze data sets from studies with repeated measures or hierarchical data structures with a wide range of distributions for the dependent variable (such as Bernoulli, Poisson, Multinomial, etc). As BANOVA aims to be a one-stop-shop for the analysis of experiments in social sciences the package includes reports of effect sizes, Bayesian p-values, tables of predictions, and a wide range of follow-up techniques, including mediation analysis with a single mediator. The package is available from the Comprehensive R Archive Network (CRAN)[1].

BANOVA allows users who have a limited amount of experience with Bayesian statistics and R programming language to analyze a variety of models in experimental and observational studies. Ease of application and rigorous underlying statistical mechanisms make this library interesting for students, academics, and practitioners in the industry. A review of current functionality in statistical software such as SAS, SPSS, Stata, and JASP, as well as available R-packages, identified a lack of an available software solution for moderation and mediation (with multiple mediators) in the context of Bayesian hierarchical models. Since mediation and moderation analysis are prevalent in social and behavioral sciences (Yuan and MacKinnon, 2009), extending BANOVA with this functionality and improving users' experience with these techniques is of great practical value.

The main goal of this project is to enrich the functionality of the package with techniques for moderation analysis and mediation with multiple mediators. The first sub-project entails a development of moderation analysis in problems where one moderating factor influences an arbitrary number of other factors. For proper interpretation of the effects of the moderated factors, the calculation of simple effects of independent factors and their interactions at all levels of the moderator is often desired. The calculation of simple effects in a Bayesian framework has not yet been addressed in the literature. This sub-project was guided by the following research question:

> *RQ1. How to calculate and implement Bayesian simple effects in the framework of the BANOVA R-package?*

The second sub-project is aimed at developing a function for the calculation of an indirect effect of a causal variable transmitted through multiple, possibly correlated mediators, which are affecting the dependent variable at the same time. This problem appears not to have been previously addressed in the literature. Therefore, the second research question is:

> *RQ2. How to calculate and implement Bayesian mediation effects in problems with multiple correlated mediators in the framework of the BANOVA R-package?*

Furthermore, there are two secondary goals that I aim to achieve with this thesis. Firstly, the objective of this paper is not only to present the new techniques but also to introduce current and potential users of the package to Bayesian statistics and techniques implemented in BANOVA. I present a review of concepts in Bayesian inference to prevent some of the potential users from forgoing the analysis with the package due to their lack of familiarity with the relevant methods. The final goal of the project is to improve the quality of the statistical analysis performed in the package by

---

[1] `https://CRAN.R-project.org/package=BANOVA`

resolving major or minor aberrations identified during the code review and development of the new functions, by fixing the problems in question of extending the functions where necessary.

The remainder of this thesis is divided into six sections. Section 2 introduces the relevant concepts in Bayesian inference, starting from the basics and finishing with hierarchical regressions and hypothesis testing. Section 3 presents a comprehensive review of the methods implemented in the package and their connection to traditional frequentist techniques, followed by a description of model estimation of the results in the latest version of BANOVA. Section 4 introduces mediation and moderation analyses by first presenting theoretical and methodological considerations, along with a brief discussion of the available functionality of the package in these domains. A description of extensions implemented in BANOVA is presented in section 5 followed by an illustration of their application in section 6. The illustration is based on data sets from previously published studies of Perfecto et al. (2019), which is experimental research with repeated measures, and Huellemann and Calogero (2020), which is a cross-sectional observational survey. The concluding remarks and discussion are presented in section 7.

2 Bayesian Inference

The goal of statistical methods is to make inferences from numerical data about the quantities of interest via probability models. The distinction of Bayesian methods is their explicit use of probability statements to quantify the uncertainty in the inference. These methods allow one to include the knowledge about observable and unobservable quantities in the underlying problem and data collection in the full probability model. By conditioning of this model on the observed data, the knowledge about the unknown quantity of ultimate interest is updated and represented via a conditional probability distribution, termed as posterior distribution in the Bayesian framework. Treatment of the unknown quantity in probabilistic terms reflects either its stochastic nature or uncertainty in knowledge about a truly fixed parameter. This perspective facilitates a common-sense interpretation of inference. The lack thereof is the main antecedent of methodological failure, which is the central driver of the replication crisis in science (Guttinger and Love, 2019).

The section introduces Bayesian concepts and techniques which are necessary for an understanding of the underlying mechanism of the BANOVA R-package. The discussion starts with an introduction of basic concepts relevant for Bayesian data analysis, followed by a brief description of the computational inference. These two pillars allow me to introduce more complex topics, such as multilevel techniques and hypothesis testing approach.

2.1 Basic concepts in Bayesian inference

The observable information about the quantity of interest is obtained based on a set of n subjects or units and can be collected in a vector $\boldsymbol{y'} = (y_1, \ldots, y_n)$[2]. These outcomes are considered to be random due to the variation of population and the sampling process. The observed values are assumed to be exchangeable, meaning that the model

---

[2] Here and in all subsequent chapters of this thesis, all random variables are denoted as vectors, using bold italic lowercase letters.

outcomes are not changed by the reordering of the data points. An exchangeable distribution is commonly modeled as an independent and identically distributed (iid) given some unknown parameters of interest collected in a vector $\boldsymbol{\theta}$.

The knowledge and beliefs about $\boldsymbol{\theta}$ before the evidence is observed is summarized in the prior distribution $p(\boldsymbol{\theta})$. Priors can be specified based on the past research information or subjective beliefs of an expert, as well as built upon the consideration of the ease of computation (i.e. conjugate priors) or other theoretical principles. Gelman (2006) is a good starting point for an interested reader who wants to delve into the vast and still growing body of academic literature on different priors, their influence, and suitability in specific models. The degree of certainty about the prior distribution determines how informative the specification is. A common uninformative, or diffuse, prior can take the form of a Normal distribution with mean zero and large variance. Informative priors can improve inference by conveying additional information not directly available in the data. However, due to their strong influence informative priors must be selected very carefully and subjected to sensitivity analysis (see Jackson et al. (2015) for more details).

Statistical conclusions about $\boldsymbol{\theta}$ are made based on probability statements conditional on the observed values $\boldsymbol{y}$ and values of a set of k explanatory variables for n subjects collected in a matrix $\mathbf{X}$[3]. The resulting posterior distribution is denoted as $p(\boldsymbol{\theta}|\boldsymbol{y})$, with the conditioning on $\mathbf{X}$ being implicit. To model the posterior one must begin with a joint probability distribution for $\boldsymbol{y}$ and $\boldsymbol{\theta}$, denoted as $p(\boldsymbol{\theta}, \boldsymbol{y})$, which is a product of a prior distribution $p(\boldsymbol{\theta})$ and the sampling distribution $p(\boldsymbol{y}|\boldsymbol{\theta})$. Conditioning of the joint probability distribution on the observed values of data $\boldsymbol{y}$ using the Bayes' rule yields the posterior distribution

$$p(\boldsymbol{\theta}|\boldsymbol{y}) = \frac{p(\boldsymbol{\theta}, \boldsymbol{y})}{p(\boldsymbol{y})} = \frac{p(\boldsymbol{\theta})p(\boldsymbol{y}|\boldsymbol{\theta})}{p(\boldsymbol{y})}. \tag{2.1}$$

The likelihood function $p(\boldsymbol{y}|\boldsymbol{\theta})$, also denoted as $L(\boldsymbol{y}|\boldsymbol{\theta})$, is considered as a function of $\boldsymbol{\theta}$ for fixed $\boldsymbol{y}$ and contains all the relevant information from the data. The denominator in equation above $p(\boldsymbol{y})$ is commonly referred to as the evidence or normalizing factor. It is the marginal likelihood of the observed data marginalized over the parameters obtained as $\int_{\mathbb{R}^d} p(\boldsymbol{\theta})p(\boldsymbol{y}|\boldsymbol{\theta})d\boldsymbol{\theta}$ in case of continuous parameters in $\boldsymbol{\theta}$. An equivalent form of equation 2.1 is commonly obtained by omitting the $p(\boldsymbol{y})$, which fixed $\boldsymbol{y}$ does not depend on $\boldsymbol{\theta}$ and thus it is considered as a constant. The unnormalized posterior density is proportional to the $p(\boldsymbol{\theta}|\boldsymbol{y})$, which is formally expressed as

$$p(\boldsymbol{\theta}|\boldsymbol{y}) \propto p(\boldsymbol{\theta})p(\boldsymbol{y}|\boldsymbol{\theta}). \tag{2.2}$$

The posterior probability distribution $p(\boldsymbol{\theta}|\boldsymbol{y})$ is of a central interest in Bayesian inference. It includes all information about $\boldsymbol{\theta}$ and describes the uncertainty about the quantities of interest after the data is considered. The most informative way to represent this information is to plot the posterior distribution of the parameters. For practical purposes the results of the analysis are commonly reported in the form of numerical point and interval summaries of the posterior distribution. These quantities are typically obtained by computing an expected value of a certain function $g(\boldsymbol{\theta})$ with respect to the posterior distribution defined as

$$\mathrm{E}_{\boldsymbol{\theta}|\boldsymbol{y}}\big(g(\boldsymbol{\theta})\big) = \int_{\mathbb{R}^d} g(\boldsymbol{\theta})p(\boldsymbol{\theta}|\boldsymbol{y})d\boldsymbol{\theta}. \tag{2.3}$$

---

[3] Here and in all subsequent chapters of the thesis, matrices are denoted using bold uppercase letters.

To compute the posterior distribution and its summaries Bayesian framework relies heavily on integration techniques. When the closed-form expression for the posterior is known the computations do not present a significant challenge. However, high dimensional integrals, which correspond to more realistic probability models, commonly do not have an analytic solution. When the exact computations are practically infeasible approximation and simulation techniques can be used to obtain the posterior distribution. A discussion of basic computational techniques such as rejection and importance sampling are beyond the scope of this thesis, and these topics are covered to a various degree in most of the popular textbooks on Bayesian statistics, such as Gelman et al. (2013) or Gill (2014).

2.2 Bayesian inference engine

A challenge of estimating complex multidimensional distributions was one of the main burdens encountered by the practitioners of Bayesian statistics. But everything changed when Markov chain Monte Carlo (MCMC) algorithms have spilled from the domain of physics (Metropolis and Ulam, 1949) into the mainstream statistics. The idea behind Markov chains (MCs) allows to construct and sample from an arbitrary target distribution, by recursively drawing values of $\boldsymbol{\theta}$ from an approximate distributions and correcting those values to better approximate the target posterior distribution. The definition of a MC is a sequence of samples $\boldsymbol{\theta}^{(1)}, \boldsymbol{\theta}^{(2)}, \ldots, \boldsymbol{\theta}^{(M)}$, whose distribution converges to a unique stationary target distribution $p(\boldsymbol{\theta}|\boldsymbol{y})$ and follows a Markov property. This property implies that distribution of the sampled draws depends only on the last draw from a MC, such that $p(\boldsymbol{\theta}^{(m)}|\boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(m-1)}) = p(\boldsymbol{\theta}^{(m)}|\boldsymbol{\theta}^{(m-1)})$. To specify a MC one must define a probability distribution of the initial value $p(\boldsymbol{\theta}^{(0)})$ and transition distributions for the subsequent draws in a form of $T(\boldsymbol{\theta}^{(m)}|\boldsymbol{\theta}^{(m-1)})$. A further introduction to MCMC sampling techniques can be found in Appendix A.

All MCMC algorithms have an iterative nature, such that at each step of the simulation the draw from a distribution becomes closer to the target density $p(\boldsymbol{\theta}|\boldsymbol{y})$ as defined in equations 2.1 and 2.2. It is important to let the simulation run long enough such that it draws values close enough to the specified posterior distribution. The iterations made early in the chain, referred to as a warm-up period, are commonly discarded, as they represent pre-convergent values. To check if the number of iterations is sufficiently large it is necessary to assess the convergence and mixing of the simulated draws after the warm-up. By definition, an MC generates samples from the posterior distribution only after that chain has mixed throughout the support of the density and converged to the target distribution. A wide palette of visual and quantitative convergence assessment techniques is available in practice. For example, see Cowles and Carlin (1996) and Brooks and Gelman (1998). Furthermore, it is common to simulate multiple sequences with dispersed starting points and to compare the results between and within simulated sequences.

2.3 Bayesian hierarchical regression techniques

The flexibility of the stochastic numerical integration techniques allows researchers to fit a wide range of complex statistical models. Among them is a class of hierarchical models, which explicitly model the relationships between multiple parameters in a given problem. These connections arise naturally when the data exhibit hierarchical or nested structures, which means that units of analyses at a lower level are nested within aggregated units at a higher level. Application of non-hierarchical models to nested data results in underfitting or overfitting to the data, depending on the number of involved parameters (Gelman et al., 2013). The hierarchical models allow to include a sufficient number of parameters to assure a good fit while accounting for their dependence structure through a population (or prior) distribution and avoid the issue of overfitting. It is important to note, that when hierarchical data is considered the assumption of exchangeability of units of analysis, mentioned in the section 2.1, breaks down at the lower level as they are not independent. The hierarchical regression model thus assumes the units to be conditionally exchangeable. This implies that the units are exchangeable after conditioning them on a categorical variable which indicates to which group the observation belongs.

In the hierarchical setting, if lower-level units belong to the aggregated level j, for example, observations for a single subject in the experiments with repeated measures design, it is expected that all these lower-level observations share a common $\boldsymbol{\theta}_j$. In turn parameters $\boldsymbol{\theta}_j$, for $j = 1, \ldots, J$, where J is the number of higher-level units, are assumed to be sampled from a common population distribution. This is modeled by conditioning the observable outcomes $\boldsymbol{y}$ on certain parameters, which themselves have a probabilistic specification. More specifically each parameter $\boldsymbol{\theta}_j$ is treated as an independent sample from a prior distribution with a hyperparameter vector $\boldsymbol{\phi}$, which has a prior distribution $p(\boldsymbol{\phi})$. The specification of the distribution for parameters $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_J$ requires the assumption of exchangeability, which implies the joint distribution of the $\boldsymbol{\theta}$s is invariant to the permutation of the indexes. The resulting joint posterior distribution is defined as follows:

$$p(\boldsymbol{\theta}, \boldsymbol{\phi}|\boldsymbol{y}) \propto p(\boldsymbol{y}|\boldsymbol{\theta})p(\boldsymbol{\theta}|\boldsymbol{\phi})p(\boldsymbol{\phi}). \tag{2.4}$$

For linear regression models, it is possible to add a hierarchical prior, which gives rise to a distinction between fixed and random effects, also referred to as population-level and group-level effects. The coefficients of random effects vary between the groups according to some statistical distribution, while fixed effects correspond to the constant regression coefficients in the whole population. These mixed-effect models allow us to find the best fit for the individual observations while accounting for the systematic unexplained variation between aggregated groups. The hierarchical or multilevel linear models are referred to as models with partial pooling, which is a compromise between complete pooling and no pooling. Methods with complete pooling completely ignore the differences between groups, and thus ignore the between-group variation, while methods with no pooling analyze the groups separately, and thereby increases the inference uncertainty due to the inefficient data use.

These models are particularly useful in the context of the linear regressions when the predictors are measured at different levels of variation. As an example, consider a case where a certain variable $\boldsymbol{y}$ is observed repeatedly for each subject j for n_j times. Suppose there is one explanatory variable $x_{i,j}$ which varies between the measurements for each subject, and one variable z_j which is constant for each individual subject. Under an assumption of normality the model can be expressed as

$$y_{ij} \sim \mathrm{N}\left(\alpha_j + \beta x_{i,j}, \sigma_y^2\right), \quad \text{for } i = 1, \ldots, n_j, j = 1, \ldots, J, \tag{2.5}$$

$$\alpha_j \sim \mathrm{N}\left(\theta_0 + \theta_1 z_j, \sigma_\alpha^2\right), \quad \text{for } j = 1, \ldots, J. \tag{2.6}$$

The error variance σ_y^2 in equation 2.5 and σ_α^2 in 2.6 represent the within-subject and between-subjects variations respectively, which are beyond what is explained by the covariates. The model above can be further extended with additional predictors at the individual and aggregated levels, and by allowing the slopes $\boldsymbol{\beta}$ vary by subject. To fit the simple model in 2.5-2.6 with Normal bi-variate prior for $\boldsymbol{\theta}$ with a mean μ_θ and a covariance matrix Σ_θ, and a uniform prior distributions on σ_y and σ_α, the following posterior density must be estimated:

$$\begin{aligned} &p\left(\boldsymbol{\alpha}, \beta, \boldsymbol{\theta}, \sigma_y, \sigma_\alpha | \boldsymbol{y}, \boldsymbol{x}, \boldsymbol{z}\right) \\ &\propto \textstyle\prod_{j=1}^{J} \prod_{i=1}^{n_j} \mathrm{N}\left(y_{ij} | \alpha_j + \beta x_{i,j}, \sigma_y^2\right) \prod_{j=1}^{J} \mathrm{N}\left(\alpha_j | \theta_0 + \theta_1 z_j, \sigma_\alpha^2\right) \mathrm{N}\left(\boldsymbol{\theta} | \mu_\theta, \Sigma_\theta\right) p(\sigma_y) p(\sigma_\alpha). \end{aligned} \tag{2.7}$$

While in this example the analytic derivation of posterior distribution is still possible it is not as convenient as a MCMC simulation technique. The joint posterior distribution contains all the information about all the unknowns in the model. However, the aim of the analysis is to draw inference about the particular parameters of interest. To obtain these marginal posterior distributions one must integrate the joint posterior density over all of the nuisance parameters that are not of immediate interest. In case a simulation technique is used to draw samples from the joint posterior distribution the marginal distribution can be obtained by considering only the samples of the parameters of interest while ignoring all the draws of nuisance parameter. The samples from the marginal posterior distribution can be used to compute posterior results, which commonly include mean, standard deviation, and quantiles of interest. In general based on M draws of a given parameter, for example β, any function $g(\beta)$ presented in equation 2.3 can be computed from the marginal posterior samples as follows

$$\mathrm{E}_{\beta|\cdot,\boldsymbol{y}}\big(g(\beta)\big) \approx \frac{\sum_m^M g\left(\beta^{(m)}\right)}{M}. \tag{2.8}$$

2.4 Bayesian hypothesis testing

The joint and marginal posterior distributions establish the basis for hypothesis testing in the Bayesian framework (Jeffreys, 1998). In the frequentist approach, where parameters are assumed to be fixed in the population, it is necessary to evoke the assumption of repeated probability samples and perform inference based on the p-value null hypothesis significance testing (pNHST). In contrast to only being able to reject the null hypothesis with pNHST, Bayes factors (BF) allow researchers to quantify support for any hypothesis H_1 relative to hypothesis H_0. These hypotheses need not be nested, and are commonly used as an approach for model selection. The comparison starts by evaluating the apriori plausibility of the hypothesis in terms of the degrees of belief of the researcher, expressed as $p(\cdot)$ through the prior odds, that is $p(H_1)/p(H_0)$. The Bayesian hypothesis testing is sensitive to the prior specification, and thus the choice of the prior distribution requires an understanding of the underlying hypothesis and careful consideration. After the data $\mathbf{D}$ are observed, the posterior odds can be computed by multiplying the prior odds and BF, where BF is computed in the following manner:

$$\mathrm{BF}_{1|0} = \frac{p(\mathbf{D}|H_1)}{p(\mathbf{D}|H_0)}, \text{ where} \tag{2.9}$$

$$p(\mathbf{D}|H_i) = \int_{R^d} p(\mathbf{D}|\boldsymbol{\theta}_i, H_i)p(\boldsymbol{\theta}_i)d\boldsymbol{\theta}_i. \tag{2.10}$$

Thus, BF evaluates the evidence by comparing the marginal distribution of the observed data under the alternative hypothesis $p(\mathbf{D}|H_1)$ versus the probability of the observed data under the null hypothesis $p(\mathbf{D}|H_0)$. The competing hypothesis correspond unique parameter vectors $\boldsymbol{\theta}_i$, with respective proper prior densities $p(\boldsymbol{\theta}_i)$ for $i = 1, 2$. The BF is a working engine of the Bayesian hypothesis testing, which can be interpreted as the amount by which the data shifts researchers' beliefs in favor of one of the hypothesis. It is possible to summarize the BF in terms of evidential strength with discrete categories ranging from no evidence for any of the hypothesis ($BF_{1|0} = 1$) to extreme evidence in favor of H_0 ($BF_{1|0} < 1/100$) or H_1 ($BF_{1|0} > 100$) (Kass and Raftery, 1995).

A special case occurs when hypothesis about a given regression parameter β is tested. Two competing hypothesis $H_1 : \beta > 0$ and $H_2 : \beta < 0$ are apriori equally likely under the Normal prior centered around zero and with a constant variance σ_β. This approach is advantageous compared to a traditional approach of comparing a directional hypothesis against a null $H_0 : \beta = 0$, which has been criticized for it's mismatch with the nature of reality, as, at least in observational studies, it is never true exactly (Cohen, 1994). In this special case Marsman and Wagenmakers (2017) have showed that BF can be directly converted to the posterior probability of of the alternative hypothesis

$$\mathrm{BF}_{12} = \frac{p(\mathbf{D}|H_1)}{p(\mathbf{D}|H_2)} = \frac{\int_0^\infty p(\beta|\boldsymbol{y})d\beta}{\int_{-\infty}^0 p(\beta|\boldsymbol{y})d\beta} = \frac{p(H_1|\mathbf{D})}{1 - p(H_1|\mathbf{D})} = \frac{1 - p(H_2|\mathbf{D})}{p(H_2|\mathbf{D})}, \tag{2.11}$$

$$p(H_1|\mathbf{D}) = \frac{BF_{1|2}}{1 + BF_{1|2}}, \quad p(H_2|\mathbf{D}) = \frac{1}{1 + BF_{1|2}}. \tag{2.12}$$

In the the special case describe above the $p(H_i|\mathbf{D})$ is a so-called one-sided Bayesian p-values, which can be calculated as

$$p(H_1|\mathbf{D}) = \int I(\beta^{(m)} > 0)p(\beta|\boldsymbol{y})d\beta \approx \frac{\sum_m^M I\left(\beta^{(m)} > 0\right)}{M}, \tag{2.13}$$

$$p(H_2|\mathbf{D}) = \int I(\beta^{(m)} < 0)p(\beta|\boldsymbol{y})d\beta \approx \frac{\sum_m^M I\left(\beta^{(m)} < 0\right)}{M}. \tag{2.14}$$

These p-values in the case described above can be used as a measure of the strength of evidence for the null or alternative hypothesis. In contrast, the p-values obtained through frequentist methods only allow to reject the null hypothesis, without quantifying the evidence in any interpretable manner. To obtain the two-sided Bayesian p-value one must obtain $p(H_1|\mathbf{D})$ and $p(H_2|\mathbf{D})$, and multiplying the smaller quantity by two.

3 BANOVA R-Package for Hierarchical Bayesian ANOVA

Dong and Wedel (2017) have first introduced the open-source R-package BANOVA as an alternative to the widely popular Analysis of variance (ANOVA) first presented by Fisher (1921). In BANOVA, the variation in the dependent variable can be explained at two measurement levels: at the within-subjects level with predictors that vary within participants, and at the between-subject level, where the explanatory variables are the same within one observational unit but vary between participants. This approach allows to account for systematic unexplained differences between the unique subjects or participants. The package capitalizes on the fact that ANOVA is a special case of the linear regression model and, in combination with Bayesian techniques, creates a more flexible substitute to the traditional technique used for the analysis of experimental data in social and behavioral sciences. This section briefly introduces the ANOVA in the frequentist framework, followed by the introduction of the Bayesian ANOVA paradigm and its advantages. Next, I present to the reader the details about the model specification and estimations and a brief discussion of the BANOVA results.

3.1 Traditional ANOVA and its assumptions

The classical approach refers to an additive decomposition of the dependent variable, which is equivalent to a linear regression model on categorical, also known as factors or index variables, and their interactions. This method allows users to test hypotheses about the differences between multiple means while keeping constant the Type I error rate, which is the probability of falsely rejecting a correct null hypothesis, which is inevitable if multiple pair-wise comparisons with t-tests are performed (Gelman and Hill, 2007).

In its simplest form, ANOVA partitions the data according to a treatment factor in k groups and compares the variation in the dependent variable between groups to the variation within groups. The null hypothesis of the test is that there are no differences between group means, which implies that the differences between groups are small relative to the variation within groups. To test this hypothesis two measures of variation are computed: sums-of-squares-between (SSB) and sums-of-squares-within (SSW). When these quantities are divided by the degrees of freedom, which are $k - 1$ and $n - k$ respectively, group mean square and means square error are computed. The fraction of these measures of variation defines the F-statistic, which follows $F(k - 1, n - k)$ distribution (Gelman and Hill, 2007). However, this result requires several assumptions. The observations of the dependent variable are assumed to follow independent (i.i.d) homoskedastic Normal distribution and the independent variables are assumed to be factors. The latter assumption is relaxed in the Analysis of covariance (ANCOVA), which allows for continuous predictors to be included in the model (Fisher, 1934). ANCOVA requires additional assumptions of independence between the covariates and the treatment factors, a linear relationship between the covariates and the dependent variable, as well as the same regression coefficients on the covariates in each treatment group.

These assumptions are restrictive and are commonly violated in academic experiments. The independence assumption is met only if one observation is obtained from each unique participant, who was randomly sampled from a distinct population or assigned to a treatment group. This is often unrecognized by researchers and can lead to formidable inflation of the Type I error rate (Scheffé, 1959). The non-independence due to group, sequence, and space commonly arises in experiments with stratified or clustered sampling techniques. Another common cause of non-independence is repeated measures designs that correspond to nested data structures, where predictors are observed at different levels of variation. A moderate violation in the assumption of homoskedasticity, or constant error variance across observations, does not lead to the substantial effect on the Type I error rate if the sizes of the compared groups are equal (i.e. balanced) (Scheffé, 1959). Furthermore, it has been shown, that moderate deviations in the dependent variable from the Normal distribution also do not lead to decreased validity of the F-test and increased Type I error rate, given a sufficiently large sample size (Pearson, 1931; Donaldson, 1968). However, a wide-spread application of traditional ANOVA and ANCOVA to categorical dependent variables leads to spurious results, which lack any meaningful interpretation (Jaeger, 2008).

