Rethinking Robust Statistics with Modern Bayesian Methods

Donald R. Williams
University of California, Davis

Stephen R. Martin
Baylor University

Developing robust statistical methods is an important goal for psychological science. Whereas classical methods (i.e., sampling distributions, p -values, etc.) have been thoroughly characterized, Bayesian robust methods remain relatively uncommon in practice and methodological literatures. Here we propose a robust Bayesian model ($B_{HS,t}$) that accommodates heterogeneous (H) variances by predicting the scale parameter on the log scale and tail-heaviness with a Student- t likelihood (S_t). Through simulations with normative and contaminated (i.e., heavy-tailed) data, we demonstrate that $B_{HS,t}$ has consistent frequentist properties in terms of type I error, power, and mean squared error compared to three classical robust methods. With a motivating example, we illustrate Bayesian inferential methods such as approximate leave-one-out cross-validation and posterior predictive checks. We end by suggesting areas of improvement for $B_{HS,t}$ and discussing Bayesian robust methods in practice.

Keywords: Bayesian, Robust Methods, Student- t , Calibration, Tail-Heaviness

1. Introduction

An important goal for research psychologists is developing quantitative methods that are robust to violating statistical assumptions (Wilcox, 1998a, 1998b), including normality and homogeneous variances. Reports of typical psychological data have shown that satisfying these assumptions is the exception and not the rule (Erceg-Hurn & Mirosevich, 2008). For example, a review that included hundreds of datasets indicated distributions are commonly skewed or heavy-tailed (Micceri, 1989). In addition, variance ratios (VR) between groups are far from the idealized 1:1 (Erceg-Hurn & Mirosevich, 2008) and often exceed what is commonly considered problematic ($VR > 4:1$; Wilcox, Charlin, & Thompson, 1986). These findings were not restricted to certain subfields in psychology, but were found in various research areas including clinical (Grissom, 2000), educational (Keselman et al., 1998), and experimental psychology (Erceg-Hurn & Mirosevich, 2008). Although classical statistical methods can be robust to violations, in that type I error rates are at nominal levels (Ramsey, 1980), they can also show substan-

tial decreases in power. This reduction in detecting a non-zero effect is especially pronounced for heavy-tailed data (i.e., outliers present; Wilcox, 1995).

Several robust statistical methods have been developed (Wilcox, 2012). The benefits of these approaches have been thoroughly described, including in terms of improved parameter estimation (Özdemir, 2010) and optimal power (Wilcox, 2001). Yet, misconceptions remain and they are underutilized in the psychological literature (Erceg-Hurn & Mirosevich, 2008). These “modern methods” consist of some computationally intensive techniques such as the bootstrapped one-step M-estimator (Özdemir, 2013). Other robust methods do not take advantage of modern computation and directly alter the data (Wilcox & Keselman, 2002). For example, Yuen’s test for trimmed means consists of removing a predetermined proportion of data from the tail-areas. Furthermore, Yuen’s method requires Winsorizing which manually adjusts extreme values to some percentile (for computing the squared error) (for computing the squared error; Luh & Guo, 2007). This approach involves manipulating the observed data to achieve nominal error rates and to maintain power.

In the present paper, we propose a Bayesian robust model that allows for estimating heavy-tailed data with heterogeneous variances. Bayesian methods readily allow for utilizing a variety of probability distributions for estimation purposes (Gelman, Carlin, et al., 2014). The inherent flexibility of Bayesian modeling allows for self-calibration to the observed data, including data otherwise considered “outliers”. These models do not require assumptions such as normality, homogeneous variances, or distortion of the observed data (Martin & Williams, 2017). However, Bayesian robust meth-

Affiliations Donald R. Williams, Department of Psychology, University of California, Davis; Stephen R. Martin, Department of Psychology and Neuroscience, Baylor University

Acknowledgments We extend our gratitude to Paul-Christian Bürkner for providing helpful feedback.

Corresponding Author Donald R. Williams, Dept. of Psychology, University of California, Davis, One Shields Ave., Davis, CA, 95616 (drwwilliams@ucdavis.edu)

ods are not common in practice and there are few methodological papers on this topic (Kruschke, 2013; Fonseca, Ferreira, & Migon, 2008; Fernández & Steel, 1998). Importantly, frequentist properties have not been explicitly characterized in comparison to classical robust methods (those that use sampling distributions, p -values, etc) in the psychological literature.

We first describe classical robust methods in the psychological literature: Yuen's test, the one-step M-Estimator, and the modified one-step M-Estimator. Next, the proposed Bayesian robust model B_{HS_t} is introduced. This section also serves to illustrate mathematical differences between classical and Bayesian methods. We then characterize frequentist performance between B_{HS_t} compared to the classical robust methods: type I error, power, and parameter estimation (e.g., mean squared error). A motivating example is provided, in which real and artificial data is analyzed with the previously described robust methods. This section also serves to highlight inferential techniques of the proposed Bayesian model such as model comparison (Gelman, Hwang, & Vehtari, 2014; Vehtari, Gelman, & Gabry, 2016) and posterior predictive checks (Gelman, Meng, & Stern, 1996). We end by discussing modern Bayesian methods in practice, and how they can increase the use of robust statistics in psychology.

2. Description of Methods

2.1. Yuen's test

Yuen (1974) described a two-sample t -test for symmetrical data based on trimmed means that accommodates heterogeneous variances and heavy-tailed distributions. It was shown that expected error rates were achieved, in addition to optimal power compared to Welch's t -test. More recently, Yuen's test has been extended with various transformations and bootstrapping procedures (Keselman, Othman, Wilcox, & Fradette, 2004). Due to being the subject of numerous methodological papers, we describe the original procedure that is implemented in the R package *WRS2* (Mair, Schoenbrodt, & Wilcox, 2017).

Yuen's method tests the null hypothesis of equal population trimmed means between two groups $H_0 : \mu_{t1} = \mu_{t2}$. The test statistic is given by

$$t_y = \frac{\bar{X}_{t1} - \bar{X}_{t2}}{\sqrt{d_1 + d_2}} \quad (1)$$

where $\bar{X}_{tj}$ is the j th group trimmed mean ($j = 2$).

$$\bar{X}_{tj} = \frac{1}{h_j} \sum_{i=g_j+1}^{n_j-g_j} X_{(i)j}. \quad (2)$$

The effective sample size h_j is defined as $n_j - 2(g_j)$. Here $g_j = \lceil \gamma n_j \rceil$ denotes the trimmed proportion (rounded) γ from

the distribution tails. Thus, h_j is the trimmed sample size for the j th group. The sampling distribution of t_y follows the Student- t distribution, but like Welch's t -test uses an approximate degrees of freedom (Welch, 1938) defined as

$$\nu_y = \frac{(d_1 + d_2)^2}{\left(\frac{d_1^2}{h_1-1}\right) + \left(\frac{d_2^2}{h_2-1}\right)}. \quad (3)$$

in which the squared error d_j for the j th group follows

$$d_j = \frac{(n_j - 1)\hat{\sigma}_{wj}^2}{h_j(h_j - 1)}. \quad (4)$$

To compute Yuen's test statistic (1) for the difference between trimmed means, a valid estimate for the squared error must be obtained with the Winsorized variance $\hat{\sigma}_{wj}^2$

$$\hat{\sigma}_{wj}^2 = \frac{1}{n_j - 1} \sum_{i=1}^{n_j} (X_{ij} - \bar{X}_{wj})^2 \quad (5)$$

and mean $\bar{X}_{wj}$ for the j th group

$$\bar{X}_{wj} = \frac{1}{h_j} \sum_{i=1}^{n_j} X_{ij} \quad (6)$$

where

$$X_{ij} = \begin{cases} Y_{(g_j+1)j} & \text{if } Y_{ij} \leq Y_{(g_j+1)j} \\ Y_{ij} & \text{if } Y_{(g_j+1)j} < Y_{ij} < Y_{(n_j-g_j+1)j} \\ Y_{(n_j-g_j+1)j} & \text{if } Y_{ij} \geq Y_{(n_j-g_j+1)j} \end{cases} \quad (7)$$

2.2. One-step M-estimator

The one-step M-estimator overcomes at least one limitation of the trimmed means approach. Yuen's method requires selection of the trimming proportion γ . This can simultaneously provide non-optimal data reduction (Wilcox, 2012) and can increase researcher degrees of freedom in that each value of γ will produce a different p -value (Simmons, Nelson, & Simonsohn, 2011). In contrast, the M-estimation approach quantitatively determines the degree of trimming which allows for accommodating a wider class of distributions (e.g., skewed).

