

Dynamic Treatment Regimes for Within- and Between-Group Interference in Clustered and Hierarchical Datasets

by

Alexandra Karin Mossman

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Statistics (Biostatistics)

Waterloo, Ontario, Canada, 2025

© Alexandra Karin Mossman 2025

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

- External Examiner: Dongling Zeng
Professor, Dept. of Biostatistics, School of Public Health
University of Michigan
- Supervisor(s): Michael Wallace
Associate Professor, Dept. of Statistics & Actuarial Science,
University of Waterloo
Yeying Zhu
Associate Professor, Dept. of Statistics & Actuarial Science,
University of Waterloo
- Internal Members: Joel Dubin
Professor, Dept. of Statistics & Actuarial Science,
University of Waterloo
Lan Wen
Assistant Professor, Dept. of Statistics & Actuarial Science,
University of Waterloo
- Internal-External Member: Helen Chen
Professor of Practice, School of Public Health Sciences,
University of Waterloo

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Precision medicine is an interdisciplinary field that aims to tailor treatments based on an individual’s unique characteristics. Dynamic treatment regimes (DTRs) formalize this process through step-by-step decision rules that utilize patient-specific information at each stage of analysis and subsequently output recommendations for the optimal course of action. To date, much of the biostatistical literature has analyzed DTR estimation under the assumption of no interference; that is, a given individual’s outcome is not affected by the treatments received by others. However, this assumption is often violated in a variety of social and spatial networks, such as households and communities, and particularly in the context of infectious diseases or resource allocation. Although some recent developments have been made for DTR estimation for couples in households and in networks where general forms of interference are taking place, it has yet to be shown how DTRs can be estimated for individuals in clustered networks where interference may occur within and between predefined groups. Specifically, our attention shifts toward hierarchical networks, which provide a framework where interference can occur both within and between groups in the same hierarchy but not across hierarchies.

This thesis contains three main projects that contribute to DTR estimation in the context of service utilization within the healthcare system and under interference networks. In Chapter 3, we show how the dynamic weighted ordinary least squares regression (dWOLS) DTR methodology can be applied to determine whether a patient would benefit from being discharged to home or admitted to a hospital using data simulated from a retrospective cohort study for patients who experienced an opioid-related overdose in British Columbia. This project was inspired from collaboration with a provincial health authority in Canada, which has resulted in the draft of a manuscript that can be submitted for publication following an ethics approval. In Chapter 4, we demonstrate an application of dWOLS to hierarchical networks where both within- and between-group interference takes place, and through a series of simulation studies, we show how incorrectly assuming that there only exists within-group interference can result in lower rates of optimal treatments assigned, particularly as the number of subgroups (and thus the potential for between-group interference) increases. Our models assess the performance of “standard” overlap weights constructed from logistic mixed models with nested random intercepts as well as network weights that are constructed from a joint propensity score. We apply our findings to the simulated opioid overdose dataset, where our goal is to illustrate how sequences of recommended dispositions for opioid overdose patients can be optimized if interference occurs within and between different facilities. Finally, in Chapter 5, we propose several sets of simulation studies to assess the extent to which interference may be taking place between individuals to warrant use of modified dWOLS methodology that accounts for interference.

Acknowledgements

I would like to thank my supervisors, Dr. Michael Wallace and Dr. Yeying Zhu, for their ongoing support, feedback, and encouragement over the past four years. You have both been, and continue to be, role models in my pursuit of biostatistics, and I appreciate you providing me with many opportunities to grow as an independent researcher. Michael, I chose biostatistics as my career path when I first attended a talk you gave on precision medicine at the University of Calgary in summer 2019. Yeying, STAT 931 was one of my favourite courses that I took in my graduate studies. I respect you both and appreciate everything you have done to help me succeed throughout my PhD.

I would also like to thank Dr. Dongling Zeng, Dr. Helen Chen, Dr. Joel Dubin, and Dr. Lan Wen for kindly serving as my committee members and for providing invaluable feedback on my thesis work. To my friends Gracia Dong, Bryn Crandles, Megan French, Luke Hagar, Minzee Kim, Mohan Wu, Xianwei Li, Kecheng Li, Tatiana Krikella, Victor Malinowski, Madison Cranstone, Ken Mawer, and everyone else I've met since coming to Waterloo: thank you for helping me take a break, laugh, get coffee, and stay motivated during some of my most difficult times in the past few years. I am incredibly fortunate to have spent my graduate studies with you all, and you've made Waterloo feel like home to me.

To Jessica Churcher and Lucas Brown: thank you for being my family, for your endless encouragement and support, for making sure we always had a well-stocked fridge and delicious home-cooked meals, and for making me laugh at all of our inside jokes in the evenings over our favourite TV shows as the three of us navigated different challenges over the past few years. I love you both and am so thankful to have lived with you for the past two years.

I would also like to thank Mary Lou Dufton, who has supported me and provided me with guidance since I started my Master's degree in the department nearly five years ago. You are a mentor and friend to everyone in M3, and I can't tell you enough how much I've appreciated you helping me stay on track in pursuit of academic success. Even on days when I felt stuck or hopeless, I knew I could never quit because you wouldn't let me leave until I achieved all my goals here in Waterloo.

I would also like to thank Dad, Mom, Grandma, and Christopher. I love you all more than you'll ever know. Thank you for always being a phone call away whenever I missed you and needed someone to talk to these last four years.

To Matthew Robbins: you are my inspiration, my joy, and my best friend, and I am thankful every day to have met you. Thank you for all the happiness, love, and kindness

you bring into my life; for never ceasing to make me smile; for your endless support as I finished my degree; and for making learning a lifelong adventure with me. I love you dearly.

Finally, I would like to thank my Grandpa. I promised you that I would complete my PhD, and I wish I could have told you in person that I was able to finish writing my thesis this summer. Your words carried me through some of the most tumultuous periods of my life that I've experienced this past year, and I love you and miss you every day.

Dedication

This thesis is dedicated to my grandparents, Mary and James Mossman, who have always shown me unconditional love and support throughout my life.

Table of Contents

Examining Committee	ii
Author's Declaration	iii
Abstract	iv
Acknowledgements	v
Dedication	vii
List of Figures	xii
List of Tables	xvi
List of Abbreviations	xviii
1 Introduction	1
1.1 Introduction to Precision Medicine and Dynamic Treatment Regimes . . .	1
1.2 Causal Inference in Observational Studies	4
1.3 Interference and Dynamic Treatment Regimes	6
1.3.1 Interference Structures	8
1.3.2 Interference Networks for Clustered Data with Mixed Effects Models	11
1.3.3 Dynamic Treatment Regimes under Interference Networks	12

1.4	Service System Utilization for Opioid Overdose Patients within and between Facilities	14
1.5	Overview of Thesis Organization	17
2	Notation and Methodology for Dynamic treatment regime (DTR)s with Network Interference	19
2.1	Causal Inference with Interference: Notation and Assumptions	19
2.2	Statistical Reinforcement Learning for Dynamic Treatment Regimes	21
2.2.1	Q-Learning with Linear Models	24
2.2.2	Q-Learning Using Observational Data	26
2.2.3	G-Estimation	26
2.3	Dynamic Weighted Ordinary Least Squares Regression (dWOLS)	28
2.4	Optimal Treatment Rules in the Presence of Interference	30
2.5	Propensity Score Weights for Network Interference	31
2.6	dWOLS in the Presence of Interference with Network Propensity Weights	33
2.7	Propensity Score Modelling for Multilevel Data	36
2.7.1	Network Weights for Multilevel Data	37
3	Optimizing Post-ED Visit Dispositions Among Opioid Overdose Patients Using Age-Tailored Guidelines with Dynamic Weighted Ordinary Least Squares	39
3.1	Introduction	39
3.1.1	Overview of the Opioid Overdose Crisis in Canada	40
3.1.2	Motivating Example: A Personalized Treatment Approach to Opioid Overdose Harm Reduction and Prevention	41
3.2	Application: Optimizing Disposition Rules for Opioid Overdose Patients with an Age-Tailored DTR	43
3.2.1	Data Pre-Processing and Methods	44
3.3	Results	47
3.4	Discussion	48

4	dynamic weighted ordinary least squares regression (dWOLS) with Mixed Effects Propensity Score Models in Hierarchical Networks under Between- and Within-Group Interference	50
4.1	Research Objective	50
4.2	Simulation Studies: One- and Two-Stage Settings	51
4.2.1	<i>Study 1a</i>	51
4.2.2	<i>Study 1b</i>	55
4.2.3	<i>Study 1c</i>	61
4.3	Application to Service Class Utilization in Opioid Overdose Patients	65
4.3.1	Data Pre-Processing and Methods	65
4.3.2	Ego and Alter Assignments Within and Between Facilities	70
4.3.3	Results	70
4.4	Discussion	72
5	Application of dWOLS with Network Propensity Weights with Varying Extents of Interference	74
5.1	Introduction	74
5.2	Simulation Studies: Varying the Overall Proportion of Egos Assigned with Alters	77
5.2.1	<i>Study 2a</i>	77
5.2.2	Results of One-Stage Settings: Varying the Overall Proportion of Egos Assigned with Alters	80
5.2.3	<i>Study 2b</i>	80
5.3	Simulation Studies: Varying the Ratio of Egos Assigned Between-Group Alters to Egos Assigned Within-Group Alters	86
5.3.1	Results of One-Stage Settings: Varying the Ratio of Egos Assigned Between-Group Alters to Egos Assigned Within-Group Alters	94
5.4	Discussion	94
6	Conclusion and Future Work	99

References	102
APPENDICES	111
A Chapter 2 Appendix	112
A.1 Linear Mixed Models	112
B Chapter 3 Appendix	114
B.1 Application of dWOLS to Simulated Opioid Overdose Data with Changes to Age in Stage 2	114
B.2 Results	115
C Chapter 4 Appendix	117
C.1 Additional Results from <i>Study 1a</i>	117
C.2 Additional Results from <i>Study 1b</i>	123
C.3 Additional Results from <i>Study 1c</i>	126

List of Figures

1.1	“Personalized Cancer Therapy” by Jeyang259, used under CC BY-SA 4.0 (https://creativecommons.org/licenses/by-sa/4.0/), via Wikimedia Commons: https://commons.wikimedia.org/wiki/File:Personalized_Cancer_Therapy.png	2
1.2	Illustration of a DAG for two individuals in a simple interference network when the Stable unit treatment value assumption (SUTVA) is violated. Here, X_{ego} , A_{ego} , and Y_{ego} represent the covariate, treatment, and outcome of the ego, respectively, while X_{alter} , A_{alter} , and Y_{alter} represent the covariate, treatment, and outcome of the alter, respectively.	6
1.3	Illustration of a partial interference structure for smoking cessation strategies within households.	14
1.4	Illustration of within- and between-group interference, where different facilities are nested within an overall region.	17
3.1	Illustration of data for a two-stage analysis using simulated data from a patient with three recorded ED visits.	44
3.2	Overview of selection criteria for 100 simulated overdose datasets.	45
4.1	DAG of simulation studies in <i>Study 1a</i> , where $\mathbf{X}_N$ denotes the individual-level covariates of the ego’s alters, $\mathbf{A}_N$ denotes the treatments of the ego’s alters, V denotes a group-level variable corresponding to the hierarchy that the ego and alter belong to, and Y_{ego} represents the ego’s outcome.	51
4.2	Hierarchical network structure, with interference occurring within and between Level 2 groups as depicted using arrows within and between individuals at Level 1.	52