3.2 Bayesian ANOVA framework

The design of the R-package for hierarchical Bayesian ANOVA relaxes some of the assumptions described above. It allows the techniques implemented in the package to reflect the nature of experimental data in social and behavioral sciences (Dong and Wedel, 2017). The hierarchical Bayesian models are at the core of BANOVA, as they arise naturally in the repeated measures designs with mixed within and between-subjects factors and covariate (Gelman et al., 2013). An introduction to the theory behind the hierarchical models is discussed in section 2.3, and a formal definition is presented in 3.3. To account for the various sources of variation in the model two levels of regressions are distinguished in BANOVA. At the first level, a generalized linear model (GLM) is used to model the variation within participants, while at the second level, the population-level differences between subjects are assessed via ANOVA techniques.

The generalization of the linear model was introduced by McCullagh and Nelder (1989) as a technique for inference with dependent variables from the exponential family. These models consist of three different parts: random and systematic components, as well as a link between them. The random component of a GLM model for a random variable $\boldsymbol{y}$ has a distribution with an expected value of $\boldsymbol{\mu}$. Optional addition distribution-specific parameters, which do not vary across the subjects in the data, are collected in a vector $\boldsymbol{u}$. For instance, it would include a variance in the case of a Normal model. In the systematic component intercept, factors and covariates in $\mathbf{X}$ contribute to the linear predictor $\boldsymbol{\eta}$ via an unobserved parameter vector $\boldsymbol{\theta}$. The final part is a canonical link function $g(\cdot)$ which connects the two parts mentioned before, such that $\boldsymbol{\eta} = g(\boldsymbol{\mu})$. This allows for the additive systematic effects to occur in a monotonic transformation of the mean, such that $\boldsymbol{\mu} = g^{-1}(\boldsymbol{\eta})$. Thus, BANOVA links the dependent variably $\boldsymbol{y}$ with an assumed distribution $f(\cdot)$ to a linear predictor $\boldsymbol{\eta}$ through an inverse function $g^{-1}(\cdot)$, which can be expressed as follows:

$$\boldsymbol{y} \sim f\left(g^{-1}(\boldsymbol{\eta}), \boldsymbol{u}\right), \text{ where } \boldsymbol{\eta} = \mathbf{X}\boldsymbol{\theta}. \tag{3.1}$$

The flexibility of the GLM approach allows BANOVA to model dependent variables with a wide range of distributions $f(\cdot)$. Continuous outcome variables can follow a Normal or a t-distribution. Binary responses are accommodated with Bernoulli and Binomial probability mass functions. Poisson distribution is available for modeling of the dependent variables which measure counts. Finally ordered categorical and multinomial responses can be analyzed with ordered and multinomial logistic regression models.

The hierarchical model used in BANOVA is also referred to as the slopes-as-outcomes model (Burstein, 1980), as the second-level model uses batches of estimated subject-specific coefficients $\hat{\boldsymbol{\theta}}$ as dependent variables or outcomes. Thus, the model follows a pattern presented in equations 2.5 and 2.6, where both intercepts and slopes are regressed on an intercept and between-subjects variables collected in a matrix $\mathbf{Z}$. Each batch is modeled as independent while following a Normal distribution, and thus it is an additional source of variation. If all subscripts are omitted for convenience, the model can be expressed as

$$\boldsymbol{\theta} = \mathbf{Z}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \text{ where } \boldsymbol{\varepsilon} \sim N\left(\mathbf{0}, \sigma_{\theta}^{2}\right), \tag{3.2}$$

which implies that a linear predictor $\boldsymbol{\eta}$ in a multilevel model is defined as $\boldsymbol{\eta} = \mathbf{X}\mathbf{Z}\boldsymbol{\beta} + \mathbf{X}\boldsymbol{\varepsilon}$.

The hierarchical nature of the model implies that within-subjects regressors are assumed to have both fixed and random components, while the effects of between-subjects explanatory variables are fixed. Furthermore, the second-level regressors influence the magnitude of (i.e. moderate) the effects of the within-subjects predictors. In this hierarchical approach, the population-level coefficients of are usually of the key interest in ANOVA. Note, that if a single-level model with a Normally distributed dependent variable is being fitted in BANOVA the outcome corresponds to traditional ANOVA.

3.3 Bayesian Hierarchical Models in BANOVA

In BANOVA Bayesian generalized hierarchical models link the dependent variable to two submodels (Dong and Wedel, 2017). At the first level a subject-level model analyzes effects of within-subject regressors, which can be continuous or factor variables. Each observation of a dependent variable y_i, for $i = 1, \ldots, N$, is assumed to follow a distribution in an exponential family, $f(y_i|\mu_i)$. The influence of the subject-level regressors on the mean μ_i is conveyed through a suitable link function as described in a section 3.2. This model is expressed as follows:

$$\mathrm{E}\left(y_i\right) = \mu_i = g^{-1}\left(\eta_i\right), \quad i = 1, \ldots, N \tag{3.3}$$

$$\eta_i = \sum_{p=0}^{P} \sum_{j=1}^{J_p} x_{i,j}^{p} \theta_{j,s}^{p} = \boldsymbol{x}_i^T \boldsymbol{\theta}_s \tag{3.4}$$

Here p indexes the within-subject variables and their interactions, where $p = 0$ indicates an intercept, and $p = 1, \ldots, P$ indexes all other explanatory variables. If a variable p is a factor there exist J_p regressors and coefficients which represent the influence of this factor. When a given variable p is continuous $J_p = 1$. Hence, $x_{i,j}^{p}$ represents an observed value of a regressor j which corresponds to a variable p for an individual observation i. These individual observations belongs to a subject s, which is reflected in the subscript of the coefficient $\theta_{j,s}^{p}$. This coefficient represents an effect of a regressor j of variable p. This means that the marginal effects of a change in regressors are subject-specific, or in other words vary across each subject s for $s = 1, \ldots, S$. All observed values of regressors for an observation i are collected in a vector $\boldsymbol{x}_i$, which is a row in a N by J matrix $\mathbf{X}$, with corresponding coefficients stored in a vector $\boldsymbol{\theta}_s$. Both vectors are of a length $\sum_{p=0}^{P} J_p$, which will be labeled as J for convenience.

A population-level model expresses the influence of between-subject regressors, which vary between subjects. This model allows one to account for an unobserved heterogeneity among subjects. Based on the assumption that subject-level coefficients follow a multivariate Normal distribution the following model is formulated for each $s = 1, \ldots, S$, $p = 1, \ldots, P$, and $j = 1, \ldots, J_p$:

$$\theta^p_{j,s} = \sum_{q=0}^{Q} \sum_{k=1}^{K_q} z^q_{s,k} \beta^p_{j,k} + \varepsilon^p_{j,s} = \boldsymbol{z}_s^T \boldsymbol{\beta}^p_j + \varepsilon^p_{j,s} \tag{3.5}$$

In equation above $\theta^p_{j,s}$ depends on a vector of subject-specific regressors $\boldsymbol{z}_s$ of length $\sum_{q=0}^{Q} K_q$, which is further denoted as K for an ease of notation. Similarly to the first level model in equation 3.4 the notation implies that Q between-subject variables, including an intercept indexed by $q = 0$, are included in the estimation. Each variable is represented by a batch of K_q regressors, with $K_q = 1$ for continuous regressors. Hyper-parameters which capture the effects of between-subject regressors on $\theta^p_{j,s}$ are collected in a vector $\boldsymbol{\beta}^p_j$ of length K. If all components of $\boldsymbol{\theta}_s$ are stacked, the model can be expressed in the following matrix notation:

$$\begin{bmatrix} \theta^1_{1,s} \\ \theta^1_{2,s} \\ \vdots \\ \theta^P_{J_P,s} \end{bmatrix} = \begin{bmatrix} \boldsymbol{z}_s^T & \mathbf{0} & \ldots & \mathbf{0} \\ \mathbf{0} & \boldsymbol{z}_s^T & \ldots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \ldots & \boldsymbol{z}_s^T \end{bmatrix} \begin{bmatrix} \boldsymbol{\beta}^1_1 \\ \boldsymbol{\beta}^1_2 \\ \vdots \\ \boldsymbol{\beta}^{J_P}_P \end{bmatrix} + \begin{bmatrix} \varepsilon^1_{1,s} \\ \varepsilon^1_{2,s} \\ \vdots \\ \varepsilon^P_{J_P,s} \end{bmatrix}$$

$$\boldsymbol{\theta_s} = \mathbf{W}_s \boldsymbol{\beta} + \boldsymbol{\varepsilon}_s, \quad s = 1, \ldots, S \tag{3.6}$$

$$\boldsymbol{\varepsilon}_s \sim \mathrm{N}\left(\mathbf{0}, \sigma^2_\theta \mathbf{I}\right) \tag{3.7}$$

$$\boldsymbol{\beta} \sim \mathrm{N}\left(\mathbf{0}, c\mathbf{I}\right) \tag{3.8}$$

$$\sigma^2_\theta \sim \mathrm{InvGamma}\left(a, b\right) \tag{3.9}$$

Here $\boldsymbol{\theta}_s$ is a vector of length J with regression coefficients for a subject s. The matrix $\mathbf{W}_s$ is a J by KJ matrix which stores regressors $\boldsymbol{z}_s$ for the respective subject. Note, that the vector $\boldsymbol{z}_s$ is a row from a matrix $\mathbf{Z}$, which is a S by K matrix of between-subject covariates for all subjects. A vector of KJ length $\boldsymbol{\beta}$ with population-level coefficients follows a multivariate Normal distribution with zero mean and diagonal covariance matrix with a hyper-parameter c as diagonal elements. The disturbances are collected in a vector $\boldsymbol{\varepsilon}_s$ of length J, which follow a multivariate Normal distribution with zero mean and a diagonal covariance matrix with σ^2_θ as diagonal elements. In turn σ^2_θ follows an Inverted Gamma distribution with hyper-parameters a and b.

3.4 Model specification and estimation in BANOVA

To fit a Bayesian hierarchical regression with BANOVA it is necessary to first load a model with a desired distribution via the `BANOVA.model()` function and compile it with `BANOVA.build()`. The compiled model is further passed as an argument in the `BANOVA.run()` function, which executes the estimation and further analysis. This function takes as an input specification of a formula for a subject-level model, and, in the case of a multilevel model, a formula for the population-level equation. The function also requires a data frame with a data set that includes a column with

IDs of individual units in the analysis. Once the `BANOVA.run()` is called continuous variables included in the model are mean-centered by default. Factor variables with k levels, where $k \geq 2$, are automatically encoded with $k - 1$ effect coded regressors. This scheme, which is also referred to as sum coding, is similar to dummy coding, except the reference level is encoded as (-1) such that the sum of the effects is constrained to zero. When $g^{-1}(\cdot)$ in equation 3.3 is an identity function this coding scheme compares the mean of the outcome variable at a given level of a factor to the overall mean of the dependent variable over all levels of the investigated factor (Cohen et al., 2013). It is also possible to encode the factors with custom contrasts, which is a powerful and appropriate tool to test specific hypotheses about differences in the group means between specific conditions or their combinations.

Currently, in BANOVA uninformative priors for parameters are assigned automatically. For all single-level models, regression parameters are restricted to have Normal priors with a mean value of zero mean and variance of ten thousand. In the current implementation, the parameters are assumed to be uncorrelated. Additionally models with Normal and t-distributed error terms assign an Inverted Gamma (1, 1) prior for the error variances. To model dependent variables that follow a t-distribution a Normal (5, 25) prior for the degrees of freedom is assigned. Furthermore, in case an ordered Multinomial model is estimated the cut-off points are assigned a Normal (5, 100) prior. In multilevel models, the population-level parameters are assumed to have a Normal (0, 10000) prior. The subject-level parameters are generated from a Normal distribution with the mean determined by the subject-specific mean and a variance, which has an Inverted Gamma (1, 1) prior. All the other distribution-specific parameters follow the same priors as in single-level models.

Model fitting in BANOVA is performed via Stan, which is a state-of-the-art platform for statistical modeling and computation with MCMC sampling (Carpenter et al., 2017). It is possible to use a Hamiltonian Monte Carlo (HMC) sampler and its extension the no-U-turn sampler (NUTS) (Hoffman and Gelman, 2014) for parameter estimation. There are several advantages of the algorithms in the HMC family. First of all, they converge quickly in high-dimensional models with conjugate and non-conjugate priors; furthermore, the algorithms can produce less auto-correlated samples, compared to the Metropolis-Hastings algorithm (Hoffman and Gelman, 2014). The drawback of these algorithms is their time-consuming nature, which arises due to the necessity to calculate the gradient of the log-posterior. In BANOVA the NUTS sampler is called by default, as it eliminates the need to tune any parameters which, as described in section A.2, are required for an efficient HMC algorithm.

Several practical considerations are important in getting reliable inference from MCMC algorithms. The starting values of unconstrained parameters in the chains are drawn from a uniform distribution on the interval $(-2, 2)$. For constrained parameters, the initialization bounds are transformed through an appropriate transformation function. The randomly distributed starting points reflect the notion that little is known about the parameters before the estimation process and the starting values need to be reasonably dispersed. The second important choice in the algorithm design is the number of iterations in the chains and the length of the warm-up period. The default number of iterations in BANOVA is set to 2000, the first half of which is discarded. This number of iterations is sufficient for fitting simple models, but a follow-up convergence assessment is always required to check the reliability of the results. Finally, a thinning value of k can be selected, which specifies to record every kth value of the chain. Thinning is routinely used in practice, as it is believed to improve the quality of estimates by reducing the autocorrelation and boosting convergence (Link and Eaton, 2012). However, it has been shown that the estimation quality can decrease with thinning, and thus a thin value other than 1, which is set by default in BANOVA, should only be used to deal with limitations in computer storage capabilities (Link and Eaton, 2012; Gill, 2014).

All of the considerations described above can be modified by the user through the specification of the appropriate arguments in the BANOVA function calls. For further details about the values of arguments relevant in the program, execution see the Stan reference manual (Stan Development Team, 2020). It is also worth noting that to ensure that the results are reproducible a seed state can be passed as arguments in the BANOVA.

3.5 Results in BANOVA

All quantities of interests and test results are constructed based on samples from the posterior distributions obtained from a defined Stan model. The results of BANOVA include tables with convergence diagnostics, sums-of-squares, effect sizes, and estimated coefficients, with corresponding Bayesian p-values and credible intervals. BANOVA also reports model fit and tables of predicted means for the groups based on the factors in the data.

To assess the convergence of the estimated parameters Geweke diagnostic (Geweke et al., 1991) and Heidelberg and Welch diagnostic (Heidelberger and Welch, 1983) are reported in BANOVA. These empirical techniques are applied to samples of individual parameters after warm-up and test whether given posterior samples are from a chain in a stationary distribution. Geweke diagnostic uses a standard Z-score to test the difference in means of the first 20% and last 50% of the draws for a given parameter, under a null hypothesis of equal means in both groups. Heidelberg and Welch diagnostic uses a Cramer-von Mises test statistic to test a null hypothesis that a given MC is from a stationary distribution. First, the procedure calculates a test statistic for all draws from the posterior distribution, and if the null is rejected the test is repeated on the last 90% of the samples for a given parameter. This approach is repeated in an iterative manner, such that each time the null is rejected the first 10% of the samples are discarded. The iterations proceed until it is not possible to reject the null or 50% of the draws are discarded. In the latter case, the test indicates that the chain has failed to converge. If the first part of the diagnostic did not reject the null hypothesis of stationarity, the non-discarded samples of the chain are considered further in a halfwidth analysis. This test finds half of a 95% credible interval for the considered samples and compares it with their mean. Only if the ratio of halfwidth to mean is lower than a certain parameter ϵ, which takes a value of 0.1 in BANOVA, the test is passed. If any of the tests are rejected the number of iterations in the MC must be increased to obtain sufficient accuracy.

To allow for balanced and unbalanced research designs, BANOVA calculates Bayesian Type III sums-of-squares (SS) (Dong and Wedel, 2017) based on the work of Yates (1934). These SS are based on the conditioning of the main effects on all other terms in the model, and thus the tests based on them are valid whether or not interactions of variables are present in the regression (Fox, 1997). To compute the SS of any factor, covariate or an interaction BANOVA selects columns corresponding to a given variable of interest form matrix $\mathbf{X}$ for a single-level model and form $\mathbf{Z}$ for a multilevel-level model. The selected columns are multiplied by relevant estimated coefficients, to obtain a prediction for the dependent variable $\boldsymbol{y}$. The predicted values based on a single variable are demeaned and used for computation of a prediction variance. To obtain the SS of the considered variable this calculation is repeated for M sample draws of the MC and then averaged. For example, if in a single level model we are interested in SS for a factor with two levels or a continuous variable, denoted as vector $\boldsymbol{x}_j$, which corresponds to a parameter θ_j, the formula for the computation is defined as

$$SS_{\theta_j} = \frac{1}{M}\sum_{m=1}^{M} \operatorname{Var}\left(\boldsymbol{x}_j \hat{\theta_j}^{(m)} - \mathbf{1}\frac{1}{N}\sum_{i=1}^{N} x_{ji}\hat{\theta_j}^{(m)}\right), \tag{3.10}$$

where $\hat{\theta_j}^{(m)}$ is the m^{th} draw of parameter θ_j, $\boldsymbol{x}_j\hat{\theta_j}^{(m)}$ are the predicted values of $\boldsymbol{y}$, $\mathbf{1}$ is a vector of ones with the same size as $\boldsymbol{y}$, and $\frac{1}{N}\sum_{i=1}^{N} x_{ji}\hat{\theta_j}^{(m)}$ is the mean of the predicted values. The calculation of the residual SS is different for single-level and multilevel models. For single-level models, where the dependent variable follows a Normal or a t-distribution the residual SS are calculated by taking the difference between the actual values of $\boldsymbol{y}$ and the predicted values $\hat{\boldsymbol{y}}$ based on the full model. For models with Poisson, Bernoulli, and Binomial dependent variables they are based on the distributions' deviance's proposed by Nelder and Wedderburn (1972). Finally, for ordered categorical and multinomial responses, the residual SS are assigned a value of $\frac{\pi^2}{6}$, which is a variance of the individual Gumbel distributed error terms with a mean value of one (Ben-Aakiva and Lerman, 1985). Residual SS in a multilevel case are calculated based on the model presented equations 3.1-3.2, by taking the difference between the values of $\boldsymbol{\theta}$ estimated in the first level and the predictions made at the second level.

A degree of association between an effect of a given variable and the dependent variable is captured in the eta-squared (η^2) measure of the effect size. The standardized nature of the metric allows one to assess the magnitude of the reported effects regardless of the measurement scale of the dependent variable. The practical advantages of the standardized effect sizes in meta-analytic studies, as well as their value in the cumulative scientific process, make them one of the most important outcomes of empirical studies (Lakens, 2013). In BANOVA generalized partial eta-squared effect sizes (η_p^2) are calculated for each of the investigated factors and covariates based on the MC draws (Dong and Wedel, 2017). The stochastic nature of the estimation allows the metric to account for the uncertainty in the parameter estimates, and illustrate the degree of uncertainty about the effect sizes in a form of credibility intervals. For a single-level model, the metric is obtained by dividing the SS for a given variable, defined in equation 3.10, by the sum of the corresponding residual SS and the SS of the given variable. For example, for a parameter θ_j, the partial eta-squared is

$$\eta^2_{p\theta_j} = \frac{SS_{\theta_j}}{SS_{\theta_j} + SS_{\varepsilon_{\theta_j}}}. \tag{3.11}$$

For the effect size calculation in a multilevel model, a distribution-specific correction term s_y for the error variance of the within-subjects model (Nakagawa and Schielzeth, 2013) is included in the denominator. For instance, if an effect size of a population-level effect β_j is of interest, the effect size is computed as follows

$$\eta^2_{p\beta_j} = \frac{SS_{\beta_j}}{SS_{\beta_j} + SS_{\varepsilon_{\beta_j}} + s_y}. \tag{3.12}$$

The correction term s_y for models with Normal or a t-distributed dependent variables is an estimated standard deviation of the $\boldsymbol{y}$. When the dependent variable follows a Poisson distribution the s_y takes the value of zero. If a dependent variable is modeled via a Bernoulli or Binomial distribution s_y is zero and can be neglected. For models with ordered categorical and multinomial responses s_y takes a value of $\frac{\pi^2}{6}$.

To provide insights about the estimates of the parameters in the hierarchical models the draws of each of the coefficients from the marginal posterior distributions are summarized. The point estimates and the uncertainties about them are computed by using equation 2.8 with $g(\cdot)$ in a form of an identity and standard deviation functions respectively. Furthermore, credible intervals (CI) report intervals which include the true unknown parameters with a 95% probability. The CIs are identified by finding the sample quantiles of the marginal posterior distributions of interest which correspond to the 2.5th and 97.5th percentiles.

Statistical significance of the results included in BANOVA is indicated via two-sided Bayesian p-values presented in section 2.4. This quantity is computed by finding fractions of the draws from the marginal posterior distribution of parameters of interest, which are above and below zero, as defined in equations 2.13-2.14, and multiplying the smaller value by two. If a one-sided directional hypothesis is tested with a prior distribution symmetric around zero, the p-values can be directly transformed into Bayes Factors by using equations in 2.11 and interpreted as evidence for or against a null hypothesis.

It is common to use a coefficient of determination R^2 to assess model's goodness-of-fit, measured as a proportion of variance in the dependent variable, which is explained by this model. BANOVA implements a generalization of the traditional R^2 to the case of multilevel models with mixed-effects (Dong and Wedel, 2017). This generalization allows one to account for different sources of variation at different levels of the models. Furthermore, it is also applicable for models with non-Normal dependent variables. In BANOVA the R^2 is computed during each iteration for a given set of coefficient drawn from the MC after the warm-up period. For a single- and multilevel models it is computed as

$$R^2_{single} = \frac{1}{M}\sum_{m=1}^{M}\left(\frac{\mathrm{Var}(\mathbf{X}\hat{\boldsymbol{\theta}}^{(m)})}{\mathrm{Var}(\mathbf{X}\hat{\boldsymbol{\theta}}^{(m)}) + s_y}\right), \tag{3.13}$$

$$R^2_{multi} = \frac{1}{M}\sum_{m=1}^{M}\left(\frac{\mathrm{Var}(\mathbf{X}\hat{\boldsymbol{\theta}}_j^{(m)})}{\mathrm{Var}(\mathbf{X}\mathbf{Z}\hat{\boldsymbol{\beta}}^{(m)}) + s_y}\right), \tag{3.14}$$

where the term s_y is the same correction factor as the one used for the η_p^2 calculation.

4 Mediation and Moderation Techniques

Scientific investigation of a certain phenomenon begins by establishing an association between relevant variables. As the relationships between variables gain evidential support, the path of the investigation pivots into the causal direction, which searches for an understanding of the underlying mechanism. By revealing how and when investigated effects are manifested scientists grow our understanding of the world and illuminate possibilities for effective interventions, thereby contributing to both fundamental and applied research. Mediation analysis is concerned with how or through which other variable or variables the effect of an independent variable is carried on a dependent variable. Moderation analysis investigates if this effect is inhibited or facilitated based on another variable or variables. These two analysis techniques are widely used in social and behavioral sciences, as well as health and medical research (Hayes, 2017).

This section first introduces moderation analysis by considering theoretical and conceptual aspects of the topic and then mentioning the functionality of the BANOVA R-package for moderation. Mediation analysis is presented in a subsequent subsection by first digging into the theoretic and methodological facets of the technique, followed by a discussion of the current capabilities of the package in this domain.

4.1 Moderation analysis

4.1.1 Theoretic framework

Moderation or interaction analysis is concerned for whom or when the effect of $\boldsymbol{x}$ on $\boldsymbol{y}$ occurs, and what is its magnitude. If size, sign, or strength of the effect of $\boldsymbol{x}$ depends on a variable $\boldsymbol{w}$, it is referred to as a moderator of $\boldsymbol{x}$. This relationship is visualized in figure 4.1.1. The goal of moderation analysis is to analyze how the relationship between $\boldsymbol{x}$ and $\boldsymbol{y}$ changes as a function of $\boldsymbol{w}$. To model this relationship an interaction of $\boldsymbol{x}$ and $\boldsymbol{w}$ must be included in the regression model (Cohen et al., 2013). Adding an interaction term to an equation gives rise to a differentiation between main and joint effects of the interacting variables. An effect of a given variable averaged across the levels of other independent variables is referred to as the main effect. The interaction effect arises due to a combination of variables influencing the dependent variable at the same time, which can result in the joint effects of the variables being above or below the sum of their separate contributions. A significant coefficient of the interaction term implies a mathematically symmetric relationship, which means that the effects of each variable are conditional on the levels of the other variable involved in the interaction.

To find evidence of moderation in the simple case with one explanatory variable $\boldsymbol{x}$ and moderator $\boldsymbol{w}$, a GLM for the outcome variable can be fit. As described in section 3.4 the expected value of $\boldsymbol{y}$ can be modeled as $\mathrm{E}(\boldsymbol{y}) = g^{-1}(\boldsymbol{\eta})$, where the linear predictor of $\boldsymbol{y}$ is defined as

$$\boldsymbol{\eta} = \theta_0 + \theta_1 \boldsymbol{x} + \theta_2 \boldsymbol{w} + \theta_3 \boldsymbol{w} \circ \boldsymbol{x}. \tag{4.1}$$

Here $\boldsymbol{x}$ and $\boldsymbol{w}$ can be either continuous or dichotomous variables, and $\circ$ is the element-wise product of the two variables corresponding to the interaction term. The effects θ_1 and θ_2 are referred to as conditional main effects, and θ_3 is the interaction effect. If the estimate $\hat{\theta_3}$ is determined to be significantly different from zero, the moderation between $\boldsymbol{x}$ and

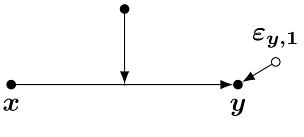

Figure 4.1.1: Conceptual graph a model in equation 4.1, where $\boldsymbol{w}$ moderates the relationship between $\boldsymbol{x}$ and $\boldsymbol{y}$.