M-estimators are a general class of estimators in which some discrepancy function is minimized, with ordinary least squares and maximum likelihood estimators being special cases (Susanti, Pratiwi, Sulistijowati, & Liana, 2014). M-estimation utilizes a function in which the inputs are robust to extreme values. Some function $\xi(X - \mu_m)$ assesses the difference between X and an unknown constant μ_m of the j th group, where ψ is the derivative with respect to μ_m . A commonly used function for ψ follows Huber (1981)

$$\psi(x) = \max\{-k, \min(k, x)\}, k = 1.28. \quad (8)$$

A measure of location μ_m then satisfies the following

$$\sum \psi\left(\frac{X_i - \hat{\mu}_m}{\hat{\zeta}}\right) = 0 \quad (9)$$

where $\hat{\mu}_m$ minimizes the function $\xi(X - \mu_m)$. The measure of scale $\hat{\zeta}$ is computed with respect to the median absolute deviation

$$MAD = \text{median}(|X_i - M_j|) \quad (10)$$

$$\hat{\zeta} = \frac{MAD}{0.675}$$

where M_j is the median of the j th group. The one-step M-estimator is then obtained with one iteration of the Newton-Raphson method. Importantly, one iteration provides asymptotic calibration and expected frequentist properties (e.g., type I error rates) (Wilcox, 2012). The null hypothesis test is then $H_0 : \mu_{m1} = \mu_{m2}$, with the j th group M-estimator as

$$\hat{\mu}_{mj} = \frac{k\hat{\zeta}(i_2 - i_1) + \sum_{i=i_1}^{n-i_2} X_{(i)}}{n - i_1 - i_2}. \quad (11)$$

Here i_1 denotes the total observation X_{nj} for the j th group in which

$$\frac{X_{nj} - M_j}{\hat{\zeta}} < -k \quad (12)$$

whereas i_2 follows

$$\frac{X_{nj} - M_j}{\hat{\zeta}} > k. \quad (13)$$

Additionally, we used the percentile bootstrap (*PBOM*) method to obtain the standard error for $\hat{\mu}_m$ that is more accurate than other methods. This procedure is described by Özdemir (2013) and implemented in the R package *WRS2* (Mair et al., 2017).

2.3. Modified One-Step M-estimator

The modified one-step M-estimator (*MOM*) is obtained by removing $k\hat{\zeta}(i_2 - i_1)$ from equation (11). This allows for averaging observations not determined to be outliers (Wilcox, 2012). The *MOM* of $\hat{\mu}_m$ is given by

$$\hat{\mu}_{mj} = \frac{k\hat{\zeta}(i_2 - i_1) + \sum_{i=i_1}^{n-i_2} X_{(i)}}{n - i_1 - i_2}. \quad (14)$$

An efficient estimate is obtained with i_1 and i_2 that denote the total observation in which $(X_{nj} - M_j)/\hat{\zeta} < -2.24$ and $(X_{nj} - M_j)/\hat{\zeta} > 2.24$. Following Özdemir (2013), we used also used the percentile bootstrap method (*PBMOM*) that is implemented in the package *WRS2* (Mair et al., 2017).

2.4. Bayesian Heteroskedastic Student-t Regression

A Bayesian model begins with specifying a joint probability distribution for the parameters to be estimated (Gelman, Carlin, et al., 2014). For the proposed robust model—Bayesian heteroskedastic Student- t regression ($B_{HS,t}$)—Bayes theorem states that

$$p(\mu, \omega, \nu|y) \propto p(y|\mu, \omega, \nu)p(\mu, \omega, \nu) \quad (15)$$

where the posterior distribution for μ , ω , and ν , given the observed data y , is proportional ($\propto$) to the likelihood of those values $p(y|\mu, \omega, \nu)$ times the product of their prior probabilities $p(\mu, \omega, \nu)$. The regression notion for $B_{HS,t}$ follows

$$y_i \sim \text{Student-}t(\nu, \mu_i, \omega_i) \quad (16)$$

$$u_i = X\beta_\mu = \begin{bmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_n \end{bmatrix} \begin{bmatrix} \beta_{\mu 1} \\ \beta_{\mu 2} \end{bmatrix}$$

$$\omega_i = \exp(X\beta_\omega) = \exp\left(\begin{bmatrix} 1 & x_1 \\ 1 & x_2 \\ \vdots & \vdots \\ 1 & x_n \end{bmatrix} \begin{bmatrix} \beta_{\omega 1} \\ \beta_{\omega 2} \end{bmatrix}\right)$$

where the location μ and the scale parameters ω of the Student- t distribution are predicted by a $n \times 2$ matrix X . When comparing two groups, β_1 is the intercept (i.e., reference group mean) and β_2 is the difference between those groups. The exponential function $\exp(X\beta_\omega)$ ensures that the scale parameter ω_i is restricted to positive values. To be clear, this model is robust in that predicting ω allows for heterogeneous variances, while estimating ν from the data accommodates tail-heaviness by self-calibrating to the expected variance

$$\text{Var}(X) = \omega \frac{\nu}{\nu - 2}, \nu > 2. \quad (17)$$

Of course, Bayesian estimation requires specifying prior distributions for all parameters. Although in practice we prefer *weakly* informative priors (Gelman, Jakulin, Pittau, & Su, 2008; Gelman, Simpson, & Betancourt, 2017), we used so-called *non-informative* priors defined as

$$\beta_{\mu 1} \sim N(0, 100) \quad (18)$$

$$\beta_{\mu 2} \sim N(0, 100)$$

$$\beta_{\omega 1} \sim \text{Student-}t^+(3, 0, 10)$$

$$\beta_{\omega 2} \sim \text{Cauchy}^+(0, 5)$$

$$\nu \sim \text{Gamma}(2, 0.1)$$

where $+$ denotes a lower boundary of zero. Following Gelman (2006), the Student- t and Cauchy distributions were assumed as priors for the scale parameters, whereas a Gamma distribution was used for ν (Kruschke, 2013). These priors were selected to have minimal influence such that the parameter estimates were dominated by the data.

3. Simulation Study Design

Bayesian statistics have become more common in psychology (Andrews & Baguley, 2013), but remain infrequently used in applied settings compared to classical approaches (Sanborn et al., 2014). Furthermore, in the psychological literature, the dominate Bayesian approach advocates for a subjectivist perspective that includes posterior model probabilities and Bayes factors (Rouder, Morey, & Wagenmakers, 2016). While building statistical models and inference requires many subjective decisions, the approach presented here aims for well calibrated models in which examining frequentist properties allow for model assessment and improvement. Type I error is not necessarily a Bayesian inferential goal (Kruschke & Liddell, 2017), but quantifying model performance when used many times allows for inferring important aspects of the model. For example, liberal type I error rates may indicate that the posterior distribution is overstating precision (i.e., credible intervals are too narrow).

We compared the proposed Bayesian model to the classical robust methods with simulations. The first set of simulations (type I error and power) assumed normality for two groups, but allowed for heterogeneous variances. This simulation characterized performance when distributional assumptions were met, which is important to consider for robust methods. In addition, numerous studies have shown that unequal variances among groups can increase type I error rates. However, this topic has remained largely unexplored with Bayesian methods, and it is unclear whether predicting ω can provide nominal error rates. We also included a Bayesian model that assumed equal variances and normality (B_{NE}) and Welch's t -test. We assumed two group sizes (40, 60), four values for σ (1, 2, 3, 4), and two mean differences (0, 0.5). Power was expected to vary between sample sizes and variances. We thus computed Yuen's effect size y_{es} for each condition with simulations. The sampling distribution mean for y_{es} is reported with the results.

The second set of simulations assumed contaminated normal distributions for each group. This can be thought of as a scale-mixture distribution, in which the means of the mixture components are equal but the variances differed. Importantly, scale-mixtures produce heavy-tailed data and are commonly examined with robust methods. Here data were generated by sampling from a normal distribution with probability ε that has a standard deviation K_j for the j th group, while the remaining observations were sampled from

a standard normal distribution $(1 - \varepsilon)$. The expected variance for the j th group follows

$$\text{Var}(X) = 1 - \varepsilon + \varepsilon K_j^2 \quad (19)$$

In addition to type I error and power, we investigated parameter estimation (e.g., mean squared error) and coverage probabilities. We assumed two contamination values ε (0.10, 0.20), three group sizes (20, 30, 40), and two values for K (10, 20). For determining power, the contaminated distribution of one group was shifted by adding a constant to each observation ($\theta = 0.8$) to each observation. For comparison purposes Welch's t -test was also included.

All simulations were completed in R with 10,000 replications for each model and condition. The package *WRS2* was used for the classical robust methods, including: Yuen's test with two trimming proportions γ (0.10, 0.20), the one-step M-estimator, and the modified one step M-estimator. Percentile bootstrapping was used for both M-estimators ($B = 599$). The Bayesian models were fitted with the R packages *brms* (Bürkner, in press) and *rstan* (Stan Development Team, 2016) that use a modern Markov chain Monte Carlo algorithm: the Hamiltonian Monte Carlo No-U-Turn sampler (Betancourt, 2017). Each model consisted of two-chains of 500 posterior samples each (excluding a 500-iteration warm-up). The posterior distributions were summarized with medians and 95% highest density intervals (*HDI*). Before beginning the simulations, we determined sampling convergence by examining model diagnostics (e.g., $\hat{R}$ values, effective sample sizes, and visual inspection) for several conditions. Expected type I error rates and coverage probabilities were 5 % and 95 % (R code is provided online: <https://osf.io/xamuq/>)

4. Simulation Results

4.1. Normally Distributed Outcome

4.1.1 Type I Error. The results for type I error rates are provided in Table 1. For the proposed robust Bayesian model (B_{HSI}), predicting ω on the log scale showed better performance than the opposing Bayesian model (B_{NE} : assuming normality and equal variances). With a variance ratio of 16:1 ($n_1 = 20$; $n_2 = 40$), for example, the *HDI*s excluded zero in 6.5 % (B_{HSI}) and 14.6 % (B_{NE}) of simulations. The opposite pattern was observed when the larger group was more variable (i.e., B_{NE} was conservative). However, B_{HSI} was also anti-conservative compared to the classical robust methods. Type I error rates exceeded 0.06 in five of twelve conditions, whereas the classical methods (combined) exceeded 0.06 in only one condition. Importantly, the error rate for B_{HSI} was at most 0.02 greater than a classical method for the same condition (Table 1[row three]: $B_{HSI} = 0.065$; $PBMOM = 0.045$).

