Published in final edited form as:

J R Stat Soc Ser C Appl Stat. 2018 January ; 67(1): 127–144. doi:10.1111/rssc.12220.

Bayesian mixed treatment comparisons meta-analysis for correlated outcomes subject to reporting bias

Yulun Liu,

Department of Biostatistics, The University of Texas Health Science Center Houston, Houston, Texas 77030, U.S.A

Stacia M. DeSantis, and

Department of Biostatistics, The University of Texas Health Science Center Houston, Houston, Texas 77030, U.S.A

Yong Chen

Department of Biostatistics and Epidemiology, University of Pennsylvania, Philadelphia, Pennsylvania, 19104, U.S.A

Summary

Many randomized controlled trials (RCTs) report more than one primary outcome. As a result, multivariate meta-analytic methods for the assimilation of treatment effects in systematic reviews of RCTs have received increasing attention in the literature. These methods show promise with respect to bias reduction and efficiency gain compared to univariate meta-analysis. However, most methods for multivariate meta-analysis have focused on pairwise treatment comparisons (i.e., when the number of treatments is two). Current methods for mixed treatment comparisons (MTC) meta-analysis (i.e., when the number of treatments is more than two) have focused on univariate or very recently, bivariate outcomes. To broaden their application, we propose a framework for MTC meta-analysis of multivariate (≥ 2) outcomes where the correlations among multivariate outcomes within- and between-studies are accounted for through copulas, and the joint modeling of multivariate random effects, respectively. We consider a Bayesian hierarchical model using Markov Chain Monte Carlo methods for estimation. An important feature of the proposed framework is that it allows for borrowing of information across correlated outcomes. We show via simulation that our approach reduces the impact of outcome reporting bias (ORB) in a variety of missing outcome scenarios. We apply the method to a systematic review of RCTs of pharmacological treatments for alcohol dependence, which tends to report multiple outcomes potentially subject to ORB.

Keywords

Bayesian model; Mixed treatment comparison; Multivariate Meta-analysis; Network meta-analysis; Publication bias; Systematic review

1. Introduction

The increased attention to evidence-based medicine has resulted in the growth of systematic reviews and meta-analytic methods in recent years. In the Cochrane Database of systematic

reviews, mixed treatment comparisons (MTC), or network structures, are commonly encountered. These types of systematic reviews form a graphical network of treatment comparisons across potentially multi-arm clinical trials. Mixed treatment comparisons meta-analysis allows for synthesis of evidence structures that contain both direct evidence about treatments (e.g., A:B trials) and “indirect” evidence about two treatments (e.g., evidence for A:B obtained from A:C and B:C via a common comparator) (Lumley, 2002; Lu and Ades, 2006). Statistical methods for MTCs applied to univariate outcomes have been well-studied in the last decade; see, for examples, Lumley (2002), Lu and Ades (2006), Caldwell et al. (2005), and Salanti et al. (2008).

Since most clinical trials report more than one outcome, MTC structures with multivariate (multiple) outcomes are commonly encountered. These structures are also frequently subject to missingness at the outcome-level, that is, not all outcomes are reported in all studies. The choice of whether to report an outcome in a given study may be influenced by the significance or direction of the effect size for that outcome, known as outcome reporting bias (ORB) (Copas and Shi, 2001; Kirkham et al., 2012; Copas, 2013). If a univariate MTC meta-analysis was performed on each outcome separately, results could be subject to the effects of publication bias (PB); i.e., it is as if studies that do not report a given outcome went unreported since no information is available about that outcome. It has been well-acknowledged that ignoring ORB or PB can lead to biased estimates of overall effect sizes in pairwise meta-analysis (Riley et al., 2007; Egger et al., 2008).

Advantages of multivariate pairwise meta-analysis in terms of reduction in bias and efficiency of pooled effect sizes have been studied (Riley et al., 2007; Jackson et al., 2011); however, multivariate methods to address challenges of ORB/PB in the MTC setting are scarce. Two very recent papers present bivariate MTC meta-analysis and apply it to a network of acute mania (Efthimiou et al., 2014, 2015). Another applies a multivariate MTC meta-analysis using a similar model to a systematic review of poison prevention strategies with 3 outcomes (Achana et al., 2014). Specifically, Efthimiou et al. (2015) thoughtfully present several extensive decompositions of the variance-covariance matrix and corresponding WinBUGS to model both treatment and outcome correlation in this complex setting. While these methods could be extended to “true” multivariate outcome vectors (i.e., 2), a data application for multivariate outcomes for which covariance matrices become more unwieldy in practice is not considered. Further, as Efthimiou et al. (2014, 2015) analyze the same data, additional substantive evidence about the practicality of application would be useful. Finally, as the multivariate methodology is likely most useful in the setting of ORB, a comprehensive simulation study to determine behavior of the method when outcomes are subject to various types of ORB is warranted.

We address the above challenges by modeling multiple (2) outcomes within-studies using copulas, and between-studies by joint modeling of multivariate study-specific random effects. The formulation relies on the fact that certain outcomes in a given subject area (such as efficacy, or discontinuation, depending on the topic area) will almost always be reported in a paper order for an RCT to be accepted for publication. This assumption is discussed further in the Application Section. Specifically, we hypothesize that modeling correlated outcomes in a multivariate MTC framework will lessen the effect of outcome reporting bias

when outcomes are missing at random (MAR) and missing not at random (MNAR), using the nomenclature of Rubin (1976).

The layout of the rest of the paper is as follows. In Section 2, we introduce the motivating data, a systematic review of RCTs for alcohol dependence where 74% of studies do not report all of the 3 primary drinking outcomes. Evidence of publication bias and outcome reporting bias is shown therein. In Section 3, we present a Bayesian multivariate MTC (MMTC) meta-analytic model, a multivariate assessment of evidence consistency, and a multivariate posterior probability-based treatment ranking procedure. Parameter estimation is powered by Markov Chain Monte Carlo (MCMC) methods. In Section 4, we perform simulation studies to assess the ability of our approach to mitigate effects of outcome reporting bias. In Section 5, we apply the proposed method to the motivating dataset and assess sensitivity of inference to prior specification for the variance-covariance matrix. In Section 6, we present a discussion.

2. Motivating study

The motivating study is a systematic review of pharmacological treatments for alcohol dependence, published in two consecutive Cochrane reviews of naltrexone (NAL) and acamprosate (ACA) (Srisurapanont and Jarusuraisin, 2005; Rösner et al., 2010). The original reviews were updated as of 2014 (DeSantis and Zhu, 2014). A detailed description of the methods undertaken for data extraction can be found in the full reviews available in the Cochrane Library and the entire data set is provided in Section E of the Supplementary Material. We note that many systematic reviews in substance dependence and mental health also give rise to multivariate MTC evidence structures, thus the methods proposed here are widely applicable.

