Behavior Research Methods
https://doi.org/10.3758/s13428-022-01860-1

Bayesian hypothesis testing of mediation: Methods and the impact of prior odds specifications

Xiao Liu[1] · **Zhiyong Zhang**[1] · **Lijuan Wang**[1]

Accepted: 11 April 2022

Abstract

Mediation analysis is widely used to study whether the effect of an independent variable on an outcome is transmitted through a mediator. Bayesian methods have become increasingly popular for mediation analysis. However, limited research has been done on formal Bayesian hypothesis testing of mediation. Although hypothesis testing using Bayes factor for a single path is readily available, how to integrate the Bayes factors of two paths (from input to mediator and from mediator to outcome) while incorporating prior beliefs on the two paths and/or mediation is under-studied. In the current study, we propose a general approach to Bayesian hypothesis testing of mediation. The proposed approach allows researchers to specify prior odds based on the substantive research context and can be used in mediation modeling with latent variables. The impact of prior odds specifications on Bayesian hypothesis test of mediation is demonstrated via both real and hypothetical data examples. Both R functions and a user-friendly R web app are provided for the implementation of the proposed approach. Our study can add to researchers' toolbox of mediation analysis and raise researchers' awareness of the importance of prior odds specifications in Bayesian hypothesis testing of mediation.

Keywords Bayes factor · Mediation analysis · Bayesian hypothesis testing

Introduction

Mediation studies are common in psychology and other social and behavioral science disciplines. Mediation analysis is a useful approach for studying the mechanisms that underlie the effect of an independent variable *X* on an outcome variable *Y*. When performing mediation analyses, hypothesis testing is often conducted to test whether the effect of *X* on *Y* is transmitted by a mediator *M* (Baron & Kenny 1986; MacKinnon 2008).

For mediation analysis, various frequentist methods have bee proposed and evaluated (e.g., MacKinnon, Fritz, Williams, & Lockwood 2007; MacKinnon, Lockwood, Hoffman, West, & Sheets 2002; Tofighi & Kelley 2020). In recent years, Bayesian methods have been developed and becoming increasingly popular for mediation analysis (e.g., Che, Jin, & Zhang 2021; Miočević, Gonzalez, Valente, & MacKinnon 2018; Wang & Preacher 2015; Yuan & MacKinnon 2009). For example, Yuan & MacKinnon (2009) proposed Bayesian approaches for estimating mediation effects for both single-level and multilevel mediation models; Wang & Preacher (2015) applied Bayesian methods to moderated mediation analysis; and Che, Jin, & Zhang (2021) utilized Bayesian estimation in mediation analysis with social network data. Compared to the frequentist methods, Bayesian methods allow researchers to incorporate prior knowledge into mediation analysis. For estimation, incorporating prior information has been shown to help reduce uncertainty in posterior distributions of mediation effects and improve precision of mediation effect estimates (e.g., Yuan & MacKinnon 2009).

A majority of the developed Bayesian methods for mediation analysis focused on estimation rather than hypothesis testing. Credible interval estimates of mediation effects have been used in Bayesian mediation analysis. While Bayesian parameter estimation is useful, Bayesian hypothesis testing can facilitate theory prediction by (1) quantifying the strength of evidence provided by the data for two hypotheses (null vs. alternative hypotheses) using Bayes factor and (2) making probabilistic statements about the hypotheses using posterior odds given both prior beliefs and data. Formal Bayesian hypothesis testing, however, is still not

✉ Xiao Liu
xliu19@nd.edu

[1] Department of Psychology, University of Notre Dame, 390 Corbett Hall, Notre Dame, IN 46556, USA

Published online: 17 May 2022

widely conducted in Bayesian mediation analysis. Yuan & MacKinnon (2009, p. 319) noted that, "One important topic we have not covered in this article is hypothesis testing. An approximation is that we may test the null hypothesis of no mediation effect on the basis of whether the 95% credible interval contains zero. However, this approach has more of the flavor of conventional frequentist hypothesis testing. Strict Bayesian hypothesis testing is based on Bayes factor ..." To reduce the research gap, we study how to conduct Bayesian hypothesis testing for mediation in this paper. Next, we first briefly introduce Bayesian hypothesis testing, and then review the only existing study, to the best of our knowledge, on Bayesian hypothesis testing of mediation.

Bayesian hypothesis testing

Formal Bayesian hypothesis testing (BHT) is formulated as comparing two or more competing models. Using null hypothesis testing as an example, there are two models: the model under the null hypothesis (H_0) and the model under the alternative hypothesis (H_1). The two models are both probable to be true. Before collecting data, a researcher evaluates the prior probabilities of the two hypotheses against each other (prior odds; $PriorOdds = p[H_1]/p[H_0]$, with $p[H_1]$ and $p[H_0]$ denoting the prior probabilities of the alternative and null hypotheses respectively) based on prior knowledge. After collecting data (D), odds between the two hypotheses are updated to form the posterior odds ($PosteriorOdds = p[H_1|D]/p[H_0|D]$) using information in the collected data, by means of the Bayes' rule. The posterior odds quantifies the relative probabilities of alternative vs. null hypotheses by integrating both prior beliefs and the data information. The Bayes' rule states that the posterior odds equals the product of the prior odds and the Bayes factor (BF):

$$PosteriorOdds = PriorOdds \times BF. \tag{1}$$

The Bayes factor (Jeffreys 1961)

$$BF = \frac{p(D|H_1)}{p(D|H_0)} \tag{2}$$

is the relative likelihood of observing the data (D) under H_1 versus under H_0.[1] BF can be interpreted as a measure of the relative strength of data evidence in favor of the alternative hypothesis compared to the null hypothesis and is a popular quantity reported in Bayesian hypothesis testing.

During the debate on frequentist null hypothesis significance testing (NHST) over the past few decades (e.g., Berger & Sellke, 1987; Harlow et al., 1997; Iverson et al., 2009; Nickerson, 2000), BHT has been promoted as a promising alternative to NHST and becoming increasingly popular in psychology (e.g., Dienes 2019; Hoijtink, Mulder, van Lissa, & Gu 2019; Ly, Verhagen, & Wagenmakers 2016; Mulder & Wagenmakers 2016; Wagenmakers et al. 2018). In particular, with BHT, researchers can obtain the posterior probability that the null or the alternative hypothesis is true given the observed data to quantify the evidence supporting H_0 or H_1 (i.e., $p[H_0|D]$ or $p[H_1|D]$). In comparison, NHST evaluates the evidence against but not in favor of H_0. The p value yielded by NHST, which is the probability of observing data as extreme or more extreme than the observed one assuming H_0 is true, has been found to be liable to misinterpretation in practice (e.g., Anderson 2020; Berger & Sellke 1987; Ioannidis 2019). Furthermore, unlike NHST, BHT allows researchers to incorporate prior information (e.g., from researchers' substantive knowledge, previous studies, or historical data) into evaluating evidence for H_1 vs. H_0. This feature makes BHT a useful approach for cumulative research.and the posterior credible interval approach

In BHT, the null hypothesis can be a point (or "nill") null hypothesis stating that the parameter of interest is precisely equal to a specific value (e.g., zero), or an interval null hypothesis stating that the value of the parameter belongs to a certain interval (e.g., an interval around zero). There is a large body of literature on the use of BHT for testing point null hypotheses (e.g., Jeffreys 1961; Kass & Raftery 1995; Wagenmakers et al. 2018; Wetzels & Wagenmakers 2012). Recent years have also seen great development of BHT approaches for testing interval null hypotheses (e.g., Hoijtink et al., 2008; Liao et al., 2020; Morey & Rouder, 2011). In the current study, we consider point null hypotheses for testing mediation effects. This is because when conducting mediation analysis, researchers are often interested in testing whether the hypothesized mediation effect is present or not.

Bayesian hypothesis testing of mediation

Despite the popularity of BHT, limited research has been done on testing mediation with BHT. Different from the association or direct effect between two variables (e.g., a correlation or regression coefficient), a mediation effect involves at least two relations among three variables: the relation between an independent variable and a mediator (path α) and the relation between the mediator and an outcome (path β) (e.g., MacKinnon 2008; Valeri & VanderWeele 2013). The mediation effect (i.e., the indirect effect transmitted by the mediator) is present when both paths α and β are present.

To the best of our knowledge, the only existing article studying BHT for mediation is Nuijten et al. (2014). In

[1] To simplify notations, throughout the article, the numerators of *PriorOdds*, *PosteriorOdds* and *BF* are for H_1 or evaluated under H_1; and the denominators are for H_0 or evaluated under H_0.

Nuijten et al. (2014), a three-step BHT procedure was proposed for testing mediation with the simple mediation model. Specifically, in the first two steps, the Bayes factors of paths α and β were separately computed and transformed to the posterior probabilities that the paths exist, assuming each path has a prior odds of 1. In the third step, the posterior probability that mediation exists was calculated as the product of the two path posterior probabilities assuming the two paths are independent, and then converted to the Bayes factor of mediation, assuming the mediation effect also has a prior odds of 1.