4.1.2. Interval Width. We also examined highest density and confidence interval width (Table 2). Here B_{NE} consistently produced the narrowest intervals when the smaller group had more variance. This precision was erroneous, as these same conditions showed substantially inflated type I error rates (Table 1). In contrast, B_{NE} produced the widest intervals when the larger group was more variable. The intervals for B_{HSI} followed the same pattern as the classical methods. For all conditions B_{HSI} provided narrower intervals than the classical methods. Although B_{HSI} had some type I error rates greater than 0.06, those closest to the nominal level ($\alpha = 0.05$) also had a narrower interval than the classical robust methods (e.g., Table 2 [row 10]).

4.1.3 Power. Table 3 shows power across a range of effect sizes (y_{es}). B_{HSI} consistently had more power than the other robust methods, and for all conditions power was comparable to Welch’s t -test. While type I error was also higher for B_{HSI} than the other methods (4.1.1. Type I error rates; Table 1), this does not necessarily explain the power differential. For example, at most B_{HSI} had 2 % more error but power was regularly between 3 – 10 % greater than the classical robust methods. The power discrepancy was especially pronounced between B_{HSI} and $PBOM$. The most comparable robust methods were B_{HSI} and $PBMOM$, although the latter never had more power. Importantly, the largest difference in type I error between these models was 0.007 (Table 1 [row 3]), but the power difference for the same condition was 0.017 (Table 3 [row 3]).

4.2. Contaminated (Heavy-Tailed) Outcome

4.2.1. Type I Error. The type I error rates for contaminated normal data are provided in Table 4. The Bayesian model B_{HSI} had error rates ranging from 0.046 – 0.069. The most comparable classical methods was $PBOM$ (0.048 – 0.062). The remaining methods were consistently conservative (<0.05). Notably, the trimming amount for Yuen’s test produced very different type I error rates. With a trimming proportion of 10 % ($\gamma = 0.10$), for example, the error rates were the most conservative of all methods. In contrast, a trimming portion of 20 % ($\gamma = 0.20$) produced error rates consistently around the nominal level ($\alpha = 0.05$).

4.2.2. Power. Power results are reported in table 5. For all methods, the larger contamination value ($\varepsilon = 0.2$) produced lower estimates of power. However, there were substantial differences between methods. For example, Welch’s t -test W_t was consistently conservative. This power differential was considerable, in that B_{HSI} had $\geq 400\%$ more power across all conditions. Yuen’s test ($\gamma = 0.10$) also had lower power when the contamination value was 0.2 which suggests that the tails were not trimmed adequately. The methods with the highest power were B_{HSI} and $PBMOM$. These methods were consistently within 3 % of one another for individual conditions, and the averages across condi-

tions (denoted with bold) were within 1 %. Importantly, B_{HSI} showed optimal power for normative and contaminated data, whereas the classical robust methods were inconsistent. That is, with normative data B_{HSI} and $PBOM$ had similar power, whereas B_{HSI} and $PBMOM$ had similar power for contaminated data.

4.2.3. Mean Squared Error. Mean squared error (MSE) is reported in Table 6. Welch’s t -test consistently had the highest MSE. Like type I error rates and power, MSE varied with the trimming proportion for Yuen’s test. For example, $\gamma = 0.10$ had at least twice as much MSE than $\gamma = 0.20$. The optimal models in terms of MSE were B_{HSI} and $PBMOM$, in that largest average MSE difference between methods was 0.008.

4.2.4. Returned Average. For all models the returned average was close to the true value ($\theta = 0.8$). The results are provided in the Table 7.

5. Motivating Example

We present two analyses that compare B_{HSI} , B_{NE} , Welch’s t -test, and the classical robust methods. The first uses the Electric dataset from the package *WRS2* in which classrooms were randomized into treatment ($N = 21$) and control groups ($N = 21$). The former was exposed to an educational TV show, but not the latter. Both groups completed a reading test at the beginning and end of the school year. For the first analysis, we included only first graders and analyzed gain scores (Figure 1A). For the second analysis group means, variances, and samples sizes from the same data were used to generate contaminated distributions (Figure 1A). In this section, we also highlight the predictive utility of the Bayesian robust model and introduce an effect size δ_{SI} .

5.1. Normative Outcome

5.1.1. Comparison Between Methods. The results are presented in Figure 1B (normative). Interval width varied between methods, and for $PBMOM$ and B_{NE} overlapped zero. Since groups sizes were equal, the previous simulation (4.1.2. Interval Width) provides some justification for interpreting the posterior distribution for B_{NE} probabilistically $p(b > 0|y) = 97.9\%$. The HDI for B_{HSI} did not include zero, which paralleled the classical robust methods (excluding $PBMOM$). Yuen’s method and $PBOM$ overestimated the observed difference in gain scores (Figure 1B: normative; black dots = observed).

5.1.2. Leave-One-Out Cross-Validation. The Bayesian models readily allow for model comparison. For example, approximate leave-one-out cross validation is computed from the expected log predictive density

Figure 1. A: Normative gain scores are the actual data. Contaminated gain scores were generated by contaminating the actual data. B: The bullseyes are point estimates obtained from each model. The black dots are the observed mean difference. The error bars denote 95% HDIs or confidence intervals.

(($elpd_{loo}$); Vehtari et al., 2016)

$$elpd_{loo} = \sum_{i=1}^n \log(y_i | y_{-i}) \quad (20)$$

in which

$$(y_i | y_{-i}) = \int p(y_i | \theta) p(\theta | y_{-i}) d\theta \quad (21)$$

denotes the predictive density with the i th data point removed. The leave-one-out information criterion (*LOOIC*) then follows

$$LOOIC = -2elpd_{loo} \quad (22)$$

where multiplying $elpd_{loo}$ by -2 converts the quantity to the deviance scale. Like related measures (e.g., AIC) lower values indicate better fit, but *LOOIC* has the advantage of providing a measure of uncertainty (standard error: SE). We compared *LOOIC* between the Bayesian models ($B_{HSt} = 340.69$, SE = 8.59; $B_{NE} = 338.74$, SE = 8.49). The *LOOIC* difference was 1.95 with a SE of 1.54. Therefore, although B_{HSt} is more complex than B_{NE} , the models did not differ substantially in terms of out-of-sample prediction and there was large uncertainty surrounding the *LOOIC* difference. Importantly, since B_{NE} is a special case of B_{HSt} ($\sigma_1^2 = \sigma_2^2$, $\nu \rightarrow \infty$) similar performance is expected with normative data ($\nu = 19.96$, 95% *HDI* = [2.94, 53.59]) in which variances are similar ($\beta_{\sigma^2_2} = 62.66$, 95% *HDI* = [-101.47, 266.05]).

5.2. Contaminated (Heavy-Tailed) Outcome

5.2.1. Comparison Between Methods. The results are provided in Figure 1B (Contaminated), and clearly

demonstrate the advantages of robust methods when outliers or tail-heaviness is present (Figure 1A: Contaminated). For example, while providing the optimal point estimate for the observed gain difference (black dot = actual; target = model based estimate), the intervals for both B_{NE} and W_t included zero. In contrast, the intervals for all robust methods excluded zero. This parallels the simulation results, in that the robust methods had substantially higher power than W_t .

However, the robust methods showed some discrepancies between the estimate and observed mean difference. Yuen's method with $\gamma = 0.1$ showed the largest difference (observed = 12.31; estimated = 14.97).

5.2.2. Posterior Predictive Checks. The Bayesian models allow for examining model fit via the posterior predictive distribution. Posterior predictive checks (*PPC*) work under the assumption that the fitted model should generate data y_{rep} that looks like the observed y

$$p(y_{rep} | y) = \int p(y_{rep} | \theta) p(\theta | y) d\theta. \quad (23)$$

Importantly, the posterior predictive distribution $p(y_{rep} | y)$ accounts for parameter uncertainty by marginalizing over the posterior distribution $p(\theta | y)$. As a measure of discrepancy, the posterior predictive p -value is defined as

$$p\text{-value} = Pr(T(y_{rep}) > T(y) | y) \quad (24)$$

where $T(y)$ is a test-statistic of interest (e.g., observed mean). The p -value is then the probability that $T(y_{rep})$ is greater than $T(y)$. Optimal p -values ≈ 0.50 , whereas extreme values (> 0.95 or < 0.05) suggest areas for model improvement. Figure 2A ($T(y)$ = dotted line; $T(y_{rep})$ = histogram) illustrates

Figure 2. A: Posterior predictive distributions. B: Posterior distributions for B_{HS_t} ($\beta_{\mu 2}$ = mean difference; $\beta_{\sigma^2 2}$ = variance difference; ν = degrees of freedom parameter; δ_{St} = standardized effect size). Black bars denotes the 95-% highest density interval (HDI), whereas the black and white triangles are the posterior medians.

posterior predictive checks for B_{HS_t} and B_{NE} . The group specific standard deviations were better described by B_{HS_t} . This was expected since B_{NE} assumed equal variances, for example, the p -value was 0.99 for the control group. Examining the minimum and maximum values allows for determining model performance in the tail regions. Here B_{HS_t}

modeled the observed data better than B_{NE} . However, B_{HS_t} was imperfect (treatment group minimum: $p = 0.94$) and could be improved. For example, we assumed a common value for ν but the model could be extended to estimate group-specific tail-heaviness.