4.3	Boxplots of ψ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 1 in <i>Study 1a</i> , where the treatment model is mis-specified via omission of the non-linear term in X_{ijk}	56
4.4	Boxplots of ψ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 2 in <i>Study 1a</i> , where the treatment-free model is mis-specified via omission of the non-linear term in X_{ijk}	57
4.5	Boxplots of ξ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 1 in <i>Study 1a</i> , where the treatment model is mis-specified via omission of the non-linear term in X_{ijk}	58
4.6	Boxplots of ξ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 2 in <i>Study 1a</i> , where the treatment-free model is mis-specified via omission of the non-linear term in X_{ijk}	59
4.7	DAG of simulation studies in <i>Study 1c</i> , where $\mathbf{X}_{1,\mathcal{N}}$ and $\mathbf{X}_{2,\mathcal{N}}$ denote the individual-level covariates of the ego's alters at stages 1 and 2, respectively, while $\mathbf{A}_{1,\mathcal{N}}$ and $\mathbf{A}_{2,\mathcal{N}}$ denote the treatments of the ego's alters at stages 1 and 2, respectively.	61
4.8	Boxplots of ψ parameter estimates for $J = 2$ in a two-stage setting in <i>Study 1c</i> , where the treatment model is mis-specified via omission of the non-linear term. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.	66
4.9	Boxplots of ξ parameter estimates for $J = 2$ in a two-stage setting in <i>Study 1c</i> , where the treatment model is mis-specified via omission of the non-linear term. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.	67
4.10	Boxplots of ψ parameter estimates for $J = 2$ in a two-stage setting in <i>Study 1c</i> , where the treatment-free model is mis-specified by using linear terms in $X_{1,ijk}$, $X_{2,ijk}$, $V_{1,i}$, and $V_{2,i}$. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.	68
4.11	Boxplots of ξ parameter estimates for $J = 2$ in a two-stage setting in <i>Study 1c</i> , where the treatment-free model is mis-specified by using linear terms in $X_{1,ijk}$, $X_{2,ijk}$, $V_{1,i}$, and $V_{2,i}$. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.	69

5.1	Blip parameter estimates when the treatment model is mis-specified under a low probability of treatment assignment, $P(A = 1 X = x, V = v) = 25\%$.	82
5.2	Blip parameter estimates when the treatment model is mis-specified under a high probability of treatment assignment, $P(A = 1 X = x, V = v) = 75\%$.	83
5.3	Blip parameter estimates when the treatment free model is mis-specified under a low probability of treatment assignment, $P(A = 1 X = x, V = v) = 25\%$.	84
5.4	Blip parameter estimates when the treatment free model is mis-specified under a high probability of treatment assignment, $P(A = 1 X = x, V = v) = 75\%$.	85
5.5	Blip parameter estimates under ME2 in <i>Study 2b</i> when 20% of egos are assigned an alter in stages one and two ((a) and (b), respectively) versus when 80% of egos are assigned an alter ((c) and (d), respectively).	87
5.6	Blip parameter estimates under ME4 in <i>Study 2b</i> when 20% of egos are assigned an alter in stages one and two ((a) and (b), respectively) versus when 80% of egos are assigned an alter ((c) and (d), respectively).	88
5.7	Spillover parameter estimates under ME2 in <i>Study 2b</i> when 20% of egos are assigned an alter in stages one and two ((a) and (b), respectively) versus when 80% of egos are assigned an alter ((c) and (d), respectively).	89
5.8	Spillover parameter estimates under ME4 in <i>Study 2b</i> when 20% of egos are assigned an alter in stages one and two ((a) and (b), respectively) versus when 80% of egos are assigned an alter ((c) and (d), respectively).	90
5.9	Comparison of mean optimal treatment rates when the treatment model is mis-specified (a) versus correctly specified (b).	95
5.10	Blip parameter estimates when the treatment model is mis-specified.	96
5.11	Blip parameter estimates when the treatment model is correctly specified.	97
C.1	Boxplots of ψ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 3, where both the treatment and treatment-free models are mis-specified via omission of the non-linear term in X_{ijk} .	119
C.2	Boxplots of ψ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 4, where both the treatment and treatment-free models are correctly specified.	120

C.3	Boxplots of ξ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 3, where both the treatment and treatment-free models are mis-specified via omission of the non-linear term in X_{ijk}	121
C.4	Boxplots of ξ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 4, where both the treatment and treatment-free models are correctly specified.	122
C.5	Boxplots of ψ parameter estimates under Scenario (a) in <i>Study 1b</i>	124
C.6	Boxplots of ψ parameter estimates under Scenario (b) in <i>Study 1b</i>	125
C.7	Boxplots of ξ parameter estimates under Scenario (a) in <i>Study 1b</i>	126
C.8	Boxplots of ξ parameter estimates under Scenario (b) in <i>Study 1b</i>	127
C.9	Boxplots of ψ parameter estimates under Scenario (a) in Table C.3.	128
C.10	Boxplots of ψ parameter estimates under Scenario (b) in Table C.3.	129
C.11	Boxplots of ξ parameter estimates under Scenario (a) in Table C.3.	132
C.12	Boxplots of ξ parameter estimates under Scenario (b) in Table C.3.	133
C.13	Boxplots of ψ parameter estimates for $J = 2$ in a two-stage setting under Scenario 2 in <i>Study 1c</i> where the treatment-free model is mis-specified by using linear terms in V_i only. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.	134
C.14	Boxplots of ξ parameter estimates for $J = 2$ in a two-stage setting under Scenario 2 in <i>Study 1c</i> where the treatment-free model is mis-specified by using linear terms in V_i only. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.	135

List of Tables

3.1	Average proportions of CTAS scores and dispositions across males and females at stages 1 and 2 of analysis from 100 simulated datasets.	47
3.2	Mean blip parameter estimates and regrets from 100 simulations.	48
4.1	Mean optimal treatment rates for Scenario 1 in <i>Study 1a</i> , where the treatment model is mis-specified via omission of the non-linear term in X_{ijk} . . .	55
4.2	Mean optimal treatment rates for Scenario 2 in <i>Study 1a</i> , where the treatment-free model is mis-specified via omission of the non-linear term in X_{ijk}	58
4.3	Overview of methods for comparison under each scenario in <i>Study 1b</i> , with a ✓ indicating which weights have been implemented in a given method and which component of the outcome model has been correctly specified, while an 7 indicates which weights have not been implemented and which component of the outcome model has been incorrectly specified.	62
4.4	Optimal treatment rates under Scenarios (a) and (b) in <i>Study 1b</i>	63
4.5	Mean blip and spillover parameter estimates from 100 different ego-alter assignments in the simulated opioid overdose dataset.	71
5.1	Mean optimal treatment rates, Treatment Model Mis-specified, $P(A = 1 X = x, V = v) = 25\%$	81
5.2	Mean optimal treatment rates, Treatment Model Mis-specified, $P(A = 1 X = x, V = v) = 75\%$	81
5.3	Mean optimal treatment rates, Treatment-Free Model Mis-specified, $P(A = 1 X = x, V = v) = 25\%$	81
5.4	Mean optimal treatment rates, Treatment-Free Model Mis-specified, $P(A = 1 X = x, V = v) = 75\%$	86

B.1	Average proportions of CTAS scores and dispositions across males and females at stages 1 and 2 of analysis from 100 simulated datasets.	115
B.2	Mean blip parameter estimates and regrets from 100 simulations.	116
C.1	Mean optimal treatment rates for Scenario 3 in <i>Study 1a</i> , where both the treatment and treatment-free models are incorrectly specified via omission of the non-linear term in X_{ijk}	118
C.2	Mean optimal treatment rates for Scenario 4 in <i>Study 1a</i> , where both the treatment and treatment-free models are correctly specified.	118
C.3	Extension of <i>Study 1b</i> where only one of the ego's, alter-W's, or alter-B's treatment-free model is correctly specified.	130
C.4	Optimal treatment rates under Scenarios (a) and (b) in <i>Study 1b</i> when only one component of the treatment-free model (the ego's or one of the alter's) is correctly specified	131

List of Abbreviations

BC British Columbia. 15, 40, 42, 43, 99

DAG Directed acyclic graph. 6

DTR Dynamic treatment regime. 2–4, 12, 13, 15–17, 19, 21, 22, 25, 26, 28, 33, 35, 36, 39, 42, 48, 49, 73, 99–101, 106, 107

dWOLS dynamic weighted ordinary least squares regression. 3, 12, 16–18, 28, 29, 33–35, 42, 43, 47, 49–51, 53, 54, 72–79, 92–94, 99, 100, 106, 108

ED Emergency department. 14–18, 39–44, 47–49, 70, 73, 99–101, 105

OD Opioid use disorder. 13–15, 39–42, 49, 99

pWOD persons with OD. 14–16, 40, 41, 49, 105

SUTVA Stable unit treatment value assumption. 6–9, 12, 14, 15, 36, 40, 99

Chapter 1

Introduction

1.1 Introduction to Precision Medicine and Dynamic Treatment Regimes

Precision medicine is a discipline that overlaps with areas of study in the health, mathematical, and computational sciences. The overall concept of precision medicine is to develop prevention and treatment strategies that account for heterogeneity in response to treatment, thereby improving patient care based on a person’s unique background and needs [16]. Increased accessibility to large-scale biologic databases, and development of computational tools for their analysis, have motivated creative approaches to precision medicine to help guide clinical practice [16], as illustrated in Figure 1.1. Indeed, as health researchers shift their focus from “standardized” treatment practices to individually-tailored ones, data have been described as the link between two “revolutions”: that of medical practice transitioning from “one-size-fits-all” to more personalized treatments, and the emergence of big data techniques [27]. As “evidence-based medicine” refers to the use of relevant data to guide medical decisions, it should be noted that precision medicine is always evidence-based, though not all evidence-based medicine is personalized; for instance, “evidence-based medicine” is frequently used in the context of research studies that investigate the response of an average patient to a particular treatment, thereby implying “one-size-fits-all” protocols [27]. With the ever-increasing accessibility and management of data, appropriate statistical methodology must be developed to provide effective analyses of large-scale datasets and, subsequently, to help inform decision-making in clinical settings.

Although precision medicine has emerged as its own research initiative fairly recently




Figure 1.1: “Personalized Cancer Therapy” by Jeyang259, used under CC BY-SA 4.0 (<https://creativecommons.org/licenses/by-sa/4.0/>), via Wikimedia Commons: https://commons.wikimedia.org/wiki/File:Personalized_Cancer_Therapy.png

[16], researchers have considered the effects of individual variation on treatment procedures well before the last decade. Sorensen [70] has described it as “intuitive” to consider a patient’s own unique characteristics and medical background when making treatment decisions in their best interest, further acknowledging how clinical trials tend to compare treatment options among patients who are overall homogeneous according to the criteria of interest. Longford and Nelder [42] have emphasized that failure to acknowledge treatment heterogeneity can result in undesirable consequences, as a large standard deviation of treatment effects corresponding to “highly variable” drugs may be associated with some patients experiencing effects that are very different from the mean treatment effect. Indeed, even aspects of more “standardized” or “one-size-fits-all” clinical practice acknowledge the importance of accounting for individual variation, such as using blood typing to guide blood transfusions [16]. As such, there is much motivation for developing appropriate methodology to implement patient-specific information into clinical guidelines that can be readily interpreted by health researchers.

A dynamic treatment regime (DTR) is a sequence of decision rules that uses individual-level covariates and treatment information (i.e. their *history*) to provide a recommended course of action as output for one or more stages of decision making [13]. Also referred to as treatment strategies, adaptive treatment strategies, and treatment policies, DTRs offer much applicability toward multi-stage decision problems in the context of chronic illnesses [13]. For example, when treating a patient with cancer, a clinician may use a DTR to help maximize disease-free survival time by monitoring the patient’s response to induction therapy and then deciding whether to maintain remission with a maintenance therapy or attempt a second-line induction therapy to induce remission in the case of

non-responsiveness [13]. A clinician may also use a DTR when treating a patient with generalized anxiety disorder by monitoring the patient’s response to a given dosage, such as improvement in symptoms or emergence of undesirable side effects, and then make any adjustments as necessary over time. DTR development constitutes a core area of study within the field of precision medicine when utilizing subject-specific characteristics to optimize patient outcome over courses of treatment, and it motivates a rigorous framework within which they can be constructed through the expertise of statisticians and health researchers.

Within a statistical framework, determining the optimal DTR over one or more stages of treatment translates to determining the sequence of decision rules that maximizes a patient’s expected outcome of interests conditional on their covariate and treatment history at each decision point [81]. In this context, Murphy [50] proposed viewing the observed outcome as being equal to some optimal outcome (not necessarily observed) minus penalty terms at stages where non-optimal treatments were taken. For instance, in the case of two stages of treatment, we may observe binary treatments at stages 1 and 2; a patient covariate that is measured at each of the two stages; the observed outcome as a function of the two received binary treatments; and the optimal outcome as a function of two binary treatments that are “optimal” in the sense that they maximize the outcome function (although these optimal treatments may or may not be the same as those observed) [50]. Under this setup, we can model the observed outcome as the optimal outcome minus two penalty functions known as *regrets*, which represent any potential loss in outcome due to receiving non-optimal treatment assignment [81]. The observed outcome can also be framed in terms of *blips*, which represent the expected difference in outcome under the observed treatment compared with under a “control” or “reference” treatment, assuming optimal treatment thereafter. To maximize the expected outcome, it follows that the optimal treatments must maximize their corresponding blip, which in turn defines the treatment rules as a function of patient covariates included in the blip functions [81]. Several well-known DTR methods to date include G-estimation [60], inverse probability of treatment weights (IPTW) [59], and dWOLS [80].