Nuijten et al. (2014) contributed a useful addition to researchers' toolbox of mediation analysis. Despite the merits, the approach proposed in Nuijten et al. (2014) has a few potential limitations. First, the prior odds of the presence of path α, path β, and mediation are all assumed to equal 1. In order for this assumption to hold, the two paths have to be either both present or both absent, and thus the scenarios where one path is present while the other is absent are not allowed to occur a priori.[2] Although this specific prior odds specification might be appropriate for some substantive studies, it could occur that researchers may have other prior beliefs about the two paths and/or mediation. In such scenarios, other methods are needed to facilitate the integration of data evidence and researchers' prior beliefs into BHT of mediation. How the prior odds specifications impact the Bayes factor and posterior odds of mediation also awaits further investigation.

Second, the approach assumes that paths α and β are independent. However, this assumption is actually violated when the assumed prior odds in Nuijten et al. (2014) are applied, as we will show later. In addition, when latent variables are used in mediation analysis, the two paths can be dependent. In this case, the three-step procedure is not applicable and BHT methods for testing mediation need further development.

The current study

In this article, we expand the use of BHT for testing mediation with two specific goals. First, we aim to develop a method for BHT of mediation that can be applied to different research contexts. We will propose a general approach to calculating the Bayes factor and posterior odds between the null and alternative mediation hypotheses. The proposed approach is applicable to (1) different substantive scenarios where researchers might have different prior beliefs regarding the two paths and/or mediation effect and (2) different mediation models including the simple mediation and latent variable mediation models.

Second, we aim to examine the impact of prior odds specifications on BHT of mediation. To fulfill this goal, we will assess how the mediation Bayes factor and posterior odds would change with different prior odds specifications, and demonstrate the impact with both real and hypothetical data examples. We hope the examination helps raise researchers' awareness of the importance of prior odds specifications in BHT of mediation.

In the rest of the article, we first propose a general approach to BHT of mediation. Then we implement the proposed approach in the simple mediation model, provide an illustrative example, and investigate the impact of the prior odds specification. Next, we apply the proposed approach to the latent mediation model, and use a simulated-data example for illustration. We end the article with discussing the implications, limitations of the current study, and future research directions.

The Proposed Approach to Bayesian Hypothesis Testing of Mediation

In this section, we propose a general approach to computing the Bayes factor and posterior odds for hypothesis testing of a mediation effect. Denote the effect of an independent variable on a mediator by α, and the effect of the mediator on an outcome conditioning on the independent variable by β. Using the product of coefficient method (Baron & Kenny 1986; MacKinnon 2008), the mediation effect is quantified by $\alpha\beta$.[3] The BHT of mediation is formulated as comparing the null hypothesis (or model) $H_0 : \alpha\beta = 0$ vs. the alternative $H_1 : \alpha\beta \neq 0$. Other key notations used in the article and their descriptions are summarized in Table 1.

Following Eq. 2, the Bayes factor for testing mediation, denoted by BF^{med}, is the ratio between the likelihoods of observing the data under the alternative vs. null hypotheses. We have

$$BF^{med} = \frac{p(D|\alpha \neq 0, \beta \neq 0)}{q_{00|0}p(D|\alpha = 0, \beta = 0) + q_{01|0}p(D|\alpha = 0, \beta \neq 0) + q_{10|0}p(D|\alpha \neq 0, \beta = 0)}. \quad (3)$$

[2] If the prior odds of the presence of path α, path β, and mediation all equal 1, we have $p(\alpha \neq 0) = p(\beta \neq 0) = p(\alpha \neq 0, \beta \neq 0) = 0.5$. It is always true that $p(\alpha \neq 0) = p(\alpha \neq 0, \beta \neq 0) + p(\alpha \neq 0, \beta = 0)$ and $p(\beta \neq 0) = p(\alpha \neq 0, \beta \neq 0) + p(\alpha = 0, \beta \neq 0)$. In order for these prior probabilities to hold, both $p(\alpha \neq 0, \beta = 0) = 0$ and $p(\alpha = 0, \beta \neq 0) = 0$ need to be true.

[3] Besides the product of coefficient method, mediation effects can also be defined using the counterfactual approach (e.g., Holland 1986; Pearl 2001; Valeri & VanderWeele 2013), especially when causal interpretation is of interest. In this article, we use the product of coefficient method because it is widely used in the practice of mediation analysis in psychology and social science research and is relatively easy for substantive researchers to interpret.

Table 1 Summary of notations

Notation	Description	Mathematical definition
$q_{00\|0}$	conditional prior probability that paths α and β are both zero if mediation is absent	$q_{00\|0} = p(\alpha = 0, \beta = 0 \| \alpha\beta = 0)$
$q_{01\|0}$	conditional prior probability that path α is zero and path β is nonzero if mediation is absent	$q_{01\|0} = p(\alpha = 0, \beta \neq 0 \| \alpha\beta = 0)$
$q_{10\|0}$	conditional prior probability that path α is nonzero and path β is zero if mediation is absent	$q_{10\|0} = p(\alpha \neq 0, \beta = 0 \| \alpha\beta = 0)$
$PriorOdds^{\alpha}$	prior odds of path α	$PriorOdds^{\alpha} = \frac{p(\alpha \neq 0)}{p(\alpha = 0)}$
$PriorOdds^{\beta}$	prior odds of path β	$PriorOdds^{\beta} = \frac{p(\beta \neq 0)}{p(\beta = 0)}$
$PriorOdds^{med}$	prior odds of the mediation effect	$PriorOdds^{med} = \frac{p(\alpha\beta \neq 0)}{p(\alpha\beta = 0)}$
BF^{α}	Bayes factor of path α	$BF^{\alpha} = \frac{p(D\|\alpha \neq 0)}{p(D\|\alpha = 0)}$
BF^{β}	Bayes factor of path β	$BF^{\beta} = \frac{p(D\|\beta \neq 0)}{p(D\|\beta = 0)}$
BF^{med}	Bayes factor of the mediation effect	$BF^{med} = \frac{p(D\|\alpha\beta \neq 0)}{p(D\|\alpha\beta = 0)}$
$PosteriorOdds^{\alpha}$	posterior odds of path α	$PosteriorOdds^{\alpha} = \frac{p(\alpha \neq 0\|D)}{p(\alpha = 0\|D)}$
$PosteriorOdds^{\beta}$	posterior odds of path β	$PosteriorOdds^{\beta} = \frac{p(\beta \neq 0\|D)}{p(\beta = 0\|D)}$
$PosteriorOdds^{med}$	posterior odds of the mediation effect	$PosteriorOdds^{med} = \frac{p(\alpha\beta \neq 0\|D)}{p(\alpha\beta = 0\|D)}$

Specifically, under the alternative hypothesis that the mediation effect exists, paths α and β are both non-zero. Thus, the numerator in Eq. 3 provides the likelihood of observing the data under the alternative hypothesis. Under the null hypothesis of no mediation, there are three possible scenarios: paths α and β are both zero; path α is zero and path β is non-zero; or path β is zero and path α is non-zero. The conditional prior probabilities of the three scenarios under the null hypothesis, $q_{00|0}$, $q_{01|0}$, and $q_{10|0}$,[4] are used to obtain the weighted average (or marginal) likelihood of observing the data under the null hypothesis (the denominator of Eq. 3). The three conditional probabilities sum up to 1 and are determined based on researchers' prior knowledge about how likely each of the three scenarios occurs if the null is true. For example, $q_{00|0}$, $q_{01|0}$, and $q_{10|0}$ are set to be 1/3 when researchers believe the three scenarios under the null are equally likely to occur before collecting data. In the next section, we will discuss how the conditional prior probabilities are related to the prior odds of α, β, and mediation when α and β are independent.

The mediation posterior odds $PosteriorOdds^{med}$, the relative posterior probabilities of the presence vs. absence of mediation, can be obtained following the Bayes' rule (Eq. 1):

$$PosteriorOdds^{med} = BF^{med} \times PriorOdds^{med} \tag{4}$$

with $PriorOdds^{med}$ being the prior odds of mediation specified to reflect researchers' prior belief about the relative probabilities of the presence vs. absence of mediation before collecting data.

Equations 3 and 4 provide a means to conduct BHT of mediation that accommodates different research contexts. First, both equations do not require that paths α and β are independent. This makes the proposed approach applicable to not only the simple mediation model where paths α and β are independent but also latent variable mediation models in which the two paths may be dependent. Furthermore, as we will demonstrate, with Eqs. 3 and 4, prior knowledge can be incorporated into computing the mediation Bayes factor and posterior odds using different ways. This allows the application of the proposed approach across different substantive scenarios in which researchers' prior beliefs about the mediation process may vary.

Next, we first examine the special case where paths α and β are independent and then consider the general case where the paths can be dependent. Both R functions (fun_BHTmed.R) and an R web app (BayesianHypothesisTesting Mediation) are developed for implementing the proposed approach, and can be downloaded from https://github.com/xliu12/BHT.med. More details about the developed tools will be provided in the illustrative examples.

Scenario 1: Independent Paths α and β

In this section, we examine the proposed BHT approach when paths α and β are independent. We use the simple mediation model where the two paths are independent (Nuijten et al., 2014) as an example for illustration. We first present the technical procedure and illustrate it through a real-data example. Then, we investigate the impact of prior odds specification on the BHT results using hypothetical data sets.

