Scuola di Dottorato
Università degli Studi di Milano-Bicocca

Department of

Economics, Management, and Statistics

Ph.D. program: **Economics and Statistics** Cycle: **XXXVI°**
Curriculum: **Statistics**

Advancements in Bayesian Methods for Personalized Medicine

Surname: **Bi**
Name: **Jiefeng**
Registration number: **868800**

Supervisor: Prof. **Bernardo Nipoti**

Academic Year: 2023-2024

Abstract

Personalized medicine is a rapidly evolving and transformative approach in healthcare. As the complexity and volume of data continue to grow, more advanced methodologies are needed to effectively implement personalized strategies. In this manuscript, we develop novel Bayesian methodologies specifically designed to handle high-dimensional data and identify complex patterns more accurately.

The thesis consists of three main contributions. Firstly, we propose a Bayesian model for optimizing dynamic treatment regimes, a practical framework within personalized medicine, to address the uncertainty in identifying optimal decision sequences and incorporate dimensionality reduction to manage high-dimensional individual covariates. To achieve this, the model is augmented to accommodate counterfactual variables. Additionally, we introduce a novel class of spike-and-slab priors for the multi-stage selection of significant factors, to favor the sharing of information across stages. The second primary contribution of this manuscript is an extension of the Bayesian model with spike-and-slab priors to accommodate frameworks with more than two stages. We use this model specification to identify significant molecular biomarkers and clinical parameters influencing drug effects on estimated Glomerular Filtration Rate (eGFR) using a type II diabetes dataset. Specifically, the dataset we analyze is derived from the PROVALID study, a prospective observational study encompassing a diverse range of chronic kidney disease stages. Lastly, we introduce mixed Bayesian networks to tackle causal structure learning across diverse data types, facilitating a more accurate representation of causal dependencies by a directed acyclic graph. We present that under this model, the causal structure between mixed functional data can be uniquely identifiable under certain mild conditions. Furthermore, we develop a Bayesian framework to infer the adjacency matrix of the directed acyclic graph with inherent uncertainty quantification.

Reading maketh a full man; conference a ready man; and writing an exact man.

— Francis Bacon, 1597

Contents

Abstract **iii**

Contents **vii**

List of Figures **xi**

List of Tables **xv**

1 Introduction **1**

1.1 Personalized medicine 1

1.2 Bayesian methods 2

1.2.1 Bayesian variable selection 3

1.2.2 Bayesian networks 5

1.3 Thesis outline 5

2 Uncertainty quantification and multi-stage variable selection for personalized treatment regimes **11**

2.1 Introduction 11

2.1.1 Our contribution 12

2.2 Bayesian augmented learning with variable selection 13

2.2.1 Regression models 16

2.2.2 Multi-stage variable selection 17

2.3 Posterior computation 19

2.4 Simulation studies . . . 20
2.4.1 First experiment . . . 21
2.4.2 Second experiment . . . 24
2.5 Clinical case study . . . 26
2.6 Conclusion . . . 29
Bibliography . . . 34
Appendices . . . 37
2.A On the validity of the equivalences in (2.3) . . . 37
2.B Derivation of the expression in (2.8) . . . 39
2.C Posterior inference . . . 40
2.D Additional results on the simulation study . . . 43
2.E Additional results on the clinical case study . . . 44

3 Dependent spike-and-slab for drug combinations in diabetic kidney disease 49
3.1 Introduction . . . 49
3.2 Data . . . 51
3.3 Methods . . . 56
3.3.1 Bayesian model and variable selection . . . 56
3.3.2 Optimal treatment prediction . . . 60
3.4 Results . . . 62
3.5 Discussion . . . 68
Bibliography . . . 71
Appendices . . . 74
3.A Posterior inference . . . 74
3.B Additional results on the case study . . . 77

4 Causal structure learning with mixed functional Bayesian networks 83
4.1 Introduction . . . 83

4.2 Method and causal identification 85
4.2.1 Notation and background 85
4.2.2 Mixed functional linear non-Gaussian Bayesian networks 87
4.2.3 Causal identifiability 89
4.3 Bayesian inference 92
4.4 Simulation studies 95
4.4.1 Data generation 95
4.4.2 Model specification and inference 97
4.4.3 Results 98
4.5 Discussion 100
Bibliography 102
Appendices 107
4.A Proof of Theorem 4.2.5 107
4.B Posterior inference 113
4.C Simulations: sensitivity analysis 117

5 Conclusion **119**

Acknowledgment **123**

List of Figures

2.1 Posterior inclusion probabilities at stage 2 (left panel) and stage 1 (right panel). Treatment options are shown on the x-axis, while clinical factors are listed on the y-axis. Specifically, 'A' represents ACEi, 'B' refers to beta-blockers, 'C' indicates CCB, and 'D' stands for diuretics. 'AB' denotes the combination of ACEi and beta-blockers, with other combinations defined similarly. The notation '1' refers to first-stage treatments, and '2' pertains to second-stage agents. For instance, if the x-axis is labeled A.1 and the y-axis is SBP, the corresponding value represents the posterior inclusion probability of the interaction between first-stage ACEi and SBP. The symbol 'S', corresponding to posterior inclusion probabilities larger than 0.5, indicates statistical significance. 28

2.2 MIMIC-III data. Posterior probability of each of the $T = 8$ treatment options being the optimal choice for Individual #1 at stage 1 (left panel), for Individual #2 at stage 2 (middle panel), and for Individual #2 at stage 1 (right panel). In each barplot, the white bin indicates the largest posterior probability. 29

2.3 MIMIC-III data. Posterior marginal probability of each of four antihypertensive agents being part of the optimal combination of treatments, for Individual #1 at stage 1 (left panel), for Individual #2 at stage 2 (middle panel), and for Individual #2 at stage 1 (right panel). In each barplot, the white bin indicates the largest posterior marginal probability. 45

3.1 Depiction of eGFR over time. The top-left panel illustrates the trajectories of eGFR for all observations, with red lines indicating critical cutoff values for assessing kidney function. The top-right panel presents a boxplot of eGFR across 4 stages, where the red values represent the mean eGFR at each stage. The bottom panel displays boxplots of eGFR, stratified by the four treatments in each stage. 54

3.2 The individual trajectories of eGFR for 16% of the dataset. 55

3.3 The False Discovery Rate (FDR) and the corresponding number of significant variables (D) across 4 stages. The x-axis represents the cutoff values, while the y-axis on the right side indicates the number of significant variables, and the y-axis on the left side indicates the false discovery rate. 63

3.4 Interaction plots and conditional interaction plots by stages between eGFR and HACR1_MESO, as well as EGF_MESO, under different treatments. . 65

3.5 Control FDR at 0.01, significant variables selected across 4 stages. The suffix '1' represents the interaction between RASi and clinical factors, while '1.1', '1.2', and '1.3' represent interactions between RASi combined with SGLT2i, MCRa, and GLP1a, respectively, with clinical factors. The bars' values represent the variable's posterior inclusion probability, while the color gradient, ranging from red to blue, indicates a progressive increase in the posterior inclusion probability. 66

3.6 The predicted optimal treatments of patients of the test set. The Y-axis denotes the patients' ID, and the text in each box means the predicted optimal treatment via DSS, while colors indicate the corresponding percentage increase of individuals in the test set. The percentage increase is calculated as the estimated optimal eGFR minus the observed eGFR, divided by the observed eGFR, and multiplied by 100. 67

3.7 Estimated improvement in patients from the training set after DSS treatments. The difference ratio is calculated as the estimated optimal eGFR minus the observed eGFR, divided by the observed eGFR, and multiplied by 100. 79

3.8 Predicted improvement in patients from the test set after DSS treatments. The difference ratio is calculated as the predicted optimal eGFR minus the observed eGFR, divided by the observed eGFR, and multiplied by 100. . . . 79

3.9 Interaction plots and conditional interaction plots by stages between eGFR and NPHS1_MESO, FGF21_MESO, as well as LCN2L_MESO, under different treatments. 80

4.1 Illustration of functions and their corresponding basis coefficients via basis expansion. The functional data consists of 50 time points, and we use 5 basis functions to perform an expansion. 85

4.2 Illustrate the Non-Injectivity of Example 1. The left panel shows the density distribution of ϵ_2, while the right panel presents its cumulative distribution function (CDF). The red segment in the right panel indicates that the CDF of ϵ_2 remains the same within the range denoted by the red line. 92

4.3 Illustrative example of causal identification. The data are simulated according to the causal graph $1 \to 2$. The colored lines in each panel represent the fitted regression lines for the variables across groups 1 through 4. The top-left panel shows a probit regression of W_2 on W_1 in the observed space, the top-right shows a linear regression of Z_2 on Z_1 in the latent space, the bottom-left shows a linear regression of W_1 on W_2 in the observed space, and the bottom-right shows a linear regression of Z_1 on Z_2 in the latent space. 93

4.4 A DAG illustrating the model hierarchy. The blue nodes represent observed variables, while the red nodes correspond to latent space coefficients that establish causal structures. 95

4.5 True latent functions of a binary function and their estimates via Mixed-FLiNG-BN. One binary function and nine individuals are randomly selected from a total of 5 binary functions and 500 samples. The trajectories of the binary function (represented by green dots) are plotted, along with the corresponding true latent function (represented by the blue solid line) and the estimated latent function (represented by the red dashed line) obtained using our proposed model. The number of time points is 50. 100

List of Tables

2.1 First experiment: simulated data with n smaller than p_2. For stage 2 and stage 1, three summaries of the accuracy in selecting significant variables (proportion of FN, proportion of FP, and F_1 score), for DSS, ISS, and QLL, across the nine scenarios obtained by specifying $(k, n) \in \{(10, 25), (20, 50), (30, 75)\}$ and $\rho^* \in \{0.3, 0.6, 0.9\}$. Results are based on 100 replicated datasets. 23

2.2 First experiment: simulated data with n smaller than p_2. For stage 2, stage 1 and overall, two measures of prediction accuracy (MRE and ER), for DSS, ISS and QLL, across the nine scenarios obtained by specifying $(k, n) \in \{(10, 25), (20, 50), (30, 75)\}$ and $\rho^* \in \{0.3, 0.6, 0.9\}$. Results are based on 100 replicated datasets. 31

2.3 Second experiment. For stage 2 and stage 1, three summaries of the accuracy in selecting significant variables (proportion of FN, proportion of FP, and F_1 score), for DSS, ISS, and QLL, across the four scenarios obtained by specifying $n \in \{200, 400)\}$ and $T \in \{4, 8\}$. The number of individual covariates is $k = 10$, while $\rho^* = 0.6$. Results are based on 100 replicated datasets. . . . 32

2.4 Second experiment. For stage 2, stage 1 and overall, two measures of prediction accuracy (MRE and ER), for DSS, ISS and QLL, across the four scenarios obtained by specifying $n \in \{200, 400)\}$ and $T \in \{4, 8\}$. The number of individual covariates is $k = 10$, while $\rho^* = 0.6$. Results are based on 100 replicated datasets. 33

2.5 First experiment: simulated data with n larger than p_2. For stage 2 and stage 1, three summaries of the accuracy in selecting significant variables (proportion of FN, proportion of FP, and F_1 score), for DSS, ISS, and QLL, across the nine scenarios obtained by specifying $(k, n) \in \{(10, 50), (20, 100), (30, 150)\}$ and $\rho^* \in \{0.3, 0.6, 0.9\}$. Results are based on 100 replicated datasets. 44

2.6 First experiment: simulated data with n larger than p_2. For stage 2, stage 1 and overall, two measures of prediction accuracy (MRE and ER), for DSS, ISS and QLL, across the nine scenarios obtained by specifying $(k, n) \in \{(10, 50), (20, 100), (30, 150)\}$ and $\rho^* \in \{0.3, 0.6, 0.9\}$. Results are based on 100 replicated datasets. 46

2.7 Frequency analysis of actual medication to patients by doctors in two stages. A refers to angiotensin-converting enzyme inhibitors (ACEi), B refers to beta-blockers, C refers to calcium channel blockers (CCB), and D refers to diuretics. 47

2.8 MIMIC-III data. Values taken by BIC and LPML for DSS, ISS and QLL. LPML is not applicable to QLL, hence the missing entries. 47

3.1 Medical interventions administered to the 223 patients across four stages. . 55

3.2 Control FDR at 0.01, the corresponding number of significant variables and cutoff values across 4 stages. 64

3.3 The percentage of patients experiencing improved healthcare in the test set after receiving the predicted optimal treatment through DSS. 68

3.4 Covariates after variable selection using DSS in the PROVALID study. . . . 78

4.1 Observations are generated using mixed functions on an evenly spaced grid. The reported average performance is based on 30 repetitions, and standard deviations are given within the parentheses. The best-performing method for each metric is highlighted in bold. 99

4.2 Sensitivity analysis with respect to the choices of hyperparameters with $(p, h, N) = (20, 50, 200)$ where five of the nodes are binary and the rest are continuous. The reported average performance is based on 50 repetitions, and standard deviations are given within the parentheses. 117

Chapter 1

Introduction

1.1 Personalized medicine

As a medical paradigm, personalized medicine, uses individual patient data systematically to optimize healthcare outcomes (Adam et al., 2007; Diamandis et al., 2010). Unlike traditional medical approaches, which often rely on standardized treatments for broad populations, commonly referred to as the one-size-fits-all strategy, personalized medicine accounts for individual variability in genetics, environment, lifestyle, and disease mechanisms. Guided by the principle of 'the right drug, with the right dose at the right time to the right patient' (Sadee and Dai, 2005), the benefits of personalized medicine include improved treatment adherence, the ability to enhance patient care through the selection of optimal therapies, reduced side effects, and a decrease in overall medical costs (Chan and Ginsburg, 2011; Jameson and Longo, 2015; Ashley, 2016; Duffy, 2016; Johnson et al., 2022). The growing interest in personalized medicine is primarily driven by the effort to make it more evidence-based or data-driven, which introduces new methodological challenges often beyond the capabilities of traditional quantitative tools.

This manuscript focuses on personalized treatments for chronic diseases, such as hypertension, diabetes, HIV infection, mental illnesses, among others. Chronic diseases, a major public health challenge (WHO, 1997; Anderson et al., 2002), often require ongoing medical intervention, complicating the personalized medicine framework due to the need for multi-stage treatments. Early statistical methods in personalized medicine focus primarily on basic frequentist models (Murphy, 2005; Murphy and Bingham, 2009; Chakraborty et al., 2010; Chakraborty, 2011), which were suitable for smaller and simpler datasets (Zhou et al., 2014). As the complexity and volume of data increase, machine learning algorithms (Laber and Murphy, 2011; Goldberg and Kosorok, 2012; Zhao et al., 2012; Cui et al., 2017; Kline

et al., 2022; Abbaoui et al., 2024) and Bayesian methods (Guo and Yuan, 2017; Guo et al., 2017; Spanbauer and Sparapani, 2021; Saha et al., 2022; Zhao et al., 2024) have been developed and extensively employed. However, variable selection in high-dimensional datasets and the identification of complex patterns continue to present pressing difficulties in this field.

Although there is a vast literature on variable selection in high-dimensional data, these methods are largely focused on selecting variables that are vital for prediction. Consequently, in personalized medicine, such methods may overlook variables that, while not predictive, are essential for treatment decision-making (Chakraborty and Moodie, 2013; Tsiatis et al., 2019). The identification of complex patterns recovered from data, enabling the identification and estimation of causal effects is called causal discovery (Spirtes and Zhang, 2016). It enables the assessment of the effects of interventions, policies, or treatments, providing a critical basis for making informed, evidence-driven decisions. Yet, establishing causal structure is complex and requires careful consideration of the data, methodology, and underlying assumptions, especially in personalized medicine, where the heterogeneity of patients, high-dimensional data, and potential confounding factors complicate the identification of robust causal-and-effect relationships (Sanchez et al., 2022; Wu et al., 2024).

To tackle the aforementioned difficulties, we propose novel Bayesian methods for variable selection and causal discovery within the context of personalized medicine. Specifically, we introduce a dependent spike-and-slab model for variable selection and develop a mixed functional Bayesian network approach for causal discovery.

1.2 Bayesian methods

The purpose of Bayesian inference is to update the probability of a hypothesis or model as new data or evidence becomes available. Bayesian methods involve specifying a probability model for the observed data y, given a vector of unknown parameters $\boldsymbol{\theta}$, leading to sampling model $p(y|\boldsymbol{\theta})$ (Hoff, 2009). We assume the parameter $\boldsymbol{\theta}$ follows a prior distribution $p(\boldsymbol{\theta})$, representing our belief about the true population characteristics. To update our belief about $\boldsymbol{\theta}$ using the posterior distribution $p(\boldsymbol{\theta}|y)$, we apply Bayes' theorem, combining the prior distribution with the sampling model:

$$p(\boldsymbol{\theta}|y) = \frac{p(y|\boldsymbol{\theta})p(\boldsymbol{\theta})}{\int_{\boldsymbol{\Theta}} p(y|\tilde{\boldsymbol{\theta}})p(\tilde{\boldsymbol{\theta}})d\tilde{\boldsymbol{\theta}}} \tag{1.1}$$

where $\boldsymbol{\Theta}$ is the set of possible parameter values.

1.2.1 Bayesian variable selection

The Bayesian approach has proven useful for estimating regression models, particularly when the key task is selecting explanatory variables from a large set of potential candidates to include in the model. Mixture priors, particularly those with spike-and-slab components, have gained substantial traction in the field of variable selection (Mitchell and Beauchamp, 1988; George and McCulloch, 1993, 1997; Ishwaran and Rao, 2005). See also Tadesse and Vannucci (2021) for a recent overview on the topic. The spike component, concentrated around zero, serves to shrink small effects towards zero, while the slab component covers a broader range of potential values for the regression coefficients, offering greater flexibility in the model. To specify spike-and-slab priors, let $\boldsymbol{\theta} = (\theta_1, \cdots, \theta_d)$ represent regression coefficients, and introduce indicator variables $\boldsymbol{\delta} = (\delta_1, \cdots, \delta_d)$ where δ_j is 1, if θ_j is associated with the slab component. The spike-and-slab priors can be formulated as follows:

$$p(\boldsymbol{\theta}|\boldsymbol{\delta}) = p_{\text{slab}}\ (\boldsymbol{\theta}_{\boldsymbol{\delta}}) \prod_{j:\delta_j=0} p_{\text{spike}}\ (\theta_j) \tag{1.2}$$

where p_{spike} and p_{slab} denote the univariate spike and the multivariate slab distribution respectively, and $\boldsymbol{\theta}_{\boldsymbol{\delta}}$ denotes the vector containing elements of $\boldsymbol{\theta}$ where $\delta_j = 1$. Two primary types of spikes have been proposed: (1) spikes defined by an absolutely continuous distribution (George and McCulloch, 1993; Ishwaran and Rao, 2003, 2005), and (2) spikes defined as a point mass at zero, known as Dirac spikes (Zellner, 1986; O'Hagan, 1995). In this thesis, we adopt the first approach, using normal priors for both the spike-and-slab components, with inverse Gamma distributions for their respective variances (Ishwaran and Rao, 2003, 2005). The prior inclusion probability $P(\delta_j = 1)$ of the regression coefficient θ_j is specified hierarchically as:

$$P(\delta_j = 1|\omega_j) = \omega_j, \quad \omega_j \sim \text{Beta}\left(a_{\omega_j}, b_{\omega_j}\right) \tag{1.3}$$

The dependence among responses associated with the same covariates is captured through the prior specification of parameters ω_j. In recent years, this idea has been widely applied in genetic association studies with multiple outcomes. For example, Jia and Xu (2007) proposed a simple specification:

$$\omega_{sj} \equiv \omega_s \sim \text{Dirichlet}(1, 1), \tag{1.4}$$

where ω_{sj} represents the inclusion probability of covariate X_s for the response y_j. In contrast, Richardson et al. (2010) choose a quite complex specification:

$$\omega_{sj} = \rho_s \omega_j, \quad \omega_j \sim \text{Beta}\left(a_j, b_j t\right), \quad \rho_s \sim \text{Gamma}\left(c_s, d_s\right), \quad 0 \leq \omega_{sj} \leq 1 \tag{1.5}$$

where the inclusion probability of X_s for the response y_j is modeled by ω_j, which is specific to that particular response, while being moderated by the parameter ρ_s, which is shared across all responses. Scott-Boyer et al. (2012) proposed an alternative variant for this prior (1.5):

$$\omega_{sj} \mid a_s, b_s, \pi_s \sim \pi_s \delta_0 + (1 - \pi_s)\, \text{Beta}\left(a_s, b_s\right), \quad \pi_s \sim \text{Beta}\left(a_0, b_0\right), \tag{1.6}$$

with $a_s \sim \text{Exp}\left(\lambda_a\right)$ and $b_s \sim \text{Exp}\left(\lambda_b\right)$. Ruffieux et al. (2017) developed a simpler model, partly motivated by practical considerations,

$$\omega_{sj} \equiv \omega_s \sim \text{Beta}\left(a_s, b_s\right), \tag{1.7}$$

where the probability that a covariate affects any response is represented by a parameter shared across all responses. The simpler specification in (1.7) offers the advantage of enabling closed-form expressions for a variational algorithm, which is not possible with the specification in (1.5). As outlined in the introduction above, prior information on individual inclusion probabilities could be incorporated by appropriately selecting the parameters of the beta distribution. Zhang et al. (2023) proposed dependent priors for ω_s to borrow information among the gene pairs, by setting $a_s = c, b_s = 1 - c$, and c is from uniform distribution $(0, 1)$. In turn, $\omega_s^{'}$s are given a beta prior distribution with global hyperparameters c.

However, the types of spike-and-slab priors mentioned above are primarily designed for panel datasets rather than longitudinal datasets. Longitudinal data requires spike-and-slab priors to have capable of handling the complex dependencies of repeated measures. To overcome this, we propose a novel class of spike-and-slab priors for multi-stage selection of significant factors, promoting the sharing of information across stages in personalized medicine. Our approach accounts for the temporal structure of the data, promoting the sharing of information across different stages, and ensuring that the model reflects the continuity and dependencies between observations overtime. By integrating information across stages, we assign consistent prior inclusion probabilities to the same variable throughout multiple stages, leading to the proposal of a dependent spike-and-slab model. Variable selection is then based on the posterior inclusion probability $P(\delta_j = 1|y)$, which is sampled using Gibbs sampling.

1.2.2 Bayesian networks

Let $\boldsymbol{Y} = (Y_1, \ldots, Y_d)^T$ denote a d-dimensional random vector. Bayesian networks (BNs) are probabilistic graphical models that represent the conditional dependencies among a set of random variables $\boldsymbol{Y}$ using a directed acyclic graph (DAG). A DAG $\mathcal{G} = (\boldsymbol{V}, \boldsymbol{E})$ consists a set of nodes $\boldsymbol{V}$ and a set of directed edges represent by a binary adjacency matrix $\boldsymbol{E} = (E_{j\ell})$ where $E_{j\ell} = 1$ if and only if $\ell \to j$ for $\ell \neq j \in \boldsymbol{V}$. For each variable Y_j, its conditional distribution depends only on its parent variable denoted by $\text{Pa}(Y_j)$ in the DAG. A BN $B = (\mathcal{G}, P)$ on $\boldsymbol{Y}$ is a probability model where the joint probability distribution P of $\boldsymbol{Y}$ factorizes with respect to $\mathcal{G}$ in the following manner,

$$P(\boldsymbol{Y}) = \prod_{j=1}^{d} P_j \left(Y_j \mid \boldsymbol{Y}_{pa(j)} \right) \tag{1.8}$$

where P_j is the conditional distribution of Y_j given $\boldsymbol{Y}_{pa(j)}$ under P. Let $\text{de}(j) = \{\ell \in \boldsymbol{V} : j \to \cdots \to \ell\}$ denote the descendants of j in $\mathcal{G}$ and let $nd(j) = \boldsymbol{V} \backslash \{de(j) \cup \{j\}\}$ denote the non-descendants of j. The local directed Markov property states that any variable is conditionally independent of its non-descendants given its parents, $\boldsymbol{Y}_j \perp \boldsymbol{Y}_{nd(j)\backslash pa(j)} \mid \boldsymbol{Y}_{pa(j)}$ for any $j \in V$. BN factorization (1.8) and local Markov property are equivalent (Lauritzen, 1996). This equivalence underpins the use of Bayesian networks for causal discovery, as it allows the identification of conditional independencies and causal relationships within the network. In our proposal, Y_j is a random function that can be either continuous or binary. We propose mixed functional Bayesian networks to address the challenge of learning causal structures fro mixed random functions, particularly in the context of personalized medicine.

1.3 Thesis outline

The rest of the dissertation is organized as follows. In Chapter 2, we propose a Bayesian model for optimizing dynamic treatment regimes, which addresses the uncertainty in identifying optimal decision sequences and incorporates dimensionality reduction to manage high-dimensional individual covariates. Chapter 3 expands on the model introduced in Chapter 2, applying it to identify key molecular biomarkers and clinical parameters at each stage of drug effects on eGFR, using a type II diabetes dataset. In Chapter 4, we introduce mixed functional Bayesian networks to address the challenges posed by diverse data types, thereby enhancing the accurate representation of causal dependencies through a directed acyclic graph.

Bibliography

Abbaoui, W., Retal, S., El Bhiri, B., Kharmoum, N., and Ziti, S. (2024). Towards revolutionizing precision healthcare: A systematic literature review of artificial intelligence methods in precision medicine. *Informatics in Medicine Unlocked*, page 101475.

Adam, W., Stephen, J., and Naylor, S. (2007). Personalised medicine. *Drug Discovery*, 19.

Anderson, G., Horvath, J., Knickman, J., Colby, D., Schear, S., and Jung, M. (2002). Chronic conditions: making the case for ongoing care.

Ashley, E. A. (2016). Towards precision medicine. *Nature Reviews Genetics*, 17(9):507–522.

Chakraborty, B. (2011). Dynamic treatment regimes for managing chronic health conditions: a statistical perspective. *American Journal of Public Healthh*, 101(1):40–45.

Chakraborty, B. and Moodie, E. E. (2013). Statistical methods for dynamic treatment regimes. *Springer-Verlag. doi*, 10:978–1.

Chakraborty, B., Murphy, S., and Strecher, V. (2010). Inference for non-regular parameters in optimal dynamic treatment regimes. *Statistical Methods in Medical Research*, 19(3):317–343.

Chan, I. S. and Ginsburg, G. S. (2011). Personalized medicine: progress and promise. *Annual Review of Genomics and Human Genetics*, 12(1):217–244.

Cui, Y., Zhu, R., and Kosorok, M. (2017). Tree based weighted learning for estimating individualized treatment rules with censored data. *Electronic Journal of Statistics*, 11(2):3927.

Diamandis, M., White, N. M., and Yousef, G. M. (2010). Personalized medicine: marking a new epoch in cancer patient management. *Molecular Cancer Research*, 8(9):1175–1187.

Duffy, D. J. (2016). Problems, challenges and promises: perspectives on precision medicine. *Briefings in Bioinformatics*, 17(3):494–504.

George, E. I. and McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889.

George, E. I. and McCulloch, R. E. (1997). Approaches for Bayesian variable selection. *Statistica Sinica*, pages 339–373.

Goldberg, Y. and Kosorok, M. R. (2012). Q-learning with censored data. *Annals of Statistics*, 40(1):529.

Guo, B. and Yuan, Y. (2017). Bayesian phase i/ii biomarker-based dose finding for precision medicine with molecularly targeted agents. *Journal of the American Statistical Association*, 112(518):508–520.

Guo, W., Ji, Y., and Catenacci, D. V. (2017). A subgroup cluster-based Bayesian adaptive design for precision medicine. *Biometrics*, 73(2):367–377.

Hoff, P. D. (2009). *A first course in Bayesian statistical methods*, volume 580. Springer.

Ishwaran, H. and Rao, J. S. (2003). Detecting differentially expressed genes in microarrays using Bayesian model selection. *Journal of the American Statistical Association*, 98(462):438–455.

Ishwaran, H. and Rao, J. S. (2005). Spike and slab variable selection: frequentist and Bayesian strategies.

Jameson, J. L. and Longo, D. L. (2015). Precision medicine—personalized, problematic, and promising. *Obstetrical & Gynecological Survey*, 70(10):612–614.

Jia, Z. and Xu, S. (2007). Mapping quantitative trait loci for expression abundance. *Genetics*, 176(1):611–623.

Johnson, K., Wei, W., Weeraratne, D., Frisse, M., Misulis, K., Rhee, K., Zhao, J., and Snowdon, J. (2022). Precision medicine, AI, and the future of personalized health care. clin transl sci. 2021; 14 (1): 86-93.

Kline, A., Wang, H., Li, Y., Dennis, S., Hutch, M., Xu, Z., Wang, F., Cheng, F., and Luo, Y. (2022). Multimodal machine learning in precision health: A scoping review. *npj Digital Medicine*, 5(1):171.

Laber, E. B. and Murphy, S. A. (2011). Adaptive confidence intervals for the test error in classification. *Journal of the American Statistical Association*, 106(495):904–913.

Lauritzen, S. (1996). *Graphical models*. Clarendon Press.

Mitchell, T. J. and Beauchamp, J. J. (1988). Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032.

Murphy, S. A. (2005). A generalization error for Q-learning.

Murphy, S. A. and Bingham, D. (2009). Screening experiments for developing dynamic treatment regimes. *Journal of the American Statistical Association*, 104(485):391–408.

O'Hagan, A. (1995). Fractional Bayes factors for model comparison. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):99–118.

Richardson, S., Bottolo, L., and Rosenthal, J. S. (2010). Bayesian models for sparse regression analysis of high dimensional data. *Bayesian Statistics*, 9:539–569.

Ruffieux, H., Davison, A. C., Hager, J., and Irincheeva, I. (2017). Efficient inference for genetic association studies with multiple outcomes. *Biostatistics*, 18(4):618–636.

Sadee, W. and Dai, Z. (2005). Pharmacogenetics/genomics and personalized medicine. *Human molecular genetics*, 14(suppl_2):R207–R214.

Saha, A., Ha, M. J., Acharyya, S., and Baladandayuthapani, V. (2022). A Bayesian precision medicine framework for calibrating individualized therapeutic indices in cancer. *The Annals of Applied Statistics*, 16(4):2055–2082.

Sanchez, P., Voisey, J. P., Xia, T., Watson, H. I., O'Neil, A. Q., and Tsaftaris, S. A. (2022). Causal machine learning for healthcare and precision medicine. *Royal Society Open Science*, 9(8):220638.

Scott-Boyer, M. P., Imholte, G. C., Tayeb, A., Labbe, A., Deschepper, C. F., and Gottardo, R. (2012). An integrated hierarchical Bayesian model for multivariate eQTL mapping. *Statistical Applications in Genetics and Molecular Biology*, 11(4).

Spanbauer, C. and Sparapani, R. (2021). Nonparametric machine learning for precision medicine with longitudinal clinical trials and Bayesian additive regression trees with mixed models. *Statistics in Medicine*, 40(11):2665–2691.

Spirtes, P. and Zhang, K. (2016). Causal discovery and inference: concepts and recent methodological advances. In *Applied Informatics*, volume 3, pages 1–28. Springer.

Tadesse, M. G. and Vannucci, M. (2021). Handbook of Bayesian variable selection.

Tsiatis, A. A., Davidian, M., Holloway, S. T., and Laber, E. B. (2019). *Dynamic treatment regimes: Statistical methods for precision medicine*. Chapman and Hall/CRC.

WHO (1997). The world health report 1997: conquering suffering; enriching humanity. In *The world health report 1997: conquering suffering; enriching humanity*, pages 162–162.

Wu, H., Shi, W., and Wang, M. D. (2024). Developing a novel causal inference algorithm for personalized biomedical causal graph learning using meta machine learning. *BMC Medical Informatics and Decision Making*, 24(1):137.

Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g-prior distributions. *Bayesian Inference and Decision Techniques*.

Zhang, W., Ma, Z., Wang, L., Fan, D., and Ho, Y.-Y. (2023). Genome-wide search algorithms for identifying dynamic gene co-expression via Bayesian variable selection. *Statistics in Medicine*, 42(30):5616–5629.

Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. (2012). Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118.

Zhao, Z., Banterle, M., Lewin, A., and Zucknick, M. (2024). Multivariate Bayesian structured variable selection for pharmacogenomic studies. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 73(2):420–443.

Zhou, X.-H., Obuchowski, N. A., and McClish, D. K. (2014). *Statistical methods in diagnostic medicine*. John Wiley & Sons.

Chapter 2

Uncertainty quantification and multi-stage variable selection for personalized treatment regimes

2.1 Introduction

Personalized medicine is a data-driven approach to treatment decisions, tailored to each patient by accounting for individual variability in characteristics and clinical history. Dynamic treatment regimes (DTRs) have emerged as a popular framework for implementing personalized care, offering adaptive medical decisions that adjust over time in response to a patient's evolving clinical state (Murphy, 2003, 2005). By exploiting patient-specific information, DTRs aim at recommending optimal treatments at each stage of care, making them particularly effective for managing chronic diseases and complex conditions. The task of constructing optimal DTRs based on patient data can be viewed as a multi-stage decision-making problem, akin to a learning agent navigating an environment by selecting actions that maximize rewards. This framework is rooted in reinforcement learning, a branch of machine learning that has provided numerous solutions to multi-stage decision problems, from dynamic programming (Bellman, 1957) to Q-learning (Watkins, 1989). Q-learning has also been widely applied in the statistical literature, where it has been adapted to address the specific challenges of DTRs (Zhao et al., 2009). Building on this concept, many contributions on DTRs (see Chakraborty and Moodie, 2013; Tsiatis et al., 2019, and references therein) have employed regression-based models, resorting on the principle of backward induction to account for the dynamic nature of the problem. The optimal treatment regime is then determined as the sequence of treatments that maximizes the overall expected response for an individual, also known as the individual value. A coherent Bayesian framework for this class of models has been introduced by Murray et al. (2018). Within this framework, we propose a modeling strategy called Bayesian augmented learning, which is designed to

effectively quantify the uncertainty involved in identifying optimal treatment decision sequences, a fundamental statistical task that, to the best of our knowledge, has been largely overlooked in existing research. A second objective of our proposal addresses the selection of significant variables within the context of DTRs. With the growing availability of prognostic factors for individual patients, identifying the key variables for determining optimal DTRs has become increasingly critical. In the context of regression-based approaches for DTRs, the majority of variable selection methods have been developed within the frequentist paradigm (Qian and Murphy, 2011; Lu et al., 2013; Song et al., 2015; Fan et al., 2016; Wallace et al., 2019; Zhu et al., 2019; Bian et al., 2021). Although less explored in the DTR literature, Bayesian methods offer a compelling alternative, as they naturally accommodate the quantification of uncertainty regarding the significance of prognostic factors. We fill this gap by combining our modeling proposal with a strategy for Bayesian variable selection that is specifically designed to deal with the dynamic nature of DTRs. This will be achieved by adapting the class of spike-and-slab priors (see Tadesse and Vannucci, 2021, for a recent review) to the context of DTRs.

2.1.1 Our contribution

In this work, we propose a Bayesian model for optimising DTRs that addresses the uncertainty involved in identifying optimal decision sequences and incorporates dimensionality reduction to handle high-dimensional individual covariate data. The goal of our analysis is to facilitate personalised care by providing clinicians and decision-makers with tools to assess the likelihood of each treatment being optimal for a specific patient, while also identifying the key prognostic factors critical for determining the optimal sequence of treatments. Our approach is motivated by the analysis of data from the MIMIC-III clinical database (Johnson et al., 2016), which features a various treatment options and a wealth of prognostic factors of potential interest. The specific contributions of this work can be summarised as follows:

i) We propose a modelling strategy called Bayesian augmented learning, which addresses the presence of counterfactual variables in DTRs through a suitable augmentation, in the spirit of Albert and Chib (1993). Counterfactual observations are treated as latent random variables, or pseudo-outcomes, and their imputation is handled seamlessly. This approach facilitates an efficient computational framework for posterior simulation, while also enabling the quantification of uncertainty in identifying optimal

treatments for individuals.

ii) We introduce a class of priors for variable selection, referred to as dependent spike-and-slab, which takes into account the association of significance of a particular covariate that is repeated in multiple stages. Our proposal favors borrowing information across the multi-stage selection of significant factors, thereby facilitating the identification of significant covariates, which is especially appealing in situations with many regressors and limited sample sizes.

iii) We conduct a comprehensive empirical study to evaluate the performance of the model obtained by combining the Bayesian augmented learning scheme with the dependent spike-and-slab prior. Our study demonstrates that this method outperforms current state-of-the-art approaches in identifying optimal treatment sequences and selecting significant variables. We also explore the impact of borrowing information across stages.

iv) We illustrate the practical application of our approach by analysing severe acute arterial hypertension data, extracted from the MIMIC-III clinical database. The goal is to identify the optimal combination of antihypertensive agents to achieve the greatest reduction in maximum systolic blood pressure (SBP) at the individual level.

The rest of the chapter is organised as follows. In Section 2.2, we outline our modelling approach. Section 2.3 describes the Gibbs sampling algorithm used to implement our method. In Section 2.4, we evaluate the finite sample performance of our approach through extensive simulation studies, while Section 2.5 demonstrates its application to the analysis of the MIMIC-III data. The chapter concludes with a discussion in Section 2.6. Additional technical proofs, details on posterior computation, and further results from the clinical case study are provided in the Appendix.