$\boldsymbol{w}$ is established. However, it is necessary to probe the interaction further. This means performing additional inferential tests for specific patterns of the effects of $\boldsymbol{x}$ on $\boldsymbol{y}$ as a function of $\boldsymbol{w}$. These conditional or simple effects can be found by rearranging equation 4.1 in the following manner:

$$\boldsymbol{\eta} = \theta_0 + \underbrace{(\theta_1 + \theta_3 \boldsymbol{w})}_{\omega_{\boldsymbol{x}}} \circ \boldsymbol{x} + \theta_2 \boldsymbol{w}. \tag{4.2}$$

From equation above it is evident that simple effect of $\boldsymbol{x}$ is a linear function of $\boldsymbol{w}$, defined as $\omega_{\boldsymbol{x}} = (\theta_1 + \theta_3 \boldsymbol{w})$. Probing the interaction, means evaluated and testing significance of the effect of $\boldsymbol{x}$ at specific values of $\boldsymbol{w}$. If the moderator is a continuous variable the simple slope of $\boldsymbol{x}$ can be evaluated and tested at some observed values of $\boldsymbol{w}$, which are usually selected to be the mean and one or two standard deviations around it. This technique is referred to as spotlight analysis and it can also be used with some meaningful user-specified values of the moderator (Jaccard et al., 1990). To avoid the selection of specific values of $\boldsymbol{w}$ for the evaluation of the effect of $\boldsymbol{x}$ floodlight analysis can be used instead (Spiller et al., 2013). This technique identifies a range of values of the continuous moderator where there is a statistically significant effect of $\boldsymbol{x}$ on $\boldsymbol{y}$. The endpoints of the range are referred to as Johnson-Neyman (JN) points. The estimation details of floodlight analysis are beyond the scope of this thesis and can be found in the book of Hayes (2017). Furthermore, for Bayesian treatment of floodlight analysis see the publication of Wedel and Dong (2019).

If $\boldsymbol{w}$ is a factor, the effects of $\boldsymbol{x}$ at different levels of $\boldsymbol{w}$ are referred to as simple effects. The calculation implies fixing the levels of the moderating variable $\boldsymbol{w}$ one by one and combining the coefficients for the variable which is being moderated. As an illustration consider a model in equations 4.1-4.2, where $\boldsymbol{x}$ and $\boldsymbol{w}$ are assumed to be effect-coded factors, with (1) indicating the first levels and (-1) the second levels of each of the variables. The simple effects $\omega_{\boldsymbol{x}}$ at different values of moderator are defined as

$$\begin{aligned} \omega_{\boldsymbol{x}}|(\boldsymbol{w} = 1) &= (\theta_1 + \theta_3)\boldsymbol{x} \\ \omega_{\boldsymbol{x}}|(\boldsymbol{w} = -1) &= (\theta_1 - \theta_3)\boldsymbol{x} \end{aligned} \tag{4.3}$$

If one or both of the variables involved in equation 4.1 are factors with more than two levels the expression changes, such that for each variable with k levels $k-1$ regressors are included in the model. A discussion about the calculation of simple effects in a general case with an arbitrary number of interacting factors and its implementation is presented in section 5.1.

It is possible to include a higher-order interaction into the estimation, which can represent a simultaneous effect of three or more independent variables on the investigated outcome. Simple effects of each of the factors and their interactions can be calculated for each level of the moderating factor. Thus in a model with variables $\boldsymbol{x}_1$, $\boldsymbol{x}_2$, and a moderator $\boldsymbol{w}$, simple effects of $\boldsymbol{x}_1$ can be computed either at zero or an average values of $\boldsymbol{x}_2$. To achieve the latter it is necessary to encode $\boldsymbol{x}_2$ with sum contrast coding or mean center the variable before estimation (Spiller et al., 2013). The calculation is performed analogously as for two-way interaction, with $\boldsymbol{x}_2$ evaluated at zero. The simple effects for $\boldsymbol{x}_2$ are computed similarly. Derivation the effect of the interaction $\boldsymbol{x}_1\boldsymbol{x}_2$ at the different levels of $\boldsymbol{w}$ is performed in the same manner as in case of only one moderated variable presented in equation 4.2.

In hierarchical models interactions can manifest within a given level of the model of between levels. In the slopes-as-outcomes models described in section 3.2 cross-level interaction is especially common. As an illustrations consider a model with one within-subject variable $\boldsymbol{x}$ and one between-subject moderator $\boldsymbol{w}$, defined as

$$\begin{aligned} \mathrm{E}(y_i|x_i, w_s) &= g_y^{-1}(\theta_0^s + \theta_1^s x_i), \\ \theta_0^s &= \beta_{00} + \beta_{01} w_s + \varepsilon_{0,s}, \\ \theta_1^s &= \beta_{10} + \beta_{11} w_s + \varepsilon_{1,s}. \end{aligned} \tag{4.4}$$

The model above can be re-expressed in the reduced form as

$$\mathrm{E}(y_i|x_i, w_s) = g_y^{-1}(\beta_{00} + \beta_{01} w_s + \beta_{10} x_i + \beta_{11} w_s x_i + \varepsilon_{0,s} + \varepsilon_{1,s} x_i). \tag{4.5}$$

From equation 4.5 it evident that the regression of the subject specific slopes on $\boldsymbol{w}$ implicitly assumes a cross-level interaction between $\boldsymbol{x}$ and $\boldsymbol{w}$. By analogy with single-level models, conditional effect $\omega_{\boldsymbol{x}}$ is defined as $(\beta_{10} + \beta_{11}\boldsymbol{w})$. For a detailed treatment of the estimation and testing of moderation effects in a frequentist framework an interested reader can consult with the publication of Bauer and Curran (2005).

4.1.2 Moderation analysis in BANOVA

In BANOVA Bayesian floodlight analysis for one continuous moderator $\boldsymbol{w}$ and a factor variable $\boldsymbol{x}$ is implemented in the function `BANOVA.floodlight()` (Wedel and Dong, 2019). Based on the fitted regression returned by the `BANOVA.run()` function it finds where in the distribution of the moderator $\boldsymbol{w}$ there is a significant difference between the effect of $\boldsymbol{x}$ at different levels. It computes a point where the difference between levels is zero and a corresponding 95% CI for, which is equivalent to JN points described in the previous section. For the values of $\boldsymbol{w}$ inside the CI, there is no difference between the levels of $\boldsymbol{x}$, the values outside of CI the difference is evidential. However, JN points should only be considered if they are within the range of $\boldsymbol{w}$ (Hayes, 2017).

4.2 Mediation analysis

4.2.1 Theoretic framework

A researcher interested in understanding mechanisms of causal effects of an independent variable or treatment $\boldsymbol{x}$ on an outcome variable $\boldsymbol{y}$ can resort to the application of mediation analysis techniques. The simplest form of this analysis assumes that one additional variable $\boldsymbol{m}$ causally links the dependent and independent variables (MacKinnon, 2008). This variable is referred to as a mediator and is hypothesized to transmit the influence of the explanatory variable on the outcome. The goal of mediation analysis is to partition the effect of the independent variable into direct effect from $\boldsymbol{x}$ to $\boldsymbol{y}$, and indirect, or mediated, effect transmitted via $\boldsymbol{m}$.

Regression-based techniques, such as ordinary least squares (OLS) or structural equation modelling (SEM), are the most commonly used approaches in mediation analysis (Hayes, 2017). These techniques are developed based on the causal step approach designed by Baron and Kenny (1986). To estimate direct and indirect effects of an independent variable in the simple case described above three-variable models, presented in the figure 4.2.2, are defined as

$$\boldsymbol{y} = \theta_{0,1} + \theta_1 \boldsymbol{x} + \varepsilon_{\boldsymbol{y},\mathbf{1}}, \tag{4.6}$$

$$\boldsymbol{y} = \theta_{0,2} + \theta_1^{'} \boldsymbol{x} + \theta_2 \boldsymbol{m} + \varepsilon_{\boldsymbol{y},\mathbf{2}}, \tag{4.7}$$

$$\boldsymbol{m} = \delta_0 + \delta_1 \boldsymbol{x} + \varepsilon_{\boldsymbol{m}} \tag{4.8}$$

A dependent variable $\boldsymbol{y}$ and a mediator $\boldsymbol{m}$ are assumed to be linearly related continuous variables, while the independent variable $\boldsymbol{x}$ can be dichotomous or continuous. Error terms $\varepsilon_{\boldsymbol{y},\mathbf{1}}, \varepsilon_{\boldsymbol{y},\mathbf{2}}$, and $\varepsilon_{\boldsymbol{m}}$ are assumed to follow a Normal distribution with zero means and finite variances. A parameter θ_1 shows the total effect of $\boldsymbol{x}$ on $\boldsymbol{y}$. The direct effects of an independent variable, or the strength of relationship between $\boldsymbol{x}$ and $\boldsymbol{y}$, while controlling for the effect of $\boldsymbol{m}$, is represented by $\theta_1^{'}$. A parameter θ_2, illustrates the strength of relationship between $\boldsymbol{m}$ and $\boldsymbol{y}$, while controlling for the effect of $\boldsymbol{x}$. A parameter δ_1, which is the strength of relationship between $\boldsymbol{x}$ and $\boldsymbol{m}$, together with θ_2 can be used to quantify the mediated effect of $\boldsymbol{x}$ on $\boldsymbol{y}$, denoted as γ. This is accomplished by taking a product of two parameters, such that $\gamma = \delta_1 \theta_2$. Alternatively the mediated effect γ can be computed based on the difference between θ_1 and $\theta_1^{'}$, which implies that controlling for $\boldsymbol{m}$ decreases the total effect, and thus quantifies the effect mediated through $\boldsymbol{m}$. These two approaches yield the same results in correctly specified Normal linear models with no missing data (MacKinnon et al., 1995). The simple models in equations 4.6-4.8 commonly include other control variables, which are omitted in equations above for convenience.

The most common approach to test the influence of the indirect effects γ is to compute the Sobel test based on the standard errors proposed by Sobel (1982), which assumes that γ is Normally distributed. However, the sampling distribution of the indirect effect often exhibits asymmetric patterns with high skewness (Bollen and Stine, 1990). Furthermore, this technique relies on a large sample approximation and commonly suffers from loss of power in studies with small sample sizes. Alternatively, MacKinnon et al. (2007) suggests using bootstrapping to find an empirical distribution of the mediated effects. In recent years there has been a growing interest in Bayesian mediation analysis introduced by Zhang et al. (2009). The MCMC based techniques allow for a straightforward computation of the standard error of the indirect effect γ, based on the draws of δ_1 and θ_2 from the respective marginal posterior distributions.

Interpretation of the effects on the mediation analysis as causal relies on several assumptions presented by MacK-

(a) (b)

Figure 4.2.1: (a) Conceptual graph of a model in equation 4.6 without mediation. (b) Conceptual graph of a model in equations 4.7-4.8, where $\boldsymbol{m}$ mediates the relationship between $\boldsymbol{x}$ and $\boldsymbol{y}$.

innon (2008). First of all, it is assumed that the correct temporal order in the mediation process is captured via an appropriate timing of the measurements. Next, the assumptions of regression models used for the estimation must be satisfied. This implies that the specified functional form of the relationships in the mediation analysis is assumed to be correct. Lastly, the error terms in equations for mediator (4.8) and dependent variable (4.6, 4.7) are assumed to be independent. This assumption implies that there must be no confounders of the effect of the treatment on the outcome and mediator, conditional on other covariates. Likewise, there must be no confounding of the mediator, conditional on the treatment and other covariates, and the outcome variable. The absence of a spurious correlation between the treatment and the outcome can be assured via a random assignment of subjects in the study to different treatment levels. In the case of the relationship between the mediator and the outcome variable, the assumption about the absence of confounds is more difficult to satisfy. Several solutions to this issue based on methodological study designs and analysis-based approaches can be found in the work of Valente et al. (2017). When mediation analysis is applied to observational studies, the strong assumptions described above are not verifiable. In this case, significant indirect effects do not prove causal relationships but make them more plausible. Without causal interpretation of the estimates, results can only be interpreted as descriptive linear associations between variables. In this case inference about causal relationships in the investigated phenomenon is usually made based on the design of research methodology and logical or theoretical analysis, rather than statistical methods described above (Hayes, 2017).

Mediation and moderation analysis described in the previous section can be combined to asses generalizability of mediated effects. A vast array of different models can be constructed based on theory about the variation in the association between the variables as a function of a one or multiple moderators. Two most common cases are referred to as mediated moderation and moderated mediation. In the first case the influence of treatment on the mediator is moderated by a third variable $\boldsymbol{w}$, while the effect of $\boldsymbol{m}$ on $\boldsymbol{y}$ is constant (see figure 4.2.2a). To incorporate mediated moderation in the model equation 4.7 must be changed such that

$$\boldsymbol{m} = \delta_0 + \delta_1 \boldsymbol{x} + \delta_3 \boldsymbol{w} + \delta_4 \boldsymbol{w} \circ \boldsymbol{x} + \boldsymbol{\varepsilon}_{\boldsymbol{m}}. \tag{4.9}$$

In models with moderated mediation $\boldsymbol{w}$ influences path between $\boldsymbol{m}$ and $\boldsymbol{y}$ (not presented here). In both cases the indirect effect of $\boldsymbol{x}$ depends on the level $\boldsymbol{w}$. Another interesting case occurs when a $\boldsymbol{x}$ itself affects the link between $\boldsymbol{m}$ and $\boldsymbol{y}$, thus moderating its own indirect effect. A conceptual graph of this model is presented in figure 4.2.2b.

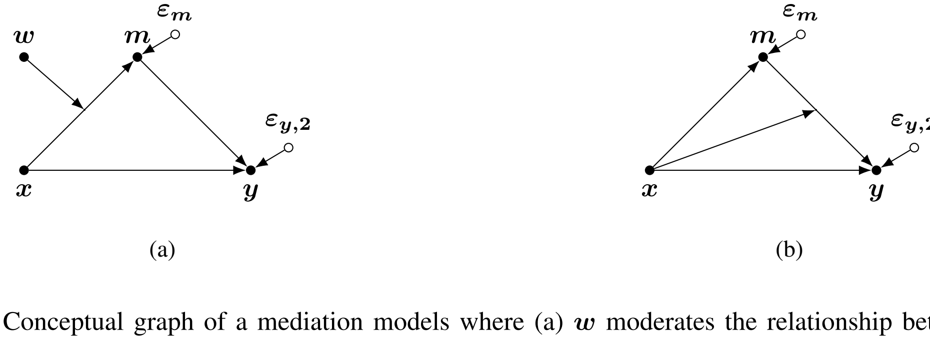

Figure 4.2.2: Conceptual graph of a mediation models where (a) $\boldsymbol{w}$ moderates the relationship between $\boldsymbol{x}$ and $\boldsymbol{m}$ (mediated moderation); (b) $\boldsymbol{x}$ moderates the relationship between $\boldsymbol{m}$ and $\boldsymbol{y}$.

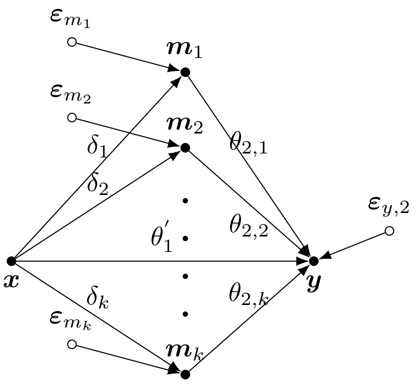

Figure 4.2.3: Conceptual graph of a model with multiple mediators $\boldsymbol{m}_1, \boldsymbol{m}_2, \ldots, \boldsymbol{m}_k$, which mediate the relationship between $\boldsymbol{x}$ and $\boldsymbol{y}$.

The models described above allow for only one mediator, however, in practice multiple mediators can be involved at the same time. Even if only one of these mediators is of interest they must be analyzed together, as omitting mediators from the analysis would violate the no confounders assumption described above. For causal interpretation of the effects in this setting the assumptions described above are generalized such that they apply for each individual mediator in the study (VanderWeele, 2015). The relationships between the mediators give rise to parallel and serial mediation models (Hayes, 2017). The latter assumes the mediators to be causally linked together in a certain order, such that the treatment variable $\boldsymbol{x}$ causes $\boldsymbol{m}_1$, which causes $\boldsymbol{m}_2$, and so on, until $\boldsymbol{m}_k$ affects $\boldsymbol{y}$. Other sources should be consulted for the introduction to this topic, for example Hayes (2017). In models with parallel mediation treatment variable $\boldsymbol{x}$ is hypothesized to influence $\boldsymbol{y}$ directly and indirectly through two or more mediators. These mediators are allowed to be correlated, but must be causally independent. A model with k mediators $\boldsymbol{m}_1, \boldsymbol{m}_2, \ldots, \boldsymbol{m}_k$ can be visualized via a conceptual graph presented in the figure 4.2.3. Further discussion about parallel mediation and its implementation is presented in the section 5.2.

Two important limitations of the classical method described above are effect constancy or "no interaction" assumption, which states that there is no interaction between casual variable and mediator, and its reliance on the linear models. The "no interaction" assumption is required to assure that direct and indirect effects sum to the total effect (Pearl, 2012). If a nonlinear model is investigated or it is assumed that the relationship between the mediator and the outcome variables varies as a function of the treatment, a causal mediation method, proposed by Pearl (2001), can be used. This approach, based on the potential outcomes framework (Rubin, 1974), is presented in the Appendix B.

4.2.2 Mediation analysis in BANOVA

BANOVA allows its users to analyze models with mediation and combination of moderation and mediation in cases when a single mediator and one casual variable are considered. The function `BANOVA.mediation()` takes as an input results of Bayesian hierarchical regressions for the dependent and mediating variables obtained via `BANOVA.run()` function. The algorithm is a generalization of the traditional casual step approach, as it estimates direct and indirect effects based on previously obtained samples of relevant regression coefficients. For all models, the mediation effects are correlational, except for liner models with an interaction between the mediator and a dummy coded causal variable. A further discussion of causal mediation analysis is presented in Appendix B.

The mediation function allows for a wide range of possible combinations of single- and multilevel models. In the latter case, the treatment and mediator can both occur at the within-subject level, between-subject levels, or at different levels. Both of these focal variables can be additionally moderated. A detailed description of possible models is presented in the publication of Wedel and Dong (2019). All models allow for the mediation analysis of a Normally distributed mediator and outcome variables that follow Normal, T, Poisson, Bernoulli, Binomial, and ordered Multinomial distributions. Furthermore, if regressions for the outcome and the mediator are multilevel the results of the mediation analysis can be reported for individual units in the analysis, such as unique subjects in studies with repeated measures.

Based on the posterior samples of a regression for the outcome variable `BANOVA.mediation()` calculates the direct effect of the treatment and effect of the mediator on the dependent variable. To compute the effect of the causal variable on the mediator the regression with the mediator as a dependent variable is used. When the algorithm is calculating the effects of a given variable it first identifies all interactions which include the investigated variable. Based on each interaction moderated coefficients are computed and summarized in a table. If an interaction includes a moderating factor, the effects are reported for each level of the moderator. If a continuous variable is involved in the interaction the effects are evaluated at its mean value, which is zero by default. This is equivalent to omitting the continuous variable from the interaction. If the investigated variable is a part of an interaction term with three or more variables, separate results are computed for each of the moderators and all of their possible combinations. This results in multiple tables reported to the user. Next, the function combines the effects of the mediator on the dependent variable with the effect of the causal variable on the mediator in a multiplicative manner to obtain the indirect effect of the treatment. If multiple tables with moderated effects of the mediator or the causal variable on mediator are obtained in the previous steps of the analysis the indirect effects are computed for each combination of these tables.

To compute mediated effects, as was mentioned before, two models described in sections 3.3-3.4 must be specified. To illustrate the calculations in mediation analysis consider two models with a binary treatment variable $\boldsymbol{x}$. While in the examples below dependent variables and corresponding mediators have the same number of levels other combinations of single- and multilevel models are possible. Furthermore, other covariates and moderators can be included in the models in a straight forward manner. In the first mediation model, presented in equations 4.10-4.11, equations for mediator and outcome variables are single-level and are defined for $i = 1, \dots, N$, where N is the total number of observations in the sample. In this case, a treatment-mediator interaction is included as a regressor.

$$\mathrm{E}(y_i|x_i, m_i) = g_y^{-1}(\theta_0 + \theta_1 x_i + \theta_2 m_i + \theta_3 x_i m_i) \tag{4.10}$$

$$\mathrm{E}(m_i|x_i) = g_m^{-1}(\delta_0 + \delta_1 x_i) \tag{4.11}$$

In the second mediation model, defined in equations 4.12-4.13, the mediator and the outcome are multilevel models, with $i = 1, \dots, N$ and $s = 1, \dots, S$, where S is the unique number of subjects in the study. In this model the mediator is at the within-subject level, and treatment is at the between-subject level, thus there is an implicit interaction between the casual and mediating variables.

$$
\begin{aligned}
\mathrm{E}(y_i|x_i, m_i) &= g_y^{-1}(\theta_0^s + \theta_1^s m_i) \\
\theta_0^s &= \beta_{00} + \beta_{01} x_s + \varepsilon^{0,s} \\
\theta_1^s &= \beta_{10} + \beta_{11} x_s + \varepsilon^{1,s}
\end{aligned}
\tag{4.12}
$$

$$
\begin{aligned}
\mathrm{E}(m_i|x_i) &= g_m^{-1}(\delta_0^s) \\
\delta_0^s &= \psi_{00} + \psi_{01} x_s + \xi_{0,s}
\end{aligned}
\tag{4.13}
$$

To calculate mediated effects as a product the mediators must be Normally distributed, and thus in both models, g_m^{-1} is an identity function. By default, BANOVA uses effect coding for all factor variables (see section 3.4 for details), and thus all the effects in the mediation analysis are reported for both levels of the treatment. The direct effects in the first and second mediation models are $\pm\theta_1$ and $\pm\beta_{01}$ respectivly. A positive sign of the effects corresponds to the first level of the treatment, while a negative sign corresponds to the second level of $\boldsymbol{x}$. The interaction with the mediator does not influence direct effects, as it is evaluated at the mean value of the mediator, which is zero by default. The effects of the mediator on the dependent variable in both models are moderated by the causal variable $\boldsymbol{x}$, and thus are computed as $\theta_2 \pm \theta_3$ in the first mediation model and $\beta_{10} \pm \beta_{11}$ in the second. The effects of the causal variables on the mediators are computed based on equation 4.11 for the first model and 4.13 in the second model, such that they are $\pm\delta_1$ and $\pm\psi_{01}$ respectively. In the final step of the algorithm the indirect effects are computed, such that in the first model the treatment has an effect of $(\theta_2 + \theta_3)(\delta_1)$ and $(\theta_2 - \theta_3)(-\delta_1)$ at the first and second levels of the causal variable. In the multilevel mediation model the indirect effects are $(\beta_{10} + \beta_{11})(\psi_{01})$ and $(\beta_{10} - \beta_{11})(-\psi_{01})$ at different levels of the treatment variable.

`BANOVA.mediation()` computes posterior means, standard deviations and CI of the effects defined above based on equation 2.8 as described in section 3.5, which takes combinations of the samples from marginal posterior distributions of the relevant parameters as an input. Furthermore, two sided p-values are calculated for each effect in the mediation analysis based on equations 2.13-2.14, as described in sections 2.4. To calculate the effect size of the indirect effect Wedel and Dong (2019) proposed a generalized partial eta-squared (η_p^2), conceptually similar to the effect sizes defined in equations 3.11-3.12. A detailed description of the metric is beyond the scope of this thesis and can be found in the publication of Wedel and Dong (2019).

5 Contributions to the BANOVA R-package

The main objective of this project was to contribute to the development of the BANOVA R-package by extending its functionality for moderation and mediation analysis. The first step in the project was to fork an original repository to my personal GitHub account[4]. This allowed me to create a copy of the original code in a separate environment where

---

[4] `https://github.com/kopyakova/BANOVA_R`

I could experiment without affecting the original package. Next, I performed an extensive review of the package to assure, that I utilize the existing codebase effectively. For the development of each of the sub-projects, I created a new branch in the forked repository. After I developed and refactored the code, it went through multiple rounds of testing. In the first stage, I assessed the accuracy of the results obtained with a new function by performing multiple analyses on empirical and synthetic data sets. These results were checked manually, as well as compared with the outcomes of other software or package, such as JASP, SPSS, brms R-package, and others. In the next stage of testing, I confirmed that new functions operate on all popular operating systems, such as Windows, macOS, and Linux. Finally, to assure that the new version of the package is compliant with various CRAN policies I performed the standard CRAN testing (The Comprehensive R Archive Network, 2020). Once all the issues and errors were resolved the new branches were merged to the master branch in the original repository and released to CRAN.

In this chapter, I describe my main contributions to the BANOVA. The functionality for mediation analysis is extended with an algorithm for the estimation of the simple effect in models with multiple factor variables. The contribution to mediation analysis consists of a function for the estimation of mediation effects in studies with multiple possibly correlated mediators. To fit multiple mediators a model with Multivariate Normal distribution is integrated into the existing function for the fitting of the Bayesian hierarchical models. While it was developed as a part of the extension of mediation functionality it can also be used as a stand-alone analysis. These new functions are accessible in the latest version of the package available on CRAN and additionally showcased in a separate repository on GitHub[5]. The documentation for the new functions is presented in the Appendix C.1.

Each of the following subsections first defines the implemented analysis, followed by a description of the developed functions and their integration in the existing codebase. It is important to note that during the first stage of testing several bugs were identified in the original code of BANOVA. To improve the quality of the package these issues were fixed and released to CRAN along with the new function. A list of resolve issues is presented in the Appendix C.3.