[4] That is, $q_{00|0} = p(\alpha = 0, \beta = 0 | H_0 : \alpha\beta = 0) = \frac{p(\alpha=0,\beta=0)}{p(\alpha\beta=0)}$, $q_{01|0} = p(\alpha = 0, \beta \neq 0 | H_0 : \alpha\beta = 0) = \frac{p(\alpha=0,\beta\neq 0)}{p(\alpha\beta=0)}$, and $q_{10|0} = p(\alpha \neq 0, \beta = 0 | H_0 : \alpha\beta = 0) = \frac{p(\alpha\neq 0,\beta=0)}{p(\alpha\beta=0)}$.

Technical details (e.g., derivations of Eqs. 7-10) are given in the supplemental materials (https://github.com/xliu12/BHT.med/blob/main/bfmed_supplement_technical.pdf).

Consider the following simple mediation model

$$M = \alpha X + \epsilon_M \tag{5}$$

$$Y = \tau' X + \beta M + \epsilon_Y \tag{6}$$

where X, M, and Y are centered for the ease of discussion, and the residuals ϵ_M and ϵ_Y are assumed to have mean zero and be independent.

Under the assumption that paths α and β are independent, the mediation Bayes factor can be obtained using the Bayes factors of paths α and β (denoted by BF^{α} and BF^{β}, respectively; see Table 1 for the mathematical definitions):

$$BF^{med} = \frac{BF^{\alpha}BF^{\beta}}{q_{00|0} + q_{01|0}BF^{\beta} + q_{10|0}BF^{\alpha}}. \tag{7}$$

Following Eq. 4, the mediation posterior odds can be obtained as the product of the mediation Bayes factor and mediation prior odds. To compute the individual Bayes factors for paths α and β, various software programs or packages (e.g., Hoijtink et al., 2019; JASP Team, 2020; Morey et al., 2018; Nuijten et al., 2014; Wetzels & Wagenmakers, 2012) have been developed and can be used.

When paths α and β are independent, the prior odds specifications for path α, path β, and the mediation effect must satisfy the following constraint

$$PriorOdds^{med} = \frac{PriorOdds^{\alpha} \times PriorOdds^{\beta}}{1 + PriorOdds^{\beta} + PriorOdds^{\alpha}} \tag{8}$$

where $PriorOdds^{\alpha}$ and $PriorOdds^{\beta}$ denote the prior odds of paths α and β, respectively.[5]

With the equality constraint under the scenario that paths α and β are independent, the mediation Bayes factor and posterior odds can be obtained by using the Bayes factors and prior odds of paths α and β:

$$BF^{med} = \frac{(1 + PriorOdds^{\beta} + PriorOdds^{\alpha})BF^{\alpha}BF^{\beta}}{1 + PriorOdds^{\beta}BF^{\beta} + PriorOdds^{\alpha}BF^{\alpha}}; \tag{9}$$

$$PosteriorOdds^{med} = \frac{BF^{\alpha}BF^{\beta} \times PriorOdds^{\alpha}PriorOdds^{\beta}}{1 + BF^{\beta}PriorOdds^{\beta} + BF^{\alpha}PriorOdds^{\alpha}}. \tag{10}$$

Here, we recommend researchers to specify prior odds for BHT of mediation using one of two following ways. Which way to use depends on the prior beliefs researchers hold about path α, path β, and the mediation effect.

The first way is to specify the prior odds for paths α and β separately. For example, suppose a researcher, Researcher A, specifies the prior odds of paths α and β as both 1, indicating the prior probability of the presence (or absence) of each path is 0.5 based on prior knowledge. With this specification, according to the constraint in Eq. 8, the mediation prior odds is 1/3 (not 1 as specified in Nuijten et al., 2014), indicating the prior probability is 0.25 for the presence of mediation and 0.75 for the absence of mediation.

The second way is to specify prior odds for the mediation effect and one individual path, for example, path α.[6] With the equality constraint, the prior odds for path β is readily solvable. For example, suppose another researcher, Researcher B, specifies the mediation prior odds as 1 to reflect that Researcher B is neutral on the presence and absence of mediation, and the prior odds of paths α and β to be the same. With the equality constraint, the prior odds of the two individual paths is $\frac{\sqrt{0.5}}{1-\sqrt{0.5}} \approx 2.41$. For each individual path (α or β), it is more likely to be present than to be absent in the prior beliefs.

Using either way, Eqs. 9 and 10 can then be used to obtain the Bayes factor and posterior odds of mediation, respectively. Note that with the independence assumption for paths α and β, the three conditional prior probabilities under the null (i.e., $q_{00|0}$, $q_{01|0}$, and $q_{10|0}$) are determined with the specified prior odds. They are all 1/3 using Researcher A's prior odds specification,[7] suggesting that the three scenarios of no mediation are equally likely to occur. Using Researcher B's prior odds specifications, under the null of no mediation, the scenarios where one path is present and the other is absent are more likely ($q_{10|0}$ and $q_{01|0}$ being 0.41) than the scenario where the two paths are both absent ($q_{00|0}$ being 0.17).[8] Therefore, different prior odds specifications for α

[5] Nuijten et al. (2014) assumed paths α and β are independent and specified the prior odds of path α, path β, and the mediation effect all as 1 (i.e., $PriorOdds^{\alpha} = PriorOdds^{\beta} = PriorOdds^{med} = 1$). Their prior odds specification, however, does not satisfy the equality constraint (Eq. 8).

[6] Under the equality constraint, the prior odds specification for mediation should not exceed that for an individual path.

[7] With the independence assumption for paths α and β, $q_{00|0} = p(\alpha = 0, \beta = 0|\alpha\beta = 0) = p(\alpha = 0)p(\beta = 0)/p(\alpha\beta = 0) = (0.5 * 0.5)/0.75 = 1/3$; and similarly, $q_{01|0} = p(\alpha = 0)p(\beta \neq 0)/p(\alpha\beta = 0) = (0.5 * 0.5)/0.75 = 1/3$ and $q_{10|0} = p(\alpha \neq 0)p(\beta = 0)/p(\alpha\beta = 0) = (0.5 * 0.5)/0.75 = 1/3$.

[8] With the independence assumption for paths α and β, $q_{10|0} = p(\alpha \neq 0)p(\beta = 0)/p(\alpha\beta = 0) = \frac{\sqrt{0.5}(1-\sqrt{0.5})}{0.5} \approx 0.41$, $q_{01|0} = p(\alpha = 0)p(\beta \neq 0)/p(\alpha\beta = 0) = \frac{(1-\sqrt{0.5})\sqrt{0.5}}{0.5} \approx 0.41$, and $q_{00|0} = p(\alpha = 0, \beta = 0|\alpha\beta = 0) = p(\alpha = 0)p(\beta = 0)/p(\alpha\beta = 0) = \frac{(1-\sqrt{0.5})^2}{0.5} \approx 0.17$.

and β have different implications of the conditional prior probabilities of the three scenarios under the null.

An illustrative example: The PHLAME firefighter study

We use the example presented in Yuan & MacKinnon (2009) and Nuijten et al. (2014), the PHLAME (Promoting Healthy Lifestyles: Alternative Models' Effects) firefighter study (Elliot et al., 2007), to illustrate our BHT procedure for the independent paths scenario. We use the prior odds specifications of Researchers A and B to gain insights on the sensitivity of BHT results to prior odds specifications. In the firefighter example, the hypothesized mediation mechanism (H_1) is that randomized exposure to an intervention (X) affects knowledge of the benefits of eating fruits and vegetables (M), which in turn affects reported eating of fruits and vegetables (Y). The Bayes factors for paths α (the exposure-mediator relation) and β (the mediator-outcome relation) were reported in Nuijten et al. (2014) as $BF^{\alpha} = 10.06$ and $BF^{\beta} = 2.68$, respectively, indicating that the data are in favor of the presence of each path.

Researcher A's BHT results can be obtained by running our developed R function *pathb.a()* as follows:

$$pathb.a(PriorOdds.a = 1, PriorOdds.b = 1, BF.a = 10.06, BF.b = 2.68).$$

In the above R code, arguments "PriorOdds.a" and "PriorOdds.b" read in the prior odds Researcher A specified for paths α and β, respectively. Arguments "BF.a" and "BF.b" read in the Bayes factors of paths α and β, respectively.

Researcher B's BHT results can be obtained by running R function *med.a()*:

$$med.a(PriorOdds.med = 1, PriorOdds.a = sqrt(0.5)/(1 - sqrt(0.5)), BF.a = 10.06, BF.b = 2.68).$$

In the above R code, arguments "PriorOdds.med" and "PriorOdds.a" read in the prior odds Researcher B specified for the mediation effect and path α, respectively.

The R functions and R code used in this example can be downloaded from https://github.com/xliu12/BHT.med. The results can also be obtained using our developed R Shiny (RStudio, Inc 2021) web app (https://xiaoliu.shinyapps.io/BayesianHypothesisTestingMediation/), and are displayed in Table 2.