2.2 Bayesian augmented learning with variable selection

We consider a DTR framework with two stages, where $a_j \in \mathcal{A}_j$ indicates the treatment, or action, assigned at the jth stage, with the index $j = 1, 2$ henceforth used to refer to specific stages. We let $Y = Y_1 + Y_2$ denote the overall payoff, with Y_j being the part of the payoff observed after the jth stage. Dealing with specific applications in which Y_1 is not observable is possible by letting $Y_1 = 0$ and $Y = Y_2$. Henceforth we use the index i, with $i = 1, \ldots, n$, to denote quantities referring to the ith statistical unit, e.g. a_{ij} denotes

the treatment assigned to the ith individual in the sample at the jth stage, similarly for Y_i and Y_{ij}. In order to stress the dependence of a payoff on the assigned treatments a_{i1} and a_{i2}, we write $Y_i(a_{i1}, a_{i2}) = Y_{i1}(a_{i1}) + Y_{i2}(a_{i1}, a_{i2})$. The observed payoffs y_i are realizations of the corresponding random variables $Y_i(a_{i1}, a_{i2})$, and the same applies to y_{i1} and y_{i2}. Throughout this work, capital letters denote random variables, while lowercase letters represent their realizations. We further introduce the notation Y_i^{OPT} to denote the overall payoff $Y_i(a_{i1}, a_{i2}^{\text{OPT}})$, where a_{i2}^{OPT} is a treatment that, given that the treatment a_{i1} was assigned at stage 1 and according to an optimal decision rule yet to be specified, is optimal at stage 2. By following the principle of backward induction (see, e.g., Tsiatis et al., 2019), we consider a sequence of two stage-specific regression models, one for the second stage payoffs and one for the optimal overall payoffs. Namely,

$$\begin{aligned} Y_{i2} \mid a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2 &\overset{\text{ind}}{\sim} f_2(y \mid a_{i1}, a_{i2}, \boldsymbol{z}_{i2}; \boldsymbol{\theta}_2) \\ Y_i^{\text{OPT}} \mid a_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1 &\overset{\text{ind}}{\sim} f_1(y \mid a_{i1}, \boldsymbol{z}_{i1}; \boldsymbol{\theta}_1), \end{aligned} \tag{2.1}$$

where $\boldsymbol{z}_{ij}$ denotes the vector of covariates available for the ith individual at the jth stage, f_1 and f_2 are regression functions, $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ are vectors of regression coefficients. To facilitate the identification of causal effects of stage-specific actions, we henceforth adopt the following standard assumptions: the stable unit treatment value assumption (Rubin, 1978) and the assumption of no unmeasured confounders (Rubin, 1980). Model (2.1) is completed by specifying an optimal decision rule: along the lines of a large body of literature on DTR optimization, we rely on an approach from dynamic programming (Bellman, 1957) and identify the optimal treatment regime as the sequence of treatments that maximizes the so-called individual value, that is the expected payoff of an individual, where, without loss of generality, we consider a framework in which larger payoffs are preferable. Accordingly, the optimal treatments at stage 2 and stage 1 are defined as

$$\begin{aligned} a_{i2}^{\text{OPT}} &= \arg\sup_{a \in \mathcal{A}_2} (\mathbb{E}[Y_{i2} \mid a_{i1}, a, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2]), \\ a_{i1}^{\text{OPT}} &= \arg\sup_{a \in \mathcal{A}_1} (\mathbb{E}[Y_i \mid a, a_{i2}^{\text{OPT}}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1]). \end{aligned} \tag{2.2}$$

For the sake of simplicity, we henceforth assume that the solution to the optimization problems in (2.2) is unique: this will be the case, with probability 1, for the specification of f_2 and f_1 considered throught the chapter. A well-known hurdle characterizing DTR models, such as model (2.1), is that, while, given the assigned treatments a_{i1} and a_{i2}, the payoff $Y_i = Y_i(a_{i1}, a_{i2})$ is an observable quantity, $Y_i^{\text{OPT}} = Y_i(a_{i1}, a_{i2}^{\text{OPT}})$ in general is not. To be more specific, Y_i^{OPT} coincides with the observable $Y_i(a_{i1}, a_{i2})$ when $a_{i2}^{\text{OPT}} = a_{i2}$, and is

otherwise counterfactual. The response variable in the second regression in (2.1) is thus not always observable, which immediately highlights the need for missing data imputation. We propose a novel strategy to conveniently deal with both observed and counterfactual payoffs. For the sake of simplicity, we illustrate it by considering binary actions, that is $\mathcal{A}_1 = \mathcal{A}_2 = \{0, 1\}$, although our approach is more general and can be extended to any finite number of treatment options, as in the illustrations of Sections 2.4.2 and 2.5. We henceforth work under the assumption that, at stage 2, observations Y_{i2} and counterfactual payoffs $\bar{Y}_{i2} = Y_{i2}(a_{i1}, 1 - a_{i2})$ are conditionally independent, given $\boldsymbol{\theta}_2$. Similarly for the overall payoffs Y_i^{OPT} and $\bar{Y}_i^{\text{OPT}} = Y_i(1 - a_{i1}, a_{i2}^{\text{OPT}})$, assumed independent conditionally on $\boldsymbol{\theta}_1$. Moreover, we assume that the regression functions f_2 and f_1 in (2.1) are such that the following equivalencies hold:

$$\begin{aligned} a_{i2}^{\text{OPT}} = a_{i2} &\iff \frac{P(Y_{i2} \geq \bar{Y}_{i2})}{P(Y_{i2} \leq \bar{Y}_{i2})} > 1, \\ a_{i1}^{\text{OPT}} = a_{i1} &\iff \frac{P(Y_i^{\text{OPT}} \geq \bar{Y}_i^{\text{OPT}})}{P(Y_i^{\text{OPT}} \leq \bar{Y}_i^{\text{OPT}})} > 1. \end{aligned} \tag{2.3}$$

The validity of such an assumption is key and will be discussed in Section 2.2.1. The double implications in (2.3) suggest an alternative strategy to implement the optimal rule (2.2): rather than maximising the expected values in (2.2), a stage-specific optimal treatment can be identified as the one that is more likely to lead to a larger payoff. In order to put this idea in action, we introduce n pairs of augmenting random vectors $(V_{i2}^{(0)}, V_{i2}^{(1)})$ and $(V_i^{(0)}, V_i^{(1)})$, henceforth referred to as stage 2 and overall pseudo-outcomes, which, conditionally on $\boldsymbol{\theta}_2$ and $\boldsymbol{\theta}_1$, are independent of and distributed as $(Y_{i2}(a_{i1}, 0), Y_{i2}(a_{i1}, 1))$ and $(Y_i(0, a_{i2}^{\text{OPT}}), Y_i(1, a_{i2}^{\text{OPT}}))$, respectively. The pseudo-outcomes can be interpreted as vectors of potential payoffs of individuals with identical covariate information as those for which observations were actually collected. We thus define the random variables $A_{i2}^{\text{OPT}} = \mathbb{1}\{V_{i2}^{(1)} > V_{i2}^{(0)}\}$ and $A_{i1}^{\text{OPT}} = \mathbb{1}\{V_i^{(1)} > V_i^{(0)}\}$, and observe the mode of A_{ij}^{OPT} is a_{ij}^{OPT}, while its expectation coincides with $P(V_{i2}^{(1)} > V_{i2}^{(0)})$ if $j = 2$, and $P(V_i^{(1)} > V_i^{(0)})$ if $j = 1$. The joint distribution of stage 2 observations Y_{i2} and stage 2 pseudo-outcomes $(V_{i2}^{(0)}, V_{i2}^{(1)})$ defines an augmented version of the stage 2 regression model in (2.1); similarly for the overall optimal payoffs Y_i^{OPT}, the overall pseudo-outcomes $(V_i^{(0)}, V_i^{(1)})$ and the second regression in (2.1). Moreover, if we assume that $Y_i^{\text{OPT}} \stackrel{\text{d}}{=} Y_{i1} + \max(V_{i2}^{(0)}, V_{i2}^{(1)})$, with the notation $\stackrel{\text{d}}{=}$ used to denote identity in distribution, and we specify a prior for the regression coefficients, we can write the following model

$$(Y_{i2}, V_{i2}^{(0)}, V_{i2}^{(1)}) \mid a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2 \stackrel{\text{ind}}{\sim} f_2(y \mid a_{i1}, a_{i2}, \boldsymbol{z}_{i2}; \boldsymbol{\theta}_2)$$

$$
\begin{aligned}
& \times f_2(v_0 \mid a_{i1}, 0, \boldsymbol{z}_{i2}; \boldsymbol{\theta}_2) \\
& \times f_2(v_1 \mid a_{i1}, 1, \boldsymbol{z}_{i2}; \boldsymbol{\theta}_2), \\
(Y_{i1}, V_i^{(0)}, V_i^{(1)}) \mid V_{i2}^{(0)}, V_{i2}^{(1)}, a_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1 & \overset{\text{ind}}{\sim} f_1(y + \max(V_{i2}^{(0)}, V_{i2}^{(1)}) \mid a_{i1}, \boldsymbol{z}_{i1}; \boldsymbol{\theta}_1) \\
& \times f_1(v_0 \mid 0, \boldsymbol{z}_{i1}; \boldsymbol{\theta}_1) \\
& \times f_1(v_1 \mid 1, \boldsymbol{z}_{i1}; \boldsymbol{\theta}_1), \\
(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) & \sim p(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2).
\end{aligned}
\tag{2.4}
$$

The introduction of the pseudo-outcomes allows us to devise an efficient computational scheme for posterior simulation with three notable properties: i) counterfactual observations are treated as latent random variables and their imputation is seemingly handled; ii) a single computational scheme is used to sample from the joint posterior distribution of the parameters involved in the two stages, thus avoiding two-step procedures such as the backward induction Gibbs sampler of Murray et al. (2018); iii) the uncertainty associated with the identification of individual optimal treatments can be quantified through the posterior distribution of A_{i2}^{OPT} and A_{i1}^{OPT}.

We will henceforth refer to the approach utilizing model (2.4) to address DTR problems as *Bayesian Augmented Learning* (BAL). The BAL model can be easily extended to accommodate more than two treatment options, i.e., when the sets $\mathcal{A}_1$ and $\mathcal{A}_2$ have cardinality $T \geq 2$ as in the illustrations of Sections 2.4.2 and 2.5. This is achieved by considering T-dimensional stage 2 and overall pseudo-outcomes, denoted as $V_{i2}^{(0:(T-1))} = (V_{i2}^{(0)}, \ldots, V_{i2}^{(T-1)})$ and $V_i^{(0:(T-1))} = (V_i^{(0)}, \ldots, V_i^{(T-1)})$, respectively. The specification of the two key components in (2.4), namely the regression functions f_2 and f_1, and the prior $p(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$, will be discussed in Sections 2.2.1 and 2.2.2.

2.2.1 Regression models

While BAL offers flexibility in specifying f_2 and f_1, in view of the application of Section 2.5, we consider continuous payoffs and model them with normal multiple linear regressions, namely

$$
\begin{aligned}
f_2(y; a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2) &= f_{\text{N}}(y; \boldsymbol{\theta}_2^{\intercal}[\boldsymbol{z}_{i2} \times (a_{i1}, a_{i2})], \sigma_2^2), \\
f_1(y; a_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1) &= f_{\text{N}}(y; \boldsymbol{\theta}_1^{\intercal}[\boldsymbol{z}_{i1} \times a_{i1}], \sigma_1^2),
\end{aligned}
\tag{2.5}
$$

where $f_{\text{N}}(y; \mu, \sigma^2)$ is the probability density function of a normal random variable with mean μ and variance σ^2, and the notation $[\boldsymbol{z} \times \boldsymbol{a}]$ indicates the vector of regressors including $\boldsymbol{z}$, $\boldsymbol{a}$, and the interactions between the covariates $\boldsymbol{z}$ and each of the binary variables in $\boldsymbol{a}$. While

for simplicity we exclude interactions between covariates and the term $a_{i1}a_{i2}$, these could be incorporated in the stage 2 regressions, if necessary. Given the continuity of the payoffs, the equivalences in (2.3) can be rewritten as

$$\begin{aligned} a_{i2}^{\text{OPT}} = a_{i2} &\iff P(Y_{i2} \geq \bar{Y}_{i2}) > 0.5, \\ a_{i1}^{\text{OPT}} = a_{i1} &\iff P(Y_i^{\text{OPT}} \geq \bar{Y}_i^{\text{OPT}}) > 0.5. \end{aligned} \tag{2.6}$$

The next proposition formalizes the validity of (2.3) and thus of (2.6).

Proposition 1. If f_2 and f_1 are normal linear regressions, specified as in (2.5), then the equivalences in (2.3) hold.

The proof follows from the symmetry of the normal distribution and is postponed to Appendix 2.A. It is worth stressing that the applicability of the BAL approach is broad and not limited to the normal regression case, as the equivalences in (2.3) hold also for other specifications of the regression functions f_2 and f_1, e.g. logistic and Poisson, relevant in various applications. More details are provided in Appendix 2.A.

2.2.2 Multi-stage variable selection

For the sake of a compact notation, we introduce $\tilde{\boldsymbol{z}}_{ij}$ to denote the p_j-dimensional vector of regressors at the jth stage for the ith individual, including both the covariates $\boldsymbol{z}_{ij}$ and the corresponding interactions with current and past treatments. The focus of our work is on data, such as the MIMIC-III clinical database, where the number of individual covariates at each stage is large. This emphasizes the need for dimensionality reduction, especially when interaction terms are incorporated into the regressions, as in (2.5). We propose to achieve this by defining a joint spike-and-slab prior for the regression coefficients $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$, which we name *dependent spike-and-slab* (DSS) prior. Our proposal takes into account that some regressors, and similarly some interactions between regressors and stage-specific actions, might appear in both $\tilde{\boldsymbol{z}}_{i1}$ and $\tilde{\boldsymbol{z}}_{i2}$. Without loss of generality, we assume the common elements are in the first $p \leq \min(p_1, p_2)$ positions of $\tilde{\boldsymbol{z}}_{i1}$ and $\tilde{\boldsymbol{z}}_{i2}$. While the value taken by these regressors may vary across stages, it seems appealing to define a joint model for the corresponding regression coefficients, the rationale being that the same regressor is likely to play similar roles within the two regression models. To this end, we design a prior with positively correlated inclusion variables for the same regressor across the two stages. Among the possible specifications of spike-and-slab priors considered in the literature (see, e.g., Tadesse and Vannucci, 2021), with the purpose of simplifying computations, we opt

for absolutely continuous spikes, that is we consider both normal spike and normal slab components, with a ratio r between the variance of the spike and that one of the slab, considerably smaller than 1. We let δ_{lj}, with $l = 1, \dots, p_j$, be an indicator variable such that $\delta_{lj} = 1$ ($\delta_{lj} = 0$) if the lth regressor is (is not) included in the jth regression, and call $w_{lj} = P(\delta_{lj} = 1)$. We assume the prior variance of the regression coefficient θ_{lj} is equal to ψ_{lj} if $\delta_{lj} = 1$, and equal to $r\psi_{lj}$ if $\delta_{lj} = 0$. We introduce the notation $\boldsymbol{\delta}_j = (\delta_{1j}, \dots, \delta_{p_j j})$, $\boldsymbol{\psi}_j = (\psi_{1j}, \cdots, \psi_{p_j j})$ and $\boldsymbol{w}_j = (w_{1j}, \dots, w_{pj})$, and define the DSS prior $p(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ as the distribution arising from the following hierarchy:

$$
\begin{aligned}
(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \mid \boldsymbol{\delta}_1, \boldsymbol{\delta}_2, \boldsymbol{\psi}_1, \boldsymbol{\psi}_2 &\sim \prod_{j=1}^{2} \prod_{l_j=1}^{p_j} \mathrm{N}\left(0, r^{1-\delta_{l_j j}} \psi_{l_j j}\right) \\
(\boldsymbol{\psi}_1, \boldsymbol{\psi}_2) &\sim \prod_{j=1}^{2} \prod_{l_j=1}^{p_j} \textsc{inv-Gamma}(\nu, Q) \\
(\boldsymbol{\delta}_1, \boldsymbol{\delta}_2) \mid \boldsymbol{w}_1, \boldsymbol{w}_2 &\sim \prod_{j=1}^{2} \prod_{l_j=1}^{p_j} \textsc{Bern}(w_{l_j j}) \qquad (2.7) \\
w_{l1} = w_{l2} &\overset{\text{iid}}{\sim} \textsc{Beta}(a, b) \qquad l = 1, \dots, p \\
w_{lj} &\overset{\text{iid}}{\sim} \textsc{Beta}(a, b) \qquad l_j = p+1, \dots, p_j, \ j = 1, 2.
\end{aligned}
$$

The distribution in (2.7) is an extension to the case of two vectors of regression coefficients of the normal mixture of inverse-Gamma spike-and-slab prior proposed by Ishwaran and Rao (2003, 2005). The indicator variables of common regressors, that is the pairs $(\delta_{l1}, \delta_{l2})$ for $l = 1, \dots, p$, are marginally dependent, which is an appealing property as it favours borrowing of information across the two stages when performing variable selection. The dependent spike-and-slab prior thus formalizes the idea that there might be positive correlation between the significance of the same regressor across stages. Following the DSS prior specification, the correlation coefficient between $\delta_{l_1 1}$ and $\delta_{l_2 2}$ is equal to

$$
\rho(\delta_{l_1 1}, \delta_{l_2 2}) = \frac{1}{1 + a + b} \qquad (2.8)
$$

if $l_1 = l_2 \in \{1, \dots, p\}$, and 0 otherwise. The proof is provided in Appendix 2.B.

2.3 Posterior computation

The joint distribution of observations $\boldsymbol{Y}_1 = (Y_{11}, \ldots, Y_{n1})$ and $\boldsymbol{Y}_2 = (Y_{12}, \ldots, Y_{n2})$, pseudo-outcomes $\boldsymbol{V}_2 = (V_{12}^{(0)}, V_{12}^{(1)}, \ldots, V_{n2}^{(0)}, V_{n2}^{(1)})$ and $\boldsymbol{V} = (V_1^{(0)}, V_1^{(1)}, \ldots, V_n^{(0)}, V_n^{(1)})$, and parameters $\varphi = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\delta}_1, \boldsymbol{\delta}_2, \boldsymbol{\psi}_1, \boldsymbol{\psi}_2, \boldsymbol{w}_1, \boldsymbol{w}_2, \sigma_1^2, \sigma_2^2)$ is described by combining the BAL model in (2.4), completed with normal regression functions as in (2.5), and DSS prior as in (2.7). Additionally, the hyperparameters σ_1^2 and σ_2^2 are assigned independent inverse gamma priors, each with shape and scale parameters set to 1/2. From this, we devise a Gibbs sampling scheme for the posterior distribution $p(\varphi, \boldsymbol{V}_2, \boldsymbol{V} \mid \boldsymbol{y}_1, \boldsymbol{y}_2)$, from which it is straightforward to estimate the posterior distribution $p(\varphi \mid \boldsymbol{y}_1, \boldsymbol{y}_2)$ and its functionals. The steps of the Gibbs sampler are summarized in Algorithm 1 and involve a sequential update of the parameters characterizing the two stage-specific regressions and the pseudo-outcomes. Thanks to the conjugate specification of the priors, the full conditional distributions involved in the algorithm are simple to sample from, as detailed in Appendix 2.C. Again, the extension to the case with T treatment options at each stage is straightforward.

Algorithm 1: Gibbs sampler for BAL with DSS

Data: Input data
Result: Final result
Initialize parameters $\varphi^{(0)}$ ▷ Superscript (r) to denote iteration r
Initialize pseudo-outcomes $(\boldsymbol{V}_2^{(0)}, \boldsymbol{V}^{(0)})$
for $r = 1$ **to** R **do**
  Update parameters $\varphi^{(r)}$;
  Update pseudo-outcomes $(\boldsymbol{V}_2^{(r)}, \boldsymbol{V}^{(r)})$;
  Compute $(A_{i2}^{\text{OPT}(r)}, A_{i1}^{\text{OPT}(r)})$, for $i = 1, \ldots, n$.
Estimate $\hat{a}_{ij}^{\text{OPT}}$ as mode in $\{A_{ij}^{\text{OPT}(R_0+1)}, \ldots, A_{ij}^{\text{OPT}(R)}\}$, for $j = 1, 2$ ▷ R_0: burn-in iterations

While the inclusion of the psuedo-outcomes $\boldsymbol{V}_2$ and $\boldsymbol{V}$ in the model was instrumental to the goal of sampling from $p(\varphi \mid \boldsymbol{y}_1, \boldsymbol{y}_2)$, Algorithm 1 also provides a means to sample from their posterior distribution. The probability $\mathbb{E}[A_{ij}^{\text{OPT}}]$ that the optimal treatment at the jth stage for the ith individual is equal to 1 can be estimated with the frequency of the event $A_{ij}^{\text{OPT}} = 1$, as defined in (2.3), in the posterior sample. The corresponding estimated individual optimal treatment $\hat{a}_{ij}^{\text{OPT}}$ is given by the mode of the posterior distribution of A_{ij}^{OPT}.

The hyper-parameters a and b control both the prior distribution of the probability of inclusion of a regressor, and, as apparent in (2.8), the correlation across the significance

of common regressors across stages. We assign them a prior distribution and discuss in Appendix 2.C a Metropolis–Hastings step to deal with their update.

2.4 Simulation studies

We conduct an extensive simulation study to assess the performance of the BAL model and to compare two of its specifications with an alternative approach from the literature. Throughout the study, we consider three distinct modeling strategies:

i) BAL with normal regression functions (2.5) and DSS prior (2.7), simply referred to as DSS;

ii) BAL with normal regression functions (2.5) and independent spike-and-slab priors (ISS), simply referred to as ISS;

iii) Q-learning with lasso penalty (Qian and Murphy, 2011), referred to as QLL.

The ISS specification is obtained by ignoring the presence of common regressors across stages in the definition of the DSS prior, that is by setting $p = 0$ in (2.7). As for DSS and ISS, we set $r = 0.001$, $\nu = 3$, $Q = 4$, and assign a and b independent inverse-gamma priors with shape and scale parameters equal to 1. The QLL approach is implemented by using the R package `DTRlearn2` (Chen et al., 2020), with the lasso penalty parameter optimised with the R package `glmnet` (Friedman et al., 2021).
Our investigation focuses on the models' ability to accurately identify significant regressors across stages, and on the effect of variable selection on identifying the optimal DTR. Throughout the study, two data generating processes are considered, both defined as minor modifications of the class of generative models developed by Chakraborty et al. (2010), and later used, for example, in Chakraborty et al. (2013) and Laber et al. (2014).
The simulation study is divided into two experiments: Section 2.4.1 considers a range of scenarios with different numbers of stage-specific regressors, and varying degrees of correlation between the significance of common regressors across stages; Section 2.4.2 focuses on the analysis of simulated datasets characterized by varying number of treatments, thus extending the framework described in Section 2.2 to sets $\mathcal{A}_1$ and $\mathcal{A}_2$ with cardinality larger than two. As far as variable selection is concerned, model performance is assessed by resorting to three evaluation indices: proportion of false negatives (FN), that is regressors that are incorrectly deemed as not significant; proportion of false positives (FP), that is regressors incorrectly identified as significant; and F_1 score, combining precision and recall

into a single metric. At the same time, the capability of a model to correctly identify the optimal treatment is assessed by means of the treatment error rate (ER) and the payoff mean relative error (MRE). When specific to the jth stage, these two measures are defined as

$$\mathrm{ER}_j = \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}\{\hat{a}_{ij}^{\mathrm{OPT}} \neq a_{ij}^{\mathrm{OPT}}\}, \quad \mathrm{MRE}_j = \frac{1}{n}\sum_{i=1}^{n} \frac{|\mathbb{E}[Y_{ij}^{\mathrm{OPT}}] - \mathbb{E}[Y_{ij}^{\widehat{\mathrm{OPT}}}]|}{\mathbb{E}[Y_{ij}^{\mathrm{OPT}}]},$$

where $\mathbb{E}[Y_{ij}^{\mathrm{OPT}}]$ and $\mathbb{E}[Y_{ij}^{\widehat{\mathrm{OPT}}}]$ are the expected payoffs at the jth stage, respectively under the assumptions that the assigned treatments are, respectively, optimal and identified as optimal. Analogously, overall measures of model performance can be defined as

$$\mathrm{ER} = \frac{1}{n}\sum_{i=1}^{n} \mathbb{1}\{(\hat{a}_{i1}^{\mathrm{OPT}}, \hat{a}_{i2}^{\mathrm{OPT}}) \neq (a_{i1}^{\mathrm{OPT}}, a_{i2}^{\mathrm{OPT}})\}, \quad \mathrm{MRE} = \frac{1}{n}\sum_{i=1}^{n} \frac{|\mathbb{E}[Y_{i}^{\mathrm{OPT}}] - \mathbb{E}[Y_{i}^{\widehat{\mathrm{OPT}}}]|}{\mathbb{E}[Y_{i}^{\mathrm{OPT}}]}.$$

While the ERs only takes into account the actual optimality of the treatment identified as optimal, the MREs also accounts for the magnitude of the variation in terms of potential payoff, when a suboptimal treatment is mistakenly identified as optimal.

2.4.1 First experiment

We let k be the number of individual covariates available at each stage and we denote by $\boldsymbol{z}_{ij} = (1, z_{ij1}, \ldots, z_{ijk})$ the $(k+1)$-dimensional vector of covariates specific to the jth stage, whose first component is set equal to 1. We assume that $z_{ijl} \in \{-1, 1\}$ for every $l = 1, \ldots, k$, and $a_{ij} \in \{0, 1\}$. A data generating process for synthetic data consisting of stage-specific covariates $\boldsymbol{z}_{ij}$, treatments a_{ij}, and stage-specific payoffs y_{ij}, is defined by means of the following chain of conditional distributions:

$$\begin{aligned}
&A_{ij} \overset{\text{iid}}{\sim} \textsc{Bern}(\{0,1\}; 0.5) && \text{(treatments)}\\
&Z_{i1l} \overset{\text{iid}}{\sim} \textsc{Bern}(\{-1,1\}; 0.5), \text{ for } l = 1, \ldots, k && \text{(covariates)}\\
&Z_{i2l} \mid z_{i1l} \overset{\text{ind}}{\sim} \textsc{Bern}(\{-1,1\}; 1/(1+\exp\{z_{i1l}\}), \text{ if } l = 1, \ldots, \lfloor k/2 \rfloor &&\\
&Z_{i2l} = z_{i1l}, \text{ if } l = \lfloor k/2 \rfloor + 1, \ldots, k && (2.9)\\
&Y_{i1} \mid a_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^* \overset{\text{ind}}{\sim} \mathrm{N}(\boldsymbol{\theta}_1^{*\intercal}[\boldsymbol{z}_{i1} \times a_{i1}] - \mathbb{E}[Y_{i2}^{\mathrm{OPT}}], 1) && \text{(payoffs)}\\
&\mathbb{E}[Y_{i2}^{\mathrm{OPT}}] = \max(\boldsymbol{\theta}_2^{*\intercal}[\boldsymbol{z}_{i2} \times (a_{i1}, 0)], \boldsymbol{\theta}_2^{*\intercal}[\boldsymbol{z}_{i2} \times (a_{i1}, 1)]) &&\\
&Y_{i2} \mid a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2^* \overset{\text{ind}}{\sim} \mathrm{N}(\boldsymbol{\theta}_2^{*\intercal}[\boldsymbol{z}_{i2} \times (a_{i1}, a_{i2})], 1), &&
\end{aligned}$$

where the notation $X \sim \text{Bern}(\{a, b\}; p)$ is used to indicate $P(X = b) = 1 - P(X = a) = p$. According to (2.9), covariates $\boldsymbol{z}_{ij}$ are generated to mimic a realistic scenario where about half remain constant across stages, while the values taken by other half evolve from stage 1 to stage 2. The data generating process in (2.9) relies on the specification of two vectors of regression coefficients $\boldsymbol{\theta}_1^*$ and $\boldsymbol{\theta}_2^*$, of dimension $p_1 = 2(k+1)$ and $p_2 = 3(k+1)$, respectively. To generate data for which the significance of predictors is correlated across stages, we define the distribution of $(\boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*)$ along the lines of the dependent spike-and-slab prior (2.7). Namely,

$$\begin{aligned}
\omega_{l1}^* = \omega_{l2}^* &\overset{\text{iid}}{\sim} \text{Beta}(a^*, b^*), \text{ for } l = 1, \ldots, p_1; \\
\omega_{l2}^* &\overset{\text{iid}}{\sim} \text{Beta}(a^*, b^*), \text{ for } l = p_1 + 1, \ldots, p_2; \\
(\delta_{l_1 1}^*, \delta_{l_2 2}^*) \mid \omega_{l_1 1}^*, \omega_{l_2 2}^* &\overset{\text{ind}}{\sim} \text{Bern}(\omega_{l_1 1}^*) \times \text{Bern}(\omega_{l_2 2}^*), \text{ for } l_j = 1, \ldots, p_j; \\
(\theta_{l_1 1}^*, \theta_{l_2 2}^*) \mid \delta_{l_1 1}^*, \delta_{l_2 2}^* &\overset{\text{ind}}{\sim} \text{N}(0, 2 \cdot 10^{-5(1-\delta_{l_1 1}^*)}) \times \text{N}(0, 2 \cdot 10^{-5(1-\delta_{l_2 2}^*)}), \text{ for } l_j = 1, \ldots, p_j.
\end{aligned} \tag{2.10}$$

We consider nine simulation scenarios, arising as the combination of $k \in \{10, 20, 30\}$ (i.e. $p_1 \in \{22, 42, 62\}$ and $p_2 \in \{33, 63, 93\}$, when counting also for interactions between covariates $\boldsymbol{z}_{ij}$ and current and past treatments), and three levels for the correlation $\rho^* = \rho(\delta_{l1}^*, \delta_{l2}^*)$, when $l \in \{1, \ldots, p_1\}$, namely weak (0.3), medium (0.6) and strong (0.9) correlation. These levels are obtained through the specification of the parameters (a^*, b^*) in (2.10): by applying (2.8), we have that $(a^*, b^*) \in \{(7/10, 49/30), (1/5, 7/15), (1/30, 7/90)\}$ implies $\rho^* = \{0.3, 0.6, 0.9\}$, respectively. These specifications for the parameters (a^*, b^*) also imply that the expected proportion of significant variables at each stage is 0.3, as given by $a^*/(a^* + b^*)$. We consider datasets simulated under two regimes, namely n smaller than p_2 and n larger than p_2. This is achieved by considering two sample sizes for each value of k: $n \in \{25, 50\}$ if $k = 10$, $n \in \{50, 100\}$ if $k = 20$ and $n \in \{75, 150\}$ if $k = 30$. While the larger sample sizes satisfy $n > p_2 > p_1$, the analysis of the smaller samples is more challenging as in that case $p_1 < n < p_2$, resulting in more regressors than observations at stage 2. For each scenario and each regime, 100 datasets are independently generated from the data generating process obtained by combining (2.9) and (2.10). Both versions of the BAL model, that is DSS and ISS, are fit to all datasets by running Algorithm 1 for 10,000 iterations, of which the first 5,000 are discarded as burn-in.

Table 2.1 summarises the performance of DSS, ISS and QLL in identifying the significant variables when analysing data simulated under the regime where n is smaller than p_2. The F_1 scores clearly demonstrate that DSS outperforms both ISS and QLL at both stages and across all scenarios. For instance, at stage 2, with strong correlation and a sample size

Table 2.1: *First experiment: simulated data with n smaller than p_2. For stage 2 and stage 1, three summaries of the accuracy in selecting significant variables (proportion of FN, proportion of FP, and F_1 score), for DSS, ISS, and QLL, across the nine scenarios obtained by specifying $(k, n) \in \{(10, 25), (20, 50), (30, 75)\}$ and $\rho^* \in \{0.3, 0.6, 0.9\}$. Results are based on 100 replicated datasets.*

Stage 2											
n	p_2	ρ^*	FN			FP			F_1 score		
			DSS	ISS	QLL	DSS	ISS	QLL	DSS	ISS	QLL
25	33	0.3	0.119	0.191	0.397	0.187	0.248	0.132	0.775	0.677	0.626
		0.6	0.070	0.226	0.352	0.113	0.286	0.133	0.824	0.707	0.636
		0.9	0.161	0.091	0.363	0.149	0.318	0.121	0.759	0.714	0.643
50	63	0.3	0.065	0.142	0.267	0.133	0.311	0.181	0.820	0.709	0.666
		0.6	0.066	0.288	0.345	0.111	0.201	0.114	0.845	0.677	0.664
		0.9	0.097	0.125	0.408	0.164	0.191	0.116	0.799	0.718	0.613
75	93	0.3	0.071	0.201	0.302	0.152	0.261	0.203	0.801	0.723	0.645
		0.6	0.078	0.058	0.273	0.096	0.432	0.204	0.858	0.685	0.665
		0.9	0.065	0.236	0.446	0.071	0.271	0.052	0.886	0.680	0.717
Stage 1											
n	p_1	ρ^*	FN			FP			F_1 score		
			DSS	ISS	QLL	DSS	ISS	QLL	DSS	ISS	QLL
25	22	0.3	0.114	0.117	0.409	0.167	0.399	0.153	0.788	0.666	0.612
		0.6	0.092	0.165	0.411	0.160	0.393	0.140	0.761	0.623	0.643
		0.9	0.030	0.375	0.326	0.147	0.143	0.149	0.804	0.667	0.665
50	42	0.3	0.046	0.151	0.383	0.197	0.321	0.160	0.783	0.652	0.618
		0.6	0.058	0.213	0.293	0.112	0.395	0.166	0.859	0.626	0.671
		0.9	0.048	0.125	0.368	0.160	0.235	0.157	0.826	0.691	0.638
75	62	0.3	0.094	0.197	0.313	0.169	0.331	0.160	0.787	0.657	0.667
		0.6	0.069	0.056	0.321	0.112	0.561	0.089	0.832	0.611	0.724
		0.9	0.056	0.180	0.393	0.061	0.299	0.091	0.907	0.690	0.650

of 75, DSS achieves an F_1 score of 0.886, compared to QLL's 0.717 and ISS's 0.680. The performance of ISS and QLL, compared in terms of F_1 score, is rather similar, with the ISS that appears to perform better when the correlation between the significance of common regressors across the two stages is weak. These differences in the behaviour of the three models are further highlighted by the FN and FP rates. DSS exhibits the lowest FN rate in most scenarios, while QLL consistently displays the highest FN rate. On the other hand, DSS and QLL achieves the similar FP rate, whereas ISS displays the largest FP rates. In comparison, DSS shows a more balanced performance in both FN and FP rates. The results of a similar analysis, this time conducted on data simulated under the regime where n exceeds p_2, are presented in Table 2.5 in the Appendix. A comparison of the results

between the two regimes reveals notable differences in the methods' performance. When $n > p_2$, DSS and ISS exhibit similar performance across both stages and all scenarios, with DSS most of the times achieving slightly higher F_1 scores, especially in scenarios with strong correlation. In contrast, QLL consistently underperforms across all three metrics, showing notably lower F_1 scores, particularly in cases with small sample sizes and weak correlations.

Table 2.2 focuses on data where n is smaller than p_2 and summarizes the empirical performance of the three methods by comparing the estimated optimal treatment regimes to the true optimal regimes. Across all scenarios and both stages, QLL consistently exhibits the weakest performance in terms of both MRE and ER. In most cases, DSS outperforms ISS on both metrics and across both stages. The overall performance measures, which account for both stages, lead to the same conclusions: DSS demonstrates clear superiority over ISS in terms of overall MRE and ER, with the difference being particularly pronounced when the sample size is large. For instance, in the strong correlation scenario with $n = 75$, DSS achieves an overall MRE of 0.006, while ISS records 0.070. Similarly, the overall ER for DSS is 0.041, compared to 0.176 for ISS. The results of a similar analysis, this time focusing on data where n exceeds p_2, are presented in Table 2.6 in the Appendix. Compared to the results in Table 2.2, the performance gap between DSS and ISS narrows, with DSS showing only a slight advantage over ISS across most scenarios, based on both stage-specific and overall performance metrics. Both methods consistently outperform QLL, which struggles especially in scenarios with smaller sample sizes.

In conclusion, comparing the results of DSS and ISS under the two regimes where n is smaller or larger than p_2 and across different levels of correlation in the data-generating process provides valuable insights. The use of the DSS prior is particularly advantageous when the sample size is small relative to the number of regressors. Furthermore, the performance difference between the two methods becomes more pronounced when the true data-generating process exhibits strong correlation in the mechanism that governs the significance of the same covariates across stages.

2.4.2 Second experiment

In the second experiment, we set $\mathcal{A}_1 = \mathcal{A}_2 = \{0, 1, \ldots, T-1\}$ and we consider four simulation scenarios, defined by varying the number of treatments $T \in \{4, 8\}$ and the number of observations $n \in \{200, 400\}$. For this part, the number k of individual covariates is fixed equal to 10, and the data are generated with $\rho^* = 0.6$. Similarly to (2.9), a data

generating process for synthetic data consisting of stage-specific covariates $\boldsymbol{z}_{ij}$, treatments a_{ij}, and stage-specific payoffs y_{ij}, is defined by means of the following chain of conditional distributions:

$$
\begin{aligned}
&A_{ij} \overset{\text{iid}}{\sim} \textsc{Unif}(\mathcal{A}_j) && \text{(treatments)}\\
&A^*_{ijt} = \mathbb{I}_{\{t\}}(A_{ij}),\ \text{for } t = 1, \dots, T-1 && \text{(dummy variables)}\\
&Z_{i1l} \overset{\text{iid}}{\sim} \textsc{Bern}(\{-1, 1\}; 0.5),\ \text{for } l = 1, \dots, k && \text{(covariates)}\\
&Z_{i2l} \mid z_{i1l} \overset{\text{ind}}{\sim} \textsc{Bern}(\{-1, 1\}; 1/(1 + \exp\{z_{i1l}\}),\ \text{if } l = 1, \dots, \lfloor k/2 \rfloor\\
&Z_{i2l} = z_{i1l},\ \text{if } l = \lfloor k/2 \rfloor + 1, \dots, k\\
&Y_{i1} \mid \boldsymbol{a}^*_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}^*_1, \boldsymbol{\theta}^*_2 \overset{\text{ind}}{\sim} \mathrm{N}(\boldsymbol{\theta}^{*\intercal}_1[\boldsymbol{z}_{i1} \times \boldsymbol{a}^*_{i1}] - \mathbb{E}[Y^{\text{OPT}}_{i2}], 1) && \text{(payoffs)}\\
&\mathbb{E}[Y^{\text{OPT}}_{i2}] = \max_{\boldsymbol{a}^* \in \mathcal{D}_{T-1}} (\boldsymbol{\theta}^{*\intercal}_2[\boldsymbol{z}_{i2} \times (\boldsymbol{a}^*_{i1}, \boldsymbol{a}^*)])\\
&Y_{i2} \mid \boldsymbol{a}^*_{i1}, \boldsymbol{a}^*_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}^*_2 \overset{\text{ind}}{\sim} \mathrm{N}(\boldsymbol{\theta}^{*\intercal}_2[\boldsymbol{z}_{i2} \times (\boldsymbol{a}^*_{i1}, \boldsymbol{a}^*_{i2})], 1),
\end{aligned}
\tag{2.11}
$$

where $\mathbb{I}_{\{t\}}(A_{ij})$ is indicator function, equal to 1 if $A_{ij} = t$ and 0 otherwise; $\textsc{Unif}(\mathcal{A}_j)$ denotes the discrete uniform distribution on $\mathcal{A}_j$, and $\boldsymbol{a}^*_{ij} = (a^*_{ij1}, \dots, a^*_{ij(T-1)}) \in \mathcal{D}_{T-1}$ indicates the vector of dummy variables encoding the treatment assigned to the ith individual at the jth stage. Moreover, $\mathcal{D}_{T-1}$ denotes the set of $(T-1)$-dimensional vectors consisting of at most one element equal to one and the remaining equal to zero. The vectors $\boldsymbol{\theta}^*_1$ and $\boldsymbol{\theta}^*_2$ are simulated by following (2.10) with $\rho^* = 0.6$ as a result of $(a^*, b^*) = (1/5, 7/15)$, and with $p_1 = T(k+1)$ and $p_2 = (2T-1)(k+1)$. Given that $k = 10$, we have $(p_1, p_2) = (44, 77)$ when $T = 4$ and $(p_1, p_2) = (88, 165)$ when $T = 8$. We thus note that all the scenarios considered for this experiment fall under the regime where $n > p_2$. For each scenario, 100 datasets are independently generated from the data-generating process defined by combining (2.9) and (2.10). Both version of the BAL model, that is DSS and ISS, are fit to all datasets by running Algorithm 1 for 10,000 iterations, discarding the first 5,000 as burn-in.

When analyzing datasets with more than two treatments, the advantages of the DSS approach become more evident, compared to ISS and QLL. Table 2.3 summarizes the performance of the three methods in selecting the significant regressors. DSS consistently outperforms the other methods, showing lower FN and FP rates in most cases and achieving higher F_1 scores across all scenarios. Additionally, as the sample size increases from $n = 200$ to $n = 400$, all methods demonstrate improved performance, reflected in higher F_1 scores and reduced FN and FP rates.

As shown in Table 2.4, DSS also appears to be the method that leads to the most accurate prediction, with notably lower MRE and ER values compared to ISS and QLL. In

contrast, QLL exhibits the highest MRE and ER across all scenarios, making it the least effective method regardless of sample size or number of treatments. Also when it comes to prediction accuracy, increasing the sample size from 200 to 400 improves the performance of all methods.

2.5 Clinical case study

We analyze a clinical case study that focuses on hospital care medications for patients in intensive care units (ICU) with severe acute arterial hypertension. Data are extracted from the MIMIC-III clinical database (Johnson et al., 2016), using free-text clinical information extraction techniques, as outlined by Zhou et al. (2022), to retrieve relevant clinical data for analysis from unstructured text in medical records. The estimation of DTRs is of particular interest because personalized antihypertensive treatments have been shown to result in more effective interventions for managing hypertension (Kotchen et al., 2016; Padmanabhan and Joe, 2017; Padmanabhan and Dominiczak, 2021). The data we analyse include patients who met the following criteria: (i) admitted to the ICU for at least 3 days, (ii) had a first-day maximum SBP exceeding 180 mmHg, and (iii) received treatment combinations that were assigned to at least 1% of the population in both stages. As a result, stage 1 involved 742 patients admitted to the ICU due to a hypertensive crisis. After the first treatment, 320 of these patients advanced to stage 2, as they continued to experience hypertension, necessitating further medication on the second day after admission. Four classes of antihypertensive agents are commonly used to treat hypertension: angiotensin-converting enzyme inhibitors (ACEi), beta-blockers, calcium channel blockers (CCB), and diuretics (Laurent, 2017). Research has shown that combination therapies, guided by clinical practice and scientific evidence, often lead to better health outcomes compared to single-drug treatments (Wald et al., 2009; Gradman, 2010; Kotchen et al., 2016). Hence, our goal extends beyond identifying the optimal single antihypertensive agent: we aim to determine the optimal combination of antihypertensive agents by estimating a two-stage dynamic DTR based on patients' SBP management.

Out of all possible combinations of ACE inhibitors, beta-blockers, CCBs, and diuretics, the dataset we analysed includes only $T = 8$ treatment options: the four antihypertensive agents administered individually and four combinations of these agents. The frequencies of these treatments are detailed in Table 2.7 in the Appendix. Additionally, a total of $k = 12$ clinical factors are taken into consideration in the two stages. Namely, daily max SBP, heart rate (HR), oxygen saturation (SpO2), age, weight, gender, race, current or

former smoker, kidney disease, diabetes, chronic obstructive pulmonary disease (COPD), and chronic hypertension (CH). Notably, this list includes potential confounders, such as SBP and chronic hypertension, in line with standard confounder adjustment techniques (see Chakraborty and Moodie, 2013). The main goal of our DTR approach is to recommend the optimal combination of antihypertensive agents for patients with severe acute arterial hypertension at stage 1, which is immediately following ICU admission, based on their clinical characteristics at that time. Furthermore, for patients whose maximum SBP exceeds 140 mmHg the following day, our analysis seeks to recommend additional adjustments to the hypertension treatment to be administered at stage 2, taking into consideration the patient's clinical factors on that day as well as the therapy given at admission.

