

Bayesian Modeling for benefit-risk balance analysis

Konstantinos Vamvourellis

with K. Kalogeropoulos and L. Phillips

Department of Statistics at LSE

BAYES:2018 Cambridge, June 21 2018



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■

Description of the Problem

Current State of Research

Proposed Bayesian Model

Results

Discussion and Future Work

Regulatory Timeline for Rosiglitazone (Avandia)

- ▶ Rosiglitazone gets approval in US (1999) and Europe (2000)
- ▶ New evidence of risks arises [see Nissen and Wolski, 2007]
- ▶ 2010 European regulators revert their recommendation
- ▶ 2011-13 US regulators impose special restrictions
- ▶ 2013 US regulators reanalyzed clinical trials data and voted to lift restrictions

No consensus on the magnitude of the risks and whether the risks outweigh the benefits.

Objective

- ▶ Principled Benefit-Risk Assessment of a drug
- ▶ Assess and Compare different treatments
- ▶ Incorporate:
 - ▶ Clinical Judgment
 - ▶ Uncertainty

Benefit-Risk Methodology Project

In 2008 European Medicines Agency (EMA) started the Benefit-Risk Methodology Project¹ with experts in decision theory from the LSE and with the University of Groningen.

identify decision-making models that could be used in the Agency's work, to make the assessment of the benefits and risks of medicines more consistent, more transparent and easier to audit.

¹http://www.ema.europa.eu/ema/index.jsp?curl=pages/special_topics/document_listing/document_listing_000314.jsp

Multi-Criteria Decision Analysis (MCDA)

- ▶ Identify the population mean μ_j of all variables of interest
- ▶ Transform effects $f(\mu_j)$ to a common scale for comparison

$$f(x) = \begin{cases} \frac{100x_{min}}{x_{min}-x_{max}} + \frac{100}{x_{max}-x_{min}}x & \text{for favourable effects} \\ \frac{100x_{max}}{x_{max}-x_{min}} + \frac{100}{x_{min}-x_{max}}x & \text{otherwise} \end{cases}$$

- ▶ Assign clinical weights w_j to each effect so that $\sum_j w_j = 1$
- ▶ Calculate the weighted average score

$$S = \sum_j w_j \cdot f_j(\mu_j)$$

Summary of Current State of Research

aggregate level data

- ▶ State of the art focus on modeling summary data
- ▶ hence does not account for correlation among variables
- ▶ Wen et al. [2014] present 2 methods to incorporate uncertainty in MCDA Benefit-Risk Score for a known covariance matrix Γ

patient level data

- ▶ When patient level data is available we need an appropriate model to incorporate correlation
- ▶ We propose a Bayesian Latent Variable Model to and introduce correlation among the latent variables
- ▶ The model is flexible enough to handle mixed type data (continuous, binary and count)

2 Approaches to Incorporate Clinical Data Uncertainty in MCDA

- ▶ δ -method to construct confidence interval of MCDA score

$$\hat{s} = \sum_j w_j \cdot f_j(\hat{\mu}_j)$$
$$s \sim N(\hat{s}, \nabla s' \Gamma \nabla s)$$

- ▶ Monte-Carlo method for confidence interval of MCDA score

$$\mu^{(i)} \sim N(\hat{\mu}, \Gamma)$$
$$s^{(i)} = \sum_j w_j \cdot f_j(\hat{\mu}_j)$$

An estimate of Γ is needed to apply this method. Note that Γ cannot be identified from aggregate level data.

Bayesian Modeling

- ▶ Phillips et al. [2015] proposed using MCDA for drug assessment
 - ▶ Bayesian model for aggregate level data
 - ▶ assumes independence of variables
 - ▶ constructed posterior predictive distribution of the MCDA score under uniform improper priors
- ▶ Wen et al. [2014] in future research section highlight the need to for a more sophisticated Bayesian model to incorporate correlations.
- ▶ We propose method to find the covariance matrix Γ with patient level data
 - ▶ we adopt the 'matrix completion' method to find the correlation matrix R among the variables
 - ▶ we extend the Talhouk et al. [2012] algorithm to account for data of mixed type (continuous, binary, counts etc.)
 - ▶ we provide a Gibbs sampler (implemented in Python) and an HMC algorithm (implemented in Stan)

Model

Data is recorded in a $N \times J$ matrix Y_{ij}

J effects possibly correlated and N independent subjects

For binary (or count) data:

$$\begin{cases} Y_{ij} \sim \text{Bernoulli}(\eta_{ij}) \text{ (} \sim \text{Poisson}(\eta_{ij}) \text{)} \\ h_j(\eta_{ij}) = \mu_j + Z_{ij}, \text{ for appropriate link function } h \end{cases}$$

For continuous variables:

$$Y_{ij} = \mu_j + Z_{ij}, \quad i = 1, \dots, N.$$

The distribution of Z is assumed² to be

$$Z_{i:} \sim N_J(0_J, \Sigma),$$

where Σ is a $J \times J$ covariance matrix, 0_J is a row J -dimensional vector with zeros and $Z_{i:}$ are independent $\forall i$.