Comparing Tables 2a and 2b, one can see that with different prior odds specifications, BHT of mediation yields different results. First, although the Bayes factors of each individual path are the same for both researchers, the mediation Bayes factors are different. For Researcher A, the mediation Bayes factor is 5.887; whereas for Researcher B, the mediation Bayes factor is 4.948.

Table 2 BHT of mediation results for the PHLAME firefighter example. The Bayes factors for paths α (the exposure-mediator relation) and β (the mediator-outcome relation) computed by Nuijten et al. (2014) are $BF^{\alpha} = 10.06$ and $BF^{\beta} = 2.68$ respectively

	Prior Odds	Bayes Factor	Posterior Odds	Posterior Probability
(a) Researcher A's results with specifying the prior odds of paths α and β as both 1. The mediation prior odds equals 1/3 due to the equality constraint under the independent path assumption.				
Path α	1.000	10.06	10.060	0.910
Path β	1.000	2.68	2.680	0.728
Mediation	0.333	5.887	1.962	0.662
(b) Researcher B's results with specifying the mediation prior odds as 1 and the prior odds of path α as $\sqrt{0.5}/(1-\sqrt{0.5})$. The prior odds of path β equals $\sqrt{0.5}/(1-\sqrt{0.5})$ due to the equality constraint under the independent path assumption.				
Path α	2.414	10.06	24.287	0.960
Path β	2.414	2.68	6.470	0.866
Mediation	1.000	4.948	4.948	0.832

Second, although Researcher A has a larger mediation Bayes factor than Researcher B, Researcher A's posterior odds of mediation is nevertheless smaller than Researcher B's. Specifically, for Researcher A who believes the odds in favor of mediation is 1/3 initially, combining the data evidence with the prior belief leads to a mediation posterior odds of 1.962, with the posterior probability of the presence of mediation being 0.662. In comparison, for Researcher B who specifies a mediation prior odds of 1, the mediation posterior odds is about 4.948, indicating the posterior probability that mediation exists given the data is 0.832. In Nuijten et al. (2014), their default prior odds specification is that the prior odds of path α, path β, and the mediation effect all equal 1. Using their method, the Bayes factor and posterior odds of the mediation effect were both computed as 1.94 for the firefighter example. As discussed earlier, the default prior odds specification assumed in Nuijten et al. (2014) is inconsistent with the independent path assumption required by their method. This inconsistency makes their BHT results difficult to interpret.

As shown in this example, although the data are the same, the mediation Bayes factors and posterior odds change with different prior odds specifications.

Impact of prior odds specifications: Graphical demonstration with hypothetical data examples

In this subsection, we assess the impact of prior odds specification on the BHT of mediation via hypothetical data

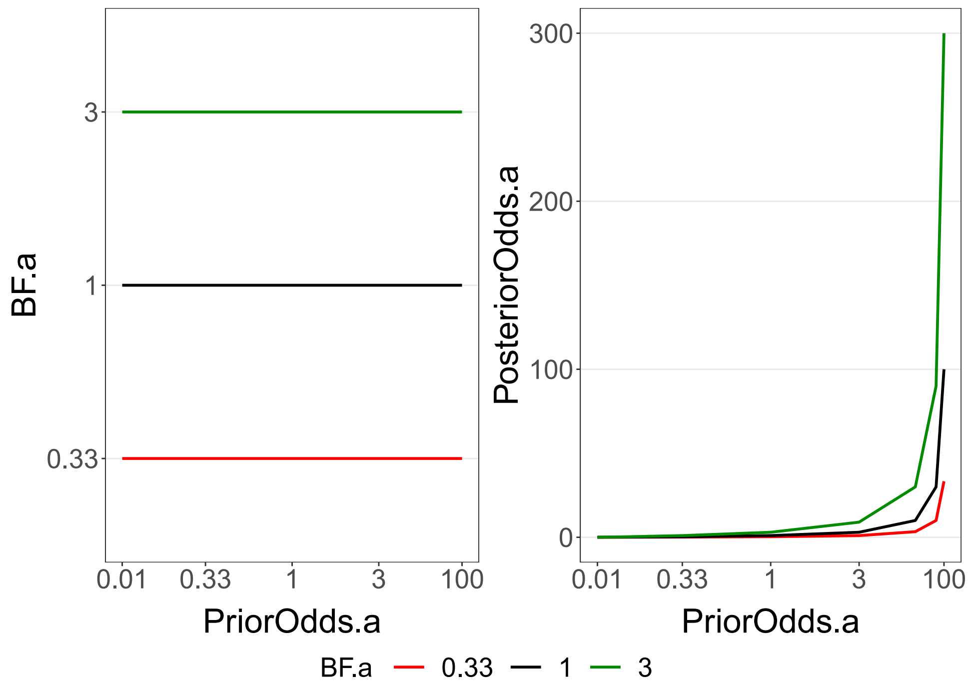

Fig. 1 Results from the Bayesian hypothesis testing (BHT) of a single relation, the regression coefficient for path α. The prior odds of path α were varied as $PriorOdds^{\alpha} = 1/100, 1/30, 1/10, 1/3, 1, 3, 10, 30,$ and 100. Given that the $PriorOdds^{\alpha}$ values are not equally spaced, we used the prior probabilities of path α being non-zero (i.e., $p[\alpha \neq 0] = PriorOdds^{\alpha}/[1 + PriorOdds^{\alpha}]$) as the breaks of the x-axes. The labels for the breaks in the x-axes are the corresponding prior odds values. The transformation from prior odds to prior probability is monotone and does not influence the increasing/decreasing patterns of the lines. The same plotting strategy is used in Fig. 2 as well

examples using a variety of prior odds specifications. Each hypothetical data set is characterized by a combination of paths Bayes factors BF^{α} (1/3, 1, or 3) and BF^{β} (1/3, 1, or 3). In total, there are $3 \times 3 = 9$ hypothetical data sets. The levels of the Bayes factors are chosen to represent varied strengths of data evidence: a Bayes factor of 1/3 indicates the data support the absence of the path, 1 means the data have no preference for the absence or presence of the path, and 3 suggests the data favor the presence of the path. For each hypothetical set, we apply the proposed BHT approach with varied prior odds specifications. The mediation Bayes factor (BF^{med}) and posterior odds ($PosteriorOdds^{med}$) are the results of interest.

First, to facilitate the comparison of the BHT for mediation vs. the BHT for a single relation, Fig. 1 demonstrates how the BHT results of path α, a single regression coefficient, change with varied prior odds specifications for path α. Specifically, in the left column, the Bayes factor of path α does not change with the prior odds of path α[9]; and in the right column, the posterior odds of path α increases monotonically as the prior odds of path α increases. The results in Fig. 1 are for a single coefficient, which are presented to serve as reference results.

Next, we assessed how the results from the BHT of mediation would change with different prior odds specifications of an individual path or the mediation effect. To do so, we (1) specified the prior odds for mediation as $PriorOdds^{med} = 1$, and varied that for path α as $PriorOdds^{\alpha} = 1, 3, 10, 30,$ and 100; or (2) fixed the prior odds of path α as $PriorOdds^{\alpha} = 1$ and varied the mediation prior odds as $PriorOdds^{med} = 1/100, 1/30, 1/10, 1/3,$ and 1.[10]

The results are demonstrated in Fig. 2. The mediation Bayes factor does not change with different prior odds specifications of an individual path or the mediation effect when the Bayes factors for individual paths α and β are both 1.

Under the scenario of fixing the mediation prior odds to be 1 (Fig. 2a), a higher prior odds for path α can lead to higher or lower values of the mediation Bayes factor and posterior odds when the individual path Bayes factors are not both 1.

Under the scenario of fixing the prior odds of path α to be 1 (Fig. 2b), the posterior odds for the mediation effect

[9] Although the varied prior odds values range from 0 to 100 in the left column of Fig. 1, this pattern holds when the prior odds of path α is larger than 100.

[10] Under the independent paths constraint (Eq. 8), the prior odds of path α must be larger than or equal to the prior odds of mediation.

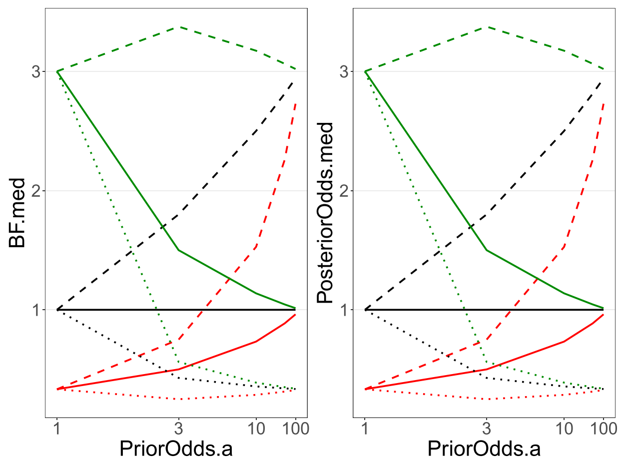

(a) Results from the BHT of mediation with varied prior odds specifications for an individual path, path α. The prior odds of the mediation effect is fixed at $PriorOdds^{med} = 1$.