In our analysis, we use the decrease in maximum SBP after treatment as the response variable. To accommodate varying numbers of patients across the two stages, we extend the BAL model by introducing a dummy variable, $\mathcal{I}_i$, which is set to 1 if the ith patient received treatment on both days, and 0 if treatment was administered only on the first day. As a result, the stage 2 augmented model is defined exclusively for patients for which $\mathcal{I}_i = 1$, and is given by

$$\begin{aligned}(Y_{i2}, V_{i2}^{(0:7)}) \mid a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2 \overset{\text{ind}}{\sim} & f_{\mathrm{N}}(y; \boldsymbol{\theta}_2^{\intercal}[\boldsymbol{z}_{i2} \times (\boldsymbol{a}_{i1}^*, \boldsymbol{a}_{i2}^*)], \sigma_2^2) \\ & \times \prod_{t=0}^{7} f_{\mathrm{N}}(v_t; \boldsymbol{\theta}_2^{\intercal}[\boldsymbol{z}_{i2} \times (\boldsymbol{a}_{i1}^*, \boldsymbol{i}_t)], \sigma_2^2),\end{aligned} \tag{2.12}$$

where $\boldsymbol{a}_{ij}^* = (a_{ij1}^*, \ldots, a_{ij7}^*) \in \mathcal{D}_7$ are vectors of dummy variables such that $a_{ijt}^* = \mathbb{I}_{\{t\}}(a_{ij})$, and $\boldsymbol{i}_t \in \mathcal{D}_7$ is the null vector if $t = 0$, and is defined as the vector of zeros with a 1 in position t, if $t \in \{1, \ldots, 7\}$. Without loss of generality, we set $V_{i2}^{(t)} = 0$, for every $t \in \{0, \ldots, 7\}$, when $\mathcal{I}_i = 0$. The stage 1 augmented model is given by

$$\begin{aligned}(Y_{i1}, V_i^{(0:7)}) \mid V_{i2}^{(0:7)}, a_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1 \overset{\text{ind}}{\sim} & f_{\mathrm{N}}(y + \max_t\{V_{i2}^{(t)}\}; \theta_0 \mathcal{I}_i + \boldsymbol{\theta}_1^{\intercal}[\boldsymbol{z}_{i1} \times \boldsymbol{a}_{i1}^*], \sigma_1^2) \\ & \times \prod_{t=0}^{7} f_{\mathrm{N}}(v_t; \theta_0 \mathcal{I}_i + \boldsymbol{\theta}_1^{\intercal}[\boldsymbol{z}_{i1} \times \boldsymbol{i}_t], \sigma_1^2).\end{aligned} \tag{2.13}$$

Finally, we assume a DSS prior for $(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ and an independent normal prior for θ_0.

The model defined through (2.12) and (2.13) includes a total of $p_2 = (2T-1)(k+1) = 195$ regressors in stage 2 and $p_1 = T(k+1)+1 = 105$ in stage 1, with identical prognostic factors used in both stages. We fitted the same three models considered in the simulation studies of Section 2.4 and concluded that the DSS outperforms its two competitors. This

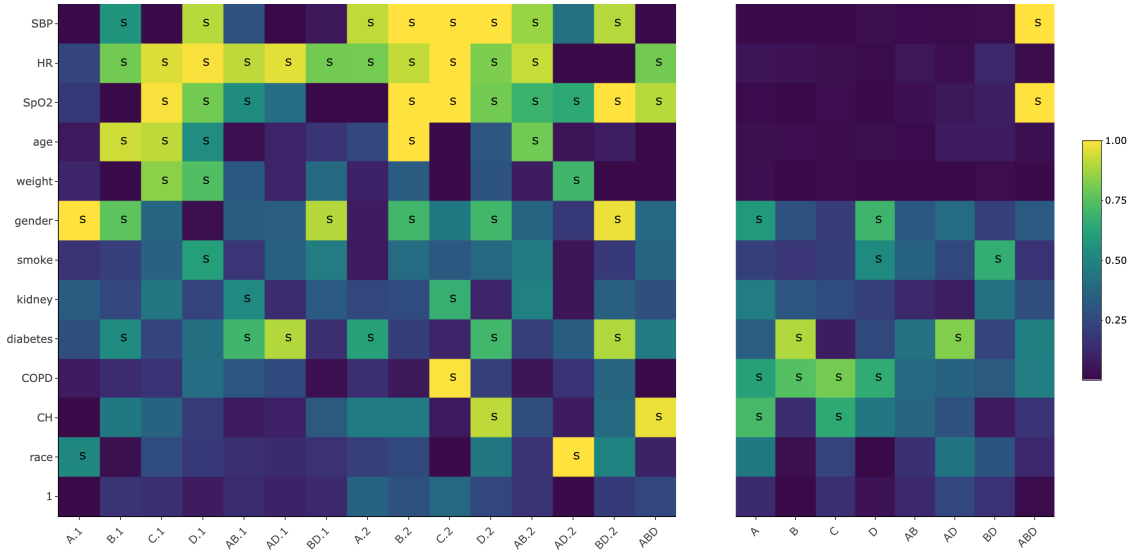

Figure 2.1: *Posterior inclusion probabilities at stage 2 (left panel) and stage 1 (right panel). Treatment options are shown on the x-axis, while clinical factors are listed on the y-axis. Specifically, 'A' represents ACEi, 'B' refers to beta-blockers, 'C' indicates CCB, and 'D' stands for diuretics. 'AB' denotes the combination of ACEi and beta-blockers, with other combinations defined similarly. The notation '1' refers to first-stage treatments, and '2' pertains to second-stage agents. For instance, if the x-axis is labeled A.1 and the y-axis is SBP, the corresponding value represents the posterior inclusion probability of the interaction between first-stage ACEi and SBP. The symbol 'S', corresponding to posterior inclusion probabilities larger than 0.5, indicates statistical significance.*

was determined by computing the Bayesian Information Criterion (BIC) and log pseudo-marginal likelihood (LPML) (see Tables 2.8 in the Appendix). These results are consistent with the findings of the simulation studies, where the DSS performed as well as or better than its competitors across all scenarios. Therefore, in the remainder of this section, we present the results of the analysis of the MIMIC-III data based on the BAL model with normal regression functions (2.5) and the DSS prior (2.7). Estimation was performed using Algorithm 1 with 10,000 iterations, discarding the first 5,000 as burn-in.

Studying interaction effects in this clinical study is of utmost importance as they shed light on how the antihypertensive agents, or their combinations, and specific clinical factors interact to influence the patient's SBP. The posterior inclusion probability of the 12 clinical factors and the interaction between the eight treatment options and these factors are illustrated in Figure 2.1. In stage 1, a limited number of variables demonstrate statistical significance, SBP and SpO2 showing notable effects, particularly in treatments involving the ABD combination. Clinical factors such as HR, age, weight, kidney function, and race

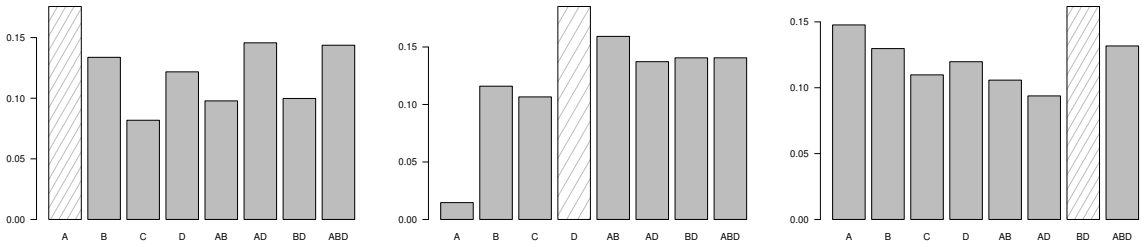

Figure 2.2: *MIMIC-III data. Posterior probability of each of the $T = 8$ treatment options being the optimal choice for Individual #1 at stage 1 (left panel), for Individual #2 at stage 2 (middle panel), and for Individual #2 at stage 1 (right panel). In each barplot, the white bin indicates the largest posterior probability.*

do not reach statistical significance at this stage. In stage 2, SBP, SpO2, gender, smoking status, diabetes, COPD, and chronic hypertension remain significant across various drug combinations, highlighting their continued relevance. Furthermore, a broader range of clinical factors, including HR, age, weight, kidney function, and race, show higher posterior inclusion probabilities when interacting with treatments. This suggests that as treatment progresses, these factors become increasingly important for optimizing therapeutic strategies.

To further demonstrate the utility of our analysis, we complete the study by considering two patients, randomly selected from the sample, henceforth referred to as Individuals #1 and #2. Figure 2.2 shows, for these two individuals, the stage-specific probabilities that each treatment option is the optimal choice, based on their individual clinical factors and prior treatments. For Individual #1, who received only one day of treatment, our DTR analysis recommends the ACEi antihypertensive agent as the most effective treatment option. For Individual #2, who received two days of treatment in the ICU, our DTR analysis suggests that the most effective combination in stage 1 is beta-blockers + diuretics, while in stage 2, diuretics are recommended. Additionally, the posterior sample produced by Algorithm 1 can be used to compute the marginal probabilities of individual antihypertensive agents being part of the optimal treatment combination. This is illustrated for the same two individuals in Figure 2.3 in the Appendix.

2.6 Conclusion

We introduced BAL, a novel Bayesian approach to DTRs that accounts for uncertainty in identifying the optimal sequence of treatments. The BAL is defined as an augmented

Bayesian model, incorporating pseudo-outcomes to seamlessly handle the counterfactual variables characterizing DTR problems. Variable selection is integrated into the BAL approach through the introduction of the DSS prior, which conveniently favors borrowing information across multiple stages of variable selection and facilitates the identification of the optimal treatment sequence. To assess the performance of BAL, we conducted extensive simulation studies and compared the results to alternative methods for variable selection in DTRs, a relatively underexplored area in the literature. The simulation studies demonstrate that BAL outperforms other approaches in identifying the optimal treatment sequence and selecting significant variables. Additionally, we applied our method to clinical trial data from patients with severe acute arterial hypertension who were admitted to the ICU and received specific hospital care medications.

The promising results presented in this work suggest various directions for expanding our current model. One direction involves refining the dependencies among the indicator variables that encode the inclusion of regressors, tailored to the specific problem at hand. For instance, it would be valuable to investigate the intra-stage dependence and treat related covariates in a given stage equally. A second direction could involve defining sequences of more flexible regression models by considering nonparametric regressions. In this context, the DSS prior could be redefined within the framework of inner spike-and-slab Bayesian nonparametric models (Canale et al., 2017, 2023). Lastly, an intriguing research direction involves extending the described approach to settings with more than two stages. In particular, as soon as the number of stages is not very small, this task may necessitate the development of new computational strategies to manage the rapidly increasing number of counterfactual treatment sequences.

Data Availability Statement

The data that support the findings of this study are available from PhysioNet. Restrictions exist regarding the accessibility of this data, which were utilized in this study under a licensing agreement. Data are available at https://physionet.org/content/mimiciii/1.4/ with the permission of PhysioNet.

Table 2.2: *First experiment: simulated data with n smaller than p_2. For stage 2, stage 1 and overall, two measures of prediction accuracy (MRE and ER), for DSS, ISS and QLL, across the nine scenarios obtained by specifying $(k, n) \in \{(10, 25), (20, 50), (30, 75)\}$ and $\rho^* \in \{0.3, 0.6, 0.9\}$. Results are based on 100 replicated datasets.*

Stage 2								
n	p_2	ρ^*	MRE			ER		
			DSS	ISS	QLL	DSS	ISS	QLL
25	33	0.3	0.003	0.143	0.172	0.040	0.064	0.123
		0.6	0.001	0.121	0.182	0.007	0.147	0.073
		0.9	0.002	0.108	0.169	0.003	0.060	0.127
50	63	0.3	0.006	0.093	0.120	0.038	0.116	0.095
		0.6	0.011	0.167	0.130	0.028	0.080	0.093
		0.9	0.024	0.117	0.143	0.073	0.090	0.110
75	93	0.3	0.004	0.051	0.121	0.039	0.110	0.064
		0.6	0.008	0.003	0.123	0.027	0.041	0.079
		0.9	0.005	0.074	0.098	0.023	0.099	0.087
Stage 1								
n	p_1	ρ^*	MRE			ER		
			DSS	ISS	QLL	DSS	ISS	QLL
25	22	0.3	0.064	0.045	0.197	0.153	0.224	0.228
		0.6	0.075	0.069	0.198	0.104	0.187	0.213
		0.9	0.032	0.063	0.142	0.100	0.210	0.137
50	42	0.3	0.007	0.059	0.117	0.060	0.120	0.140
		0.6	0.012	0.061	0.114	0.048	0.120	0.101
		0.9	0.002	0.063	0.126	0.040	0.160	0.129
75	62	0.3	0.011	0.059	0.100	0.041	0.110	0.093
		0.6	0.001	0.027	0.102	0.040	0.103	0.101
		0.9	0.007	0.065	0.103	0.019	0.104	0.101
Overall								
n	(p_1, p_2)	ρ^*	MRE			ER		
			DSS	ISS	QLL	DSS	ISS	QLL
25	(22, 33)	0.3	0.034	0.094	0.184	0.187	0.256	0.303
		0.6	0.038	0.095	0.190	0.111	0.302	0.280
		0.9	0.017	0.085	0.156	0.100	0.260	0.257
50	(42, 63)	0.3	0.007	0.076	0.118	0.096	0.216	0.223
		0.6	0.012	0.114	0.122	0.076	0.167	0.182
		0.9	0.013	0.040	0.134	0.105	0.200	0.229
75	(62, 93)	0.3	0.007	0.055	0.110	0.077	0.197	0.152
		0.6	0.005	0.015	0.112	0.063	0.120	0.172
		0.9	0.006	0.070	0.100	0.041	0.176	0.180

Table 2.3: *Second experiment. For stage 2 and stage 1, three summaries of the accuracy in selecting significant variables (proportion of FN, proportion of FP, and F_1 score), for DSS, ISS, and QLL, across the four scenarios obtained by specifying $n \in \{200, 400)\}$ and $T \in \{4, 8\}$. The number of individual covariates is $k = 10$, while $\rho^* = 0.6$. Results are based on 100 replicated datasets.*

Stage 2											
T	p_2	n	FN			FP			F_1 score		
			DSS	ISS	QLL	DSS	ISS	QLL	DSS	ISS	QLL
4	99	200	0.175	0.173	0.167	0.014	0.056	0.188	0.884	0.838	0.735
		400	0.139	0.209	0.201	0.013	0.036	0.126	0.909	0.846	0.781
8	187	200	0.228	0.263	0.255	0.059	0.076	0.172	0.817	0.777	0.688
		400	0.134	0.173	0.174	0.020	0.055	0.119	0.905	0.851	0.791
Stage 1											
T	p_1	n	FN			FP			F_1 score		
			DSS	ISS	QLL	DSS	ISS	QLL	DSS	ISS	QLL
4	55	200	0.159	0.270	0.275	0.022	0.040	0.172	0.884	0.810	0.698
		400	0.155	0.288	0.302	0.009	0.013	0.119	0.901	0.813	0.707
8	99	200	0.224	0.276	0.475	0.087	0.187	0.106	0.793	0.727	0.569
		400	0.165	0.267	0.361	0.032	0.014	0.083	0.873	0.829	0.689

Table 2.4: *Second experiment. For stage 2, stage 1 and overall, two measures of prediction accuracy (MRE and ER), for DSS, ISS and QLL, across the four scenarios obtained by specifying* $n \in \{200, 400)\}$ *and* $T \in \{4, 8\}$*. The number of individual covariates is* $k = 10$*, while* $\rho^* = 0.6$*. Results are based on 100 replicated datasets.*

Stage 2								
T	p_2	n	MRE			ER		
			DSS	ISS	QLL	DSS	ISS	QLL
4	99	200	1.013	1.289	1.334	0.031	0.080	0.214
		400	0.785	1.132	1.177	0.031	0.082	0.173
8	187	200	1.072	1.329	2.272	0.158	0.179	0.372
		400	0.953	1.106	1.371	0.071	0.175	0.302
Stage 1								
T	p_1	n	MRE			ER		
			DSS	ISS	QLL	DSS	ISS	QLL
4	55	200	1.093	1.575	1.707	0.132	0.182	0.248
		400	0.813	1.285	1.504	0.053	0.102	0.258
8	99	200	1.178	1.942	2.623	0.208	0.272	0.456
		400	1.090	1.491	1.663	0.084	0.242	0.413
Overall								
T	(p_1, p_2)	n	MRE			ER		
			DSS	ISS	QLL	DSS	ISS	QLL
4	(55,99)	200	1.053	1.432	1.521	0.158	0.258	0.411
		400	0.799	1.209	1.341	0.081	0.180	0.384
8	(99, 187)	200	1.125	1.636	2.448	0.338	0.412	0.665
		400	1.022	1.299	1.517	0.150	0.393	0.599

Bibliography

Albert, J. H. and Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422):669–679.

Bellman, R. (1957). Dynamic programming, princeton univ. *Press Princeton, New Jersey*.

Bian, Z., Moodie, E. E., Shortreed, S. M., and Bhatnagar, S. (2021). Variable selection in regression-based estimation of dynamic treatment regimes. *Biometrics*.

Canale, A., Lijoi, A., Nipoti, B., and Prünster, I. (2017). On the pitman–yor process with spike and slab base measure. *Biometrika*, 104(3):681–697.

Canale, A., Lijoi, A., Nipoti, B., and Prünster, I. (2023). Inner spike and slab Bayesian nonparametric models. *Econometrics and Statistics*, 27:120–135.

Chakraborty, B., Laber, E. B., and Zhao, Y. (2013). Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme. *Biometrics*, 69(3):714–723.

Chakraborty, B. and Moodie, E. E. (2013). Statistical methods for dynamic treatment regimes. *Springer-Verlag. doi*, 10:978–1.

Chakraborty, B., Murphy, S., and Strecher, V. (2010). Inference for non-regular parameters in optimal dynamic treatment regimes. *Statistical Methods in Medical Research*, 19(3):317–343.

Chen, Y., Liu, Y., Zeng, D., and Wang, Y. (2020). Statistical learning methods for optimizing dynamic treatment regimes in subgroup identification. *Design and Analysis of Subgroups with Biopharmaceutical Applications*, pages 271–297.

Chib, S. and Greenberg, E. (1995). Understanding the metropolis-hastings algorithm. *The American Statistician*, 49(4):327–335.

Fan, A., Lu, W., and Song, R. (2016). Sequential advantage selection for optimal treatment regime. *The Annals of Applied Statistics*, 10(1):32.

Friedman, J., Hastie, T., Tibshirani, R., Narasimhan, B., Tay, K., Simon, N., and Qian, J. (2021). Package 'glmnet'. *CRAN R Repositary*, 595.

Gelman, A., Gilks, W. R., and Roberts, G. O. (1997). Weak convergence and optimal scaling of random walk metropolis algorithms. *The Annals of Applied Probability*, 7(1):110–120.

Gradman, A. H. (2010). Rationale for triple-combination therapy for management of high blood pressure. *The Journal of Clinical Hypertension*, 12(11):869–878.

Ishwaran, H. and Rao, J. S. (2003). Detecting differentially expressed genes in microarrays using Bayesian model selection. *Journal of the American Statistical Association*, 98(462):438–455.

Ishwaran, H. and Rao, J. S. (2005). Spike and slab variable selection: frequentist and Bayesian strategies. *The Annals of Statistics*, 33:730–773.

Johnson, A. E., Pollard, T. J., Shen, L., Lehman, L.-w. H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Anthony Celi, L., and Mark, R. G. (2016). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3(1):1–9.

Kotchen, T. A., Cowley, A. W., and Liang, M. (2016). Ushering hypertension into a new era of precision medicine. *JAMA*, 315(4):343–344.

Laber, E. B., Lizotte, D. J., Qian, M., Pelham, W. E., and Murphy, S. A. (2014). Dynamic treatment regimes: Technical challenges and applications. *Electronic Journal of Statistics*, 8(1):1225.

Laurent, S. (2017). Antihypertensive drugs. *Pharmacological research*, 124:116–125.

Lu, W., Zhang, H. H., and Zeng, D. (2013). Variable selection for optimal treatment decision. *Statistical Methods in Medical Research*, 22(5):493–504.

Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355.

Murphy, S. A. (2005). An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*, 24(10):1455–1481.

Murray, T. A., Yuan, Y., and Thall, P. F. (2018). A Bayesian machine learning approach for optimizing dynamic treatment regimes. *Journal of the American Statistical Association*, 113(523):1255–1267.

Padmanabhan, S. and Dominiczak, A. F. (2021). Genomics of hypertension: the road to precision medicine. *Nature Reviews Cardiology*, 18(4):235–250.

Padmanabhan, S. and Joe, B. (2017). Towards precision medicine for hypertension: a review of genomic, epigenomic, and microbiomic effects on blood pressure in experimental rat models and humans. *Physiological Reviews*, 97(4):1469–1528.

Qian, M. and Murphy, S. A. (2011). Performance guarantees for individualized treatment rules. *Annals of Statistics*, 39(2):1180.

Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, pages 34–58.

Rubin, D. B. (1980). Bias reduction using Mahalanobis-metric matching. *Biometrics*, pages 293–298.

Song, R., Wang, W., Zeng, D., and Kosorok, M. R. (2015). Penalized Q-learning for dynamic treatment regimens. *Statistica Sinica*, 25(3):901.

Tadesse, M. G. and Vannucci, M. (2021). Handbook of Bayesian variable selection.

Tsiatis, A. A., Davidian, M., Holloway, S. T., and Laber, E. B. (2019). *Dynamic treatment regimes: Statistical methods for precision medicine*. Chapman and Hall/CRC.

Wald, D. S., Law, M., Morris, J. K., Bestwick, J. P., and Wald, N. J. (2009). Combination therapy versus monotherapy in reducing blood pressure: meta-analysis on 11,000 participants from 42 trials. *The American Journal of Medicine*, 122(3):290–300.

Wallace, M. P., Moodie, E. E., and Stephens, D. A. (2019). Model selection for G-estimation of dynamic treatment regimes. *Biometrics*, 75(4):1205–1215.

Watkins, C. J. C. H. (1989). Learning from delayed rewards.

Zhao, Y., Kosorok, M. R., and Zeng, D. (2009). Reinforcement learning design for cancer clinical trials. *Statistics in Medicine*, 28(26):3294–3315.

Zhou, N., Brook, R. D., Dinov, I. D., and Wang, L. (2022). Optimal dynamic treatment regime estimation using information extraction from unstructured clinical text. *Biometrical Journal*, 64(4):805–817.

Zhu, W., Zeng, D., and Song, R. (2019). Proper inference for value function in high-dimensional Q-learning for dynamic treatment regimes. *Journal of the American Statistical Association*, 114(527):1404–1417.

Appendices

These appendices are organized as follows. In Section 2.A, we discuss the validity of the equivalences presented in (2.3) and provide proofs for specific cases. Section 2.B outlines the derivation of the expression in (2.8), which pertains to the correlation between the indicator variables that encode the significance of the same regressor across the two stages. Section 2.C offers additional details regarding posterior computation. Finally, Sections 2.D and 2.E present supplementary results on the simulation study and the analysis of the MIMIC-III data, respectively.

2.A On the validity of the equivalences in (2.3)

We discuss and prove the validity of the equivalences in (2.3) for some cases of interest. As in Section 3.3.1, Y_{i2} and $\bar{Y}_{i2}$ are assumed independent and distributed according to f_2, as specified in (2.1). Similarly, Y_i^{OPT} and $\bar{Y}_i^{\text{OPT}}$ are assumed independent and distributed according to f_1. We start by proving Proposition 1, which refers to the normal specification of f_2 and f_1, described in Section 2.2.1 and implemented in the chapter. Propositions 2 and 3 showcase that the same approach used in this work can be used for other specifications, such as logistic and Poisson regressions.

Proof of Proposition 1. Given the continuity of the payoffs Y_{i2} and $\bar{Y}_{i2}$, proving (2.3) is equivalent to proving (2.6). Combining the assumptions on Y_{i2} and $\bar{Y}_{i2}$, we obtain that $Y_{i2} - \bar{Y}_{i2} \sim \text{N}(\mu_1 - \mu_2, 2\sigma_2^2)$, where $\mu_1 = \boldsymbol{\theta}_2^\intercal[\boldsymbol{z}_{i2} \times (a_{i1}, a_{i2})]$ and $\mu_2 = \boldsymbol{\theta}_2^\intercal[\boldsymbol{z}_{i2} \times (a_{i1}, 1 - a_{i2})]$. Thus, we can write

$$
\begin{aligned}
P(Y_{i2} - \bar{Y}_{i2} \geq 0) > 0.5 &\iff \Phi\left(\frac{\mu_1 - \mu_2}{\sqrt{2\sigma_2^2}}\right) > 0.5 \\
&\iff \mu_1 > \mu_2 \\
&\iff \mathbb{E}[Y_{i2} \mid a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2] > \mathbb{E}[Y_{i2} \mid a_{i1}, 1 - a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2] \\
&\overset{\text{(a)}}{\iff} a_{i2}^{\text{OPT}} = a_{i2},
\end{aligned}
$$

where (a) follows from (2.2).
Similarly for the equivalence involving Y_i^{OPT} and $\bar{Y}_i^{\text{OPT}}$. □

Proposition 2. If f_2 and f_1 are logistic regressions, specified as

$$\begin{aligned} f_2(y; a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2) &= f_{\text{BER}}(y; 1/(1+\exp\{-\boldsymbol{\theta}_2^\intercal[\boldsymbol{z}_{i2} \times (a_{i1}, a_{i2})]\})), \\ f_1(y; a_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1) &= f_{\text{BER}}(y; 1/(1+\exp\{-\boldsymbol{\theta}_1^\intercal[\boldsymbol{z}_{i1} \times a_{i1}]\})), \end{aligned} \tag{14}$$

where $f_{\text{BER}}(y; \pi)$ denotes the probability mass function of a Bernoulli random variable with mean π, then the equivalences in (2.3) hold.

Proof. We set $\pi_1 = 1/(1+\exp\{-\mu_1\})$ and $\pi_2 = 1/(1+\exp\{-\mu_2\})$, with $\mu_1 = \boldsymbol{\theta}_2^\intercal[\boldsymbol{z}_{i2} \times (a_{i1}, a_{i2})]$ and $\mu_2 = \boldsymbol{\theta}_2^\intercal[\boldsymbol{z}_{i2} \times (a_{i1}, 1-a_{i2})]$. Combining the assumptions on Y_{i2} and $\bar{Y}_{i2}$, we find that $Y_{i2} - \bar{Y}_{i2}$ can take values $\{-1, 0, 1\}$, with probabilities respectively equal to $\{(1-\pi_1)\pi_2, \pi_1\pi_2 + (1-\pi_1)(1-\pi_2), \pi_1(1-\pi_2)\}$. Thus we have

$$\begin{aligned} \frac{P(Y_{i2} \geq \bar{Y}_{i2})}{P(Y_{i2} \leq \bar{Y}_{i2})} > 1 &\iff \frac{p_1(1-p_2) + p_1p_2 + (1-p_1)(1-p_2)}{(1-p_1)p_2 + p_1p_2 + (1-p_1)(1-p_2)} > 1 \\ &\iff \frac{1 + \exp\{\mu_1\} + \exp\{\mu_1+\mu_2\}}{1 + \exp\{\mu_2\} + \exp\{\mu_1+\mu_2\}} > 1 \\ &\iff \mu_1 > \mu_2 \\ &\iff \mathbb{E}[Y_{i2} \mid a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2] > \mathbb{E}[Y_{i2} \mid a_{i1}, 1-a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2] \\ &\overset{\text{(a)}}{\iff} a_{i2}^{\text{OPT}} = a_{i2}, \end{aligned}$$

where (a) follows from (2.2).
Similarly for the equivalence involving Y_i^{OPT} and $\bar{Y}_i^{\text{OPT}}$. □

Proposition 3. If f_2 and f_1 are Poisson regressions, specified as

$$\begin{aligned} f_2(y; a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2) &= f_{\text{POI}}(y; \boldsymbol{\theta}_2^\intercal[\boldsymbol{z}_{i2} \times (a_{i1}, a_{i2})]), \\ f_1(y; a_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1) &= f_{\text{POI}}(y; \boldsymbol{\theta}_1^\intercal[\boldsymbol{z}_{i1} \times a_{i1}]), \end{aligned} \tag{15}$$

where $f_{\text{POI}}(y; \lambda)$ denotes the probability mass function of a Poisson random variable with mean λ, then the equivalences in (2.3) hold.

Proof. Combining the assumptions on Y_{i2} and $\bar{Y}_{i2}$, we obtain that $Y_{i2} - \bar{Y}_{i2} \sim \text{SKELLAM}(\mu_1, \mu_2)$, where $\mu_1 = \boldsymbol{\theta}_2^\intercal[\boldsymbol{z}_{i2} \times (a_{i1}, a_{i2})]$ and $\mu_2 = \boldsymbol{\theta}_2^\intercal[\boldsymbol{z}_{i2} \times (a_{i1}, 1-a_{i2})]$. Thus, if we let $f_{\text{SKE}}(y; \mu_1, \mu_2)$ denote the probability mass function of a Skellam random variable with parameters μ_1 and

μ_2, we can write

$$\begin{aligned}
\frac{P(Y_{i2} \geq \bar{Y}_{i2})}{P(Y_{i2} \leq \bar{Y}_{i2})} > 1 &\iff \frac{\sum_{\ell=0}^{\infty} f_{\text{SKE}}(\ell; \mu_1, \mu_2)}{\sum_{\ell=-\infty}^{0} f_{\text{SKE}}(\ell; \mu_2, \mu_1)} > 1. \\
&\iff \frac{\sum_{\ell=0}^{\infty} f_{\text{SKE}}(\ell; \mu_1, \mu_2)}{\sum_{\ell=0}^{\infty} f_{\text{SKE}}(-\ell; \mu_2, \mu_1)} > 1 \\
&\overset{\text{(a)}}{\iff} \mu_1 > \mu_2 \\
&\iff \mathbb{E}[Y_{i2} \mid a_{i1}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2] > \mathbb{E}[Y_{i2} \mid a_{i1}, 1 - a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2] \\
&\overset{\text{(b)}}{\iff} a_{i2}^{\text{OPT}} = a_{i2},
\end{aligned}$$

where (a) is implied by the fact that, for any $\ell > 0$,

$$\frac{f_{\text{SKE}}(\ell; \mu_1, \mu_2)}{f_{\text{SKE}}(-\ell; \mu_2, \mu_2)} = \left(\frac{\mu_1}{\mu_2}\right)^{\ell},$$

and (b) follows from (2.2).
Similarly for the equivalence involving Y_i^{OPT} and $\bar{Y}_i^{\text{OPT}}$. □

2.B Derivation of the expression in (2.8)

We show that $\rho(\delta_{l_1 1}, \delta_{l_2 2}) = 1/(1 + a + b)$ if $l_1 = l_2 \in \{1, \ldots, p\}$, and 0 otherwise. $\delta_{l_1 1}$ and $\delta_{l_1 2}$ are Bernoulli random variables with parameters $w_{l_1 1}$ and $w_{l_2 2}$ respectively, and $w_{l_1 1} = w_{l_2 2} = w \sim \text{BETA}(a, b)$ for $l_1 = l_2 \in \{1, \ldots, p\}$.

Since $\delta_{l_1 1}$ and $\delta_{l_2 2}$ are conditionally independent given w, we can express their joint distribution as:

$$P(\delta_{l_1 1} = \alpha, \delta_{l_2 2} = \beta \mid w) = w^{\alpha+\beta}(1 - w)^{2-(\alpha+\beta)},$$

where $\alpha, \beta \in \{0, 1\}$. To find the unconditional joint probability, we marginalize with respect to w and obtain

$$P(\delta_{l_1 1} = \alpha, \delta_{l_2 2} = \beta) = \int_0^1 P(\delta_{l_1 1} = \alpha, \delta_{l_2 2} = \beta \mid w) f(w) \mathrm{d}w,$$

where $f(w)$ is the PDF of the Beta(a, b) distribution, that is

$$f(w) = \frac{w^{a-1}(1 - w)^{b-1}}{B(a, b)},$$

and $B(a, b)$ is the Beta function. This leads to:

$$P(\delta_{l_1 1} = \alpha, \delta_{l_2 2} = \beta) = \frac{B(\alpha + \beta + a, 2 - (\alpha + \beta) + b)}{B(a, b)}$$

Marginally, we have

$$\mathbb{E}[\delta_{l_1,1}] = P(\delta_{l_1 1} = 1) = \frac{B(a + 1, b)}{B(a, b)} = \frac{a}{a + b},$$

and similarly for $\mathbb{E}[\delta_{l_2,2}]$. From this we get

$$\mathbb{E}[\delta_{l_1 1}\delta_{l_2 2}] = P(\delta_{l_1 1} = 1, \delta_{l_2 2} = 1) = \frac{B(a + 2, b)}{B(a, b)} = \frac{a(a + 1)}{(a + b)(a + b + 1)}.$$

We thus can write the covariance between $\delta_{l_1 1}$ and $\delta_{l_2 2}$ as

$$\mathrm{Cov}(\delta_{l_1 1}, \delta_{l_2 2}) = \mathbb{E}[\delta_{l_1 1}\delta_{l_2 2}] - \mathbb{E}[\delta_{l_1 1}]\mathbb{E}[\delta_{l_2 2}] = \frac{ab}{(a + b)^2(a + b + 1)}.$$

The variance of $\delta_{l_1 1}$ is given by

$$\mathrm{Var}(\delta_{l_1 1}) = \mathbb{E}[\delta_{l_1 1}^2] - (\mathbb{E}[\delta_{l_1 1}])^2 = \frac{a}{a + b}\left(1 - \frac{a}{a + b}\right) = \frac{ab}{(a + b)^2},$$

which coincides with the variance for $\delta_{l_2 2}$.
As a result, we can write the correlation as

$$\rho(\delta_{l_1 1}, \delta_{l_2 2}) = \frac{\frac{ab}{(a+b)^2(a+b+1)}}{\sqrt{\frac{ab}{(a+b)^2} \cdot \frac{ab}{(a+b)^2}}} = \frac{1}{1 + a + b}.$$

Given that $a > 0$ and $b > 0$, this result implies that $\rho(\delta_{l_1 1}, \delta_{l_2 2})$ is always positive. However, we note that in certain applied contexts, it may be convenient to consider a more flexible model able to induce a negative correlation between the inclusion variables $\delta_{l_1 1}$ and $\delta_{l_2 2}$.

2.C Posterior inference

We next describe the full conditional distributions involved in the definition of Algorithm 1.

Full conditional of $\boldsymbol{\theta}_1$*:*

$$\boldsymbol{\theta}_1 \mid \ldots \sim \mathrm{N}_{p_1}\left(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1\right),$$

where

$$\boldsymbol{\Sigma}_1^{-1} = \frac{\boldsymbol{X}_1^T \boldsymbol{X}_1}{\sigma_1^2} + D_1^{-1}$$

and

$$\boldsymbol{\mu}_1 = \frac{\boldsymbol{\Sigma}_1 \boldsymbol{X}_1^T \boldsymbol{Y}^{\mathrm{OPT}}}{\sigma_1^2}.$$

D_1 is a diagonal matrix with entries $r^{\left(1-\delta_{l_1 1}\right)}\psi_{l_1 1},\ l_1 = 1, \ldots, p_1$, and p_1 is the dimension of $\boldsymbol{\theta}_1$,

$$\boldsymbol{X}_1 = \begin{pmatrix} \boldsymbol{z}_{11} \times a_{11} \\ \ldots \\ \boldsymbol{z}_{n1} \times a_{n1} \end{pmatrix}, \quad \boldsymbol{Y}^{\mathrm{OPT}} = \begin{pmatrix} Y_{11} + \max(V_{12}^{(0)}, V_{12}^{(1)}) \\ \ldots \\ Y_{11} + \max(V_{n2}^{(0)}, V_{n2}^{(1)}) \end{pmatrix}.$$

Full conditional of $\boldsymbol{\theta}_2$*:*

$$\boldsymbol{\theta}_2 \mid \ldots \sim \mathrm{N}_{p_2}\left(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2\right),$$

where

$$\boldsymbol{\Sigma}_2^{-1} = \frac{\boldsymbol{X}^{*T} \boldsymbol{X}^*}{\sigma_2^2} + D_2^{-1}$$

and

$$\boldsymbol{\mu}_2 = \frac{\boldsymbol{\Sigma}_2 \boldsymbol{X}^{*T} \boldsymbol{Y}^*}{\sigma_2^2}.$$

D_2 is a diagonal matrix with entries $r^{\left(1-\delta_{l_2 2}\right)}\psi_{l_2 2},\ l_2 = 1, \ldots, p_2$, and p_2 is the dimension of $\boldsymbol{\theta}_2$. Moreover,

$$\boldsymbol{X}^* = \begin{pmatrix} \boldsymbol{z}_{12} \times (a_{11}, a_{12}) \\ \ldots \\ \boldsymbol{z}_{n2} \times (a_{n1}, a_{n2}) \\ \boldsymbol{z}_{12} \times (a_{11}, 1 - a_{12}) \\ \ldots \\ \boldsymbol{z}_{n2} \times (a_{n1}, 1 - a_{n2}) \end{pmatrix}, \quad \boldsymbol{Y}^* = \begin{pmatrix} Y_{12} \\ \ldots \\ Y_{n2} \\ V_{12}^{(1-a_{12})} \\ \ldots \\ V_{n2}^{(1-a_{n2})} \end{pmatrix}$$

Full conditional of σ_1^2*:*

$$\sigma_1^2 \mid \ldots \sim \text{INV-GAMMA}\left(\frac{n+1}{2}, \beta + \frac{1}{2}\right),$$

where $\beta = \frac{1}{2}\sum_{i=1}^{n}\left(Y_{i1} - \left(\boldsymbol{\theta}_1^T[\boldsymbol{z}_{i1} \times a_{i1}] - \max\left(V_{i2}^{(0)}, V_{i2}^{(1)}\right)\right)\right)^2$ is the scale parameter of the inverse gamma distribution.

Full conditional of σ_2^2:

$$\sigma_2^2 \mid \ldots \sim \text{inv-Gamma}\left(n + \frac{1}{2}, \beta + \frac{1}{2}\right),$$

where

$$\beta = \frac{1}{2}\sum_{i=1}^{n}\left(\left(Y_{i2} - \boldsymbol{\theta}_2^T[\boldsymbol{z}_{i2} \times (a_{i1}, a_{i2})]\right)^2 + \left(V_{i2}^{(1-a_{i2})} - \boldsymbol{\theta}_2^T[\boldsymbol{z}_{i2} \times (a_{i1}, 1-a_{i2})]\right)^2\right)$$

is the scale parameter of the inverse gamma distribution.

Full conditionals for $\delta_{l_1 1}$ and $\delta_{l_2 2}$: for any $l_1 = 1, \ldots, p_1$

$$\delta_{l_1 1} \mid \ldots \sim \text{Bern}\left(\frac{w_{l_1 1} f_{\text{N}}(\theta_{l_1 1}; 0, \psi_{l_1 1})}{w_{l_1 1} f_{\text{N}}(\theta_{l_1 1}; 0, \psi_{l_1 1}) + (1 - w_{l_1 1}) f_{\text{N}}(\theta_{l_1 1}; 0, r\psi_{l_1 1})}\right).$$

Similarly, for any $l_2 = 1, \ldots, p_2$

$$\delta_{l_2 2} \mid \ldots \sim \text{Bern}\left(\frac{w_{l_2 2} f_{\text{N}}(\theta_{l_2 2}; 0, \psi_{l_2 2})}{w_{l_2 2} f_{\text{N}}(\theta_{l_2 2}; 0, \psi_{l_2 2}) + (1 - w_{l_2 2}) f_{\text{N}}(\theta_{l_2 2}; 0, r\psi_{l_2 2})}\right).$$

Full conditionals for $\psi_{l_1 1}$ and $\psi_{l_2 2}$: for any $l_1 = 1, \ldots, p_1$

$$\psi_{l_1 1} \mid \ldots \sim \text{inv-Gamma}\left(v + \frac{1}{2}, Q + \frac{\theta_{l_1 1}^2}{2r^{(1-\delta_{l_1 1})}}\right).$$

Similarly, for any $l_2 = 1, \ldots, p_2$

$$\psi_{l_2 2} \mid \ldots \sim \text{inv-Gamma}\left(v + \frac{1}{2}, Q + \frac{\theta_{l_2 2}^2}{2r^{(1-\delta_{l_2 2})}}\right).$$

Full conditional for $w_{l1} = w_{l2}$: for every $l = 1, \ldots, p$

$$w_{l1} \mid \ldots \sim \text{Beta}\left(a + \delta_{l1} + \delta_{l2}, b + 2 - \delta_{l1} - \delta_{l2}\right).$$

Full conditionals for $w_{l_1 1}$ and $w_{l_2 2}$: for any $l_1 = p+1, \ldots, p_1$:

$$w_{l_1 1} \mid \ldots \sim \text{Beta}\left(a + \delta_{l_1 1}, b + 1 - \delta_{l_1 1}\right).$$

Similarly, for any $l_2 = p+1, \ldots, p_2$

$$w_{l_2 2} \mid \ldots \sim \text{Beta}\left(a + \delta_{l_2 2}, b + 1 - \delta_{l_2 2}\right).$$

Full conditional for the pseudo-outcomes $\boldsymbol{V}_2$ *and* $\boldsymbol{V}$*:* for $i = 1, \ldots, n$, we have

$$\begin{aligned}
V_{i2}^{(0)} \mid \ldots &\sim \text{N}(\boldsymbol{\theta}_2^T[\boldsymbol{z}_{i2} \times (a_{i1}, 0)], \sigma_2), \\
V_{i2}^{(1)} \mid \ldots &\sim \text{N}(\boldsymbol{\theta}_2^T[\boldsymbol{z}_{i2} \times (a_{i1}, 1)], \sigma_2), \\
V_{i}^{(0)} \mid \ldots &\sim \text{N}(\boldsymbol{\theta}_1^T[\boldsymbol{z}_{i1} \times 0], \sigma_1), \\
V_{i}^{(1)} \mid \ldots &\sim \text{N}(\boldsymbol{\theta}_1^T[\boldsymbol{z}_{i1} \times 1], \sigma_1).
\end{aligned}$$

Full conditionals for the hyper-parameters a *and* b*:*

$$a \mid \ldots \sim \left(\frac{\Gamma(a+b)}{\Gamma(a)}\right)^{p_1+p_2-p} \left(\prod_{l_1=1}^{p_1} w_{l_1 1} \prod_{l_2=p+1}^{p_2} w_{l_2 2}\right)^{(a-1)} p(a),$$

$$b \mid \ldots \sim \left(\frac{\Gamma(a+b)}{\Gamma(b)}\right)^{p_1+p_2-p} \left(\prod_{l_1=1}^{p_1} (1 - w_{l_1 1}) \prod_{l_2=p+1}^{p_2} (1 - w_{l_2 2})\right)^{(b-1)} p(b).$$

where $p(a)$ and $p(b)$ represent the prior distributions of a and b, in this work assumed as inverse gamma with both shape and scale parameters equal to 1. To sample from the full conditional distributions of a and b, we utilize a random-walk Metropolis-Hastings algorithm (Chib and Greenberg, 1995), with uniform proposal centered at the current value of the hyper-parameter, with variances chosen to achieve optimal acceptance rates (Gelman et al., 1997).