A distinct aspect of Bayesian methods is that unknowns

are expressed as probability distributions. The posterior distributions for B_{HSI} are presented in Figure 2B ($\beta_{\mu 2}$ = mean difference; $\beta_{\sigma^2 2}$ = variance difference; ν = degrees of freedom parameter; δ_{St} = standardized effect size). The treatment group mean was larger than the control group ($\beta_{\mu 2}$ = 1027.75, 95%-HDI = [291.52, 2436.06]), but was also more variable $p(\beta_{\sigma^2 2} > 0|y) \approx 100\%$. The degrees of freedom parameter ν was estimated from the data and indicated tail-heaviness ($\nu = 5.06$, 95%-HDI = [1.01, 23.51]). The effect size δ_{St} (delta Student- t) was obtained from the posterior estimates. First, the reference group (intercept) variance follows

$$\sigma_1^2 = (\exp[\beta_{\omega 1}])^2 \frac{\hat{\nu}}{\hat{\nu} - 2}, \text{ if } \hat{\nu} > 2 \quad (25)$$

where the posterior estimates are exponentiated and $\hat{\nu}$ is the posterior median. The treatment group variance is given by

$$\sigma_2^2 = (\exp[\beta_{\omega 1} + \beta_{\omega 2}])^2 \frac{\hat{\nu}}{\hat{\nu} - 2}, \text{ if } \hat{\nu} > 2 \quad (26)$$

The proposed effect size is then computed as

$$\delta_{St} = \frac{\beta_{\mu 2}}{\sqrt{\frac{\sigma_1^2 + \sigma_2^2}{2}}} \quad (27)$$

in which $\beta_{\mu 2}$ is the posterior estimates for the mean difference. We then summarized $p(\delta_{St}|y)$ with the posterior median and highest density interval ($\delta_{St} = 0.60$, 95%-HDI = [0.03, 1.20]). This effect size is like Cohen's d , but is computed with respect to ν and can be extended to more complex methods such as multilevel models (Williams, Carlsson, & Bürkner, 2017).

6. Discussion

The present study is the first to characterize frequentist properties of a robust Bayesian model compared classical robust methods. We also demonstrated Bayesian inferential methods such as approximate leave-one-out cross validation and posterior predictive checks. With respect to frequentist calibration, we found that B_{HSI} often had superior performance compared to classical methods. For example, B_{HSI} had similar power to Welch's t -test for normally distributed data. The percentile bootstrap one-step M-estimator ($PBOM$) had the best performance among classical methods, but never had more power than B_{HSI} . However, for heavy-tailed data $PBOM$ had considerably less power (7.6 – 9.9 %) than B_{HSI} . Here, the percentile bootstrap modified one-step M-estimator $PBMOM$ had almost identical power to B_{HSI} . Additionally, type I error rates, coverage probabilities, and mean square error showed consistent performance across data types. This demonstrates that B_{HSI} can provide optimal estimation for normative and heavy-tailed data.

6.1. Comparison to Robust Bayesian Methods

Some studies have investigated robust Bayesian methods. For example, Gelman, Carlin, et al. (2014) described using a Student- t likelihood for robust estimation. In the psychological literature, Kruschke (2013) introduced a Bayesian model for unequal variances and outliers that sought to provide an alternative to null hypothesis significance testing. We built upon this paper in several important ways. First, we compared B_{HSI} to well-established robust methods with simulations. In contrast, Kruschke (2013) compared one dataset between the Bayesian model and Welch's t -test. This approach does not allow for determining which model has optimal performance when used many times. Second, the present model is a linear regression that can be extended to include continuous predictors and multi-factorial interactions. Kruschke (2013) estimated group specific posterior distributions and then computed differences with subtraction. This parametrization works for two-groups (or one-way "ANOVA"), but does not readily allow for more complex models. Third, the proposed effect size did not consider the degrees of freedom parameter ν . In other words, the scale parameter and not the expected variance was used to compute the effect size. This can lead to overestimation that is directly related to tail-heaviness ($\nu \rightarrow 0$, $\sigma > \omega$).

6.2. Comparison to Classical Robust Methods

There is a large literature on classical robust methods that includes several classes of estimators. For example, S-estimators that generalize least median squares by minimizing the expected residual scale from M-estimation (Çetin & Toka, 2011). The robust methods included in the present paper were selected for many reasons. First, they have been thoroughly characterized with simulations (Wilcox, Keselman, & Kowalchuk, 1998a). This allowed for comparing the proposed Bayesian model to established classical methods. Second, we focused on methods that have been described in the psychological literature (Wilcox, 2012). While this allowed for comparing methods that would be familiar to psychologists, this decision also presented some limitations. For example, the *WRS2* package does not provide information criterion (e.g., AIC). Thus, we could not compare classical methods to the Bayesian model in terms of model comparison and selection. It should be noted that M-estimators can provide an index of model fit (Ronchetti, 1985; Wang et al., 2017). To our knowledge, however, this cannot be achieved with the bootstrap versions ($PBOM$ and $PBMOM$) that have optimal frequentist properties (Wilcox, 1993).

6.3. Bayesian Predictive Utility

We demonstrated two important tools for examining fitted Bayesian models. Approximate leave-one-out cross validation was used to assess expected out-of-sample perfor-

mance (5.1.2. Leave-One-Out Cross-Validation), while posterior predictive checks allowed for inferring model fit to the observed data (5.2.2. Posterior Predictive Checks). To be clear, these methods are not necessarily for rejecting or accepting a model but for guiding model improvement (Gelman, 2013). For example, *LOOIC* provides a measure of out-of-sample performance across the outcome variable. However, predicting certain aspects of the sample may be important to consider (treatment vs. control group). In this case, out-of-sample performance can be evaluated with pointwise *LOO* values and used for model comparison (Gabry, Simpson, Vehtari, Betancourt, & Gelman, 2017).

6.4. Frequentist Properties of Bayesian Models

Bayesian methods are often promoted in stark contrast to classical methods (Kruschke, 2013). Indeed, some authors explicitly state that inferential goals differ between Bayesian and frequentist approaches. We agree, but a distinction must be made between investigating frequentist properties of our models and making frequentist inferences. Examining frequentist properties of Bayesian models has a long tradition in statistical oriented fields (Rubin, 1984). In simulation, we know the truth and random sampling allows for investigating contexts under which our model has sub-optimal performance (Gelman, 2011). Our results demonstrated that posterior density is influenced by unequal variances and tail-heaviness. Accordingly, simulations allow for inferring aspects of not only posterior distributions but also Bayes factors. Since Bayes factors can be computed as the ratio of probability density between the posterior and prior at $\delta = 0$ (i.e., Savage-Dickey ratio) (Wagenmakers, Lodewyckx, Kuriyal, & Grasman, 2010), for example, the obtained value will be sensitive to posterior width. Thus, we view examining our models' expected performance as extremely useful for model improvement and for validating Bayesian evidential quantities (Rubin, 1984).

To be clear, we are not advocating for classical inferences such as null hypothesis significance testing. For the proposed Bayesian robust model B_{HSt} , the estimated posterior distributions can be interpreted probabilistically as plausible values for the true effect (Harrell & Shih, 2001). This can be accomplished with highest density intervals, in addition to directional probabilities $p(b > 0|y)$. Of course, these probabilities are conditional on the prior distribution, all other model assumptions, and the observed data. It should be noted that priors are often thought to express belief, and that probabilistic inference depends on this narrow view (Morey, Romeijn, & Rouder, 2013). In contrast, we view the prior distribution as an assumption that can be verified like the assumed error distribution (Andrews & Baguley, 2013). A proper (i.e., integration to one) posterior distribution can always be interpreted probabilistically. The more pressing question is whether the posterior allows for demonstrably

valid inferences, and this can be determined with simulations and/or posterior predictive checks.

6.5. Bayesian Robust Methods in Practice

Despite demonstrating that B_{HSt} provides more consistent performance than classical methods, we view the flexibility of Bayesian methods as the most important practical advantage. B_{HSt} provide a general framework that can be extended to more complex models. It was recently suggested that Welch's t -test should be the default in psychology (Delacre, Lakens, & Leys, 2017). However, in a classical framework, solutions to achieve optimal frequentist properties are often method specific and do not readily generalize. For example, Welch's t -test applies a degrees of freedom correction for two groups when variances are heterogeneous, but cannot be used for other contexts (e.g., ANOVA) and other model-specific heterogeneity corrections are necessary (e.g., robust standard errors). In contrast, B_{HSt} and Bayesian modeling in general overcomes these challenges through flexible model assumptions and disregard to sampling distributions for inference. Although B_{HSt} overcomes these challenges and provides a general method that can be applied to common research designs, it should not be considered a default for psychology. Future studies should compare B_{HSt} to a skew-normal model (Martin & Williams, 2017), or a combination of the two (e.g., a skew-T distribution), from which a general default may be recommended.