5.1 Simple effects

In the function `BANOVA.simple()` I implemented a program for probing interaction effects in models where both moderator and explanatory variables are factors with arbitrary number of levels. Consider a model with a single interaction where $\boldsymbol{x}$ and $\boldsymbol{w}$ have P and Q levels respectively. As mentioned in section 3.4, in BANOVA these factors are included in the regression as $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{P-1}$ and $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_{Q-1}$ effect coded regressors. The generalization of the model presented in equation 4.1 is defined as

$$\begin{aligned}\boldsymbol{\eta} &= \theta_0 + \sum_{i=1}^{P-1} \theta_{1,i}\boldsymbol{x}_i + \sum_{j=1}^{Q-1} \theta_{2,j}\boldsymbol{w}_j + \sum_{i=1}^{P-1}\sum_{j=1}^{Q-1} \theta_{3,i,j}\boldsymbol{w}_j \circ \boldsymbol{x}_i \\ &= \theta_0 + \sum_{i=1}^{P-1}(\theta_{1,i} + \sum_{j=1}^{Q-1} \theta_{3,i,j}\boldsymbol{w}_j) \circ \boldsymbol{x}_i + \sum_{j=1}^{Q-1} \theta_{2,j}\boldsymbol{w}_j,\end{aligned} \tag{5.1}$$

---

[5] `https://github.com/kopyakova/New_functions_in_BANOVA`

where $\theta_{1,i}$ and $\theta_{2,j}$ are conditional main effects of $\boldsymbol{x}_i$ and $\boldsymbol{w}_j$, $\theta_{3,i,j}$ is the interaction effect between $\boldsymbol{x}_i$ and $\boldsymbol{w}_j$, for $i = 1, \ldots, P-1$ and $j = 1, \ldots, Q-1$. From the second line of equation 5.1 it is evident that simple effects of $\boldsymbol{x}$ is a set of relative conditional effects $\boldsymbol{x}_i$, which depend on the level of $\boldsymbol{w}$. Thus, the effect of the i^{th} level of $\boldsymbol{x}$ at the j^{th} level of the base factor $\boldsymbol{w}$ is defined as

$$\omega_{\boldsymbol{x}}|(\boldsymbol{x}=i, \boldsymbol{w}=j) = \theta_{1,i} + \theta_{3,i,j}, \text{ for } i = 1, \ldots, P-1, \text{ for } j = 1, \ldots, Q-1. \tag{5.2}$$

With the effects coding scheme it is also possible to evaluate the effects at the last levels of $\boldsymbol{x}$ and $\boldsymbol{w}$. The effect of the last level of a given factor can be obtained by evaluating all of the regressors associated with this factor at value of (-1). The effects of the edge cases are computed as:

$$\begin{aligned}
\omega_{\boldsymbol{x}}|(\boldsymbol{x}=P, \boldsymbol{w}=j) &= -\sum_{i=1}^{P-1}(\theta_{1,i} + \theta_{3,i,j}), \text{ for } j = 1, \ldots, Q-1; \\
\omega_{\boldsymbol{x}}|(\boldsymbol{x}=i, \boldsymbol{w}=Q) &= \theta_{1,i} - \sum_{j=1}^{Q-1}\theta_{3,i,j}, \text{ for } i = 1, \ldots, P-1; \\
\omega_{\boldsymbol{x}}|(\boldsymbol{x}=P, \boldsymbol{w}=Q) &= -\sum_{i=1}^{P-1}\theta_{1,i} + \sum_{i=1}^{P-1}\sum_{j=1}^{Q-1}\theta_{3,i,j}.
\end{aligned} \tag{5.3}$$

The implemented function can also calculate simple effects for a variable with user-defined contrasts. In this case, the effects are reported only for the defined regressors as specified in equation 5.2, because the edge cases defined above might not have a meaningful interpretation.

`BANOVA.simple()` calculates the effects defined in equations 5.2-5.3 by taking as an input a fitted model of interest obtained with `BANOVA.run()` function, passed via an argument `BANOVA_output`. The function evaluates conditional effects of all moderated factors and their combinations relative to the levels on a single moderator specified as a `base` parameter. Therefore, in a model with multiple possible moderators, only one of them can be evaluated in a single call of `BANOVA.simple()`. For each of the interactions which include the `base`, this function prints a table with moderated effects of a given factor or an interaction. The algorithm in `BANOVA.simple()` finds all possible combinations of levels of the investigated interacting factors. Based on these cases, the posterior samples of the selected regression coefficients are combined to calculate the whole set of simple effects $\omega_{\boldsymbol{x}}$. Since $\omega_{\boldsymbol{x}}$ are linear combinations of estimated parameters, the function draws an inference about them by evaluating combined samples of relevant marginal posterior distributions. For each simple effect, it computes a summary of the results, which contains the mean, standard deviation, CI, and two-sided Bayesian p-value. The first three quantities are computed based on equation 2.8 where $g(\cdot)$ can be identity, standard deviation, and sample quantiles functions. For example, for simple effects presented in equation 5.2 these quantities are obtained as

$$\widehat{g(\omega_{\boldsymbol{x}})}|(\boldsymbol{x}=i, \boldsymbol{w}=j) = \frac{\sum_m^M g\left(\theta_{1,i}^{(m)} + \theta_{3,i,j}^{(m)}\right)}{M}. \tag{5.4}$$

Two-sided Bayesian p-values are computed in the same way as described in section 3.5. The CIs by default report central 95% intervals, which can be adjusted via the `quantiles` argument. `BANOVA.simple()` automatically returns a list of all tables previously printed to the user's console. Furthermore, posterior samples of simple effects can be returned as well if the argument `return_posterior_samples` is set to `TRUE`.

The procedure described above applies for models with Normal, T, Poisson, Bernoulli, Binomial, and ordered Multinomial distributions. A special case of the analysis occurs if simple effects in a Multinomial model are of interest. Multinational regression models a nominal outcome variable with more than two levels, which are referred to as alternatives. Given a set of explanatory variables this model predicts the probabilities of possible outcomes of the dependent variable. The previously described procedure for calculation of simple effects in multi-level models is applied iteratively for each alternative because in this case the intercept and between-subject regressors have outcome-specific coefficients. To run this special case of the analysis `BANOVA.simple()` requires a user to specify the label of the dependent variable as the `dep_var_name` argument. Furthermore, the function works with the Multivariate Normal models introduced later in section 5.3.

5.2 Mediation analysis with parallel mediators

A new function `BANOVA.multi.mediation()` extends the functionality of BANOVA described in section 4.2.2 to the case with multiple, possibly correlated, parallel mediators. These mediators, as was mentioned in section 4.2.1, are assumed to have no causal influence on each other. Consider a single-level model with one causal variable $\boldsymbol{x}$ and L mediators $\boldsymbol{m}_1, \boldsymbol{m}_2, \dots, \boldsymbol{m}_L$, which can be defined as

$$\mathrm{E}(\boldsymbol{y}) = g_{\boldsymbol{y}}^{-1}(\theta_0 + \theta_1 \boldsymbol{x} + \sum_{l=1}^{l} \theta_{2,l} \boldsymbol{m}_l), \tag{5.5}$$

$$\boldsymbol{m}_l = \delta_{0,l} + \delta_l \boldsymbol{x} + \boldsymbol{\varepsilon}_{\boldsymbol{m},l}, \text{ for } l = 1, \dots, L. \tag{5.6}$$

In the models above the dependent variable $\boldsymbol{y}$ can follow a Normal, T, Poisson, Bernoulli, Binomial, and ordered Multinomial distributions. The L mediators are assumed to be Normal variables. Under the assumptions presented in section 4.2.1, θ_1 represents the direct effect of $\boldsymbol{x}$ on $\boldsymbol{y}$, while controlling for all mediators. Parameter δ_l represents the effect of $\boldsymbol{x}$ on $\boldsymbol{m}_l$, and $\theta_{2,l}$ is the effect of $\boldsymbol{m}_l$ on $\boldsymbol{y}$, while controlling for $\boldsymbol{x}$ and other $l-1$ mediators. A product of the parameters, which correspond to the same mediator l, represents one of the l specific indirect effects of the treatment. Total indirect effect (TIE) of $\boldsymbol{x}$ on $\boldsymbol{y}$ is defined as a sum of all specific indirect effects:

$$TIE = \sum_{l=1}^{l} \theta_{2,l} \delta_l. \tag{5.7}$$

Mediation effects in multilevel models are obtained similarly by combining population-level coefficients.

To perform parallel mediation analysis with multiple mediators `BANOVA.multi.mediation()` takes as an input two sets of results obtained with `BANOVA.run()`. Fitted models for the dependent variable and for the mediators are passed as arguments `sol_1` and `sol_2` respectively. To fit multiple Normal mediators a Multivariate Normal model introduced in the next subsection must be used. Furthermore, it is necessary for the user to specify the name of the causal variable, as an argument `xvar`, and a vector with labels of the analyzed mediators in an argument `mediators`. `BANOVA.multi.mediation()` adapts the output of the Multivariate Normal model for the mediators to work with an existing function `BANOVA.mediation()`. The details about the calculation of the mediation effects in `BANOVA.mediation()` are presented in section 4.2.2. After `BANOVA.mediation()` was called for each of the analyzed mediators the TIE of the causal variable, defined in equation 5.7, is computed based on the posterior samples of the specific indirect effects. To obtain these effects `BANOVA.mediation()` was modified to return posterior draws of mediated effects if an argument `return_posterior_samples` is set to `TRUE`. Next, the posterior samples of the TIE are summarized and reported to the user.

The new function reports to users multiple tables with mediated effects. Tables with direct effects of the causal variable and mediators on the outcome variable, as well as direct effects of the causal variable on the mediators include a posterior mean and a 95% CI of the effects. Next, the function displays on the console tables with specific indirect effects and effect sizes of the mediators, followed by the TIE of the causal variable. These tables include the mean, 95% CI, and two-sided Bayesian p-values. All of these results are also returned by the function as a list. If models in `sol_1` and `sol_2` are multilevel the results can be reported for individual units in the analysis by setting the argument `individual` to `TRUE`.

5.3 Multivariate Normal model in BANOVA

To model multiple correlated dependent variables I integrated a Multivariate Normal distribution in the `BANOVA.run()` function. The fitting of these models is performed with seemingly unrelated regressions (SUR) introduced by Zellner (1962). This model is suitable for L dependent variables with shared predictors and an explicit covariance structure of error terms. The proposed covariance structure assumes that the disturbances for the multiple outcomes are correlated within a single observation unit, and are independent between the observation units (Cameron and Trivedi, 2005). A multivariate regression model assumes that a set of within- and between-subject continuous and factor predictors are shared between L dependent variables. The regression for the l^{th} of L equations observation for the i^{th} of N observation is as follows:

$$y_{l,i} = \boldsymbol{x}_i^T \boldsymbol{\theta}_{l,s} + u_{l,i}, \quad l = 1, \dots, L, \quad i = 1, \dots, N. \tag{5.8}$$

In equation above $y_{l,i}$ is related to a vector $\boldsymbol{x}_i$ of length J which includes within-subject regressors and an intercept. The subject-specific effects of the regressors on the dependent variable l are stored in the J-vector $\boldsymbol{\theta}_{l,s}$. If all L equations for a single observation are stacked a multivariate hierarchical linear model can be formulated:

$$\begin{bmatrix} y_{1,i} \\ y_{2,i} \\ \vdots \\ y_{L,i} \end{bmatrix} = \begin{bmatrix} \boldsymbol{x}_i^T & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \boldsymbol{x}_i^T & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \boldsymbol{x}_i^T \end{bmatrix} \begin{bmatrix} \boldsymbol{\theta}_{1,s} \\ \boldsymbol{\theta}_{2,s} \\ \vdots \\ \boldsymbol{\theta}_{L,s} \end{bmatrix} + \begin{bmatrix} u_{1,i} \\ u_{2,i} \\ \vdots \\ u_{L,i} \end{bmatrix}$$

$$\boldsymbol{y_i} = \mathbf{X}_i \boldsymbol{\theta}_s + \boldsymbol{u}_i, \quad i = 1, \dots, N \tag{5.9}$$

$$\boldsymbol{u}_i \sim \mathrm{N}\left(\mathbf{0}, \boldsymbol{\Sigma}\right) \tag{5.10}$$

Here $\boldsymbol{y_i}$ and $\boldsymbol{u}_i$ are vectors of length L, $\mathbf{X}_i$ is a L by JL matrix with regressors for the i^{th} observation, and $\boldsymbol{\theta}_s$ is a vector of length JL with subject-specific coefficients for all the dependent variables. Vector $\boldsymbol{u}_i$ follows a Multivariate Normal distribution with a L dimensional mean vector of zeros and a L by L covariance matrix $\boldsymbol{\Sigma}$, which is the assumed to be the same for all observations.

In the hierarchical model, $\boldsymbol{\Sigma}$ itself has a prior. Following an approach proposed by Barnard et al. (2000) the covariance matrix is decomposed in terms of its correlation matrix and corresponding standard deviations. Given the L by L correlation matrix $\boldsymbol{\Omega}$ and a vector with standard deviations $\boldsymbol{\tau}$ of length L, $\boldsymbol{\Sigma}$ can be expressed as specified in equation 5.11. The prior on each element of $\boldsymbol{\tau}$ is weakly informative and follows a half-Cauchy distribution with a scale of 2.5. The correlation matrix $\boldsymbol{\Omega}$ follows a LKJ (η) prior distribution (Lewandowski et al., 2009), which is a density proportional to the determinant of the matrix raised to the power η. The value of η is selected to be 1, which corresponds to the uniform distribution over correlation matrices. This implies that the marginal distribution of the entries in $\boldsymbol{\Omega}$ are distributed according to Beta distribution scaled between -1 and 1, and centered at 0.

$$\boldsymbol{\Sigma} = diag(\boldsymbol{\tau})\boldsymbol{\Omega} diag(\boldsymbol{\tau}) \tag{5.11}$$

$$\tau_l \sim \text{Cauchy}\,(0, 2.5)\, I\,(0, \infty)\,, \quad l = 1, \ldots, L \tag{5.12}$$

$$\boldsymbol{\Omega} \sim \text{LKJCorr}\,(1) \tag{5.13}$$

The influence of between-subject regressors is captured in the the population-level model presented in equation 5.14. A detailed description of the second level of the models in BANOVA is presented in the section 3.3.

$$\boldsymbol{\theta_s} = \mathbf{W}_s\boldsymbol{\beta} + \boldsymbol{\varepsilon}_s, \quad s = 1, \ldots, S \tag{5.14}$$

$$\boldsymbol{\varepsilon}_s \sim \text{N}\left(\mathbf{0}, \sigma_\theta^2 \mathbf{I}\right) \tag{5.15}$$

$$\boldsymbol{\beta} \sim \text{N}\,(\mathbf{0}, 10000\mathbf{I}) \tag{5.16}$$

$$\sigma_\theta^2 \sim \text{InvGamma}\,(1, 1) \tag{5.17}$$

In the model above $\boldsymbol{\theta}_s$ is explained in terms of the population-level regressors for subject s stored in the JL by KJ matrix $\mathbf{W}_s$. Vector $\boldsymbol{\varepsilon}_s$ of length JL follows a Multivariate Normal distribution with zero mean vector of size JL and a diagonal covariance matrix of the respective size with σ_θ^2 as diagonal elements, which in turn follows an Inverted Gamma $(1, 1)$ distribution. A KJ-vector $\boldsymbol{\beta}$ has a multivariate Norma distribution with a zero mean and the variances of a ten thousand.

To run an analysis with a Multivariate Normal model one must first call the function `BANOVA.model()` with a `"multiNormal"` model name, which loads Stan code for a single-level or multilevel SUR models presented in code snippets C.1-C.2 in Appendix C.2. After a fitted model was obtained, as described in section 3.4, `BANOVA.run()` can be called in the same way as for other models. The input data, however, must be specially prepared: a set of dependent variables should be combined in a single matrix, which is next added to a data frame used in the analysis. The name of the matrix with multiple outcomes is expected to be passed as the name of the dependent variable in the model specification. Once `BANOVA.run()` has been called it performs multiple checks of the dependent variables, followed by MCMC sampling. A separate internal function `results.BANOVA.mlvNormal()` is called to prepare the results of the analysis described in section 3.5. This function first analyzes posterior samples that correspond to different outcome variables one by one and then combines the results to be reported together. A summary of the results additionally reports a correlation matrix, as well as tables with summaries of the posterior distributions of standard deviations and correlations of the dependent variables. The posterior means, standard deviations, 95% CI, and two-sided Bayesian p-values are computed as elsewhere in the package.

6 Applications

In this section, I demonstrate how to compute Bayesian simple effects and Bayesian mediation effects in problems with parallel mediators in the framework of the BANOVA R-package. The illustrations are based on data from previously published experimental research of Perfecto et al. (2019) and observational survey of Huellemann and Calogero (2020). For each application, I present a brief description of the data and the hypotheses of interest, followed by the model specification and relevant R code. Next, I discuss the results and compare them with the outcomes of the original studies. This chapter aims to familiarize the users with the new functions and showcase the flexibility of BANOVA.

6.1 Simple effects in the study by Perfecto, Donnelly, and Critcher (2019)

Data and hypotheses To demonstrate `BANOVA.simple()` in action I reanalyze data from the second experiment in the study by Perfecto et al. (2019), which investigated a relationship between a mental simulation and perceptual judgments in volume assessment. In this study, two hundred fifty-seven participants were asked to estimate the volume of 24 digital cups, which resulted in 6006 observations. A sample of the data is presented in code snippet D.1 in the Appendix D.1. The authors suggested that people often simulate how much they could pour into a cup to estimate its volume. Since the pouring happens from an opening of a cup to a base, orientation, and cavern effects hypotheses about the influence of mental simulation on the perception of volume were formulated. The orientation effect, which implies that a cup seems larger when it is in an upright position, is likely to occur due to the relative cognitive ease of imagining filling in a container with an opening on top. The cavern effect postulates that a low top-to-base ratio of a cup creates a sense that it is cavernous, and thus has a higher volume.

To test these two hypotheses the digital cups were given orientation and shape features. The first factor determined whether a cup was presented with the right side up or upside down, with the corresponding levels being contrasted by encoding them as (1) and (-1) respectively. The second factor differentiated between wide base, wide top, or cylindrical shape, with a contrast between the wide-based (1) and a wide top (-1) containers, while cylindrical cups were encoded as (0). To investigate the nature of orientation and cavern effects, participants' responses to the cognitive simulations were analyzed. The participants were randomly assigned to two groups with varying simulation assignments in each group: in the first group they were asked to imagine filling an empty cup (encoded as (1)), and in the second group they were instructed to imagine pouring out all liquid from the cup (encoded as (-1)). The authors hypothesized that the emptying simulation condition would reduce or eliminate the orientation and cavern effects, as the ease of imagining fulling an upright cup with a narrow opening would not hold.

The analyzed dependent variable is the estimated volume of the digital cups on a scale from 0 to 30 units, with an increment of 0.1. The participants were instructed to evaluate the volume of cups relative to a reference cup of 15 units. Following the original study observations with an estimated volume of zero are excluded from the analysis, resulting in 5999 observations. The histogram and the Q-Q plot of this variable, presented in figure D.1.1 in the Appendix D.1, suggest that normality of the dependent variable is a reasonable assumption, even though its domain is bounded.

Model specification In the second experiment of the Perfecto et al. (2019) study, each participant was assigned to a single simulation condition and then was asked to evaluate the volume of multiple digital cups. Therefore, the levels of shape and orientation factors are nested within the simulation conditions. In BANOVA this data set can be analyzed via a hierarchical model with a Normal dependent variable. Details about the mathematical form of the model and corresponding priors can be found in sections 3.3-3.4. Following the original study, orientation and shape, as well as their interaction, are modeled as subject-level variables with fixed and random intercepts and slopes. Furthermore, the simulation condition is included in the model as a moderating between-subject factor, along with an interaction between all three factors. The authors did not discuss which of the interacting subject-level variables is considered as a second moderator and which one is the independent variable. Therefore, figure 6.1.1 presents the most likely form a conceptual graph which correspond to the specified model. As was mentioned in section 4.2.1 the distinction between moderating and moderated variables is purely theoretical and the interaction effect is mathematically symmetric. Thus, there are other possible ways to illustrate the relationships between the variables. The original model in the Perfecto et al. (2019) study also includes a random cup effect, which is omitted from the reanalysis, as it is currently not supported in the BANOVA R-package.

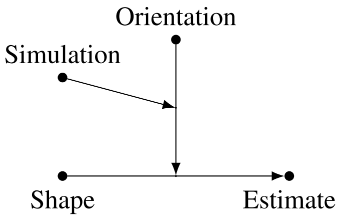

Figure 6.1.1: Conceptual graph of the model in the second experiment of the Perfecto et al. (2019) study, where the relations between the estimated volume of digital cups and its shape is moderated by orientation and participants' simulation conditions.

To test orientation and cavern effect hypotheses within a single model all of the factors are encoded with different contrasts. This approach differs from the one in the Perfecto et al. (2019) study, as it does not require running a second model to test the cavern effects with cylindrical cups excluded from the analysis. To assess the difference between effects of the orientation levels the upright position is encoded as $(\frac{1}{2})$ and upside-down position as $(-\frac{1}{2})$. Furthermore, since the shape factor has three levels it corresponds to two regressors which evaluate differences in the estimated volume of cups with different shapes. I specify the first regressor such that it estimates the difference between wide-based and wide-topped cups, while the second regressor compares cylindrical cups with the average of the wide-based and wide-topped cups. To encode this factor a variation of Helmert coding (Wendorf, 2004) is used. The first regressor contrasts wide-based and wide-topped cups by encoding them as $(\frac{1}{2})$ and $(-\frac{1}{2})$, while cylindrical containers are labeled as (0). The second regressor encodes cylindrical, wide-based and wide-topped cups as $(\frac{2}{3})$, $(-\frac{1}{3})$, and $(-\frac{1}{3})$ respectively. Lastly, the simulation condition is encoded with the default effect coding scheme.

Results To test the hypotheses of interest it is necessary to fit a hierarchical Bayesian model described above. The code for loading the data, specifying factor levels, building and running the model is presented in the code snippet 6.1.

```
R> library(BANOVA)
R> data_perfecto  <- read.csv("perfecto_data.csv", header=T)
R> data_perfecto$simulation  <- factor(data_perfecto$simulation, levels = c(1,-1), labels = c("1", "-1"))
R> data_perfecto$orientation <- factor(data_perfecto$orientation, levels = c(1,-1), labels = c("1", "-1"))
R> data_perfecto$shape       <- factor(data_perfecto$shape, levels = c(1, 0, -1), labels = c("1", "0", "-1"))
R> # Exclude observations with estimated volume of zero
R> data_perfecto  <- data_perfecto[data_perfecto$estimate>0, ]
R> # Build a model
R> model  <- BANOVA.model("Normal")
R> Normal <- BANOVA.build(model)
R> # Run the analysis
R> out1   <- BANOVA.run(l1_formula = estimate ~ orientation * shape, l2_formula = ~simulation,  fit = Normal,
                   data = data_perfecto, id = "id", iter = 30000, thin = 3, chains = 2, cores = 2, seed = 123,
                   contrast = list(orientation = c(1/2,-1/2), shape = cbind(c(1/2, 0, -1/2), c(-1/3, 2/3, -1/3))))
```

Code snippet 6.1: Code for fitting the model in the second experiment of the Perfecto et al. (2019) study.

Table 6.1 presents the obtained results and includes point estimates, 95% CI, and Bayesian p-values for the investigated parameters. The convergence diagnostics, along with a full summary of the results is presented in code snippet D.2 in the Appendix D.1. As indicated by the Geweke and the Heidelberg & Welch convergence diagnostics all chains have converged to the target densities. Consistent with the orientation and cavern effect hypothesis there is very strong evidence for the main effect of orientation ($p = 0.001$) and first shape regressor ($p < 0.0001$), which compares the volume estimate of the wide-based and wide-topped cups. The results also indicate that there is strong evidence for the moderating effect of the simulation when the interactions of the orientation ($p = 0.006$), first shape regressor ($p < 0.0001$), and the interaction of these regressors ($p = 0.014$) are considered. Furthermore, there is substantial evidence, which suggests that the second regressor corresponding to the shape variable is also moderated by the simulation with $p = 0.061$. In this situation, it is not meaningful to interpret the main effects of the orientation and shape. To obtain the effects of the factors moderated by the simulation variable `BANOVA.simple()` was called via the following command:

```
R> simple_effects  <- BANOVA.simple(out1, base = "simulation")
Simple effects of orientation at each level of simulation factor
 simulation  orientation     Simple effect         SD  Quantile 0.025  Quantile 0.975   p-value
          1 orientation1            0.8523     0.2067          0.4443          1.2561   <0.0001
         -1 orientation1            0.0438     0.1909         -0.3347          0.4165    0.8078
```

```
Simple effects of shape at each level of simulation factor
 simulation  shape     Simple effect         SD  Quantile 0.025  Quantile 0.975   p-value
          1 shape1            1.9703     0.2508          1.4771          2.4677   <0.0001
          1 shape2           -1.7835     0.2193         -2.2133         -1.3559   <0.0001
         -1 shape1            0.5643     0.2340          0.0972          1.0266    0.0196
         -1 shape2           -1.2295     0.2086         -1.6385         -0.8220   <0.0001

Simple effects of orientation:shape at each level of simulation factor
 simulation  orientation  shape     Simple effect         SD  Quantile 0.025  Quantile 0.975   p-value
          1 orientation1 shape1            4.6518     0.5889          3.4630          5.7880   <0.0001
          1 orientation1 shape2           -1.8782     0.5490         -2.9643         -0.8222    0.0004
         -1 orientation1 shape1            0.7330     0.5447         -0.3239          1.8193    0.1678
         -1 orientation1 shape2           -1.1540     0.4946         -2.1304         -0.1960    0.0178
```

Code snippet 6.2: Function call of `BANOVA.simple()` and the corresponding console output

The function reported simple effects of all variables which are interacting with the simulation, along with their standard deviations, 95% CI, and Bayesian p-values. Since user-defined contrasts were used to encode orientation and shape, simple effects were calculated with respect to the specified regressors and not underlying factor levels. The top panel of the output indicates that there is very strong evidence ($p < 0.0001$) that the participants who were asked to imagine filling an empty cup when estimating its volume judged it to be larger when the cup was right side up compared to the upside-down position with an average difference of 0.85 units. This effect seized to exist ($p = 0.808$) at the second level of the simulation. Thus, asking the participants to imagine emptying the cup prevents

Regressor	**point estimate**	**95% CI**	**p-vlaue**
intercept	17.92	(17.59, 18.26)	<0.001
simulation	0.12	(-0.22, 0.46)	0.498
orientation	0.45	(0.17, 0.73)	0.001
shape1	1.27	(0.92, 1.60)	<0.001
shape2	-1.51	(-1.81, -1.21)	<0.001
orientation:simulation	0.40	(0.12, 0.68)	0.006
shape1:simulation	0.70	(0.37, 1.04)	<0.001
shape2:simulation	-0.28	(-0.57, 0.01)	0.060
orientation:shape1	0.98	(0.32, 1.63)	0.004
orientation:shape2	-0.46	(-1.07, 0.12)	0.129
orientation:shape1:simulation	0.85	(0.19, 1.50)	0.014
orientation:shape2:simulation	-0.49	(-1.07, 0.11)	0.107

shape1: wide-based vs wide-topped cups

shape2: cylindrical cups vs the average of wide-based and wide-topped cups

Table 6.1: Results in the second experiment of Perfecto et al. (2019). study

the orientation hypothesis from holding. The next panel displays simple effects of the shape factor at the different simulation conditions. The participants who imagined filling the cup judged the containers to be larger when they had a narrow top and a wide bottom compared to wide topped cups with an average difference of 1.97 units ($p < 0.0001$). Thus, under the filling simulation condition, there is strong evidence for the cavern effect hypothesis. For the participants who estimated the volume by emptying the cup, the difference was reduced to 0.56 ($p = 0.02$). These findings are in line with the results of the original paper and the hypothesis about an increased complexity associated with an emptying simulation condition, which decreases the strength of the cavern effect and eliminates the orientation effect.