2.D Additional results on the simulation study

The next tables report the results of the same analysis displayed in Section 2.4, this time for data simulated under the regime with $n > p_2$.

Table 2.5: *First experiment: simulated data with n larger than p_2. For stage 2 and stage 1, three summaries of the accuracy in selecting significant variables (proportion of FN, proportion of FP, and F_1 score), for DSS, ISS, and QLL, across the nine scenarios obtained by specifying $(k, n) \in \{(10, 50), (20, 100), (30, 150)\}$ and $\rho^* \in \{0.3, 0.6, 0.9\}$. Results are based on 100 replicated datasets.*

Stage 2											
n	p_2	ρ^*	FN			FP			F_1 score		
			DSS	ISS	QLL	DSS	ISS	QLL	DSS	ISS	QLL
50	33	0.3	0.067	0.017	0.179	0.056	0.052	0.197	0.933	0.933	0.759
		0.6	0.037	0.063	0.219	0.106	0.061	0.160	0.881	0.902	0.789
		0.9	0.053	0.089	0.286	0.065	0.091	0.044	0.919	0.869	0.787
100	63	0.3	0.037	0.043	0.185	0.087	0.064	0.121	0.875	0.909	0.787
		0.6	0.019	0.078	0.147	0.059	0.053	0.097	0.939	0.911	0.808
		0.9	0.027	0.074	0.143	0.030	0.036	0.092	0.953	0.918	0.824
150	93	0.3	0.019	0.075	0.138	0.061	0.038	0.115	0.931	0.917	0.815
		0.6	0.042	0.078	0.150	0.038	0.049	0.078	0.927	0.917	0.837
		0.9	0.030	0.080	0.133	0.042	0.028	0.068	0.941	0.926	0.856
Stage 1											
n	p_1	ρ^*	FN			FP			F_1 score		
			DSS	ISS	QLL	DSS	ISS	QLL	DSS	ISS	QLL
50	22	0.3	0.049	0.132	0.162	0.067	0.028	0.151	0.933	0.896	0.732
		0.6	0.025	0.192	0.251	0.100	0.046	0.122	0.907	0.852	0.780
		0.9	0.031	0.078	0.210	0.079	0.035	0.029	0.915	0.911	0.848
100	42	0.3	0.020	0.104	0.208	0.101	0.014	0.104	0.911	0.929	0.790
		0.6	0.060	0.055	0.212	0.079	0.048	0.075	0.887	0.921	0.809
		0.9	0.002	0.100	0.219	0.047	0.034	0.076	0.947	0.906	0.794
150	62	0.3	0.039	0.186	0.157	0.129	0.014	0.066	0.874	0.897	0.845
		0.6	0.020	0.049	0.184	0.146	0.060	0.057	0.877	0.918	0.851
		0.9	0.037	0.044	0.189	0.040	0.023	0.049	0.938	0.949	0.847

2.E Additional results on the clinical case study

We report additional information on the MIMIC-III data (see Table 2.7) and on the goodness-of-fit of DSS, ISS and QLL models for the same data (see Table 2.8). We also present the marginal probabilities of individual antihypertensive agents being part of the optimal treatment combination for the same two patients in the main manuscript (see Figure 2.3).

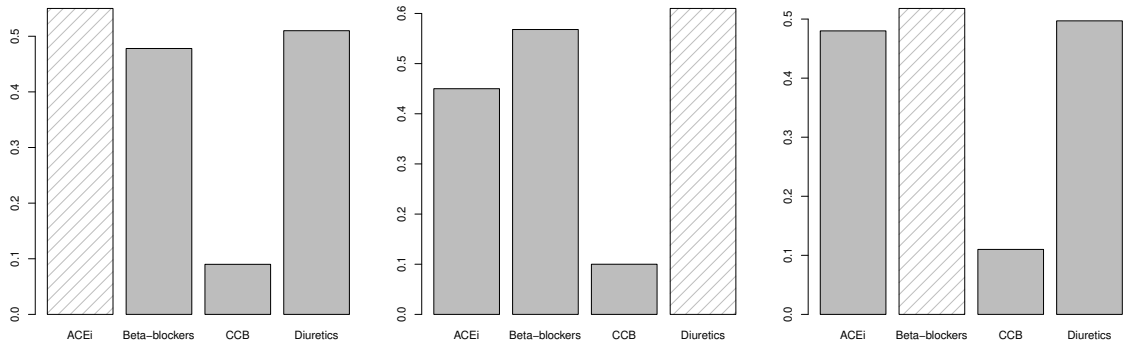

Figure 2.3: *MIMIC-III data. Posterior marginal probability of each of four antihypertensive agents being part of the optimal combination of treatments, for Individual #1 at stage 1 (left panel), for Individual #2 at stage 2 (middle panel), and for Individual #2 at stage 1 (right panel). In each barplot, the white bin indicates the largest posterior marginal probability.*

Table 2.6: *First experiment: simulated data with n larger than p_2. For stage 2, stage 1 and overall, two measures of prediction accuracy (MRE and ER), for DSS, ISS and QLL, across the nine scenarios obtained by specifying $(k, n) \in \{(10, 50), (20, 100), (30, 150)\}$ and $\rho^* \in \{0.3, 0.6, 0.9\}$. Results are based on 100 replicated datasets.*

Stage 2								
n	p_2	ρ^*	MRE			ER		
			DSS	ISS	QLL	DSS	ISS	QLL
50	33	0.3	0.001	0.002	0.150	0.014	0.000	0.108
		0.6	0.001	0.004	0.144	0.008	0.048	0.084
		0.9	0.000	0.000	0.196	0.013	0.004	0.036
100	63	0.3	0.002	0.002	0.125	0.010	0.030	0.048
		0.6	0.000	0.002	0.120	0.004	0.028	0.052
		0.9	0.000	0.001	0.120	0.009	0.020	0.053
150	93	0.3	0.000	0.000	0.096	0.007	0.013	0.044
		0.6	0.001	0.002	0.102	0.015	0.023	0.043
		0.9	0.001	0.001	0.101	0.010	0.012	0.043
Stage 1								
n	p_1	ρ^*	MRE			ER		
			DSS	ISS	QLL	DSS	ISS	QLL
50	22	0.3	0.002	0.003	0.169	0.028	0.100	0.146
		0.6	0.000	0.002	0.158	0.032	0.076	0.052
		0.9	0.001	0.000	0.165	0.020	0.032	0.104
100	42	0.3	0.001	0.003	0.104	0.067	0.042	0.057
		0.6	0.003	0.000	0.114	0.016	0.006	0.085
		0.9	0.002	0.003	0.111	0.023	0.030	0.070
150	62	0.3	0.005	0.001	0.086	0.047	0.019	0.045
		0.6	0.001	0.000	0.096	0.032	0.027	0.041
		0.9	0.005	0.001	0.089	0.033	0.031	0.058
Overall								
n	p_2	ρ^*	MRE			ER		
			DSS	ISS	QLL	DSS	ISS	QLL
50	(22, 33)	0.3	0.002	0.003	0.160	0.038	0.100	0.228
		0.6	0.000	0.003	0.151	0.040	0.124	0.132
		0.9	0.000	0.000	0.180	0.032	0.036	0.130
100	(42, 63)	0.3	0.002	0.002	0.114	0.077	0.070	0.103
		0.6	0.002	0.001	0.117	0.020	0.032	0.130
		0.9	0.001	0.002	0.115	0.032	0.050	0.122
150	(62, 93)	0.3	0.002	0.000	0.091	0.047	0.031	0.085
		0.6	0.001	0.001	0.099	0.045	0.050	0.082
		0.9	0.003	0.001	0.095	0.043	0.043	0.100

Table 2.7: *Frequency analysis of actual medication to patients by doctors in two stages. A refers to angiotensin-converting enzyme inhibitors (ACEi), B refers to beta-blockers, C refers to calcium channel blockers (CCB), and D refers to diuretics.*

	A	B	C	D	AB	AD	BD	ABD	Total
Stage 2	23	60	15	143	28	14	31	6	320
(%)	7.2	18.8	4.7	44.7	8.8	4.4	9.7	1.9	
Stage 1	108	165	41	239	66	28	71	24	742
(%)	14.6	22.2	5.5	32.2	8.9	3.8	9.6	3.2	

Table 2.8: *MIMIC-III data. Values taken by BIC and LPML for DSS, ISS and QLL. LPML is not applicable to QLL, hence the missing entries.*

	Stage 2			Stage 1		
	DSS	ISS	QLL	DSS	ISS	QLL
BIC	3301.565	3336.514	3561.146	6729.134	6760.931	7163.631
LPML	-1563.793	-1564.403	n.a.	-3532.095	-3565.184	n.a.

Chapter 3

Dependent spike-and-slab for drug combinations in diabetic kidney disease

3.1 Introduction

Diabetic kidney disease (DKD) is a chronic kidney condition linked to poorly controlled diabetes (Thomas et al., 2015). If the individual is dealing with prediabetes, taking action to prevent type II diabetes is an important step in preventing kidney disease. DKD leads to kidney damage, such as filtration issues due to high blood sugar levels, and it is a leading cause of End-Stage Renal Disease (ESRD) (Cameron, 2006). Risk factors include modifiable and non-modifiable factors that contribute to DKD's progression. Currently, modifiable factors include hypertension, glycemic level management, dyslipidemia, and smoking, while non-modifiable factors are race, age, gestation, genetic profile, and sex (Lizicarova et al., 2014; Natesan and Kim, 2021). Due to its complex and the individual's unique characteristics, personalized medicine is crucial in managing DKD. By analyzing an individual's clinical history, personalized medicine allows for effective therapies and hospitalizations. DTRs represent an essential aspect of personalized medicine, offering a systematic approach to adapt treatment strategies over time based on an individual's evolving health responses. By incorporating the principles of DTRs into personalized medicine, healthcare practitioners can optimize treatment outcomes by adjusting interventions according to the patient's specific needs and ongoing health changes. This dynamic approach acknowledges that patients may respond differently to treatments at various stages, emphasizing the importance of tailoring medical decisions to real-time, individualized data. A common approach in DTRs is Q-learning, which focuses on regression-based models (Chakraborty and Murphy, 2014). As pointed out from Chakraborty and Murphy (2014), Q-learning is a version of the optimal structural nested mean model developed in the causal inference literature (Robins, 2004).

Several studies have proposed various Q-learning approaches in DTRs to address complex data challenges (Wallace and Moodie, 2015; Fan et al., 2019; Bhattacharya et al., 2023; Lyu et al., 2023).

In this chapter, we extend the Bayesian augmented learning approach proposed in Chapter 2 and published in Bi et al. (2024). Within the framework of Q-learning to advance the goal of personalized medicine in the management of DKD, BAL addresses the uncertainty in identifying optimal decision sequences and incorporates dimensionality reduction to manage high-dimensional individual covariates. Additionally, we use a dependent spike-and-slab (DSS) prior for BAL as in Chapter 2, which takes into account the association between the significance of covariates across multiple stages.

More specifically, the contributions of this chapter can be summarized as follows:

(i) We analyze an extremely rich dataset about DKD, and more precisely, we use data from the Prospective Observational Study in Patients with type II diabetes Mellitus for Validation of Biomarkers (PROVALID) (Mayer et al., 2016; Eder et al., 2018, 2019; Thöni et al., 2021). This study encompasses a diverse spectrum of chronic kidney disease (CKD) states, providing a comprehensive foundation for assessing and validating biomarkers in the context of personalized medicine for DKD management. Here, the focus is on drug responses to DKD at the individual level, explicitly considering the estimated Glomerular Filtration Rate. These data were previously used by Jones et al. (2023) and Abebe et al. (2024). Jones et al. (2023) based their study on a subset of variables that may influence type II diabetes. More specifically, nine variables are selected based on expert knowledge. Differently, in the work of Abebe et al. (2024), the number of variables considered in the analysis is 30, and they are selected through a Bayesian network approach. Our work considers 100 variables that may impact the type II diabetes course during the various stages. We also consider interactions among the variables, leading to an increase in the complexity of the problem.

(ii) We extend the BAL approach from two to four stages. The ability to consider more than two stages in a dynamic treatment regime problem opens up the method's application to more complex studies by allowing the optimal treatment to be determined even in the most advanced stages of the problem.

(iii) We leverage False Discovery Rate (FDR) (Müller et al., 2004; Ma et al., 2023) for variable selection purposes. Controlling FDR is a way to identify as many significant

variables as possible while incurring a relatively low proportion of false positives. FDR allows the selection of the significant molecular biomarkers and clinical parameters on drug effect to eGFR and considers the association of significance of a particular variable repeated in multiple stages.

(iv) We propose a predictive model to determine the optimal drug combination to enhance healthcare outcomes for renal pathologies at the individual level with type II diabetes.

The chapter's organization is outlined as follows: in Section 3.2, we introduce the dataset derived from the distinct prospective observational study known as PROVALID. In Section 3.3, we present the Bayesian Augmented Learning model for a four stage dynamic treatment regime. Furthermore, we introduce the use of false discovery rate for significant variable selection, and we describe the predictive model for optimal decision-making. Section 3.4 summarizes the main results obtained on the PROVALID data. The Section 3.5 engages in a discussion of the results while also outlining future considerations aimed at complementing and extending the current model. Details on the posterior computation and additional results from the case study are presented in the Appendix.

3.2 Data

Our research utilizes data focused on individuals diagnosed with a specific form of diabetic kidney disease (DKD). This data is obtained from a prospective observational study known as PROVALID (PROspective cohort study in patients with type II diabetes mellitus for VALIDation of biomarkers) (Mayer et al., 2016; Eder et al., 2018, 2019; Thöni et al., 2021), including a diverse spectrum of chronic kidney disease (CKD) stages.

The PROVALID study collected data on 4,000 individuals diagnosed with type II diabetes mellitus (DM2). The data includes information on the individual's medical history, physical status, laboratory measurements, and medication use, being taken care of at the primary level of healthcare in 5 European countries (Austria, Hungary, Netherlands, Poland, and Scotland). The data was collected prospectively, meaning it was collected over time as the participants' conditions progressed (Gregorich et al., 2021). Patient data was collected during consecutive visits for at least five years, and only one treatment was given by the doctor from the four treatment choices available during a visit. The four drug treatments are:

1. Renin-Angiotensin-System-inhibitor (RASi) only treatment;

2. A combination of sodium-glucose transporter 2 inhibitors (SCLT2i) and the RASi treatment;

3. A combination of mineralocorticoid Receptor Antagonist (MCRa) and the RASi treatment;

4. A combination of glucagon-like peptide 1 receptor agonist (GLP1a) and the RASi treatment.

RASi protects the kidneys and prevents other cardiovascular diseases. Furthermore, in End-Stage Renal Disease (ESRD), patients should use RASi antihypertensive drugs and control various risk factors, including low blood sugar, high blood cholesterol, high blood pressure, overweight, and smoking. Sodium-glucose cotransporter 2 inhibitors (SGLT2i) and glucagon-like peptide-1 receptor agonists (GLP1a), in combination with RASi, are used to reduce the risk of kidney disease (Yang et al., 2023). More precisely, SGLT2 inhibitor demonstrated effectiveness in treating DKD in the context of the Canagliflozin and Renal Events in Diabetes study (Yamazaki et al., 2021). GLP1a is used to stabilize the global Glomerular Filtration Rate (Hu and Du, 2019; Wang et al., 2021) by contrasting vasoconstrictive response triggered by glomerular feedback. The combination of MCRa and RASi is considered to address cases of hypertension and heart failure (HF) characterized by a diminished ejection fraction (Sica, 2015).

Our analysis study uses sampled patients from PROVALID dataset, which is longitudinal data, to build the optimal treatment decision model for DKD. When determining the most optimal drug combinations, it is essential to consider each individual's medical history and characteristics rigorously.

The sampled dataset includes 223 patients, including 113 males and 110 females, and the average age for them to enroll in the first-time therapy was 53.6. Selected patients are those who have at least 4 consecutive visits. The patients' history was constructed using 100 clinical factors, including the individual's medical history, physical status, laboratory measurements, and medication use from this longitudinal dataset. The model incorporates 404 predictors, including 100 clinical factors, 4 treatments represented by 3 dummy variables, interaction terms between the 100 clinical factors and these 3 dummy variables, and an intercept.

In such a situation, statistical techniques capable of extracting information efficiently from limited data sets are necessary. Our proposal relies on a variable selection method based on Bayesian correlated spike-and-slab regression, which reduces dimensionality and

selects important variables from 404 predictors to make optimal decisions.

After each visit, we use a patient's estimated Glomerular Filtration Rate (eGFR) value as the final payoff in each stage, and eGFR is the longitudinal outcome. An eGFR greater than 90 indicates normal kidney function, 60-89 reflects kidney damage with mild loss of function, 30-59 suggests moderate loss of function, and values below 29 indicate severe loss of kidney function. As shown in Figure 3.1, the top-left panel illustrates the eGFR trajectories for all individuals, with red lines representing cutoff values for different levels of kidney function. This panel reveals that most individuals with type II diabetes exhibit some degree of kidney dysfunction at each stage. The top-right panel provides boxplots of eGFR across all individuals. The mean eGFR values, highlighted in red, decline from 63.35 in the first stage to 62.61 in the second, 59.81 in the third, and 57.60 in the fourth stage, indicating a progressive deterioration in kidney function over time. The bottom panel of Figure 3.1 presents eGFR boxplots stratified by treatment at each stage, where the average eGFR for the RASi & GLP1a and RASi & SGLT2i treatment groups are higher than that for the RASi-only and RASi & MCRa groups. We also present the individual eGFR trajectories for 36 randomly selected samples in Figure 3.2. Despite receiving the prescribed medications to improve kidney function, some patients did not fully restore their eGFR to the normal range. The trends among these patients vary, with some demonstrating an upward trajectory, others exhibiting a downward trend, and a portion maintaining stability in their eGFR.

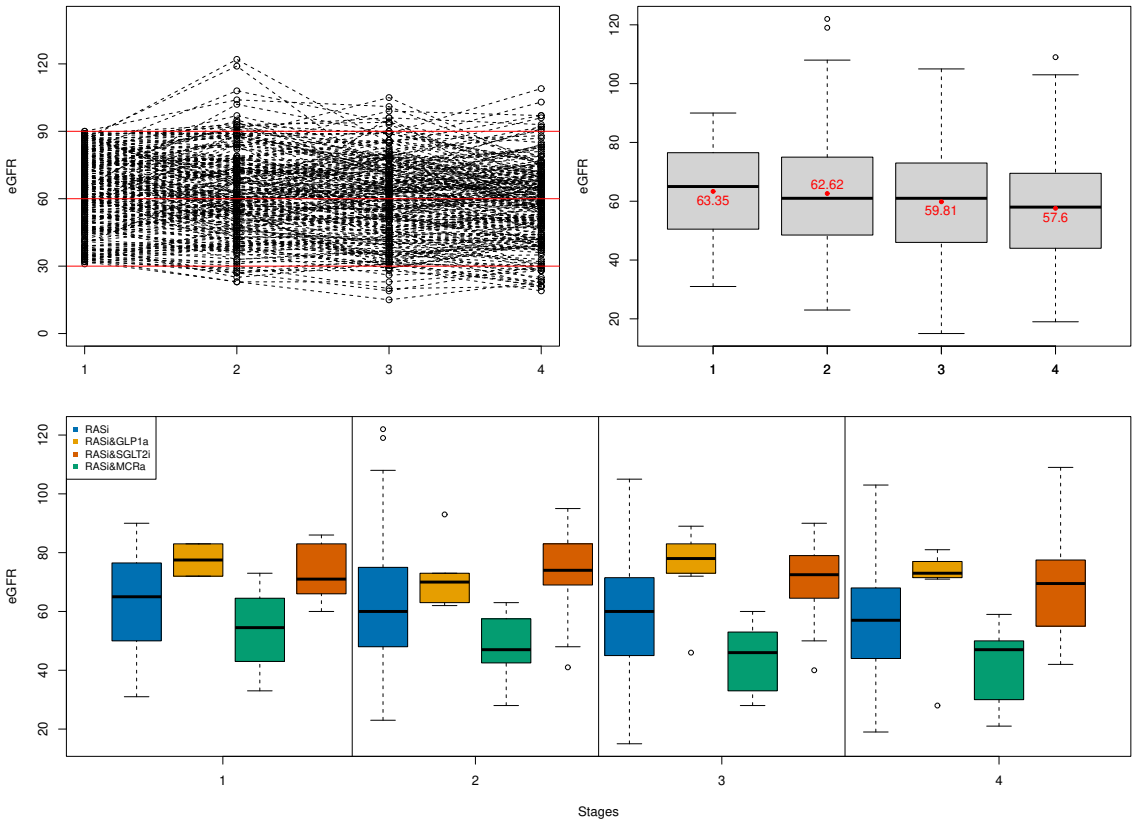

Figure 3.1: *Depiction of eGFR over time. The top-left panel illustrates the trajectories of eGFR for all observations, with red lines indicating critical cutoff values for assessing kidney function. The top-right panel presents a boxplot of eGFR across 4 stages, where the red values represent the mean eGFR at each stage. The bottom panel displays boxplots of eGFR, stratified by the four treatments in each stage.*

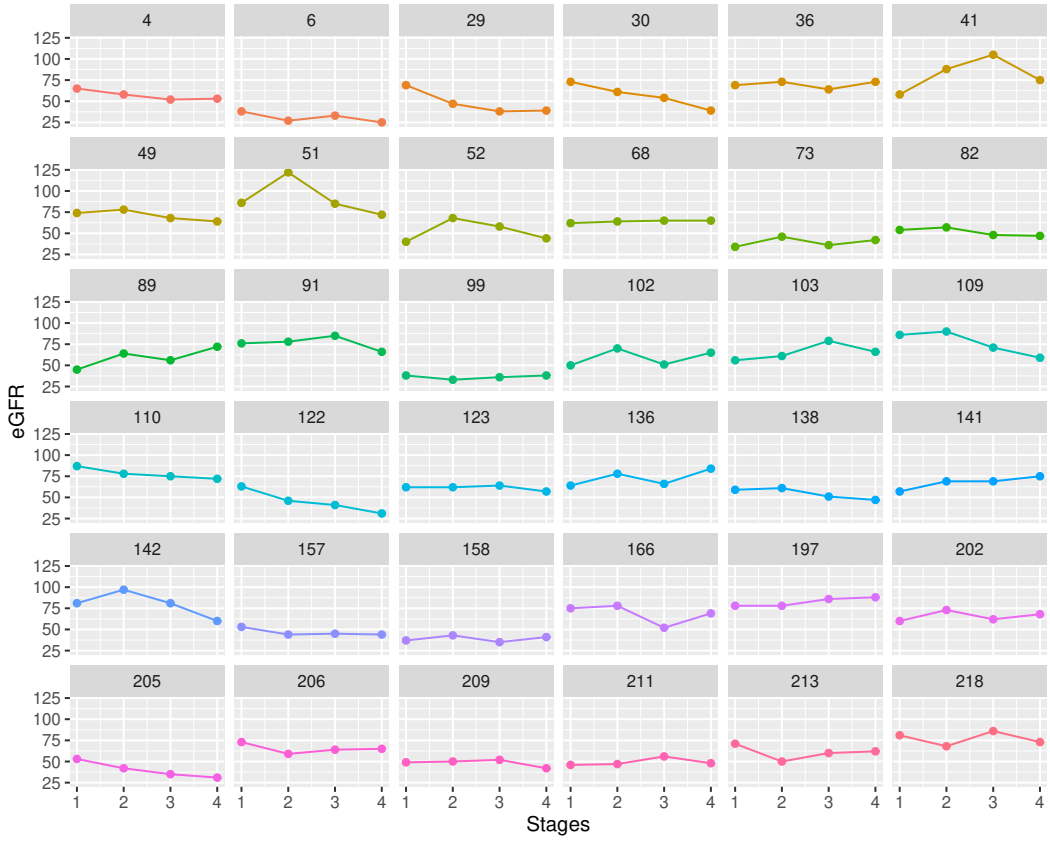

Figure 3.2: *The individual trajectories of eGFR for 16% of the dataset.*

The medical interventions administered to the 223 patients at each stage are summarized in Table 3.1. Notably, more than 85% of patients consistently received RASi treatment alone at each stage. Additionally, there is potential to introduce supplementary medications, such as SGLT2i, GLP1a, or MCRa, in combination with RASi. Furthermore, transitioning between different drug combinations may mitigate further patient eGFR decline.

	RASi	RASi & SGLT2i	RASi & MCRa	RASi & GLP1a
Stage 1	212 (95.07%)	5 (2.24%)	4 (1.79%)	2 (0.90%)
Stage 2	193 (86.55%)	16 (7.17%)	8 (3.59%)	6 (2.69%)
Stage 3	191 (85.65%)	16 (7.17%)	9 (4.04%)	7 (3.14%)
Stage 4	191 (85.65%)	16 (7.17%)	9 (4.04%)	7 (3.14%)

Table 3.1: *Medical interventions administered to the 223 patients across four stages.*

3.3 Methods

We implement the approach proposed in Chapter 2, the so-called Bayesian Augmented Learning (BAL) method, incorporating a dependent spike and slab prior to the regression model coefficients. In this work, we extend the approach from two to four stages. Additionally, we introduce a way of controlling the number of selected significant variables by exploiting the false discovery rate. In the context of our study utilizing the PROVALID dataset, where decisions are based on a comprehensive set of 404 predictors at each stage, the importance of variable selection cannot be overstated. Furthermore, we propose a predictive model for optimal treatment prediction.

3.3.1 Bayesian model and variable selection

The extension of BAL to a DTR framework with four stages, requires the introduction of new notation. We let $A_j \in \mathcal{A}_j = \{0, 1, \ldots, T-1\}$, with $T = 4$, indicate the treatment, or action, assigned at the jth stage, with the index $j = 1, 2, 3, 4$ henceforth used to refer to specific stages. Moreover, $Y = Y_1 + Y_2 + Y_3 + Y_4$ denotes the overall payoff, with Y_j being the part of the payoff observed after the jth stage. Henceforth we use the index i, with $i = 1, \ldots, n$, to denote quantities referring to the ith statistical unit, e.g. A_{ij} denotes the treatment assigned to the ith individual in the sample at the jth stage, similarly for Y_{ij}. Also in this case, in order to stress the dependence of a payoff on the assigned sequence of treatments, we write $Y_i(a_{i1}, a_{i2}, a_{i3}, a_{i4}) = Y_{i1}(a_{i1}) + Y_{i2}(a_{i1}, a_{i2}) + Y_{i3}(a_{i1}, a_{i2}, a_{i3}) + Y_{i4}(a_{i1}, a_{i2}, a_{i3}, a_{i4})$. The observed payoffs y_i are realizations of the corresponding random variables Y_i. We let $\boldsymbol{z}_{ij}$ denote the set of covariates available for the ith individual at the jth stage and we denote the observed data by $\mathcal{D} = \{(y_{i1}, y_{i2}, y_{i3}, y_{i4}, a_{i1}, a_{i2}, a_{i3}, a_{i4}, \boldsymbol{z}_{i1}, \boldsymbol{z}_{i2}, \boldsymbol{z}_{i3}, \boldsymbol{z}_{i4}) : i = 1, \ldots, n\}$. We further introduce the notations $Y_{i3:4}^{\text{OPT}}, Y_{i2:4}^{\text{OPT}}$, and $Y_{i1:4}^{\text{OPT}}$ to denote the sequence of cumulative payoffs defined as

$$\begin{aligned} Y_{i3:4}^{\text{OPT}} &= Y_{i3}(a_{i3}) + Y_{i4}(a_{i4}^{\text{OPT}}) \\ Y_{i2:4}^{\text{OPT}} &= Y_{i2}(a_{i2}) + Y_{i3}(a_{i3}^{\text{OPT}}) + Y_{i4}(a_{i4}^{\text{OPT}}) \\ Y_{i1:4}^{\text{OPT}} &= Y_{i1}(a_{i1}) + Y_{i2}(a_{i2}^{\text{OPT}}) + Y_{i3}(a_{i3}^{\text{OPT}}) + Y_{i4}(a_{i4}^{\text{OPT}}), \end{aligned} \tag{3.1}$$

where $a_{i2}^{\text{OPT}}, a_{i3}^{\text{OPT}}$, and a_{i4}^{OPT} are the treatments that, given the treatments assigned in previous stages, are optimal at stages 2, 3, and 4, respectively. That is, for each $j = 1, 2, 3$, $Y_{ij:4}^{\text{OPT}}$ represents the cumulative payoffs for the last $4 - j + 1$ stages, if we assume that the treatment assigned at the jth stage is a_{ij}, while the treatments assigned at later stages are

optimal. As in Chapter 2, we follow the principle of backward induction (see, e.g., Tsiatis et al., 2019), and we consider a sequence of four stage-specific regression models, one for the fourth stage payoffs Y_{i4} and the others for the optimal cumulative payoffs $Y_{i1:4}^{\text{OPT}}$, $Y_{i2:4}^{\text{OPT}}$ and $Y_{i3:4}^{\text{OPT}}$. Namely,

$$\begin{aligned} Y_{i4} \mid a_{i4}, \boldsymbol{z}_{i4}, \boldsymbol{\theta}_4 &\overset{\text{ind}}{\sim} f_4(y \mid a_{i4}, \boldsymbol{z}_{i4}; \boldsymbol{\theta}_4) \\ Y_{i3:4}^{\text{OPT}} \mid a_{i3}, \boldsymbol{z}_{i3}, \boldsymbol{\theta}_3 &\overset{\text{ind}}{\sim} f_3(y \mid a_{i3}, \boldsymbol{z}_{i3}; \boldsymbol{\theta}_3) \\ Y_{i2:4}^{\text{OPT}} \mid a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2 &\overset{\text{ind}}{\sim} f_2(y \mid a_{i2}, \boldsymbol{z}_{i2}; \boldsymbol{\theta}_2) \\ Y_{i1:4}^{\text{OPT}} \mid a_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1 &\overset{\text{ind}}{\sim} f_1(y \mid a_{i1}, \boldsymbol{z}_{i1}; \boldsymbol{\theta}_1), \end{aligned} \tag{3.2}$$

where f_j is the regression function, $\boldsymbol{\theta}_j$ is the vector of regression coefficients for $j = 1, 2, 3, 4$. Compared to the corresponding model specification (2.1) for the two-stage framework, (3.2) introduces a convenient simplification to handle the presence of four stages. Specifically, each regression only accounts for the treatment administered at a specific stage, rather than incorporating the entire history of treatments given up to that point. Given that we also consider interactions between covariates and treatments, this simplification has the advantage of preventing the total number of regressors from growing exponentially in the stage index j. Similarly to Chapter 2, we introduce an optimal decision rule in which larger payoffs are preferable. We propose to identify the optimal treatment regime as the sequence of treatments that is most likely to maximize the individual cumulative payoffs. Namely,

$$\begin{aligned} a_{i4}^{\text{OPT}} &= \arg\sup_{a \in \mathcal{A}_4} (\mathbb{E}[Y_{i4} \mid a, \boldsymbol{z}_{i4}, \boldsymbol{\theta}_4]), \\ a_{i3}^{\text{OPT}} &= \arg\sup_{a \in \mathcal{A}_3} (\mathbb{E}[Y_{i3:4}^{\text{OPT}} \mid a, a_{i4}^{\text{OPT}}, \boldsymbol{z}_{i3}, \boldsymbol{\theta}_3]), \\ a_{i2}^{\text{OPT}} &= \arg\sup_{a \in \mathcal{A}_2} (\mathbb{E}[Y_{i2:4}^{\text{OPT}} \mid a, a_{i3}^{\text{OPT}}, a_{i4}^{\text{OPT}}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2), \\ a_{i1}^{\text{OPT}} &= \arg\sup_{a \in \mathcal{A}_1} (\mathbb{E}[Y_{i1:4}^{\text{OPT}} \mid a, a_{i2}^{\text{OPT}}, a_{i3}^{\text{OPT}}, a_{i4}^{\text{OPT}}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1]), \end{aligned} \tag{3.3}$$

While Model (3.2) is flexible as for the form in which covariates and treatment actions are included as regressors, as in Chapter 2, we consider continuous payoffs and model them with normal multiple linear regressions, that is

$$f_j(y \mid a_{ij}, \boldsymbol{z}_{ij}; \boldsymbol{\theta}_j) = f_{\text{N}}(y; \boldsymbol{\theta}_j^{\intercal}[\boldsymbol{z}_{ij} \times \boldsymbol{a}_{ij}^*], \sigma_j^2), \tag{3.4}$$

for $j = 1, \ldots, 4$, where $f_{\text{N}}(y; \mu, \sigma^2)$ is the probability density function of a normal random variable with mean μ and variance σ^2, $\boldsymbol{a}_{ij}^* = (a_{ij1}^*, \ldots, a_{ij3}^*) \in \mathcal{D}_3$ are vectors of dummy variables such that $a_{ijt}^* = \mathbb{I}_{\{t\}}(a_{ij})$, and the notation $[\boldsymbol{z} \times \boldsymbol{a}]$ indicates the vector of regressors

including $\boldsymbol{z}$, $\boldsymbol{a}$, and all the interactions between the covariates $\boldsymbol{z}$ and the binary variables in $\boldsymbol{a}$. We extend the approach proposed in Chapter 2 for identifying optimal treatment sequences to a four-stage framework. This extension requires defining 4-dimensional cumulative pseudo-outcomes, denoted by $\boldsymbol{V}_{i4}^{(0:3)} = (V_{i4}^{(0)}, \ldots, V_{i4}^{(3)})$ and $\boldsymbol{V}_{ij:4}^{(0:3)} = (V_{ij:4}^{(0)}, \ldots, V_{ij:4}^{(3)})$ for $j = 1, 2, 3$. The logic guiding the model setup for observations and pseudo-outcomes is the same as the one illustrated in Chapter 2, though the notation is inevitably more complex. Specifically, the model is defined as follows:

$$
\begin{aligned}
(Y_{i4}, \boldsymbol{V}_{i4}^{(0:3)}) \mid a_{i4}, \boldsymbol{z}_{i4}, \boldsymbol{\theta}_4 &\sim f_{\mathrm{N}}(y; \boldsymbol{\theta}_4^{\intercal}[\boldsymbol{z}_{i4} \times \boldsymbol{a}_{i4}^*], \sigma_4^2) \prod_{t=0}^{3} f_{\mathrm{N}}(v_t; \boldsymbol{\theta}_4^{\intercal}[\boldsymbol{z}_{i4} \times \boldsymbol{i}_t^*], \sigma_4^2) \\
(Y_{i3}, \boldsymbol{V}_{i3:4}^{(0:3)}) \mid \boldsymbol{V}_{i4}^{(0:3)}, a_{i3}, \boldsymbol{z}_{i3}, \boldsymbol{\theta}_3 &\sim f_{\mathrm{N}}(y + \max_t\{V_{i4}^{(t)}\}; \boldsymbol{\theta}_4^{\intercal}[\boldsymbol{z}_{i3} \times \boldsymbol{a}_{i3}^*], \sigma_3^2) \\
&\quad \times \prod_{t=0}^{3} f_{\mathrm{N}}(v_t; \boldsymbol{\theta}_3^{\intercal}[\boldsymbol{z}_{i3} \times \boldsymbol{i}_t^*], \sigma_3^2) \\
(Y_{i2}, \boldsymbol{V}_{i2:4}^{(0:3)}) \mid \boldsymbol{V}_{i3:4}^{(0:3)}, a_{i2}, \boldsymbol{z}_{i2}, \boldsymbol{\theta}_2 &\sim f_{\mathrm{N}}(y + \max_t\{V_{i3:4}^{(t)}\}; \boldsymbol{\theta}_2^{\intercal}[\boldsymbol{z}_{i2} \times \boldsymbol{a}_{i2}^*], \sigma_2^2) \\
&\quad \times \prod_{t=0}^{3} f_{\mathrm{N}}(v_t; \boldsymbol{\theta}_2^{\intercal}[\boldsymbol{z}_{i2} \times \boldsymbol{i}_t^*], \sigma_2^2) \\
(Y_{i1}, \boldsymbol{V}_{i1:4}^{(0:3)}) \mid \boldsymbol{V}_{i2:4}^{(0:3)}, a_{i1}, \boldsymbol{z}_{i1}, \boldsymbol{\theta}_1 &\sim f_{\mathrm{N}}(y + \max_t\{V_{i2:4}^{(t)}\}; \boldsymbol{\theta}_1^{\intercal}[\boldsymbol{z}_{i1} \times \boldsymbol{a}_{i1}^*], \sigma_1^2) \\
&\quad \times \prod_{t=0}^{3} f_{\mathrm{N}}(v_t; \boldsymbol{\theta}_1^{\intercal}[\boldsymbol{z}_{i1} \times \boldsymbol{i}_t^*], \sigma_1^2),
\end{aligned}
\tag{3.5}
$$

where $\boldsymbol{i}_t \in \mathcal{D}_3$ is a null vector if $t = 0$ and is a vector of zeros with a 1 in position t if $t \in \{1, \ldots, 3\}$. The introduction of the cumulative pseudo-outcomes allows us to devise an efficient computational scheme for posterior simulation with similar advantages to those discussed in Chapter 2 for the two-stage framework. We observe that the total number of pseudo-outcomes in model (3.5) grows linearly with both the number of stages and the number of treatments T.

3.3.1.1 Variable selection

We extend the DSS prior proposed in Section 2.2.2 from two to four stages, more specifically, in Section 2.2.2, we estimate two groups of parameters across two stages, considering their correlation, whereas here we estimate four groups of parameters across four stages, considering correlations among all four stages. We retain the same notations as in Section 2.2.2.

Let $\tilde{\boldsymbol{z}}_j$ denote the vector of regressors at the jth stage, including both the covariates $\boldsymbol{z}_j$ and the corresponding interactions. We let p_j be the dimension of $\tilde{\boldsymbol{z}}_j$. Our proposal takes into account that the vectors $\tilde{\boldsymbol{z}}_j, j = 1, 2, 3, 4$ might share some common covariates and, similarly, interactions between common covariates and stage-specific actions, say a total of $p \leq \min(p_1, p_2, p_3, p_4)$ elements. While the value taken by these regressors might vary across stages, defining a joint model for the corresponding regression coefficients seems appealing, as the same regressor might play similar roles within the four regression models. To this end, we design a prior p_0 with positively correlated inclusion variables for the same regressor across the four stages. That is we consider both normal spike and normal slab components, with a ratio r between the variance of the spike and that one of the slab, considerably smaller than 1. Let δ_{lj}, with $j = 1, 2, 3, 4$ and $l = 1, \ldots, p_j$, be an indicator variable such that $\delta_{lj} = 1$ ($\delta_{lj} = 0$) if the lth regressor is (is not) included in the jth regression, and call $w_{lj} = P(\delta_{lj} = 1)$. We assume the prior variance of the regression coefficient θ_{lj} is equal to ψ_{lj} if $\delta_{lj} = 1$, and equal to $r\psi_{lj}$, with $r \ll 1$, if $\delta_{lj} = 0$. We introduce the notation $\boldsymbol{\delta}_j = (\delta_{1j}, \ldots, \delta_{p_j j})$, $\boldsymbol{\psi}_j = (\psi_{1j}, \cdots, \psi_{p_j j})$ and $\boldsymbol{w}_j = (w_{1j}, \ldots, w_{p_j j})$, and define the dependent spike-and-slab prior $p_0(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3, \boldsymbol{\theta}_4)$ as the distribution arising from the following hierarchy:

$$
\begin{aligned}
(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3, \boldsymbol{\theta}_4) \mid \boldsymbol{\delta}_1, \boldsymbol{\delta}_2, \boldsymbol{\delta}_3, \boldsymbol{\delta}_4, \boldsymbol{\psi}_1, \boldsymbol{\psi}_2, \boldsymbol{\psi}_3, \boldsymbol{\psi}_4 &\overset{\text{ind}}{\sim} \prod_{j=1}^{4} \prod_{l_j=1}^{p_j} \mathrm{N}\left(0, r^{1-\delta_{l_j j}} \psi_{l_j j}\right) \\
(\boldsymbol{\psi}_1, \boldsymbol{\psi}_2, \boldsymbol{\psi}_3, \boldsymbol{\psi}_4) &\overset{\text{iid}}{\sim} \prod_{j=1}^{4} \prod_{l_j=1}^{p_j} \text{INV-GAMMA}(\nu, Q) \\
(\boldsymbol{\delta}_1, \boldsymbol{\delta}_2, \boldsymbol{\delta}_3, \boldsymbol{\delta}_4) \mid \boldsymbol{w}_1, \boldsymbol{w}_2, \boldsymbol{w}_3, \boldsymbol{w}_4 &\overset{\text{ind}}{\sim} \prod_{j=1}^{4} \prod_{l_j=1}^{p_j} \text{BERN}(w_{p_j j}) \\
w_{l1} = w_{l2} = w_{l3} = w_{l4} &\overset{\text{iid}}{\sim} \text{BETA}(a, b) \text{ for } l = 1, \ldots, p \\
w_{l_j j} &\overset{\text{iid}}{\sim} \text{BETA}(a, b) \text{ for } l_j = p+1, \ldots, p_j \text{ and } j = 1, 2, 3, 4.
\end{aligned}
\tag{3.6}
$$

The prior defined in (3.6) is an extension to the case of two vectors of regression coefficients of the normal mixture of inverse-Gamma spike-and-slab prior proposed by Ishwaran and Rao (2003, 2005). Interestingly, the indicator variables of common regressors, that is, the pairs $(\boldsymbol{\delta}_{l1}, \boldsymbol{\delta}_{l2}, \boldsymbol{\delta}_{l3}, \boldsymbol{\delta}_{l4})$ for $l = 1, \ldots, p$, are marginally dependent, which is an appealing property as it favors borrowing of information across the four stages when performing variable selection, the idea being that, if the effect of one regressor is significant at one stage, it is likely to be significant also at others.