Consistent outcome reporting for alcohol dependence has not been adopted across clinical trials making meta-analysis difficult (Shirley et al., 2010; DeSantis et al., 2013). Table 1 shows the number of RCTs missing each of 3 primary outcomes of interest in the current analysis (RH = return to heavy drinking, RD = return to drinking, and DIS = discontinuation where “discontinuation” refers to discontinuing the assigned medication in the course of a trial). If one were to analyze RH alone, only 13 studies would be included in the analysis. However, since DIS is almost always reported in trials of alcohol dependence (38/41 trials) and as we will show that it is not subject to publication bias, it provides unbiased outcome information from which to borrow. Thus, a joint analysis of RH with the other two outcomes allows all 41 studies to provide information about the multivariate pooled effect sizes; this joint analytic strategy can potentially reduce the impact of ORB. We note that the scenario of publication bias in this multivariate analysis would only arise when none of the outcomes is reported (i.e., the entire study is not published).

Fig. 1 presents the network structures of clinical trials for each of the three binary outcomes. The network reveals that active treatments (i.e., ACA and NAL) have been most often compared to placebo, and that few direct comparisons of active treatments exist. We assess the extent of publication bias using funnel plots and Egger’s test (Egger et al., 1997) for each of the drinking outcomes. Supplementary Figure 1S reveals publication bias in the

reporting of some outcomes. In particular, for RH in naltrexone vs placebo studies, there is a clear asymmetry in the plot with large values missing from the bottom quadrant (Egger's test $p = 0.025$). For RD in naltrexone vs placebo studies, there is also asymmetry, with small values missing (Egger's test $p = 0.045$). For RD in acamprosate vs placebo studies, the large values in the bottom right quadrant are missing (Egger's test $p = 0.070$). In addition to this pairwise series of funnel plots, we present the 'comparison-adjusted' funnel plots excluding head-to-head comparisons, as the choice of comparator would be arbitrary. (Chaimani et al., 2013) (Fig. 2). In the 'comparison-adjusted' funnel plot, the horizontal axis represents the difference between the observed log odds ratio of the newer versus older treatments from the direct summary estimate (i.e., the comparison specific pooled estimates from the pairwise meta-analysis based on the fixed-effects model (Chaimani and Salanti, 2012)), and the vertical axis represents the standard error of the observed log odd ratios. An asymmetric 'comparison-adjusted' funnel plot may indicate the existence of small-study effects (Chaimani and Salanti, 2012; Chaimani et al., 2013). As suggested in Chaimani et al. (2013), treatments should be ordered in a meaningful way. We assume the placebo is the oldest followed by NAL, ACA and NAL+ACA treatments. Hence, small studies on the right side of the funnel plot provide evidence that small studies tend to favor the old treatment rather than the new one. As shown in Fig. 2, there is asymmetry for outcomes RH and RD, indicating that smaller studies do favor the active treatments (i.e., NAL or ACA) as more effective than placebo. These results also show evidence of publication bias for RH and RD (but not discontinuation) when outcomes are considered univariately. Thus it seems sensible to borrow information across outcomes for network meta-analysis, especially as the (almost) always reported outcome, discontinuation, likely contains information about unreported RH and RD.

3. Statistical Methodology

3.1. Model Description

Let $\{(y_{ikl}, n_{ikl}), i = 1, 2, \dots, N; k \in T_i; l \in L_i\}$ represent the binary summary data obtained from the randomized controlled trials, where y_{ikl} and n_{ikl} are the number of positive (or negative) responses and subjects, respectively, for the l th outcome of the k th treatment in the i th study. Let T_i be the set of all treatments compared in study i and L_i denote the set of outcomes reported in study i . For example, if outcome 2 is missing in study i , then the vector $y_{ik} = \{y_{ik1}, y_{ik3}\}$. We assume that within a study, all treatment arms report the same set of outcomes. This assumption is met in our motivating alcohol dependence study and would probably hold for the majority of networks.

We assume $y_{ikl} \sim \text{Binom}(n_{ikl}, p_{ikl})$, where $\text{Binom}(\cdot; \cdot)$ denotes the binomial distribution, and p_{ikl} is the probability of the l th outcome being observed in the k th treatment of the i th study. Let $\text{logit}(p_{iXl}) = \mu_{il}$ where X is the placeholder for the baseline treatment in study i , and $\text{logit}(p_{iYl}) = \mu_{il} + \delta_{(X:Y)l}$ for $Y \neq X$. We specify the model through a contrast based approach as in Lu and Ades (2006). Specifically, $\delta_{(X:Y)l}$ is the log odds ratio (referred as LOR) for treatment Y relative to treatment X in study i for outcome l . Similar to the univariate MTC (UMTC) setting, we define *basic parameters* as the effect sizes comparing treatments with placebo, and the *functional parameters* as the effect sizes comparing active

treatments. Under the consistency assumption, the functional parameters can be written as functions of basic parameters.

For simplicity of notation, we consider the simplest setting with a total of two outcomes ($I = 1, 2$) and three treatment arms, treatments A and B and placebo P, i.e., $k \in \{P, A, B\}$. If study i includes the placebo arm and both outcomes are reported, the vector of LOR for study i is written as $\delta_i = (\delta_{(P:A)1}^i, \delta_{i(P:B)1}^i, \delta_{i(P:A)2}^i, \delta_{i(P:B)2}^i)^T$. In this case, it is easy to see that the probability p_{ikl} can be written as

$$p_{ikl} = \text{logit}^{-1}(\mu_{il} + X_{ik} \delta_{(P:k)l}^i), \quad (1)$$

where $X_{ik} = 1$ if $k = A$ or B , and 0 if $k = P$. If only outcome 1 or 2 is reported, equation (1) holds for $l \in L_i$.

If the study i does not include the placebo arm (i.e., $T_i = \{A, B\}$), the probability p_{ikl} can be written as

$$p_{ikl} = \text{logit}^{-1}(\mu_{il} + X_{ik} \delta_{(b(i):k)l}^i), \quad (2)$$

where $k \in T_i$, $l \in L_i$, $b(i)$ is the specified baseline treatment of study i , and $X_{ik} = 1$ if $k = b(i)$ and 0 otherwise. It is easy to see that equation (2) is a generalization of equation (1).