6.6. Areas for Improvement

According to Bradley (1978) all error rates for B_{HSt} were acceptable if following the more liberal guidelines (0.025 – 0.075). However, the proposed Bayesian robust model can be improved in several ways. First, while the differences were not large, type I error rates were consistently higher than the classical robust methods. This could be due to predicting the scale parameter ω on the logarithmic scale, which may underestimate variability and produce *HDI*s that are too narrow. Future work should investigate this possibility, particularly whether B_{HSt} can accurately estimate group specific variances. Second, the proposed effect size δ_{St} can be improved. We assumed a common value for ν , which can bias group specific variances and thus the effect size estimate. Importantly, B_{HSt} can be extended to allow ν to differ between groups. This would improve upon the proposed effect size, and potentially address the inflated error rates. We used the posterior median for ν to compute δ_{St} . Ideally, the full posterior should be used by restricting ν to be greater than 2. Future work could also define parameters not as fixed point estimates, but as probability distributions in which random draws are used to generate data (Andradóttir & Bier, 2000). This would account for uncertainty by characterizing performance measures across a hypothesized distribution. It should also be noted that B_{HSt}

is explicitly for reasonably symmetric data and investigating performance for skewed data is an important future direction (Martin & Williams, 2017).

6.7. Concluding Remarks

In conclusion, we demonstrated that Bayesian robust methods can have adequate frequentist calibration, optimal performance across data types (normative vs. contaminated), and allow for model evaluation. Bayesian methods have been limited in practice due to computational demand and requiring specialized software. The methods described in this paper can be implemented with as much programming ability as the classical methods (Appendix B: R-code). Furthermore, with modern computers fully Bayesian methods can now be characterized with simulation. Together, this work provides a foundation from which Bayesian robust methods can be used in practice and further characterized in methodological literatures.

References

- Andradóttir, S. & Bier, V. M. (2000). Applying Bayesian ideas in simulation. *Simulation Practice and Theory*, 8(3-4), 253–280. doi:10.1016/S0928-4869(00)00025-2
- Andrews, M. & Baguley, T. (2013). Prior approval: The growth of Bayesian methods in psychology. *British Journal of Mathematical and Statistical Psychology*, 66(1), 1–7. doi:10.1111/bmsp.12004
- Betancourt, M. (2017, January). A Conceptual Introduction to Hamiltonian Monte Carlo. arXiv: 1701.02434. Retrieved from <http://arxiv.org/abs/1701.02434>
- Bradley, J. V. (1978). Robustness? *British Journal of Mathematical and Statistical Psychology*, 31(2), 144–152. doi:10.1111/j.2044-8317.1978.tb00581.x
- Buning, H. (1997). Robust analysis of variance. *Journal of Applied Statistics*, 24(3), 319–332. doi:10.1080/03610927708827539
- Bürkner, P.-C. (in press). brms: an R package for bayesian multilevel models using stan. *Journal of Statistical Software*.
- Çetin, M. & Toka, O. (2011). The comparing of S-estimator and M-estimators in linear regression. *Gazi University Journal of Science*, 24(4), 747–752.
- Daszykowski, M., Kaczmarek, K., Vander Heyden, Y., & Walczak, B. (2007). Robust statistics in data analysis — A review: Basic concepts. *Chemometrics and Intelligent Laboratory Systems*, 85(2), 203–219. doi:10.1016/j.chemolab.2006.06.016
- Delacre, M., Lakens, D., & Leys, C. (2017). Why Psychologists Should by Default Use Welch's t-test Instead of Student's t-test. *International Review of Social Psychology*, 30(1), 92–101. doi:10.5334/irsp.82
- Erceg-Hurn, D. M. & Miroseovich, V. M. (2008). Modern robust statistical methods: an easy way to maximize the accuracy and power of your research. *American Psychologist*, 63(7), 591–601. doi:10.1037/0003-066X.63.7.591
- Fernández, C. & Steel, M. F. J. (1998). On Bayesian Modeling of Fat Tails and Skewness. *Journal of the American Statistical Association*, 93(441), 359–371. doi:10.1080/01621459.1998.10474117
- Fonseca, T. C. O., Ferreira, M. A. R., & Migon, H. S. (2008, February). Objective Bayesian analysis for the Student-t regression model. *Biometrika*, 95(2), 325–333. doi:10.1093/biomet/asn001
- Gabry, J., Simpson, D., Vehtari, A., Betancourt, M., & Gelman, A. (2017). Visualization in Bayesian workflow. arXiv: 1709.01449. Retrieved from <http://arxiv.org/abs/1709.01449>
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (Comment on Article by Browne and Draper). *Bayesian Analysis*, 1(3), 515–534. doi:10.1214/06-BA117A
- Gelman, A. (2011). Bayesian Statistical Pragmatism. *Statistical Science*, 26(1), 10–11. doi:10.1214/10-STS337
- Gelman, A. (2013). Two simple examples for understanding posterior p-values whose distributions are far from uniform. *Electronic Journal of Statistics*, 7(1), 2595–2602. doi:10.1214/13-EJS854. arXiv: 0000000 [math.PR]
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, B. D., Vehtari, A., & Rubin, D. B. (2014). *Bayesian Data Analysis* (3rd ed.). Boca Raton: CRC Press.
- Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 24(6), 997–1016. doi:10.1007/s11222-013-9416-2. arXiv: 1307.5928
- Gelman, A., Jakulin, A., Pittau, M. G., & Su, Y.-S. (2008). A Weakly Informative Default Prior Distribution for Logistic and Other Regression Models. *The Annals of Applied Statistics*, 2(4), 1360–1383. doi:10.1214/08-AOAS
- Gelman, A., Meng, X.-L., & Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. Vol.6, No.4. *Statistica Sinica*, 6(4), 733–807. doi:10.1.1.142.9951
- Gelman, A., Simpson, D., & Betancourt, M. (2017). The prior can generally only be understood in the context of the likelihood. *ArXiv preprint*. arXiv: 1708.07487. Retrieved from <http://arxiv.org/abs/1708.07487>
- Grissom, R. J. (2000). Heterogeneity of variance in clinical data. *Journal of Consulting and Clinical Psychology*, 68(1), 155–165. doi:10.1037/0022-006X.68.1.155
- Harrell, F. E. & Shih, Y. C. T. (2001). Using full probability models to compute probabilities of actual interest to decision makers. *International journal of technology*