Sampling from the posterior distribution is implemented via Augmented Gibbs sampling, with hyperparameters $\nu = 3, Q = 4$. In addition, we use Metropolis-Hastings to update hyperparameters a and b. Details are reported in Appendix 3.A.

3.3.1.2 False discovery rate control

False Discovery Rate (FDR) arises as a performance of employing Bayesian analysis to address the variable selection problem within a hierarchical linear regression model, which is the expected proportion of falsely selected variables. Writing $D = \sum d_i$ for the number of significant predictors, the posterior expected False Discovery Rate (FDR) denoted by $\overline{\text{FDR}}(d, y)$ (Müller et al., 2004) as follows,

$$\overline{\text{FDR}}(d, y) = \frac{\sum d_i(1 - v_i)}{D + \epsilon} \tag{3.7}$$

where $v_i = P(\delta_i = 1|y)$ represents the marginal posterior probability for the ith variable, and $d_i = \text{I}(v_i \geqslant t)$, where t is the cutoff, and the additional term ϵ avoids a zero denominator.

In alignment with the recommendations by Emery and Keich (2019) and Ma et al. (2023) that advocate for practical FDR thresholds to be ≤ 0.3, we control the $\overline{\text{FDR}}$ at 0.01 to determine the cutoff value t at each stage. The threshold value of $\overline{\text{FDR}}$ was chosen in relation to the number of significant variables that are selected such that a reasonable number of variables are obtained to allow practical interpretation.

We implement $\overline{\text{FDR}}$ after completing Gibbs sampling. Once the posterior samples are obtained, $\overline{\text{FDR}}$ is applied as a post-processing step to adjust the selection threshold for variables. These identified significant variables are subsequently utilized to make predictions, enhancing our predictive models' robustness and reliability. By adhering to this FDR control strategy, we ensure a balance between identifying meaningful predictors and limiting the inclusion of false positives, ultimately improving the performance and interpretability of our hierarchical linear regression model.

3.3.2 Optimal treatment prediction

We use the $\mathscr{S}_i$ to denote the set of significant variables selected by control FDR at 0.01 in jth stage and $\boldsymbol{\theta}$ to denote the posterior mean of regression coefficients obtained by Gibbs sampler. Then we can get posterior predictive distributions for cumulative payoff across 4

stages,

$$Y_{1:4}^{\text{OPT}} \mid \tilde{\boldsymbol{z}}_{\mathscr{S}_1} \sim \text{N}(\widehat{\boldsymbol{\theta}}_{\mathscr{S}_1}\tilde{\boldsymbol{z}}_{\mathscr{S}_1}, \widehat{\sigma}_1) \tag{3.8}$$

$$Y_{2:4}^{\text{OPT}} \mid \tilde{\boldsymbol{z}}_{\mathscr{S}_2} \sim \text{N}(\widehat{\boldsymbol{\theta}}_{\mathscr{S}_2}\tilde{\boldsymbol{z}}_{\mathscr{S}_2}, \widehat{\sigma}_2) \tag{3.9}$$

$$Y_{3:4}^{\text{OPT}} \mid \tilde{\boldsymbol{z}}_{\mathscr{S}_3} \sim \text{N}(\widehat{\boldsymbol{\theta}}_{\mathscr{S}_3}\tilde{\boldsymbol{z}}_{\mathscr{S}_3}, \widehat{\sigma}_3) \tag{3.10}$$

$$Y_4 \mid \tilde{\boldsymbol{z}}_{\mathscr{S}_4} \sim \text{N}(\widehat{\boldsymbol{\theta}}_{\mathscr{S}_4}\tilde{\boldsymbol{z}}_{\mathscr{S}_4}, \widehat{\sigma}_4) \tag{3.11}$$

where $\tilde{\boldsymbol{z}}_{\mathscr{S}j}$ and $\widehat{\boldsymbol{\theta}}_{\mathscr{S}j}$ represent the design matrix and the posterior mean of regression coefficients, respectively, for the selected significant variables in the jth stage. Additionally, Y_j^{OPT} denotes the optimal payoff in the jth stage, while $\widehat{\sigma_j}$ signifies the posterior mean of σ_j for that corresponding stage.

In jth stage, consider the design matrix $\tilde{\boldsymbol{z}}_{\mathscr{S}j}$, which consists of three components: clinical factors ($\boldsymbol{z}$), treatment (A), and the interaction ($\boldsymbol{z} \cdot A$). The treatment (A) comprises four options: RASi, RASi & SGLT2i, RASi & MCRa, and RASi & GLP1a. Therefore, we can get four potential design matrices ($\tilde{\boldsymbol{z}}_{1\mathscr{S}_i}$, $\tilde{\boldsymbol{z}}_{2\mathscr{S}_i}$, $\tilde{\boldsymbol{z}}_{3\mathscr{S}_i}$, $\tilde{\boldsymbol{z}}_{4\mathscr{S}_i}$). By putting these design matrices into formulas (3.8 $\sim$ 3.11) one by one, it is straightforward to determine the posterior distributions for the cumulative potential payoffs. Based on the above discussion, the exposition of prediction details is depicted in Algorithm 2.

Algorithm 2: Prediction for Optimal Treatments from Individual Level

Data: Potential design matrix $\tilde{\boldsymbol{z}}_{1\mathscr{S}_j}$, $\tilde{\boldsymbol{z}}_{2\mathscr{S}_j}$, $\tilde{\boldsymbol{z}}_{3\mathscr{S}_j}$, $\tilde{\boldsymbol{z}}_{4\mathscr{S}_j}$ for a patient, $j = 1, 2, 3, 4$ (stage); the observed payoff Y_1, Y_2, Y_3, Y_4.

Result: Predicted Optimal Treatment $(\widehat{A}_1^{\text{OPT}}, \widehat{A}_2^{\text{OPT}}, \widehat{A}_3^{\text{OPT}}, \widehat{A}_4^{\text{OPT}})$

foreach *stage* $j \in \{1, 2, 3, 4\}$ *for one patient* **do**

- Generate n samples for each of the potential cumulative payoff distributions: $Y_{(j:4)1m}^{\text{OPT}}, Y_{(j:4)2m}^{\text{OPT}}, Y_{(j:4)3m}^{\text{OPT}}, Y_{(j:4)4m}^{\text{OPT}}$ for $m = 1, \cdots, n$;
- $\widehat{A}_{jm}^{\text{OPT}} \leftarrow \arg\max(Y_{(j:4)1m}^{\text{OPT}}, Y_{(j:4)2m}^{\text{OPT}}, Y_{(j:4)3m}^{\text{OPT}}, Y_{(j:4)4m}^{\text{OPT}})$;
- **foreach** *treatment* $t \in \{1, 2, 3, 4\}$ **do**
  - Calculate the probability of each treatment being optimal: $P(A_{jt} \text{ is the optimal treatment}) \leftarrow \frac{\#\{A_{jm}^{\text{OPT}} = t\}}{n}$;
- Find the estimated optimal treatment: $\widehat{A}_j^{\text{OPT}} \leftarrow \arg\max_{t \in \{1,2,3,4\}}\{P(A_{jt} \text{ is the optimal treatment})\}$;

return $\widehat{A}_j^{OPT}$, $j = 1, 2, 3, 4$.

3.4 Results

We implement Gibbs sampling for 10,000 iterations to approximate the posterior distribution of the model parameters, discarding the first 5,000 iterations as a burn-in period. The dataset is partitioned into two distinct subsets: 80% is allocated to the training set for model estimation, while the remaining 20% is used as the test set for the execution of predictions (Algorithm 2) and the evaluation of model performance. This section will show the model's performance on the training set, highlighting the significant variables selected by controlling the false discovery rate (FDR) at 0.01. Furthermore, we will demonstrate the predictions of the optimal treatment across 4 stages for each patient in the test set and discuss how we can enhance patient healthcare. The primary goal is to achieve a higher eGFR value when patients receive the recommended optimal treatment through dependent spike-and-slab (DSS).

As a preliminary point, our focus lies on the difference between the estimated optimal eGFR and the observed eGFR to see how much kidney function has increased under recommended optimal treatment via DSS. If this difference ratio (calculated as the estimated optimal eGFR minus the observed eGFR, divided by the observed eGFR, multiplied by 100) is greater than 0, it indicates that the estimated optimal treatment enhances the individual's current kidney function. As illustrated in Figure 3.7 in the Appendix, the density distribution of these differences is concentrated at values larger than 0 across all stages. The boxplot also demonstrates that over 75% of the difference, values are larger than 0. These findings imply that a large proportion of patients in the training set can experience improved healthcare outcomes by following the estimated optimal treatments obtained through DSS. Furthermore, we observe that the variance of the difference is larger in the initial stages compared to later stages due to the overestimation phenomenon inherent in Q-learning (Hasselt, 2010; Sabry and Khalifa, 2019). Overestimation can lead to higher variance, particularly early in the learning process. The max operator in Q-learning tends to choose the highest (and often overestimated) values, further exacerbating the impact of random fluctuations in the initial stages, thereby increasing the variance. Furthermore, due to the design of our model, more information is utilized when modeling Y_1, incorporating data from stages 2, 3, and 4, compared to Y_3, which incorporates information from stages 3 and 4, or Y_2, which only incorporates data from stage 4. In contrast, Y_4 is modeled using data from the final stage alone. Exploiting this information leads to more room for improvement and, hence, to a larger variance of differences.

To control the false discovery rate (FDR) to select significant variables across the four

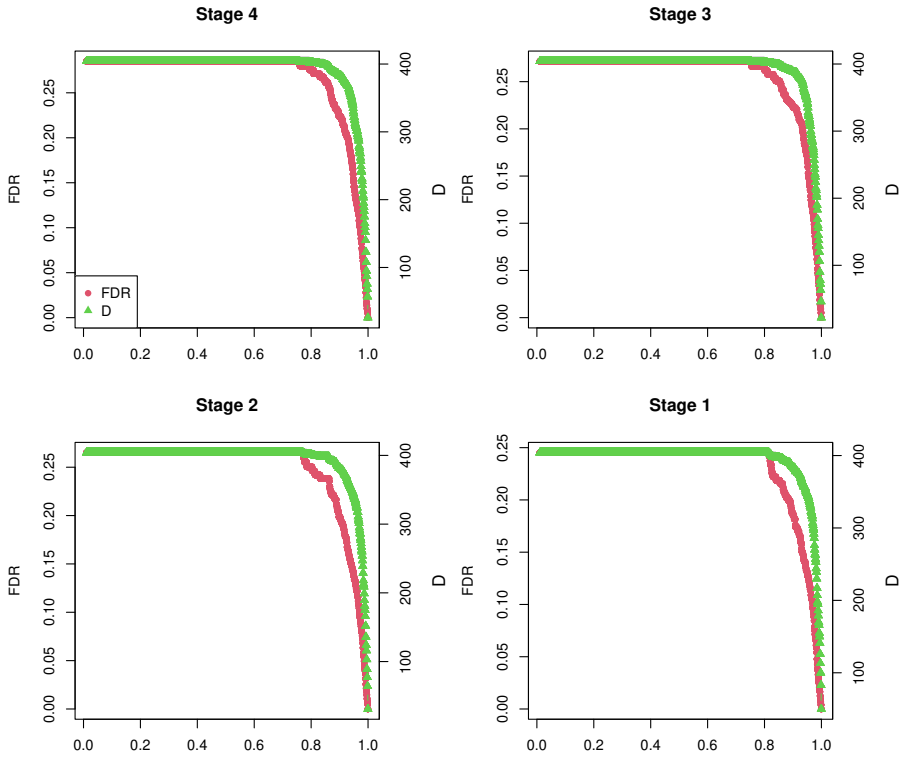

Figure 3.3: *The False Discovery Rate (FDR) and the corresponding number of significant variables (D) across 4 stages. The x-axis represents the cutoff values, while the y-axis on the right side indicates the number of significant variables, and the y-axis on the left side indicates the false discovery rate.*

	Cutoff	Number of Discoveries	FDR
Stage 4	0.951	19	0.01
Stage 3	0.956	17	0.01
Stage 2	0.947	24	0.01
Stage 1	0.950	27	0.01

Table 3.2: *Control FDR at 0.01, the corresponding number of significant variables and cutoff values across 4 stages.*

stages presented in section 3.3.1.2, we need to determine the cutoff value for each stage. The FDR and the number of significant variables decrease with increasing cutoff values shown in Figure 3.3. In Table 3.2, we control FDR at 0.01 and identify the corresponding number of significant variables along with their respective cutoff values across 4 stages. We select significant variables whose posterior inclusion probability exceeds the specified thresholds in Table 3.2. For instance, in the 4th stage, when we apply an FDR control of 0.01, the corresponding cutoff value is set at 0.951. Under this specific setting, we select variables with a posterior inclusion probability greater than 0.951. As a result, we identified 19 significant variables that meet this criterion in the 4th stage among 404 variables.

The specific significant variables chosen from the total 404 predictors are presented in Figure 3.5. Several variables, such as the interaction between RASi and FGF21_MESO_num (denoted as FGF21_MESO_num_1), exhibit statistical significance across all four stages of analysis. Furthermore, with interaction effects with treatments, specific biomarkers — HACR1_MESO, EGF_MESO, FGF21_MESO, NPHS1_MESO, and LCN2L_MESO—are consistently present across stages, although their relative importance fluctuates. This indicates their involvement throughout the process, albeit with varying degrees of influence. Notably, these biomarkers are associated with urinary functions, as identified in Jones et al. (2023). We present interaction plots and conditional interaction plots by stages to explain the treatment's contribution to HACR1_MESO and EGF_MESO, shown in Figure 3.4, revealing that the interactions differ significantly depending on the treatment and stage. We also show interaction plots between eGFR and FGF21_MESO, NPHS1_MESO, and LCN2L_MESO under treatment, shown in Figure 3.9 in the Appendix, displaying a similar pattern. On the contrary, certain variables, such as AHDT, demonstrate significance in only particular stages, with AHDT being significant solely in stage 4. Details of the selected variables are provided in Table 3.4 in the Appendix. The suffix 'norm' indicates that the variable is normalized by UCREA, while the suffix 'num' indicates that NA values below the minimum and above the maximum in the variable are replaced by 0.5 times the minimum and 1.5 times the maximum, respectively.

Figure 3.4: *Interaction plots and conditional interaction plots by stages between eGFR and HACR1_MESO, as well as EGF_MESO, under different treatments.*

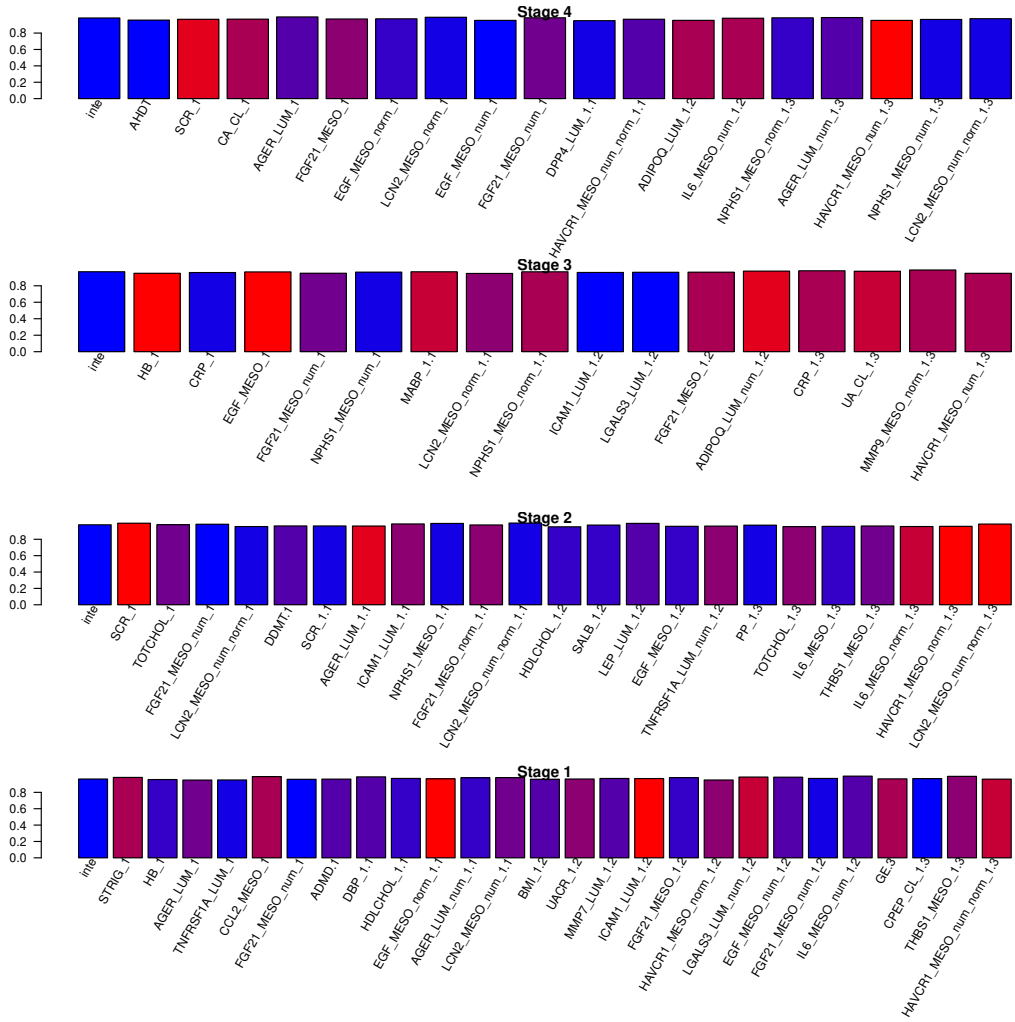

Figure 3.5: *Control FDR at 0.01, significant variables selected across 4 stages. The suffix '1' represents the interaction between RASi and clinical factors, while '1.1', '1.2', and '1.3' represent interactions between RASi combined with SGLT2i, MCRa, and GLP1a, respectively, with clinical factors. The bars' values represent the variable's posterior inclusion probability, while the color gradient, ranging from red to blue, indicates a progressive increase in the posterior inclusion probability.*

We use significant variables shown in Figure 3.5 to build prediction models following section 3.3.2. The predicted optimal treatment sequences for each patient in the test set are

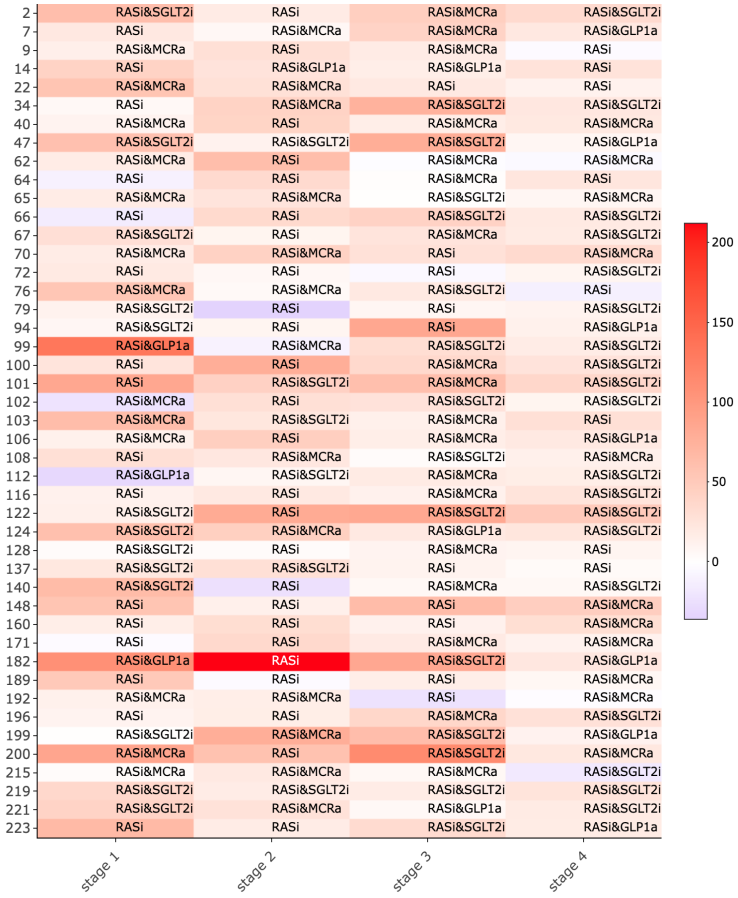

Figure 3.6: *The predicted optimal treatments of patients of the test set. The Y-axis denotes the patients' ID, and the text in each box means the predicted optimal treatment via DSS, while colors indicate the corresponding percentage increase of individuals in the test set. The percentage increase is calculated as the estimated optimal eGFR minus the observed eGFR, divided by the observed eGFR, and multiplied by 100.*

depicted in Figure 3.6. For instance, in stage 1, the DSS suggested drug combination for the first individual in test set is RASi & SGLT2i; in stage 2 is RASi, 3 is RASi & MCRa, 4 is RASi & SGLT2i. To visually represent the suggested drug combination's impact on enhancing an individual's kidney function, we employ color coding to signify the percentage increase between the estimated optimal eGFR and the observed eGFR within the test set. Observing the color, it becomes evident that the majority of boxes are shaded in red, indicating that the recommended drug combination via DSS has the potential to enhance

	Stage 1	Stage 2	Stage 3	Stage 4
Improved Percentage	88.89%	91.11%	93.33%	88.89%

Table 3.3: *The percentage of patients experiencing improved healthcare in the test set after receiving the predicted optimal treatment through DSS.*

the eGFR for the majority of individuals. Regarding the original dataset information, in the first stage, 95.07% of patients only received RASi under doctor medication, indicating a preference for this single drug regimen at the initial treatment phase. However, in the subsequent stages, the percentage decreased to 86.55% in the second stage and 85.65% in both the third and fourth stages. This trend suggests that, as treatment progresses, doctors are more likely to prescribe a combination of drugs, such as alongside the foundational RASi, integrating drugs SGLT2, MCRa, or GLP1a. This situation is also evident from Figure 3.6, where the DSS recommends more patients to receive drug combinations rather than relying solely on a single drug RASi. Additionally, the improvement situation for the test set, as shown in Figure 3.8 in the Appendix, reveals that the predicted optimal sequence treatment using DSS can enhance the healthcare outcomes for most patients in the test set, at least over 75% of patients.

Furthermore, a frequency table is provided (Table 3.3) to assess the improvement in patient care across the four stages for patients in the test set. In this table, a patient who gets improved healthcare can be considered one case if the difference between the predicted optimal eGFR and the observed eGFR is greater than 0. The results from Table 3.3 reveal that during the 1st stage, 88.89% of patients in the test set experience improved healthcare when following the estimated optimal treatment via DSS. This percentage slightly increases to 91.11% of patients during the 2nd stage. In the 3rd stage, 93.33% of patients demonstrate improved healthcare, and in the 4th stage, 88.89% of patients exhibit improved healthcare outcomes. Our model could lead to a larger predicted improvement at all stages.

3.5 Discussion

In this study, we extend the Bayesian Augmented Learning (BAL) method for DTRs proposed by in Chapter 2. This method addresses counterfactual variables in DTRs through a suitable augmentation. In BAL, a class of priors for variable selection, called dependent spike-and-slab, is used to consider the association between the significance of the same covariates across multiple stages.

The motivating issue that led to the extension of the BAL method is related to personalized medicine. More specifically, it concerns a prospective observational study in patients with type II diabetes mellitus for validation of biomarkers (PROVALID). The goal is to provide empirical support for the optimal sequences of drug combinations to prevent or mitigate kidney impairment progression at the individual level. The sampled dataset contains 223 patients across 4 visiting times. Each patient is described by 100 predictors. Four treatments can be administered to each patient: RASi only, SCLT2i & RASi, MCRa & RASi, and GLP1a & RASi. To improve understanding of treatment-patient dynamics, all interactions between patient status and each treatment were also included, bringing the final number of variables to 404 (including the intercept). The payoff is the estimated Glomerular Filtration Rate (eGFR).

Given this complex scenario, in addition to identifying the best treatment, another goal is to reduce dimensionality by selecting the most significant variables at each stage. We achieve this by borrowing information across multiple stages through the dependent spike-and-slab method proposed in BAL and leveraging the False Discovery Rate (FDR). At each step, FDR is used to select the most significant variables based on a pre-selected threshold. In this way, our approach involves the identification of significant molecular biomarkers and clinical parameters associated with drug effects on estimated Glomerular Filtration Rate (eGFR). Instead, to address the need to identify the optimal treatment, we introduce a prediction model to determine the optimal drug combination to improve kidney function for renal pathologies at the individual level with type II diabetes.

We find evidence that several variables, such as the interaction between RASi and FGF21_MESO_num (denoted as FGF21_MESO_num_1), exhibit statistical significance across all four stages of analysis. Furthermore, with interaction effects with treatments, specific biomarkers—HACR1_MESO, EGF_MESO, FGF21_MESO, NPHS1_MESO, and LCN2L_MESO—are consistently present across stages, although their relative importance fluctuates. This indicates their involvement throughout the process, albeit with varying degrees of influence. Notably, these biomarkers are associated with urinary functions, as identified in Jones et al. (2023). We present interaction plots and conditional interaction plots by stages to explain the treatment's contribution to the above mentioned biomarkers, revealing that the interactions differ significantly depending on the treatment and stage. Moreover, our results demonstrate that integrating drugs such as SGLT2 inhibitors, mineralocorticoid receptor antagonists (MCRa), and GLP-1 receptor agonists (GLP1a) alongside the foundational renin-angiotensin-system inhibitors (RASi) can offer additive benefits for safeguarding kidney function. Our prediction optimal treatment sequences analysis of the

percentage increase improvement situation indicates that the anticipated optimal treatment sequences, as predicted by the correlated spike-and-slab approach, have the potential to significantly enhance healthcare outcomes for a substantial majority, with a benefit extending to at least 75% of the patient population in each stage.

We are exploring various potential directions to expand upon our current research. Firstly, we may enhance our approach by sharing information not only within the same variable across multiple stages but also spanning both the same stage and multiple stages, which would consequently alter the prior inclusion probability of variables. Secondly, we can investigate the substitution of linear regression with nonlinear regression with correlated spike-and-slab priors to model the flexible relationship between patient outcomes and molecular biomarkers as well as clinical parameters, with the aim of developing a predictive treatment tree. Lastly, we presently control the False Discovery Rate (FDR) at a higher threshold to select significant molecular biomarkers and clinical parameters, therefore, there is potential in future endeavors to reduce FDR levels, thus minimizing Type I errors and increasing the statistical power of the test.

Bibliography

Abebe, S., Poli, I., Jones, R. D., and Slanzi, D. (2024). Learning optimal dynamic treatment regime from observational clinical data through reinforcement learning. *Machine Learning and Knowledge Extraction*, 6(3):1798–1817.

Bhattacharya, I., Ertefaie, A., Lynch, K. G., McKay, J. R., and Johnson, B. A. (2023). Nonparametric Bayesian Q-learning for optimization of dynamic treatment regimes in the presence of partial compliance. *Statistical Methods in Medical Research*, 32(9):1649–1663.

Bi, J., Borrotti, M., and Nipoti, B. (2024). Uncertainty quantification and multi-stage variable selection for personalized treatment regimes. *arXiv preprint arXiv:2411.02123*.

Cameron, J. S. (2006). The discovery of diabetic nephropathy: from small print to centre stage. *Journal of Nephrology*, 19:S75–87.

Chakraborty, B. and Murphy, S. A. (2014). Dynamic treatment regimes. *Annual Review of Statistics and Its Application*, 1:447–464.

Chib, S. and Greenberg, E. (1995). Understanding the metropolis-hastings algorithm. *The American Statistician*, 49(4):327–335.

Eder, S., Leierer, J., Kerschbaum, J., Rosivall, L., Wiecek, A., de Zeeuw, D., Mark, P. B., Heinze, G., Rossing, P., Heerspink, H. L., et al. (2018). A prospective cohort study in patients with type 2 diabetes mellitus for validation of biomarkers (PROVALID)–study design and baseline characteristics. *Kidney and Blood Pressure Research*, 43(1):181–190.

Eder, S., Leierer, J., Kerschbaum, J., Rosivall, L., Wiecek, A., de Zeeuw, D., Mark, P. B., Heinze, G., Rossing, P., Heerspink, H. L., and Mayer, G. (2019). Guidelines and clinical practice at the primary level of healthcare in patients with type 2 diabetes mellitus with and without kidney disease in five european countries. *Diabetes and Vascular Disease Research*, 16:47–56.

Emery, K. and Keich, U. (2019). Controlling the FDR in variable selection via multiple knockoffs. *arXiv preprint arXiv:1911.09442*.

Fan, Y., He, M., Su, L., and Zhou, X.-H. (2019). A smoothed Q-learning algorithm for estimating optimal dynamic treatment regimes. *Scandinavian Journal of Statistics*, 46(2):446–469.

Gregorich, M., Heinzel, A., Kammer, M., Meiselbach, H., Böger, C., Eckardt, K.-U., Mayer, G., Heinze, G., and Oberbauer, R. (2021). A prediction model for the decline in renal function in people with type 2 diabetes mellitus: study protocol. *Diagnostic and Prognostic Research*, 5:1–9.

Hasselt, H. (2010). Double Q-learning. *Advances in neural information processing systems*, 23.

Hu, J. and Du, Y. (2019). Managing chronic kidney disease in diabetes patients with the latest chemical therapies. *Expert Review of Clinical Pharmacology*, 12(1):53–60.

Ishwaran, H. and Rao, J. S. (2003). Detecting differentially expressed genes in microarrays using Bayesian model selection. *Journal of the American Statistical Association*, 98(462):438–455.

Ishwaran, H. and Rao, J. S. (2005). Spike and slab variable selection: frequentist and Bayesian strategies. *The Annals of Statistics*, 33:730–773.

Jones, R. D., Abebe, S., Distefano, V., Mayer, G., Poli, I., Silvestri, C., and Slanzi, D. (2023). Candidate composite biomarker to inform drug treatments for diabetic kidney disease. *Frontiers in Medicine*, 10:1271407.

Lizicarova, D., Krahulec, B., Hirnerova, E., Gaspar, L., and Celecova, Z. (2014). Risk factors in diabetic nephropathy progression at present. *Bratislavske lekarske listy*, 115(8):517–521.

Lyu, L., Cheng, Y., and Wahed, A. S. (2023). Imputation-based Q-learning for optimizing dynamic treatment regimes with right-censored survival outcome. *Biometrics*, 79(4):3676–3689.

Ma, T., Cai, H., Qi, Z., Shi, C., and Laber, E. B. (2023). Sequential knockoffs for variable selection in reinforcement learning. *arXiv preprint arXiv:2303.14281*.

Mayer, G., Eder, S., Rosivall, L., Voros, P., Heerspink, H. L., de Zeeuw, D., Czerwienska, B., Wiecek, A., Hillyard, D., Mark, P., et al. (2016). Baseline data from the multinational prospective cohort study for validation of biomarkers (PROVALID). *Nephrology Dialysis Transplantation*, 31:1482–1482.

Müller, P., Parmigiani, G., Robert, C., and Rousseau, J. (2004). Optimal sample size for multiple testing: the case of gene expression microarrays. *Journal of the American Statistical Association*, 99(468):990–1001.

Natesan, V. and Kim, S.-J. (2021). Diabetic nephropathy–a review of risk factors, progression, mechanism, and dietary management. *Biomolecules & Therapeutics*, 29(4):365.

Robins, J. M. (2004). Optimal structural nested models for optimal sequential decisions. In *Proceedings of the second seattle Symposium in Biostatistics*, pages 189–326. Springer.

Sabry, M. and Khalifa, A. (2019). On the reduction of variance and overestimation of deep Q-learning. *arXiv preprint arXiv:1910.05983*.

Sica, D. A. (2015). Mineralocorticoid receptor antagonists for treatment of hypertension and heart failure. *Methodist DeBakey Cardiovascular Journal*, 11(4):235.

Thomas, M. C., Brownlee, M., Susztak, K., Sharma, K., Jandeleit-Dahm, K. A., Zoungas, S., Rossing, P., Groop, P.-H., and Cooper, M. E. (2015). Diabetic kidney disease. *Nature reviews Disease primers*, 1(1):1–20.

Thöni, S., Keller, F., Denicolo, S., Eder, S., Rosivall, L., Wiecek, A. J., Mark, P., Rossing, P., Heerspink, H. L., and Mayer, G. (2021). Mo514 cardiorenal outcomes and mortality in patients with type 2 diabetes mellitus: a multinational prospective cohort study (PROVALID). *Nephrology Dialysis Transplantation*, 36:1–2.

Tsiatis, A. A., Davidian, M., Holloway, S. T., and Laber, E. B. (2019). *Dynamic treatment regimes: Statistical methods for precision medicine*. Chapman and Hall/CRC.

Wallace, M. P. and Moodie, E. E. (2015). Doubly-robust dynamic treatment regimen estimation via weighted least squares. *Biometrics*, 71(3):636–644.

Wang, J., Xiang, H., Lu, Y., Wu, T., and Ji, G. (2021). New progress in drugs treatment of diabetic kidney disease. *Biomedicine & Pharmacotherapy*, 141:111918.

Yamazaki, T., Mimura, I., Tanaka, T., and Nangaku, M. (2021). Treatment of diabetic kidney disease: current and future. *Diabetes & metabolism journal*, 45(1):11–26.

Yang, A., Shi, M., Lau, E. S., Wu, H., Zhang, X., Fan, B., Kong, A. P., Luk, A. O., Ma, R. C., Chan, J. C., et al. (2023). Clinical outcomes following discontinuation of renin-angiotensin-system inhibitors in patients with type 2 diabetes and advanced chronic kidney disease: A prospective cohort study. *EClinicalMedicine*, 55.

Appendices

3.A Posterior inference

With consideration of the specifications provided, as well as data augmentation, we rewrite the joint distribution $p(\boldsymbol{y}_4, \boldsymbol{y}^{\text{OPT}_3}, \boldsymbol{y}^{\text{OPT}_2}, \boldsymbol{y}^{\text{OPT}_1}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3, \boldsymbol{\theta}_4 \mid \tilde{\boldsymbol{z}}_1, \tilde{\boldsymbol{z}}_2, \tilde{\boldsymbol{z}}_3, \tilde{\boldsymbol{z}}_4, \boldsymbol{a}_1, \boldsymbol{a}_2, \boldsymbol{a}_3, \boldsymbol{a}_4)$ as

$$p(\boldsymbol{y}_1, \boldsymbol{y}_{21}, \boldsymbol{y}_{22}, \boldsymbol{y}_{23}, \boldsymbol{y}_{24}, \boldsymbol{y}_{31}, \boldsymbol{y}_{32}, \boldsymbol{y}_{33}, \boldsymbol{y}_{34}, \boldsymbol{y}_{41}, \boldsymbol{y}_{42}, \boldsymbol{y}_{43}, \boldsymbol{y}_{44}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3, \boldsymbol{\theta}_4 \mid \ldots)$$

Therefore, our focus is on the joint distribution with the augmentation incorporated into the spike-and-slab prior definition for $(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3, \boldsymbol{\theta}_4)$:

$$\begin{aligned} p(\boldsymbol{y}_1, \boldsymbol{y}_{21}, \boldsymbol{y}_{22}, \boldsymbol{y}_{23}, \boldsymbol{y}_{24}, \boldsymbol{y}_{31}, \boldsymbol{y}_{32}, \boldsymbol{y}_{33}, \boldsymbol{y}_{34}, \boldsymbol{y}_{41}, \boldsymbol{y}_{42}, \boldsymbol{y}_{43}, \boldsymbol{y}_{44}, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3, \boldsymbol{\theta}_4 \mid \ldots) = \\ \prod_{j=1}^{4} p(\boldsymbol{\theta}_j \mid \boldsymbol{\delta}_j, \boldsymbol{\psi}_j) p(\boldsymbol{\delta}_j \mid \boldsymbol{w}) p(\boldsymbol{\psi}_j) p(\boldsymbol{w}) \prod_{j=2}^{4} \prod_{t=1}^{4} p(\boldsymbol{y}_{jl} \mid \ldots) p(\boldsymbol{y}_1 \mid \ldots) \end{aligned} \quad (12)$$

where $\boldsymbol{y}_{jl}$ denotes the stage j pseudo-outcomes when taken treatment t.
The steps of the Gibbs sampler are summarized in Algorithm 3 and involve a sequential update of the parameters characterizing the four stage-specific regressions and the pseudo-outcomes.

Algorithm 3: Augmented Gibbs sampler

Initialize parameters $\varphi^{(0)} = (\boldsymbol{\theta}_j^{(0)}, \boldsymbol{\delta}_j^{(0)}, \boldsymbol{\psi}_j^{(0)}, \boldsymbol{w}_j^{(0)}, \sigma_j^{2(0)}, j = 1, 2, 3, 4)$;
Initialize pseudo-outcomes $(\boldsymbol{y}_{21}^{(0)}, \boldsymbol{y}_{22}^{(0)}, \boldsymbol{y}_{23}^{(0)}, \boldsymbol{y}_{24}^{(0)}, \ldots, \boldsymbol{y}_{41}^{(0)}, \boldsymbol{y}_{42}^{(0)}, \boldsymbol{y}_{43}^{(0)}, \boldsymbol{y}_{44}^{(0)})$;
For $r = 1$ to R
  Update parameters $\varphi^{(r)}$;
  Update pseudo-outcomes $(\boldsymbol{y}_{21}^{(r)}, \boldsymbol{y}_{22}^{(r)}, \boldsymbol{y}_{23}^{(r)}, \boldsymbol{y}_{24}^{(r)}, \ldots, \boldsymbol{y}_{41}^{(r)}, \boldsymbol{y}_{42}^{(r)}, \boldsymbol{y}_{43}^{(r)}, \boldsymbol{y}_{44}^{(r)})$;
  Compute $(A_{i1}^{\text{OPT}(r)}, A_{i2}^{\text{OPT}(r)}, A_{i3}^{\text{OPT}(r)}, A_{i4}^{\text{OPT}(r)})$, for $i = 1, \ldots, n$.
End
Estimate $\hat{a}_{ij}^{\text{OPT}}$ as mode in $\{A_{ij}^{\text{OPT}(R_0+1)}, \ldots, A_{ij}^{\text{OPT}(R)}\}$, for $j = 1, 2, 3, 4$ and R_0 is burn-in iterations.

The parameters are all reasonably simple since we have utilized the conjugate priors for all parameters, enabling the posterior distributions to be obtained with ease. Detailed information about the full conditional distributions is provided below.

Full conditional of $\boldsymbol{\theta}_1$:

$$\boldsymbol{\theta}_1 \mid \ldots \sim \mathrm{N}\left(b_1, B_1\right),$$

where $B_1^{-1} = \frac{1}{\sigma_1^2}\tilde{\boldsymbol{z}}_1^T\tilde{\boldsymbol{z}}_1 + D_1^{-1}$ and

$$b_1 = \frac{B_1\tilde{\boldsymbol{z}}_1^T Y_1^{\mathrm{OPT}}}{\sigma_1^2}$$

D_1 is a diagonal matrix with entries $r\left(\delta_{p1}\right)\psi_{p1},\ p = 1, 2, \cdots, p_1,$

$$Y^{\mathrm{OPT}_1} = Y_1 + Y_2^{\mathrm{OPT}} + Y_3^{\mathrm{OPT}} + Y_4^{\mathrm{OPT}}$$

where $Y_j^{\mathrm{OPT}} = \max(Y_{p1}, Y_{p2}, Y_{p3}, Y_{p4})$ and Y_{pj} represents the corresponding potential payoff.
Full conditional of $\boldsymbol{\theta}_j,$ for $j = 2, 3, 4$:

$$\boldsymbol{\theta}_j \mid \ldots \sim \mathrm{N}\left(b_j, B_j\right),$$

where $B_j^{-1} = \frac{1}{\sigma_j^2}Z_j^{*T}Z_j^* + D_j^{-1}$ and

$$b_j = \frac{B_j Z_j^{*T} Y_j^*}{\sigma_j^2}$$

D_j is a diagonal matrix with entries $r\left(\delta_{p_j j}\right)\psi_{p_j j},\ p = 1, 2, \cdots, p_j,$

$$Z^* = \begin{pmatrix} \tilde{\boldsymbol{z}}_{p1} \\ \tilde{\boldsymbol{z}}_{p2} \\ \tilde{\boldsymbol{z}}_{p3} \\ \tilde{\boldsymbol{z}}_{p4} \end{pmatrix} = \begin{pmatrix} [\boldsymbol{z}_p \times A_{p1}] \\ [\boldsymbol{z}_p \times A_{p2}] \\ [\boldsymbol{z}_p \times A_{p3}] \\ [\boldsymbol{z}_p \times A_{p4}] \end{pmatrix}, \quad Y_p^* = \begin{pmatrix} Y_{p1} \\ Y_{p2} \\ Y_{p3} \\ Y_{p4} \end{pmatrix},$$

where $A_{pj},\ j = 1, 2, 3, 4$ denote the 4 kinds of drugs combination, Y_{pj} represents the corresponding potential payoff.
Full conditional of σ_1^2:

$$\sigma_1^2 \mid \ldots \sim \textsc{Inv-Gamma}\left(\frac{n+1}{2}, \beta_1 + \frac{1}{2}\right),$$

where $\beta_1 = \frac{[Y_1 - (\boldsymbol{\theta}_1^T\tilde{\boldsymbol{z}}_1 - \sum_{j=2}^4 Y_j^{\mathrm{OPT}})]^\top [Y_1 - (\boldsymbol{\theta}_1^T\tilde{\boldsymbol{z}}_1 - \sum_{j=2}^4 Y_j^{\mathrm{OPT}})]}{2}$ is the scale parameter of the inverse gamma distribution.