DTRs can be constructed using data from longitudinal observational studies or a special class of randomized trials known as sequential multiple assignment randomized trials (SMARTs) [13]. In a usual randomized trial, it is of interest to evaluate or confirm the efficacy of a newly developed treatment, rather than develop a new treatment in and of itself; thus, the purpose of a SMART is to develop an optimal DTR by initially randomizing patients to possible treatments, collecting information, and then re-randomizing some or all of the patients to another treatment at subsequent stages [13]. However, DTRs can also be constructed using increasingly-available observational data under several critical

assumptions and with methodologies that account for complexities such as time-varying confounding variables [13]. An example of a study aimed at estimating optimal DTRs from observational data is shown in the estimation of optimal dynamic anticoagulation regimes by Rosthøj et al. [63], where observational data from hospital records were analyzed to estimate optimal warfarin dosing strategies to maximize the percentage of study time in which subjects’ clotting tendencies of blood was within an acceptable range. Other studies have analyzed data from The Promotion of Breastfeeding Intervention Trial (PROBIT) to determine how tailoring variables such as an infant’s weight can be used to determine whether a mother should continue to breastfeed in order to optimize their infant’s growth [36, 48, 80].

1.2 Causal Inference in Observational Studies

Causal inference plays an important role in applications of precision medicine, such as in clinical genetics where it is of interest to establish a causal link between diseases and genetic variants so that the variant may be useful as a diagnostic test or therapeutic target [3]. While much discussion - and enthusiasm - has centered around machine learning-powered approaches to precision medicine, these approaches have several important limitations. First, they are limited to making predictions, whereas causal methods estimate not only just what is most likely to happen (i.e. prediction), but also the effects under different conditions, known as *counterfactual predictions* [83]. Furthermore, causal inference requires several assumptions that can be satisfied through study design or external contextual knowledge from human input, but that cannot be discovered solely from observational data [83]. Indeed, in the counterfactual causal inference framework, health researchers can formulate causal questions, encode them in algorithms, and identify average causal effects, even if the datasets at hand are observational rather than experimental (provided several identifiability conditions hold) [57]. The causal inference framework with potential outcomes therefore offers an arguably stronger appeal than prediction-based approaches for large datasets by not only defining counterfactual causal effects, but also by specifying the relationships between variables representing treatment, outcome, confounders, mediators, and colliders [57]. Methods for establishing causality in readily-accessible electronic health records, health databases, and other datasets can aid health researchers in their attempts to tailor treatments by better understanding the relationship between variables of interest without the limitations of prediction-only based approaches.

In classical randomized experiments, the treatment assignment mechanism is known in advance and is individualistic, probabilistic, and unconfounded - that is, it is independent of other units’ covariates and potential outcomes; each individual has a positive probability

of receiving either the control or treatment; and the covariates are balanced between groups [29]. However, many research questions cannot be feasibly conducted through randomized experiments due to various time, cost, or ethical constraints. Instead, researchers will often mimic randomization in observational studies through means such as matching, which broadly refers to any method that aims to balance the covariate distribution between the different treatment (or treatment and control) groups so that comparisons can be made when estimating treatment effects [74]. Indeed, there is a growing area of research that has focused on exploring matching methods and other appropriate methodology for quantifying causal effects within the constraints of observational data.

An important concept in causal inference is the propensity score, defined as the conditional probability of being assigned to the treatment group ($A = 1$) given an individual's covariates X ; it can be used to match treated and control units in observational datasets so that direct comparisons are meaningful [62]. More specifically, the propensity score $e(X) = P(A = 1|X)$ is defined as the coarsest possible balancing score $b(X)$ such that the conditional distribution of X given $b(X)$ is the same for treated and control units where $A = 1$ and $A = 0$, respectively, i.e.

$$X \perp A|b(X).$$

In a randomized trial, $e(X)$ is a known function with one accepted specification, while in observational datasets, it must be estimated through methods such as logistic regression [62]. Many matching methods quantify similarity between treated and control units using distance metrics on the estimated propensity scores, which in turn replicates a randomized experiment with respect to the observed covariates [74].

There are several key assumptions that are often made in the context of causal inference methodology for observational data. The assumption of strongly ignorable treatment assignment states that treatment assignment A is independent of an individual's potential outcomes $Y(0), Y(1)$ (i.e. the outcomes that an individual would have under $A = 0$ and $A = 1$, respectively) given all observed covariates X , and the positivity assumption states that there is a positive probability of receiving each treatment for all covariates [62]. That is,

$$(Y(0), Y(1)) \perp A|X, \quad 0 < Pr(A = 1|V) < 1$$

for all possible vectors V of covariates X [62]. If $0 < Pr(A = 1|V) < 1$ for all vectors of covariates V , then treatment assignment is strongly ignorable even if treatment assignment is a deterministic function of V and other (possibly unobserved) covariates U provided that




Figure 1.2: Illustration of a DAG for two individuals in a simple interference network when the SUTVA is violated. Here, X_{ego} , A_{ego} , and Y_{ego} represent the covariate, treatment, and outcome of the ego, respectively, while X_{alter} , A_{alter} , and Y_{alter} represent the covariate, treatment, and outcome of the alter, respectively.

the additional covariates are unrelated to the response at all levels of X [61]. It has also been established that if treatment assignment is strongly ignorable given X , then it is strongly ignorable given the propensity score $e(X)$ even when this must be estimated from the data at hand [62].

1.3 Interference and Dynamic Treatment Regimes

Another key assumption in the causal inference literature is the SUTVA, which states that an individual’s outcome is independent of the treatments received by other individuals [64]. This assumption consists of two subassumptions: an individual’s potential outcomes $Y(0)$ and $Y(1)$ are not affected by the actions received by any other individuals [17], and there are no hidden versions of treatment, meaning that no matter how an individual received a given treatment, their corresponding outcomes would be the same [65]. This assumption ensures that the causal effect for each individual is stable; that is, it is hindered neither by how the individual came to be exposed nor by the exposures of others [67]. Violations of the SUTVA results in instability of the causal effects, meaning that there is no longer a unique potential outcome for each individual under each level of treatment, but rather, there are instead “multiple versions of treatment” [67]. Figure 1.2 illustrates a Directed acyclic graph (DAG) for the simplest case where a person (referred to as the *ego* [32]) has their outcome affected by both their own treatment and covariate as well as those of another individual (referred to as the *alter* [32]).

Nonetheless, there are many settings where the SUTVA is likely to be violated. A classic

example has been outlined by Halloran and Struchiner [24], who note an epidemiology of infectious diseases needs to account for dependence in “happenings” or “occurrences” of the event of interest. They differentiate between interventions in *individuals*, such as vaccination, chemotherapy, or behavioural education, and intervention programs in the *population*, which may be in the environment and not received directly at the individual-level (such as mosquito control programs to reduce malaria transmission) [24]. In the context of dependent happenings, they define the following types of effects of treatment:

- The direct effect of an intervention is the difference between the outcome for an individual with the intervention and what the outcome would have been without the intervention, with all other factors held constant.
- The indirect effect is the effect of an intervention program had the individual not personally received intervention themselves.
- The total effect is the combined effect of an intervention program and the effect of an intervention that an individual personally receives.
- The overall effect is the difference between the outcome of an average individual in a population with an intervention program and the outcome of an individual in a comparable population without an intervention program.

To illustrate the concept of indirect effects when the SUTVA fails to hold, an individual may be less likely to become infected if their close contacts have been vaccinated and, as such, one must now consider the indirect effects of other people receiving treatment on the given individual’s outcome, rather than just the direct effect of the treatment taken by the individual themselves [25]. By considering neighbouring individuals’ treatments on a unit’s outcome, it follows that the magnitude of these indirect effects can vary depending on the fraction of the neighbouring population receiving treatment [25]. Evaluating and comparing direct, indirect, total, and overall effects is of key interest across many cluster-level trials for infectious disease research, such as the SEARCH study evaluating community-based strategies for HIV prevention and treatment and the SHINE study testing strategies to reduce *Staphylococcus aureus* infection within households [5]. Studies such as these demonstrate the need to consider different types of effects that are relevant in interference networks, as well as the extent to which they can occur, in order to optimize treatment allocation strategies.

1.3.1 Interference Structures

Much of the literature to date has provided examples where interference is likely to occur through social interaction or spatial proximity in networks that can be partitioned into different groups or clusters. An example of interference occurring through social interaction was seen in the “Moving to Opportunity” (MTO) demonstration sponsored by the U.S. Department of Housing and Urban Development, where eligible residents across five different U.S. cities may or may not have received different forms of assistance to relocate [69]. A family was thought to be more likely to move if all families in an interaction set were assigned to, and moved to, the suburbs, but they may also have been less likely to relocate in a tight rental market if many other families were also assigned to relocate at the same time [69]. In this example, Sobel [69] notes that social interaction is a common form of interference with human subjects, and in the case of MTO, each family can be considered as an individual unit, with multiple families (perhaps within the same neighbourhood) comprising a group in which interference can occur. Failing to account for interference can result in misleading conclusions, as the difference between treatment and control group means is no longer the estimate of a causal parameter, but rather, the difference between the average effect of treatment over a collection of policies and the spillover effect on the untreated [69]. Considerations such as these motivate the need for improvised methodology when exploring different treatment effects for causal inference in situations where the SUTVA is not likely to be met.

When considering interference structures in social or spatial networks, it is important to determine whether treatments are applied at an individual or group level. Much of the literature on causal inference with interference to date has been formalized in a setting of one level of observational unit, such as in the case of an individual’s vaccination status affecting the infection risk of their family members or in social networks among friends where interventions can both directly affect an individual and indirectly affect their peers as well [87]. Bipartite causal inference with interference considers interventions (or treatments) that are applied at a group-level, with the outcomes of interest being at the individual-level within those groups, which can depend on the treatments assigned to (possibly many) other units [87]. Such settings likely arise from underlying scientific phenomena that cannot be described by unit-to-unit outcome dependencies common in infectious disease settings or social networks; rather, it is the consequence of complex exposure dependencies that describe how a particular treatment effect propagates across units [87].

One example of bipartite causal inference with interference was outlined by Hong and Raudenbush [26], which considered the policy of retaining low-achieving children in kindergarten when a child’s potential outcomes depend both on which school they attend in addition to the treatment assignments of other children. They used multilevel propensity

score stratification to approximate a two-stage experiment, where schools were randomly assigned to a policy of retaining comparatively more or fewer children in kindergarten after being blocked on covariates in the first stage and “at-risk” students within schools were assigned at random to be retained (after being blocked on covariates) in the second stage [26]. However, Zigler and Papadogeorgou [87] highlight that clustering units may be distinguishable from interventional units, as a clustering unit may consist of more than one interventional unit. For example, when investigating the effect of a teacher intervention on student outcomes, both the teachers (interventional units) and students (outcome units) may be clustered within schools and therefore comprise the overall clustering units. [87]. Thus, it is important to consider the clustered structure of the data and the level at which an intervention is applied in order to clearly identify and quantify the different effects of interference taking place.

Although a sequential two-stage randomization design is one framework for analyzing interference in clustered datasets, there are cases where randomization may only occur at the individual or group level [55]. Tchetgen Tchetgen and VanderWeele [76] proposed estimators for treatment effects in the presence of interference without requiring randomization of individuals or groups, which use weighted averages of the observed responses based on the inverse of the group-level propensity score; that is, the joint probability a group of individuals receives a particular vector of treatments [62]. The group-level propensity score is based on group-level generalizations of the usual individual-level assumptions of no unmeasured confounders and positivity made in analysis of observational studies where the SUTVA is assumed to be met [55]. These inverse-probability weighted (IPW) estimators were shown to be unbiased when the group-level propensity score was known [76, 55], and Perez-Heydrich et al. [55] explored an extension of this approach to assess the finite-sample properties of these estimators when the group-level propensity score must be estimated. If each individual in a group is thought to be receiving treatment independently of one another, then the group-level propensity score may be estimated by taking the product of each person’s estimated probability of treatment [55]. Indeed, the joint propensity score for a group of individuals can be used in IPW estimators for any general form of interference where no assumption is made on the structure or form of interference present, provided the interference sets for each individual are known a priori [41].