The random effects, δ_i , where placebo is set as the baseline are assumed to follow a multivariate normal distribution

$$\delta_i \sim MVN(\mathbf{d}, \mathbf{V}), \quad (3)$$

where $\mathbf{d} = (d_{(P:A)1}, d_{(P:B)1}, d_{(P:A)2}, d_{(P:B)2})^T$ are the pooled effect sizes across studies comparing treatments A and B with placebo for outcomes 1 and 2, and are referred to as the basic parameters. The covariance matrix, $\mathbf{V}$, is defined as

$$\mathbf{V} = \begin{pmatrix} \sigma_{(P:A)1}^2 & \rho_{(P:AB)11} \sigma_{(P:A)1} \sigma_{(P:B)1} & \rho_{(P:AA)12} \sigma_{(P:A)1} \sigma_{(P:A)2} & \rho_{(P:AB)12} \sigma_{(P:A)1} \sigma_{(P:B)2} \\ & \sigma_{(P:B)1}^2 & \rho_{(P:AB)12} \sigma_{(P:B)1} \sigma_{(P:A)2} & \rho_{(P:BB)12} \sigma_{(P:B)1} \sigma_{(P:B)2} \\ & & \sigma_{(P:A)2}^2 & \rho_{(P:AB)22} \sigma_{(P:A)2} \sigma_{(P:B)2} \\ & & & \sigma_{(P:B)2}^2 \end{pmatrix},$$

where the diagonal elements of $\mathbf{V}$ account for between-study heterogeneity in the LOR comparing the risk of the l th outcome for the k th treatment with placebo ($l = 1, 2$ and $k = A, B$), and the off-diagonal parameters, $\rho_{(P:AA)l'l'}$, $\rho_{(P:AB)l'l'}$, $\rho_{(P:BB)l'l'}$ ($l, l' = 1, 2$),

characterize the between-study correlations among different treatments and outcomes. Importantly, these between-study correlations allow for borrowing of information between-studies across different outcomes and across treatments.

Under the consistency assumption, the functional parameters, $(d_{(A:B)1}, d_{(A:B)2})$, are written as functions of the basic parameters,

$$d_{(A:B)l} = d_{(P:B)l} - d_{(P:A)l} \quad \text{for } l=1, 2. \quad (4)$$

Also, the random effects in equation (2) are typically assumed to follow the multivariate normal distribution

$$\begin{pmatrix} \delta_{(A:B)1}^i \\ \delta_{(A:B)2}^i \end{pmatrix} \sim MVN \left(\mathbf{X}_i \mathbf{d}, \mathbf{X}_i \mathbf{V} \mathbf{X}_i^T \right), \quad (5)$$

where $\mathbf{X}_i$ is the design matrix for the i th study, and defined as

$$\mathbf{X}_i = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 1 \end{pmatrix}.$$

To account for the *within-study* correlations among multivariate outcomes, we use a Clayton copula model. Given the study-specific random effects, the multivariate outcome has a joint distribution of $C(u_1, u_2; \theta) = (u_1^{-\theta} + u_2^{-\theta} - 1)^{-1/\theta}$ is the Clayton copula (Clayton, 1978); $\theta_j \in (0, \infty)$ accounts for the within-study correlations, and

$F_l(y_{ikl}) = \sum_{t=0}^{\lfloor y_{ikl} \rfloor} \binom{n_{ikl}}{t} p_{ikl}^t (1-p_{ikl})^{n_{ikl}-t}$ is the cumulative probability mass function of y_{ikl} , and p_{ikl} is the probability of the l th outcome being positive in the k th treatment of the i th study. By Sklar's Theorem (Sklar, 1959), any joint distribution with any type of correlation structure, e.g., positive or negative, has a uniquely determined copula constructed from the marginal distribution. Our formulation is similar to the copulas proposed in Kuss et al. (2014) but other choices of copulas may also be used.

Given the above, we specify the following Bayesian hierarchical MMTC model,

Level 1—The conditional likelihood of the observed correlated data based on the bivariate Clayton copula model given the random effects, δ_{ib} , is

$$\begin{aligned}
& \prod_{i=1}^N \prod_{k \in T_i} f(y_{ik1}, y_{ik2}) \\
&= \prod_{i=1}^N \prod_{k \in T_i} c\{F_1(y_{ik1}), F_2(y_{ik2}); \theta_i\} \prod_{l=1}^2 f_l(y_{ikl}) \\
&= \prod_{i=1}^N \prod_{k \in T_i} \left\{ (1+\theta_i)(u_{ik1}u_{ik2})^{-\theta_i-1} (u_{ik1}^{-\theta_i} + u_{ik2}^{-\theta_i} - 1)^{-2-1/\theta_i} \right\} \prod_{l=1}^2 f_l(y_{ikl}),
\end{aligned}$$

where $c(u_{ik1}, u_{ik2}; \theta_i) = \frac{\partial^2 C(u_{ik1}, u_{ik2}; \theta_i)}{\partial u_{ik1} \partial u_{ik2}}$ is the density corresponding to the Clayton copula $C(u_{ik1}, u_{ik2}; \theta_i)$ and θ_i accounts for within-study correlation.

Level 2—Denote δ_{ib} as the vector of random effects $\delta_{(b(i):k)l}^i$ for pre-specified $b(i)$, $k \in T_i$ and $l \in L_i$ where the distribution of δ_{ib} is specified in equation (5).

Level 3—Complete the specification with choice of priors for the parameters and hyperparameters.

Noninformative, vague through proper prior distributions are chosen; $\pi(\mu_{ij}) \sim N(0, 10000)$, $\pi(d_j) \sim N(0, 10000)$, and $\theta_j \sim \text{Uniform}[1, 10]$. The basic requirement for specifying a prior distribution for V is that it should be nonnegative definite. We assume an unstructured variance-covariance matrix and use an inverse-Wishart prior distribution as described by Wei and Higgins (2013a) for multivariate pairwise meta-analysis; $\pi(V) \sim IW(\Omega, r)$, where IW denotes the inverse-Wishart distribution with scale matrix Ω which is a 4×4 known positive definite matrix, with degrees of freedom, $r = 4$. There are several possible options for priors but when little or no information regarding within- and between-study correlation exists, a sensible option is to perform a sensitivity analysis to assess whether and how conclusions depend on the choice of prior, for example, degrees of freedom of the Wishart distribution. We perform such a sensitivity analysis in Section 5.