A side by side comparison of the results described above and the effects reported in the original study is presented in table 6.2. It is evident that the point estimates of the effects are almost the same, but the Bayesian credible intervals are wider than the frequentist confidence intervals. Even though these two intervals have a different interpretation, in practice they are used as a measure of the degree of uncertainty about the estimated parameter. Therefore, the frequentist confidence intervals are underestimating the uncertainly about the obtained results, which is due to the inability of the underlying techniques to explicitly model all relevant variation in multi-level models.

The effects of three-way interactions between the orientation, shape, and simulation are explored in the bottom panel of the output. The results indicate that under a filling simulation condition there is an average difference of 4.65 ($p < 0.0001$) in the estimated volume of upright wide-based cups and upside down wide-topped cups. There is no evidence for this difference under an emptying simulation condition ($p = 0.168$). This result offers additional evidence in support of the previous findings. While the second shape regressor does not contribute to the investigation of the orientation and cavern effects hypotheses it is still interesting to discuss the obtained results. As can be seen in the table of coefficients there is very weak evidence ($p = 0.107$) for the effect of the three-way interaction between the orientation, simulation, and the second shape regressor. However, there is substantial evidence that the simple effects of the orientation and the second shape regressor are moderated by the simulation condition. Their simple effects compare

Simple effect	**Simulation**	**Results of BANOVA**			**Original results**		
		point estimate	***95% CI***	***p-value***	***point estimate***	***95% CI****	***p-value***
Upright vs upside-down	*Fill*	0.85	(0.44, 1.26)	<0.001	0.86	(0.58, 1.13)	<0.001
cups (Orientation)	*Empty*	0.04	(-0.33, 0.42)	0.808	0.04	(-0.21, 0.30)	0.730
Wide-based vs wide-topped	*Fill*	1.97	(1.47, 2.47)	<0.001	1.93	(1.58, 2.27)	<0.001
cups (Shape1)	*Empty*	0.56	(0.10, 1.03)	0.020	0.56	(0.24, 0.88)	0.001

* frequentist confidence interval

Table 6.2: Comparison of the results of the second experiment reported in the Perfecto et al. (2019) study and the reanalysis with BANOVA.

the estimated volume of the upright cylindrical cups with an average estimated volume of upside-down wide-based and wide-topped cups at both simulation levels. It is evident that the difference is decreasing from -1.88 ($p = 0.0004$) at the first level of simulation to -1.15 ($p = 0.018$) at the second level.

Conclusion This application illustrates how the calculation of simple effects based on the output of a regression model with interacting variables simplifies the interpretation of the estimated coefficients. Despite some differences in the analytical approach, the reanalysis yielded the same conclusions as the original study of Perfecto et al. (2019). The replication confirmed that cognitive simulation plays an important role in the translation of the perceptual input into perceptual judgment. The orientation effect was identified only under the filling simulation condition. The cavern effect was found to be stronger when participants were instructed to imagine filling a cup compared to the emptying simulation condition. Additionally, the simple effects of the three-way interactions, which were omitted from the discussion in the original study, offered further insights into the nature of the explored phenomena. These results indicated that when orientation and cavern effects are evaluated together the evidence for their influence is present only under the filling simulation condition.

The analysis with BANOVA has several advantages compared to the techniques used in the Perfecto et al. (2019) study. First of all, it permits the evaluation of both hypotheses of interest based on the output of a single regression, rather than running two separate models. Secondly, `BANOVA.simple()` offers an easy and convenient approach for the calculation of simple effects in problems with an arbitrary number of interacting variables. Lastly, wider 95% CIs associated with Bayesian simple effects do not overestimate the certainty about the estimates in multi-level models, unlike the frequentist approach used in the original study.

6.2 Multiple mediators in the study by Huellemann and Calogero (2020)

Data and hypotheses To illustrate the results of `BANOVA.multi.mediation()` I reanalyze data from a recent observational study published by Huellemann and Calogero (2020). The authors looked at the link between self-compassion and body checking behavior, mediated through stigmatizing self-perceptions. Body checking, which is a habit of excessive monitoring of weight and size of one's body, is associated with a high risk of eating disorder (Walker et al., 2018). A sample of the data, collected from 177 female participants, is displayed in code snippet D.3 in the Appendix D.2. The authors built their hypothesis based on a self-deficit/emotional regulation model, which suggests that weight bias internalization, self-objectification, and body shame are stigmatizing beliefs, which lead to critical body judgments and potential disordered eating (Woodward et al., 2019). The main goal of the survey was to examine if self-compassion is an antecedent which may inhibit the stigmatizing perceptions. To test this hypothesis the authors investigate the effect of self-compassion on body checking exhibited directly, as well as indirectly via parallel mediators. Thus, in this study self-compassion is the independent variable, body checking is the investigated outcome, ans weight bias internalization, self-objectification, as well as body shame are the mediators.

To measure each of the relevant constructs study participants responded to several online questionnaires. Self-

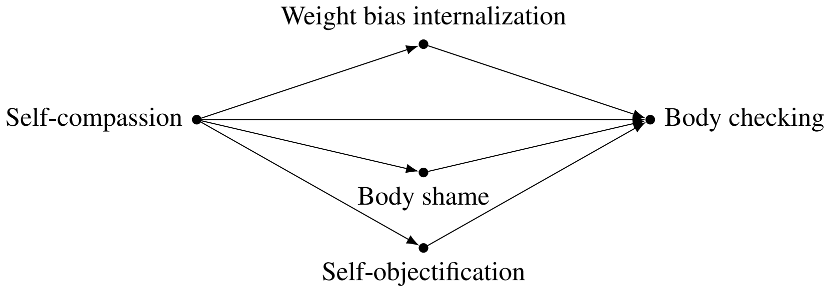

Figure 6.2.1: Conceptual graph the model in the study by Huellemann and Calogero (2020), with the relations between self-compassion and body checking mediated by weight-bias internalization, self-objectification, and body shame.

compassion and body-checking were measured on Likert scales ranging from 1 to 5, with 12 and 23-item questionnaires respectively. Weight bias internalization, which is characterized by the endorsement of narrow beauty ideals, was measured on a 1 to 7 Likert scale with 11 items. Self-objectification, measured via 14 items on a 1 to 5 Likert scale, referrers to the degree of self-evaluation which is due to one's appearance, compared to the internal attributes and experiences. Body shame, which is also measured on a 1 to 5 Likert scale with 18 items, defines to what extent a person feels ashamed of his or her body. In all the scales a higher score indicates a greater propensity to exhibit a certain behavior or hold a high level of a given belief. All scale items were averaged to construct the corresponding unidimensional scales, and thus can be assumed to follow an approximately Normal distribution. This assumption was confirmed for body checking scores and stigmatizing self-perceptions via an investigation of histograms and Q-Q plots presented in figures D.2.1-D.2.4 in the Appendix D.2.

At the end of the survey participants responded to a questionnaire which assessed their age, gender, and socioeconomic status. The moderating effect of socioeconomic status was not tested in the original study. However, here I hypothesize that the indirect association between self-compassion and body checking transmitted via all the stigmatizing self-perceptions is stronger for higher socioeconomic classes. To test this hypothesis I re-coded the levels of socioeconomic status such that such that working and lower middle classes are labeled as ("lower"), middle and upper middle classes are encoded as ("middle"), and upper middle and upper class correspond to an ("upper") label. It is important to note that there were 9% of the participant were in the "lower" class, 50% in the "middle" class, and 41% in the "upper" class.

Specification of the models To investigate the direct and mediated effects of self-compassion on body checking, illustrated in figure 6.2.1, it is necessary to fit two models as described in section 5.2. Since the data for each participant were collected at a single point in time it not possible to fit individual-level models. A single-level population model was presented in the equations 3.3-3.4 and described in section 3.4. In the first model, body checking is regressed on self-compassion and mediators, namely, weight bias internalization, self-objectification, and body shame. The mediators in the study are positively correlated (see figure D.2.5 in the Appendix D.2), and thus, should be modeled together via a Multivariate Normal model. The second model investigates the association between stigmatizing self-perceptions and respondents' self-compassion. In the original study, Huellemann and Calogero (2020) used Model 4 in the PROCESS macro in SPSS 26.0 Hayes (2017). To test hypotheses about the indirect effects of self-compassion the authors considered bias-corrected 95% bootstrap confidence intervals.

To test the hypothesis about the moderating effect of socioeconomic status and demonstrate an analysis with moderated mediation the model described above is reanalyzed with aforementioned factor included as a moderating variable. It is assumed to moderate every link in the relationship between self-compassion, stigmatizing self-perceptions, and body-checking. The conceptual graph of this model is presented in figure D.2.6 in Appendix D.2. The moderator is encoded with the default effect coding scheme.

Results The R commands for loading the data, as well as building and running a Normal model for the body shame as an outcome variable are presented in the code snippet below.

```
R> library(BANOVA)
R> data_huelemann <- read.csv("Huellemann_data.csv")
R> #Build a model for the outcome variable
R> model.1       <- BANOVA.model("Normal", single_level = T)
R> Normal.single <- BANOVA.build(model.1)
R> #Run the analysis of the outcome variable
R> out2.1 <- BANOVA.run(body_checking ~ self_compassion + weight_bias + self_objectification + body_shame,
                        fit = Normal.single, data = data_huelemann, id = "id", iter = 10000,
                        chains = 2, cores = 2, seed = 123)
```

Code snippet 6.3: Building and running a model for the outcome variable.

Table 6.3 contains point estimates, 95% CI, and Bayesian p-values obtained with after fitting the regression. In the Appendix D.2 a full summary of the results is presented in code snippet D.4, including outcomes of the Geweke and the Heidelberg & Welch convergence diagnostics. These tests indicate that the chains have converged to the target densities. From table 6.3 it is evident that there is no direct association between self-compassion and body checking ($p = 0.856$), when controlling for stigmatizing perceptions. However, there is very strong evidence for a positive association between each of the stigmatizing beliefs and body checking behavior. When controlling for all other stigmatizing perceptions one-point increase in the weight bias internalization score, body shame score, and self-objectification score are associated with an increase in the body checking score by 0.21 ($p < 0.0001$), 0.37

Regressor	**point estimate**	**95% CI**	**p-value**
intercept	2.25	(2.17, 2.33)	<0.0001
self-compassion	0.01	(-0.14, 0.17)	0.8556
weight bias internalization	0.21	(0.14, 0.29)	<0.0001
self-objectification	0.37	(0.22, 0.51)	<0.0001
body shame	0.23	(0.10, 0.35)	0.0002

Table 6.3: Results for the outcome model in the study by Huellemann and Calogero (2020).

($p < 0.0001$), and 0.23 ($p = 0.0002$) respectively.

To model the relationship between correlated mediators and self-compassion I have loaded, compiled, and run a Multivariate Normal model, as presented in code snippet 6.4. A full summary with Geweke and the Heidelberg & Welch convergence diagnostics, BANOVA analysis, correlation matrix, and tables of prediction is presented in code snippet D.5 in the Appendix D.2. As all chains have converged the presented outcomes can be interpreted further. Table 6.4 includes point estimates, 95% CI, and Bayesian p-values for each of the mediators in a single table. The results indicate that there exists very strong evidence for a negative association between self-compassion and all of the stigmatizing beliefs. One point increase in the self-compassion score is on associated with an average decrease of 0.93 ($p < 0.0001$) in the weight bias internalization score, 0.42 ($p < 0.0001$) reduction in the self-objectification score, and 0.70 ($p < 0.0001$) decline in the body shame.

```
R> #Build a model for the multiple mediatirs
R> model.2 <- BANOVA.model("multiNormal", single_level = T)
R> multiNormal.single <- BANOVA.build(model.2)
R> #Prepare mediators
R> mediators           <- cbind(data_huelemann$weight_bias,
                             data_huelemann$self_objectification,
                             data_huelemann$body_shame)
R> colnames(mediators) <- c("weight_bias", "self_objectification", "body_shame")
R> data_huelemann$mediators      <- mediators
R> #Run the analysis of the mediators
R> out2.2 <- BANOVA.run(mediators ~ self_compassion, fit = multiNormal.single, iter = 10000,
                    data = data_huelemann, id = "id", chains = 2, cores = 2, seed = 123)
```

Code snippet 6.4: Building and running the model for the mediator variables.

Dependent variable	**Regressor**	**point estimate**	**95% CI**	**p-value**
Weight bias internalization	*Intercept*	3.37	(3.17, 3.57)	<0.0001
	Self-compassion	-0.93	(-1.23, -0.64)	<0.0001
Self-objectification	*Intercept*	2.94	(2.85, 3.03)	<0.0001
	Self-compassion	-0.42	(-0.56, -0.28)	<0.0001
Body shame	*Intercept*	1.99	(1.88, 2.11)	<0.0001
	Self-compassion	-0.70	(-0.86, -0.53)	<0.0001

Table 6.4: Results for the model with mediators as dependent variables in the study by Huellemann and Calogero (2020).

To analyze the mediation model proposed in the study the fitted models for the outcome variable (`out2.1`) and the mediators (`out2.2`) were combined in the `BANOVA.multi.mediation()`. The function call and resulting console output are presented in code snippet 6.5. The output displays direct effects of self-compassion and stigmatizing beliefs on the body shame, as well as the effect of self-compassion on mediators. Next specific indirect effects of each mediator are presented to the user. It is apparent that there is very strong evidence for the negative indirect effect of self-compassion on body checking through each of the mediators. The strongest effect is observed via weight bias internalization with a size of -0.20 ($p < 0.0001$), followed by the indirect effect via body shame -0.16 ($p < 0.0001$), and self-objectification -0.13 ($p < 0.0001$). Finally, the total indirect effect of the self-compassion via all of the mediators in the study is reported to be -0.51 ($p < 0.0001$). Note that since the data comes from an observational study the effects are correlational, and not causal.

Code snippet 6.5: Function call of `BANOVA.multi.mediation()` and the corresponding console output.

```
R> mediated_effects <- BANOVA.multi.mediation(sol_1 = out2.1, sol_2 = out2.2, xvar = "self_compassion",
                                            mediators = c("weight_bias", "self_objectification","body_shame"))
----------------------------------------------------------------------------------------------------
Direct effects of the causal variable self_compassion on the outcome variable
   mean    2.5%   97.5%
 0.0139 -0.1378  0.1660
----------------------------------------------------------------------------------------------------
Direct effects of mediators weight_bias and self_objectification and body_shame on the outcome variable
Direct effects of weight_bias
  mean   2.5%  97.5%
0.2137 0.1363 0.2873

Direct effects of self_objectification
  mean   2.5%  97.5%
0.3656 0.2150 0.5143

Direct effects of body_shame
  mean   2.5%  97.5%
0.2310 0.1032 0.3578
----------------------------------------------------------------------------------------------------
Direct effects of the causal variable self_compassion on the mediator variables
Direct effects of self_compassion on weight_bias
   mean    2.5%   97.5%
-0.9341 -1.2346 -0.6417

Direct effects of self_compassion on self_objectification
   mean    2.5%   97.5%
-0.4209 -0.5608 -0.2824

Direct effects of self_compassion on body_shame
   mean    2.5%   97.5%
-0.6967 -0.8645 -0.5285
----------------------------------------------------------------------------------------------------
Indirect effects of the causal variable self_compassion on the outcome variables
Indirect effects of self_compassion via weight_bias
    mean    2.5%   97.5% p.value
 -0.1996 -0.3042 -0.1127 <0.0001

Indirect effects of self_compassion via self_objectification
    mean   2.5%    97.5% p.value
 -0.1539 -0.242 -0.0804 <0.0001

Indirect effects of self_compassion via body_shame
    mean    2.5%   97.5% p.value
 -0.1609 -0.2653 -0.0693 <0.0001
```

```
------------------------------------------------------------------------------------------------------
Total indirect effects of the causal variable self_compassion on the outcome variables
    mean    2.5%   97.5% p.value
 -0.5144 -0.6677 -0.3753 <0.0001
```

A side by side comparison of the mediation analysis obtained with BANOVA and frequentist approach used in the original study are presented in table 6.5. The table compares point estimates and 95% CIs of the effects of self-compassion. The p-values were not presented in the original study, as the authors used bias-corrected 95% confidence intervals to assess the significance of the effects. It is apparent from the table that results obtained with both analytic techniques are very similar.

Effect of self-compassion	**Results of BANOVA**		**Original results**	
	point estimate	***95% CI***	***point estimate***	***95% CI****
Direct effect	0.01	(-0.14, 0.17)	0.01	(-0.14, 0.16)
Total indirect effect	-0.51	(-0.67, -0.38)	-0.51	(-0.68, -0.37)
Indirect effect via weight bias internalization	-0.20	(-0.30, -0.11)	-0.20	(-0.31, -0.11)
Indirect effect via self-objectification	-0.15	(-0.24, -0.08)	-0.15	(-0.25, -0.08)
Indirect effect via body shame	-0.16	(-0.26, -0.07)	-0.16	(-0.28, -0.05)

* frequentist confidence interval

Table 6.5: Comparison of the results in the Huellemann and Calogero (2020) study and the reanalysis with BANOVA.

Next, I discuss the results of a model with socioeconomic status as a moderating variable. The code for running the analysis along with the corresponding console output is presented in code snippets D.6-D.8 in Appendix D.2. Table 6.6 presents indirect effects of self-compassion via stigmatizing beliefs for different groups in terms of socioeconomic status. From the table we can see that strong evidence for the negative indirect association of self-compassion and body checking via weight bias internalization and self-objectification is observed only for participans in the middle ($p = 0.0004, p = 0.0044$) and upper ($p = 0.0034, p < 0.0001$) socioeconomic categories. Furthermore, the indirect association of self-compassion and body checking is transmitted via body shame only for women in the upper socioeconomic category ($p = 0.0002$). The combination of these results indicates that there is very strong evidence for an indirect association between self-compassion and body checking transmitted via all three stigmatizing self-perceptions observed at the middle ($p < 0.0001$) and upper ($p < 0.0001$) levels of socioeconomic status. Furthermore, the effect appears to be the largest for the participants with the highest level of socioeconomic status.

Conclusion This application illustrates how to estimate mediated effects in models with multiple parallel mediators. The results indicate that there exists an indirect negative association between self-compassion and body checking behavior. The absence of a direct association between these two variables suggests that the effect is facilitated through a decrease in stigmatizing self-perceptions, specifically weight bias internalization, self-objectification, and body shame.

Effect of self-compassion	**Lower SES**		**Middle SES**		**Upper SES**	
	point estimate	*95% CI*	*point estimate*	*95% CI*	*point estimate*	*95% CI*
Direct effect	0.17	(-0.37, 0.72)	-0.20	(-0.42, 0.01)	0.24	(-0.01, 0.49)
Total IE	-0.36	(-1.02, 0.18)	-0.35	(-0.56, -0.16)	-0.72	(-1.01, -0.46)
IE via weight bias internalization	-0.06	(-0.38, 0.15)	-0.21	(-0.36, -0.08)	-0.20	(-0.38, -0.06)
IE via self-objectification	-0.08	(-0.43, 0.16)	-0.12	(-0.23, -0.03)	-0.25	(-0.42, -0.11)
IE via body shame	-0.22	(-0.81, 0.19)	-0.02	(-0.16, 0.12)	-0.26	(-0.43, -0.12)

Note: IE = indirect effect

Table 6.6: Results in the Huellemann and Calogero (2020) study with a socioeconomic status as a moderator.

The reanalysis yielded virtually same estimates of the direct and indirect effects of self-compassion as in the original study. This was to be expected, as analysis of between-subjects designs with relatively large sample size and approximately Normal dependent variables yield the same results in Bayesian and frequentist approaches. However, in BANOVA no bias correction was necessary to yield the correct CIs.

Furthermore, the data was reanalyzed with socioeconomic status included as a moderating variable. The outcomes of the model with moderated mediation support the postulated hypothesis and indicate, that the transmission of the indirect association between self-compassion and body checking happens only for participants in the middle and upper levels of socioeconomic statuses. This outcome offers further insights into the relationship between self-compassion and body checking and demonstrates the versatility of BANOVA in the domain of mediation and moderation analysis.

7 Discussion and Conclusion

The goal of this thesis was to extend the functionality of the BANOVA R-package in the domains of moderation and mediation analyses. Therefore, I have derived and implemented calculations of Bayesian simple effects and mediation analysis with multiple correlated mediators. The latter sub-project required me to develop a model with a Multivariate Normal dependent variable, which can also be used as a stand-alone multivariate regression. Considering the popularity of these methods, as well as the growing interest in Bayesian techniques among social scientists, these extensions are necessary for BANOVA to achieve its goal of being a go-to package for the analysis of experiments in the social sciences. Newly developed functions and their descriptions make both practical and theoretic contributions. Naturally, the practical value comes from the possibility for the students, academics, and practitioners to perform an advanced data analysis and derive insights, without extensive experience with programming or Bayesian statistics. The theoretical contribution comes in the form of a discussion of simple effects and mediation effects in the context of Bayesian hierarchical models, which to the best of my knowledge have not been developed before.

Furthermore, several other secondary goals were achieved throughout this project. First of all, this paper presents a detailed description of BANOVA, along with a review of relevant concepts in Bayesian inference, which are useful for the users who are interested in an in-depth understanding of the analysis performed by the package. Secondly, I demonstrate how to use the new functions and interpret their results by applying them to data sets from previously published studies. The calculation of the Bayesian simple effects is illustrated based on experimental research with repeated measures of Perfecto et al. (2019). Furthermore, I reanalyzed the data from the study of Huellemann and Calogero (2020) to illustrate the calculation of the Bayesian mediated effects transmitted through multiple mediators in a non-experimental, cross-sectional setting. Extending the latter study, an additional model with moderated mediation was analyzed. This reanalysis offers a deeper insight into the nature of the observed phenomenon and illustrates the flexibility of BANOVA in the domain of mediation and moderation. Lastly, I have contributed to the quality and variety of the statistical analysis performed in BANOVA, by correcting several issues and bugs and extending the original code.

The new functions for mediation and moderation have several limitations, which should be addressed in the future. Currently, the calculation of Bayesian simple effects can only be performed with a single moderating variable as the simple-effect basis. In studies that investigate interactions between three and more variables, it would be useful to allow for multiple moderators to be considered as bases in the calculation of the simple effects. In the scope of mediation analysis, several possible improvements could be made. First of all, mediators are currently assumed to be Normal. As often categorical variables play a mediating role in social sciences experiments, future efforts should be directed towards relaxing this assumption. Secondly, the function evaluates the conditional effects of an independent variable, which provides only correlational evidence for the investigated mediation model. With the growing popularity of causal studies, it is worthwhile to develop an option to obtain causal effects of an independent variable via evaluation of its marginal direct and indirect effect, as performed in the mediation R-package (Tingley et al., 2014).

Furthermore, to assure that BANOVA is an industry-standard in the domain of Bayesian data analysis for social sciences, several other extensions of the package should be made in the future. First of all, user-specified priors are of foremost importance, as they will allow researchers to explicitly incorporate their beliefs about the investigated models. Furthermore, a function for mediation with multiple sequential mediators should be developed to offer a full spectrum of mediation techniques for users. Lastly, although in this thesis BANOVA was extended to include an analysis of truncated Normal variables, there is a need for a wider range of distributions of the dependent variable that could be analyzed with BANOVA. Namely, models with Exponential, Gamma, and Beta distributed dependent variables, as well as time-to-event and time-series models should be developed in the future versions of the package.

I sincerely hope that this thesis has made a significant contribution to the development of the BANOVA R-package, and consequentially contribute to research in different areas of social sciences, for example, psychology, consumer research, management research, and political science.

References

Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing*, 18(4):343–373.

Barnard, J., McCulloch, R., and Meng, X.-L. (2000). Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica*, 10(4):1281–1311.

Baron, R. M. and Kenny, D. A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6):1173–1182.

Bauer, D. J. and Curran, P. J. (2005). Probing interactions in fixed and multilevel regression: Inferential and graphical techniques. *Multivariate Behavioral Research*, 40(3):373–400.

Ben-Aakiva, M. and Lerman, S. (1985). *Discrete choice analysis: theory and application to travel demand*. The MIT Press.

Betancourt, M. (2017). A conceptual introduction to hamiltonian monte carlo. *arXiv preprint arXiv:1701.02434*.

Betancourt, M., Byrne, S., and Girolami, M. (2014). Optimizing the integrator step size for hamiltonian monte carlo. *arXiv preprint arXiv:1411.6669*.

Betancourt, M., Byrne, S., Livingstone, S., Girolami, M., et al. (2017). The geometric foundations of hamiltonian monte carlo. *Bernoulli*, 23(4A):2257–2298.

Bollen, K. A. and Stine, R. (1990). Direct and indirect effects: Classical and bootstrap estimates of variability. *Sociological Methodology*, 20:115–140.

Brooks, S. and Gelman, A. (1998). Some issues for monitoring convergence of iterative simulations. *Computing Science and Statistics*, pages 30–36.

Burstein, L. (1980). Chapter 4: The analysis of multilevel data in educational research and evaluation. *Review of Research in Education*, 8(1):158–233.

Cameron, A. C. and Trivedi, P. K. (2005). *Microeconometrics: methods and applications*. Cambridge University Press.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., and Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1).

Cohen, J. (1994). The earth is round (p<.05). *American Psychologist*, 49(12):997–1003.

Cohen, J., Cohen, P., West, S. G., and Aiken, L. S. (2013). *Applied multiple regression/correlation analysis for the behavioral sciences*. Routledge.

Cohen, S. and Wills, T. A. (1985). Stress, social support, and the buffering hypothesis. *Psychological Bulletin*, 98(2):310–357.

Cowles, M. K. and Carlin, B. P. (1996). Markov chain monte carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91(434):883–904.

Dixson, D. D., Keltner, D., Worrell, F. C., and Mello, Z. (2018). The magic of hope: Hope mediates the relationship between socioeconomic status and academic achievement. *The Journal of Educational Research*, 111(4):507–515.

Donaldson, T. S. (1968). Robustness of the F-test to errors of both kinds and the correlation between the numerator and denominator of the F-ratio. *Journal of the American Statistical Association*, 63(322):660–676.

Dong, C. and Wedel, M. (2017). BANOVA: an R package for hierarchical bayesian ANOVA. *Journal of Statistical Software*, 81(1):1–46.