When analyzing interference networks, it is often helpful to visualize the connections or “links” between individuals as edges of a pre-specified graph [1, 4, 20, 19, 34]. Aronow and Samii [1] considered an estimation framework for interference characterized by three components: the probability distribution of treatments assigned; an exposure mapping that relates treatments assigned to exposures received; and a set of causal estimands selected based on the primary questions of interest to the researchers. The work by Forastiere,

Mealli, and Airolidi [20] introduced a group propensity score that balanced individual and neighbourhood covariates across units under different levels of individual treatment assignment and exposure to treatments from an individual’s neighbours, with adjustments performed using a penalized spline regression. Similarly, Giffin et al. [23] proposed a method to specifically address general forms of spatial interference in which exposures at nearby locations are considered more influential than exposures at locations further apart. More specifically, their approach assumes that a spillover effect from nearby neighbours is proportional to a kernel of a distance, and is therefore able to simultaneously estimate the spatial range and weighting of those individuals [23].

Considering treatments of multiple neighbours on a given individual can become computationally burdensome in large networks if no limit is placed on the connections between an individual and possible neighbours whose treatments may affect that individual’s outcome. Indeed, under general exposure interference in a network with n units, there are 2^n unit-level potential outcomes to consider, and the problem quickly becomes intractable [23]. To address this, researchers often adopt the partial interference assumption, which assumes that interference is restricted to occur only within some groups of individuals, but not between groups [69]. This assumption is also known as clustered interference, where clusters may refer to households, classrooms, or other grouping structures [8], and it is often reasonable when the groups can be considered sufficiently separate in space or time [28]. However, there may be scenarios when this assumption does not hold, whereby an individual j ’s neighbours consist of individuals both within their own group i and from another group i' for a variety of reasons. An immediate violation of this assumption may occur when several groups are in close spatial proximity of each other; for instance, individuals within a household in a given neighbourhood may be more likely to interact with adjacent households than with households from other neighbourhoods in far-away communities. As such, it may be of interest to consider between-group interference in a more restricted setting where clustered data are hierarchical, as a cluster comprised of several smaller subgroups accounts for both subgroup-specific differences in addition to some overall degree of similarity at the overall-group level. One may then be able to account for possible interference between subgroups within a given hierarchy while assuming that interference does not occur between other hierarchies. Indeed, Li et al. [38] describe scenarios in which one can consider hierarchical clustering within a hospital’s administration, geographical area, or health plan, while Vaughan et al. [78] note that an analysis of student achievement may include student-level variables (such as socio-economic status and prior performance) in addition to cluster-level, school-specific variables (such as faculty size and annual funding). Thus, considering hierarchical interference networks can provide flexibility beyond the assumption of partial interference, as hierarchical networks provide a useful framework where interference may occur both within and between some, but not

all, other groups.

1.3.2 Interference Networks for Clustered Data with Mixed Effects Models

As outlined in the previous section, much of the literature on interference to date has considered general, spatial, or partial interference structures. However, studies on within- and between-group interference in hierarchical networks have been limited. Zigler, Dominic, and Wang [86] discussed a limitation of the partial interference assumption in the analysis of causal effects of the 1990 Clean Air Act Amendments (CAAA) among Medicare enrollees living within 6 miles of a pollution monitor, as assuming no interference between locations in different counties may not account for observations near county borders and thus be too restrictive [86]. Tchetgen Tchetgen and Vanderweele [76] acknowledged that association between the exposure levels of individuals belonging to different groups may occur if those groups are in close proximity and suggested that future developments may consider a framework to potentially allow for both within-cluster and between-cluster forms of interference. When estimating the effectiveness of power plant emission reduction technologies on ambient ozone pollution, Papadogeorgou, Mealli, and Zigler [53] again emphasize that the partial interference assumption may be a simplification, since pollution from one power plant can spread across longer distances. Examples such as these highlight the possibility of between-group interference occurring in various settings as a result of spatial proximity and environmental factors, and they motivate ongoing work toward developing causal inference methods for interference structures where the partial interference assumption may not hold.

One possible solution for accounting for both within- and between-group interference is to consider hierarchical-structured networks, as they provide a framework for assessing both individual- and group-level characteristics that may determine whether interference occurs between individuals. Perez et al. [55] proposed hierarchical models as a means of analyzing different levels of interference, such as interference taking place both between individuals within a given household as well as between individuals from other houses but restricted to the same neighbourhood. Indeed, Li, Zaslavsky, and Landrum [38] note that both measured and unmeasured factors at the cluster-level can create subgroup-specific variation in quality of treatment and, subsequently, affect the outcomes of interest. Mixed effects models have previously been discussed in the context of interference networks as a way of accounting for both types of characteristics in the group propensity score [76, 55, 41]. Having nested random effects at the lower and higher grouping levels can be used to model unobserved heterogeneity both between and within groups [56], and indeed,

if the relationship between explanatory variables and the response of interest differs by group, then this can be accounted for through random coefficients for each grouping level [58]. By considering nested hierarchical structures, one can consider factors at different grouping levels that may affect treatment assignment, thereby accounting for both within- and between-group interference as represented through group-level covariates and random effects within groups.

To date, causal inference methods for interference have yet to examine propensity scores with mixed effects in the context of addressing possible between-group interference. Hong and Raudenbush [26] have previously explored relaxing SUTVA through use of a two-stage propensity score model when estimating effects of retaining low-achieving students in kindergarten rather than promoting them to the first grade, but while they included student-level and school-level covariates in their propensity score model, they assumed no interference between schools and did not incorporate random effects [26]. Li, Zaslavsky, and Landrum [38] have proposed propensity-score-weighted estimators for multilevel data using random effects models, noting that while accounting for clustering in propensity score analyses can also account for violations of the SUTVA at a given cluster level, more than one level of clustering may also be relevant. These previous works motivate interest in considering between-group interference under hierarchical settings, and in particular, how this type of interference can be implemented in DTR methods.

1.3.3 Dynamic Treatment Regimes under Interference Networks

DTRs have only recently been considered in the context of interference. Su et al. [75] derived Q-learning and A-learning methods that account for interaction between outcomes and treatments in a network; by considering a network mean as the outcome of interest, and by not including interactions between others’ treatment assignments and one’s own treatment assignment, the optimal treatment rules depend only on a given individual’s covariates and their degree in the network (number of neighbours) but not on their neighbours’ treatments. Their proposed model and methods were applied to a mobile game social network where one’s behavior might affect one’s friends, and a given individual may be more likely to consider the game that their friends invite them to play [75]. Jiang et al. [32] proposed the network propensity function and the concept of network weights to establish an implementation of dWOLS that remains doubly robust under interference associated with network links. Their methods were applied to data from the Population Assessment of Tobacco and Health (PATH) study to investigate cigarette dependence between couples in household networks, with an “ego” randomly selected from each couple for whom a DTR was to be estimated and the other individual designated as an “alter”

whose treatments could affect the ego’s outcome. The cigarette-dependence outcome was a continuous score that was obtained from combining responses to a questionnaire, with larger values of the score indicating greater improvement between the first and last stages of follow-up, and the optimal treatment rule included tailoring variables such as age, an individual’s plan to quit, and their household member’s use of e-cigarettes (treatment status) [32].

In their PATH analysis, Jiang et al. [32] assumed that interference was occurring between members of the same household, but not between members of different households. By considering each household as its own group, the authors proposed an application of DTRs to interference networks where the partial interference assumption held, that is, interference only occurred within groups (households) but not between groups (Figure 1.3). However, an interesting extension of this work would be to consider whether between-household interference could also take place between households that were within close proximity, such as between households that were nested within the same neighbourhood or community. Thus, a hierarchical network of houses nested within neighbourhoods, perhaps in turn nested within some other region, would present an opportunity to assess DTRs under interference without being restricted by the partial interference assumption. Realistically, it may be difficult to determine which neighbours from other households that a given individual may be interacting with in addition to their own household members without explicit relationships being recorded in a dataset. Nonetheless, if some form of between-group interference is occurring in a hierarchical network through the interactions of individuals within and between groups, then assuming that interference only occurs within the same households is a mis-specification of the form of interference occurring, as it underestimates the exposure to other individuals on a given person’s outcome and will possibly result in the estimated DTR rules being inaccurate. For instance, suppose the optimal treatment rule specifies that: if a 25 year old has an alter whose treatment is pre-determined to be strategy A, then this individual should adopt strategy A to quit smoking and otherwise, they should also be assigned to strategy B. If the individual has an alter who is from a different household who has taken strategy A, but no alters within their own household, then absence of a “within-group” member with treatment strategy A may result in the individual being incorrectly assigned strategy B if their “between-group” alters are not accounted for in their treatment rule. Examples such as these elucidate real-world applicability to implementations of DTRs in hierarchical interference networks.




Figure 1.3: Illustration of a partial interference structure for smoking cessation strategies within households.

1.4 Service System Utilization for Opioid Overdose Patients within and between Facilities

Opioid use disorder (OUD) is a chronic, relapsing disorder that is detrimental to an individual’s physical, social, and psychological well-being [66]. As the rate of opioid prescribing and use continues to rise, it has become a public health crisis in North America, with its effects seen in the form of increased mortality and healthcare utilization [66]. The multifaceted impact of OUD makes it difficult to define an outcome measure for clinical trialists, healthcare services, and providers when developing treatment programs, and patient goals for medication-assisted treatment (MAT) tend to be overlooked in clinical trials and guidelines [66].

Emergency departments (EDs) are a primary access point to the health care system for interventions to reduce harm from OUD, yet studies recommend that ED visits be coupled with community-based follow-up and support to ensure longer-term retention[33]. Furthermore, the rapid increase in ED visits related to substance use has created numerous challenges for emergency clinicians and the healthcare system, such as increased length of stay, higher service delivery costs, and higher rates of hospital admissions [9]. As such, it is of interest to consider how personalized treatment strategies can be used to identify which persons with OUD would benefit the most from ED and post-ED visit services so that they can be strategically allocated.

Post-ED visits among persons with OUD (pWOUD) provide an interesting setting under which the SUTVA is likely to be violated. While ED visits are an important opportunity for healthcare providers to connect pWOUD to treatment, many factors can influence

patients’ engagement in OUD treatment following an ED visit, such as the type of treatment received in the ED (behavioural counseling, initiation of buprenorphine treatment, and so forth); stigma and misconceptions about treatment; and referrals to treatment at discharge, among others [14]. If living arrangements of pWOUD provide opportunities for close contact with other pWOUD, such as through shared housing or receiving treatment at facilities within the same community, then an individual’s engagement with treatment and need for ongoing care may be influenced by the extent to which peers within their social or spatial networks have also successfully engaged with treatment at discharge. Indeed, previous studies have shown that negative social ties may reinforce addiction, while interventions around positive social network support can help augment treatment for OUD [37]. Management of post-ED treatment among pWOUD to improve patient outcomes and reduce strain on scarce hospital resources is a challenging problem that motivates solutions within the DTR framework that account for possible violations of the SUTVA that may occur within social and spatial networks.

While there are many types of treatments and outcomes for pWOUD that can be considered in a DTR framework that accounts for interference, we present a motivating example for this thesis in the context of patient dispositions following an ED visit, with data simulated from demographic summaries of a retrospective cohort study by Moe et al. [46] using the Provincial Overdose Cohort, which contains data for patients who experienced an opioid-related overdose in British Columbia (BC). This project was inspired from collaboration on a project with a Canadian provincial health authority, which is currently undergoing a delayed ethics clearance before the resulting manuscript can be submitted for publication. While the demographic summaries by Moe et al. [46] do not identify relationships among individuals or provide living arrangement information that could directly support violation of the SUTVA, to demonstrate the applicability of our methods throughout this thesis, we simulate cluster-level variables that represent different facilities nested within broader provincial regions and assume that individuals within the same region can interact with each other through the same or different facility. In this way, we provide a framework where interference can occur within and between broadly-defined groups in a hierarchical structure. In our analyses, we consider the “treatment” received to be a binary representation of the patient’s disposition following an ED visit, set to be 1 if the patient was admitted to hospital or transferred and 0 if they were discharged home or to another residence. For simplicity in our analyses, we only consider patients who have not transitioned between different regions between their stages of treatment, thereby restricting the extent of between-group interference that may take place.