3.2. Test of multivariate consistency

The consistency assumption in equation (4) can be extended to test for multivariate evidence consistency. Specifically, for the l th outcome ($l = 1, 2$), we denote the inconsistency parameters, $\omega_{(P:AB)l}$ for $l = 1, 2$ as

$$d_{(A:B)l} = d_{(P:B)l} - d_{(P:A)l} + \omega_{(P:AB)l}, \quad \text{for } l=1, 2. \quad (6)$$

Replacing the formula in equation (4) by equation (6), we can obtain estimates of the inconsistency parameters, $(\omega_{(P:AB)1}, \omega_{(P:AB)2})$, denoted as $\hat{\omega} = (\hat{\omega}_{(P:AB)1}, \hat{\omega}_{(P:AB)2})$, and their covariance matrix, $\hat{\Sigma}_{\omega}$. We propose a multivariate inconsistency measure (MICM) and an outcome-specific inconsistency measure (OICM) defined as,

$$\text{MICM} = \hat{\omega}^T \left(\sum_{\omega} \right)^{-1} \hat{\omega}, \quad \text{and} \quad \text{OICM} = \hat{\omega}_l^T \left(\sum_{\omega}^{(l,l)} \right)^{-1} \hat{\omega}_l,$$

where $\hat{\omega}_{(l)}$ is the subvector of $\hat{\omega}$ corresponding to the l th outcome, and $\sum_{\omega}^{(l,l)}$ is the corresponding submatrix of $\hat{\Sigma}_{\omega}$. By standard asymptotic theory (Casella and Berger, 2002) and multivariate normal results (Graybill, 2000), it can be shown that MICM and OICM approximately follow a chi-squared distribution with $(K-1)(L-1)$ and $(L-1)$ degrees of freedom when the number of studies is large, respectively. The corresponding p-values can be used to test the null hypothesis that direct and indirect evidence are directionally consistent across outcomes.

To develop a multivariate ranking procedure, suppose P_{kl} is the marginal posterior probability of a particular event of outcome, l , under treatment, k , which is modeled from equation (2) using the posterior of d_{kl} and the posterior mean of μ_l across studies. A loss function is defined as $T_{kl} = P_{kl}$ if the event is a positive outcome and $T_{kl} = 1 - P_{kl}$ if the event is a negative outcome, such that treatment with the smallest loss is the best treatment. Then, $\Pr(k \text{ is the best treatment for the } l\text{th outcome} | \text{Data}) = \Pr\{\text{rank}(T_{kl}) = 1 | \text{Data}\}$, where $\text{rank}(T_{kl})$ is the ordinal number of T_{kl} . Therefore, following the univariate procedure, the multivariate ranking probability of the best treatment combination for L outcomes can be defined as $\Pr\{(k_{(1)}, \dots, k_{(L)}) \text{ is the best treatment combination of } L \text{ outcomes}\} = \Pr\{\text{rank}(T_{k_{(1)1}}) \times \dots \times \text{rank}(T_{k_{(L)L}}) = 1 | \text{Data}\}$, where $k_{(l)}$ denotes as the best treatment for outcome l for $l \in \{1, \dots, L\}$.

4. Simulation Study

We assess the finite sample performance of the proposed Bayesian MMTC method under different settings of missingness. For all settings, we consider three treatment arms (i.e., $K = 3$) and assume each study reports a maximum of two outcomes (i.e., $L = 2$). We generate complete data from a two-step procedure. In the first step, we generate the random effects using equation (5). In the second step, we generate correlated binary outcomes in each treatment group from a bivariate Clayton copula model using the “claytonCopula” function in the *copula* package in R (Yan et al., 2007; Kojadinovic et al., 2010; Hofert et al., 2015). For the Clayton copula model, we must specify the parameter value for θ , which is used to measure the correlation between two marginals via the use of Kendall’s tau, τ , i.e., $\theta = 2\tau / (1 - \tau)$. In particular, the parameter, θ , allows one to model a positive or negative correlation in the “claytonCopula” function, where $-1 \leq \tau \leq 1$. As we described in Level 1 of the Bayesian hierarchical MMTC model, in the most flexible case, the parameter θ is allowed to vary across studies; however, for model parsimony in this simulation study, we consider θ fixed as 2 or 8 in each study, corresponding to a correlation between two marginal distributions as 0.5 or 0.8, respectively. These represent moderate and strong within-study correlations, respectively. We consider a moderate size meta-analysis with $N = 50$ studies, where the number of subjects in each study is randomly drawn from a discrete uniform distribution [50,200]. We specify parameters for the pooled effect sizes, which are the weighted average of study-specific effect sizes across studies for treatments 1–3 for

outcomes 1 and 2 as $d_1 = (d_{11}, d_{21}, d_{31}) = (0.3, 0.4, 0.5)$ and $d_2 = (d_{12}, d_{22}, d_{32}) = (0.4, 0.6, 0.4)$, respectively. Recalling that $LOR_{(1:2)1}$ refers to the log odds ratio comparing treatment 2 against treatment 1 for outcome 1, the above parameter specification leads to the following set of log odds ratios: $LOR_{(1:2)1} = 0.442$, $LOR_{(1:3)1} = 0.847$, $LOR_{(2:3)1} = 0.405$, $LOR_{(1:2)2} = 0.811$, $LOR_{(1:3)2} = 0$, and $LOR_{(2:3)2} = -0.811$. The interpretation of other LORs is similar. To evaluate the impact of the true correlation structure on model fitting, we let the correlation parameter take on values of 0.5 and 0.8 to represent moderate and strong correlations between outcome 1 and outcome 2, within treatments 1–3. To generate missing outcomes, we assume outcome 1 is always reported (mirroring the data example) and then construct 30% missing data in the complete data under MCAR, MAR, and MNAR mechanisms. For the MCAR setting, the probability of missingness for outcome 2 is not associated with observed or any unobserved outcomes. For the MAR setting, the probability of missingness for outcome 2 is dependent on the p-values of reporting outcome 1. For the MNAR setting, the probability of missingness for outcome 2 is dependent on the p-values of reporting both outcome 1 and outcome 2. Specifically, we set cut-points of p-values for both outcome 1 and outcome 2 to 0.01, 0.05, 0.10 and 1.00, and let the probability of reporting outcome 2 be dependent on the significance of both outcomes. For example, the probability of reporting outcome 2 is 0.7 if the p-value for outcome 1 is greater than 0.1 and the p-value for outcome 2 is marginally significant (between 0.05 and 0.1). Full details of MNAR data generation are described in Section A of the Supplementary Material. The latter two missingness scenarios result in outcome reporting bias with MAR outcomes being ignorable and MNAR outcomes being nonignorable, respectively.

For each simulation setting, we generate 200 data sets and apply the UMTC and MMTc methods to the resulting networks. As described in Section 3, we set priors μ_{ij} and d_{ij} to be $N(0, 10000)$. For UMTC, we assign the common variance component, σ , which follows a uniform (0, 2) prior and is same for all pairwise comparisons. For MMTc, the specification of the prior distribution for the covariance matrix, V , is an inverse-Wishart distribution, with degrees of freedom $r = 4$. The diagonal elements of the scale matrix, Ω , are set to one and the off-diagonal elements are set to 0.8. We assess a subset of the simulations and do not observe sensitivity to this specification. Two MCMC chains are each run for 20,000 iterations with the first 10,000 discarded as burn-in. To compare the two methods we report the empirical bias of the LOR (average difference between the estimate and the true effect size) and the coverage probability (CP) of the 95% credible intervals for the LORs (Spiegelhalter et al., 2003). The WinBUGS program used to fit the model is presented in Supplementary Material Section D. There we simulate one data set as an example, and as mentioned, assume for the sake of model parsimony that the within-study correlations across studies are fixed.