²other options are available, e.g. a multivariate t

Model

- ▶ Parametrisation according to covariance is non likelihood identifiable
- ▶ Gibbs sampler is adapted from Talhouk et al. [2012] targets conditionals $p(\Sigma|Z, \mu)$ and $p(\mu|Z, \Sigma)$. Uses Metropolis within Gibbs step for $p(Z|\Sigma, \mu)$
- ▶ HMC sampler is able to sample from $p(Z, \Sigma, \mu|Y)$ simultaneously using information from the gradient of the parameter space
- ▶ We use appropriately wide priors as suggested in relevant literature

With posterior samples from $p(\mu^{(g)}|Y^{(g)})$ for $g = \{C, T\}$ we are able to simulated any metric of interest, such as the distribution of final scores $p(s^{(g)})$ or the probability of the treatment being better $P(s^T > s^C|Y)$.

Simulations

Simulated datasets for the efficacy and adverse effects of a hypothetical drug. We created two datasets, Treatment (T) and Control (C) and calculated Benefit-Risk scores s^T and s^C respectively. We compare the two models

- ▶ *Model 1* Independent Model
- ▶ *Model 2* Latent Variable model that learns the correlation matrix R

Compared cases between datasets generated with $R = I$ and $R \neq 0$ where sampled correlation values

$$R' = \begin{bmatrix} 1 & u & u & 0 & 0 \\ u & 1 & u & 0 & 0 \\ u & u & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & v \\ 0 & 0 & 0 & v & 1 \end{bmatrix}$$

- ▶ $u \sim U(0.5, 0.9)$ among the continuous effects
- ▶ $v \sim U(0.2, 0.6)$ among the binary effects

Results

Correlation matters

- ▶ The posterior distribution $p_{M_1}(\mu|Y)$ has lower variance than $p_{M_2}(\mu|Y)$
- ▶ As a result $P_{M_1}(s^T > s^C|y)$ overestimates the true probability $P(s^T > s^C|y)$

The proposed free model is relatively robust against overfitting and is able to retrieve the correct values even when the data has no correlation.

Results

We generate two synthetic datasets: correlation $R = 1$ (Dataset A), and correlation $R = R'$ (Dataset B).


We estimate the probability that treatment is better than the control $P(s^T > s^C|y)$ with both models 1 and 2.

Fully Bayesian	Model 1	Model 2
Dataset		
A	94%	93%
B	93%	91%

App. Normal	Model 1	Model 2
Dataset		
A	91%	91%
B	92%	88%

Application to real data

- ▶ We applied our model to a patient level dataset for 3 treatments for type 2 Diabetes
- ▶ 4 adverse binary variables (Diarrhoea, Nausea/Vomiting, Dyspepsia, Oedema) and 2 efficacy continuous variables (Haemoglobin and Glucose levels)
- ▶ We did discovered strong correlations only between efficacy variables
- ▶ We confirmed that the results are very similar between Model 1 and Model 2



Application to real data

Fully Bayesian	Model 1	Model 2
Treatment		
RSG - AVM	93%	93%
RSG - MET	99%	99%

App. Normal	Model 1	Model 2
Treatment		
RSG - AVM	92%	94%
RSG - MET	99%	99%

Discussion

- ▶ Currently working on assessing the effect of priors on the posterior mean and variance
- ▶ Current inference methods (Gibbs and HMC) provide reasonable agreement between the true parameter values and their posterior distributions.
- ▶ HMC is more powerful than Gibbs but potentially more computationally expensive
- ▶ There is room to improve MCMC. Possible solution includes Pseudo-Marginal Likelihood method to integrate out latent variables.

Discussion

- ▶ There is still the question of how to choose a parsimonious model
- ▶ Neither of the two inference methods provides estimates of marginal likelihood for Bayesian model choice
- ▶ Possible solution includes Pseudo-Marginal Likelihood method to integrate out latent variables.
- ▶ Future work includes Sequential Monte Carlo methods that address all the above limitations

Thank you!



THE LONDON SCHOOL
OF ECONOMICS AND
POLITICAL SCIENCE ■

Konstantinos Vamvourellis
Department of Statistics
k.vamvourellis@lse.ac.uk
github.com/bayesways
personal.lse.ac.uk/vamourel

References I

- Steven E. Nissen and Kathy Wolski. Effect of Rosiglitazone on the Risk of Myocardial Infarction and Death from Cardiovascular Causes. *New England Journal of Medicine*, 356(24):2457–2471, jun 2007. ISSN 0028-4793. doi: 10.1056/NEJMoa072761. URL <http://www.nejm.org/doi/abs/10.1056/NEJMoa072761>.
- Lawrence Phillips, Billy Amzal, Alex Asiimwe, Edmond Chan, Chen Chen, Diana Hughes, Juhaeri Juhaeri, Alain Micallef, Shahrul Mt-Isa, and Becky Noel. Wave 2 Case Study Report: Rosiglitazone. 2015.
- Aline Talhouk, Arnaud Doucet, and Kevin Murphy. Efficient Bayesian Inference for Multivariate Probit Models With Sparse Inverse Correlation Matrices. *Journal of Computational and Graphical Statistics*, 21(3):739–757, jul 2012. ISSN 1061-8600. doi: 10.1080/10618600.2012.679239. URL <http://www.tandfonline.com/doi/abs/10.1080/10618600.2012.679239>.

References II

Shihua Wen, Lanju Zhang, and Bo Yang. Two Approaches to Incorporate Clinical Data Uncertainty into Multiple Criteria Decision Analysis for Benefit-Risk Assessment of Medicinal Products. *Value in Health*, 17(5):619–628, jul 2014.