Since we assume that individuals within the same region can interact with each other, but we do not have a designated ego-alter mapping to identify these potential interactions,

we choose to randomly sample egos from each region so that each ego is guaranteed to have at least one alter who also “belongs” to the same facility or possibly a different facility within the same region. We again emphasize that this is for proof of concept due to limitations with the available data, and we consider the time from the j -th ED visit until just before the $j + 1$ visit as a “stage” of treatment in the DTR framework while assuming that an alter’s disposition was pre-determined at the ego’s stage of analysis. Our outcome of interest is the total number of physician visits made since the last recorded ED visit, and we use the patient’s age at each stage as a tailoring variate to determine whether the ego may benefit from being admitted to a hospital or being discharged (Figure 1.4). Available demographics summarized in the retrospective cohort study by Moe et al. [46] include age, sex, presence of a psychiatric comorbidity and CTAS score.

In this thesis, we address the following objectives:

- Assuming no interference, implement dWOLS to estimate the optimal DTR for the simulated opioid overdose cohort, in terms of their age, to minimize the total number of physician visits made since the last recorded ED visit.
- Assuming interference in which the dispositions undertaken by alters are fixed, implement dWOLS to estimate the optimal DTR for the simulated opioid overdose cohort, in terms of their age and the alters’ dispositions, to minimize the total number of physician visits made since the last recorded ED visit.

Specifically, we illustrate how a patient’s age can be used to determine whether they should be admitted to hospital or transferred as the “treatment” of interest ($A = 1$) versus being discharged to home or another residence ($A = 0$), depending on the fixed “treatment” received by their designated alter. We emphasize that although we have arbitrarily selected the “egos” for which we estimate optimal DTRs and their “alters” whose disposition can affect the egos’ outcomes for proof of concept, this set-up nonetheless has several important implications for future work in this area. First, previous research has already supported the idea that social and spatial networks are relevant to the problem of the opioid overdose crisis, as households and social networks are thought to play a role in how pWOUD live their life. An analysis of self-reported data from the Alberta Health and Drug Use Survey revealed a significant association between housing instability and hospital use, with unstably housed individuals being nearly twice as likely to require hospital care [45]. Additionally, it has been suggested that there is a need to tailor interventions based on dynamics within housing environments [18]. Given that much emphasis is placed on the scarcity of hospital resources, assessing the disposition of an individual within one’s immediate social or spatial network may influence a health expert’s opinion as to whether




Figure 1.4: Illustration of within- and between-group interference, where different facilities are nested within an overall region.

the patient of interest may also benefit from the same or a different disposition. These findings can, in turn, provide insights into optimizing hospital resource allocation for other pWOD.

1.5 Overview of Thesis Organization

It is of interest to consider the dWOLS method with network propensity weights in settings where network data takes on a hierarchical structure, as it may provide a means of estimating the optimal DTR in the presence of both within- and between-group interference. Previous work has shown that ignoring the multilevel structure in propensity score weighting analysis can result in biased estimates, as more than one level of clustering may be relevant when one attempts to account for clustering in propensity score analyses [38]. Given that previous studies have discussed using mixed effects models as a way of accounting for both individual- and group-level characteristics in propensity score applications to interference networks [76, 55, 41], it is intuitive to consider incorporating this methodology into a modified version of dWOLS. Furthermore, it provides an opportunity for assessing the modified DTR methodology under interference proposed by Jiang et al. [32], specifically for estimating the optimal DTR in a hierarchical setting with interference. This will offer a promising approach to DTR estimation across a variety of clustered networks in which there is flexibility to consider between-group interference in addition to within-group interference.

In Chapter 2, we will provide the theoretical framework behind DTRs such as dWOLS, the network propensity function in the context of DTRs with interference, and generalized linear mixed-effects models. In Chapter 3, we illustrate how dWOLS can be used to estimate the optimal DTR for the simulated opioid overdose cohort with an age-tailored

approach to minimize the number of physician visits between ED visits. Chapter 4 presents simulation studies that assess the performance of dWOLS with mis-specified interference forms and different weighting methods in a hierarchical network where both within- and between-group interference is taking place. We conclude this chapter with an application of these methods to the simulated opioid overdose cohort to illustrate how within- and between-group interference can be incorporated in estimation of the optimal treatment rule when recommending the disposition of a patient following an ED visit. In Chapter 5, we introduce several sets of simulation studies that determine the extent to which interference may be taking place in an observational dataset to warrant use of modified dWOLS methodology that accounts for interference versus a “standard” implementation of dWOLS that does not account for any form of interference. We conclude with a discussion of our findings and future work in Chapter 6.

Chapter 2

Notation and Methodology for DTRs with Network Interference

In this chapter, we first provide an overview of commonly-used notation in causal inference with interference. In section 2.2, we then introduce the theoretical framework for DTRs such as Q-learning and G-estimation as outlined in Chakraborty and Moodie [13], which provide the foundation for the methodological development of dynamic weighted ordinary least squares regression (dWOLS) [80]. Finally, we refer to the work by Jiang et al. [32] to establish an implementation of dWOLS under network interference.

2.1 Causal Inference with Interference: Notation and Assumptions

Throughout this section, we will use notation as outlined in [28], [80], and [75]. We consider $i = 1, \dots, I$ groups of individuals, with n_i denoting the number of individuals in group i . The vector $\mathbf{A}_i = (A_{i1}, \dots, A_{in_i})$ represents the binary treatments of each individual, where A_{ij} denotes the treatment of individual j in group i and can take on a value of 0 or 1. Thus, for group i , there are 2^{n_i} possible configurations for the treatment vector $\mathbf{A}_i$, which may be considered as possible treatment programs. Let $\mathbf{a}_i$ and a_{ij} represent the realizations of $\mathbf{A}_i$ and A_{ij} , respectively. We can also consider the vector $\mathbf{A}_{i(j)} = (A_{i1}, \dots, A_{ij-1}, A_{ij+1}, \dots, A_{in_i})$, which is the $\mathbf{A}_i$ treatment vector with the j -th entry omitted. Finally, let $Y_{ij}(\mathbf{a}_i)$ represent the potential outcome of individual j in group i under treatment program $\mathbf{a}_i$. Under the partial inference assumption, an individual j in group i may be affected by the treatments of other individuals in group i , but not those in some other group i' .

Interference is often discussed in the context of networks, where a network of individuals $\mathcal{G}$ consists of a pair of nodes $\mathcal{N}$ with edges $(j, k) = (k, j) \in \mathcal{E}$ representing links between individuals [19]. For a given individual j , the triplet $(j, \mathcal{N}_j, \mathcal{N}_{-j})$ represents a partition of $\mathcal{N}$ where $\mathcal{N}_j$ consists of all nodes (or neighbours) of individual j connected with an edge in $\mathcal{E}$ and $\mathcal{N}_{-j}$ consists of all nodes or individuals not in $\mathcal{N}_j$ [19]. Such a partition helps to identify the neighbours of a given individual in a network, and the mechanism used to generate this partition should be determined at the discretion of the researcher in the context of the data at hand. For instance, two nodes (or individuals) may be considered “linked” in a network as a result of spatial proximity, social contact, or some other metric. Both $\mathcal{N}_j$ and $\mathcal{N}_{-j}$ groups of individuals have respective treatment vectors $\mathbf{A}_j, \mathbf{A}_{-j}$ and respective outcome vectors $\mathbf{Y}_j, \mathbf{Y}_{-j}$, while the covariates of individual j , $\mathbf{X}_j$, can be partitioned into individual-specific covariates $\mathbf{X}_j^{\text{ind}}$ and neighbourhood covariates $\mathbf{X}_j^{\text{neigh}}$ [19]. Here, $\mathbf{X}_j^{\text{neigh}}$ may consist of variables representing the structure of the neighborhood (such as the degree or topology), network properties at the individual level representing the position of the individual’s neighborhood in the graph (e.g. centrality, betweenness, the number of shared neighbors, etc.), and aggregational covariates that summarize individual-level covariates in the neighborhood (e.g. proportion of females) [19]. We further define $\mathcal{X}_j$ as the covariate set of all neighbours of individual j , i.e. $\mathcal{X}_j = \{\mathbf{X}_k : k \in \mathcal{N}_j\}$ [75].

Forastiere, Airolidi, and Mealli [19] introduce the stable unit treatment on neighbourhood value assumption (SUTNVA) to expand on the potential outcomes framework in the presence of neighbourhood interference, which consists of the following assumptions:

- **Assumption 1:** No multiple versions of treatment (consistency)

$$\text{If } \mathbf{A} = \mathbf{a} \text{ then } Y_j = Y_j(\mathbf{a})$$

- **Assumption 2:** Neighbourhood interference

Given a function $t_j : \{0, 1\}^{N_j} \rightarrow \mathcal{T}_j, \forall j \in \mathcal{N}, \forall \mathbf{A}_{-j}, \mathbf{A}'_{-j}$, and $\forall \mathbf{A}_{\mathcal{N}_j}, \mathbf{A}'_{\mathcal{N}_j}$:
 $t_j(\mathbf{A}_{\mathcal{N}_j}) = t_j(\mathbf{A}'_{\mathcal{N}_j})$, the following equality holds:

$$Y_j(A_j, \mathbf{A}_{\mathcal{N}_j}, \mathbf{A}_{\mathcal{N}_{-j}}) = Y_j(A_j, \mathbf{A}'_{\mathcal{N}_j}, \mathbf{A}'_{\mathcal{N}_{-j}}).$$

- **Assumption 3:** Unconfoundedness of individual and neighbourhood treatment

$$Y_j(a, \mathbf{a}_{\mathcal{N}_j}) \perp A_j, \mathbf{A}_{\mathcal{N}_j} | \mathbf{X}_j, \mathcal{X}_j, \forall a \in \{0, 1\}, \mathbf{a}_{\mathcal{N}_j} \in \mathcal{T}_j, \forall j$$

- **Assumption 4:** Positivity

$P(A_j = a_j, \mathbf{A}_{\mathcal{N}_j} = \mathbf{a}_{\mathcal{N}_j} | \mathbf{X}_j, \mathbf{X}_{\mathcal{N}_j}) > 0$ for all $(A_j, \mathbf{A}_{\mathcal{N}_j})$ such that $P(\mathbf{X}_j, \mathbf{X}_{\mathcal{N}_j}) > 0$ [32]

The first assumption states that different mechanisms used to assign treatment do not constitute different treatments; that is, if $\mathbf{A} = \mathbf{a}$, then $Y_j = Y_j(\mathbf{a})$ [19]. In other words, the observed outcome is the same as the potential outcome under that received treatment. The second assumption allows the outcome Y_j of individual j to depend on the treatments received by their neighbours through a function t_j , but not on other individuals [19]. In this case, the order of individuals being treated in $\mathbf{A}_j, \mathbf{A}'_j$ does not matter provided that $t_j(\mathbf{A}_j) = t_j(\mathbf{A}'_j)$, such as when t_j represents the total number of neighbours of individual j who are treated. For example, if $\mathbf{A}_j = (0, 1, 1)$, $\mathbf{A}'_j = (1, 0, 1)$, and $t_j(\cdot) = \sum_{k \in \mathcal{N}_j} A_k$, then $t_j(\mathbf{A}_j) = t_j(\mathbf{A}'_j) = 2$. The function $t_j(\cdot)$ is known as an “exposure mapping” that provides a low-dimensional summary of neighbours’ treatment assignments and is assumed to fully characterize the interference taking place [1, 32]. Although the exposure mapping is application-specific and must be determined by the researcher, we will see that in the DTR framework, $t_j(\cdot)$ cannot take only one value for every configuration of neighbours’ treatments, that is, it cannot be true that $t_j(a_{\mathcal{N}_j}) = c$ for every $a_{\mathcal{N}_j}$ where c is a fixed value, to ensure identifiability of the interference term in the model [32].