Full conditional of σ_2^2:

$$\sigma_2^2 \mid \ldots \sim \text{Inv-Gamma}\left(2n + \frac{1}{2}, \beta_2 + \frac{1}{2}\right),$$

where $\beta_2 = \sum_{l=1}^{4} \frac{[Y_{2l} - (\boldsymbol{\theta}_2^\top \tilde{\boldsymbol{z}}_{2l} - Y_3^{\text{OPT}} - Y_4^{\text{OPT}})]^\top [Y_{2l} - (\boldsymbol{\theta}_2^\top \tilde{\boldsymbol{z}}_{2l} - Y_3^{\text{OPT}} - Y_4^{\text{OPT}})]}{2}$ is the scale parameter of the inverse gamma distribution.

Full conditional of σ_3^2:

$$\sigma_3^2 \mid \ldots \sim \text{Inv-Gamma}\left(2n + \frac{1}{2}, \beta_3 + \frac{1}{2}\right),$$

where $\beta_3 = \sum_{l=1}^{4} \frac{[Y_{3l} - (\boldsymbol{\theta}_3^\top \tilde{\boldsymbol{z}}_{3l} - Y_4^{\text{OPT}})]^\top [Y_{3l} - (\boldsymbol{\theta}_3^\top \tilde{\boldsymbol{z}}_{3l} - Y_4^{\text{OPT}})]}{2}$ is the scale parameter of the inverse gamma distribution.

Full conditional of σ_4^2:

$$\sigma_4^2 \mid \ldots \sim \text{Inv-Gamma}\left(2n + \frac{1}{2}, \beta_4 + \frac{1}{2}\right),$$

where $\beta_4 = \sum_{l=1}^{4} \frac{[Y_{4l} - \boldsymbol{\theta}_4^\top \tilde{\boldsymbol{z}}_{4l}]^\top [Y_{4l} - (\boldsymbol{\theta}_4^\top \tilde{\boldsymbol{z}}_{4l}]}{2}$ is the scale parameter of the inverse gamma distribution.

Full conditional for $\delta_{l_j j}$, for any $l_j = 1, \ldots, p_j,\ j = 1, 2, 3, 4$:

$$\delta_{l_j j} \mid \ldots \sim \text{Bern}\left(\frac{w_{l_j j} f_{\text{N}}(\theta_{l_j j}; 0, \psi_{l_j j})}{w_{l_j j} f_{\text{N}}(\theta_{l_j j}; 0, \psi_{l_j j}) + (1 - w_{l_j j}) f_{\text{N}}(\theta_{l_j j}; 0, r\psi_{l_j j})}\right).$$

Full conditional for $\psi_{l_j j}$, for any $l_j = 1, \ldots, p_j,\ j = 1, 2, 3, 4$:

$$\psi_{l_j j} \mid \ldots \sim \text{Inv-Gamma}\left(v + \frac{1}{2}, Q + \frac{\theta_{l_j j}^2}{2r(\delta_{l_j j})}\right).$$

Full conditional for $w_{l1} = w_{l2} = w_{l3} = w_{l4} = w_l$, for every $l = 1, \ldots, p$:

$$w_l \mid \ldots \sim \text{Beta}\left(a + \sum_{j=1}^{4} \delta_{lj}, b + 4 - \sum_{j=1}^{4} \delta_{lj}\right),$$

Similarly, for $w_{l_j j}$, for any $l_j = p+1, \ldots, p_j$, $j = 1, 2, 3, 4$:

$$w_{l_j j} \mid \ldots \sim \text{Beta}\left(a + \delta_{l_j j}, b + 1 - \delta_{l_j j}\right),$$

◇ *Full conditionals for the hyper-parameters a and b:*

$$a \mid \ldots \sim \left(\frac{\Gamma(a+b)}{\Gamma(a)}\right)^{\sum_{j=1}^{4} p_j - 4p} \left(\prod_{l_1=1}^{p} w_{l_1 1} \prod_{j=1}^{4} \prod_{l_j=p+1}^{p_j} w_{l_j j}\right)^{(a-1)} p(a),$$

$$b \mid \ldots \sim \left(\frac{\Gamma(a+b)}{\Gamma(b)}\right)^{\sum_{j=1}^{4} p_j - 4p} \left(\prod_{l_1=1}^{p} (1 - w_{l_1 1}) \prod_{j=1}^{4} \prod_{l_j=p+1}^{p_j} (1 - w_{l_j j})\right)^{(b-1)} p(b).$$

The prior distributions of a and b, denoted as $p(a)$ and $p(b)$, are specified as inverse gamma distributions with shape and scale parameters both set to 1. To obtain samples from the full conditional distributions of a and b, a random-walk Metropolis-Hastings algorithm is employed (Chib and Greenberg, 1995).

3.B Additional results on the case study

We report the differences between the estimated optimal eGFR and the observed eGFR for both the training and test sets (see Figure 3.7 and Figure 3.8, respectively). Additionally, we provide a detailed description of selected variables in the main manuscript (see Table 3.4). Furthermore, we display interaction plots to show how the effects of specific variables (NPHS1_MESO, FGF21_MESO, as well as LCN2L_MESO) on eGFR differs across treatment groups (see Figure 3.9).

Variables acronyms	Variable description
ADIPOQ_LUM	Adiponectin, C1Q And Collagen Domain Containing concentration in serum
AGER	Advanced Glycosylation End-Product Specific Receptor concentration
AHDT	Age at HT diagnosis
BMI	Body Mass Index
CA_CL	Calcium concentration in serum
CCL2_MESO	C-C motif chemokine ligand 2 concentration in urine
CPEP_CL	C-peptide concentration in serum
DBP	Diastolic BP
DDMT	Duration of DM2 pharmacological treatment at first visit in PROVALID
DPP4_LUM	Dipeptidyl Peptidase 4 concentration in serum
EGF_MESO	epidermal growth factor concentration in urine
FGF21_MESO	fibroblast growth factor 21 concentration in urine
GE	Gender
HAVCR1_MESO	hepatitis A virus cellular receptor 1 concentration in urine
HB	Hemoglobin
HDLCHOL	Serum cholesterol (HDL)
ICAM1_LUM	Intercellular Adhesion Molecule 1 concentration in serum
IL6_MESO	Interleukin 6 concentration in urine
LCN2_MESO	lipocalin-2 concentration in urine
LEP_LUM	Leptin concentration in serum
LGALS3_LUM	Galectin 3 concentration in serum
MABP	Mean arterial blood pressure
MMP7_LUM	Matrix Metallopeptidase 7 concentration in serum
MMP9_MESO	matrix metallopeptidase 9 concentration in urine
NPHS1	NPHS1 adhesion molecule, nephrin concentration
NPHS1_MESO	NPHS1 adhesion molecule, nephrin concentration in urine
PP	Pulse pressure
SALB	Serum albumin
SCR	Serum creatinine
STRIG	Serum triglycerides
THBS1_MESO	thrombospondin 1 concentration in urine
TNFRSF1A_LUM	TNF Receptor Superfamily Member 1A concentration in serum
TOTCHOL	Serum cholesterol (total)
UA_CL	Uric Acid concentration in serum

Table 3.4: *Covariates after variable selection using DSS in the PROVALID study.*

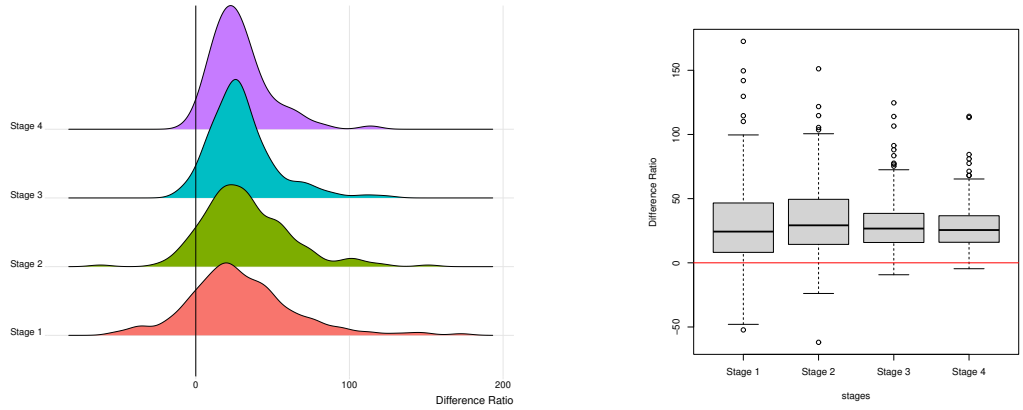

(a) *Density distribution across 4 stages* ***(b)*** *Boxplot across 4 stages*

Figure 3.7: *Estimated improvement in patients from the training set after DSS treatments. The difference ratio is calculated as the estimated optimal eGFR minus the observed eGFR, divided by the observed eGFR, and multiplied by 100.*

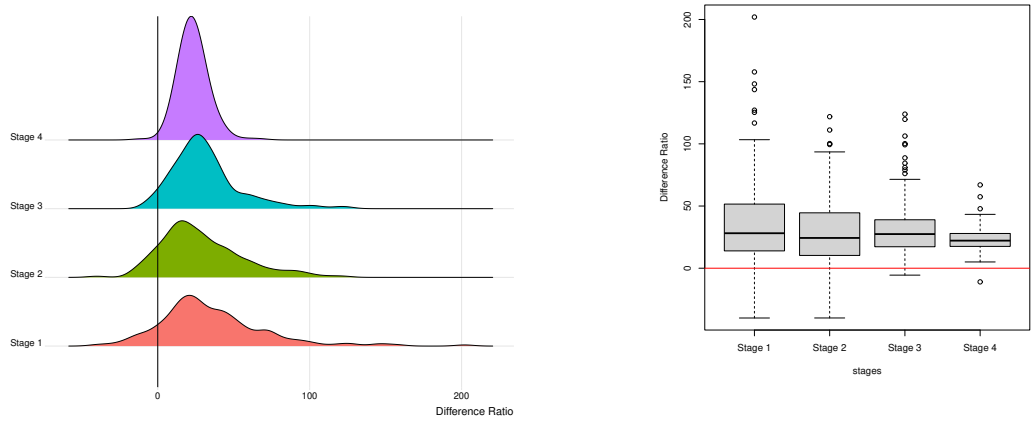

(a) *Density distribution across 4 stages* ***(b)*** *Boxplot across 4 stages*

Figure 3.8: *Predicted improvement in patients from the test set after DSS treatments. The difference ratio is calculated as the predicted optimal eGFR minus the observed eGFR, divided by the observed eGFR, and multiplied by 100.*

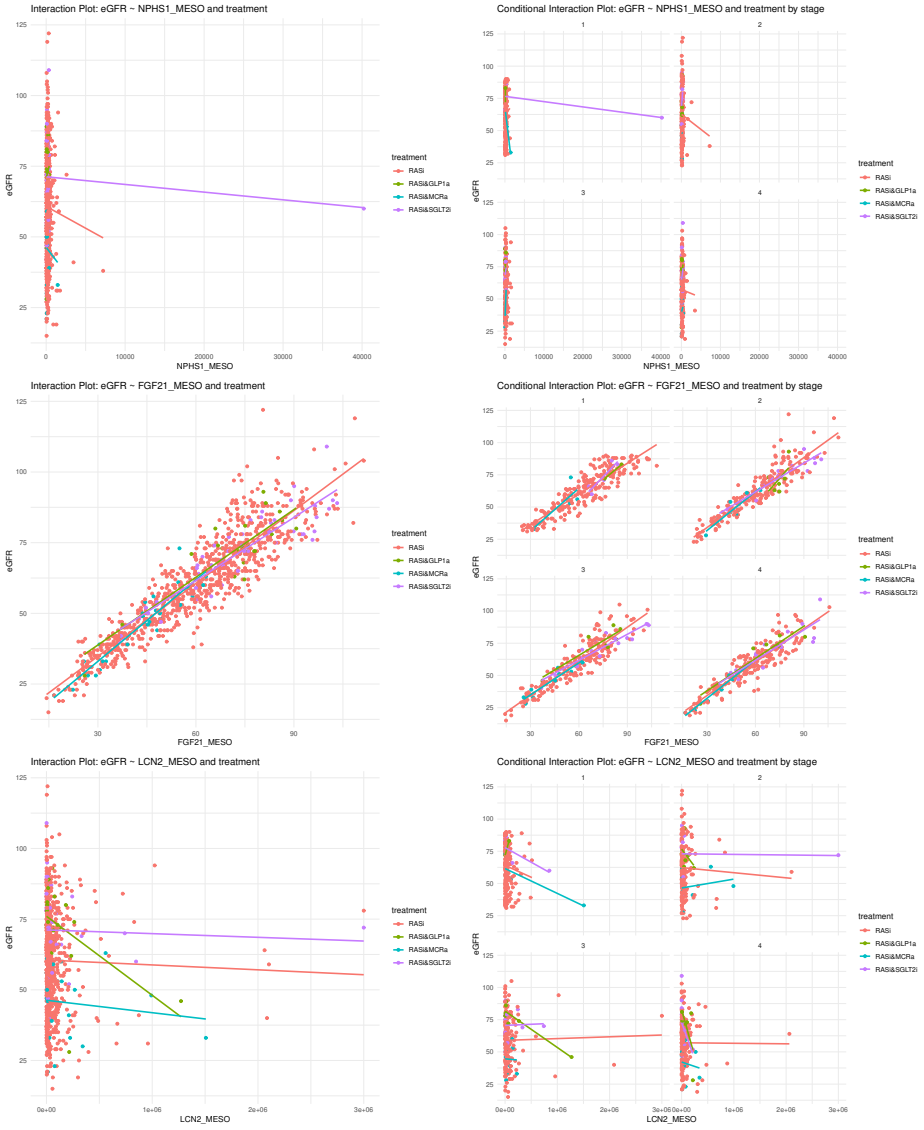

Figure 3.9: *Interaction plots and conditional interaction plots by stages between eGFR and NPHS1_MESO, FGF21_MESO, as well as LCN2L_MESO, under different treatments.*

Chapter 4

Causal structure learning with mixed functional Bayesian networks

4.1 Introduction

Multivariate functional data refers to a type of data where multiple related functions are observed continuously over a shared domain (like time, space, or frequency), meaning that each variable in the dataset is represented as a function rather than a single value, and these functions are observed simultaneously, capturing the dynamic relationships between them across the domain. Such multivariate functional data come into prevalent in various fields, including biomedical research (Wei and Li (2008); Ieva and Paganoni (2013); Zhang et al. (2023); Moindjié et al. (2024)), environmental science (Bouveyron et al. (2022); Carroll and Müller (2023)), economics (Kowal et al. (2017); Zhu (2023)), and sociology science (López-Pintado et al. (2014); Li et al. (2022); Golovkine et al. (2022)). A key objective is to learn the causal structure among multivariate functional data. This can help in understanding the underlying mechanisms driving the observed functional data, thereby enabling researchers and practitioners to improve predictive models and design more effective interventions.

While most existing research focuses on causal discovery for continuous multivariate functional data, real-world data often present challenges by including a mix of continuous and binary functions. The challenges of mixed data stem from the complexity and computational demands of modeling functional binary data. Unlike functional continuous data, binary data provide limited information and are influenced by temporal dependencies, nonlinearity, noise, and data sparsity. Consequently, performing causal discovery on binary functions requires specialized approaches to effectively capture their unique structures and behaviors (Van Der Linde, 2009; Zhong et al., 2023). Even in cross-sectional

studies, Glymour et al. (2019) also mentioned that in discrete case, determining the causal direction from observed data is particularly complex, especially when the variables have low cardinality due to limited information available.

Functional undirected graphical models were initially developed for continuous multivariate functional data (Zhu et al., 2016; Li and Solea, 2018; Qiao et al., 2019; Solea and Li, 2022). Recently, functional directed graphical models have also been advanced for continuous multivariate functional data. Lee and Li (2022) employed functional structural equation models to estimate directional relations from multivariate functional data. Zhou et al. (2023) developed a Bayesian network model that encodes conditional independencies and causal structure using a directed acyclic graph for functional data. Later, Roy et al. (2024) extended the Bayesian network model from a directed acyclic graph to a directed cyclic graph for functional data. Yang et al. (2024) proposed the functional linear Non-Gaussian acyclic model for functional data. The aforementioned literature predominantly focuses on continuous cases. Rogovchenko et al. (2024) newly introduced a scalar-function causal model to analyze the relationship between a continuous functional variable and a binary scalar variable, unfortunately, rather than a binary functional variable. To date, several studies have addressed causal discovery for mixed data, yet they primarily focus on scale continuous and binary variables (Tsagris et al., 2018; Raghu et al., 2018; Wenjuan et al., 2018; Yamayoshi et al., 2020; Zeng et al., 2022; Cai et al., 2022). To resolve the challenges of learning causal structures from mixed functional data, we propose mixed functional Bayesian networks, where each observed function is derived from a continuous latent function through a transformation process involving a link function.

Mixed functional linear non-Gaussian Bayesian networks (*mixed-FLiNG-BN*) proposed in this chapter is a mixed probabilistic graphical model that represents a set of functions and their conditional dependencies via a directed acyclic graph (DAG). By employing orthogonal basis expansions, we transform mixed random functions from the functional space to an isometrically isomorphic space represented by basis coefficients, as depicted in Figure (4.1). This approach allows for the convenient construction of Markov distributions and hyper Markov laws (Zhu et al., 2016). Furthermore, we reframe the proposed mixed Bayesian network, originally defined in the functional space, into an equivalent network expressed in the basis coefficient space through this expansion method. In what follows, we will provide a deeper exploration of mixed-FLiNG-BN. Firstly, each observed function is generated from a latent function through a one-to-one correspondence. For continuous functions, this correspondence means that the observed function is exactly the latent function. For binary functions, it is obtained by dichotomizing the latent function at 0. Subsequently, we

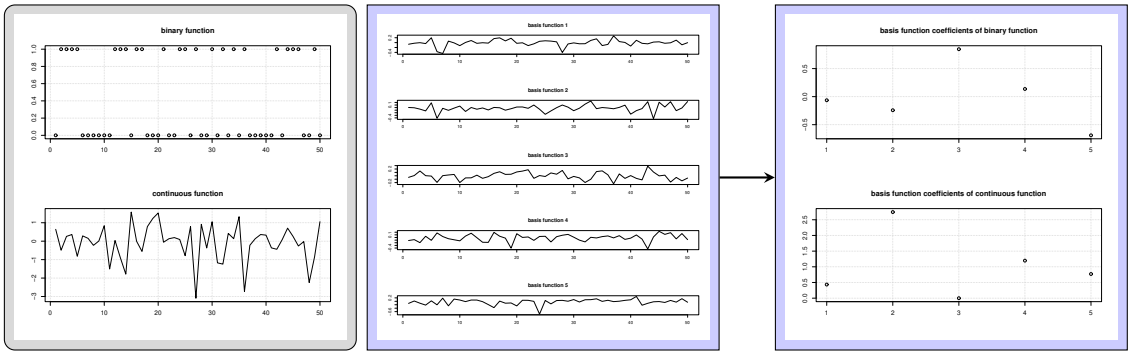

Figure 4.1: *Illustration of functions and their corresponding basis coefficients via basis expansion. The functional data consists of 50 time points, and we use 5 basis functions to perform an expansion.*

represent these latent functions using orthonormal basis expansions, which provide both interpretability and computational efficiency (Van den Hof et al., 2000). Secondly, we encode the unknown causal structure using a structural equation model on the basis coefficients of latent functions, which are then transformed back to the original random functions. To ensure causal identification of mixed-FLiNG-BN, we assume that our latent functions are generated from a mixture of Gaussians. We also present a theoretical proof that the causal DAG of the proposed model is uniquely identifiable under mild conditions. Additionally, we employ a Bayesian hierarchical framework with appropriate prior distributions on mixed-FLiNG-BN to quantify uncertainty with limited data, and posterior inferences are conducted using Markov Chain Monte Carlo (MCMC) methods.

The chapter is structured as follows. In Section 4.2, we introduce the proposed model and causal identification. Section 4.3 presents the model inference within a Bayesian framework. In Section 4.4, we demonstrate the finite sample performance of our method through extensive simulation studies. The chapter concludes with a discussion in Section 4.5. We provide the technical proof, detailed MCMC, and additional case study results in the Appendix.

4.2 Method and causal identification

4.2.1 Notation and background

Let $[p] = \{1, 2, \ldots, p\}$ denote the set of integers from 1 to p for any positive integer p. We then begin by introducing some fundamental notations for DAGs and mixed BNs. A DAG

$\mathcal{G} = (\boldsymbol{V}, \boldsymbol{E})$ comprises of a set of nodes $\boldsymbol{V} = [p]$ and a set of directed edges $\boldsymbol{E} \subset \boldsymbol{V} \times \boldsymbol{V}$ representing the causal relationships between the nodes $\boldsymbol{V}$ without cycles. In the context of a directed edge $j \to k \in \boldsymbol{E}$, where $j, k \in \boldsymbol{V}$, node j is called a parent of k, and conversely, the node k is termed a child of j. The set of parents of node j in $\mathcal{G}$ is denoted by $pa(j)$. We initially consider a BN for mixed random variables, then proceed to a BN for mixed random functions. Let $\boldsymbol{X} = (X_1, \ldots, X_p)$ denote a set of p random variables of which some are binary and some are continuous. So in this chapter, a DAG $\mathcal{G} = (\boldsymbol{V}, \boldsymbol{E})$ with variables $\boldsymbol{V} = \boldsymbol{\Delta} \bigcup \boldsymbol{\Gamma}$, where $\boldsymbol{\Delta}$ and $\boldsymbol{\Gamma}$ are the sets of binary and continuous variables, respectively. A BN for $\boldsymbol{X}$ is $\mathcal{B} = (\mathcal{G}, P)$ with the joint distribution $p(\cdot)$ factorizing over $\mathcal{G}$ as follows:

$$P(\boldsymbol{X}) = \prod_{j \in \boldsymbol{\Delta}} P_b\left(X_j = x_j \mid \boldsymbol{X}_{pa(j)}\right) \prod_{j \in \boldsymbol{\Gamma}} P_c\left(X_j \mid \boldsymbol{X}_{pa(j)}\right),$$

where P_b and P_c denote the conditional distribution of binary and continuous variables given its corresponding parents, respectively.

Next, we consider a BN for mixed random functions. Denote the space of square-integrable functions on a domain $\mathscr{D}$ with respect to a measure μ as $L^2(\mathscr{D}, \mu)$, defined by: $L^2(\mathscr{D}, \mu) = \{h : \int_{\mathscr{D}} |h(\omega)|^2 \, d\mu(\omega) < \infty\}$ Here, h is a function defined on $\mathscr{D}$ such that the integral of $|h(\omega)|^2$ with respect to μ over $\mathscr{D}$ is finite where the domain $\mathscr{D}$ is a subset of $\mathbb{R}$ and the measure μ is typically the Lebesgue measure. Let $\boldsymbol{W} = (W_1, W_2, \ldots, W_p) \in L^2(\mathscr{D}_1) \times \cdots \times L^2(\mathscr{D}_p)$ denote a set of p random functions. The domain of $\boldsymbol{W}$ is denoted by $\bigsqcup_{j=1}^{p} \mathscr{D}_j$, where $\bigsqcup$ denotes the disjoint union. For $\boldsymbol{V} = [p]$ and $\boldsymbol{A} \subset \boldsymbol{V}$, denote by $\boldsymbol{W}_A$ the subset of $\boldsymbol{W}$ with domain $\mathscr{D}_A = \bigsqcup_{j \in A} \mathscr{D}_j$. Let $\left(L^2(\mathscr{D}_j), \mathcal{B}\left(L^2(\mathscr{D}_j)\right), P\right)$ be a probability space, where $\mathcal{B}\left(L^2(\mathscr{D}_j)\right)$ is the Borel σ-algebra on $L^2(\mathscr{D}_j)$. When the random functions are continuous, the measurable set in $L^2(\mathscr{D}_j)$ is denoted by D_j. When the random functions are binary, the possible values for $L^2(\mathscr{D}_j)$ is $\{0, 1\}$. A DAG $\mathcal{G} = (\boldsymbol{V}, \boldsymbol{E})$ with functions $\boldsymbol{V} = \boldsymbol{\Delta} \bigcup \boldsymbol{\Gamma}$, where $\boldsymbol{\Delta}$ and $\boldsymbol{\Gamma}$ are the sets of binary and continuous functions, respectively. In the following, we give the definition of a mixed functional BN for $\boldsymbol{W}$.

Definition 4.2.1 (Mixed Functional Bayesian Networks)**.** Consider the mixed random functions $\boldsymbol{W} = (W_1, W_2, \cdots, W_p)$. We say $\mathcal{B} = (\mathcal{G}, P)$ is a functional Bayesian network for

a set of mixed random functions $\boldsymbol{W}$ if P factorizes with respect to DAG $\mathcal{G}$,

$$
\begin{aligned}
&P\left(\bigcap_{j\in\boldsymbol{\Delta}}\{W_j = w_j\} \cap \bigcap_{j\in\boldsymbol{\Gamma}}\{W_j \in D_j\}\right) = \\
&\prod_{j\in\boldsymbol{\Delta}} P_b\left(W_j = w_j \mid \boldsymbol{W}_{pa(j)} = \boldsymbol{w}_{pa(j)}, \boldsymbol{W}_{pa(j)} \in D_{pa(j)}\right) \\
&\times \prod_{j\in\boldsymbol{\Gamma}} P_c\left(W_j \in D_j \mid \boldsymbol{W}_{pa(j)} = \boldsymbol{w}_{pa(j)}, \boldsymbol{W}_{pa(j)} \in D_{pa(j)}\right),
\end{aligned}
$$

for $w_j \in \{0,1\}$, $\forall j \in \boldsymbol{\Delta}$ and any other measurable sets $D_j \subseteq L^2(\mathcal{D}_j)$, $j \in \boldsymbol{\Gamma}$, where $\boldsymbol{\Delta}$ and $\boldsymbol{\Gamma}$ are the sets of binary and continuous functions, respectively. P_b and P_c denotes the conditional probability distribution of binary and continuous functions given its corresponding parents, respectively.

In the subsequent section, we introduce latent continuous space functions W^* corresponding to each binary function W. These latent functions allow us to better capture the underlying continuous nature of the data while maintaining the binary representation in W. The causal relationships encoded in the DAG $\mathcal{G}$ for W^* directly map to those of W, ensuring that the structural dependencies and causal inference remain consistent across both representations. This approach provides a more flexible and comprehensive framework for modeling complex systems where binary and continuous variables are intertwined.

4.2.2 Mixed functional linear non-Gaussian Bayesian networks

For modeling convenience, we use an orthonormal basis expansion of random functions to redefine the functional Bayesian Network (BN) in the space of basis coefficients, leveraging the orthogonality and normalization properties of the orthonormal basis. Let $\{\phi_{jk}\}_{k=1}^{\infty}$ be a sequence of orthonormal basis functions for $L^2(\mathcal{D}_j)$, and let $e_j(\cdot)$ represent measurement error. We expand $W_j(\cdot)$ from the perspective of data augmentation, i.e., latent variables, as described in Albert and Chib (1993). Specifically, the model is defined as follows:

$$
W_j(\cdot) = \begin{cases} \mathbf{I}\left(W_j^*(\cdot) > 0\right), & j \in \boldsymbol{\Delta} \\ W_j^*(\cdot), & j \in \boldsymbol{\Gamma} \end{cases}, \tag{4.1}
$$

where

$$
W_j^*(\cdot) = \sum_{k=1}^{\infty} Z_{jk}\phi_{jk}(\cdot) + e_j(\cdot),\ \forall j \in [p], \tag{4.2}
$$

and $e_j(\cdot) \sim \mathrm{N}(0, \sigma_j^2)$. Here, $\boldsymbol{\Delta}$ and $\boldsymbol{\Gamma}$ represent the sets of binary and continuous random functions, respectively.

The resulting coefficient sequence $\boldsymbol{Z}_j = (Z_{jk})_{k=1,\cdots,\infty}$ lies in the space of square-summable sequences $\mathscr{L}_j^2 = \{h_j : \sum_{k=1}^{\infty} |h_{jk}|^2 < \infty\}$. We call our proposed model, **mixed-FLiNG-BN**. The mixed-FLiNG-BN model assumes that the random variable $\boldsymbol{Z}$ follows a linear Structural Equation Model (SEM):

$$\boldsymbol{Z}_j = \sum_{\ell=1}^{p} \boldsymbol{B}_{j\ell} \boldsymbol{Z}_\ell + \boldsymbol{\epsilon}_j, \ \forall j \in [p], \tag{4.3}$$

where $\boldsymbol{\epsilon}_j$ is an infinite-dimensional exogenous vector, $\boldsymbol{B}_{j\ell} = (B_{j\ell}(k_j, k_\ell))_{k_j=1,k_\ell=1}^{\infty,\infty}$ is an infinite-dimensional direct causal effect matrix from $\boldsymbol{Z}_\ell$ to $\boldsymbol{Z}_j$. The notation $\ell \to j$ indicates that $\boldsymbol{Z}_\ell$ is a direct cause of $\boldsymbol{Z}_j$ if there exist k_ℓ and k_j such that $B_{j\ell}(k_j, k_\ell) \neq 0$. Both the causal effects and the causal graph are considered unknown and will be inferred from observational data. Additionally, we assume that $\boldsymbol{\epsilon}_1, \cdots, \boldsymbol{\epsilon}_p$ are independent, thereby eliminating the potential interference of latent confounders. One crucial element of the mixed-FLiNG-BN for discovering causality is the specification of the probability distribution of the exogenous variable $\boldsymbol{\epsilon}_j = (\epsilon_{jk})_{k=1}^{\infty}$ in (4.3). The exogenous variables introduce noise into the model, and their distribution affects how well the model can capture the underlying causal relationships. The constraint that enables causal identification, as we will demonstrate in the following section, is the specification of a non-Gaussian distribution for the exogenous variables $\boldsymbol{\epsilon}_j$. Specifically, we assume that ϵ_{jk} follows a finite mixture of Gaussian distributions. This can be mathematically expressed as $\epsilon_{jk} \sim \sum_{m=1}^{M_{jk}} \pi_{jkm} \mathrm{N}(\mu_{jkm}, \tau_{jkm})$, where each component $\mathrm{N}(\mu_{jkm}, \tau_{jkm})$ represents a Gaussian distribution with mean μ_{jkm} and variance τ_{jkm}. The coefficients π_{jkm} are the mixing proportions for each component, and they sum to 1. The parameter M_{jk} denotes the number of mixture components and is greater than 1, ensuring that the distribution is indeed a mixture rather than a single Gaussian distribution.

Finite-sample inference is necessary because, in practical applications, we often encounter issues such as computational efficiency, noise reduction, and overfitting. Handling infinite-dimensional data directly is computationally prohibitive, so truncating the basis functions significantly reduces complexity and makes computations feasible. To achieve finite-sample inference, we reduce the dimensionality of the functional data by truncating the orthonormal basis functions at level K, resulting in $\{\phi_{jk}\}_{k=1}^{K}$, as commonly practiced in existing functional data analysis literature (Yao et al., 2005; Hyndman and Ullah, 2007).

As a result, the latent continuous function W_j^* in (4.2) can be expressed as

$$W_j^*(\cdot) = \sum_{k=1}^{K} Z_{jk}\phi_{jk}(\cdot) + e_j(\cdot),\ \forall j \in [p], \tag{4.4}$$

where $e_j(\cdot)$ denotes independent white noises. By truncating the series at K, we focus on the most significant components of the functional data, making the model more computationally efficient while still capturing the essential features of the data.

4.2.3 Causal identifiability

In this section, we demonstrate that the graph structure of the proposed mixed-FLiNG-BN model is identifiable for mixed functional data under several causal assumptions. We begin by defining causal identifiability and outlining our assumptions.

Definition 4.2.2 (Causal Identifiability)**.** Assuming $\boldsymbol{W}$ represents the observed noisy random function within the mixed-FliNG-BN framework $\mathcal{B} = (\mathcal{G}, P)$, we say that the causal DAG of mixed-FLiNG-BN is identifiable from $\boldsymbol{W}$ if there does not exist another BN $\mathcal{B}' = (\mathcal{G}', P')$ where $\mathcal{G} \neq \mathcal{G}'$ such that the induced distribution on $\boldsymbol{W}$, P'_W, is equivalent to P_W, denoted as, $P_W(\boldsymbol{W}) \equiv P'_W(\boldsymbol{W})$.

Put simply, a causal graph is identifiable if no other graph induces the same joint distributions. Afterwards, we'll classify and discuss two assumptions to establish the causal identifiability of the proposed model.

Assumption 4.2.3 (Non-Gaussian)**.** The disturbances ϵ_{jk} are from independent mixture of Gaussian distributions with the number of mixture components is greater than 1, i.e., $\epsilon_{jk} \sim \sum_{m=1}^{M_{jk}} \pi_{jkm}\mathrm{N}(\mu_{jkm}, \tau_{jkm})$, where M_{jk} is the number of mixture components and $M_{jk} > 1$.

The assumption of non-Gaussian on the disturbances ϵ_{jk} induces model identifiability in the linear Structural Equation Modeling (SEM) framework (Shimizu et al., 2006; Shimizu, 2014) for continuous random variables. It is necessary also for binary variables (Kong et al., 2022) since with Gaussian distribution, the model is not identifiable in the probit case. When a continuous random function $\boldsymbol{W}$ follows to a Gaussian structural model, the model loses its identifiability. In the context of probit regression, the binary random function $\boldsymbol{W}$ can be derived from dichotomizing a continuous latent random function that conforms to a Gaussian distribution. However, there is no evidence to suggest that dichotomization

enhances identifiability. Therefore, it is also necessary for the disturbances to follow a non-Gaussian distribution in the case of binary random functions. For simplicity, $\mathrm{N}(\mu_m, \tau_m)$ can be considered as elements of a mixture distribution. However, in the more general case, we consider $\mathrm{N}(\mu_{jkm}, \tau_{jkm})$.

Assumption 4.2.4 (Non-Injective)**.** For a pair of mixed functions $(\boldsymbol{W}_1, \boldsymbol{W}_2)$, each measured on n time points $\boldsymbol{W}_1 = (W_1(t_1), \ldots, W_1(t_n))^T \in \{0,1\}^n$, $\boldsymbol{W}_2 = (W_2(t_1), \ldots, W_2(t_n))^T \in \mathcal{R}^n$, there exists values $\boldsymbol{w}_2', \boldsymbol{w}_2''$ of $\boldsymbol{W}_2$ such that $\boldsymbol{w}_2' \neq \boldsymbol{w}_2''$, while

$$P(\boldsymbol{W}_1 = \boldsymbol{w}_1 \mid \boldsymbol{W}_2 = \boldsymbol{w}_2') = P(\boldsymbol{W}_1 = \boldsymbol{w}_1 \mid \boldsymbol{W}_2 = \boldsymbol{w}_2'').$$

This assumption implies us that even though $\boldsymbol{w}_2'$ and $\boldsymbol{w}_2''$ are different, the likelihood of observing the binary $\boldsymbol{w}_1$ in $\boldsymbol{W}_1$ does not change, we call the function between $\boldsymbol{w}_1$ and $\boldsymbol{w}_2$ are non-injective. Zhang and Hyvarinen (2012); Peters et al. (2014); Tagasovska et al. (2018) stated that identifiable additive noise models can be obtained when the relationship between two variables is not injective. Non-injective mappings are commonly found in empirical data, reflecting the complexity and multifaceted nature of real-world phenomena where multiple distinct causes can produce similar observable outcomes (Peters et al., 2017).

Theorem 4.2.5 (Causal Identifiability)**.** *Under Assumptions 4.2.3-4.2.4, the causal DAG of mixed-FLiNG-BN is identifiable.*

The proof addresses the bivariate identifiability of mixed-FLiNG-BN in the continuous-binary and continuous-continuous scenarios, respectively. The binary-binary case is not considered due to the absence of relevant research at the time of submission. While a formal proof is available in Appendix 4.A, we provide a brief intuition here. For the continuous-continuous case, the model simplifies to a form of functional Bayesian networks, whose identifiability has been thoroughly demonstrated in previous studies (Zhou et al., 2023). For the mixed continuous-binary case, identifiability is established based on the properties of real analytic functions. In the following, we present example 1 to illustrate the results of continuous-binary causal identification.

Example 1: Consider a true functional causal graph $1 \to 2$ and the corresponding data generating model

$$W_1 = Z_1,$$

with $Z_1 = \epsilon_1 \sim 0.5\mathrm{N}(0, 0.5) + 0.5\mathrm{N}(0.5, 1)$ and

$$W_2 = \begin{cases} 1 & : Z_2 \geq 0 \\ 0 & : \text{else} \end{cases}$$

$Z_2 = Z_1 + \epsilon_2$ with $\epsilon_2 \sim 0.5\mathrm{N}(0, 0.5) + 0.5\mathrm{N}(0, 1)$.

Example 1 meets the criteria of both non-Gaussianity and non-injectivity. The non-Gaussian nature is clearly evident. We will now provide a detailed explanation of why Example 1 also satisfies the non-injective condition. If we want to find w_1' and w_1'' such that $w_1' \neq w_1''$, making

$$P(W_2 = w_2 \mid W_1 = w_1') = P(W_2 = w_2 \mid W_1 = w_1'').$$

Here we set $w_2 = 1$, then

$$\begin{aligned} & P(W_2 = 1 \mid W_1 = w_1') = P(W_2 = 1 \mid W_1 = w_1'') \\ \Rightarrow & P(Z_2 \geq 0 \mid Z_1 = z_1') = P(Z_2 \geq 0 \mid Z_1 = z_1'') \\ \Rightarrow & P(Z_1' + \epsilon_2 \geq 0 \mid Z_1 = z_1') = P(Z_1'' + \epsilon_2 \geq 0 \mid Z_1 = z_1'') \\ \Rightarrow & P(\epsilon_2 \geq -z_1') = P(\epsilon_2 \geq -z_1'') \end{aligned} \tag{4.5}$$

We plot the density and cumulative distribution function (CDF) of ϵ_2, as shown in Figure (4.2). The red segment in the right panel of Figure (4.2) indicates that the CDF of ϵ_2 remains the same within the specified range. As a result, we can identify z_1' and z_1'' in the neighborhood of -1,000, such that $P(\epsilon_2 \geq -z_1') = P(\epsilon_2 \geq -z_1'')$.

We sample $n = 10,000$ instances from this model, and define four groups of instances based on the combination of variances of ϵ_1 and ϵ_2,
$G_1 = \{i : \mathrm{Var}(\epsilon_{i1}) = 0.5 \text{ and } \mathrm{Var}(\epsilon_{i2}) = 0.5\}$, $G_2 = \{i : \mathrm{Var}(\epsilon_{i1}) = 0.5 \text{ and } \mathrm{Var}(\epsilon_{i2}) = 1\}$, $G_3 = \{i : \mathrm{Var}(\epsilon_{i1}) = 1 \text{ and } \mathrm{Var}(\epsilon_{i1}) = 0.5\}$, $G_4 = \{i : \mathrm{Var}(\epsilon_{i1}) = 1 \text{ and } \mathrm{Var}(\epsilon_{i1}) = 1\}$, for $i = 1, \cdots, 10,000$. We fit regression models separately for the instances in each of the four groups, considering both the true causal direction ($1 \to 2$, i.e., logistic regression of W_2 on W_1 and linear regression of Z_2 on Z_1) and the anti-causal direction ($2 \to 1$, i.e., linear regression of W_1 on W_2 and linear regression of Z_1 on Z_2). Our results show that the fitted regression coefficients are nearly identical in the causal direction across all groups in both the observed and latent spaces. In contrast, the coefficients exhibit significant variation across groups in the anti-causal direction in both spaces, as depicted in Figure 4.3.

This finding suggests that only the true causal graph produces consistent regression coefficients across all groups. Consequently, global regression in the anti-causal direction fails

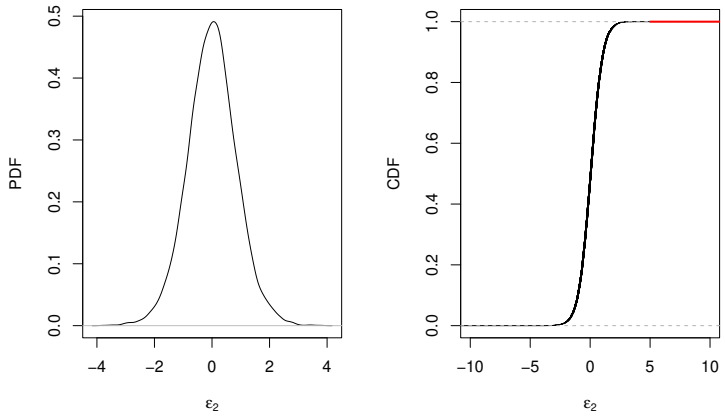

Figure 4.2: *Illustrate the Non-Injectivity of Example 1. The left panel shows the density distribution of ϵ_2, while the right panel presents its cumulative distribution function (CDF). The red segment in the right panel indicates that the CDF of ϵ_2 remains the same within the range denoted by the red line.*

to be optimal for each local component, indicating that the causal model is more suitable than the anti-causal model. This leads to the identification of the causal relationship.