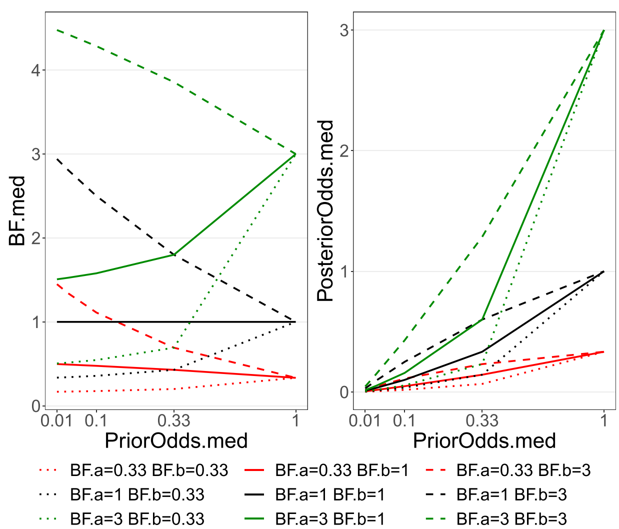

(b) Results from the BHT of mediation with varied prior odds specifications for mediation. The prior odds of path α is fixed at $PriorOdds^{\alpha} = 1$.

Fig. 2 The impact of prior odds specifications on the results from the Bayesian hypothesis testing (BHT) of mediation. Each line in a subplot is plotted for a hypothetical data example with a specific combination of the Bayes factors for paths α and β (the "BF.a" and "BF.b" in the legend)

increases as the mediation prior odds increases. This is similar to the change pattern of the posterior odds for a single relation (the right column of Fig. 1). However, when the individual path Bayes factors are not both 1, the mediation Bayes factor can be higher or lower as the mediation prior odds increases, which is different from the change pattern of the Bayes factor for a single relation (the left column of Fig. 1).

In summary, through the hypothetical data examples, we showed that in the BHT of mediation, both the prior odds of an individual path α (or β) and the prior odds of the mediation effect can impact the results of BHT. Unlike the BHT for a single relation (e.g., a regression coefficient), the prior odds of the mediation effect can influence not only the mediation posterior odds, but also the mediation Bayes factor. Based on the findings, we emphasize the importance of prior odds specifications in the BHT of mediation. We recommend that researchers should evaluate their prior belief about the mediation process carefully when specifying the prior odds, and provide the prior odds specifications of the mediation effect and of paths α and β when reporting the results.

Scenario 2: Dependent Paths α and β

In this section, we apply the proposed BHT approach to the general case where paths α and β can be dependent. We use a latent mediation model for illustration, as it has been shown that paths α and β in mediation models involving latent variables can be dependent (e.g., MacKinnon 2008; Valente, Gonzalez, Miočević, & MacKinnon 2016). We first provide the key BHT idea and then illustrate the implementation through a simulated-data example. Technical details are available in the supplemental material (https://github.com/xliu12/BHT.med/blob/main/bfmed_supplement_technical.pdf).

Consider the following latent mediation model:

$$\mathrm{x} = \Lambda_x \xi_x + \delta_x \tag{11}$$

$$\mathrm{m} = \Lambda_m \eta_m + \epsilon_m \tag{12}$$

$$\mathrm{y} = \Lambda_y \eta_y + \epsilon_y \tag{13}$$

$$\begin{pmatrix} \eta_m \\ \eta_y \end{pmatrix} = \begin{pmatrix} \alpha \\ \tau' \end{pmatrix} \xi_x + \begin{pmatrix} 0 & 0 \\ \beta & 0 \end{pmatrix} \begin{pmatrix} \eta_m \\ \eta_y \end{pmatrix} + \begin{pmatrix} \zeta_m \\ \zeta_y \end{pmatrix} \tag{14}$$

where ξ_x, η_m and η_y are the latent independent variable, mediator, and outcome variable respectively, and x, m, and y are the corresponding indicators. In Eqs. 11-13, the measurement errors δ_x, ϵ_m, and ϵ_y are assumed to have mean zero, be independent of one another and of the latent variables. In Eq. 14, the path coefficient α denotes the effect of ξ_x on η_m, the path coefficient β represents the effect of η_m on η_y conditioning on ξ_x; and the structural residuals ζ_y and ζ_m are assumed to have mean zero, be independent of one another and of ξ_x and the measurement errors. Figure 3 shows the structural relationships among ξ_x, η_m, and η_y in the latent mediation model.

For this model, we use the Bayesian information criterion (BIC) approximation approach to obtain the mediation

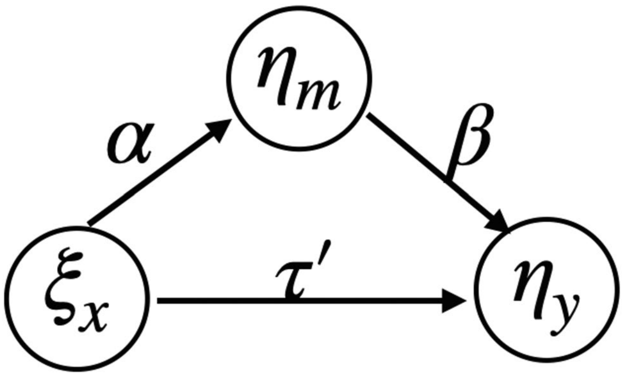

Fig. 3 The structural relationships among the latent independent variable ξ_x, mediator η_m and outcome variable η_y in the latent mediation model (Eqs. 11 to 14)

Bayes factor. For complex models such as the latent mediation model, the likelihoods in the Bayes factor formula (Eq. 3) are challenging to compute (e.g., Tendeiro & Kiers 2019; Bollen, Harden, Ray, & Zavisca 2014). For example, to compute the likelihood in the numerator, one needs to integrate the likelihood function over the entire space of the parameters under the alternative hypothesis. For the latent mediation model, the parameters include both those in the measurement model (Eqs. 11-13) and those in the structural model (Eq. 14). The complexity of such integration in latent variable models makes it challenging to obtain exact values of the likelihoods in Eq. 3. In the BHT literature, a frequently used approach to approximating integrated likelihoods is the BIC approximation (Kass & Raftery 1995; Raftery 1995), which has been shown to be useful for obtaining Bayes factors for models with latent variables (e.g., Bollen, Harden, Ray, & Zavisca 2014; Merkle & Wang 2016). Approximating the Bayes factor in Eq. 3 using BIC also makes it easy for researchers to conduct complex BHT of mediation. Let BIC$\{\mathbf{M}\}$ denote the BIC value of model $\mathbf{M}$. The mediation Bayes factor can be approximated as

$$BF^{med} \approx \frac{e^{-\mathrm{BIC}\{\mathbf{M}_{11}:\alpha\neq0,\beta\neq0\}/2}}{q_{00|0}e^{-\mathrm{BIC}\{\mathbf{M}_{00}:\alpha=0,\beta=0\}/2} + q_{01|0}e^{-\mathrm{BIC}\{\mathbf{M}_{01}:\alpha=0,\beta\neq0\}/2} + q_{10|0}e^{-\mathrm{BIC}\{\mathbf{M}_{10}:\alpha\neq0,\beta=0\}/2}} \tag{15}$$

where $\mathbf{M}_{11}$, $\mathbf{M}_{01}$, $\mathbf{M}_{10}$, and $\mathbf{M}_{00}$ denote respectively the unconstrained mediation model, the model where path α is constrained at zero, the model where path β is constrained at zero, and the model where both paths are constrained at zero. The BIC values can be obtained from, often in the output of, commonly used structural equation modeling software, such as Mplus (Muthén & Muthén 2017) and the lavaan R package (Rosseel et al., 2020). The mediation posterior odds can then be obtained following the Bayes' rule (Eq. 1).

With Eq. 15, researchers can incorporate their prior belief into the BHT of mediation in two ways. First, researchers can specify the conditional prior probabilities $q_{00|0}, q_{01|0}$, and $q_{10|0}$ for the three scenarios of no mediation under the constraint that they sum up to 1 to obtain the mediation Bayes factor, and specify a prior odds for the mediation effect to obtain the mediation posterior odds following the Bayes' rule. For example, suppose a researcher, Researcher C, specifies the three conditional prior probabilities under the null hypothesis to all equal 1/3 to reflect the prior belief that the three scenarios of no mediation are equally likely to occur, and specifies the mediation prior odds as 1 to give the presence and absence of mediation equal prior probabilities. With this specification, the prior odds of paths α and β both equal 2, implying the prior belief that for each of the two paths, its presence is more likely than its absence.[11] Furthermore, the two paths are not independent: the joint prior probability of paths α and β being both present under this specification is 1/2, higher than that under the independent paths assumption, which is 4/9.[12]

Second, researchers can specify prior odds for the two paths and the mediation effect under the inequality constraint that the mediation prior odds must not exceed the prior odds of any path. Then, the conditional prior probabilities $q_{00|0}, q_{01|0}$, and $q_{10|0}$ for the three scenarios of no mediation are readily obtainable.[13] Equation 15 and the Bayes' rule (Eq. 1) can then be used to obtain the mediation Bayes factor and posterior odds, respectively. In particular, for the default prior odds specification assumed in Nuijten et al. (2014) (i.e., the prior odds of path α, path β, and the mediation effect all equal 1), $q_{01|0} = q_{10|0} = 0$ and $q_{00|0} = 1$. That is, the scenarios where one path is present and the other is absent are not allowed to occur a priori; and the only possible scenario under H_0 is that the two paths are both absent. While this specification does not meet the independent paths constraint, it can occur in the dependent paths scenario.