2.2 Statistical Reinforcement Learning for Dynamic Treatment Regimes

Many DTRs are constructed using concepts from reinforcement learning (RL), whereby a *learning agent* interacts with its *environment* at one or more *stages* by choosing an action from a possible set of actions based on the observed *state* of the environment. In RL, a *policy* is defined as the actions that the agent takes based on the current state at any given stage, and this can be either *deterministic* (represented as a vector of mappings, with each vector mapping from the state space to the action space for a given stage) or *stochastic* (a vector of mappings from the state space at a given stage to the space of probability distributions over the stage-specific action space). The *value function* (or value) at a given state with respect to a policy measures the total expected future reward of the agent, starting with that state, and following the given policy to choose subsequent actions. The RL framework can be compared to constructing DTRs in a medical setting through the following analogies:

- the clinician’s decision support system takes the role of the agent;

- the population of subjects of interest plays the role of the environment;
- the patient’s pre-treatment observations (medical history, age, smoking status, etc.) constitute the state;
- the treatment (drug received, medication dosage, intervention type, etc.) serves as the action; and
- the measure of the patient’s outcome (improvement in symptoms, time spent without substance abuse, longevity, etc.) is the value of interest.

It should be noted that, despite these analogies, there are some differences between the RL and DTR frameworks, one noteworthy one being the non-Markovian set-up of DTRs. Many applications of RL algorithms assume that the state transition probabilities depend only on the immediate past as opposed to the entire history, whereas this assumption is rarely met in a medical setting where a patient’s distant history (past states and treatments) can influence future states and, subsequently, treatment.

We now outline a probabilistic framework for RL which will be used to help build the theoretical framework for DTR construction. Let $\bar{O}_T = (O_1, \dots, O_T)$ denote the vector of observed states from stages 1 through T , where for $1 \leq t \leq T$, each $O_t \in \mathcal{O}_t$ is a vector of discrete or continuous variables belonging to the state space $\mathcal{O}_t$. Let $\bar{A}_T = (A_1, \dots, A_T)$ denote the vector of observed actions (or treatments) from stages 1 through T , where for $1 \leq t \leq T$, each $A_t \in \mathcal{A}_t$ is discrete and belongs to the action space $\mathcal{A}_t$. Building on these, we can define the history at stage t as $H_t = (\bar{O}_t, \bar{A}_{t-1})$, noting that H_t, O_t , and A_t are random variables with their realizations denoted h_t, o_t , and a_t , respectively. The outcome at stage t can be expressed in terms of H_t, A_t , and O_{t+1} , i.e.

$$Y_t = Y_t(H_t, A_t).$$

When constructing DTRs, one must use data from a sample of n subjects; specifically, n finite-horizon trajectories of the form $\{O_1, A_1, O_2, \dots, A_T, O_{T+1}\}$ (in a setting of T stages). It is assumed that these subjects are sampled at random according to a probability distribution P_π which is composed of the unknown distribution of each O_t conditional on (H_{t-1}, A_{t-1}) and a fixed exploration policy $\pi = (\pi_1, \dots, \pi_T)$ for generating the actions. In SMARTs, this exploration policy is known by design; in observed datasets, this can be estimated by the propensity score. The likelihood under P_π of the trajectory $\{o_1, a_1, o_2, \dots, a_T, o_{T+1}\}$ is

$$f_1(o_1)\pi_1(a_1|o_1)\Pi_{t=2}^T f_j(o_t|h_{t-1}, a_{t-1})\pi_t(a_t|h_t)f_{T+1}(o_{T+1}|h_T, a_T),$$

and if an arbitrary deterministic policy $d = (d_1, \dots, d_T)$ is used, the likelihood under P_d becomes

$$f_1(o_1)I[a_1 = d(o_1)]\Pi_{t=2}^T f_t(o_t|h_{t-1}, a_{t-1})I[a_t = d(h_t)]f_{T+1}(o_{T+1}|h_T, a_T),$$

where $f_1, \dots, f_T$ are the unknown conditional distributions of $O_1, \dots, O_T$, respectively, and $\pi(a_t|h_t) > 0$ for all $a_t \in \mathcal{A}_t$ and for each possible value h_t .

Letting E_d denote the expectation with respect to the distribution P_d , we want to learn the optimal policy d^* from n finite-horizon trajectories which may not be optimal, i.e. find the policy within a class of policies that has the greatest value. The value function for a state o_1 is

$$V^d(o_1) = E_d[\sum_{t=1}^T Y_t(H_t, A_t, O_{t+1})|O_1 = o_1],$$

and the stage t value function for a given history h_t is given by

$$V_t^d(h_t) = E_d[\sum_{u=t}^T Y_u(H_u, A_u, O_{u+1})|H_u = h_u].$$

It can be shown that

$$V_t^d(h_t) = E_d[Y_t(H_t, A_t, O_{t+1}) + V_{t+1}^d(H_{t+1})|H_t = h_t], \quad 1 \leq t \leq T,$$

and the optimal stage t value function for a history h_t is

$$V_t^{opt}(h_t) = \max_{d \in D} V_t^d(h_t).$$

Furthermore, the optimal value functions satisfy the Bellman equation [10],

$$V_t^{opt}(h_t) = \max_{a_t \in A_t} E_d[Y_t(H_t, A_t, O_{t+1}) + V_{t+1}^d(H_{t+1})|H_t = h_t, A_t = a_t], \quad 1 \leq t \leq T$$

when all observations and actions are discrete.

The marginal value of a policy d is the average value function under that policy taken over all possible initial observations O_1 , i.e.

$$V^d = E_{O_1}[V^d(O_1)] = E_d[\sum_{t=1}^T Y(H_t, A_t, O_{t+1})].$$

To estimate a policy that maximizes the value (i.e. to find the *optimal policy*), there are two approaches: one must either specify a policy space and estimate the value of each policy in that space to find the best one, or one can work with an *action-value function*, or *Q-function* (where “Q” stands for the “quality of action”). The stage t Q-function for policy d is the total expected future reward starting from a history h_t at stage t , taking an action a_t , and following the policy d thereafter. Thus,

$$Q_t^d(h_t, a_t) = E_d[Y_t(H_t, A_t, O_{t+1}) + V_{t+1}^d(H_{t+1}) | H_t = h_t, A_t = a_t]$$

and the optimal stage t Q-function is

$$Q_t^{opt}(h_t, a_t) = E_d[Y_t(H_t, A_t, O_{t+1}) + V_{t+1}^{opt}(H_{t+1}) | H_t = h_t, A_t = a_t].$$

A disadvantage of searching through the policy space and estimating the value of each policy is that there may be high variability in the function estimates, and resulting high variability in the estimated optimal policies. On the other hand, by modelling the entire multivariate distribution of the data and employing dynamic programming methods, the entire likelihood of the data may not be relevant for choosing optimal actions, which may result in biased estimators of the value function and subsequently the optimal policy if the model specification is incorrect. Q-learning is a method that lies between these two extremes whereby only part of the entire likelihood, i.e. the Q-functions, are modelled via regression.

2.2.1 Q-Learning with Linear Models

Since the true Q-functions are not known in practice, they must be estimated from the data, e.g. with regression models. In a two-stage setting, we define the optimal Q-functions for each stage as

$$Q_2^{opt}(H_2, A_2) = E[Y_2 | H_2, A_2]$$

and

$$Q_1^{opt}(H_1, A_1) = E[Y_1 + \max_{a_2} Q_2^{opt}(H_2, A_2) | H_1, A_1].$$

Using a backwards induction argument (as in dynamic programming), the optimal DTR (d_1^{opt}, d_2^{opt}) would be

$$d_t^{opt}(h_t) = \arg \max_{a_t} Q_t^{opt}(h_t, a_t), \quad 1 \leq t \leq 2.$$

The stage $t, t = 1, 2$ Q-function can be modelled as

$$Q_t^{opt}(H_t, A_t; \beta_t, \psi_t) = \beta_t^T H_{t0} + (\psi_t^T H_{t1}) A_t,$$

where H_{t0} is the “main effect of history” while H_{t1} is the “treatment effect of history” containing *prescriptive* or *tailoring* variables. The Q-learning algorithm involves the following steps:

- Stage 2 regression: $(\hat{\beta}_2, \hat{\psi}_2) = \arg \min_{\beta_2, \psi_2} \frac{1}{n} \sum_{i=1}^n (Y_{2i} - Q_2^{opt}(H_{2i}, A_{2i}; \beta_2, \psi_2))^2$
- Stage 1 pseudo-outcome: $\hat{Y}_{1i} = Y_{1i} + \max_{a_2} Q_2^{opt}(H_{2i}, a_2; \hat{\beta}_2, \hat{\psi}_2)$
- Stage 1 regression: $(\hat{\beta}_1, \hat{\psi}_1) = \arg \min_{\beta_1, \psi_1} \frac{1}{n} \sum_{i=1}^n (\hat{Y}_{1i} - Q_1^{opt}(H_{1i}, A_{1i}; \beta_1, \psi_1))^2$

The estimated optimal DTR, $(\hat{d}_1^{opt}, \hat{d}_2^{opt})$, is given by

$$\hat{d}_t^{opt}(h_t) = \arg \max_{a_t} Q_t^{opt}(h_t, a_t; \hat{\beta}_t, \hat{\psi}_t), \quad t = 1, 2,$$

and the above process can be generalized to $T > 2$ stages. By defining $Q_{T+1}^{opt} = 0$ and

$$Q_t^{opt}(H_t, A_t) = E[Y_t + \max_{a_{t+1}} Q_{t+1}^{opt}(H_{t+1}, A_{t+1}) | H_t, A_t], \quad t = 1, \dots, T,$$

and by parameterizing the stage-specific Q-functions as before, i.e.

$$Q_t^{opt}(H_t, A_t; \beta_t, \psi_t) = \beta_t^T H_{t0} + (\psi_t^T H_{t1}) A_t, \quad t = 1, \dots, T,$$

then the stage t regression parameters can be estimated as

$$\arg \min_{\beta_t, \psi_t} \frac{1}{n} \sum_{i=1}^n (Y_{ti} + \max_{a_{t+1}} Q_{t+1}^{opt}(H_{t+1i}, a_{t+1}; \hat{\beta}_{t+1}, \hat{\psi}_{t+1}) - Q_t^{opt}(H_{ti}, A_{ti}; \beta_t, \psi_t))^2,$$

where $Y_{ti} + \max_{a_{t+1}} Q_{t+1}^{opt}(H_{t+1i}, a_{t+1}; \hat{\beta}_{t+1}, \hat{\psi}_{t+1})$ is the stage t pseudo-outcome. The corresponding t -th stage optimal treatment rule is then given by

$$\hat{d}_t^{opt}(h_t) = \arg \max_{a_t} Q_t^{opt}(h_t, a_t; \hat{\beta}_t, \hat{\psi}_t), \quad t = 1, \dots, T.$$

2.2.2 Q-Learning Using Observational Data

Although Q-learning had previously been studied and implemented in settings where the exposure was randomized, much work has shifted to implementing Q-learning in observational datasets. As explained in Chapter 1, the propensity score can be used for adjustment in methods such as matching and inverse probability of treatment weighting so that comparisons can be made between the outcomes of control and treated groups. To achieve unbiasedness from adjustment with the propensity score, the propensity score must have all confounding variables measured and included in its model (no unmeasured confounding), and the propensity score must not be 0 or 1 to satisfy the positivity assumption. Confounder adjustments in Q-learning implementations using the propensity score can be done through several methods, such as by including the propensity score as a linear term in the Q-function, caliper matching on the propensity score, or inverse probability of treatment weighting with the propensity score. Simulations by Chakraborty and Moodie [77] showed that for large samples, inverse probability of treatment weighting or including covariates as linear terms in the model for the Q-function provided the most improvement for treatment effect estimates in terms of bias.