- assessment in health care*, 17(1), 17–26. doi:[10.1017/S0266462301104034](https://doi.org/10.1017/S0266462301104034)
- Huber, P. J. (1981). *Robust Statistics*. New York: Wiley.
- Keselman, H. J., Huberty, C. J., Lix, L. M., Olejnik, S., Cribbie, R. A., Donahue, B., ... Levin, J. R. (1998). Statistical practices of educational researchers: An analysis of their ANOVA, MANOVA, and ANCOVA analyses. *Review of Educational Research*, 68(3), 350–386. doi:[10.3102/00346543068003350](https://doi.org/10.3102/00346543068003350)
- Keselman, H. J., Othman, A. R., Wilcox, R. R., & Fradette, K. (2004). The new and improved two-sample t test. *Psychological Science*, 15(1), 47–51. doi:[10.1111/j.0963-7214.2004.01501008.x](https://doi.org/10.1111/j.0963-7214.2004.01501008.x)
- Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General*, 142(2), 573–603. doi:[10.1037/a0029146](https://doi.org/10.1037/a0029146)
- Kruschke, J. K. & Liddell, T. M. (2017). The Bayesian New Statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychonomic Bulletin & Review*. doi:[10.3758/s13423-016-1221-4](https://doi.org/10.3758/s13423-016-1221-4)
- Luh, W.-M. & Guo, J.-H. (2007). Approximate sample size formulas for the two-sample trimmed mean test with unequal variances. *British Journal of Mathematical and Statistical Psychology*, 60(1), 137–146. doi:[10.1348/000711006X100491](https://doi.org/10.1348/000711006X100491)
- Mair, P., Schoenbrodt, F., & Wilcox, R. R. (2017). WRS2: Wilcox robust estimation and testing. 0.9-2.
- Martin, S. R. & Williams, D. R. (2017). Outgrowing the Procrustean Bed of Normality : The Utility of Bayesian Modeling for Asymmetrical Data Analysis, 1–22. doi:[10.17605/OSF.IO/26M49](https://doi.org/10.17605/OSF.IO/26M49). PsyArXiv:26m49
- Micceri, T. (1989). The Unicorn, the Normal Curve, and Other Improbable Creatures. *Psychological Bulletin*, 105(1), 156–166. doi:[10.1037/0033-2909.105.1.156](https://doi.org/10.1037/0033-2909.105.1.156)
- Morey, R. D., Romeijn, J. W., & Rouder, J. N. (2013). The humble Bayesian: Model checking from a fully Bayesian perspective. *British Journal of Mathematical and Statistical Psychology*, 66(1), 68–75. doi:[10.1111/j.2044-8317.2012.02067.x](https://doi.org/10.1111/j.2044-8317.2012.02067.x)
- Özdemir, A. F. (2010). Comparing measures of location when the underlying distribution has heavier tails than normal. *İstatistikçiler Dergisi*, 3, 8–16.
- Özdemir, A. F. (2013). Comparing two independent groups: A test based on a one-step M-estimator and bootstrap-t. *British Journal of Mathematical and Statistical Psychology*, 66(2), 322–337. doi:[10.1111/j.2044-8317.2012.02053.x](https://doi.org/10.1111/j.2044-8317.2012.02053.x)
- Ramsey, P. H. (1980). Exact type I error rates for robustness and Student's t test with unequal variances. *Journal of Educational and Behavioral Statistics*, 5(4), 337–349. doi:[10.3102/10769986005004337](https://doi.org/10.3102/10769986005004337)
- Ronchetti, E. (1985). Robust model selection in regression. *Statistics and Probability Letters*, 3(1), 21–23. doi:[10.1016/0167-7152\(85\)90006-9](https://doi.org/10.1016/0167-7152(85)90006-9)
- Rouder, J. N., Morey, R. D., & Wagenmakers, E.-J. (2016, May). The Interplay between Subjectivity, Statistical Practice, and Psychological Science. *Collabra*, 2(1), 1–12. doi:[10.1525/collabra.28](https://doi.org/10.1525/collabra.28)
- Rubin, D. B. (1984). Bayesianly Justifiable and Relevant Frequency Calculations for the Applied Statistician. *The Annals of Statistics*, 12(4), 1151–1172. doi:[10.1214/aos/1176346785](https://doi.org/10.1214/aos/1176346785). arXiv: [arXiv:1306.3979v1](https://arxiv.org/abs/1306.3979v1)
- Sanborn, A. N., Hills, T. T., Dougherty, M. R., Thomas, R. P., Yu, E. C., & Sprenger, A. M. (2014). Reply to Rouder (2014): Good frequentist properties raise confidence. *Psychonomic Bulletin & Review*, 21(2), 309–311. doi:[10.3758/s13423-014-0607-4](https://doi.org/10.3758/s13423-014-0607-4)
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-Positive Psychology. *Psychological Science*, 22(11), 1359–1366. doi:[10.1177/0956797611417632](https://doi.org/10.1177/0956797611417632)
- Stan Development Team. (2016). RStan: the R interface to Stan. R package version 2.14.1. Retrieved from <http://mc-stan.org/>
- Susanti, Y., Pratiwi, H., Sulistijowati, S., & Liana, T. (2014). M estimation, S estimation, And MM estimation in robust regression. *International Journal of Pure and Applied Mathematics*, 91(3), 349–360. doi:[10.12732/ijpam.v91i3.7](https://doi.org/10.12732/ijpam.v91i3.7)
- Vehtari, A., Gelman, A., & Gabry, J. (2016). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5), 1–20. doi:[10.1007/s11222-016-9696-4](https://doi.org/10.1007/s11222-016-9696-4). arXiv: [1507.04544](https://arxiv.org/abs/1507.04544)
- Wagenmakers, E. J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage-Dickey method. *Cognitive Psychology*, 60(3), 158–189. doi:[10.1016/j.cogpsych.2009.12.001](https://doi.org/10.1016/j.cogpsych.2009.12.001)
- Wang, J., Zamar, R., Marazzi, A., Yohai, V., Salibian-Barrera, M., Maronna, R., ... Konis, K. (2017). *Robust: port of the s+ "robust library"*. R package version 0.4-18. Retrieved from <https://CRAN.R-project.org/package=robust>
- Welch, B. L. (1938). The Significance of the Difference Between Two Means when the Population Variances are Unequal. *Biometrika*, 29(3/4), 350. doi:[10.2307/2332010](https://doi.org/10.2307/2332010)
- Wilcox, R. R. (1993). Comparing one-step M-estimators of location when there are more than two groups. *Psychometrika*, 58(1), 71–78. doi:[10.1007/BF02294471](https://doi.org/10.1007/BF02294471)
- Wilcox, R. R. (1995). ANOVA: A Paradigm for Low Power and Misleading Measures of Effect Size? *Review of Educational Research*, 65(1), 51–77. doi:[10.3102/00346543065001051](https://doi.org/10.3102/00346543065001051)

- Wilcox, R. R. (1996). A note on testing hypotheses about trimmed means. *Biometric Journal*, 38(2), 173–180. doi:[10.1002/bimj.4710380205](https://doi.org/10.1002/bimj.4710380205)
- Wilcox, R. R. (1998a). How many discoveries have been lost by ignoring modern statistical methods? *American Psychologist*, 53(3), 300–314. doi:[10.1037/0003-066X.53.3.300](https://doi.org/10.1037/0003-066X.53.3.300)
- Wilcox, R. R. (1998b). The goals and strategies of robust methods. *British Journal of Mathematical and Statistical Psychology*, 51(1), 1–39. doi:[10.1111/j.2044-8317.1998.tb00659.x](https://doi.org/10.1111/j.2044-8317.1998.tb00659.x)
- Wilcox, R. R. (2001). *Fundamentals of Modern Statistical Methods*. New York, NY: Springer New York. doi:[10.1007/978-1-4757-3522-2](https://doi.org/10.1007/978-1-4757-3522-2)
- Wilcox, R. R. (2012). *Introduction to Robust Estimation and Hypothesis Testing*. Academic Press. doi:[10.1198/tech.2005.s334](https://doi.org/10.1198/tech.2005.s334). arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3)
- Wilcox, R. R. (2016). Comparing dependent robust correlations. *British Journal of Mathematical and Statistical Psychology*, 69(3), 215–224. doi:[10.1111/bmsp.12069](https://doi.org/10.1111/bmsp.12069)
- Wilcox, R. R., Charlin, V. L., & Thompson, K. L. (1986). New monte carlo results on the robustness of the anova f, w and f statistics. *Communications in Statistics - Simulation and Computation*, 15(February 2015), 933–943. doi:[10.1080/03610918608812553](https://doi.org/10.1080/03610918608812553)
- Wilcox, R. R. & Keselman, H. J. (2002). Power analyses when comparing trimmed means. *Journal of Modern Applied Statistical Methods*, 1(1). doi:[10.22237/jmasm/1020254820](https://doi.org/10.22237/jmasm/1020254820)
- Wilcox, R. R. & Keselman, H. J. (2003). Repeated measures one-way ANOVA based on a modified one-step M-estimator. *British Journal of Mathematical and Statistical Psychology*, 56(1), 15–25. doi:[10.1348/000711003321645313](https://doi.org/10.1348/000711003321645313)
- Wilcox, R. R., Keselman, H. J., & Kowalchuk, R. K. (1998a). Can tests for treatment group equality be improved?: The bootstrap and trimmed means conjecture. *British Journal of Mathematical and Statistical Psychology*, 51(1), 123–134. doi:[10.1111/j.2044-8317.1998.tb00670.x](https://doi.org/10.1111/j.2044-8317.1998.tb00670.x)
- Wilcox, R. R., Keselman, H. J., & Kowalchuk, R. K. (1998b). Can tests for treatment group equality be improved?: The bootstrap and trimmed means conjecture. *British Journal of Mathematical and Statistical Psychology*, 51(1), 123–134. doi:[10.1111/j.2044-8317.1998.tb00670.x](https://doi.org/10.1111/j.2044-8317.1998.tb00670.x)
- Williams, D. R., Carlsson, R., & Bürkner, P.-C. (2017). Between-litter variation in developmental studies of hormones and behavior: Inflated false positives and diminished power. *Frontiers in Neuroendocrinology*, (August), –1. doi:[10.1016/j.yfrne.2017.08.003](https://doi.org/10.1016/j.yfrne.2017.08.003)
- Yuen, K. K. (1974). Two-sample trimmed t for unequal population variances. *Biometrics Bulletin*, 61(1), 165–170. doi:[10.1093/biomet/61.1.165](https://doi.org/10.1093/biomet/61.1.165)

Appendix A
Tables

Table 1

Type I error rates ($\alpha = 0.05$) for normally distributed data: the Bayesian Heteroskedastic Student-t regression (B_{HS_t}), a Bayesian model assuming normality and equal variances (B_{NE}), Welch's t -test (W_t), Yuen's test ($\gamma = 0.1$ and 0.2), the percentile bootstrap one-step M -estimator ($PBOM$), and the percentile bootstrap modified one-step M -estimator ($PBMOM$).

n_1, n_2	σ_1, σ_2	B_{HS_t}	B_{NE}	W_t	Yuen ($\gamma = 0.1$)	Yuen ($\gamma = 0.2$)	$PBOM$	$PBMOM$
20,40	1,1	0.054	0.055	0.052	0.052	0.055	0.054	0.040
20,40	2,1	0.057	0.112	0.047	0.051	0.052	0.054	0.038
20,40	3,1	0.065	0.146	0.051	0.052	0.057	0.058	0.045
20,40	4,1	0.063	0.162	0.050	0.052	0.057	0.060	0.047
$\bar{x}$		0.060	0.119	0.050	0.052	0.055	0.057	0.043
30,30	1,1	0.054	0.052	0.049	0.050	0.053	0.052	0.039
30,30	2,1	0.058	0.056	0.052	0.051	0.053	0.054	0.040
30,30	3,1	0.061	0.056	0.051	0.051	0.051	0.055	0.042
30,30	4,1	0.060	0.057	0.051	0.052	0.054	0.057	0.042
$\bar{x}$		0.058	0.055	0.051	0.051	0.053	0.055	0.041
40,20	1,1	0.056	0.053	0.052	0.054	0.057	0.056	0.042
40,20	2,1	0.052	0.021	0.050	0.051	0.054	0.051	0.038
40,20	3,1	0.055	0.014	0.055	0.049	0.052	0.049	0.038
40,20	4,1	0.060	0.012	0.053	0.054	0.055	0.057	0.043
$\bar{x}$		0.056	0.025	0.053	0.052	0.055	0.053	0.040