Fig. 3 summarizes the bias and CPs of LORs under the moderate correlation scenario. Panels (a), (b), (c), (g), (h) and (i) in Fig. 3 indicate both the UMTC and MMTc methods that give approximately unbiased estimates under MCAR. Under MAR, the UMTC method exhibits severely biased LOR estimates for outcome 2 (e.g., $LOR_{(1:2)2}$ and $LOR_{(1:3)2}$) with CPs ranging from 87% to 95%. Compared to the UMTC method, the MMTc method provides substantially less biased estimates. For example, the bias in the estimation of

$LOR_{(1:2)2}$ is reduced by 1/3 (from 0.142 to 0.049). Under MNAR, we do not expect either UMTC or MMTC methods to produce unbiased estimates; however, the MMTC method does reduce the bias as seen in panels (g) and (h). Since the simulation results are based on only 200 simulations, there is non-negligible sampling variation in the estimated CPs estimated to be $\pm 1.96 \sqrt{0.95 \times 0.05/200} = \pm 3\%$. Even given the sampling variation, the CPs of the credible intervals from the MMTC method are consistently closer to the nominal level than those from the UMTC method. Simulation results with strong LOR correlations (i.e., 0.8) are provided in Table S1 of the Supplementary Material and demonstrate even more compelling reduction in bias as within-study outcome correlation increases to 0.8. We note that in some cases, binary outcomes may be negatively correlated. Table S2 in the Supplementary Material shows results generated under a negative correlation structure, where similar trends in bias reduction using MMTC versus UMTC are seen. Such a model involves a change to the copula formulation. It is worth mentioning that the empirical standard errors of estimates from both methods are similar, suggesting that there is little efficiency loss resulting from applying the UMTC versus MMTC method. A similar finding has been reported in a different multivariate meta-analytic setting (Chen et al., 2015). These simulation studies suggest that the main advantage of MMTC over UMTC methods is a reduction in the bias of parameter estimates across a variety of ORB patterns.

In addition to parameter estimation, identification of the “best” treatment is often of interest in MTC framework. In our simulation, we know the true best treatment for outcome 1 is treatment 3, while the best treatment for outcome 2 is treatment 2. Table 2 presents the mean posterior probability of selecting the best treatment combination for the bivariate outcome. The larger the posterior probability of selecting the true best treatment combination, the better the performance. As expected, the joint posterior probability of recovering the correct treatment rankings is larger using the MMTC method compared to the UMTC method. The posterior probabilities of selecting the correct best treatment combination under the two correlation structures are similar.

5. Application

We apply the proposed method to the motivating alcohol dependence study described in Section 2. First, we assess the multivariate evidence consistency using MICM and OICM as described in Section 3. The MICM χ^2 test statistic of the null hypothesis of no deviation from consistency of effect sizes is 6×10^{-5} with p -value close to 1, suggesting consistent directionality of evidence for multivariate outcomes; as this test may be underpowered given the large degree of freedom, such a result should be interpreted with caution. However, the univariate evidence consistency previously was established for this network in DeSantis and Zhu (2014). Since the alcohol dependence network is relatively sparse, we assumed that each study has the common within-study correlation to avoid the issue of weak identifiability. We run two MCMC chains for 50,000 iterations, discarding the first 20,000 as a burn-in to ensure convergence of the chains. Convergence of the resulting samples is assessed using trace plots and history plots as well as Gelman-Rubin diagnostics (provided in Section B of the Supplementary Materials).

As noted earlier, we assess the sensitivity to prior specification. We consider the non-informative inverse-Wishart prior distribution for the covariance matrix V (i.e., $\pi(V) \sim IW(\Omega, r)$) in equation (3) with degrees of freedom of $r = 9, 10, 12, 15$. Table 3 reports the posterior mean of the diagonal elements of scale matrix, Ω , standard errors, and 95% credible intervals for these four model fits. To identify the best fit to the data, we compute the deviance information criterion (DIC), which is defined as the sum of both the posterior mean deviance and the model complexity (Spiegelhalter et al., 2002; Gelman et al., 2004; Congdon, 2005). A difference larger than 10 provides evidence in favor of the model with lower DIC. As shown in Table 3, the posterior means of the diagonal elements, σ^2 , of Ω leads to smaller posterior means with narrower credible intervals with increasing (i.e., more informative) degrees of freedom of the Inverse-Wishart distribution. Since the marginal correlations between effect estimates are all positive, misspecification of the within-study correlation structures had little influence on the results. The DIC favors the model with $r = 9$ coinciding with the smallest possible degrees of freedom to reflect vague prior knowledge. We further assess the sensitivity to the choice of prior for θ by assuming a $U(1, 10)$ prior in addition to assuming it takes on very small fixed values $-0.2, 0.001, 0.2$, corresponding to nearly 0 correlation. These results are summarized in Figure S4 of the Supplementary Material, which shows that when θ is fixed at 0.001 (i.e., the correlation between the two outcomes is close to zero), the inference for some comparisons is affected.

To further assess the validity of model fit with $r = 9$, we consider the conditional predictive ordinate (CPO) (Gelfand et al., 1992; Dey et al., 1995; Chen et al., 2000). The CPO_i is a Monte Carlo approximation based on the harmonic mean of the likelihood of study, i (Dey et al., 1997). It is calculated as $CPO_i^{-1} \approx T^{-1} \sum_{t=1}^T \{L(y_{ikl} | \beta^{(t)})\}^{-1}$, where $\beta^{(t)}$ is parameter vector sampled at iteration $t = 1, \dots, T$, and $L(y_{ikl} | \beta^{(t)})$ is the likelihood of the i th study evaluated at iteration t . As a guideline in Congdon (2005), the harmonic mean estimate of CPO will be stable as long as individual log-likelihoods generally exceed -10 in value. Fig. 4 shows the ordered $-\log(CPO_i)$ values, by study on the x-axis. More than 90% of study-specific log-likelihood values exceed -10 , indicating stability in the harmonic mean estimate. Fig. 4 further flags those individual studies that appear to be outliers in the context of the model fit.