To illustrate the BHT of mediation in the dependent paths scenario, we simulated a sample of size $n = 100$ from the latent mediation model (Eqs. 11-14) with the path coefficients $\alpha = 0$, $\beta = .39$ and $\tau' = 0$. That is, the null hypothesis of no mediation is true in the data generation. The latent input variable, mediator, and outcome are each measured

[11] This can be seen by computing the prior probabilities of the models $\mathbf{M}_{11}$, $\mathbf{M}_{01}$, $\mathbf{M}_{10}$, and $\mathbf{M}_{00}$, which are $p_{11} = p(\alpha \neq 0, \beta \neq 0) = PriorOdds^{med}/(1 + PriorOdds^{med}) = 1/2$, $p_{01} = (1 - p_{11})q_{01|0} = 1/6$, $p_{10} = (1 - p_{11})q_{10|0} = 1/6$ and $p_{00} = (1 - p_{11})q_{00|0} = 1/6$, respectively. Thus, the prior probabilities of paths α and β are $p(\alpha \neq 0) = p_{11} + p_{10} = 2/3, p(\beta \neq 0) = p_{11} + p_{01} = 2/3$, respectively. The prior odds of the two paths are thus $PriorOdds^{\alpha} = PriorOdds^{\beta} = \frac{2/3}{1-2/3} = 2$.

[12] See the supplemental material for technical details.

[13] See the supplemental material for technical details.

by three indicators with loadings 0.9, 0.8, and 0.7; and the variances of the indicators and latent variables are set as 1.

For the simulated sample, we obtained the BIC values of the unconstrained mediation model ($\text{BIC}\{\mathbf{M}_{11}\} \approx 2157$), the model where path α is constrained at zero ($\text{BIC}\{\mathbf{M}_{01}\} \approx 2153$), the model where path β is constrained at zero ($\text{BIC}\{\mathbf{M}_{10}\} \approx 2160$), and the model where both paths are constrained at zero ($\text{BIC}\{\mathbf{M}_{00}\} \approx 2155$) using the lavaan package (Rosseel et al., 2020) in R.

The BHT results of the simulated data using Researcher C's prior specifications can be obtained by running our developed R function *med.qs()*as follows:

$$med.qs(PriorOdds.med = 1, H0.q10 = 1/3, H0.q01 = 1/3, H0.q00 = 1/3,$$
$$BIC11 = 2157, BIC01 = 2153, BIC10 = 2160, BIC00 = 2155).$$

In the above R code, argument "PriorOdds.med" reads in the prior odds specified for the mediation effect. Arguments "H0.q10", "H0.q01", and "H0.q00" read in the specified values of $q_{10|0}$, $q_{01|0}$, and $q_{00|0}$ (the conditional prior probabilities of the three scenarios of no mediation under the null hypothesis), respectively. Note that the three specified values need to sum up to 1. Arguments "BIC11", "BIC01", "BIC10", and "BIC00" read in the BIC values of the unconstrained mediation model, the model where path α is constrained at zero, the model where path β is constrained at zero, and the model where both paths are constrained at zero, respectively.

The BHT results of specifying the prior odds of the two paths and the mediation effect to all equal 1 can be obtained by running our developed R function *med.ab()* :

$$med.ab(PriorOdds.med = 1, PriorOdds.a = 1, PriorOdds.b = 1,$$
$$BIC11 = 2157, BIC01 = 2153, BIC10 = 2160, BIC00 = 2155).$$

In the above R code, arguments "PriorOdds.med", "PriorOdds.a", and "PriorOdds.b" read in the specified prior odds for the mediation effect, path α, and path β, respectively. Note that these prior odds are all 1, which are the default prior odds in Nuijten et al. (2014) and can be used when paths α and β are dependent.

The results are displayed in Table 3. The R functions and R code used in this example can be downloaded from https://github.com/xliu12/BHT.med. The results can also be obtained using our developed R Shiny (RStudio, Inc, 2021) web app (https://xiaoliu.shinyapps.io/BayesianHypothesisTestingMediation/).[14]

In Table 3a, the mediation Bayes Factor is 0.216, indicating that the likelihood of observing the data is approximately 5 times as large if mediation is absent than if mediation is present. Given the evidence in the data, the odds between the presence and absence of mediation is lowered from 1 to about 0.216, with the posterior probability of the presence of mediation being 0.178.

Table 3 BHT of mediation results for the simulated-data example. The data ($n = 100$) were simulated from the latent mediation model (Eqs. 11-14) with $\alpha = 0$, $\beta = .39$ and $\tau' = 0$. The BIC values of models $\mathbf{M}_{11}$, $\mathbf{M}_{01}$, $\mathbf{M}_{10}$, and $\mathbf{M}_{00}$ are 2157, 2153, 2160 and 2155, respectively

	Prior Odds	Bayes Factor	Posterior Odds	Posterior Probability
(a) Results with prior odds specifications as $q_{00|0} = q_{01|0} = q_{10|0} = 1/3$ and $PriorOdds^{med} = 1$. With these specifications, $PriorOdds^{\alpha} = PriorOdds^{\beta} = 2$.				
Path a	2.000	0.124	0.248	0.199
Path b	2.000	1.668	3.335	0.769
Mediation	1.000	0.216	0.216	0.178
(b) Results with prior odds specifications as $PriorOdds^{\alpha} = PriorOdds^{\beta} = PriorOdds^{med} = 1$. With these specifications, $q_{01|0} = q_{10|0} = 0$ and $q_{00|0} = 1$.				
Path a	1.000	0.282	0.282	0.220
Path b	1.000	0.282	0.282	0.220
Mediation	1.000	0.282	0.282	0.220

In Table 3b, the Bayes Factor and posterior odds for the mediation effect are 0.282 and 0.220, respectively, which are the same as those for either individual path (α or β). This makes sense given the specified prior odds. Specifically, by specifying the same prior odds for the two paths and mediation, the BHT for either individual path and the BHT for the mediation effect are comparing the same two models: the model where the two paths are both present vs. the model where they are both absent, because the models where one path is present and the other path is absent are impossible to be true given the prior odds specifications.

Discussion

Mediation analysis plays an important role in theory development and testing in many research fields. To test mediation, many frequentist methods have been proposed and evaluated. Despite the growing popularity of applying Bayesian methods to estimating mediation effects, Bayesian approaches to testing mediation effects are under-developed and under-evaluated. In this article, we studied the use of

[14] In Table 3a, the prior odds, Bayes factors and posterior odds of the paths α and β are computed by first computing the prior probabilities of the models $\mathbf{M}_{11}$, $\mathbf{M}_{01}$, $\mathbf{M}_{10}$, and $\mathbf{M}_{00}$ as $p_{11} = PriorOdds^{med}/(1 + PriorOdds^{med})$, $p_{01} = (1 - p_{11})q_{01|0}$, $p_{10} = (1 - p_{11})q_{10|0}$ and $p_{00} = (1 - p_{11})q_{00|0}$, respectively. Then the prior odds, Bayes factor, and posterior odds for a path, for example, path α are computed as $PriorOdds^{\alpha} = \frac{p_{11}+p_{10}}{p_{01}+p_{00}}$, $BF^{\alpha} \approx \frac{\frac{p_{11}}{p_{11}+p_{10}} e^{-\text{BIC}\{\mathbf{M}_{11}:\alpha\neq 0,\beta\neq 0\}/2} + \frac{p_{10}}{p_{11}+p_{10}} e^{-\text{BIC}\{\mathbf{M}_{10}:\alpha\neq 0,\beta=0\}/2}}{\frac{p_{01}}{p_{01}+p_{00}} e^{-\text{BIC}\{\mathbf{M}_{01}:\alpha=0,\beta\neq 0\}/2} + \frac{p_{00}}{p_{01}+p_{00}} e^{-\text{BIC}\{\mathbf{M}_{00}:\alpha=0,\beta=0\}/2}}$ and $PosteriorOdds^{\alpha} = PriorOdds^{\alpha} BF^{\alpha} \approx \frac{p_{11} e^{-\text{BIC}\{\mathbf{M}_{11}:\alpha\neq 0,\beta\neq 0\}/2} + p_{10} e^{-\text{BIC}\{\mathbf{M}_{10}:\alpha\neq 0,\beta=0\}/2}}{p_{01} e^{-\text{BIC}\{\mathbf{M}_{01}:\alpha=0,\beta\neq 0\}/2} + p_{00} e^{-\text{BIC}\{\mathbf{M}_{00}:\alpha=0,\beta=0\}/2}}$, respectively.

BHT for mediation testing. Our study can add to the literature in the following ways.