Fisher, R. A. (1921). On the "probable error" of a coefficient of correlation deduced from a small sample. *Metron*, 1:3–32.

Fisher, R. A. (1934). *Statistical Methods for Research Workers*. Edinburgh: Oliver and Boyd.

Fox, J. (1997). *Applied regression analysis, linear models, and related methods.* Sage Publications, Inc.

Gelman, A. (2006). Prior distribution. *Encyclopedia of econometrics*, 4.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian data analysis*. CRC Press.

Gelman, A. and Hill, J. (2007). *Data analysis using regression and multilevel hierarchical models*, volume 1. Cambridge University Press.

Geweke, J. et al. (1991). *Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments*, volume 196. Federal Reserve Bank of Minneapolis, Research Department Minneapolis, MN.

Gill, J. (2014). *Bayesian methods: A social and behavioral sciences approach*, volume 20. CRC Press.

Guttinger, S. and Love, A. C. (2019). Characterizing scientific failure: Putting the replication crisis in context. *EMBO Reports*, 20(9):e48765.

Hayes, A. F. (2017). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. Guilford Publications.

Heidelberger, P. and Welch, P. D. (1983). Simulation run length control in the presence of an initial transient. *Operations Research*, 31(6):1109–1144.

Hoffman, M. D. and Gelman, A. (2014). The No-U-Turn sampler: Adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, 15(1):1593–1623.

Huellemann, K. L. and Calogero, R. M. (2020). Self-compassion and body checking among women: the mediating role of stigmatizing self-perceptions. *Mindfulness*, 11:2121–2130.

Imai, K., Keele, L., and Tingley, D. (2010a). A general approach to causal mediation analysis. *Psychological Methods*, 15(4):309–334.

Imai, K., Keele, L., and Yamamoto, T. (2010b). Identification, inference and sensitivity analysis for causal mediation effects. *Statistical Science*, 25(1):51–71.

Jaccard, J., Wan, C. K., and Turrisi, R. (1990). The detection and interpretation of interaction effects between continuous variables in multiple regression. *Multivariate Behavioral Research*, 25(4):467–478.

Jackson, C. H., Jit, M., Sharples, L. D., and De Angelis, D. (2015). Calibration of complex models through bayesian evidence synthesis: a demonstration and tutorial. *Medical Decision Making*, 35(2):148–161.

Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, 59(4):434–446.

Jeffreys, H. (1998). *The theory of probability*. OUP Oxford.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795.

Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and anovas. *Frontiers in Psychology*, 4:863.

Lewandowski, D., Kurowicka, D., and Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, 100(9):1989–2001.

Link, W. A. and Eaton, M. J. (2012). On thinning of chains in MCMC. *Methods in Ecology and Evolution*, 3(1):112–115.

MacKinnon, D. P. (2008). *Introduction to statistical mediation analysis*. Routledge.

MacKinnon, D. P., Fairchild, A. J., and Fritz, M. S. (2007). Mediation analysis. *Annu. Rev. Psychol.*, 58:593–614.

MacKinnon, D. P., Valente, M. J., and Gonzalez, O. (2020). The correspondence between causal and traditional mediation analysis: the link is the mediator by treatment interaction. *Prevention Science*, 21(2):147–157.

MacKinnon, D. P., Warsi, G., and Dwyer, J. H. (1995). A simulation study of mediated effect measures. *Multivariate Behavioral Research*, 30(1):41–62.

Marsman, M. and Wagenmakers, E.-J. (2017). Three insights from a bayesian interpretation of the one-sided p value. *Educational and Psychological Measurement*, 77(3):529–539.

McCullagh, P. and Nelder, J. (1989). *Generalized Linear Models, Second Edition*. Chapman and Hall/CRC Monographs on Statistics and Applied Probability Series. Chapman & Hall.

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092.

Metropolis, N. and Ulam, S. (1949). The monte carlo method. *Journal of the American Statistical Association*, 44(247):335–341.

Nakagawa, S. and Schielzeth, H. (2013). A general and simple method for obtaining R2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 4(2):133–142.

Neal, R. M. et al. (2011). MCMC using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11):2.

Nelder, J. A. and Wedderburn, R. W. (1972). Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3):370–384.

Pearl, J. (2001). Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pages 411–420.

Pearl, J. (2012). The causal mediation formulaâa guide to the assessment of pathways and mechanisms. *Prevention Science*, 13(4):426–436.

Pearson, E. S. (1931). The analysis of variance in cases of non-normal variation. *Biometrika*, 23(1):114–133.

Perfecto, H., Donnelly, K., and Critcher, C. R. (2019). Volume estimation through mental simulation. *Psychological Science*, 30(1):80–91.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701.

Scheffé, H. (1959). The analysis of variance. *New York: John Wiley & Sons*, pages 25–50.

Sobel, M. E. (1982). Asymptotic confidence intervals for indirect effects in structural equation models. *Sociological Methodology*, 13:290–312.

Spiller, S. A., Fitzsimons, G. J., Lynch Jr, J. G., and McClelland, G. H. (2013). Spotlights, floodlights, and the magic number zero: Simple effects tests in moderated regression. *Journal of Marketing Research*, 50(2):277–288.

Stan Development Team (2020). Stan reference manual, version 2.24. `https://mc-stan.org/docs/2_24/reference-manual/index.html` [Accessed: 20-09-2020].

The Comprehensive R Archive Network (2020). Checklist for cran submissions. `https://cran.r-project.org/web/packages/submission_checklist.html` [Accessed: 24-11-2020].

Tingley, D., Yamamoto, T., Hirose, K., Keele, L., and Imai, K. (2014). Mediation: R package for causal mediation analysis. *Journal of Statistical Software*, 59(5):1–38.

Valente, M. J., Pelham III, W. E., Smyth, H., and MacKinnon, D. P. (2017). Confounding in statistical mediation analysis: What it is and how to address it. *Journal of Counseling Psychology*, 64(6):659–671.

VanderWeele, T. (2015). *Explanation in causal inference: methods for mediation and interaction*. Oxford University Press.

Walker, D. C., White, E. K., and Srinivasan, V. J. (2018). A meta-analysis of the relationships between body checking, body image avoidance, body image dissatisfaction, mood, and disordered eating. *International Journal of Eating Disorders*, 51(8):745–770.

Wedel, M. and Dong, C. (2019). Banova: Bayesian analysis of experiments in consumer psychology. *Journal of Consumer Psychology*, 30(1):3–23.

Wendorf, C. A. (2004). Primer on multiple regression coding: Common forms and the additional case of repeated contrasts. *Understanding Statistics*, 3(1):47–57.

Woodward, K., McIlwain, D., and Mond, J. (2019). Feelings about the self and body in eating disturbances: The role of internalized shame, self-esteem, externalized self-perceptions, and body shame. *Self and Identity*, 18(2):159–182.

Yates, F. (1934). The analysis of multiple classifications with unequal numbers in the different classes. *Journal of the American Statistical Association*, 29(185):51–66.

Yuan, Y. and MacKinnon, D. P. (2009). Bayesian mediation analysis. *Psychological Methods*, 14(4):301–322.

Zellner, A. (1962). An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association*, 57(298):348–368.

Zhang, J., Wedel, M., and Pieters, R. (2009). Sales effects of attention to feature advertisements: a bayesian mediation analysis. *Journal of Marketing Research*, 46(5):669–681.

A Markov Chain Monte Carlo methods

A.1 Metropolis algorithms

A basic building block of Bayesian estimation techniques is the Metropolis (Metropolis et al., 1953) algorithm. It ensures a convergence to the target distribution through adaptation of a random walk with a rejection rule. The algorithm initiates the chains with values $\boldsymbol{\theta}_0$ from a starting distribution $p(\boldsymbol{\theta}_0)$ and proceeds in an iterative manner. During the mth iteration it samples $\boldsymbol{\theta}^*$ from a symmetric proposal distribution $q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{m-1})$, which is accepted as the next value $\boldsymbol{\theta}^m$ with probability

$$\alpha = \min(1, r), \text{ where } r = \frac{p(\boldsymbol{\theta}^*|\boldsymbol{y})}{p(\boldsymbol{\theta}^{m-1}|\boldsymbol{y})}. \tag{A.1}$$

Since $p(\cdot|\boldsymbol{y})$ is the target posterior distribution the rule in equation A.1 implies that the proposed value $\boldsymbol{\theta}^*$ is always accepted if it increases the posterior density, otherwise it is accepted with probability equal to the density ratio r. If the proposed $\boldsymbol{\theta}^*$ is rejected $\boldsymbol{\theta}^{m-1}$ is set as $\boldsymbol{\theta}^m$.

In its simplest form this method and it's generalization the Metropolis-Hastings algorithm, despite their immense popularity, have several issues. In practice, it is difficult to find an appropriate candidate distribution $q(\cdot)$, which can result in a failure of the sampler. Secondly, there exists the so-called autocorrelation issue, which is the within-sequence correlation between the samples drawn from an MCMC chain. Most importantly their random walk behavior decreases the efficiency of the chain mixing and speed of convergence to the target distribution.

A.2 Hamiltonian Monte Carlo and No-U-Turn Sampler

Also known as hybrid Monte Carlo, the Hamiltonian Monte Carlo (HMC) algorithm is a recent generalization of the Metropolis algorithm, which uses Hamiltonian dynamics, in which both momentum and position in a given system evolve in continuous time, to sample draws from a posterior distribution (Neal et al., 2011). In other words, HMC uses an auxiliary momentum variable $\boldsymbol{\gamma}$ of the same dimension as $\boldsymbol{\theta}$, which helps to achieve a faster mixing of the chains, in particular in high dimensional problems. The momentum variable $\boldsymbol{\gamma}$, which largely determines the proposal distribution $q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{m-1})$, is commonly assumed to be independent of $\boldsymbol{\theta}$ and to follow a multivariate Normal distribution with mean zero and a symmetric positive definite covariance matrix $\mathbf{M}$. The negative logarithm of their joint posterior distribution $p(\boldsymbol{\gamma}, \boldsymbol{\theta}|\boldsymbol{y}) = p(\boldsymbol{\gamma})p(\boldsymbol{\theta}|\boldsymbol{y})$ defines an invariant Hamiltonian function $H(\boldsymbol{\gamma}, \boldsymbol{\theta}|\boldsymbol{y})$, which captures the geometry of the typical set of the density. The typical set indicates an area in a density, which has the highest contribution to an evaluated integral and is characterized by a large density and volume of associated high dimensional distribution (Betancourt, 2017).

To propose a new $\boldsymbol{\theta}^*$ HMC simulates the Hamiltonian dynamics, by solving Hamiltonian equations. These differential equations are obtained by taking a gradient of $H(\boldsymbol{\gamma}, \boldsymbol{\theta}|\boldsymbol{y})$ with respect to continuous time t. This approach allows the HMC to create a trajectory for an efficient exploration of the parameter space of the target posterior distribution (Betancourt, 2017). During each iteration m HMC starts by drawing a new value of $\boldsymbol{\gamma}^m$, independent across iterations.

Next the solution to Hamiltonian equations are commonly approximated by a so-called leapfrog integrator. In each iteration this methods uses L steps to simultaneously update $\boldsymbol{\theta}$ and $\boldsymbol{\gamma}$ through a discrete approximation of Hamiltonian equations using a step size ϵ, by alternating between the half-step updates of the momentum and full-step updates of the position. The starting values for the leapfrog integrator $\boldsymbol{\gamma}_0$ and $\boldsymbol{\theta}_0$ are $\boldsymbol{\gamma}^m$ and $\boldsymbol{\theta}^{m-1}$ respectively. The L leapfrog steps of the integrator proceed via the following scheme:

$$\boldsymbol{\gamma}_{t+\frac{\epsilon}{2}} = \boldsymbol{\gamma}_t + \frac{\epsilon}{2}\frac{d\log p(\boldsymbol{\theta}_t|\boldsymbol{y})}{d\boldsymbol{\theta}}, \quad \boldsymbol{\theta}_{t+\epsilon} = \boldsymbol{\theta}_t + \epsilon\mathbf{M}^{-1}\boldsymbol{\gamma}_{t+\frac{\epsilon}{2}}, \quad \boldsymbol{\gamma}_{t+\epsilon} = \boldsymbol{\gamma}_{t+\frac{\epsilon}{2}} + \frac{\epsilon}{2}\frac{d\log p(\boldsymbol{\theta}_{t+\epsilon}|\boldsymbol{y})}{d\boldsymbol{\theta}}. \tag{A.2}$$

At the end of the L steps values $\boldsymbol{\theta}^* = \boldsymbol{\theta}_{L\epsilon}$ and $\boldsymbol{\gamma}^* = \boldsymbol{\gamma}_{L\epsilon}$ are generated, but $\boldsymbol{\gamma}^*$ is discarded as it is not of interest in itself. Since the leapfrog integrator is not perfect numerically it is necessary to apply Metropolis acceptance step for the proposed value to account for possible errors during integration. Then $\boldsymbol{\theta}^*$ is accepted as $\boldsymbol{\theta}^m$ based on the rule presented in equation A.1, with

$$r = \frac{p(\boldsymbol{\theta}^*|\boldsymbol{y})p(\boldsymbol{\gamma}^*)}{p(\boldsymbol{\theta}^{m-1}|\boldsymbol{y})p(\boldsymbol{\gamma}^{m-1})}. \tag{A.3}$$

The sampling efficiency of the HMC algorithm depends on several parameters, namely the number of leapfrog steps L, the leapfrog step size ϵ, and the covariance matrix $\mathbf{M}$ in the probability distribution for the momentum variable $\boldsymbol{\gamma}$. The number of steps L must be carefully selected to trade off the undesirable random walk behavior corresponding to a low value of the parameter, and a long-running time of the algorithm which comes with a large value of L. If the value of ε is large the proposals $\boldsymbol{\theta}^*$ is inaccurate and a corresponding acceptance rate r is small, and when the value is small the leapfrog will require too many leapfrog steps L. The goal of the parameter tuning is to obtain a rate of r in a 0.6-0.9 interval, which Betancourt et al. (2014) showed to be theoretically optimal. It is desired that the scale of $\mathbf{M}$ to approximate the inverse covariance matrix of the posterior distribution $p(\boldsymbol{\theta}|\boldsymbol{y})$ (Betancourt et al., 2017). Since this is not known before the start of estimation, $\mathbf{M}$ can either be selected to be assumed to be constant and estimated along with other parameters. However, a poorly suited covariance matrix $\mathbf{M}$ requires a decrease in the step size ε and an increase in the number of steps L.

These parameters are the main determinants of the speed of the iterations and their sampling effectiveness (Neal et al., 2011). While it is possible to use adaptive MCMC techniques, for example, see a review by Andrieu and Thoms (2008), they usually require multiple computationally costly tuning runs. The solution in the form of the No-U-Turn Sampler (NUTS) was presented by Hoffman and Gelman (2014), which is at least as efficient as a well tuned traditional HMC. This algorithm eliminates the need to select parameter L and comes along with adaptive tuning of ϵ and $\mathbf{M}$, which is performed during the warm-up and used for all sampling iterations. The estimation begins in the same way as HMC, but during the leapfrog discretization step, a sophisticated tree building algorithm is used. The length of the trajectory of the space exploration in the NUTS algorithm is accumulated during each iteration m by doubling the value of parameter L. This growth proceeds until the trajectory diverges by making a "u-turn", or reaches the maximum number of doublings, referred to as maximum tree depth. The "u-turn" in the trajectory is reflected in an infinite value of the Hamiltonian function. Further details of this recursive algorithm are beyond the scope of this paper, and an interested reader should refer to the original paper by Hoffman and Gelman (2014).

B Causal mediation analysis

The potential outcome framework, developed by Rubin (1974), is based on a notion of a binary treatment variable $\boldsymbol{t}$, which takes a value of 1 or 0, depending on whether a participant in the study received the treatment or not. When i.i.d observations of the dependent variable collected in the vector $\boldsymbol{y}$ are considered, the potential outcomes for the individual i is labeled as $y_i(t_i)$, where $t_i = 0, 1$ and $i = 1, \ldots, n$. A realization of the outcome variable for this individual is defined as $y_i = y_i(t_i) = t_i y_i(1) + (1 - t_i) y_i(0)$, which implies that only one of the potential outcomes can be observed. The unobserved potential outcomes are referred to as counterfactual. In this setting the treatment effect for the individual i, denoted by $y_i(1) - y_i(0)$, can not be observed. Instead, this technique focuses on estimation of average treatment effect, defined as $\mathrm{E}(y_i(1) - y_i(0))$ taken with respect to the random sample from the population of interest.

Under a random assignment of treatment to the participants in the study, the potential outcomes are said to be statistically independent of the treatment, which is denoted as $\big(y_i(1), y_i(0)\big) \perp\!\!\!\perp t_i$. This is referred to as ignorability or unconfoundedness assumption (A1), which assures those potential outcomes are comparable across treatment and control groups, and thus the average difference in outcomes is due to the causal variable. In case this assumption is satisfied the average causal effect is calculated as the difference in the observed means of the treatment and control groups, $\mathrm{E}(y_i|t_i = 1) - \mathrm{E}(y_i|t_i = 0)$ (VanderWeele, 2015, p. 460). This derivation requires also an assumption of no interference between the subject in the study, such that the potential outcomes for an individual do not depend on t_j, for $j \neq i$ (Imai et al., 2010a).

To define causal mediation effects in this framework the mediator $\boldsymbol{m}$ also has two potential outcomes. For each individual in the study i the observed value of mediator is $m_i = m_i(t_i)$. Thus, the potential outcomes of the dependent variable are a function of treatment and mediator, and for the i_{th} observation they are defined as $y_i(t_i, m_i)$, with the observed outcome equal to $y_i(t_i, m_i(t_i))$. The assumption of no interference between the participants in the study generalizes to imply that the potential values of mediator for each individual are independent of the treatment status of other individuals, and his or her potential outcomes do not depend on the values of mediators and treatment of other participants in study (Imai et al., 2010a).

Based on the counterfactual framework, Pearl (2001) defines the total effect, as well as the natural direct and indirect effects of a causal variable. The total effect (TE) is the overall effect of the treatment. The natural direct effect (NDE) is the difference in the outcome of the dependent variable when the mediator is held constant at the value which corresponds to the actual treatment condition, but the treatment status is changed. The natural indirect effect (NIE) is the effect of the causal variable on the outcome which occurs via the mediator. In other words, NIE measures the change in outcome when the treatment status is held constant, but the value of the mediator under the control condition is changed to the values under the treatment condition. For the indirect effect to be different from zero, the causal variable must influence the mediator, the change in which must propagate to the change in the outcome. These effects are defined at the unit level, but since only one of the potential outcomes is observed for each individual the average population effects, defined in equations B.1 - B.3, are estimated instead. The NDE and NIE depend on the treatment status t_i, which means the effects may differ for different levels of the causal variable. These definitions allow us to decompose the total effect on natural direct and indirect effects without making any assumptions about the model specification. Furthermore, these definitions can be generalized for non-dichotomous causal variables, by setting the t_i to any of the possible treatment levels.

$$TE = \mathrm{E}\big(y_i(1, m_i(1)) - y_i(0, m_i(0))\big) \tag{B.1}$$

$$NDE(t_i) = \mathrm{E}\big(y_i(1, m_i(t_i)) - y_i(0, m_i(t_i))\big), \text{ where } t_i = 0, 1 \tag{B.2}$$

$$NIE(t_i) = \mathrm{E}\big(y_i(t_i, m_i(1)) - y_i(t_i, m_i(0))\big), \text{ where } t_i = 0, 1 \tag{B.3}$$

Similar to the traditional mediation techniques, models with mediation, where both outcome and mediator are linear, can be estimated via SEM in the potential outcome framework. Unlike the traditional method, the interaction between the causal variable and the mediator must be included in the estimation, and contrasts (dummy coding) should be used to encode a non-continuous causal variable. MacKinnon et al. (2020) showed that simple indirect effects in the treatment and control groups correspond to NIE(1) and NIE(0), while simple direct effects in the treatment and control groups are equivalent to NDE(1) and NDE(0) respectively. However, this correspondence seizes to exist in models with non-linear relationships and not Normal outcomes and mediators. Imai et al. (2010a) proposed a general algorithm for estimation in mediation analysis applicable to parametric or non-parametric models, under "sequential ignorability" assumptions described below. This technique is suitable for dependent variables and mediators fitted with generalized additive and linear models, as well as quantile and survival regressions.

Several assumptions must be fulfilled to assign a causal interpretation to the effects in the potential outcome framework. This interpretation arises only if the parameters, which represent the effects of interest, are identified. To identify total effects the ignorability assumption (A1) mentioned above must be fulfilled. In observational studies, this assumption is often violated, so instead, a strong ignorability is assumed (A1.2). Under this assumption the potential outcomes are independent within a strata of covariates $\mathbf{X}$, such that $\big(y_i(1), y_i(0)\big) \perp\!\!\!\perp t_i|\boldsymbol{x_i}$. Indirect and direct effects require additional assumptions referred to as "sequential ignorability" (A2) (Imai et al., 2010b). This implies that treatment is statistically independent of potential outcomes and potential mediators, such that $\big(y_i(t_i, m_i), m_i(t_i)\big) \perp\!\!\!\perp t_i|\boldsymbol{x_i}$ (A2.1). Furthermore, the mediator, conditional on treatment and covariates, must be statistically independent of the potential outcome $y_i(t_i, m_i) \perp\!\!\!\perp, m_i|, t_i, \boldsymbol{x_i}$ (A2.2). The latter assumption implies that there must not be any covariates omitted from the model that affect both $\boldsymbol{m}$ and $\boldsymbol{y}$. Imai et al. (2010a) proved that "sequential ignorability" assumption is equivalent to the independence between error terms in equations for mediator (4.8) and dependent variable (4.6, 4.7). The first condition (A2.1) can be satisfied by randomly assigning subjects in the study to different treatment levels or via conditioning on the observed covariates. The assumption A2.2 cannot be assured unless the mediator levels are random assigned to the participants. Imai et al. (2010b) proposed a formal sensitivity analysis to examines the sensitivity of the results to violation of the identification assumptions.

C Appendix for Contributions to the BANOVA R-package

C.1 Documentation for `BANOVA.simple()` and `BANOVA.multi.mediation()`

`BANOVA.simple` *Simple effects calculation*

Description

`BANOVA.simple` is a function for probing interaction effects in models where both moderator and explanatory variables are factors with an arbitrary number of levels. The function estimates and tests simple or partial effects, also known as simple main or conditional effects. Both single-level and multi-level models with any of the distributions accommodated in the package can be analyzed.

Usage

```
BANOVA.simple(BANOVA_output, base = NULL, quantiles = c(0.025, 0.975),
dep_var_name = NULL, return_posterior_samples = FALSE)
```

Arguments

`BANOVA_output`	an object of class "BANOVA" returned by BANOVA.run function with an outcome of the hierarchical Bayesian ANOVA analysis.
`base`	a character string which specifies the name of the mediator variable used as a base for calculation.
`quantiles`	a numeric vector with quantiles for the posterior interval of the simple effects. Must include two elements with values between 0 and 1 in ascending order, default c(0.025, 0.975)
`dep_var_name`	a character string with a name of the dependent variable, for the Multinomial model only, default NULL.
`return_posterior_samples`	logical indicator of whether samples of the posterior simple effects distributions should be returned, default `FALSE`.

Details

The function identifies all factors and their combinations that are interacting with a moderating of "base" variable. For each interaction, it determines all possible level combinations of the involved regressors, which are further used to combine the posterior samples of the selected regression coefficients to calculate simple effects.

When the default effect coding scheme is used the simple effects are calculated for all levels of the interacting variables, as specified in the data. If a user specifies different contrasts for any of the interacting variables the simple effects for these variables are reported for the user-defined regressors. This distinction is reflected in the labels of the reported results: in the default case labels from the original factors are displayed; in the case of user-defined contrasts, the name of the regressor is displayed instead.

The summary of the posterior distribution of each simple effect contains the mean, standard deviation, posterior interval, which by default reports a central 95% interval, but can also be specified by the user, and a two-sided Bayesian p-value.

Note that for a Multinomial model intercepts and between-subject regressors have choice specific coefficients and thus simple effects are reported for each possible choice outcome. To perform the calculation for a Multinomial model an additional argument `dep_var_name` with a name of the dependent variable must be specified.

Value

Returns a list with the summary tables of the results; optionally returns the samples drawn from the posterior simple effects distributions.

`results_summary`	a list of tables with summaries of the posterior simple effects distributions for all factors and their combinations that are interacting with a moderating variable.
`samples_simple_effects`	if `return_posterior_samples` is set to `TRUE` a list of tables with samples of the posterior simple effects is returned. The tables include results for all levels of all factors and their combinations that are interacting with a moderating variable.

Author(s)

Anna Kopyakova

Examples

```
# Use the colorad data set
data(colorad)

# Build and analyze the model
model <- BANOVA.model('Binomial')
banova_model <- BANOVA.build(model)
res_1 <- BANOVA.run(y ~ typic, ~ color*blurfac, fit = banova_model,
                    data = colorad, id = 'id', num_trials = as.integer(16),
```

```
                    iter = 2000, thin = 1, chains = 2)
# Calculate simple effects with "blurfac" as a moderating vriable
simple_effects <- BANOVA.simple(BANOVA_output = res_1, base = "blurfac")
```

`BANOVA.multi.mediation`

Mediation analysis with multiple possibly correlated mediators

---

Description

`BANOVA.multi.mediation` is a function for analysis of multiple possibly correlated mediators. These mediators are assumed to have no causal influence on each other. Both single-level and multi-level models can be analyzed.

Usage

```
BANOVA.multi.mediation(sol_1, sol_2, xvar, mediators, individual = FALSE)
```

Arguments

`sol_1`	an object of class "BANOVA" returned by BANOVA.run function with a fitted model for an outcome variable regressed on a causal variable, a mediator, and, possibly, moderators and control variables. The outcome variable can follow Normal, T, Poisson, Bernoulli, Binomial, Truncated Normal, and ordered Multinomial distributions.
`sol_2`	an object of class "BANOVA" returned by BANOVA.run function, which contains an outcome of the analysis for multiple Multivariate Normal mediators regressed on a casual variable and other possible moderators and control variables.
`xvar`	a character string that specifies the name of the causal variable used in both models.
`mediators`	a vector with character strings, which specifies the names of the mediator variables used in the models.
`individual`	logical indicator of whether to output effects for individual units in the analysis (TRUE or FALSE). This analysis requires a multilevel `sol_1`.

Details

The function extends `BANOVA.mediation` to the case with multiple possibly correlated mediators. For details about mediation analysis performed in BANOVA see the help page for the BANOVA.mediation.