Fig. 5 presents the posterior geometric mean OR and 95% CrIs resulting from the MMTC method (solid line) and UMTC method (dashed line) for treatments versus placebo. The results of the two approaches are generally in the same direction, as would be expected under evidence consistency; however, the CrIs for the MMTC method are smaller, a byproduct of utilizing more data for estimation. The posterior mean estimate [95% CrI] of θ is 3.58 [1.15, 1.57], which is consistent with a moderate to high level of correlation among the three outcomes in the multivariate network. Defining Bayesian significance as the 95% CrI excluding 1, the ORs comparing acamprosate vs placebo and naltrexone vs placebo both show a significant reduction in RH by the UMTC method, but not by the MMTC method. The ORs for naltrexone+acamprosate versus placebo show significant reduction in RH for the MMTC but not the UMTC method. Both MMTC and UMTC methods indicate that naltrexone, acamprosate, and naltrexone+acamprosate show significant reduction in RD vs placebo. The ORs for naltrexone vs placebo and acamprosate vs placebo show greater

discontinuation under the active medications using MMTC but not the UMTc method. The finding that those on active treatments have greater odds of discontinuing their medication during the study than those on placebo using MMTC meta-analysis is an important one. The results demonstrate that the univariate and multivariate methods may result in different inference, although directionality of findings is very consistent. In general, the comparisons of active treatments have wider 95% CRIs because there are few direct comparisons in this network. Due to few direct (head-to-head) comparisons, MTC results should be interpreted with caution.

We report the posterior probabilities of being the k th best treatment combination for all outcomes considered; the posterior probability of the combined treatment being the best treatment for all 3 outcomes is 0.30. The closest competitor is naltrexone with with posterior probability of 0.16. We note that placebo is never the best treatment for the set of multivariate outcomes. Of course, since posterior probabilities are nowhere near 1 for the “best” treatment, results are inconclusive for treatment ranking and should in no way inform treatment decisions.

6. Discussion

In this paper, we developed a Bayesian multivariate mixed treatment comparisons meta-analytic model by accommodating the complex covariance structure that results from correlations among multivariate outcomes and multiple treatments in networks of evidence resulting from systematic reviews. Via simulation, we showed that the proposed framework effectively reduces the impact of outcome reporting bias for a given univariate outcome. We applied the proposed MMTC method to an alcohol dependence study where publication bias and outcome reporting bias are present.

Inference resulting from MMTC is different from that obtained through UMTc which ignores the correlation structure among outcomes. In this particular network, the variable return to heavy drinking is nested within return to drinking. As pointed out by Wei and Higgins (2013b), the correlation may arise due to nesting. In the nested case, the correlation can be inferred from the data as it is derived from a multinomial distribution. However, most mental health networks would not typically report nested outcomes (alcohol dependence is a unique example that does). For example, mental health treatment trials tend to report outcomes such as efficacy (on multiple non-nested scales), discontinuation, and serious adverse events. The method we propose would be applicable to those types of studies even though the current example reports two nested outcomes. We utilize the bivariate Clayton copula because we assume all pairwise correlations in our trivariate data example are the same; thus in our formulation, only one parameter θ need be specified. However, the trivariate Gaussian copula may also be used if it is expected that the correlation differs for different pairs of outcomes.

This research is timely and important because comparative effectiveness research (CER) and medical decisions increasingly rely on systematic reviews and meta-analyses, and MTCs are becoming prevalent data structures. Since RCTs typically report multiple primary outcomes, multivariate MTCs with missing outcomes are becoming increasingly common. The

proposed joint modeling framework can reduce the influence of outcome reporting bias on pooled effect sizes, thereby providing important information for prioritizing new research under CER initiatives.

As widely discussed in the literature, MTC meta-analysis is more assumption-heavy than pairwise meta-analysis. Tests of evidence consistency between direct and indirect evidence can be conducted before MTC meta-analysis is applied, though caution should be exercised as this “2-stage” procedure can influence the type I error rate of the analysis.

There are several limitations of the proposed framework that must be addressed. While more information is available in the multivariate setting, the test for evidence consistency is probably underpowered due to the large degrees of freedom. While in practice, it is probably sufficient to test consistency of each outcome separately, a joint statistical test based on the full covariance structure would be ideal. Future research on such a test is warranted, relying on other properties of the of the covariance matrix for which more powerful tests of evidence consistency could be constructed. Another limitation of the current framework is that the proposed Bayesian procedure is optimal under the missing at random outcome assumption (although we observe bias reduction under MNAR as well). Additionally, when multivariate correlations are very weak the majority of bias remains unadjusted, and therefore extensions of the current methods are warranted. For example, one could incorporate a modified Copas selection model Copas and Shi (2001) into the hierarchical modeling approach to directly model the publication and outcome reporting processes. In our analysis, all arms of all studies are considered in a pairwise fashion; since only 3 studies had more than 2 arms, we did not apply the extended method of Franchini et al. (2012) for maintaining randomization in the multi-arm context. Results would likely only be marginally different had we used this approach. Finally, the proposed framework accommodates binary outcomes - extension to other categorical outcomes requires careful consideration of covariance matrices, which we plan to investigate in the future.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by NIH/NIAAA grant AA020648 and NIH/NIMH grant MH110110.

References

- Achana FA, Cooper NJ, Bujkiewicz S, Hubbard SJ, Kendrick D, Jones DR, Sutton AJ. Network meta-analysis of multiple outcome measures accounting for borrowing of information across outcomes. *BMC medical research methodology*. 2014; 14(1):92. [PubMed: 25047164]
- Caldwell DM, Ades A, Higgins J. Simultaneous comparison of multiple treatments: combining direct and indirect evidence. *BMJ: British Medical Journal*. 2005; 331(7521):897–900. [PubMed: 16223826]
- Casella, G., Berger, RL. *Statistical inference*. Vol. 2. Duxbury; Pacific Grove, CA: 2002.
- Chaimani A, Higgins J, Mavridis D, Spyridonos P, Salanti G. Graphical tools for network meta-analysis in stata. *PloS one*. 2013; 8(10):e76654. [PubMed: 24098547]