Table 2

Interval width ($\theta = 0$; Bayesian = 95-% highest density; Classical = 95-% confidence) for normally distributed data: the Bayesian Heteroskedastic Student-t regression (B_{HS_t}), a Bayesian model assuming normality and equal variances (B_{NE}), Welch's t -test (W_t), Yuen's test ($\gamma = 0.1$ and 0.2), the percentile bootstrap one-step M -estimator ($PBOM$), and the percentile bootstrap modified one-step M -estimator ($PBMOM$). The parentheses include interval width standard deviations.

n_1, n_2	σ_1, σ_2	B_{HS_t}	B_{NE}	W_t	Yuen ($\gamma = 0.1$)	Yuen ($\gamma = 0.2$)	$PBOM$	$PBMOM$
20, 40	1,1	1.102 (0.13)	1.086 (0.11)	1.102 (0.13)	1.144 (0.16)	1.206 (0.21)	1.104 (0.15)	1.249 (0.20)
20,40	2,1	1.884 (0.28)	1.528 (0.18)	1.941 (0.29)	2.024 (0.36)	2.136 (0.47)	1.911 (0.33)	2.170 (0.42)
20,40	3,1	2.713 (0.44)	2.053 (0.29)	2.822 (0.44)	2.939 (0.55)	3.127 (0.98)	2.755 (0.50)	3.126 (0.65)
20,40	4, 1	3.585 (0.59)	2.623 (0.39)	3.737 (0.58)	3.901 (0.73)	4.148 (0.98)	3.642 (0.68)	4.127 (0.86)
30,30	1,1	1.033 (0.11)	1.026 (0.10)	1.032 (0.10)	1.068 (0.12)	1.116 (0.15)	1.039 (0.12)	1.172 (0.15)
30,30	2, 1	1.617 (0.18)	1.618 (0.186)	1.639 (0.18)	1.701 (0.22)	1.778 (0.29)	1.638 (0.21)	1.846 (0.26)
30,30	3, 1	2.276 (0.29)	2.287 (0.29)	2.331 (0.28)	2.421 (0.35)	2.535 (0.45)	2.314 (0.33)	2.605 (0.41)
30,30	4, 1	2.962 (0.39)	2.978 (0.39)	3.047 (0.38)	3.163 (0.48)	3.322 (0.61)	3.014 (0.45)	3.395 (0.56)
40,20	1, 1	1.100 (0.13)	1.086 (0.11)	1.101 (0.13)	1.142 (0.16)	1.201 (0.21)	1.101 (0.15)	1.247 (0.20)
40,20	2,1	1.155 (0.16)	1.888 (0.21)	1.548 (0.14)	1.603 (0.18)	1.673 (0.23)	1.561 (0.18)	1.757 (0.23)
40,20	3,1	2.088 (0.22)	2.743 (0.31)	2.097 (0.20)	2.171 (0.25)	2.267 (0.32)	2.102 (0.25)	2.362 (0.31)
40,20	4, 1	2.644 (0.29)	3.608 (0.42)	2.687 (0.28)	2.781 (0.34)	2.911 (0.44)	2.689 (0.33)	3.009 (0.41)

Table 3

Power for normally distributed data: the Bayesian Heteroskedastic Student-tregression ($B_{HS,t}$), a Bayesian model assuming normality and equal variances (B_{NE}), Welch's t -test (W_t), Yuen's test ($\gamma = 0.1$ and 0.2), the percentile bootstrap one-step M -estimator ($PBOM$), and the percentile bootstrap modified one-step M -estimator ($PBMOM$). The effect size y_{es} is the sampling distribution mean (determined via simulations).

n_1, n_2	σ_1, σ_2	y_{es}	$B_{HS,t}$	W_t	Yuen ($\gamma = 0.1$)	Yuen ($\gamma = 0.2$)	$PBOM$	$PBMOM$
20,40	1,1	0.52	0.808	0.809	0.782	0.744	0.789	0.702,
20,40	2,1	0.39	0.392	0.367	0.347	0.312	0.373	0.293,
20 40	3,1	0.34	0.226	0.197	0.190	0.181	0.209	0.162,
20 40	4,1	0.31	0.152	0.130	0.128	0.121	0.142	0.107
$\bar{x}$			0.395	0.376	0.362	0.340	0.378	0.316
30,30	1,1	0.52	0.862	0.865	0.842	0.803	0.842	0.766
30,30	2,1	0.39	0.484	0.473	0.449	0.413	0.460	0.377
30,30	3,1	0.34	0.293	0.277	0.257	0.239	0.270	0.218
30,30	4,1	0.31	0.198	0.179	0.172	0.164	0.182	0.139
$\bar{x}$			0.459	0.449	0.460	0.405	0.439	0.375
40,20	1,1	0.52	0.805	0.807	0.775	0.734	0.784	0.694
40,20	2,1	0.39	0.530	0.526	0.500	0.466	0.507	0.412
40,20	3,1	0.34	0.336	0.319	0.299	0.278	0.305	0.239
40,20	4,1	0.31	0.230	0.216	0.207	0.199	0.214	0.171
$\bar{x}$			0.475	0.467	0.445	0.419	0.453	0.379

Table 4

Type I error rates ($\alpha = 0.05$) for contaminated normal data $CN(\varepsilon, k_j)$: the Bayesian Heteroskedastic Student- t regression (B_{HS_t}), Welch's t -test (W_t), Yuen's test ($\gamma = 0.1$ and 0.2), the percentile bootstrap one-step M -estimator ($PBOM$), and the percentile bootstrap modified one-step M -estimator ($PBMOM$).

n_1, n_2	ε	k_1, k_2	B_{HS_t}	W_t	Yuen ($\gamma = 0.1$)	Yuen ($\gamma = 0.2$)	$PBOM$	$PBMOM$
220,40	0.1	10,10	0.054	0.035	0.041	0.049	0.053	0.037
20,40	0.2	10,10	0.047	0.038	0.024	0.043	0.051	0.036
20,40	0.1	20,10	0.059	0.022	0.040	0.048	0.053	0.042
20,40	0.2	20,10	0.043	0.032	0.019	0.041	0.056	0.038
20,40	0.1	10,20	0.065	0.032	0.041	0.052	0.055	0.040
20,40	0.2	10,20	0.049	0.044	0.019	0.042	0.053	0.038
20,40	0.1	20,20	0.066	0.026	0.039	0.052	0.057	0.042
20,40	0.2	20,20	0.043	0.038	0.014	0.038	0.050	0.036
$\bar{x}$			0.053	0.033	0.030	0.046	0.054	0.039
30,30	0.1	10,10	0.053	0.037	0.040	0.049	0.053	0.041
30,30	0.2	10,10	0.043	0.044	0.025	0.043	0.050	0.035
30,30	0.1	20,10	0.067	0.026	0.044	0.053	0.057	0.046
30,30	0.2	20,10	0.050	0.039	0.017	0.044	0.053	0.037
30,30	0.1	10,20	0.061	0.028	0.040	0.050	0.056	0.041
30,30	0.2	10,20	0.046	0.043	0.017	0.041	0.049	0.040
30,30	0.1	20,20	0.069	0.028	0.039	0.051	0.056	0.046
30,30	0.2	20,20	0.046	0.044	0.014	0.039	0.050	0.039
$\bar{x}$			0.054	0.036	0.030	0.046	0.053	0.041
40,20	0.1	10,10	0.062	0.031	0.043	0.049	0.057	0.041
40,20	0.2	10,10	0.045	0.039	0.025	0.041	0.052	0.038
40,20	0.1	20,10	0.065	0.034	0.043	0.050	0.054	0.042
40,20	0.2	20,10	0.051	0.041	0.018	0.041	0.048	0.039
40,20	0.1	10,20	0.063	0.024	0.046	0.056	0.062	0.045
40,20	0.2	10,20	0.045	0.029	0.020	0.040	0.053	0.039
40,20	0.1	20,20	0.066	0.025	0.038	0.049	0.056	0.044
40,20	0.2	20,20	0.048	0.041	0.014	0.042	0.051	0.039
$\bar{x}$			0.056	0.033	0.031	0.046	0.054	0.041

Table 5

Power ($\theta = 0.8$) for contaminated normal data $CN(\varepsilon, k_j)$: the Bayesian Heteroskedastic Student- t regression (B_{HSI}), Welch's t -test (W_t), Yuen's test ($\gamma = 0.1$ and 0.2), the percentile bootstrap one-step M -estimator ($PBOM$), and the percentile bootstrap modified one-step M -estimator ($PBMOM$).