2.2.3 G-Estimation

Before outlining G-estimation, we first introduce the concept of structural nested mean models (SNMMs). Rather than modelling the conditional means themselves through Q-learning, one can instead model contrasts of conditional mean outcomes for semi-parametric estimation of optimal DTRs - in this way, only difference in outcomes under different treatment regimes are parameterized, rather than the full outcome distribution. We may define the *optimal blip-to-reference-function*, $\gamma_t(h_t, a_t)$, as the expected difference in outcome for persons with the same history h_t who used a treatment a_t at stage t and subsequently receive the optimal regime $\underline{d}_{t+1}^{opt}$ as opposed to a reference regime $\underline{d}_j^{ref}$:

$$\gamma_t(\mathbf{h}_t, a_t) = E[\mathbf{Y}^{\bar{\mathbf{a}}_t, \underline{\mathbf{a}}_{t+1}^{opt}} - \mathbf{Y}^{\bar{\mathbf{a}}_t, a_t^{ref}, \underline{\mathbf{a}}_{t+1}^{opt}} | \mathbf{H}_t = \mathbf{h}_t].$$

Unlike the blip function, the stage t *regret function* considers the difference in outcome from a patient who has been prescribed treatment a_t at stage t compared to a patient with identical history who has received the optimal treatment a_t^{opt} , with all future treatments from stage $t + 1$ onwards being optimal for both patients:

$$\mu_t(\mathbf{h}_t, a_t) = E[\mathbf{Y}^{\bar{\mathbf{a}}_t, \underline{\mathbf{a}}_{t+1}^{opt}} - \mathbf{Y}^{\bar{\mathbf{a}}_t, a_t^{opt}, \underline{\mathbf{a}}_{t+1}^{opt}} | \mathbf{H}_t = \mathbf{h}_t].$$

It is easy to show the relationship between the regret and blip functions as follows:

$$\mu_t(\mathbf{h}_t, a_t) = \max_{a_t} \gamma_t(\mathbf{h}_t, a_t) - \gamma_t(\mathbf{h}_t, a_t)$$

and

$$\gamma_t(\mathbf{h}_t, a_t) = \mu_t(\mathbf{h}_t, d_t^{ref}) - \mu_t(\mathbf{h}_t, a_t).$$

Let the blip and regret functions be parametrized by ψ so that if the true form of the SNMM and the true values of ψ were known, then

$$d_t^{opt}(h_t; \psi) = \arg \max_{a_t} \gamma_t(\mathbf{h}_t, a_t; \psi)$$

or

$$d_t^{opt}(h_t; \psi) = a_t \text{ such that } \mu_t(\mathbf{h}_t, a_t; \psi) = 0.$$

To estimate ψ from the blip or regret functions, Robins [60] proposed G-estimation as a generalization of estimating equations designed to account for time-varying covariates and general treatment regimes.

$$\begin{aligned} G_t(\psi) &= Y + \sum_{u=t}^T [\gamma_t(h_u, d_u^{opt}; \psi) - \gamma_t(h_u, a_u; \psi)] \\ &= Y + \sum_{u=t}^T E[\{Y(\bar{a}_{u-1}, \underline{d}_u^{opt}) - Y(\bar{a}_{u-1}, 0_u, \underline{d}_{u+1}^{opt})\} - \{Y(\bar{a}_u, \underline{d}_{u+1}^{opt}) - Y(\bar{a}_{u-1}, 0_k, \underline{d}_{u+1}^{opt})\} | H_u = h_u] \\ &= Y + \sum_{u=t}^T E[Y(\bar{a}_{u-1}, \underline{d}_u^{opt}) - Y(\bar{a}_u, \underline{d}_{u+1}^{opt}) | H_u = h_u] \end{aligned}$$

Thus, $G_t(\psi)$ represents a person's observed outcome adjusted by the expected difference between the average outcome for someone who received a_t versus someone who received optimal treatment at this stage, where both people had the same history to the start of stage $t - 1$ and were treated optimally from stage $t + 1$ onwards.

By letting $S_t(A_t) = s_t(H_t, A_t)$ be a vector-valued function that contains variables that interact with treatment to effect a difference in the outcome, and modelling the probability of receiving treatment a_t by $p_t(A_{t=1}|H_t; \alpha)$ (e.g. with a logistic propensity score), then one can construct

$$U(\psi, \alpha) = \sum_{t=1}^T G_t(\psi) \{S_t(A_t) - E[S_t(A_t)|H_t; \alpha]\},$$

where $E[U(\psi, \alpha)] = 0$ is an unbiased estimating equation from which ψ can be estimated. This approach was further refined by Robins [60] to use ordinary least squares to model $E[G_t(\psi)|H_t; \beta_t]$, and by letting

$$U(\psi, \beta_t(\psi_t), \alpha) = \sum_{t=1}^T (G_t(\psi) - E[G_t(\psi)|H_t; \beta_t]) \{S_t(A_t) - E[S_t(A_t)|H_t; \alpha]\},$$

then the estimates $\hat{\psi}$ are doubly robust, i.e. the estimators $\hat{\psi}$ for ψ are consistent provided that either $E[G_t(\psi)|H_t; \beta_t]$ or $p_t(A_{t=1}|H_t; \alpha)$ is correctly specified. We denote $E[G_t(\psi)|H_t; \beta_t]$ as the *treatment-free* outcome model and $p_t(A_{t=1}|H_t; \alpha)$ the *treatment* model.

2.3 Dynamic Weighted Ordinary Least Squares Regression (dWOLS)

Wallace and Moodie [80] proposed a DTR method known as dynamic weighted ordinary least squares regression (dWOLS), which estimates the blip parameters ψ for determining the optimal treatment rules. It bears much resemblance to the previously-described Q-learning and G-estimation methods, offering ease of implementation as with the former while maintaining double robustness as with the latter [80]. The dWOLS method involves specification of blip and treatment-free models, as well as a treatment model from which weights are formed and used in a weighted ordinary least squares regression at each stage of treatment [81]. As with G-estimation, estimates of the ψ parameters are consistent provided that at least one of the treatment-free outcome or treatment models are correctly specified, lending the dWOLS method its double-robustness property [80].

We begin by revisiting commonly-used DTR notation in a multi-stage setting as outlined in [80]. Throughout this section, we assume treatment decisions are made in a sequential manner over the course of several follow-up appointments, or stages, for the patient.

In this setting, we let y denote a patient’s continuous outcome such that larger values are generally preferred; examples may include longevity or physical well-being scores. Let a_t denote a binary realization of a patient’s treatment at stage t , $t = 1, 2, \dots, T$, with 0 or 1 representing a “control” or “treatment,” respectively. We define $\mathbf{x}_t$ as a patient’s collection of non-treatment variates at time t ; non-treatment variates may refer to characteristics or measurements such as age, smoking status, blood pressure, and so forth. We further define

$\mathbf{h}_t$ as the patient's t -th stage history, which may include previous treatments a_1 through a_{t-1} in addition to the patient's non-treatment covariates. When expressing a sequence of past treatment stages, $\bar{\mathbf{a}}_t = (a_1, \dots, a_t)$ represents the sequence of treatments up to and including stage t , while $\underline{\mathbf{a}}_{t+1} = (a_{t+1}, \dots, a_T)$ denotes the sequence of future treatments at stage $t + 1$ and onwards.

In order to maximize the expected outcome for a patient given their full history $\mathbf{h}$, we consider the expected mean (possibly counterfactual) outcome Y under treatment regimen $\bar{a}_T$ in terms of the previously-defined blip functions as follows:

$$E[\mathbf{Y}^{\bar{a}_T} | \mathbf{H} = \mathbf{h}] = f(\mathbf{h}_0; \beta) + \sum_{t=1}^T \gamma_t(\mathbf{h}_t, a_t; \psi_t),$$

where $Y^{\bar{a}_T}$ refers to the outcome under the sequence of treatments $\bar{a}_T$ and $f(\mathbf{h}_0; \beta)$ is the *treatment-free* function with covariates $\mathbf{h}_0$ being irrespective to treatment. In order to optimize the t -th treatment decision under this model, we would prescribe treatment $a_t = 1$ if $\gamma_t(\mathbf{h}_t, 1; \psi_t) > 0$ and treatment $a_t = 0$ otherwise. As such, it is of primary interest to estimate the blip parameters ψ in each of the blip functions $\gamma_t(\mathbf{h}_t, a_t; \psi_t), t = 1, \dots, T$. An alternative form of the mean outcome model can also be written in terms of regrets, i.e.

$$E[\mathbf{Y}^{\bar{\mathbf{a}}_T} | \mathbf{H}_T = \mathbf{h}_T] = E[\mathbf{Y}^{\text{opt}} | \mathbf{H}_T = \mathbf{h}_T] - \sum_{j=1}^T \mu_t(\mathbf{h}_t, a_t; \psi),$$

In the dWOLS approach, each stage t of analysis involves the following steps:

1. Let $\tilde{y}_T = y$, the observed outcome at the last stage $t = T$ of analysis. For all previous stages $t < T$, the *pseudo-outcome* is defined as $\tilde{y}_t = y + \sum_{k=t+1}^T \mu_k(\mathbf{h}_k, a_k; \hat{\psi}_k)$, with the intuition that the observed outcome can be viewed as an optimal outcome minus the sum of regrets at each stage. Thus, adding “future” regrets to the observed outcome removes any effects of non-optimality from future treatments, allowing $\tilde{y}_t$ to estimate the stage t blip parameters that will be used in the corresponding optimal stage t treatment rule.
2. Propose a treatment, or propensity score, model $E[\mathbf{A}_t | \mathbf{X}_t = \mathbf{x}_t; \alpha_t] = P(\mathbf{A}_t = 1 | \mathbf{X}_t = \mathbf{x}_t; \alpha_t)$ to obtain estimates of parameters α_t and construct weights that satisfy

$$\pi(\mathbf{x}_t)w(1, \mathbf{x}_t; \alpha_t) = (1 - \pi(\mathbf{x}_t))w(0, \mathbf{x}_t; \alpha_t),$$

where $\pi(\mathbf{x}_t) = P(A_t = 1 | \mathbf{x}_t)$. One possible choice of weights are absolute value weights of the form $w(a_t, \mathbf{x}_t; \alpha_t) = |a_t - E[\mathbf{A}_t | \mathbf{X}_t = \mathbf{x}_t; \alpha_t]|$.

3. Perform a weighted regression of $\tilde{y}_t$, $t = 1, \dots, T$, on $\mathbf{x}_j^\psi$ and $\mathbf{x}_j^\beta$, $j = 1, \dots, N$, two (potentially identical) subsets of variables contained in the collection of explanatory variables for patient j , through a proposed model for the blip function and treatment-free outcome model such as $E[\tilde{\mathbf{Y}}_t^{a_t} | \mathbf{H}_t = \mathbf{h}_t; \beta_t, \psi_t] = \beta_t^T \mathbf{h}_t^\beta + \psi_t^T a_t \mathbf{h}_t^\psi$.
4. Use the blip parameter estimates $\hat{\psi}_t$ at stage t to determine the corresponding optimal treatment rule, for instance, " $a_t^{opt} = 1$ if $\hat{\psi}_{t0} + \hat{\psi}_{t1}x_t > 0$ ".
5. Construct the pseudo-outcome $\hat{\psi}_t$ for the previous stage using the blip parameter estimates $\hat{\psi}_t$ and repeat the above steps as necessary.

By utilizing weights in the serial regressions, possible confounding effects of treatment on blip parameter estimates are removed. Additionally, these parameter estimates are *doubly-robust*, i.e. they are consistent provided that at least one of the treatment model and treatment-free model are correctly specified.