4.3 Bayesian inference

In this section, we put forward a Bayesian inference method for jointly estimating basis coefficient sequences and the structure of DAGs. The prior specifications for the model parameters are detailed as follows:

Priors on the parameters of the Gaussian mixture distribution. The exogenous variable ϵ_{jk} follows a finite location mixture of Gaussian distributions,

$$\epsilon_{jk} \sim \sum_{m=1}^{M_{jk}} \pi_{jkm} \mathrm{N}\left(\mu_{jkm}, \tau_{jkm}\right),$$

where $\sum_{m=1}^{M_{jk}} \pi_{jkm} = 1$ and M_{jk} is the number of mixture components. We choose conjugate priors on (π_{jkm}, μ_{jkm} and τ_{jkm}),

$$\boldsymbol{\pi_{jk}} = (\pi_{jk1}, \ldots, \pi_{jkM_{jk}}) \sim \textsc{Dirichlet}(\alpha, \ldots, \alpha),$$

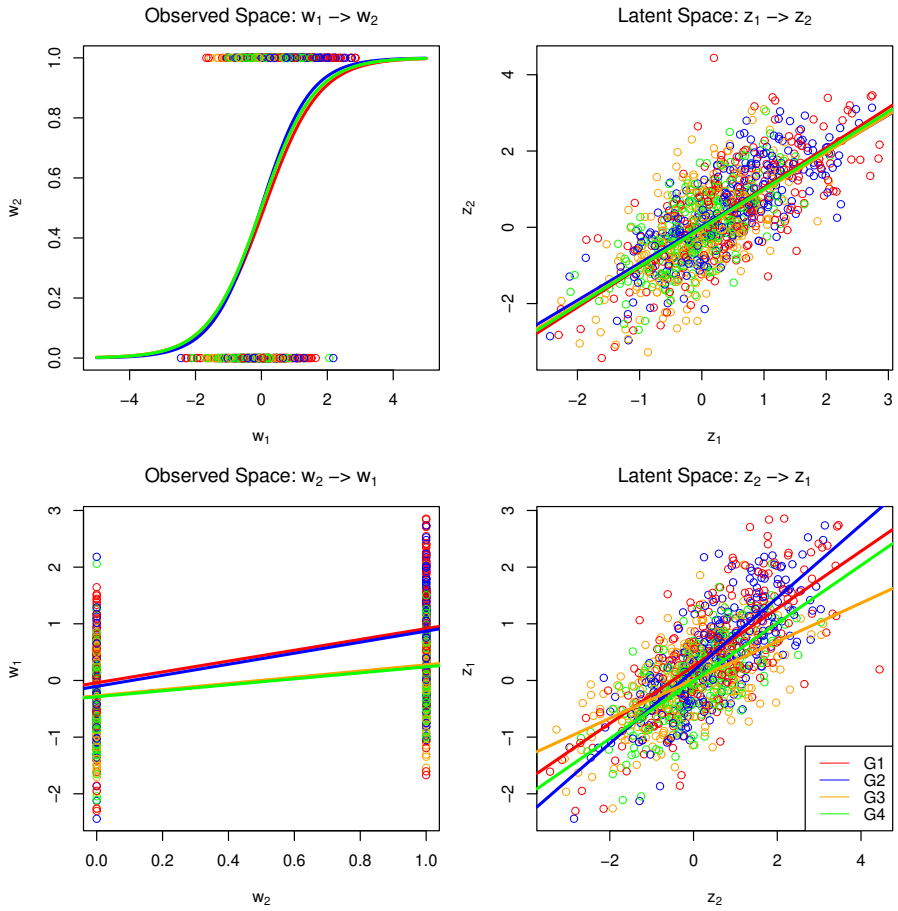

Figure 4.3: *Illustrative example of causal identification. The data are simulated according to the causal graph* $1 \rightarrow 2$*. The colored lines in each panel represent the fitted regression lines for the variables across groups 1 through 4. The top-left panel shows a probit regression of* W_2 *on* W_1 *in the observed space, the top-right shows a linear regression of* Z_2 *on* Z_1 *in the latent space, the bottom-left shows a linear regression of* W_1 *on* W_2 *in the observed space, and the bottom-right shows a linear regression of* Z_1 *on* Z_2 *in the latent space.*

$$\mu_{jkm}|\tau_{jkm} \sim N(a_\mu, b_\mu \tau_{jkm}), \quad \tau_{jkm} \sim \text{INV-GAMMA}(a_\tau, b_\tau), \quad \forall j \in [p], k \in [K], m \in [M_{jk}].$$

We set hyperparameters as follows: $\alpha = 1$, $(a_\mu, b_\mu) = (0, 1)$, and $(a_\tau, b_\tau) = (1, 1)$.

Priors on the adjacency matrix $\boldsymbol{E}$ and causal effect matrix $\boldsymbol{B}$. Let $\boldsymbol{E}$ represent the adjacency matrix of the DAG, $\boldsymbol{E} = (E_{jl})$ which is a $p \times p$ matrix, ($E_{jl} = 1$ if and only if $l \to j$). In a DAG, because it is acyclic, there won't be any cycles, there will be no 1 is in the main diagonal of the adjacency matrix $\boldsymbol{E}$. We uses Beta-Bernoulli prior on the adjacency matrix $\boldsymbol{E}$,

$$P(\boldsymbol{E}|r) \propto \prod_{j \neq l} r^{E_{jl}} (1-r)^{1-E_{jl}} I(\text{G is a DAG}), \quad r \sim \text{BETA}(a_r, b_r).$$

where r controls the probability of the edge occurring. We set hyperparameters as follows: $(a_r, b_r) = (1, 1)$.

Given $\boldsymbol{E}$, we use the independent spike and slab prior on the causal effect matrix $\boldsymbol{B}$, $\boldsymbol{B} = (\boldsymbol{B}_{jl})$ which is a $pK \times pK$ matrix,

$$\boldsymbol{B}_{jl}|E_{jl}, \gamma \sim (1 - E_{jl})\delta_{\boldsymbol{0}}(\boldsymbol{B}_{jl}) + E_{jl}\text{MN}_{K\times K}(\boldsymbol{B}_{jl}|\boldsymbol{0}, \gamma\boldsymbol{I}, \boldsymbol{I}), \quad \gamma \sim \text{INV-GAMMA}(a_\gamma, b_\gamma).$$

where $\delta_{\boldsymbol{0}}(\cdot)$ is a point mass at a $K \times K$ zero matrix and $\text{MN}_{K\times K}(\cdot|\boldsymbol{0}, \gamma\boldsymbol{I}, \boldsymbol{I})$ is a centered matrix normal distribution with row and column covariance matrix $\gamma\boldsymbol{I}$ and $\boldsymbol{I}$ where $\boldsymbol{I}$ is a $K \times K$ identity matrix. In the row covariance matrix, γ scales the identity matrix $\boldsymbol{I}$, determining the covariance structure along the rows and it also indicates the causal effect size. We set hyperparameters as follows: $(a_\gamma, b_\gamma) = (1, 1)$.

Priors on the parameters of orthonormal basis functions ϕ_k. For each $h \in [n]$, $k \in [K]$, $\phi_k(t_h) = \sum_{l=1}^{L} A_{kl} b_l(t_h)$, where $\boldsymbol{b} = (b_1, \cdots, b_L)^T$ is a set of cubic B-spline basis functions and $\boldsymbol{A}_k = (A_{k1}, \cdots, A_{kL})^T$ are B-spline coefficients. We focus on transformed B-spine coefficients $\widetilde{\boldsymbol{A}}_k$,

$$\phi_k(t_h) = \sum_{l=1}^{L} A_{kl} b_l(t_h) = \sum_{l=1}^{L} \widetilde{A}_{kl} \widetilde{b}_l(t_h), \text{ with } \widetilde{b}_l(t_h) = (1, t_h, \boldsymbol{b}^T(t_h)\boldsymbol{U}_{(L-2)\times(L-2)}\boldsymbol{D}_{L\times(L-2)}^{-1/2})^T,$$

where $\int \boldsymbol{b}''(t_h)\left[\boldsymbol{b}''(t_h)\right]^T dt_h = \boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^T$, $\boldsymbol{U}_{(L-2)\times(L-2)}$ and $\boldsymbol{D}_{L\times(L-2)}$ are submatrix of $\boldsymbol{U}$ and $\boldsymbol{D}$, respectively. Then we choose priors on $\widetilde{\boldsymbol{A}}_k$,

$$\widetilde{\boldsymbol{A}}_k \sim \text{MN}(\boldsymbol{0}, \boldsymbol{\Omega}_k), \text{ with } \Omega_k = \text{diag}(\infty, \infty, \lambda_k^{-1}, \cdots, \lambda_k^{-1})$$
$$\lambda_k \sim \text{UNIF}(L_k, U_k), \text{ with } L_k = \lambda_{k+1}, \text{ and } U_k = \lambda_{k-1}, \text{specifically}, L_1 = 10^8, L_K = 10^{-8}.$$

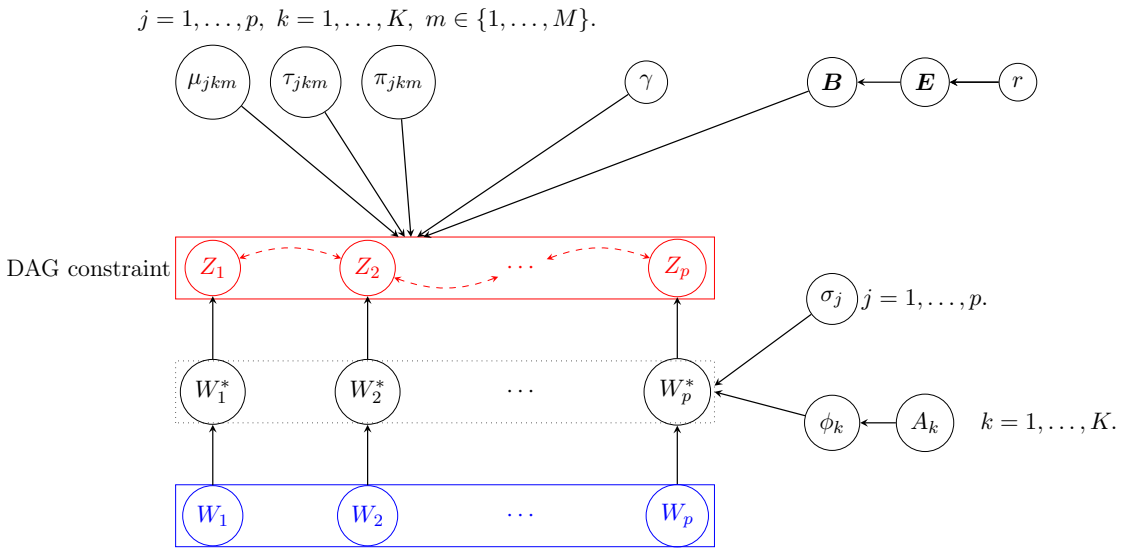

Figure 4.4: *A DAG illustrating the model hierarchy. The blue nodes represent observed variables, while the red nodes correspond to latent space coefficients that establish causal structures.*

Prior on the variance of noise σ_j. For each $j \in [p]$, we choose conjugate inverse-gamma prior to the variance of noise, $\sigma_j \sim \text{INV-GAMMA}(a_\sigma, b_\sigma)$. We set hyperparameters as follows: $(a_\sigma, b_\sigma) = (0.1, 0.1)$.

Figure 4.4 provides an overview of the proposed Bayesian hierarchical model. We simulate posterior samples through Markov Chain Monte Carlo (MCMC). Details are given in Appendix 4.B.

4.4 Simulation studies

4.4.1 Data generation

We simulate data from mixed FLiNG-BN with all possible combinations of sample sizes $N \in (50, 100, 200)$, time points $h \in (50, 100)$ and numbers of functions $p \in (10, 20, 30)$. The true graph $\mathcal{G}$ with p nodes which is simulated using the Erdős-Rényi model, with a connection probability of $2/p$, while adhering to the acyclicity constraint. Each element of each block of non-zero direct causal effects $\boldsymbol{B}_{jl}$ is independently drawn from a standard normal distribution. Following $\boldsymbol{Z} = (\boldsymbol{I} - \boldsymbol{B})^{-1}\boldsymbol{\epsilon}$, we generate the basis coefficient sequences

$\boldsymbol{Z}$, where the exogenous variables $\boldsymbol{\epsilon}$ are sampled from a centered Laplace distribution with a scale parameter $b = 0.5$.

Setting the true number of basis function to $K = 5$, we generate non-orthonormal functions $\phi'_k, \forall k \in [K]$ from a set of L=6 cubic B-spline basis functions on the interval $[0, 1]$ with evenly spaced knots, $\phi'_k = \sum_{\ell=1}^{L} A_{k\ell} b_\ell$, where $A_{k\ell}$ are generated from a standard normal distribution. We then use the singular value decomposition (SVD) to produces orthonoimal basis function ϕ_k. Finally, the first five nodes are specified as binary, while the subsequent nodes are set as continuous. Noisy observations are then generated, employing a signal-to-noise ratio set to 5. Therefore,

- for binary functions $W_j^{(i)}(t_h)$, for $j = 1, \cdots, 5$, $i = 1, \cdots, N$ and $h = 1, \cdots, n$:

$$W_j^{(i)}(t_h) \sim \mathrm{I}\left(W_j^{*(i)}(t_h) > 0\right), \quad W_j^{*(i)}(t_h) = \sum_{k=1}^{K} Z_{jk}\phi_{jk}^{(i)}(t_h) + e_j, \quad e_j \sim \mathrm{N}(0, \sigma_j).$$

- for continuous functions $W_j^{(i)}(t_h)$, for $j = 6, \cdots, p$, $i = 1, \cdots, N$ and $h = 1, \cdots, n$:

$$W_j^{(i)}(t_h) = \sum_{k=1}^{K} Z_{jk}\phi_{jk}^{(i)}(t_h) + e_j, \quad e_j \sim \mathrm{N}(0, \sigma_j).$$

where the notation $W_j^{(i)}(t_h)$ refers to the observation of the jth variable for the ith individual at the t_h time point, and σ_j is the mean values of $0.2\times \mid \sum_{k=1}^{K}\sum_{h=1}^{n} Z_{jk}\phi_{jk}(t_h) \mid$.

The data generation process described above satisfies the two assumptions mentioned in section 4.2.3. Each $W_j^*(t_h)$ is drawn from a Laplace distribution with a location parameter $\mu = e_j$ and a scale parameter $b = 0.5\left[\boldsymbol{\phi}(\mathbf{I} - \mathbf{B})^{-1}\right]_{(j)}$, where $\left[\boldsymbol{\phi}(\mathbf{I} - \mathbf{B})^{-1}\right]_{(j)}$ represents the j-th element of $\left[\boldsymbol{\phi}(\mathbf{I} - \mathbf{B})^{-1}\right]$. Here, $\boldsymbol{\phi}$ is a $p \times pK$ block diagonal matrix with diagonal elements $\boldsymbol{\phi}_j = (\phi_{j1}, \ldots, \phi_{jK})$, and $\{\phi_{jk}\}_{k=1}^{K}$ is a sequence of basis functions. If Y follows a Laplace distribution Laplace(μ, b), then Y can be represented as a mixture of normals (Ding and Blitzstein, 2018): $Y = \mu + b\sqrt{W}Z$, where W follows an exponential distribution with a mean of 1, and Z follows a standard normal distribution. Additionally, W and Z are independent. One way to interpret this equation is by considering different instruments, each having the same mean but differing in precision. Given that $P(e_j = 0) = 0$, the marginal distribution of $W_j(t_h)$ cannot be identical across each component of the latent continuous variable $W_j^*(t_h)$.

4.4.2 Model specification and inference

We apply the proposed mixed-FLiNG-BN to 30 replicate data sets. In each simulation, the MCMC algorithm is run for 10,000 iterations, using the first 5,000 iterations as burn-in, and retaining every 10th iteration. In addition, to implement mixed-FLiNG-BN, we specify the number of mixture components of the exogenous variable to $M = 5$ and the number of B-spline basis functions to $L = 20$ (the true number of B-spline basis functions is $L = 5$).

Methods for Comparative Analysis. We compare the proposed mixed-FLiNG-BN with three different methodologies relying on two-step estimation procedures. In the first step, we extract basis coefficients obtained from functional PCA using the package ***registr*** (Wrobel et al., 2019), here for continuous functions, we use functional principal components analysis via variational EM (Tipping and Bishop, 1999); while, for binary functions, we use binary functional principal components analysis (Jaakkola and Jordan, 1997). Additionally, we fix the number of B-spline basis functions at $L = 5$. In the second step, given each of the estimated basis coefficients, we construct L causal graphs using the Linear non-Gaussian Acyclic Models (LiNGAM) algorithm (Shimizu et al., 2006) (referred to as mixed FPCA-LiNGAM), the PC algorithm (Spirtes et al., 2000; Ramsey et al., 2012) (referred to as mixed FPCA-PC), and the Greedy Equivalence Search (GES) algorithm (Chickering, 2002) (referred to as mixed FPCA-GES). When any of the L causal graphs indicate that ith node and jth node have a relationship, we acknowledge the existence of a causal relationship between i and j. To implement LiNGAM, PC, and GES algorithms, we use the R package ***pcalg*** (Kalisch et al., 2012; Hauser and Bühlmann, 2012).

Evaluation Metrics. The causal graph $\mathcal{G}$ is estimated by thresholding the posterior probability of inclusion at 0.5 (i.e., the median probability model). This thresholding method serves to delineate the presence or absence of edges in the graph. To compare the graph recovery performance, we calculate True Positive Rate (TPR), False Discovery Rate (FDR), and Matthews Correlation Coefficient (MCC) (Chicco and Jurman, 2020),

$$TPR = \frac{TP}{TP + FN}, \quad FDR = \frac{FP}{FP + TP},$$

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}.$$

where TP, TN, FP, FN denote the number of True Positives, True Negatives, False Positives and False Negatives, respectively.

TPR measures the proportion of correctly identified causal relationships among all true positives; FDR quantifies the ratio of falsely identified causal relationships to the total number of positive identifications; and MCC offers a balanced evaluation by considering both true and false positives, as well as true and false negatives. The range of MCC is from -1 to 1, here values closer to 1 suggest higher predictive accuracy, while values closer to 0 indicate poorer performance or random choice.

4.4.3 Results

The results, derived from 30 repeated simulations and summarized in Table 4.1, demonstrate the significant superiority of the proposed mixed FLiNG-BN over its competitors—mixed FPCA-LiNGAM, mixed FPCA-PC, and mixed FPCA-GES—across all combinations of p, d, and n. In each scenario, the TPR of Mixed FLiNG-BN exceeds 0.6, with more than half of them surpassing 0.7. Additionally, the majority of its FDR remains below 0.15, resulting in a high MCC, the majority of which are greater than 0.75. Conversely, upon examining mixed FPCA-LiNGAM, mixed FPCA-PC, and mixed FPCA-GES, it is evident that they exhibit very high TPR along with high FDR. Consequently, this leads to a considerably lower MCC, around 0.4. In certain scenarios, mixed FPCA-PC even demonstrates lower TPR and higher FDR. This outcome is unsurprising given that the competing methods, namely mixed FPCA-LiNGAM, mixed FPCA-PC, and mixed FPCA-GES, adhere to a two-step procedure where the effectiveness of the functional PCA coefficients extracted in the initial step to accurately capture functional dependencies in the subsequent step is highly questionable. In stark contrast to our proposed approach, none of the rival methods account for false discovery control. Consequently, these methods suffer from high false discovery rates.

Alongside this, we perform a comparative analysis between the true latent functions of binary functions and the estimated latent functions derived through the mixed-FLiNG-BN model, as depicted in Figure 4.5. This comparison reveals that, for the nine randomly selected individuals, the estimated latent functions exhibit a high degree of accuracy and closely align with the true latent functions of the binary functions. This close correspondence underscores the effectiveness of the mixed-FLiNG-BN model in accurately estimating latent functions in binary data.

Sensitivity analysis. The proposed mixed-FLiNG-BN involves several hyperparameters α, (a_μ, b_μ), (a_τ, b_τ), (a_r, b_r), (a_γ, b_γ), and (a_σ, b_σ). We conduct sensitivity analyses on these hyperparameters with $(p, h, N) = (20, 50, 200)$ where five of the nodes are binary and the

			Mixed FLiNG-BN			Mixed FPCA-LiNGAM			Mixed FPCA-PC			Mixed FPCA-GES		
p	d	n	TPR	FDR	MCC	TPR	FDR	MCC	TPR	FDR	MCC	TPR	FDR	MCC
10	50	200	0.72 (0.07)	**0.12 (0.01)**	**0.77 (0.06)**	**0.81 (0.07)**	0.73 (0.001)	0.42 (0.02)	0.47 (0.04)	0.83 (0.03)	0.15 (0.05)	0.69 (0.08)	0.80 (0.07)	0.25 (0.09)
10	50	500	0.76 (0.09)	**0.07 (0.08)**	**0.82 (0.08)**	**0.94 (0.08)**	0.66 (0.01)	0.51 (0.04)	0.76 (0.06)	0.72 (0.01)	0.37 (0.05)	0.90 (0.03)	0.78 (0.002)	0.35 (0.01)
10	50	1000	0.80 (0.05)	**0.11 (0.01)**	**0.82 (0.03)**	**0.98 (0.01)**	0.71 (0.06)	0.48 (0.09)	0.78 (0.06)	0.71 (0.06)	0.39 (0.03)	0.90 (0.03)	0.79 (0.01)	0.33 (0.01)
10	50	2000	0.87 (0.06)	**0.13 (0.06)**	**0.85 (0.06)**	0.99 (0.02)	0.65 (0.06)	0.54 (0.09)	0.78 (0.04)	0.72 (0.03)	0.38 (0.02)	**0.99 (0.01)**	0.65 (0.08)	0.53 (0.08)
10	100	200	0.76 (0.07)	**0.13 (0.01)**	**0.76 (0.05)**	**0.90 (0.08)**	0.66 (0.04)	0.51 (0.10)	0.66 (0.02)	0.74 (0.01)	0.32 (0.002)	0.88 (0.03)	0.74 (0.002)	0.39 (0.002)
10	100	500	0.72 (0.09)	**0.08 (0.06)**	**0.79 (0.07)**	0.95 (0.05)	0.72 (0.07)	0.47 (0.06)	0.68 (0.02)	0.71 (0.004)	0.36 (0.02)	**0.99 (0.03)**	0.73 (0.02)	0.43 (0.05)
10	100	1000	0.79 (0.09)	**0.06 (0.07)**	**0.83 (0.05)**	0.98 (0.01)	0.69 (0.03)	0.50 (0.02)	0.70 (0.01)	0.70 (0.01)	0.37 (0.03)	**0.99 (0.03)**	0.73 (0.03)	0.44 (0.06)
10	100	2000	0.84 (0.08)	**0.12 (0.05)**	**0.84 (0.08)**	**0.99 (0.01)**	0.67 (0.01)	0.54 (0.003)	0.77 (0.03)	0.72 (0.01)	0.38 (0.02)	**0.99 (0.01)**	0.71 (0.02)	0.47 (0.02)
20	50	200	0.63 (0.04)	**0.13 (0.004)**	**0.73 (0.02)**	0.65 (0.08)	0.67 (0.09)	0.42 (0.06)	0.76 (0.07)	0.76 (0.05)	0.39 (0.05)	**0.90 (0.02)**	0.82 (0.01)	0.37 (0.03)
20	50	500	0.72 (0.05)	**0.13 (0.05)**	**0.78 (0.03)**	**0.92 (0.04)**	0.79 (0.01)	0.39 (0.02)	0.78 (0.09)	0.76 (0.02)	0.40 (0.08)	0.91 (0.02)	0.80 (0.01)	0.38 (0.02)
20	50	1000	0.76 (0.08)	**0.13 (0.04)**	**0.80 (0.04)**	0.99 (0.02)	0.79 (0.004)	0.41 (0.004)	0.85 (0.08)	0.75 (0.04)	0.42 (0.01)	**0.99 (0.01)**	0.73 (0.02)	0.45 (0.03)
20	50	2000	0.78 (0.06)	**0.14 (0.05)**	**0.80 (0.04)**	**0.99 (0.01)**	0.80 (0.02)	0.40 (0.04)	0.82 (0.07)	0.76 (0.05)	0.41 (0.01)	**0.99 (0.01)**	0.77 (0.02)	0.46 (0.02)
20	100	200	0.62 (0.08)	**0.04 (0.01)**	**0.76 (0.06)**	0.92 (0.04)	0.79 (0.01)	0.39 (0.02)	0.70 (0.08)	0.74 (0.01)	0.40 (0.04)	**0.99 (0.02)**	0.82 (0.002)	0.39 (0.01)
20	100	500	0.66 (0.06)	**0.19 (0.03)**	**0.72 (0.03)**	0.90 (0.05)	0.80 (0.04)	0.38 (0.07)	0.73 (0.09)	0.76 (0.03)	0.39 (0.08)	**0.99 (0.02)**	0.79 (0.01)	0.42 (0.04)
20	100	1000	0.69 (0.05)	**0.11 (0.05)**	**0.77 (0.03)**	**0.99 (0.02)**	0.78 (0.01)	0.43 (0.01)	0.85 (0.08)	0.74 (0.04)	0.45 (0.04)	**0.99 (0.02)**	0.77 (0.03)	0.46 (0.05)
20	100	2000	0.77 (0.08)	**0.14 (0.02)**	**0.80 (0.05)**	**0.99 (0.01)**	0.79 (0.01)	0.41 (0.01)	0.84 (0.09)	0.77 (0.06)	0.41 (0.04)	**0.99 (0.01)**	0.76 (0.01)	0.46 (0.02)
30	50	200	0.63 (0.06)	**0.12 (0.02)**	**0.74 (0.04)**	0.65 (0.02)	0.71 (0.03)	0.41 (0.09)	0.48 (0.05)	0.83 (0.01)	0.25 (0.03)	**0.90 (0.01)**	0.86 (0.02)	0.32 (0.01)
30	50	500	0.66 (0.04)	**0.15 (0.02)**	**0.74 (0.02)**	0.81 (0.04)	0.81 (0.08)	0.34 (0.02)	0.45 (0.03)	0.86 (0.01)	0.22 (0.02)	**0.95 (0.04)**	0.85 (0.003)	0.35 (0.01)
30	50	1000	0.70 (0.09)	**0.16 (0.05)**	**0.76 (0.06)**	**0.92 (0.01)**	0.85 (0.01)	0.31 (0.02)	0.55 (0.02)	0.85 (0.02)	0.26 (0.03)	0.90 (0.02)	0.86 (0.001)	0.32 (0.002)
30	50	2000	0.73 (0.04)	**0.18 (0.01)**	**0.76 (0.03)**	0.97 (0.01)	0.82 (0.04)	0.37 (0.06)	0.55 (0.01)	0.79 (0.004)	0.31 (0.004)	**0.99 (0.01)**	0.85 (0.03)	0.35 (0.003)
30	100	200	0.60 (0.09)	**0.04 (0.02)**	**0.75 (0.08)**	0.69 (0.08)	0.77 (0.003)	0.35 (0.04)	0.57 (0.09)	0.79 (0.02)	0.31 (0.06)	**0.90 (0.004)**	0.86 (0.001)	0.32 (0.01)
30	100	500	0.67 (0.04)	**0.19 (0.04)**	**0.72 (0.03)**	0.80 (0.06)	0.82 (0.01)	0.32 (0.04)	0.60 (0.08)	0.81 (0.02)	0.30 (0.05)	**0.95 (0.02)**	0.83 (0.01)	0.38 (0.004)
30	100	1000	0.70 (0.05)	**0.21 (0.06)**	**0.73 (0.04)**	0.98 (0.01)	0.82 (0.02)	0.37 (0.04)	0.62 (0.02)	0.82 (0.02)	0.30 (0.02)	**0.99 (0.01)**	0.85 (0.01)	0.36 (0.01)
30	100	2000	0.71 (0.05)	**0.19 (0.04)**	**0.74 (0.03)**	0.98 (0.01)	0.82 (0.04)	0.36 (0.06)	0.65 (0.06)	0.79 (0.01)	0.34 (0.04)	**0.99 (0.04)**	0.84 (0.02)	0.36 (0.05)

Table 4.1: *Observations are generated using mixed functions on an evenly spaced grid. The reported average performance is based on 30 repetitions, and standard deviations are given within the parentheses. The best-performing method for each metric is highlighted in bold.*

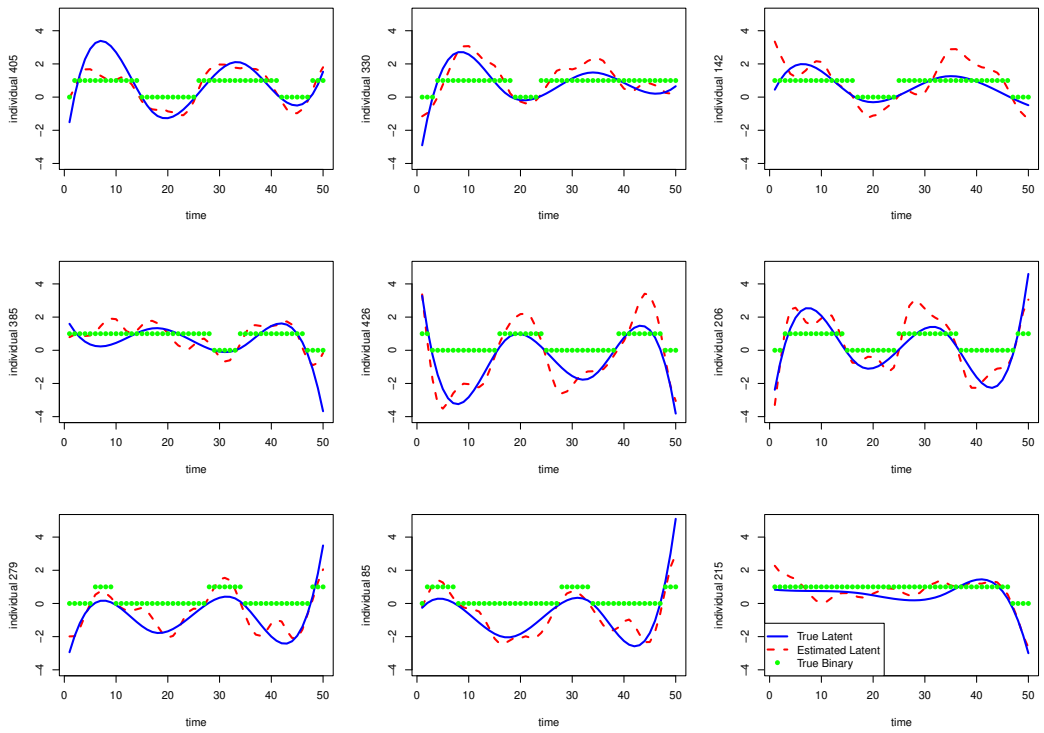

Figure 4.5: *True latent functions of a binary function and their estimates via Mixed-FLiNG-BN. One binary function and nine individuals are randomly selected from a total of 5 binary functions and 500 samples. The trajectories of the binary function (represented by green dots) are plotted, along with the corresponding true latent function (represented by the blue solid line) and the estimated latent function (represented by the red dashed line) obtained using our proposed model. The number of time points is 50.*

rest are continuous, revealing that our model exhibits relative robustness across the tested ranges. The results are shown in Appendix 4.C.

4.5 Discussion

In this chapter, we propose mixed functional Bayesian networks to address challenges associated with diverse functional data types, facilitating a more accurate representation of causal dependencies by a directed acyclic graph. We also present that under this model, the causal structure between mixed functional data can be uniquely identifiable under some mild conditions. Furthermore, we develop a Bayesian framework to infer the adjacency matrix of the directed acyclic graph with inherent uncertainty quantification. The performance

of the proposed model is demonstrated by extensive simulation studies.

We outline several potential directions for extending our current research. First, we aim to establish theoretical proofs for the casual identification of binary functions to binary functions within the framework of mixed functional Bayesian networks. Second, while our current work primarily addresses mixed datasets comprising continuous and binary functional data, an interesting extension would be to generalize binary functions to discrete functions. Lastly, as our focus thus far has been predominantly on theoretical developments and simulations, it would be valuable to apply the proposed model to real-world data for empirical validation.

Bibliography

Albert, J. H. and Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422):669–679.

Bouveyron, C., Jacques, J., Schmutz, A., Simoes, F., and Bottini, S. (2022). Co-clustering of multivariate functional data for the analysis of air pollution in the south of france. *The Annals of Applied Statistics*, 16(3):1400–1422.

Burkardt, J. (2014). The truncated normal distribution. *Department of Scientific Computing Website, Florida State University*, 1(35):58.

Cai, Z., Xi, D., Zhu, X., and Li, R. (2022). Causal discoveries for high dimensional mixed data. *Statistics in Medicine*, 41(24):4924–4940.

Carroll, C. and Müller, H.-G. (2023). Latent deformation models for multivariate functional data and time-warping separability. *Biometrics*, 79(4):3345–3358.

Chicco, D. and Jurman, G. (2020). The advantages of the matthews correlation coefficient (mcc) over f1 score and accuracy in binary classification evaluation. *BMC genomics*, 21:1–13.

Chickering, D. M. (2002). Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3(Nov):507–554.

Ding, P. and Blitzstein, J. K. (2018). On the gaussian mixture representation of the laplace distribution. *The American Statistician*, 72(2):172–174.

Glymour, C., Zhang, K., and Spirtes, P. (2019). Review of causal discovery methods based on graphical models. *Frontiers in Genetics*, 10:524.

Golovkine, S., Klutchnikoff, N., and Patilea, V. (2022). Clustering multivariate functional data using unsupervised binary trees. *Computational Statistics & Data Analysis*, 168:107376.

Hauser, A. and Bühlmann, P. (2012). Characterization and greedy learning of interventional markov equivalence classes of directed acyclic graphs. *The Journal of Machine Learning Research*, 13(1):2409–2464.

Hyndman, R. J. and Ullah, M. S. (2007). Robust forecasting of mortality and fertility rates: A functional data approach. *Computational Statistics & Data Analysis*, 51(10):4942–4956.

Ieva, F. and Paganoni, A. M. (2013). Depth measures for multivariate functional data. *Communications in Statistics-Theory and Methods*, 42(7):1265–1276.

Jaakkola, T. S. and Jordan, M. I. (1997). A variational approach to Bayesian logistic regression models and their extensions. In *Sixth International Workshop on Artificial Intelligence and Statistics*, pages 283–294. PMLR.

Kalisch, M., Mächler, M., Colombo, D., Maathuis, M. H., and Bühlmann, P. (2012). Causal inference using graphical models with the R package pcalg. *Journal of Statistical Software*, 47:1–26.

Kong, D., Yang, S., and Wang, L. (2022). Identifiability of causal effects with multiple causes and a binary outcome. *Biometrika*, 109(1):265–272.

Kowal, D. R., Matteson, D. S., and Ruppert, D. (2017). A Bayesian multivariate functional dynamic linear model. *Journal of the American Statistical Association*, 112(518):733–744.

Lee, K.-Y. and Li, L. (2022). Functional structural equation model. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(2):600–629.

Li, B. and Solea, E. (2018). A nonparametric graphical model for functional data with application to brain networks based on fMRI. *Journal of the American Statistical Association*, 113(524):1637–1655.

Li, Y., Qiu, Y., and Xu, Y. (2022). From multivariate to functional data analysis: Fundamentals, recent developments, and emerging areas. *Journal of Multivariate Analysis*, 188:104806.

López-Pintado, S., Sun, Y., Lin, J. K., and Genton, M. G. (2014). Simplicial band depth for multivariate functional data. *Advances in Data Analysis and Classification*, 8:321–338.

Mityagin, B. (2015). The zero set of a real analytic function. *arXiv preprint arXiv:1512.07276*.

Moindjié, I.-A., Dabo-Niang, S., and Preda, C. (2024). Classification of multivariate functional data on different domains with partial least squares approaches. *Statistics and Computing*, 34(1):5.

Peters, J., Janzing, D., and Schölkopf, B. (2017). *Elements of causal inference: foundations and learning algorithms*. The MIT Press.

Peters, J., Mooij, J. M., Janzing, D., and Schölkopf, B. (2014). Causal discovery with continuous additive noise models.

Qiao, X., Guo, S., and James, G. M. (2019). Functional graphical models. *Journal of the American Statistical Association*, 114(525):211–222.

Raghu, V. K., Ramsey, J. D., Morris, A., Manatakis, D. V., Sprites, P., Chrysanthis, P. K., Glymour, C., and Benos, P. V. (2018). Comparison of strategies for scalable causal discovery of latent variable models from mixed data. *International Journal of Data Science and Analytics*, 6:33–45.

Ramsey, J., Zhang, J., and Spirtes, P. L. (2012). Adjacency-faithfulness and conservative causal inference. *arXiv preprint arXiv:1206.6843*.

Rogovchenko, V., Sibu, A., and Ni, Y. (2024). Scalar-function causal discovery for generating causal hypotheses with observational wearable device data. In *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, volume 29, page 201. NIH Public Access.

Roy, S., Wong, R. K., and Ni, Y. (2024). Directed cyclic graph for causal discovery from multivariate functional data. *Advances in Neural Information Processing Systems*, 36.

Shimizu, S. (2014). LiNGAM: Non-gaussian methods for estimating causal structures. *Behaviormetrika*, 41:65–98.

Shimizu, S., Hoyer, P. O., Hyvärinen, A., Kerminen, A., and Jordan, M. (2006). A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(10).

Solea, E. and Li, B. (2022). Copula gaussian graphical models for functional data. *Journal of the American Statistical Association*, 117(538):781–793.

Spirtes, P., Glymour, C. N., and Scheines, R. (2000). *Causation, prediction, and search*. MIT press.

Tagasovska, N., Vatter, T., and Chavez-Demoulin, V. (2018). Nonparametric quantile-based causal discovery. *arXiv preprint arXiv:1801.10579*.

Tipping, M. E. and Bishop, C. M. (1999). Probabilistic principal component analysis. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 61(3):611–622.

Tsagris, M., Borboudakis, G., Lagani, V., and Tsamardinos, I. (2018). Constraint-based causal discovery with mixed data. *International Journal of Data Science and Analytics*, 6:19–30.

Van den Hof, P., Wahlberg, B., Heuberger, P., Ninness, B., Bokor, J., and e Silva, T. O. (2000). Modelling and identification with rational orthogonal basis functions. *IFAC Proceedings Volumes*, 33(15):445–455.

Van Der Linde, A. (2009). A Bayesian latent variable approach to functional principal components analysis with binary and count data. *AStA Advances in Statistical Analysis*, 93:307–333.

Wei, Z. and Li, H. (2008). A hidden spatial-temporal markov random field model for network-based analysis of time course gene expression data.

Wenjuan, W., Lu, F., and Chunchen, L. (2018). Mixed causal structure discovery with application to prescriptive pricing. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 5126–5134.

Wrobel, J., Zipunnikov, V., Schrack, J., and Goldsmith, J. (2019). Registration for exponential family functional data. *Biometrics*, 75(1):48–57.

Yamayoshi, M., Tsuchida, J., and Yadohisa, H. (2020). An estimation of causal structure based on latent LiNGAM for mixed data. *Behaviormetrika*, 47:105–121.

Yang, T.-L., Lee, K.-Y., Zhang, K., and Suzuki, J. (2024). Functional linear non-gaussian acyclic model for causal discovery. *Behaviormetrika*, pages 1–22.

Yao, F., Müller, H.-G., and Wang, J.-L. (2005). Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*, 100(470):577–590.

Zeng, Y., Shimizu, S., Matsui, H., and Sun, F. (2022). Causal discovery for linear mixed data. In *Conference on Causal Learning and Reasoning*, pages 994–1009. PMLR.

Zhang, J., Siegle, G. J., Sun, T., D'andrea, W., and Krafty, R. T. (2023). Interpretable principal component analysis for multilevel multivariate functional data. *Biostatistics*, 24(2):227–243.

Zhang, K. and Hyvarinen, A. (2012). On the identifiability of the post-nonlinear causal model. *arXiv preprint arXiv:1205.2599*.

Zhong, R., Liu, S., Li, H., and Zhang, J. (2023). Sparse logistic functional principal component analysis for binary data. *Statistics and Computing*, 33(1):15.

Zhou, F., He, K., Wang, K., Xu, Y., and Ni, Y. (2023). Functional Bayesian networks for discovering causality from multivariate functional data. *Biometrics*, 79(4):3279–3293.

Zhu, H., Strawn, N., and Dunson, D. B. (2016). Bayesian graphical models for multivariate functional data. *Journal of Machine Learning Research*, 17(204):1–27.

Zhu, T. (2023). The general linear hypothesis testing problem for multivariate functional data with applications. *arXiv preprint arXiv:2312.02518*.