First, we proposed a general approach to performing BHT of mediation (i.e., obtaining mediation Bayes factor and posterior odds). Our approach is applicable to different research contexts and can facilitate the integration of data evidence and researchers' prior beliefs. Specifically, the proposed approach can be applied to testing mediation effects in both the independent and dependent paths scenarios. For the independent paths scenario, we provided two ways to specify prior odds (for the two paths α and β; or for one path and the mediation effect) for implementing the proposed approach, and demonstrated the implementation with the simple mediation model. For the dependent paths scenario, we considered the latent mediation model and implemented the proposed approach with two ways of prior odds/probability specification (specify the mediation prior odds and the conditional prior probabilities under H_0; specify the mediation prior odds and two path prior odds). Both R functions (fun_BHTmed.R) and an R web app (BayesianHypothesisTestingMediation) were developed for implementing the techniques presented in the two scenarios, and can be downloaded from https://github.com/xliu12/BHT.med.

Second, we evaluated the sensitivity of BHT of mediation to the prior odds specification. Specifically, through nine hypothetical data examples in the independent paths scenario, we demonstrated that the prior odds of an individual path and the prior odds of the mediation effect can both change the mediation Bayes factor and posterior odds. In particular, we found that the mediation Bayes factor can be influenced by the mediation prior odds, which is unlike the Bayes factor for a single relation (e.g., a regression coefficient). We hope our demonstration can provide insights on the prior odds specification in BHT of mediation.

Recommendations

Based on the findings, we recommend researchers to pay attention to specifying the prior odds/probabilities for BHT of mediation. The specification should be based on their prior knowledge about the mediation process. When there is no prior knowledge available, in the scenario of independent paths a researcher may consider specifying a prior odds of 1 for each path (i.e., $PriorOdds^{\alpha} = PriorOdds^{\beta} = 1$) to reflect that the researcher is "in a state of genuine uncertainty" (Freedman 1987, p. 141) regarding the presence or absence of each path. Note that with this specification, the prior odds in favor of mediation is 1/3. That is, the researcher is conservative about mediation being present.

When the two paths are dependent, a researcher might consider specifying the prior odds of mediation as 1 (i.e., $PriorOdds^{med} = 1$), so as to be neutral about the null and alternative mediation hypotheses to be tested. If a researcher has no prior knowledge about which of the three scenarios under the null hypothesis are more/less likely to occur, the researcher may consider specifying the three conditional prior probabilities to be all equal 1/3 (i.e., $q_{00|0} = q_{01|0} = q_{10|0} = 1/3$). Note that with this specification, the prior odds of each path is 2 (i.e., $PriorOdds^{\alpha} = PriorOdds^{\beta} = 2$). In other words, the researcher believes that each path is more likely to be present than be absent a priori.

Furthermore, we recommend researchers to conduct sensitivity analysis to examine how the mediation Bayes factor and posterior odds would change with a list of reasonable prior odds/probabilities specifications. If the results are relatively stable, the researcher may corroborate his/her posterior belief about the presence vs. absence of mediation. Nonetheless, we suggest that the researcher should report the results from sensitivity analysis and provide a rationale for the chosen prior odds/probability specification.

Last but not least, we echo previous literature (e.g., Tendeiro & Kiers 2019) and recommend that researchers should report both the Bayes factor and posterior odds of mediation, and should distinguish between their interpretations. Specifically, the mediation Bayes factor quantifies the relative likelihood of observing the data if mediation is present vs. if mediation is absent; whereas the mediation posterior odds measures the relative posterior probability of the presence vs. absence of mediation given the observed data. We also recommend researchers to provide the Bayes factors of the two paths when the paths are independent, or the BIC values of the four models (i.e., models $\mathbf{M}_{11}$, $\mathbf{M}_{01}$, $\mathbf{M}_{10}$, and $\mathbf{M}_{00}$) when the paths can be dependent. With the information, readers can evaluate the Bayes factor and posterior odds of mediation based on their own prior beliefs.

Limitations and future directions

The current study can be extended in the future. First, although the ability to provide a continuous measure of the relative data evidence and posterior probabilities is argued to be an advantage of the proposed BHT approach, in practice, researchers may be interested in making a binary decision on the presence of mediation. In the NHST framework, a binary decision of rejecting or not rejecting the null is often made by comparing the p-value to a nominal Type I error rate (e.g., 0.05). While formal BHT has no concept of Type I errors, such a binary decision can be made by comparing Bayes factors with a certain cutoff value. For testing a single relation (e.g., a regression coefficient), commonly used Bayes factor cutoffs include the evidence categories proposed by Jeffreys (1961) and Kass & Raftery (1995), and the cutoff value 3 for Bayes factors is often used as the counterpart of the cutoff 0.05 for p-values (Hoijtink et al., 2019; Kelter, 2020). However, the Bayes factor cutoffs for testing a single relation may not be suitable for mediation testing, especially considering the impact of prior odds specifications on the mediation Bayes factor. It is our important future direction to examine what cutoffs for the mediation Bayes factor can be used to maintain certain nominal Type I error rates (e.g., 0.05) given a set of prior odds specifications.

Furthermore, with the determined cutoff values, it is of interest to examine the performance of the proposed BHT in terms of statistical power, and compare the BHT approach with other mediation testing methods such as the frequentist tests (e.g., MacKinnon, Lockwood, Hoffman, West, & Sheets 2002) and the posterior credible interval approach (e.g., Yuan & MacKinnon 2009). In addition, previous studies have found the power stagnation issue associated with frequentist tests of mediation (e.g., Fritz, Taylor, & MacKinnon 2012).[15] We are interested in investigating whether the BHT approach also has the power stagnation issue.

Second, in complex mediation analyses, there could be multiple purported mediators and researchers may be interested in testing the total mediation effects (e.g., Selig & Preacher 2009), or in testing contrasts between two or more mediation effects (e.g., MacKinnon 2000). How to extend the proposed BHT approach to testing more complex mediation hypotheses in multiple mediator models is a worthwhile future direction.

Third, in addition to testing a point null hypothesis, BHT methods for testing interval null hypotheses (e.g., Morey & Rouder 2011) and for testing informative hypotheses (e.g., Hoijtink, Klugkist, & Boelen 2008) have been proposed. It is of interest to explore the feasibility and adaptation of these methods for mediation analysis.

Fourth, in the literature of BHT, it has been shown that the Bayes factor can be sensitive to the within-model prior distributions (e.g., Du, Edwards, & Zhang 2019; Liu & Aitkin 2008; Tendeiro & Kiers 2019). It is an important future direction to examine how the mediation Bayes factor and posterior odds would change with the within-model prior distributions, given a prior odds/probability specification.

Fifth, BHT has been shown to be useful for sequential data analysis (e.g., Hoijtink, Mulder, van Lissa, & Gu 2019; Rouder 2014). We are interested in utilizing the proposed BHT to develop sequential data analysis procedures for mediation studies.

Conclusion

We hope the current study can contribute a useful BHT approach to researchers' toolbox of mediation analysis and provide insights on the impact of prior odds specifications in the BHT of mediation.

[15] That is, as the magnitude of path α increases, there can be a stagnation or reduction in power of testing mediation.

Acknowledgements Lijuan Wang and Zhiyong Zhang are grateful for the financial support from IES grant R305D210023 during the study. Lijuan Wang is also grateful for the financial support from NIH grants R01MD014737, R01HD091235, R01HD088482, and R01HD087319 during the study.

Availability of data and material The data in this study are based on simulation. None of the data or materials are related to any experiments.

References

Anderson, S. F. (2020). Misinterpreting p: The discrepancy between p values and the probability the null hypothesis is true, the influence of multiple testing, and implications for the replication crisis. *Psychological Methods, 25*(5), 596–609. https://doi.org/10.1037/met0000248.

Baron, R. M., & 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), 1173–1882. https://doi.org/10.1037/0022-3514.51.6.1173.

Berger, J. O., & Sellke, T. (1987). Testing a point null hypothesis: The irreconcilability of p values and evidence. *Journal of the American Statistical Association,82*(397), 112–122. http://www.jstor.org/stable/2289131

Bollen, K. A., Harden, J. J., Ray, S., & Zavisca, J. (2014). BIC and alternative bayesian information criteria in the selection of structural equation models. *Structural Equation Modeling: A Multidisciplinary Journal, 21*(1), 1–19. https://doi.org/10.1080/10705511.2014.856691.

Che, C., Jin, I. H., & Zhang, Z. (2021). Network mediation analysis using model-based eigenvalue decomposition. *Structural Equation Modeling: A Multidisciplinary Journal, 28*(1), 148–161. https://doi.org/10.1080/10705511.2020.1721292.

Dienes, Z. (2019). How do I know what my theory predicts? *Advances in Methods and Practices in Psychological Science, 2*(4), 364–377. https://doi.org/10.1177/2515245919876960.

Du, H., Edwards, M. C., & Zhang, Z. (2019). Bayes factor in one-sample tests of means with a sensitivity analysis: A discussion of separate prior distributions. *Behavior Research Methods, 51*(5), 1998–2021. https://doi.org/10.3758/s13428-019-01262-w.