`BANOVA.multi.mediation` estimates and tests specific indirect effects of the causal variable conveyed through each mediator. Furthermore, the total indirect effect of the causal variables are computed as a sum of the specific indirect effects.

The function prints multiple tables with mediated effects. Tables with direct effects of the causal variable and mediators on the outcome variable, as well as direct effects of the causal variable on the mediators include a posterior mean and 95% credible intervals of the effects. Next, the function displays on the console tables with specific indirect effects and effect sizes of the mediators, followed by the TIE of the causal variable. These tables include the mean, 95% credible intervals, and two-sided Bayesian p-values.

Value

Returns an object of class `"BANOVA.multi.mediation"`. The returned object is a list containing:

`dir_effects`	table or tables with the direct effect.
`individual_direct`	is returned if `individual` is set to `TRUE` and the causal variable is a within-subject variable. Contains a table or tables of the direct effect at the individual levels of the analysis
`m1_effects`	a list with tables of the effects of the mediator on the outcome
`m2_effects`	a list with tables of the effect of the causal variable on the mediator
`indir_effects`	tables of the indirect effect
`individual_indirect`	is returned if `individual` is set to `TRUE` and the mediator is a within-subject variable. Contains the table or tables with the indirect effect
`effect_sizes`	a list with effect sizes on individual mediators
`total_indir_effects`	table or tables with the total indirect effect of the causal variable
`xvar`	the name of the causal variable
`mediators`	the names of the mediating variables
`individual`	the value of the argument individual (TRUE or FALSE)

Author(s)

Anna Kopyakova

Examples

```
# Use the colorad data set
data(colorad)
# Add a second mediator to the data set
colorad$blur_squared <- (colorad$blur)^2
# Prepare mediators to be analyzed in the Multivariate Normal model
mediators <- cbind(colorad$blur, colorad$blur_squared)
colnames(mediators) <- c("blur", "blur_squared")
colorad$mediators <- mediators

# Build and analyze the model for the outcome variable
model <- BANOVA.model('Binomial')
banova_binom_model <- BANOVA.build(model)
res_1 <- BANOVA.run(y ~ typic, ~ color + blur + blur_squared, fit = banova_binom_model,
                    data = colorad, id = 'id', num_trials = as.integer(16),
                    iter = 2000, thin = 1, chains = 2)
# Build and analyze the model for the mediators
model <- BANOVA.model('multiNormal')
banova_multi_norm_model <- BANOVA.build(model)
res_2 <- BANOVA.run(mediators ~ typic, ~ color, fit = banova_multi_norm_model,
                    data = colorad, id = 'id', iter = 2000, thin = 1, chains = 2)

# Calculate (moderated) effects of "typic" mediated by "blur" and "blur_squared"
results <- BANOVA.multi.mediation(res_1, res_2, xvar='typic',
                                  mediators=c("blur", "blur_squared"))
```

C.2 Stan code for Multivariate Bayesian Hierarchical Models

Code snippet C.1: Stan code for single-level SUR models

```
data {
  int<lower=0> N; // number of observations
  int<lower=1> L; // number of dependent variables
  int<lower=0> J; // number of subject-level effects
  vector[J] X[N]; // array of size N which contains J-long vectors
  vector[L] y[N]; // response variable
}
parameters {
  matrix[L, J] beta1;              // subject-level effects
  cholesky_factor_corr[L] L_Omega; // Cholesky factor of the correlation matrix
  vector<lower=0>[L] L_sigma;      // standard Deviations of y
}
transformed parameters {
  vector[L] y_hat[N];        //mean of the distribution of y
  matrix[L, L] L_Sigma;     // covariance matrix
   for (n in 1:N) {
     for (l in 1:L) {
       y_hat[n, l] = dot_product(X[n], beta1[l]);
     }
   }
   L_Sigma = diag_pre_multiply(L_sigma, L_Omega);
}
model {
  //Priors
  to_vector(beta1) ~ normal(0, 100);
  L_Omega ~ lkj_corr_cholesky(1);
  L_sigma ~ cauchy(0, 2.5);

  y ~ multi_normal_cholesky(y_hat, L_Sigma);
}
generated quantities {
  real var_f[L];
  real r_2[L];
  matrix[L, L] Omega; // Correlation matrix
  matrix[L, L] Sigma; // Covariance matrix
  Omega = multiply_lower_tri_self_transpose(L_Omega);
  Sigma = quad_form_diag(Omega, L_sigma);
  for (l in 1:L){
    var_f[l] = variance(y_hat[,l]);
    r_2[l] = var_f[l]/(var_f[l] + Sigma[l,l]);
  }
}
```

Code snippet C.2: Stan code for multilevel SUR models

```
data {
  int<lower=0> N; // number of observations
  int<lower=1> L; // number of dependent variables
  int<lower=0> J; // number of subject-level effects
  int<lower=0> M; // number of unique subjects
  int<lower=0> K; // number of population-level effects
  vector[J] X[N]; // subject level regressors
  vector[K] Z[M]; // population level regressors
  vector[L] y[N]; // response variables
  int<lower=0> id[N];
}
parameters {
  matrix[L, J] beta1[M]; // subject-level effects
  matrix[K, J] beta2[L]; // population-level effects
  cholesky_factor_corr[L] L_Omega;   // Cholesky factor of the correlation matrix of y
  vector<lower=0>[L] L_sigma;         // standard Deviations of y
  matrix<lower=0>[L, J] tau_beta1Sq; // variance of beta1
}
transformed parameters {
  vector[L] y_hat[N];              // means for the distribution in y
  matrix[L, J] mu_beta1[M];        // means of elements in beta1
  vector[L] y_pred[N];             // predictions for the dependent variables in y
  matrix[L, L] L_Sigma;            // Choleski decomposition of the covariance matrix
  matrix<lower=0>[L, J] tau_beta1; // standard deviation for beta1
  for (n in 1:N) {
     for (l in 1:L) {
       y_hat[n, l] = dot_product(X[n], beta1[id[n],l]);
     }
   }
  for (m in 1:M){
    for (j in 1:J){
       for (l in 1:L){
         mu_beta1[m, l, j] = dot_product(Z[m], beta2[l, 1:K, j]);
      }
    }
  }
  for (n in 1:N) {
    for (l in 1:L) {
      y_pred[n, l] = dot_product(X[n], mu_beta1[id[n],l]);
    }
  }
  L_Sigma = diag_pre_multiply(L_sigma, L_Omega);
  tau_beta1 = sqrt(tau_beta1Sq);
}
model {
  // Priors
  for (m in 1:M){
    for (j in 1:J){
       for (l in 1:L){
         beta1[m, l, j] ~ normal(mu_beta1[m, l, j], tau_beta1[l,j]);
      }
    }
  }
```

```
  L_Omega ~ lkj_corr_cholesky(1);
  L_sigma ~ cauchy(0, 2.5);
  to_vector(tau_beta1Sq) ~ inv_gamma(1, 1);
  for(l in 1:L){
    to_vector(beta2[l]) ~ normal(0, 100);
  }
  for (n in 1:N){
    y[n] ~ multi_normal_cholesky(y_hat[n], L_Sigma);
  }
}
generated quantities {
  real r_2[L];
  matrix[L, L] Omega; // Correlation matrix
  matrix[L, L] Sigma; // Covariance matrix
  Omega = multiply_lower_tri_self_transpose(L_Omega);
  Sigma = quad_form_diag(Omega, L_sigma);
  for (l in 1:L){
    r_2[l] =  variance(y_pred[1:N, l])/(variance(y_hat[1:N, l]) + Sigma[l,l]);
  }
}
```

C.3 Other contributions

Correct labels for effect coded variables. The BANOVA R-package automatically converts factor variables as effect coded regressors. However, an issue occurred if a given explanatory variable is already encoded as (1), (0), and (-1) and passed as a numeric variable in the data. In this case, when `BANOVA.run()` creates effect coded regressors, it uses ascending order to encode the levels, such that the original value of (-1) is the first level and (1) is the last level. While the encoding does not influence estimation, it matters for the interpretation of the coefficients. Users unaware of the switch would interpret the coefficients as the difference in the effects between a given level and the last one, originally encoded as (-1), while it is the difference between a given level and the first one, originally encoded as (1). A similar issue occurred with numeric variables encoded as (0) and (1), as 0 is treated as the first level. To solve this issue I have implemented an internal auxiliary function, which evaluates the encoding of numeric variables. In the variable is effect or dummy coded the levels are assigned in descending order, following the expectations of the user. This issue has occurred in other functions in BANOVA, which I fixed as well.

Fixing bugs in `BANOVA.mediation()`. Several issues were fixed in the `BANOVA.mediation()` function.

- The function did not work for multilevel models with a between-subject mediating variable.
- If an interaction between a causal variable and the mediator is included in the model the moderated effects were not accurately computed for different levels of the moderating factor. This issue also occurred if mediating and causal variables were at different levels (subject-level or population-level) in the regression model.
- In some edge cases effects for the individual units in the analysis could not be computed due to an error in an internal auxiliary function.

Another issue in this function could not be classified as a bug, but still required an improvement. As described in section 4.2.2, in models with multiple moderators of causal or mediating variables, multiple tables with results are reported to the user. However, there was no indication of which moderating variables were used in the estimation of the results. I have implemented a new subroutine in the `BANOVA.mediation()` and its printing methods, such that each reported table has a clear name in the list and a corresponding title.

Truncated Normal model in BANOVA. To broaden the functionality of the BANOVA, I extended the package with a hierarchical regression model for a dependent variable that follows a Truncated Normal distribution. The values of a truncated variable are observed only if they lie within a certain range. This extension was first developed in Stan (Stan Development Team, 2020, Chapter 11) and then integrated into the `BANOVA.run()` function. To run an analysis with Truncated Normal dependent variable a user should load the model with `"truncNormal"` argument and compile it using the approach described in section 3.4. Next, the analysis can be run in the same manner as for a Normal model, with two additional parameters. Arguments `y_lowerBound` and `y_upperBound`, which are minus and plus infinity by default, should be used to specify the desired lower, upper, or both bounds for the dependent variable.

D Appendix for Applications

D.1 Perfecto, Donnelly, and Critcher (2019)

Code snippet D.1: Sample of the data

```
R> head(data_perfecto)
  id estimate simulation orientation shape
1  1     13.9          1           1     0
2  1     14.0          1           1     0
3  1     14.0          1          -1     0
4  1     14.1          1          -1     0
5  1     14.6          1           1     0
6  1     14.8          1           1    -1
```

Figure D.1.1: Visualization of the estimated cup volume

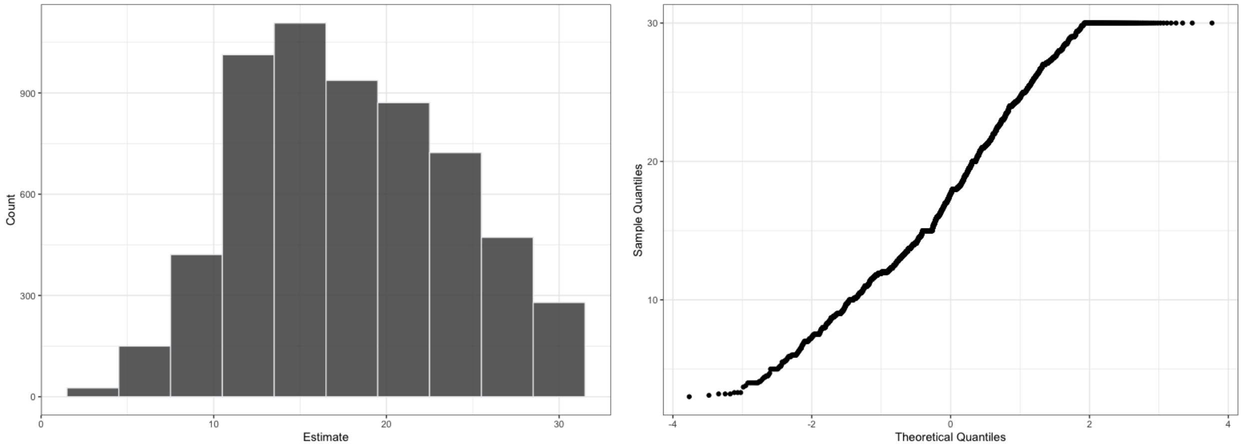

(a) Histogram

(b) Q-Q plot

Code snippet D.2: Summary of Bayesian hierarchical analysis of estimated cup volume

```
R> summary(out1)
> summary(out1)
Call:
BANOVA.run(l1_formula = estimate ~ orientation * shape, l2_formula = ~simulation, fit = Normal,
           data = data_perfecto, id = "id", iter = 30000, thin = 3, chains = 2, cores = 2, seed = 123,
           contrast = list(orientation = c(1/2, -1/2),
                           shape = cbind(c(1/2, 0, -1/2), c(-1/3, 2/3, -1/3))))

Convergence diagnostics:
Geweke Diag. & Heidelberger and Welch’s Diag.
                                    Geweke stationarity test Geweke convergence p value
(Intercept)  :  (Intercept)                           passed                     0.1767
(Intercept)  :  simulation1                           passed                     0.6629
orientation1  :  (Intercept)                          passed                     0.2037
orientation1  :  simulation1                          passed                     0.1087
shape1  :  (Intercept)                                passed                     0.0064
```

```
shape1  :  simulation1                                passed                     0.6349
shape2  :  (Intercept)                                passed                     0.6798
shape2  :  simulation1                                passed                     0.0030
orientation1:shape1  :  (Intercept)                   passed                     0.8852
orientation1:shape1  :  simulation1                   passed                     0.2383
orientation1:shape2  :  (Intercept)                   passed                     0.0361
orientation1:shape2  :  simulation1                   passed                     0.4275
                                    H. & W. stationarity test H. & W. convergence p value
(Intercept)  :  (Intercept)                            passed                      0.1268
(Intercept)  :  simulation1                            passed                      0.9708
orientation1  :  (Intercept)                           passed                      0.4102
orientation1  :  simulation1                           passed                      0.3081
shape1  :  (Intercept)                                 passed                      0.3473
shape1  :  simulation1                                 passed                      0.6106
shape2  :  (Intercept)                                 passed                      0.4322
shape2  :  simulation1                                 passed                      0.2062
orientation1:shape1  :  (Intercept)                    passed                      0.6209
orientation1:shape1  :  simulation1                    passed                      0.1534
orientation1:shape2  :  (Intercept)                    passed                      0.1216
orientation1:shape2  :  simulation1                    passed                      0.7401

The Chain has converged.

Table of sum of squares & effect sizes:

Table of sum of squares:
                    (Intercept) simulation Residuals      Total
(Intercept)          82405.6303     0.0000 1635.8037 84039.7911
orientation1            50.1112    46.2688   53.3956   149.7755
shape1                 387.3813   132.7645   63.8619   584.0077
shape2                 573.6218    24.6531   80.3173   678.5922
orientation1:shape1    244.7460   212.0537  129.1387   585.9385
orientation1:shape2     69.1553    82.9514  109.8860   261.9928

Table of effect sizes (95% credible interval):
                             (Intercept)            simulation
(Intercept)           0.9024 (0.9,0.905) 0.0000 (-0.003,0.002)
orientation1        0.0068 (0.001,0.016)  0.0063 (0.001,0.015)
shape1              0.0499 (0.027,0.077)  0.0177 (0.005,0.036)
shape2                0.0721 (0.047,0.1)      0.0033 (0,0.011)
orientation1:shape1 0.0315 (0.003,0.079)   0.0275 (0.001,0.07)
orientation1:shape2     0.0092 (0,0.035)      0.0110 (0,0.037)

Table of p-values (Multidimensional):
                  (Intercept) simulation
(Intercept)           <0.0001     0.4978
orientation            0.0012     0.0062
shape                 <0.0001    <0.0001
orientation:shape      0.0040     0.0142

Table of coefficients:
                                       mean     SD Quantile0.025 Quantile0.975 p.value Signif.codes
(Intercept)  :  (Intercept)         17.9151 0.1726       17.5874       18.2579 <0.0001          ***
(Intercept)  :  simulation1          0.1183 0.1750       -0.2220        0.4558  0.4978
```

```
orientation1  :  (Intercept)        0.4481 0.1403        0.1743        0.7267  0.0012           **
orientation1  :  simulation1        0.4042 0.1412        0.1227        0.6787  0.0062           **
shape1  :  (Intercept)              1.2673 0.1712        0.9244        1.6026 <0.0001          ***
shape1  :  simulation1              0.7030 0.1718        0.3675        1.0387 <0.0001          ***
shape2  :  (Intercept)             -1.5065 0.1529       -1.8064       -1.2062 <0.0001          ***
shape2  :  simulation1             -0.2770 0.1498       -0.5699        0.0116  0.0610            .
orientation1:shape1  :  (Intercept)  0.9771 0.3380        0.3171        1.6321  0.0040           **
orientation1:shape1  :  simulation1  0.8522 0.3326        0.1991        1.5001  0.0142            *
orientation1:shape2  :  (Intercept) -0.4577 0.3015       -1.0675        0.1248  0.1288
orientation1:shape2  :  simulation1 -0.4894 0.3001       -1.0690        0.1050  0.1070
---
Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1

Multiple R-squared:  0.0305

Table of predictions:

Grand mean:
17.9151
    2.5%   97.5%
 17.5874 18.2579

 orientation mean    2.5%    97.5%
 1           18.1391 17.7818 18.5078
 -1          17.691  17.3288 18.0673

 shape mean    2.5%    97.5%
 1     19.0509 18.6585 19.4435
 0     16.9107 16.5251 17.3137
 -1    17.7836 17.4018 18.1696

 simulation mean    2.5%    97.5%
 1          18.0334 17.5292 18.526
 -1         17.7967 17.3362 18.2585

 orientation simulation mean    2.5%    97.5%
 1           1          18.4595 17.9173 18.9901
 1           -1         17.8186 17.3162 18.3164
 -1          1          17.6072 17.0652 18.1463
 -1          -1         17.7748 17.2756 18.2814

 shape simulation mean    2.5%    97.5%
 1     1          19.613  19.0492 20.182
 1     -1         18.4887 17.9565 19.0224
 0     1          16.8444 16.2598 17.4151
 0     -1         16.977  16.4438 17.5217
 -1    1          17.6428 17.0664 18.2158
 -1    -1         17.9244 17.3954 18.4632

 orientation shape mean    2.5%    97.5%
 1           1     19.5954 19.137  20.0506
 -1          1     18.5063 18.0478 18.962
 1           0     16.9822 16.5272 17.4507
 -1          0     16.8392 16.3806 17.3175
```

```
1           -1    17.8396 17.3903 18.2902
-1          -1    17.7275 17.2755 18.1806

orientation shape simulation mean    2.5%    97.5%
1           1     1          20.6543 19.9753 21.3161
1           1     -1         18.5365 17.913  19.1608
-1          1     1          18.5717 17.9034 19.2313
-1          1     -1         18.4408 17.8224 19.0619
1           0     1          16.9549 16.2646 17.6392
1           0     -1         17.0095 16.3823 17.6451
-1          0     1          16.7339 16.0574 17.3999
-1          0     -1         16.9446 16.307  17.5841
1           -1    1          17.7694 17.1087 18.4314
1           -1    -1         17.9098 17.275  18.5207
-1          -1    1          17.5161 16.8441 18.1988
-1          -1    -1         17.939  17.3204 18.5728
```

D.2 Huellemann and Calogero (2020)

Code snippet D.3: Sample of the data

```
R> head(data_huelemann)
  id self_compassion weight_bias self_objectification body_shame body_checking    ses
1  1        3.000000         4.5             3.142857   1.944444      2.173913 middle
2  2        2.000000         5.1             3.000000   1.888889      2.347826 middle
3  3        2.916667         3.7             3.428571   1.444444      1.782609  upper
4  4        2.666667         5.7             3.857143   1.777778      3.000000 middle
5  5        2.666667         1.7             2.571429   1.333333      1.347826  upper
6  6        1.583333         4.7             3.285714   4.666667      3.695652 middle
```

Figure D.2.1: Visualization of body-checking scores

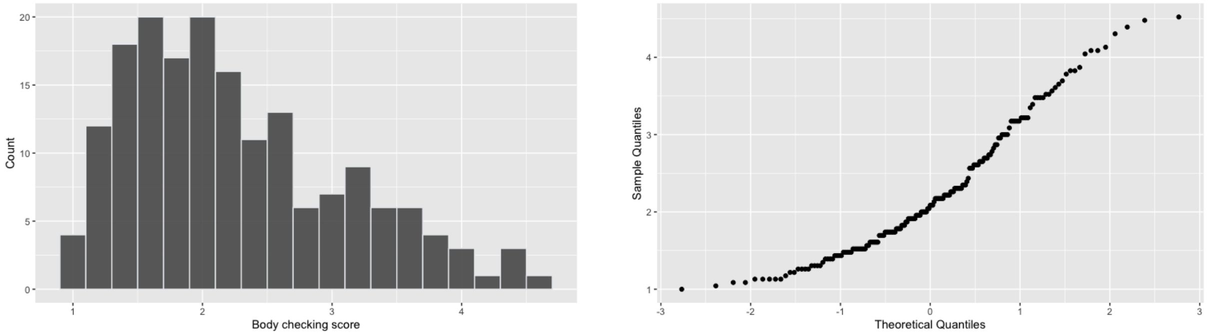

(a) Histogram

(b) Q-Q plot

Figure D.2.2: Visualization of weight bias internalization scores

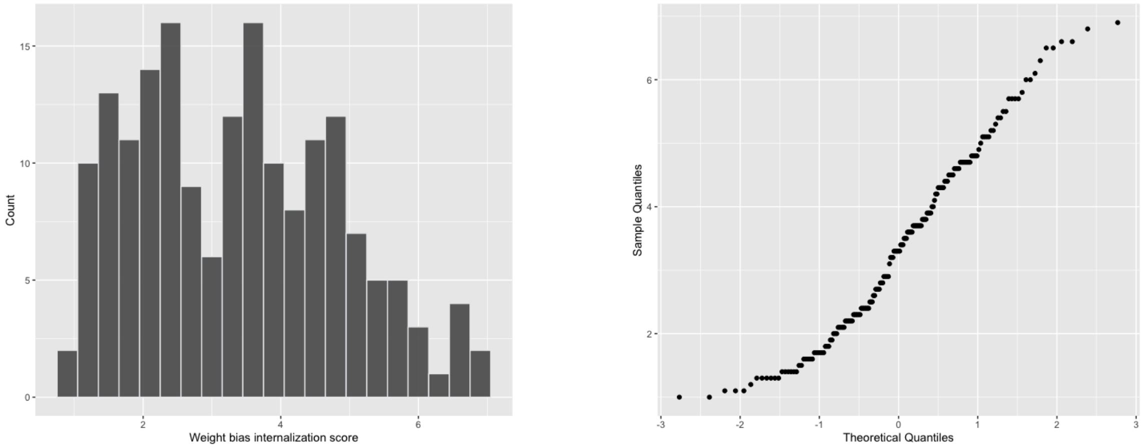

(a) Histogram

(b) Q-Q plot

Figure D.2.3: Visualization of self-objectification scores

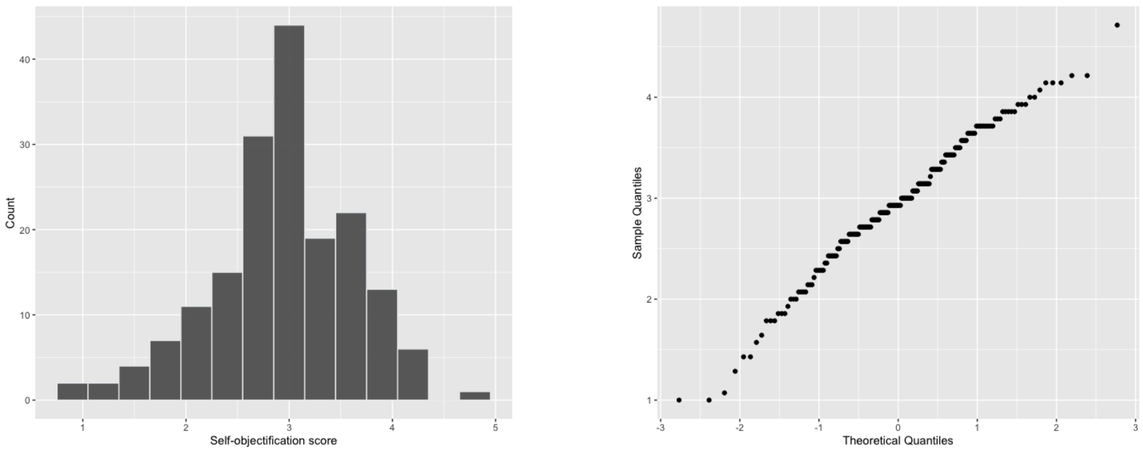

(a) Histogram

(b) Q-Q plot

Figure D.2.4: Visualization of body shame scores

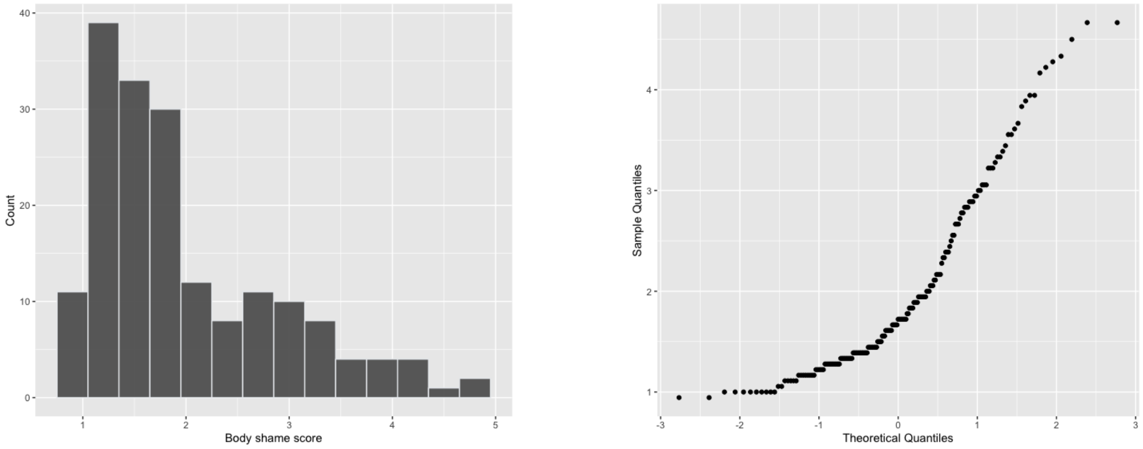

(a) Histogram

(b) Q-Q plot

Figure D.2.5: Correlation matrix of the data set

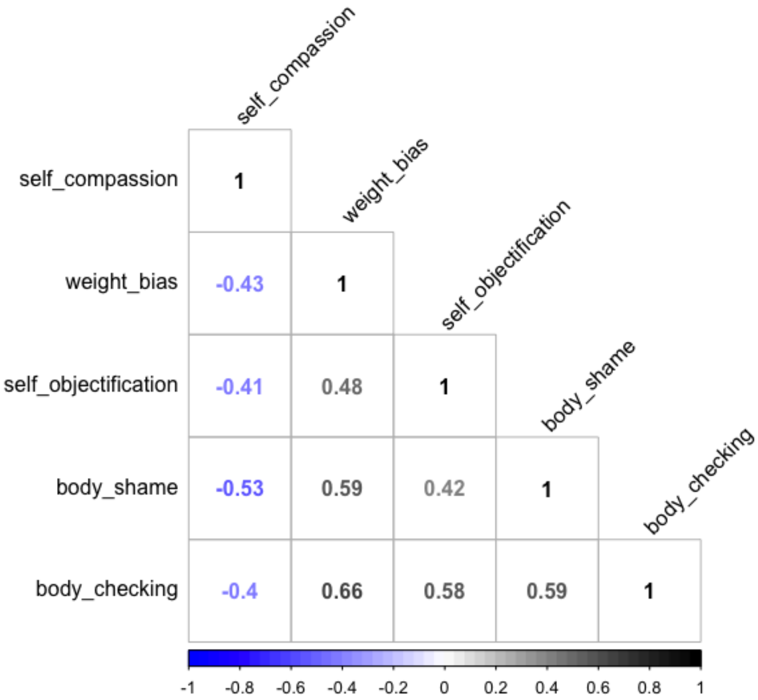

	self_compassion	weight_bias	self_objectification	body_shame	body_checking
self_compassion	1				
weight_bias	-0.43	1			
self_objectification	-0.41	0.48	1		
body_shame	-0.53	0.59	0.42	1	
body_checking	-0.4	0.66	0.58	0.59	1

Figure D.2.6: Conceptual graph a model in the study by Huellemann and Calogero (2020), where the relations between self-compassion and body checking mediated by weight-bias internalization, self-objectification, and body shame, are moderated by socioeconomic status.