Appendices

4.A Proof of Theorem 4.2.5

Proof. We respectively prove the bivariate identifiability of mixed-FLiNG-BN in continuous-binary case, and continuous-continuous case.

Theorem 4.A.1. *Given a pair of mixed functions $(\boldsymbol{W}_1, \boldsymbol{W}_2)$, each measured on n time points $\boldsymbol{W}_1 = (W_1(t_1), \ldots, W_1(t_n))^T \in \{0,1\}^n$, $\boldsymbol{W}_2 = (W_2(t_1), \ldots, W_2(t_n))^T \in \mathcal{R}^n$, under assumption 4.2.4, the mixed-FLiNG-BN is bivariate identifiable.*

Proof. We initiate the proof in the scalar scenario, setting $n = 1$, then extending it to the longitudinal case where $n > 1$. Note for simplicity, we assume that the number of basis functions is $K = 1$ and the error is 0, accordingly, assume these exist two models,
$W_1 \to W_2$:

$$\begin{aligned} W_1^* =& \epsilon_1 \\ W_2^* =& b_{12} W_1^* + \epsilon_2 \\ \epsilon_1 \sim \sum_{i=1}^{M_1} \pi_i^{(1)} \mathrm{N}(\mu_i^{(1)}, \sigma_i^{(1)}),& \;\; \epsilon_2 \sim \sum_{i=1}^{M_2} \pi_i^{(2)} \mathrm{N}(\mu_i^{(2)}, \sigma_i^{(2)}) \\ W_1 = \mathrm{I}(W_1^* > 0),& \;\; W_2 = W_2^*. \end{aligned} \tag{6}$$

$W_2 \to W_1$:

$$\begin{aligned} W_1^* =& b_{21} W_2^* + \widetilde{\epsilon}_1 \\ W_2^* =& \widetilde{\epsilon}_2 \\ \widetilde{\epsilon}_1 \sim \sum_{i=1}^{M_1} \widetilde{\pi}_i^{(1)} \mathrm{N}(\widetilde{\mu}_i^{(1)}, \widetilde{\sigma}_i^{(1)}),& \;\; \widetilde{\epsilon}_2 \sim \sum_{i=1}^{M_2} \widetilde{\pi}_i^{(2)} \mathrm{N}(\widetilde{\mu}_i^{(2)}, \widetilde{\sigma}_i^{(2)}) \\ W_1 = \mathrm{I}(W_1^* > 0),& \;\; W_2 = W_2^*. \end{aligned} \tag{7}$$

Suppose $w_1 \in \{0,1\}$ and $w_2 \in \mathcal{R}$. Consider two competing causal models $P_{W_1 \to W_2}(W_1, W_2 | b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $P_{W_2 \to W_1}(W_1, W_2 | b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}})$. We will show that these two causal models are in general not equivalent, i.e., $P_{W_1 \to W_2}(W_1 = w_1, W_2 = w_2 | b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \neq P_{W_2 \to W_1}(W_1 = w_1, W_2 = w_2 | b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}})$ for some $w_1 \in \{0,1\}$ and $w_2 \in \mathcal{R}$. We prove it

by contradiction. Suppose for any $w_1 \in \{0,1\}$ and $w_2 \in \mathcal{R}$,

$$P_{W_1 \to W_2}(W_1 = w_1, W_2 = w_2 | b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = P_{W_2 \to W_1}(W_1 = w_1, W_2 = w_2 | b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) \quad (8)$$

The left-hand size of (8) is given by

$$\begin{aligned} & P_{W_1 \to W_2}(W_1 = w_1, W_2 = w_2 | b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \\ =& f_{W_1 \to W_2}(W_2 = w_2 | W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) P_{W_1 \to W_2}(W_1 = w_1 | b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \\ =& f_{W_1 \to W_2}(W_2 = w_2 | W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \pi_{w_1} \end{aligned} \quad (9)$$

Similarly, the right-hand size of (8) is given by

$$\begin{aligned} & P_{W_2 \to W_1}(W_1 = w_1, W_2 = w_2 | b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) \\ =& P_{W_2 \to W_1}(W_1 = w_1 | W_2 = w_2, b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) P_{W_2 \to W_1}(W_2 = w_2 | b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) \\ =& P_{W_2 \to W_1}(W_1 = w_1 | W_2 = w_2, b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) f(w_2) \end{aligned} \quad (10)$$

Therefore, (8) leads to

$$f_{W_1 \to W_2}(W_2 = w_2 | W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \pi_{w_1} = P_{W_2 \to W_1}(W_1 = w_1 | W_2 = w_2, b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) f(w_2) \quad (11)$$

Summing up both sides of (11) over w_1 from 0 to 1, we have

$$\sum_{w_1=0}^{1} f_{W_1 \to W_2}(W_2 = w_2 | W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \pi_{w_1} = f(w_2) \quad (12)$$

Plug (12) into (11),

$$\begin{aligned} & f_{W_1 \to W_2}(W_2 = w_2 | W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \pi_{w_1} \\ & = P_{W_2 \to W_1}(W_1 = w_1 | W_2 = w_2, b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) \cdot \sum_{w_1=0}^{1} f_{W_1 \to W_2}(W_2 = w_2 | W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \pi_{w_1} \end{aligned} \quad (13)$$

we can have

$$P_{W_2 \to W_1}(W_1 = w_1 | W_2 = w_2, b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) = \frac{f_{W_1 \to W_2}(W_2 = w_2 | W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \pi_{w_1}}{\sum_{w_1=0}^{1} f_{W_1 \to W_2}(W_2 = w_2 | W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \pi_{w_1}} \quad (14)$$

Based on given w_2' and w_2'', such that $w_2' \neq w_2''$, while the following equality holds:

$$P_{W_2 \to W_1}(W_1 = w_1 \mid W_2 = w_2', b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) = P_{W_2 \to W_1}(W_1 = w_1 \mid W_2 = w_2'', b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}).$$

therefore, we can get

$$\frac{f_{W_1 \to W_2}(W_2 = w_2'|W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})\pi_{w_1}}{\sum_{w_1=0}^{1} f_{W_1 \to W_2}(W_2 = w_2'|W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})\pi_{w_1}} = \frac{f_{W_1 \to W_2}(W_2 = w_2''|W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})\pi_{w_1}}{\sum_{w_1=0}^{1} f_{W_1 \to W_2}(W_2 = w_2''|W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})\pi_{w_1}} \tag{15}$$

Without loss of generality, we mainly focus on $w_1 = 1$, then the (15) can be rewritten

$$\frac{f_{W_1 \to W_2}(W_2 = w_2'|W_1 = 1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})\pi_1}{\sum_{w_1=0}^{1} f_{W_1 \to W_2}(W_2 = w_2'|W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})\pi_{w_1}} = \frac{f_{W_1 \to W_2}(W_2 = w_2''|W_1 = 1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})\pi_1}{\sum_{w_1=0}^{1} f_{W_1 \to W_2}(W_2 = w_2''|W_1 = w_1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})\pi_{w_1}} \tag{16}$$

In the following, we need to consider the computation of the following conditional probabilities: $f_{W_1 \to W_2}(W_2 = w_2|W_1 = 0, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$, $f_{W_1 \to W_2}(W_2 = w_2|W_1 = 1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ for any $w_2 \in \mathcal{R}$, taking $f_{W_1 \to W_2}(W_2 = w_2|W_1 = 0, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ as an example,

$$f_{W_1 \to W_2}(W_2 = w_2|W_1 = 0, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \tag{17}$$

$$= \int_{w_1^*} f_{W_1 \to W_2}(W_2 = w_2|W_1 = 0, W_1^* = w_1^*, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \cdot P_{W_1 \to W_2}(W_1^* = w_1^*|W_1 = 0, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) dw_1^* \tag{18}$$

$$= \int_{w_1^*} f_{W_1 \to W_2}(w_2 = b_{12}w_1^* + \epsilon_2|W_1^* = w_1^*, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \cdot P_{W_1 \to W_2}(W_1^* = w_1^*|W_1 = 0, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) dw_1^* \tag{19}$$

$$= \int_{w_1^*} f_{W_1 \to W_2}(\epsilon_2 = w_2 - b_{12}w_1^*|W_1^* = w_1^*, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \cdot P_{W_1 \to W_2}(W_1^* = w_1^*|W_1 = 0, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) dw_1^* \tag{20}$$

$$= E_{w_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)\right) \tag{21}$$

Similarly, we can obtain

$$f_{W_1 \to W_2}(W_2 = w_2|W_1 = 1, b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = E_{w_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)\right). \tag{22}$$

Define the left side of (22) $:= f(b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$, let

$$\Theta = \{(b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) \mid f(b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = 0\},$$

We will show the Lebesgue measure of Θ is 0, that is $m(\Theta) = 0$.

Definition 4.A.2. (Real Analytic Function) A real function is said to be analytic if it is infinitely differentiable and matches its Taylor series in a neighborhood of every point.

Then, if we want to prove $f(b_{12}, \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is an analytic function, we only need to prove

$$\begin{aligned} &\frac{\pi_1 E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2' - b_{12}w_1^*\right)\right)}{\pi_0 E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2' - b_{12}w_1^*\right)\right) + \pi_1 E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2' - b_{12}w_1^*\right)\right)} \\ -&\frac{\pi_1 E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2'' - b_{12}w_1^*\right)\right)}{\pi_0 E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2'' - b_{12}w_1^*\right)\right) + \pi_1 E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2'' - b_{12}w_1^*\right)\right)} \end{aligned} \tag{23}$$

is an analytic function. Accordingly, if we aim to demonstrate that (23) is an analytic function, we simply need to establish $E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)\right)$ and $E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)\right)$ for any w_2 are analytic functions. In light of our previous discussion, $E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)\right)$ and $E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)\right)$ are related to the joint distribution of $(\epsilon_1, \epsilon_2)^T$,

$$\begin{pmatrix}\epsilon_1\\ \epsilon_2\end{pmatrix} \sim \sum_{i=1}^{M_1} \pi_i^{(1)} \mathrm{N}(\mu_i^{(1)}, \sigma_i^{(1)}) \sum_{i=1}^{M_2} \pi_i^{(2)} \mathrm{N}(\mu_i^{(2)}, \sigma_i^{(2)})$$

Hence, $E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)\right) = \sum_{i=1}^{M} \pi_i E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)|i=m\right)$, where $M = M_1 \cdot M_2$ and $\pi_i = \pi_i^{(1)}\pi_i^{(2)}$. $E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)|i=m\right)$ means the expected value of m-th component of $(\epsilon_1, \epsilon_2)^T$. The $f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)$ of m-th component of $(\epsilon_1, \epsilon_2)^T$ denoted by $f_{\epsilon_2}^{(m)}\left(w_2 - b_{12}w_1^*\right)$ can be found by using a Taylor series

$$\begin{aligned} f_{\epsilon_2}^{(m)}\left(w_2 - b_{12}w_1^*\right) &= \frac{1}{\sqrt{2\pi}\sigma_m} \sum_{k=0}^{\infty} (-1)^k \frac{\left(\frac{w_2 - b_{12}w_1^* - \mu_m^{(2)}}{\sqrt{\sigma_m^{(2)}}}\right)^{2k}}{2^k k!} \\ &= \frac{1}{\sqrt{2\pi}\sigma_m} \sum_{k=0}^{\infty} (-1)^k \frac{\left(w_2 - b_{12}w_1^* - \mu_m^{(2)}\right)^{2k}}{2^k k! (\sigma_m^{(2)})^k} \end{aligned}$$

We can expand $\left(w_2 - b_{12}w_1^* - \mu_m^{(2)}\right)^{2k}$ using the binomial theorem:

$$\left(w_2 - b_{12}w_1^* - \mu_m^{(2)}\right)^{2k} = \sum_{j=0}^{2k} \binom{2k}{j} (-1)^{2k-j} (b_{12}w_1^*)^{2k-j} (w_2 - \mu_m^{(2)})^j$$

Here, $\binom{2k}{j}$ represents a binomial coefficient, denoting the number of ways to choose j

items from a set of $2k$ distinct items.

$$f_{\epsilon_2}^{(m)}\left(w_2 - b_{12}w_1^*\right) = \frac{1}{\sqrt{2\pi}\sigma_m}\sum_{k=0}^{\infty}(-1)^k \frac{\sum_{j=0}^{2k}\binom{2k}{j}(-1)^{2k-j}(b_{12}w_1^*)^{2k-j}(w_2-\mu_m^{(2)})^j}{2^k k!(\sigma_m^{(2)})^k}$$

Taking expected value of $f_{\epsilon_2}^{(m)}\left(w_2 - b_{12}w_1^*\right)$ on m-th component of (ϵ_1, ϵ_2) over $W_1^* > 0$ denoted by
$E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)|i=m\right)$

$$\begin{aligned}
&E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)|i=m\right)\\
&= \frac{1}{\sqrt{2\pi}\sigma_m}\sum_{k=0}^{\infty}(-1)^{3k-j} \frac{\sum_{j=0}^{2k}\binom{2k}{j}b_{12}^{2k-j}(w_2-\mu_m^{(2)})^j E_{W_1^*|W_1=1}\left((w_1^*)^{2k-j}|i=m\right)}{2^k k!(\sigma_m^{(2)})^k}
\end{aligned}$$

Next, we consider $E_{W_1^*|W_1=1}\left((w_1^*)^{2k-j}|i=m\right)$, which is non-central moments of the truncated normal distribution (Burkardt, 2014):

$$E_{W_1^*|W_1=1}\left((w_1^*)^{2k-j}|i=m\right) = \sum_{r=0}^{2k-j}\binom{2k-j}{r}(\mu_m^{(1)})^{2k-j-r}(\sigma_m^{(1)})^{r/2}I_r \tag{24}$$

where $I_0 = 1$ and $I_{r-1} = \frac{\phi(h)}{1-\Phi(h)}$ and $I_r = \frac{h^{r-1}\phi(h)}{1-\Phi(h)} + (r-1)\ I_{r-2}$, $h = \frac{-\mu_m^{(1)}}{(\sigma_m^{(1)})^{1/2}}$. $\phi(\cdot)$ is the density of standard normal distribution, and $\Phi(\cdot)$ is standard normal cumulative distribution function. $\phi(\cdot)$ and $\Phi(\cdot)$ are analytic functions. Based on the property of analytic function, as long as $\Phi(h) \neq 1$, I_r is also analytic. Thus, when $\Phi(h) \neq 1$ for each m, $E_{W_1^*|W_1=1}\left((w_1^*)^{2k+1-j}|i=m\right)$ is an analytic function for each m. Furthermore, $E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(-b_{12}w_1^*\right)\right)$ is analytic. In a similar way, we want to prove
$E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)\right)$ is also analytic.

$E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)\right) = \sum_{i=1}^{M}\pi_i E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)|i=m\right)$, where $M = M_1 \cdot M_2$ and $\pi_i = \pi_i^{(1)}\pi_i^{(2)}$. Then

$$\begin{aligned}
&E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2 - b_{12}w_1^*\right)|i=m\right)\\
&= \frac{1}{\sqrt{2\pi}\sigma_m}\sum_{k=0}^{\infty}(-1)^{3k-j} \frac{\sum_{j=0}^{2k}\binom{2k}{j}b_{12}^{2k-j}(w_2-\mu_m^{(2)})^j E_{W_1^*|W_1=0}\left((w_1^*)^{2k-j}|i=m\right)}{2^k k!(\sigma_m^{(2)})^k}
\end{aligned}$$

Next, we consider $E_{W_1^*|W_1=0}\left((w_1^*)^{2k-j}|i=m\right)$, which is non-central moments of the trun-

cated normal distribution:

$$E_{W_1^*|W_1=0}\left((w_1^*)^{2k-j}|i=m\right)=\sum_{r=0}^{2k-j}\binom{2k-j}{r}(\mu_m^{(1)})^{2k-j-r}(\sigma_m^{(1)})^{r/2}I_r \tag{25}$$

where $I_0=1$ and $I_{r-1}=-\frac{\phi(h)}{\Phi(h)}$ and $I_r=-\frac{h^{r-1}\phi(h)}{\Phi(h)}+(r-1)\,I_{r-2}$, $h=\frac{-\mu_m^{(1)}}{(\sigma_m^{(1)})^{1/2}}$. $\phi(\cdot)$ is the density of standard normal distribution, and $\Phi(\cdot)$ is standard normal cumulative distribution function. $\phi(\cdot)$ and $\Phi(\cdot)$ are analytic functions. Based on the property of analytic function, as long as $\Phi(h)\neq 0$, I_r is also analytic. Thus, when $\Phi(h)\neq 0$ for each m, $E_{W_1^*|W_1=0}\left((w_1^*)^{2k+1-j}|i=m\right)$ is an analytic function for each m. Furthermore, $E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2-b_{12}w_1^*\right)\right)$ is also analytic. In summary, when $0<\Phi(h)<1$ for each m, both $E_{W_1^*|W_1=1}\left(f_{\epsilon_2}\left(w_2-b_{12}w_1^*\right)\right)$ and $E_{W_1^*|W_1=0}\left(f_{\epsilon_2}\left(w_2-b_{12}w_1^*\right)\right)$ are analytic. The sums, products, and compositions of analytic functions are analytic. Consequently, we finally prove that (23) is an analytic function. Therefore its zero set must have Lebesgue measure zero (Mityagin, 2015). To conclude, we have proven that the two causal models are not equivalent for almost all $(b_{12},\boldsymbol{\pi},\boldsymbol{\mu},\boldsymbol{\Sigma})$ with respect to the Lebesgue measure. Now let's study the relationship between $0<\Phi(h)<1$ for each m, and $P_{W_1\to W_2}(W_1=1)$, recall

$$P_{W_1\to W_2}(W_1=1)=\sum_{i=1}^{M_1}\pi_i^{(1)}P_{W_1\to W_2}(W_1^*>0|\mu_i^{(1)},\sigma_i^{(1)})=1-\sum_{i=1}^{M_1}\pi_i^{(1)}\Phi(h_i)$$

where $h_i=\frac{-\mu_i^{(1)}}{(\sigma_i^{(1)})^{1/2}}$ and $0<\Phi(h_i)<1$, for each i.

- $\diamond$ $\Phi(h_i)=0$, for each $i \iff P_{W_1\to W_2}(W_1=1)=1$,
- $\diamond$ $\Phi(h_i)=1$, for each $i \iff P_{W_1\to W_2}(W_1=1)=0$. It is straightforward to obtain $\Phi(h_i)=1$ for each i, which implies $P_{W_1\to W_2}(W_1=1)=0$. In another direction, $P_{W_1\to W_2}(W_1=1)=0\Rightarrow P_{W_1\to W_2}(W_1=0)=1$. Suppose there exists some i $\in A$, $\Phi(h_i)<1$, therefore,

  $$P_{W_1\to W_2}(W_1=0)=\sum_{i\in A^c}\pi_i^{(1)}+\sum_{i\in A}\pi_i^{(1)}\Phi(h_i)\leq\sum_{i\in A^c}\pi_i^{(1)}+\sum_{i\in A}\pi_i^{(1)}\max\left(\Phi(h_i)\right)<1,$$

  which contradicts with $P_{W_1\to W_2}(W_1=0)=1$. So comes $P_{W_1\to W_2}(W_1=1)=0\Rightarrow$ $\Phi(h_i)=1$, for each i.

Finally, under $0<P_{W_1\to W_2}(W_1=1)<1$, the *mixed-FLiNG-BN* is identifiable in the binary-continuous case of the scalar scenario.

Now, let's consider the case when $n > 1$. Drawing from the scalar scenario, it becomes evident that under given $\boldsymbol{w}_2$, there exists a vector $\boldsymbol{\delta}$ such that for any $\boldsymbol{w}_2'$ and $\boldsymbol{w}_2''$ in the vector space interval $(\boldsymbol{w}_2, \boldsymbol{w}_2 + \boldsymbol{\delta})$, the following equality holds:

$$P_{\boldsymbol{W}_2 \to \boldsymbol{W}_1}(\boldsymbol{W}_1 = \boldsymbol{w}_1 \mid \boldsymbol{W}_2 = \boldsymbol{w}_2', b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}) = P_{\boldsymbol{W}_2 \to \boldsymbol{W}_1}(\boldsymbol{W}_1 = \boldsymbol{w}_1 \mid \boldsymbol{W}_2 = \boldsymbol{w}_2'', b_{21}, \widetilde{\boldsymbol{\pi}}, \widetilde{\boldsymbol{\mu}}, \widetilde{\boldsymbol{\Sigma}}).$$

$0 < P(W_1(t_h) = 1) < 1$ for $h \in \{1, 2, \ldots, n\}$, the *mixed-FLiNG-BN* is identifiable in the binary-binary longitudinal scenario. This is due to the fact that we do not consider the relationship within $\boldsymbol{W}_i = (W_i(t_1), \ldots, W_i(t_n))^T \in \{0, 1\}^n$, for $i = 1, 2$. □

Theorem 4.A.3. *Given a pair of mixed functions $(\boldsymbol{W}_1, \boldsymbol{W}_2)$, each measured on n time points $W_i = (W_i(t_1), \ldots, W_i(t_n))^T \in \mathcal{R}^n$, under the certain assumption 4.2.3, the mixed-FLiNG-BN is bivariate identifiable.*

Proof. The identifiability has been well proved by Zhou et al. (2023). □

□

4.B Posterior inference

Update parameters of the mixture distribution. For each $j \in [p]$, $k \in [K]$, we introduce a latent variable $\xi_{jk}^{(i)} = m$ if observation $i \in [N]$ came from cluster m, and $P(\xi_{jk}^{(i)} = m) = \pi_{jkm}$.

1. Update $\xi_{jk}^{(i)}$ from its multinomial conditional posterior with

$$\mathrm{P}\left(\xi_{jk}^{(i)} = m \mid -\right) = \frac{\pi_{jkm} \mathrm{N}\left(\epsilon_{jk}^{(i)} \mid \mu_{jkm}, \tau_{jkm}\right)}{\sum_{m=1}^{M_{jk}} \pi_{jkm} \mathrm{N}\left(\epsilon_{jk}^{(i)} \mid \mu_{jkm}, \tau_{jkm}\right)}, \quad m = 1, \ldots, M_{jk}.$$

2. Update (μ_{jkm}, τ_{jkm}) from its conditional posterior distribution,

$$\mu_{jkm} \mid - \sim \mathrm{N}\left(\widehat{\mu}_{jkm}, \widehat{b}_\mu \tau_{jkm}\right), \quad \widehat{b}_\mu = \left(b_\mu^{-1} + n_m\right)^{-1}, \quad \widehat{\mu}_{jkm} = \widehat{b}_\mu \left(b_\mu^{-1} a_\mu + n_m \bar{\epsilon}_{jk(m)}\right),$$
$$\tau_{jkm} \mid - \sim \text{INV-GAMMA}\left(\hat{a}_\tau, \hat{b}_\tau\right), \quad \hat{a}_\tau = a_\tau + \frac{n_m}{2},$$
$$\hat{b}_\tau = b_\tau + \frac{1}{2}\left(\sum_{i:\xi_{jk}^{(i)}=m} \left(\epsilon_{jk}^{(i)} - \bar{\epsilon}_{jk(m)}\right)^2 + \left(\frac{n_m}{1 + b_\mu n_m}\right)\left(\bar{\epsilon}_{jk(m)} - a_\mu\right)^2\right),$$

where $n_m = \sum_{i=1}^{n} I_{\xi_{jk}^{(i)}=m}$ denotes the number of items allocated to component m in the current iteration, and $\bar{\epsilon}_{jk(m)} = n_m^{-1} \sum_{i:\xi_{jk}^{(i)}=m} \epsilon_{jk}^{(i)}$ is the mean within component m.

3. Update $\boldsymbol{\pi_{jk}} = (\pi_{jk1}, \ldots, \pi_{jkM_{jk}}) \mid - \sim \textsc{Dirichlet}\left(\alpha + n_1, \ldots, \alpha + n_{M_{jk}}\right)$.

Update parameters associated with the adjacency matrix $\boldsymbol{E}$ and the causal effect matrix $\boldsymbol{B}$. For each $l, j \in [p]$, $\boldsymbol{E} = (E_{jl}), \boldsymbol{B} = (\boldsymbol{B}_{jl})$.

1. Update the edge probability r from its conditional posterior distribution,

$$r \mid - \sim \textsc{Beta}(\widehat{a}_r, \widehat{b}_r), \quad \widehat{a}_r = a_r + \sum_{j=1,l=1}^{p} E_{jl}, \quad \widehat{b}_r = b_r - p + \sum_{j=1,l=1}^{p} (1 - E_{jl}),$$

   where p is the number of nodes.

2. Update the causal effect size γ with,

$$\gamma \mid - \sim \textsc{inv-gamma}(\widehat{a}_\gamma, \widehat{b}_\gamma),$$
$$\widehat{a}_\gamma = a_\gamma + K^2/2 \times \sum_{j=1,l=1}^{p} E_{jl}, \quad \widehat{b}_\gamma = 1 + \sum_{j=1,l=1}^{p} \sum_{s=1,t=1}^{K} B_{jl}^2(s,t)/2,$$

   where $B_{jl}(s,t)$ represents the element located in the s-th row and t-th column of the matrix $\boldsymbol{B}_{jl}$.

3. Update the causal effect matrix $\boldsymbol{B}$ from its conditional posterior distribution, for each $l, j \in [p]$,

   - if $E_{jl} = 0$, then $\boldsymbol{B}_{jl} = \boldsymbol{0}$.
   - if $E_{jl} = 1$, then assign l to $\mathscr{S}_j$ which is the set of indices satisfying $E_{jl} = 1$. Denote $\boldsymbol{B}_{j\mathscr{S}_j} = (\boldsymbol{B}_{jl})_{l \in \mathscr{S}_j}$, and $\boldsymbol{B}_{j\mathscr{S}_j}(k, \cdot)$ is the kth row of $\boldsymbol{B}_{j\mathscr{S}_j}$. For each $k = 1, \cdots, K$,

$$\boldsymbol{B}_{j\mathscr{S}_j}(k, \cdot) \mid - \sim \mathrm{MN}(\boldsymbol{\mu}_{jk}, \boldsymbol{\Omega}_{jk}^{-1}),$$
$$\boldsymbol{\Omega}_{jk} = \gamma^{-1}\boldsymbol{I} + \sum_{i=1}^{n} \left[\tau_{jk\xi_{jk}^{(i)}}^{-1} \cdot \boldsymbol{Z}_{\mathscr{S}_j}^{(i)} \cdot (\boldsymbol{Z}_{\mathscr{S}_j}^{(i)})^T\right], \quad \boldsymbol{\mu}_{jk} = \boldsymbol{\Omega}_{jk}^{-1} \sum_{i=1}^{n} \left[\tau_{jk\xi_{jk}^{(i)}}^{-1} \cdot Z_{jk}^{(i)} \cdot \boldsymbol{Z}_{\mathscr{S}_j}^{(i)}\right]$$

where $\boldsymbol{Z}_{\mathscr{S}_j} = (\boldsymbol{Z}_l)_{l\in\mathscr{S}_j}$, and Z_{jk} is the kth element of $\boldsymbol{Z}_j$.

4. Update the adjacency matrix $\boldsymbol{E}$. For each E_{jl}, we perform add-delete-swap move such that we generate a new candidate $E'_{jl} = 1 - E_{jl}$ with probability p, or swap candidates $E'_{jl} = E_{lj}$ and $E'_{lj} = E_{jl}$ with probability $1-p$. If the resulting $\boldsymbol{E}'$ is not acyclic, proceed directly to the next step. Otherwise, calculate

$$V_{jk}^{(i)} = \gamma \times I\left(E_{j\ell} = 1\right)\left(z_{\ell k}^{(i)}\right)^2, \quad V_j^{(i)'} = \gamma \times I\left(E'_{j\ell} = 1\right)\left(z_{\ell k}^{(i)}\right)^2,$$

$$\alpha = \sum_{i=1}^{n}\left[\mathrm{LN}\left(z_{jk}^{(i)}; \mu_{jk}^{(i)}, \tau_{jk}^{(i)} + V_j^{(i)}\right) - \mathrm{LN}\left(z_{jk}^{(i)}; \mu_{jk}^{(i)}, \tau_{jk}^{(i)} + V_{jk}^{(i)'}\right)\right] + (1 - 2E_{j\ell})\log\left(r^{-1}(1-r)\right),$$

where $\mathrm{LN}(x; \mu, s)$ denotes the log-normal density at x with mean μ and variance s. Then sample $u \sim U(0,1)$ and take the new proposal if $\alpha < \log(u)$. Here, we set the probability $p = 0.6$.

Update basis function coefficients sequences $\boldsymbol{Z}$.

For each $i \in [N]$, define $\boldsymbol{Z}^{(i)} = \left(\underbrace{\left(Z_{11}^{(i)}, \cdots, Z_{1K}^{(i)}\right)}_{\boldsymbol{z}_1^{(i)}}, \cdots, \underbrace{\left(Z_{p1}^{(i)}, \cdots, Z_{pK}^{(i)}\right)}_{\boldsymbol{z}_p^{(i)}}\right)$, then conditional posterior distribution of $\boldsymbol{Z}^{(i)}$ is

$$\boldsymbol{Z}^{(i)} \mid - \sim \mathrm{MN}(\boldsymbol{\mu}_i, \boldsymbol{\Omega}_i^{-1}),$$
$$\boldsymbol{\Omega}_i = \boldsymbol{\Omega} + (\boldsymbol{I} - \boldsymbol{B})\boldsymbol{T}_i^{-1}(\boldsymbol{I} - \boldsymbol{B})^T, \quad \boldsymbol{\mu}_i = \boldsymbol{\Omega}_i^{-1}\boldsymbol{\mu},$$

where $\boldsymbol{\Omega}$ is a diagonal matrix, with entries $\Omega(j) = \sum_{h=1}^{n} \boldsymbol{\phi}_{\boldsymbol{j}}^{(i)}(t_h)\boldsymbol{\phi}_{\boldsymbol{j}}^{(i)}(t_h)^T/\sigma_j$ and $\boldsymbol{T}_i$ is a block diagonal matrix, with entries $\boldsymbol{\tau}_j^{(i)} = \left(\tau_{j1\xi_{j1}^{(i)}}, \cdots, \tau_{jK\xi_{jK}^{(i)}}\right)$, $\boldsymbol{\mu} = \left(\sum_{h=1}^{n} \boldsymbol{W}_1^{*(i)}(t_h)\boldsymbol{\phi}_{\boldsymbol{j}}^{(i)}(t_h)^T/\sigma_1, \cdots, \sum_{h=1}^{n} \boldsymbol{W}_p^{*(i)}(t_h)\boldsymbol{\phi}_{\boldsymbol{p}}^{(i)}(t_h)^T/\sigma_p\right)^T$.

Update parameters associated with orthonormal basis functions ϕ_k.

1. Update the parameter λ_k from its conditional posterior distribution, for each $k \in [K]$

$$\lambda_k \mid - \sim \textsc{Truncated-Gamma}(\alpha, \beta_k; L_k, U_k),\ \alpha = (L-2)/2,\ \beta_k = 2^{-1}\sum_{l=1}^{L} A_{kl}^2.$$

2. Update the transformed B-spline coefficients $\widetilde{\boldsymbol{A}}_k$ with

$$\widetilde{\boldsymbol{A}}_k^U \mid - \sim \mathrm{MN}(\boldsymbol{\mu}_k, \boldsymbol{Q}_k^{-1}),$$

$$\boldsymbol{Q}_k = \left\{ \sum_{i=1}^{n} \sum_{j=1}^{p} \frac{\left(Z_{jk}^{(i)}\right)^2}{\sigma_j} \sum_{h=1}^{n} \boldsymbol{b}(t_h)\boldsymbol{b}(t_h)^T \right\} + \boldsymbol{\Omega}_k^{-1},$$
$$\boldsymbol{\mu}_k = \boldsymbol{Q}_k^{-1} \left[\sum_{i=1}^{N} \sum_{j=1}^{p} \sum_{h=1}^{n} \frac{Z_{jk}^{(i)}}{\sigma_j} \overline{W}_{jh,-k}^{*(i)} \boldsymbol{b}(t_h) \right].$$

Now, denote $\boldsymbol{P}_k = \boldsymbol{J}\tilde{\boldsymbol{A}}_{-k}$, where $\boldsymbol{J} = \int \boldsymbol{b}(t_h)\boldsymbol{b}^T(t_h)dt_h$. Finally, transform and normalize the unconstrained sample to $\tilde{\boldsymbol{A}}_k^N = \tilde{\boldsymbol{A}}_k^U - \boldsymbol{Q}_k\boldsymbol{P}_k \left(\boldsymbol{P}_k^T\boldsymbol{Q}_k\boldsymbol{P}_k\right)^{-1} \boldsymbol{P}_k\tilde{\boldsymbol{A}}_k^U$ and $\tilde{\boldsymbol{A}}_k = \tilde{\boldsymbol{A}}_k^N \times \left(\left[\tilde{\boldsymbol{A}}_k^N\right]^T \boldsymbol{J}\tilde{\boldsymbol{A}}_k^N \right)^{-1/2}$.

Update the variance of noise σ_j. For each $j \in [p]$, $i \in [N]$ and $h \in [n]$, introduce $\overline{W}_j^{*(i)}(t_h) = W_j^{*(i)}(t_h) - \sum_{k=1}^{K} Z_{jk}^{(i)}\phi_k^{(i)}(t_h)$. The conditional posterior distribution of σ_j is as follows,

$$\sigma_j \mid - \sim \textsc{inv-gamma}(\widehat{a}_\sigma, \widehat{b}_\sigma),$$
$$\widehat{a}_\sigma = a_\sigma + 2^{-1} \cdot nN, \quad \widehat{b}_\sigma = b_\sigma + 2^{-1} \sum_{i=1}^{N} \sum_{h=1}^{n} \left(\overline{W}_j^{*(i)}(t_h)\right)^2.$$

Update the latent continuous function $\boldsymbol{W}_j^*$. Given each binary function $\boldsymbol{W}_j$ measured on n time points $\boldsymbol{W}_j = (W_j(t_1), \ldots, W_j(t_n))^T \in \{0,1\}^n$, the associated latent continuous function $\boldsymbol{W}_j^* = (W_j^*(t_1), \ldots, W_j^*(t_n))^T \in \mathcal{R}^n$. For each binary function j, and time point $h \in [n]$, update $W_j^*(t_h)$ from the truncated normal distribution,

$$W_j^*(t_h) \mid \boldsymbol{\phi}_j, \boldsymbol{Z}_j, \sigma_j^2 \sim \left\{ \begin{array}{ll} \mathrm{TN}_{[0,\infty)}\left(\sum_{k=1}^{K} Z_{jk}\phi_{jk}, \sigma_j^2\right), & W_j(t_h) = 1 \\ \mathrm{TN}_{(-\infty,0)}\left(\sum_{k=1}^{K} Z_{jk}\phi_{jk}, \sigma_j^2\right), & W_j(t_h) = 0 \end{array} \right\},$$

where $\mathrm{TN}_{\mathscr{A}}$ denotes a truncated normal density in the interval $\mathscr{A}$.

4.C Simulations: sensitivity analysis

The proposed mixed-FLiNG-BN involves several hyperparameters α, (a_μ, b_μ), (a_τ, b_τ), (a_r, b_r), (a_γ, b_γ), and (a_σ, b_σ). We conduct sensitivity analyses on these hyperparameters with $(p, h, N) = (20, 50, 200)$ where five of the nodes are binary and the rest are continuous, revealing that our model exhibits relative robustness across the tested ranges.

	α	(a_μ, b_μ)	(a_τ, b_τ)	(a_r, b_r)	(a_γ, b_γ)	(a_σ, b_σ)
	0.1	(0, 0.1)	(0.1, 0.1)	(0.5, 0.5)	(0.1, 0.1)	(0.1, 0.1)
TPR	0.68 (0.05)	0.65 (0.03)	0.67 (0.07)	0.64 (0.02)	0.69 (0.08)	0.62 (0.07)
FDR	0.16 (0.06)	0.12 (0.05)	0.20 (0.06)	0.17 (0.04)	0.17 (0.07)	0.22 (0.08)
MCC	0.74 (0.05)	0.74 (0.02)	0.71 (0.07)	0.71 (0.02)	0.74 (0.05)	0.70 (0.05)
	0.5	(0, 0.5)	(0.1, 1)	(0.2, 0.8)	(1, 2)	(0.01, 0.01)
TPR	0.72 (0.07)	0.72 (0.06)	0.63 (0.08)	0.64 (0.03)	0.63 (0.07)	0.63 (0.06)
FDR	0.19 (0.09)	0.11 (0.05)	0.17 (0.02)	0.17 (0.08)	0.15 (0.08)	0.18 (0.08)
MCC	0.75 (0.05)	0.78 (0.04)	0.70 (0.04)	0.71 (0.06)	0.71 (0.02)	0.70 (0.05)
	2	(0, 1)	(10, 1)	(0.9, 0.1)	(1, 0.1)	(0.001, 0.001)
TPR	0.68 (0.03)	0.66 (0.04)	0.69 (0.07)	0.66 (0.09)	0.65 (0.08)	0.65 (0.08)
FDR	0.22 (0.03)	0.12 (0.05)	0.21 (0.06)	0.14 (0.07)	0.18 (0.07)	0.19 (0.05)
MCC	0.71 (0.03)	0.74 (0.02)	0.72 (0.06)	0.74 (0.06)	0.71 (0.07)	0.71 (0.04)
	3	(0, 2)	(20, 5)	(0.1, 0.1)	(2, 1)	(1, 1)
TPR	0.73 (0.05)	0.65 (0.08)	0.76 (0.06)	0.71 (0.08)	0.66 (0.08)	0.64 (0.06)
FDR	0.17 (0.08)	0.12 (0.08)	0.17 (0.03)	0.20 (0.07)	0.16 (0.07)	0.18 (0.08)
MCC	0.76 (0.07)	0.74 (0.05)	0.78 (0.02)	0.74 (0.05)	0.73 (0.02)	0.70 (0.05)

Table 4.2: *Sensitivity analysis with respect to the choices of hyperparameters with* $(p, h, N) = (20, 50, 200)$ *where five of the nodes are binary and the rest are continuous. The reported average performance is based on 50 repetitions, and standard deviations are given within the parentheses.*

Chapter 5

Conclusion

In this dissertation, we develop advanced Bayesian methods that tackled the fundamental challenges in personalized medicine, specifically variable selection in high dimensional data and causal structure learning in mixed functional data. Through extensive simulation studies and applications to real-world clinical datasets, we demonstrate the clinical utility of these Bayesian approaches in aiding physicians to make more effective treatment decisions.

In Chapter 2, we propose a Bayesian model for optimizing dynamic treatment regimes, a practical framework within personalized medicine, to address the uncertainty in identifying optimal decision sequences and incorporate dimensionality reduction to manage high-dimensional individual covariates. To achieve this, the model is augmented to accommodate counterfactual variables. Additionally, we introduce a novel class of spike-and-slab priors for the multi-stage selection of significant factors, to favor the sharing of information across stages.

In Chapter 3, we extend the proposed class of spike-and-slab priors, particularly from two stages to four stages, enabling the identification of significant molecular biomarkers and clinical parameters influencing drug effects on estimated Glomerular Filtration Rate (eGFR) using type II diabetes dataset. This dataset is derived from the PROVALID study, a prospective observational study encompassing a diverse range of chronic kidney disease stages.

In Chapter 4, we introduce mixed Bayesian networks to tackle causal structure learning across diverse data types, facilitating a more accurate representation of causal dependencies by a directed acyclic graph. We present that under this model, the causal structure between mixed functional data can be uniquely identifiable under certain mild conditions. Furthermore, we develop a Bayesian framework to infer the adjacency matrix of the directed acyclic graph with inherent uncertainty quantification.

There are several avenues for future work and improvement. First, in Bayesian augmented learning, we refine the dependencies among the indicator variables that encode the inclusion of regressors, tailored to the specific problem at hand. It would be valuable to investigate the potential for sharing information not only within the same variable across multiple stages but also among different variables, both within and across those stages. Secondly, we apply Bayesian augmented learning to analyze type II diabetes using a dataset comprising four stages. It would be interesting to implement dynamic variable selection; for example, the correlation of prior inclusion probabilities for the same regressor may be stronger between stage 1 and stage 2 compared to stage 1 and stage 4. Finally, further exploration is required for numerical studies and theoretical proofs of causal identification between binary functions within the mixed functional Bayesian networks proposed.

Acknowledgement

First, I would like to express my heartfelt gratitude to my supervisor and tutor, Professor Bernardo Nipoti, for his invaluable guidance and unwavering support throughout my PhD journey. His remarkable scientific insight and dedication to research have been a constant source of inspiration, and his brilliant approach to academic life has profoundly impact on my own aspirations in academia.

Next, I am grateful to Professor Matteo Borrotti, who has been instrumental in guiding me through my dissertation research, offering boundless support with exceptional humility. I also want to express my appreciation to Professor Yanxun Xu, for motivating me to pursue my research goals and hosting me in the U.S. for a year. I am also thankful to Professor Yang Ni, for his patient and thorough guidance in mathematical aspects. My appreciation also extends to Dr. Wei Jin, who offers valuable insights and ideas on my research. I am thankful as well to my collaborators, Dr. Seyum Abebe and Professor Irene Poli.

Furthermore, I extend my gratitude to all faculty and staff in the Department of Economics, Management and Statistics (DEMS) of the University of Milano-Bicocca for fostering a friendly and enriching learning environment. I am thankful to my friends in Milan, the U.S., and China for their companionship and support throughout this journey.

Lastly, I would like to express my deepest gratitude to my family, especially my parents, for their unconditional love and unwavering support throughout my life.