Elliot, D. L., Goldberg, L., Kuehl, K. S., Moe, E. L., Breger, R. K. R., & Pickering, M. A. (2007). The PHLAME (Promoting Healthy Lifestyles: Alternative Models' Effects) firefighter study: Outcomes of two models of behavior change. *Journal of Occupational and Environmental Medicine, 49*(2), 204–213. https://doi.org/10.1097/jom.0b013e3180329a8d.

Freedman, B. (1987). Equipoise and the ethics of clinical research. *New England Journal of Medicine, 317*(3), 141–145. https://doi.org/10.1056/nejm198707163170304.

Fritz, M. S., Taylor, A. B., & MacKinnon, D. P. (2012). Explanation of two anomalous results in statistical mediation analysis. *Multivariate Behavioral Research, 47*(1), 61–87. https://doi.org/10.1080/00273171.2012.640596.

Harlow, L. L., Mulaik, S. A., & Steiger, J. H. (1997). *What if there were no significance tests?* Mahwah, NJ: Lawrence Erlbaum Associates Publishers.

Hoijtink, H., Klugkist, I., & Boelen, P. A. (Eds.). (2008). *Bayesian evaluation of informative hypotheses*. New York: Springer.

Hoijtink, H., Mulder, J., van Lissa, C., & Gu, X. (2019). A tutorial on testing hypotheses using the Bayes factor. *Psychological Methods, 24*(5), 539–556. https://doi.org/10.1037/met0000201.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American statistical Association, 81*(396), 945–960. https://doi.org/10.1080/01621459.1986.10478354.

Ioannidis, J. P. A. (2019). What have we (not) learnt from millions of scientific papers with p values? *The American Statistician, 73*(sup1), 20–25. https://doi.org/10.1080/00031305.2018.1447512.

Iverson, G. J., Lee, M. D., Zhang, S., & Wagenmakers, E.-J. (2009). prep: An agony in five fits. *Journal of Mathematical Psychology, 53*(4), 195–202. https://doi.org/10.1016/j.jmp.2008.09.004.

JASP Team (2020). JASP (Version 0.14.1)[Computer software]. https://jasp-stats.org/

Jeffreys, H. (1961) Theory of probability. Oxford, UK: Oxford University Press, 3rd ed.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association,90*(430), 773–795. http://www.jstor.org/stable/2291091

Kelter, R. (2022). Power analysis and type I and type II error rates of Bayesian nonparametric two-sample tests for location-shifts based on the Bayes factor under Cauchy priors. *Computational Statistics & Data Analysis, 165,* 1–20. https://doi.org/10.1016/j.csda.2021.107326.

Liao, J. G., Midya, V., & Berg, A. (2020). Connecting and contrasting the Bayes factor and a modified ROPE procedure for testing interval null hypotheses. *The American Statistician, 75*(3), 1–19. https://doi.org/10.1080/00031305.2019.1701550.

Liu, C. C., & Aitkin, M. (2008). Bayes factors: Prior sensitivity and model generalizability. *Journal of Mathematical Psychology, 52*(6), 362–375. https://doi.org/10.1016/j.jmp.2008.03.002.

Ly, A., Verhagen, J., & Wagenmakers, E.-J. (2016). Harold Jeffreys's default Bayes factor hypothesis tests: Explanation, extension, and application in psychology. *Journal of Mathematical Psychology, 72,* 19–32. https://doi.org/10.1016/j.jmp.2015.06.004.

MacKinnon, D. P. (2000). Contrasts in multiple mediator models. In J. S. Rose, L. Chassin, C. C. Presson, & S. J. Sherman (Eds.), *Multivariate applications in substance use research: New methods for new questions* (pp. 141–160). Lawrence Erlbaum Associates Publishers.

MacKinnon, D. P. (2008). *Introduction to statistical mediation analysis*. New York, NY: Taylor & Francis Group/Lawrence Erlbaum Associates.

MacKinnon, D. P., Fritz, M. S., Williams, J., & Lockwood, C. M. (2007). Distribution of the product confidence limits for the indirect effect: Program prodclin. *Behavior Research Methods, 39*(3), 384–389. https://doi.org/10.3758/BF03193007.

MacKinnon, D. P., Lockwood, C. M., Hoffman, J. M., West, S. G., & Sheets, V. (2002). A comparison of methods to test mediation and other intervening variable effects. *Psychological Methods, 7*(1), 83. https://doi.org/10.1037/1082-989X.7.1.83.

Merkle, E. C., & Wang, T. (2016). Bayesian latent variable models for the analysis of experimental psychology data. *Psychonomic Bulletin & Review, 25*(1), 256–270. https://doi.org/10.3758/s13423-016-1016-7.

Miočević, M., Gonzalez, O., Valente, M. J., & MacKinnon, D. P. (2018). A tutorial in Bayesian potential outcomes mediation analysis. *Structural Equation Modeling: A Multidisciplinary Journal, 25*(1), 121–136. https://doi.org/10.1080/10705511.2017.1342541.

Morey, RD, Rouder, JN, Jamil, T. (2018) BayesFactor: Computation of Bayes Factors for common designs. R package version 0.9. 12-4.2. [Computer software]. Retrieved from https://CRAN.R-project.org/package=BayesFactor

Morey, R. D., & Rouder, J. N. (2011). Bayes factor approaches for testing interval null hypotheses. *Psychological Methods, 16*(4), 406–419. https://doi.org/10.1037/a0024377.

Mulder, J., & Wagenmakers, E.-J. (2016). Editors' introduction to the special issue "Bayes factors for testing hypotheses in psychological research: Practical relevance and new developments''. *Journal of Mathematical Psychology, 72,* 1–5. https://doi.org/10.1016/j.jmp.2016.01.002.

Muthén, B., Muthén, L. (2017) Mplus. Chapman and Hall/CRC

Nickerson, R. S. (2000). Null hypothesis significance testing: A review of an old and continuing controversy. *Psychological Methods, 5*(2), 241–301. https://doi.org/10.1037/1082-989x.5.2.241.

Nuijten, M. B., Wetzels, R., Matzke, D., Dolan, C. V., & Wagenmakers, E.-J. (2014). A default Bayesian hypothesis test for mediation. *Behavior Research Methods, 47*(1), 85–97. https://doi.org/10.3758/s13428-014-0470-2.

Pearl, J. (2001) Direct and indirect effects. In: Proceedings of the seventeenth conference on uncertainty in artificial intelligence, UAI'01 (pp. 411–420). San Francisco, CA, USA: Morgan Kaufmann Publishers Inc. http://dl.acm.org/citation.cfm?id=2074022.2074073

Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology,25,* 111–163. http://www.jstor.org/stable/271063

Rosseel, Y., Jorgensen, T.D., Rockwood, N., Oberski, D., Byrnes, J., Vanbrabant, L., Savalei, V., ... et al. (2020). Package 'lavaan'.

Rouder, J. N. (2014). Optional stopping: No problem for bayesians. *Psychonomic Bulletin & Review, 21*(2), 301–308. https://doi.org/10.3758/s13423-014-0595-4.

RStudio, Inc (2021) Web application framework for R. [Computer software manual]. https://shiny.rstudio.com/.

Selig, J. P., & Preacher, K. J. (2009). Mediation models for longitudinal data in developmental research. *Research in Human Development, 6*(2–3), 144–164. https://doi.org/10.1080/15427600902911247.

Tendeiro, J. N., & Kiers, H. A. L. (2019). A review of issues about null hypothesis Bayesian testing. *Psychological Methods, 24*(6), 774–795. https://doi.org/10.1037/met0000221.

Tofighi, D., & Kelley, K. (2020). Improved inference in mediation analysis: Introducing the model-based constrained optimization procedure. *Psychological Methods, 25*(4), 496–515. https://doi.org/10.1037/met0000259.

Valente, M. J., Gonzalez, O., Miočević, M., & MacKinnon, D. P. (2016). A note on testing mediated effects in structural equation models: Reconciling past and current research on the performance of the test of joint significance. *Educational and Psychological Measurement, 76*(6), 889–911. https://doi.org/10.1177/0013164415618992.

Valeri, L., & VanderWeele, T. J. (2013). Mediation analysis allowing for exposure-mediator interactions and causal interpretation: Theoretical assumptions and implementation with SAS and SPSS macros. *Psychological Methods, 18*(2), 137. https://doi.org/10.1037/a0031034.

Wagenmakers, E. J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., et al. (2018). Bayesian inference for psychology. part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin & Review, 25*(1), 35–57. https://doi.org/10.3758/s13423-017-1343-3.

Wang, L., & Preacher, K. J. (2015). Moderated mediation analysis using Bayesian methods. *Structural Equation Modeling: A Multidisciplinary Journal, 22*(2), 249–263. https://doi.org/10.1080/10705511.2014.935256.

Wetzels, R., & Wagenmakers, E.-J. (2012). A default Bayesian hypothesis test for correlations and partial correlations. *Psychonomic Bulletin & Review, 19*(6), 1057–1064. https://doi.org/10.3758/s13423-012-0295-x.

Yuan, Y., & MacKinnon, D. P. (2009). Bayesian mediation analysis. *Psychological Methods, 14*(4), 301. https://doi.org/10.1037/a0016972.