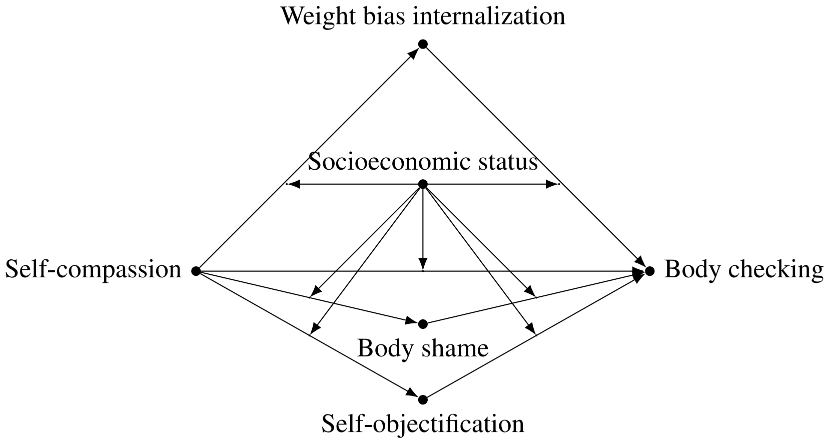

Code snippet D.4: Summary of Bayesian hierarchical analysis of body-checking as the dependent variable

```
R> summary(out2.1)
Call:
BANOVA.run(l1_formula = body_checking ~ self_compassion + weight_bias +
    self_objectification + body_shame, fit = Normal.single, data = data_huelemann,
    id = "id", iter = 10000, chains = 2, cores = 2, seed = 123)

Convergence diagnostics:
Geweke Diag. & Heidelberger and Welch’s Diag.
               Geweke stationarity test Geweke convergence p value
(Intercept)                      passed                     0.4329
self_compassion                  passed                     0.7242
weight_bias                      passed                     0.4079
self_objectification             passed                     0.2081
body_shame                       passed                     0.8000
               H. & W. stationarity test H. & W. convergence p value
(Intercept)                       passed                      0.4474
self_compassion                   passed                      0.9899
weight_bias                       passed                      0.7579
self_objectification              passed                      0.4338
body_shame                        passed                      0.5577
The Chain has converged.

Table of sum of squares & effect sizes:

Table of sum of squares:
  (Intercept) self_compassion  weight_bias   self_objectification body_shame Residuals     Total
     898.5738          0.0000      17.3952.               10.9045     7.4414   54.8267 1021.8733

Table of effect sizes (95% credible interval):
          (Intercept)      self_compassion          weight_bias self_objectification           body_shame
  0.9425 (0.94,0.944) 0.0000 (-0.027,0.012) 0.2385 (0.152,0.319) 0.1647 (0.102,0.227) 0.1177 (0.043,0.199)

Table of p-values (Multidimensional):
  (Intercept) self_compassion weight_bias self_objectification body_shame
      <0.0001          0.8556     <0.0001              <0.0001    <0.0001

Table of coefficients:
                       mean     SD Quantile0.025 Quantile0.975 p.value Signif.codes
(Intercept)          2.2520 0.0432        2.1677        2.3374 <0.0001          ***
self_compassion      0.0139 0.0778       -0.1377        0.1660  0.8556
weight_bias          0.2137 0.0386        0.1363        0.2873 <0.0001          ***
self_objectification 0.3656 0.0757        0.2152        0.5143 <0.0001          ***
body_shame           0.2310 0.0646        0.1033        0.3578  0.0002          ***
---
Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1
Multiple R-squared:  0.5545

Table of predictions:
Grand mean:
2.252
   2.5%  97.5%
 2.1677 2.3374
```

Code snippet D.5: Summary of Bayesian hierarchical analysis of weight bias internalization, self objectification, and body shame as the multivariate dependent variable

```
R> summary(out2.2)
Call:
BANOVA.run(l1_formula = mediators ~ self_compassion, fit = multiNormal.single,
    data = data_huelemann, id = "id", iter = 10000, chains = 2,
    cores = 2, seed = 123)

Convergence diagnostics:
Geweke Diag. & Heidelberger and Welch’s Diag.
                                     Geweke stationarity test Geweke convergence p value
weight_bias (Intercept)                                passed                     0.3940
weight_bias self_compassion                            passed                     0.2128
self_objectification (Intercept)                       passed                     0.8461
self_objectification self_compassion                   passed                     0.9011
body_shame (Intercept)                                 passed                     0.6133
body_shame self_compassion                             passed                     0.8091
                                     H. & W. stationarity test H. & W. convergence p value
weight_bias (Intercept)                                 passed                       0.3346
weight_bias self_compassion                             passed                       0.2863
self_objectification (Intercept)                        passed                       0.3408
self_objectification self_compassion                    passed                       0.4321
body_shame (Intercept)                                  passed                       0.6506
body_shame self_compassion                              passed                       0.6495

The Chain has converged.

Table of sum of squares & effect sizes:

Table of sum of squares:
                       (Intercept) self_compassion Residuals     Total
weight_bias              2007.7056         68.6577  315.4167 2391.7800
self_objectification     1531.5151         13.8937   69.2034 1614.6122
body_shame                702.8115         38.5680  102.0186  843.3981

Table of effect sizes (95% credible interval):
                                (Intercept)      self_compassion
weight_bias            0.8642 (0.861,0.865) 0.1788 (0.159,0.184)
self_objectification   0.9568 (0.956,0.957) 0.1672 (0.148,0.172)
body_shame             0.8732 (0.871,0.874) 0.2743 (0.257,0.279)

Table of p-values (Multidimensional):
                       (Intercept) self_compassion
weight_bias                <0.0001         <0.0001
self_objectification       <0.0001         <0.0001
body_shame                 <0.0001         <0.0001

Table of coefficients:
                                        mean     SD Quantile0.025 Quantile0.975 p.value Signif.codes
weight_bias (Intercept)               3.3665 0.1008        3.1691        3.5657 <0.0001          ***
weight_bias self_compassion          -0.9341 0.1515       -1.2341       -0.6417 <0.0001          ***
self_objectification (Intercept)      2.9413 0.0472        2.8495        3.0331 <0.0001          ***
self_objectification self_compassion -0.4209 0.0708       -0.5605       -0.2824 <0.0001          ***
```

```
body_shame (Intercept)                1.9920 0.0589        1.8751        2.1068 <0.0001          ***
body_shame self_compassion           -0.6967 0.0866       -0.8645       -0.5285 <0.0001          ***
---
Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1

Multiple R-squared:  0.1819 0.1703 0.2741

Correlation matrix:
                     weight_bias self_objectification body_shame
weight_bias               1.0000               0.3618     0.4714
self_objectification      0.3618               1.0000     0.2583
body_shame                0.4714               0.2583     1.0000
                                         Mean     SD Quantile 0.025 Quantile 0.975 p-value
Corr(weight_bias,self_objectification) 0.3618 0.0646         0.2308         0.4839 <0.0001
Corr(weight_bias,body_shame)           0.4714 0.0588         0.3493         0.5820 <0.0001
Corr(self_objectification,body_shame)  0.2583 0.0701         0.1162         0.3919  0.0002

Standard deviations of dependent variables:
                          Mean     SD Quantile 0.025 Quantile 0.975 p-value
weight_bias SD          1.8371 0.2010         1.4837         2.2794 <0.0001
self_objectification SD 0.4048 0.0440         0.3279         0.4981 <0.0001
body_shame SD           0.5961 0.0654         0.4803         0.7365 <0.0001

Table of predictions:

Predictions for weight_bias

Grand mean:
3.3665
   2.5%  97.5%
 3.1691 3.5657

Predictions for self_objectification

Grand mean:
2.9413
   2.5%  97.5%
 2.8495 3.0331

Predictions for body_shame

Grand mean:
1.992
   2.5%  97.5%
 1.8751 2.1068
```

Code snippet D.6: Summary of Bayesian hierarchical analysis of body-checking as the dependent variable and socioeconomic status as the moderator

```
R> out2.1.2 <- BANOVA.run(body_checking ~ ses*(self_compassion + weight_bias + self_objectification + body_shame),
R>                        fit = Normal.single, data = data_huelemann, id = "id",
R>                        iter = 10000, chains = 2, cores = 2, seed = 123)
R> summary(out2.1.2)
Call:
BANOVA.run(l1_formula = body_checking ~ ses * (self_compassion + weight_bias + self_objectification + body_shame),
        fit = Normal.single,  data = data_huelemann, id = "id", iter = 10000, chains = 2, cores = 2, seed = 123)
Convergence diagnostics:
Geweke Diag. & Heidelberger and Welch’s Diag.
                          Geweke stationarity test Geweke convergence p value
(Intercept)                                 passed                     0.7503
ses1                                        passed                     0.4885
ses2                                        passed                     0.8128
self_compassion                             passed                     0.8327
weight_bias                                 passed                     0.8531
self_objectification                        passed                     0.4147
body_shame                                  passed                     0.0740
ses1:self_compassion                        passed                     0.8471
ses2:self_compassion                        passed                     0.9794
ses1:weight_bias                            passed                     0.3517
ses2:weight_bias                            passed                     0.9000
ses1:self_objectification                   passed                     0.9979
ses2:self_objectification                   passed                     0.5798
ses1:body_shame                             passed                     0.0595
ses2:body_shame                             passed                     0.1489
                          H. & W. stationarity test H. & W. convergence p value
(Intercept)                                  passed                      0.2446
ses1                                         passed                      0.2607
ses2                                         passed                      0.7860
self_compassion                              passed                      0.8126
weight_bias                                  passed                      0.7581
self_objectification                         passed                      0.5775
body_shame                                   passed                      0.2249
ses1:self_compassion                         passed                      0.3904
ses2:self_compassion                         passed                      0.5572
ses1:weight_bias                             passed                      0.9052
ses2:weight_bias                             passed                      0.8470
ses1:self_objectification                    passed                      0.8583
ses2:self_objectification                    passed                      0.4377
ses1:body_shame                              passed                      0.2624
ses2:body_shame                              passed                      0.3215
The Chain has converged.

Table of sum of squares & effect sizes:

Table of sum of squares:
  (Intercept)      ses self_compassion weight_bias self_objectification body_shame ses:self_compassion
     898.5738 227.9532          0.6766     14.5742              13.7770    10.3023             84.3972
  ses:weight_bias ses:self_objectification ses:body_shame Residuals     Total
          15.6442                  29.8999        36.8185   53.8179 1021.8733
```

```
Table of effect sizes (95% credible interval):
           (Intercept)                  ses       self_compassion        weight_bias self_objectification
  0.9435 (0.939,0.946) 0.7484 (0.328,0.912) 0.0118 (-0.021,0.075) 0.2043 (0.05,0.37) 0.1955 (0.046,0.356)
            body_shame  ses:self_compassion       ses:weight_bias ses:self_objectification
  0.1489 (0.001,0.361) 0.5562 (0.164,0.789) 0.1967 (-0.002,0.496)     0.2965 (0.004,0.663)
        ses:body_shame
  0.3708 (0.079,0.634)

Table of coefficients:
                             mean     SD Quantile0.025 Quantile0.975 p.value Signif.codes
(Intercept)                2.1583 0.2591        1.6460        2.6733 <0.0001          ***
ses1                      -0.2142 0.8164       -1.8128        1.3860  0.7978
ses2                       1.1718 0.5545        0.0959        2.2567  0.0334            *
self_compassion            0.0676 0.1070       -0.1417        0.2772  0.5322
weight_bias                0.1900 0.0575        0.0788        0.3047  0.0008          ***
self_objectification       0.3971 0.1237        0.1552        0.6388  0.0012           **
body_shame                 0.2523 0.1175        0.0225        0.4856  0.0308            *
ses1:self_compassion       0.1025 0.1937       -0.2785        0.4867  0.5926
ses2:self_compassion      -0.2725 0.1247       -0.5210       -0.0317  0.0280            *
ses1:weight_bias          -0.0712 0.1059       -0.2802        0.1371  0.4956
ses2:weight_bias           0.0811 0.0660       -0.0512        0.2100  0.2094
ses1:self_objectification  0.0114 0.2286       -0.4360        0.4586  0.9600
ses2:self_objectification -0.0755 0.1375       -0.3515        0.1927  0.5818
ses1:body_shame            0.0901 0.2208       -0.3436        0.5331  0.6814
ses2:body_shame           -0.2242 0.1291       -0.4841        0.0309  0.0830            .
---
Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1
Multiple R-squared:  0.5822

Table of predictions:
Grand mean:
2.1583
   2.5%  97.5%
 1.6460 2.6733
 ses    mean   2.5%    97.5%
 lower  1.9441 -0.1616 4.0547
 middle 3.3301 2.5006  4.156
 upper  1.2008 0.3264  2.0578
```

Code snippet D.7: Summary of Bayesian hierarchical analysis of weight bias internalization, self objectification, and body shame as the multivariate dependent variable and socioeconomic status as the moderator

```
R> out2.2.2 <- BANOVA.run(mediators ~ self_compassion*ses, fit = multiNormal.single, iter = 10000,
R>                        data = data_huelemann, id = "id", chains = 2, cores = 2, seed = 123)
R> summary(out2.2.2)
Call:
BANOVA.run(l1_formula = mediators ~ self_compassion * ses, fit = multiNormal.single, data = data_huelemann,
            id = "id", iter = 10000, chains = 2, cores = 2, seed = 123)
Convergence diagnostics:
Geweke Diag. & Heidelberger and Welch’s Diag.
                                          Geweke stationarity test Geweke convergence p value
weight_bias (Intercept)                                     passed                     0.8781
weight_bias self_compassion                                 passed                     0.6820
weight_bias ses1                                            passed                     0.5808
weight_bias ses2                                            passed                     0.8102
weight_bias self_compassion:ses1                            passed                     0.6117
weight_bias self_compassion:ses2                            passed                     0.7387
self_objectification (Intercept)                            passed                     0.1242
self_objectification self_compassion                        passed                     0.9136
self_objectification ses1                                   passed                     0.1810
self_objectification ses2                                   passed                     0.2594
self_objectification self_compassion:ses1                   passed                     0.1808
self_objectification self_compassion:ses2                   passed                     0.2518
body_shame (Intercept)                                      passed                     0.1506
body_shame self_compassion                                  passed                     0.1442
body_shame ses1                                             passed                     0.2051
body_shame ses2                                             passed                     0.9139
body_shame self_compassion:ses1                             passed                     0.0941
body_shame self_compassion:ses2                             passed                     0.9166
                                          H. & W. stationarity test H. & W. convergence p value
weight_bias (Intercept)                                      passed                      0.2793
weight_bias self_compassion                                  passed                      0.9845
weight_bias ses1                                             passed                      0.3119
weight_bias ses2                                             passed                      0.5924
weight_bias self_compassion:ses1                             passed                      0.2680
weight_bias self_compassion:ses2                             passed                      0.4226
self_objectification (Intercept)                             passed                      0.2769
self_objectification self_compassion                         passed                      0.8712
self_objectification ses1                                    passed                      0.3287
self_objectification ses2                                    passed                      0.2457
self_objectification self_compassion:ses1                    passed                      0.3139
self_objectification self_compassion:ses2                    passed                      0.2045
body_shame (Intercept)                                       passed                      0.3018
body_shame self_compassion                                   passed                      0.4902
body_shame ses1                                              passed                      0.5335
body_shame ses2                                              passed                      0.8236
body_shame self_compassion:ses1                              passed                      0.3097
body_shame self_compassion:ses2                              passed                      0.8831

The Chain has converged.

Table of sum of squares & effect sizes:
```

```
Table of sum of squares:
                       (Intercept) self_compassion     ses self_compassion:ses Residuals     Total
weight_bias              2007.7056         52.1749 84.1314            124.2618  309.2580 2391.7800
self_objectification     1531.5151         10.9830 39.1547             31.7539   68.8339 1614.6122
body_shame                702.8115         36.7427 21.5751             21.3070  104.1691  843.3981

Table of effect sizes (95% credible interval):
                                (Intercept)      self_compassion                   ses   self_compassion:ses
weight_bias            0.8665 (0.861,0.869) 0.1424 (0.066,0.225) 0.1842 (-0.006,0.492)  0.2479 (0.002,0.564)
self_objectification   0.9570 (0.955,0.958) 0.1357 (0.057,0.218)  0.3186 (0.021,0.619)   0.2741 (0.007,0.59)
body_shame             0.8709 (0.866,0.874)  0.2582 (0.17,0.341) 0.1488 (-0.009,0.437) 0.1473 (-0.008,0.435)

Table of p-values (Multidimensional):
                       (Intercept) self_compassion    ses self_compassion:ses
weight_bias                <0.0001         <0.0001 0.4786              0.3990
self_objectification       <0.0001           4e-04 0.2240              0.3028
body_shame                 <0.0001         <0.0001 0.7402              0.7640

Table of coefficients:
                                             mean     SD Quantile0.025 Quantile0.975 p.value Signif.codes
weight_bias (Intercept)                    3.1293 0.3444        2.4627        3.8132 <0.0001          ***
weight_bias self_compassion               -0.8071 0.1986       -1.1974       -0.4141 <0.0001          ***
weight_bias ses1                          -0.6974 1.0208       -2.6835        1.3575  0.4786
weight_bias ses2                           0.0947 0.6638       -1.2275        1.3984  0.8818
weight_bias self_compassion:ses1           0.2894 0.3552       -0.4314        0.9754  0.3990
weight_bias self_compassion:ses2           0.0340 0.2333       -0.4260        0.4974  0.8964
self_objectification (Intercept)           2.7594 0.1664        2.4255        3.0868 <0.0001          ***
self_objectification self_compassion      -0.3701 0.0954       -0.5562       -0.1813  0.0004          ***
self_objectification ses1                 -0.5910 0.4900       -1.5673        0.3759  0.2240
self_objectification ses2                 -0.0158 0.3201       -0.6350        0.6162  0.9542
self_objectification self_compassion:ses1  0.1744 0.1705       -0.1606        0.5139  0.3028
self_objectification self_compassion:ses2 -0.0005 0.1122       -0.2229        0.2199  0.9964
body_shame (Intercept)                     1.9597 0.2006        1.5644        2.3448 <0.0001          ***
body_shame self_compassion                -0.6763 0.1166       -0.9055       -0.4495 <0.0001          ***
body_shame ses1                           -0.0844 0.5921       -1.2365        1.0640  0.8826
body_shame ses2                            0.1326 0.3896       -0.6321        0.8949  0.7402
body_shame self_compassion:ses1            0.0405 0.2060       -0.3612        0.4420  0.8320
body_shame self_compassion:ses2           -0.0404 0.1369       -0.3116        0.2304  0.7640
---
Signif. codes: 0 *** 0.001 ** 0.01 * 0.05 . 0.1 1

Multiple R-squared:  0.2254 0.2046 0.2833

Correlation matrix:
                     weight_bias self_objectification body_shame
weight_bias               1.0000               0.3830     0.4746
self_objectification      0.3830               1.0000     0.2689
body_shame                0.4746               0.2689     1.0000

                                         Mean     SD Quantile 0.025 Quantile 0.975 p-value
Corr(weight_bias,self_objectification) 0.3830 0.0652         0.2491         0.5064 <0.0001
Corr(weight_bias,body_shame)           0.4746 0.0600         0.3514         0.5879 <0.0001
Corr(self_objectification,body_shame)  0.2689 0.0714         0.1225         0.4045  0.0004
```

```
Standard deviations of dependent variables:
                          Mean     SD Quantile 0.025 Quantile 0.975 p-value
weight_bias SD          1.8039 0.2005         1.4526         2.2417 <0.0001
self_objectification SD 0.4025 0.0439         0.3243         0.4985 <0.0001
body_shame SD           0.6086 0.0667         0.4903         0.7511 <0.0001

Table of predictions:

Predictions for weight_bias

Grand mean:
3.1293
   2.5%  97.5%
 2.4627 3.8132


 ses    mean   2.5%    97.5%
 lower  2.4319 -0.2048 5.1283
 middle 3.2241 2.2797  4.1681
 upper  3.7319 2.7011  4.7455

Predictions for self_objectification

Grand mean:
2.7594
   2.5%  97.5%
 2.4255 3.0868


 ses    mean   2.5%   97.5%
 lower  2.1684 0.8662 3.4587
 middle 2.7436 2.3063 3.1957
 upper  3.3662 2.8874 3.8393

Predictions for body_shame

Grand mean:
1.9597
   2.5%  97.5%
 1.5644 2.3448


 ses    mean   2.5%   97.5%
 lower  1.8753 0.3398 3.399
 middle 2.0923 1.5478 2.6475
 upper  1.9114 1.3331 2.4958
```

Code snippet D.8: Function call of `BANOVA.multi.mediation()` and the corresponding console output with moderating socioeconomic status

```
----------------------------------------------------------------------------------------------------
Direct effects of the causal variable self_compassion on the outcome variable
Direct effects of self compassion moderated by ses
 ses    mean    2.5%  97.5%
   1  0.1701 -0.3727 0.7226
   2 -0.2049 -0.4178 0.0071
   3  0.2376 -0.0122 0.4889
----------------------------------------------------------------------------------------------------
Direct effects of mediators weight_bias and self_objectification and body_shame on the outcome variable
Direct effects of weight bias moderated by ses
 ses   mean    2.5%  97.5%
   1 0.1187 -0.1806 0.4229
   2 0.2711  0.1654 0.3771
   3 0.1801  0.0601    0.3

Direct effects of self objectification moderated by ses
 ses   mean    2.5%  97.5%
   1 0.4085 -0.2541 1.0666
   2 0.3216  0.1113 0.5308
   3 0.4613  0.2393 0.6821

Direct effects of body shame moderated by ses
 ses   mean    2.5%  97.5%
   1 0.3425 -0.2901 0.9854
   2 0.0282 -0.1595 0.2149
   3 0.3863  0.2123 0.5644
----------------------------------------------------------------------------------------------------
Direct effects of the causal variable self_compassion on the mediator variables
Direct effects of self compassion moderated by ses
 ses    mean    2.5%   97.5%
   1 -0.5177 -1.5283  0.4656
   2 -0.7731 -1.1901 -0.3529
   3 -1.1306 -1.6086 -0.6665

Direct effects of self compassion moderated by ses
 ses    mean    2.5%   97.5%
   1 -0.1957 -0.6804  0.2882
   2 -0.3706 -0.5662 -0.1727
   3  -0.544 -0.7595 -0.3269

Direct effects of self compassion moderated by ses
 ses    mean    2.5%   97.5%
   1 -0.6358 -1.2156 -0.0586
   2 -0.7167  -0.961 -0.4794
   3 -0.6765 -0.9478 -0.4093
----------------------------------------------------------------------------------------------------
Indirect effects of the causal variable self_compassion on the outcome variables
Indirect effects of self compassion moderated by ses through weight bias moderated by ses
 ses    mean    2.5%   97.5% p.value
   1 -0.0611 -0.3799  0.1452  0.6066
   2 -0.2095 -0.3656  -0.084   4e-04
```

```
   3 -0.2035 -0.3796 -0.0595  0.0034
effect size: 0.032 (0.001,0.105)

Indirect effects of self compassion moderated by ses through self objectification moderated by ses
 ses    mean    2.5%   97.5% p.value
   1 -0.0784  -0.432  0.1582  0.5468
   2 -0.1192  -0.233 -0.0323  0.0044
   3 -0.2509 -0.4222 -0.1092 <0.0001
effect size: 0.035 (0.002,0.121)

Indirect effects of self compassion moderated by ses through body shame moderated by ses
 ses    mean    2.5%   97.5% p.value
   1 -0.2183 -0.8288  0.1946  0.3042
   2 -0.0202 -0.1638  0.1162    0.77
   3 -0.2614 -0.4386 -0.1174   2e-04
effect size: 0.073 (0.006,0.273)
----------------------------------------------------------------------------------------------------
Total indirect effects of the causal variable self_compassion on the outcome variables
 ses    mean    2.5%   97.5% p.value
   1 -0.3578 -1.0109  0.1672  0.1862
   2 -0.3489 -0.5547 -0.1618 <0.0001
   3 -0.7158 -1.0060 -0.4539 <0.0001
```