n_1, n_2	ε	k_1, k_2	B_{HSI}	W_t	Yuen ($\gamma = 0.1$)	Yuen ($\gamma = 0.2$)	$PBOM$	$PBMOM$
20,40	0.1	10,10	0.673	0.218	0.616	0.637	0.647	0.663
20,40	0.2	10,10	0.558	0.112	0.360	0.504	0.466	0.537
20,40	0.1	20,10	0.680	0.145	0.586	0.632	0.637	0.680
20,40	0.2	20,10	0.553	0.068	0.289	0.487	0.436	0.556
20,40	0.1	10,20	0.685	0.139	0.594	0.631	0.636	0.665
20,40	0.2	10,20	0.571	0.077	0.294	0.496	0.452	0.545
20,40	0.1	20,20	0.684	0.098	0.568	0.625	0.629	0.682
20,40	0.2	20,20	0.555	0.060	0.243	0.480	0.434	0.565
$\bar{x}$			0.620	0.115	0.444	0.562	0.521	0.612
30,30	0.1	10,10	0.725	0.218	0.669	0.693	0.700	0.725
30,30	0.2	10,10	0.614	0.119	0.396	0.565	0.512	0.604
30,30	0.1	20,10	0.739	0.147	0.652	0.697	0.699	0.735
30,30	0.2	20,10	0.624	0.080	0.310	0.547	0.484	0.623
30,30	0.1	10,20	0.742	0.135	0.651	0.694	0.695	0.737
30,30	0.2	10,20	0.617	0.080	0.315	0.548	0.489	0.618
30,30	0.1	20,20	0.735	0.105	0.626	0.692	0.685	0.745
30,30	0.2	20,20	0.611	0.068	0.255	0.528	0.467	0.634
$\bar{x}$			0.676	0.119	0.484	0.621	0.591	0.678
40,20	0.1	10,10	0.673	0.216	0.612	0.634	0.646	0.660
40,20	0.2	10,10	0.552	0.114	0.361	0.506	0.468	0.537
40,20	0.1	20,10	0.685	0.136	0.590	0.628	0.635	0.661
40,20	0.2	20,10	0.574	0.077	0.291	0.497	0.454	0.550
40,20	0.1	10,20	0.676	0.141	0.578	0.624	0.627	0.673
40,20	0.2	10,20	0.545	0.068	0.290	0.485	0.439	0.548
40,20	0.1	20,20	0.678	0.100	0.567	0.623	0.627	0.677
40,20	0.2	20,20	0.556	0.061	0.240	0.474	0.429	0.562
$\bar{x}$			0.617	0.114	0.441	0.559	0.541	0.609

Table 6

Mean squared error ($\theta = 0.8$) over 10,000 iterations for contaminated normal data $CN(\varepsilon, k_j)$: the Bayesian Heteroskedastic Student- t regression (B_{HSI}), Welch's t -test (W_t), Yuen's test ($\gamma = 0.1$ and 0.2), the percentile bootstrap one-step M-estimator ($PBOM$), and the percentile bootstrap modified one-step M-estimator ($PBMOM$).

n_1, n_2	ε	k_1, k_2	B_{HSI}	W_t	Yuen ($\gamma = 0.1$)	Yuen ($\gamma = 0.2$)	$PBOM$	$PBMOM$
20,40	0.1	10,10	0.105	0.800	0.135	0.113	0.109	0.102
20,40	0.2	10,10	0.132	1.542	0.376	0.173	0.175	0.129
20,40	0.1	20,10	0.109	2.345	0.224	0.116	0.118	0.103
20,40	0.2	20,10	0.128	4.583	0.945	0.227	0.192	0.118
20,40	0.1	10,20	0.109	1.614	0.155	0.113	0.113	0.101
20,40	0.2	10,20	0.130	3.056	0.554	0.173	0.181	0.124
20,40	0.1	20,20	0.106	3.069	0.237	0.116	0.113	0.097
20,40	0.2	20,20	0.126	5.922	1.070	0.233	0.202	0.123
$\bar{x}$			0.118	2.866	0.462	0.158	0.150	0.112
30,30	0.1	10,10	0.096	0.728	0.117	0.098	0.100	0.090
30,30	0.2	10,10	0.111	1.362	0.303	0.149	0.150	0.109
30,30	0.1	20,10	0.095	1.695	0.141	0.101	0.100	0.089
30,30	0.2	20,10	0.113	3.514	0.592	0.156	0.158	0.106
30,30	0.1	10,20	0.098	1.758	0.143	0.099	0.102	0.092
30,30	0.2	10,20	0.117	3.427	0.586	0.154	0.164	0.109
30,30	0.1	20,20	0.098	2.774	0.165	0.101	0.102	0.088
30,30	0.2	20,20	0.116	5.557	0.874	0.172	0.172	0.104
$\bar{x}$			0.106	2.602	0.365	0.129	0.131	0.100
40,20	0.1	10,10	0.108	0.846	0.146	0.109	0.116	0.106
40,20	0.2	10,10	0.128	1.525	0.377	0.168	0.173	0.126
40,20	0.1	20,10	0.111	1.561	0.153	0.113	0.115	0.102
40,20	0.2	20,10	0.129	3.093	0.545	0.180	0.182	0.122
40,20	0.1	10,20	0.109	2.369	0.231	0.117	0.118	0.102
40,20	0.2	10,20	0.132	4.501	0.926	0.236	0.194	0.123
40,20	0.1	20,20	0.111	3.038	0.233	0.115	0.117	0.101
40,20	0.2	20,20	0.128	6.157	1.104	0.233	0.207	0.117
$\bar{x}$			0.120	2.886	0.464	0.159	0.153	0.112

Table 7

Average $\hat{\theta}$ ($\theta = 0.8$) over 10,000 iterations for contaminated normal data $CN(\varepsilon, k_j)$: the Bayesian Heteroskedastic Student- t regression (B_{HSI}), Welch's t -test (W_t), Yuen's test ($\gamma = 0.1$ and 0.2), the percentile bootstrap one-step M -estimator ($PBOM$), and the percentile bootstrap modified one-step M -estimator ($PBMOM$).

n_1, n_2	ε	k_1, k_2	B_{HSI}	W_t	Yuen ($\gamma = 0.1$)	Yuen ($\gamma = 0.2$)	$PBOM$	$PBMOM$
20,40	0.1	10,10	0.800	0.804	0.800	0.801	0.802	0.801
20,40	0.2	10,10	0.801	0.799	0.801	0.795	0.799	0.800
20,40	0.1	20,10	0.804	0.820	0.808	0.800	0.805	0.804
20,40	0.2	20,10	0.800	0.788	0.791	0.801	0.797	0.801
20,40	0.1	10,20	0.803	0.795	0.802	0.802	0.801	0.801
20,40	0.2	10,20	0.800	0.798	0.804	0.808	0.798	0.800
20,40	0.1	20,20	0.797	0.806	0.804	0.800	0.798	0.795
20,40	0.2	20,20	0.799	0.805	0.801	0.802	0.803	0.802
$\bar{x}$			0.801	0.802	0.801	0.801	0.800	0.801
30,30	0.1	10,10	0.799	0.794	0.799	0.802	0.799	0.800
30,30	0.2	10,10	0.798	0.808	0.804	0.801	0.799	0.798
30,30	0.1	20,10	0.801	0.833	0.807	0.802	0.804	0.800
30,30	0.2	20,10	0.799	0.778	0.784	0.807	0.795	0.801
30,30	0.1	10,20	0.805	0.772	0.803	0.799	0.803	0.805
30,30	0.2	10,20	0.802	0.822	0.805	0.795	0.803	0.801
30,30	0.1	20,20	0.796	0.790	0.790	0.795	0.793	0.795
30,30	0.2	20,20	0.799	0.787	0.796	0.808	0.799	0.799
$\bar{x}$			0.800	0.798	0.799	0.801	0.799	0.800
40,20	0.1	10,10	0.801	0.802	0.799	0.803	0.799	0.800
40,20	0.2	10,10	0.801	0.798	0.800	0.801	0.799	0.800
40,20	0.1	20,10	0.800	0.795	0.799	0.798	0.799	0.801
40,20	0.2	20,10	0.800	0.811	0.807	0.804	0.802	0.801
40,20	0.1	10,20	0.799	0.789	0.795	0.801	0.797	0.798
40,20	0.2	10,20	0.799	0.815	0.809	0.800	0.800	0.800
40,20	0.1	20,20	0.797	0.824	0.802	0.795	0.798	0.794
40,20	0.2	20,20	0.800	0.797	0.801	0.805	0.799	0.800
$\bar{x}$			0.800	0.804	0.802	0.801	0.799	0.799

Appendix B

R-code

R-code for fitting the Bayesian and classical robust models

1) Bayesian Heterogeneous Student-t Regression:

(note: fitted with the package *brms*)

```
bhst <- brm(bf(y ~ group, sigma ~ group), family = student(), data = dat)
```

2) Yuen's test:

(note: fitted with the package *WRS2*)

```
y10 <- yuen(y ~ group, data = dat, tr = 0.10)
```

```
y20 <- yuen(y ~ group, data = dat, tr = 0.20)
```

3) One-Step M-Estimator:

(note: fitted with the package *WRS2*)

```
pbom <- pb2gen(y ~ group, data = dat, est = "onestep", nboot = 599)
```

4) Modified One-Step M-Estimator:

(note: fitted with the package *WRS2*)

```
pbmom <- pb2gen(y ~ group, data = dat, est = "mom", nboot = 599)
```