- Chaimani A, Salanti G. Using network meta-analysis to evaluate the existence of small-study effects in a network of interventions. *Research Synthesis Methods*. 2012; 3(2):161–176. [PubMed: 26062088]
- Chen, M-H., Shao, Q-M., Ibrahim, JG. Monte Carlo methods in Bayesian computation. Springer; New York: 2000.
- Chen Y, Hong C, Riley RD. An alternative pseudolikelihood method for multivariate random-effects meta-analysis. *Statistics in Medicine*. 2015; 34(3):361–380. [PubMed: 25363629]
- Clayton DG. A model for association in bivariate life tables and its application in epidemiological studies of familial tendency in chronic disease incidence. *Biometrika*. 1978; 65(1):141–151.
- Congdon P. Bayesian predictive model comparison via parallel sampling. *Computational statistics & data analysis*. 2005; 48(4):735–753.
- Copas J, Shi J. A sensitivity analysis for publication bias in systematic reviews. *Statistical Methods in Medical Research*. 2001; 10(4):251–265. [PubMed: 11491412]
- Copas JB. A likelihood-based sensitivity analysis for publication bias in meta-analysis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*. 2013; 62(1):47–66.
- DeSantis SM, Bandyopadhyay D, Baker NL, Randall PK, Anton RF, Prisciandaro JJ. Modeling longitudinal drinking data in clinical trials: An application to the combine study. *Drug and alcohol dependence*. 2013; 132(1):244–250. [PubMed: 23566774]
- DeSantis SM, Zhu H. A bayesian mixed treatment comparisons meta-analysis of treatments for alcohol dependence and implications for planning future trials. *Medical Decision Making*. 2014 In press.
- Dey DK, Chen M-H, Chang H. Bayesian approach for nonlinear random effects models. *Biometrics*. 1997:1239–1252.
- Dey DK, Kuo L, Sahu SK. A bayesian predictive approach to determining the number of components in a mixture distribution. *Statistics and Computing*. 1995; 5(4):297–305.
- Efthimiou O, Mavridis D, Cipriani A, Leucht S, Bago P, Salanti G. An approach for modelling multiple correlated outcomes in a network of interventions using odds ratios. *Statistics in medicine*. 2014; 33(13):2275–2287. [PubMed: 24918246]
- Efthimiou O, Mavridis D, Riley RD, Cipriani A, Salanti G. Joint synthesis of multiple correlated outcomes in networks of interventions. *Biostatistics*. 2015; 16(1):84–97. [PubMed: 24992934]
- Egger, M., Smith, GD., Altman, D. Systematic reviews in health care: meta-analysis in context. John Wiley & Sons; 2008.
- Egger M, Smith GD, Schneider M, Minder C. Bias in meta-analysis detected by a simple, graphical test. *BMJ: British Medical Journal*. 1997; 315(7109):629–634. [PubMed: 9310563]
- Franchini A, Dias S, Ades A, Jansen J, Welton N. Accounting for correlation in network meta-analysis with multi-arm trials. *Research Synthesis Methods*. 2012; 3(2):142–160. [PubMed: 26062087]
- Gelfand, A., Det, D., Chang, H. Bayesian statistics. Bernardo, JM; 1992. p. 147-159.
- Gelman A, Carlin JB, Stern HS, Rubin DB. Bayesian data analysis. texts in statistical science series. 2004
- Graybill, FA. Theory and application of the linear model. Cengage Learning; 2000.
- Hofert M, Kojadinovic I, Maechler M, Yan J, Maechler MM, Suggests M. Package ?copula?. 2015
- Jackson D, Riley R, White IR. Multivariate meta-analysis: Potential and promise. *Statistics in Medicine*. 2011; 30(20):2481–2498. [PubMed: 21268052]
- Kirkham JJ, Riley RD, Williamson PR. A multivariate meta-analysis approach for reducing the impact of outcome reporting bias in systematic reviews. *Statistics in Medicine*. 2012; 31(20):2179–2195. [PubMed: 22532016]
- Kojadinovic I, Yan J, et al. Modeling multivariate distributions with continuous margins using the copula r package. *Journal of Statistical Software*. 2010; 34(9):1–20.
- Kuss O, Hoyer A, Solms A. Meta-analysis for diagnostic accuracy studies: a new statistical model using beta-binomial distributions and bivariate copulas. *Statistics in medicine*. 2014; 33(1):17–30. [PubMed: 23873593]
- Lu G, Ades A. Assessing evidence inconsistency in mixed treatment comparisons. *Journal of the American Statistical Association*. 2006; 101(474):447–459.
- Lumley T. Network meta-analysis for indirect treatment comparisons. *Statistics in Medicine*. 2002; 21(16):2313–2324. [PubMed: 12210616]

- Riley R, Abrams K, Lambert P, Sutton A, Thompson J. An evaluation of bivariate random-effects meta-analysis for the joint synthesis of two correlated outcomes. *Statistics in Medicine*. 2007; 26(1):78–97. [PubMed: 16526010]
- Rösner S, Hackl-Herrwerth A, Leucht S, Leherer P, Vecchi S, Soyka M. Acamprosate for alcohol dependence. *Cochrane Database Syst Rev*. 2010:9.
- Rubin DB. Inference and missing data. *Biometrika*. 1976; 63(3):581–592.
- Salanti G, Higgins JP, Ades A, Ioannidis JP. Evaluation of networks of randomized trials. *Statistical Methods in Medical Research*. 2008; 17(3):279–301. [PubMed: 17925316]
- Shirley KE, Small DS, Lynch KG, Maisto SA, Oslin DW, et al. Hidden markov models for alcoholism treatment trial data. *The Annals of Applied Statistics*. 2010; 4(1):366–395.
- Sklar, M. Fonctions de répartition à n dimensions et leurs marges. Université Paris; 1959. p. 8
- Spiegelhalter, D., Thomas, A., Best, N., Lunn, D. Winbugs user manual. Cambridge: MRC Biostatistics Unit; 2003.
- Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2002; 64(4):583–639.
- Srisurapanont M, Jaruraisin N. Naltrexone for the treatment of alcoholism: a meta-analysis of randomized controlled trials. *The International Journal of Neuropsychopharmacology*. 2005; 8(02):267–280. [PubMed: 15850502]
- Wei Y, Higgins J. Bayesian multivariate meta-analysis with multiple outcomes. *Statistics in Medicine*. 2013a; 32(17):2911–2934. [PubMed: 23386217]
- Wei Y, Higgins J. Estimating within-study covariances in multivariate meta-analysis with multiple outcomes. *Statistics in Medicine*. 2013b; 32(7):1191–1205. [PubMed: 23208849]
- Yan J, et al. Enjoy the joy of copulas: with a package copula. *Journal of Statistical Software*. 2007; 21(4):1–21.

Fig. 1.

Network structures of alcohol dependence RCTs for the three binary outcomes. The thickness of the edge reflects the number of treatment comparisons for each outcome and the vertex size reflects the study sizes. RH=return to heavy drinking, RD=return to drinking, DIS=discontinuation of treatment.

Fig. 2.

Comparison-adjusted funnel plots for the 3 outcomes (Upper panel: RH; middle panel: RD; bottom panel: DIS). The label $y_{iX|Y}$ denotes the observed effect size for treatments X and Y in study i , and $\mu_{X|Y}$ denotes the direct summary estimate based on the fixed-effects model. The red line represents the null hypothesis that the study-specific effect sizes do not differ from the comparison-specific pooled effect estimates. The green line is the fitted regression line. Different colored dots represent the different treatment comparisons.

Fig. 3. Bias and coverage for estimated LORs under moderate outcome correlation (0.5, 0.5, 0.5) resulting from the UMTC (solid lines) and the MMTC (dashed lines) methods. The true values are $(LOR_{(1:2)1}, LOR_{(1:3)1}, LOR_{(2:3)1}) = (0.442, 0.847, 0.405)$, with bias (a)–(c) and the corresponding coverage probabilities (d)–(f); $(LOR_{(1:2)2}, LOR_{(1:3)2}, LOR_{(2:3)2}) = (0.811, 0, -0.811)$, with bias (g)–(i) and the corresponding coverage (j)–(l). The x-axis represents three missingness mechanisms.

Fig. 4. Ordered $-\log(\text{CPO})$ values for each study based on the proposed MMTC method. The y-axis denotes each study and the x-axis denotes the value of $-\log(\text{CPO})$.

Fig. 5.

Geometric mean posterior ORs and 95% credible intervals resulting from the MMTC (solid lines) and UMTC (dashed lines) methods. The vertical line represents OR = 1.

Table 1

41 alcohol dependence studies are summarized by outcome missingness scenarios (rows). ✓: reported; and ✕: missing.

Missing scenarios	Outcomes			# of trials
	RH	RD	DIS	
1	✓	✓	✓	11
2	✕	✓	✓	20
3	✓	✕	✓	2
4	✕	✕	✓	5
5	✕	✓	✕	3
6	✕	✕	✕	NA
# of trials	13	34	38	41

Table 2

Summary of the posterior probability of being the best treatment over 200 simulations.

(ρ_1, ρ_2, ρ_3)	Probability _{best} (outcome1, outcome2)	UMTC			MMTC		
		MCAR	MAR	MNAR	MCAR	MAR	MNAR
(0.5, 0.5, 0.5)	Pr _{best} (treatment1, treatment1)	0.000	0.001	0.000	0.001	0.002	0.001
	Pr _{best} (treatment1, treatment2)	0.004	0.004	0.007	0.004	0.006	0.008
	Pr _{best} (treatment1, treatment3)	0.000	0.001	0.000	0.000	0.000	0.000
	Pr _{best} (treatment2, treatment1)	0.009	0.005	0.002	0.005	0.006	0.001
	Pr _{best} (treatment2, treatment2)	0.133	0.120	0.161	0.119	0.125	0.134
	Pr _{best} (treatment2, treatment3)	0.008	0.010	0.003	0.002	0.002	0.000
	Pr _{best} (treatment3, treatment1)	0.049	0.042	0.019	0.052	0.062	0.032
	Pr _{best} (treatment3, treatment2)	0.736	0.721	0.769	0.752	0.743	0.799
(0.8, 0.8, 0.8)	Pr _{best} (treatment3, treatment3)	0.061	0.095	0.039	0.064	0.054	0.024
	Pr _{best} (treatment1, treatment1)	0.001	0.001	0.000	0.003	0.003	0.002
	Pr _{best} (treatment1, treatment2)	0.003	0.004	0.004	0.006	0.005	0.008
	Pr _{best} (treatment1, treatment3)	0.000	0.000	0.000	0.000	0.000	0.000
	Pr _{best} (treatment2, treatment1)	0.007	0.005	0.001	0.001	0.003	0.000
	Pr _{best} (treatment2, treatment2)	0.139	0.124	0.150	0.111	0.147	0.122
	Pr _{best} (treatment2, treatment3)	0.005	0.008	0.002	0.000	0.000	0.000
	Pr _{best} (treatment3, treatment1)	0.049	0.037	0.015	0.038	0.040	0.018
	Pr _{best} (treatment3, treatment2)	0.736	0.710	0.785	0.805	0.744	0.829
	Pr _{best} (treatment3, treatment3)	0.059	0.111	0.043	0.037	0.058	0.020

Abbreviations: MCAR, missing completely at random; MAR, missing at random; MNAR, missing not at random; UMTC, univariate MTC; MMTC, multivariate MTC.

Table 3

Summary of sensitivity analysis to prior specification for the covariance matrix with the alcohol dependence data.

Parameters	Wishart prior, d.f. = 9			Wishart prior, d.f. = 10			Wishart prior, d.f. = 12			Wishart prior, d.f. = 15		
	Post mean (SE)	95% CrI		Post mean (SE)	95% CrI		Post mean (SE)	95% CrI		Post mean (SE)	95% CrI	
$\sigma^2_{(NAL:PLB)1}$	0.231 (0.101)	(0.098, 0.483)		0.214 (0.089)	(0.096, 0.435)		0.189 (0.077)	(0.085, 0.382)		0.152 (0.060)	(0.069, 0.297)	
$\sigma^2_{(ACA:PLB)1}$	0.204 (0.093)	(0.083, 0.435)		0.198 (0.085)	(0.086, 0.413)		0.168 (0.073)	(0.072, 0.350)		0.137 (0.060)	(0.060, 0.289)	
$\sigma^2_{(NAL+ACA:PLB)1}$	0.261 (0.130)	(0.100, 0.596)		0.215 (0.102)	(0.085, 0.473)		0.201 (0.094)	(0.081, 0.438)		0.148 (0.061)	(0.065, 0.298)	
$\sigma^2_{(NAL:PLB)2}$	0.172 (0.063)	(0.082, 0.325)		0.162 (0.059)	(0.079, 0.304)		0.148 (0.052)	(0.074, 0.272)		0.125 (0.044)	(0.063, 0.230)	
$\sigma^2_{(ACA:PLB)2}$	0.227 (0.098)	(0.097, 0.467)		0.212 (0.084)	(0.096, 0.422)		0.186 (0.071)	(0.084, 0.359)		0.149 (0.057)	(0.069, 0.286)	
$\sigma^2_{(NAL+ACA:PLB)2}$	0.218 (0.105)	(0.083, 0.485)		0.191 (0.099)	(0.076, 0.442)		0.160 (0.072)	(0.067, 0.342)		0.131 (0.054)	(0.059, 0.265)	
$\sigma^2_{(NAL:PLB)3}$	0.296 (0.139)	(0.110, 0.642)		0.268 (0.147)	(0.098, 0.633)		0.214 (0.101)	(0.085, 0.472)		0.163 (0.070)	(0.072, 0.342)	
$\sigma^2_{(ACA:PLB)3}$	0.297 (0.162)	(0.101, 0.721)		0.252 (0.131)	(0.098, 0.597)		0.216 (0.105)	(0.086, 0.482)		0.169 (0.077)	(0.071, 0.367)	
$\sigma^2_{(NAL+ACA:PLB)3}$	0.241 (0.133)	(0.089, 0.596)		0.205 (0.098)	(0.082, 0.454)		0.187 (0.091)	(0.073, 0.424)		0.144 (0.062)	(0.064, 0.301)	
DIC	891.321			916.202			937.946			963.018		