2.4 Optimal Treatment Rules in the Presence of Interference

Su et al. [75] have previously derived the optimal individualized treatment regime using Q-learning and A-learning approaches and a network-based regression model that includes interference effects when modelling the response. Using the same notation as before, let A_j denote a random variable for the binary treatment of individual j , and let $\mathbf{S}_j = \mathbf{A}_{\mathcal{N}_j}$ be the vector of treatment assignments for individuals in the neighbourhood of j . Let $\mathbf{X}_j$ be a p -dimensional vector of covariates for individual j and let the covariate set of their neighbours be defined as $\mathcal{X}_j = \{\mathbf{x}_k : k \in \mathcal{N}_j\}$. Finally, let $Y_j^*(a_j, \mathbf{s}_j)$ and $Y_j(a_j, \mathbf{s}_j)$ denote the potential and observed outcomes for individual j , respectively. We assume the following: (1) $Y_j(a_j, \mathbf{s}_j) = I(A_j = a_j, \mathbf{S}_j = \mathbf{s}_j)Y_j^*(a_j, \mathbf{s}_j)$ and (2) $Y_j^*(a_j, \mathbf{s}_j)$ is independent of $(A_j, \mathbf{S}_j)$ conditional on $(\mathbf{x}_j, \mathcal{X}_j)$. Let $h_j(\mathbf{x}_j, \mathcal{X}_j)$ denote the optimal treatment rule for individual j and let $\mathcal{H}_j : \{h_k(\mathbf{x}_k, \mathcal{X}_k) : k \in \mathcal{N}_j\}$ be the optimal treatment decision set of their neighbours. To obtain the optimal treatment rule, one must maximize the following equation:

$$\frac{1}{N} \sum_{j=1}^N E\{Y_j^*(h(\mathbf{x}_j, \mathcal{X}_j), \mathcal{H}_j) | \mathbf{x}_j, \mathcal{X}_j\},$$

which reduces to maximizing

$$\frac{1}{N} \sum_{j=1}^N E\{Y_j^*(a_j, \mathbf{s}_j) | \mathbf{x}_j, \mathcal{X}_j, a_j = h(\mathbf{x}_j, \mathbf{s}_j = \mathcal{X}_j), \mathcal{H}_j\}.$$

To characterize the effect of interference on an individual's outcome, Su et al. [75] proposed the following network-based regression model:

$$Y_j(a_j, \mathbf{s}_j) = \alpha + \beta' \mathbf{x}_j + \gamma_1 \sum_{k \neq j}^N W_{kj} \beta' \mathbf{x}_k + \eta a_j + \gamma_2 \sum_{k \neq j}^N W_{kj} \eta a_k + a_j \theta' \mathbf{x}_j + \gamma_3 \sum_{k \neq j}^N W_{kj} a_k \theta' \mathbf{x}_k + \epsilon_j$$

where ϵ_j is a random error term with mean 0 and $\mathbf{W} = (W_{kj}) \in \{0, 1\}^{N \times N}$ is the adjacency matrix with $W_{jk} = 1$ if $k \in \mathcal{N}_j$ and 0 otherwise. Using some algebraic manipulations, the authors showed that the outcome model is equivalent to the following:

$$\frac{1}{N} \sum_{j=1}^N (\alpha + \beta' \mathbf{x}_j + \gamma_1 \sum_{k \neq j}^N W_{kj} \beta' \mathbf{x}_k) + \frac{1}{N} \sum_{j=1}^N h(\mathbf{x}_j, \mathcal{X}_j) \{ (1 + \gamma_2 \sum_{k \neq j}^N W_{kj}) \eta + (1 + \gamma_3 \sum_{k \neq j}^N W_{kj}) \theta' \mathbf{x}_j \},$$

with the optimal treatment rule determined as follows:

$$h^{\text{opt}}(\mathbf{x}_j, \mathcal{X}_j) = h^{\text{opt}}(\mathbf{x}_j) = I\{((1 + \gamma_2 \sum_{k \neq j}^N W_{kj}) \eta + (1 + \gamma_3 \sum_{k \neq j}^N W_{kj}) \theta' \mathbf{x}_j) > 0\},$$

where $I(\cdot)$ is the indicator function. Thus, the optimal treatment rule takes on a similar form as outlined in [80], and while it depends on the number of neighbours relative to individual j (as evident through the W_{kj} terms), it only depends on the given individual's covariates $\mathbf{x}_j$.

2.5 Propensity Score Weights for Network Interference

Propensity score models specifying treatment for both a given individual and their neighbours have previously been discussed in the literature on causal inference under network interference. In the case of partial interference, we may consider N groups as before, with n_i individuals in cluster i and $\mathbf{a}_i \in \mathcal{A}(n_i)$ where $\mathcal{A}(n_i)$ consists of all 2^{n_i} different treatment vectors $\mathbf{a}_i$. Rather than simply comparing a policy (or strategy) where all individuals receive treatment to one where no individual receives treatment, it may be of interest to consider strategies where the probability of treatment conditional on pre-treatment covariates varies across individuals, allowing for different overall treatment coverages [8]. In

such settings, we define α as the marginal probability of receiving treatment for a particular policy, i.e. an average over all individuals in the population [8]. We then define $\omega(a_i, n_i, \alpha) = Pr_\alpha(\mathbf{A}_i = \mathbf{a}_i | n_i)$ as the marginal probability of cluster i receiving treatment vector $\mathbf{a}_i$ under treatment program α [8].

It should be noted that the term $P_\alpha(\cdot)$ is not specified, and its form is dependent on the type of treatment allocation program taking place [53]. In a counterfactual world where all groups with the same number of treated units occur with equal probability (regardless of which those units are), we may have $P_{\alpha, X}(\mathbf{a}_i | \mathbf{X}_i) = \prod_{j=1}^{n_i} \alpha^{a_{ij}} (1 - \alpha)^{1-a_{ij}}$ [53]. However, it may be more realistic for the treatment allocation program to consider individual- and group-specific covariates or incorporate correlation between treatment of units within a group, modelled with a specification such as $\text{logit} P_{\alpha, X}(A_{ij} = 1 | X_{ij}) = \xi_i^\alpha + \sigma X_{ij}$, where σ and ξ_i^α are fixed such that $n_i^{-1} \sum_{j=1}^{n_i} \text{expit}(\xi_i^\alpha + \sigma X_{ij}) = \alpha$ [53]. With such a specification, α parametrizes within-cluster treatment that depends on covariates X_{ij} , but is on average equal to α [53].

Joint propensity scores have also been used to construct weights for causal effect estimators in settings with network interference. Liu et al. [41] have proposed inverse-probability weighted estimators in the context of general interference when no assumption is made on the structure or form of interference present. They denote the propensity scores of individual j and of individual j and their group of neighbours as $f(a_j | \mathbf{x}_j) = Pr(A_j = a_j | \mathbf{x}_j)$ and $f(a_j, s_j | \mathbf{x}_j, \mathcal{X}_j) = Pr(A_j = a_j, S_j = s_j | \mathbf{x}_j, \mathcal{X}_j)$, respectively. The inverse probability-weighted estimator for treatment a under allocation strategy α is defined as follows:

$$\hat{Y}^{\text{ipw}}(a, \alpha) = N^{-1} \sum_j \frac{y_j(A_j, S_j) I(A_j = a) \pi(S_j; \alpha)}{f(A_j, S_j | \mathbf{x}_j, \mathcal{X}_j)},$$

while the inverse probability-weighted marginal estimator under strategy α may be defined as

$$\hat{Y}^{\text{ipw}}(a, \alpha) = N^{-1} \sum_j \frac{y_j(A_j, S_j) \pi(S_j; \alpha)}{f(A_j, S_j | \mathbf{x}_j, \mathcal{X}_j)},$$

where $\pi(s_j; \alpha) = \alpha^{\sum s_j} (1 - \alpha)^{|s_j| - \sum s_j}$.

Under the assumption of partial interference with N different groups under consideration, the above estimators can be easily adapted such that for group i ,

$$\hat{Y}_i^{\text{ipw}}(a, \alpha) = n_i^{-1} \sum_{j=1}^{n_i} \frac{y_{ij}(A_j, S_j) I(A_j = a) \pi(S_j; \alpha)}{f(A_j, S_j | \mathbf{x}_j, \mathcal{X}_j)},$$

and

$$\hat{Y}_i^{\text{ipw}}(\alpha) = n_i^{-1} \sum_{j=1}^{n_i} \frac{y_{ij}(A_j, S_j) \pi(S_j; \alpha)}{f(A_j, S_j | \mathbf{x}_j, \mathcal{A}_j)},$$

as shown in [76]. Although the inverse probability-weighted estimators proposed by Tchetgen and VanderWeele [76] are developed under the assumption of within-group interference, the authors propose a mixed model for the propensity score which allows for the parameters of interest in the model to remain consistent for non-null associations between exposure levels of individuals belonging to a fixed number of different groups.

2.6 dWOLS in the Presence of Interference with Network Propensity Weights

To date, the biostatistical literature on applications of dWOLS and other DTR methods to interference networks has been limited. Previous work by Jiang et al. [32] has explored implementations of dWOLS in networks with interference, showing that a failure to account for interference can lead to biased estimators of treatment parameters and, in turn, badly mis-identified treatment rules. When considering dWOLS in the context of a network, it is useful to refer to the *ego* as the person of primary interest for whom a treatment decision is being made, and those to whom the ego is linked may be referred to as the *alters* or contacts [32]. Assuming that egos and alters are observed from a network where each individual is treated independently, the *network propensity function* for individual k may be defined as:

$$\pi_{k, a_k, s_{k,a}}(\mathbf{x}_k, \mathcal{N}_k, \mathbf{X}_{\mathcal{N}_k}) = \mathbb{P}(a_k, s_{k,a} | \mathbf{x}_k, \mathcal{N}_k, \mathbf{X}_{\mathcal{N}_k}),$$

where $a_k = 1$ indicates the treatment status of individual k , $s_{k,a}$ represents the neighbours of individual k with treatment status 1, $\mathbf{x}_k$ represents the collection of covariates for individual k , $\mathcal{N}_k$ represents the collection of neighbours for individual k , and $\mathbf{X}_{\mathcal{N}_k}$ represents the collection of covariates for individual k 's neighbours. By defining the standard propensity as $\pi_k = \mathbb{P}(A_k = 1 | \mathbf{x}_k)$, the network propensity function may be expressed as follows:

$$\begin{aligned} \pi_{k, a_k, s_{k,a}}(\mathbf{x}_k, \mathcal{N}_k, \mathbf{X}_{\mathcal{N}_k}) &= \mathbb{P}(a_k, s_{k,a} | \mathbf{x}_k, \mathcal{N}_k, \mathbf{X}_{\mathcal{N}_k}), \\ &= \pi_k(\mathbf{x}_k)^{a_k} (1 - \pi_k(\mathbf{x}_k))^{1-a_k} \times \prod_{l \in s_{k,a}} \pi(\mathbf{x}_l) \prod_{l \in N_k, l \notin s_{k,a}} (1 - \pi(\mathbf{x}_l)). \end{aligned}$$

Using the network propensity function, Jiang et al. [32] showed that the double robustness of dWOLS can be recovered under network interference in two ways using the following theorems:

Theorem 2.6.1. *When the true outcome model satisfies $E[Y|\mathbf{x}, a, t(a_{\mathcal{N}})] = f(\mathbf{x}^\beta, t(a_{\mathcal{N}}); \beta) + \gamma(\mathbf{x}^\psi, t(a_{\mathcal{N}}); \psi)$, a weighted ordinary least squares regression based on the corresponding linear model will yield consistent estimators of ψ if either at least one of the treatment and treatment-free models is correctly specified, and the weights satisfy*

$$w_k(\mathbf{x}_k, \mathcal{N}_k, \mathbf{x}_{\mathcal{N}_k}, a_k, s_{k,a}) \pi_{k,a_k,s_{k,a}}(\mathbf{x}_k, \mathcal{N}_k, \mathbf{X}_{\mathcal{N}_k}) = \mathcal{C}(\mathbf{x}_k, \mathbf{x}_{\mathcal{N}_k}, |\mathcal{N}_k|).$$

where $\mathcal{C}$ is a function that depends only on the variables $(\mathbf{x}_k, \mathbf{x}_{\mathcal{N}_k})$ and the degree $|\mathcal{N}_k|$.

Theorem 2.6.2. *When the true outcome model satisfies $E[Y|\mathbf{x}, a, t(a_{\mathcal{N}})] = f(\mathbf{x}^\beta; \beta) + \gamma^\xi(t(a_{\mathcal{N}}), \mathbf{x}^\xi; \xi) + \gamma^\psi(\mathbf{x}^\psi, t(a_{\mathcal{N}}); \psi)$, a weighted ordinary least squares regression based on the corresponding linear model will yield consistent estimators of ψ as well as ξ if either at least one of the treatment and treatment-free models is correctly specified, and the weights satisfy $\pi(\mathbf{x})w(1, \mathbf{x}) = (1 - \pi(\mathbf{x}))w(0, \mathbf{x})$ or its sufficient but not necessary condition*

$$w_k(\mathbf{x}_k, \mathcal{N}_k, \mathbf{x}_{\mathcal{N}_k}, a_k, s_{k,a}) \pi_{k,a_k,s_{k,a}}(\mathbf{x}_k, \mathcal{N}_k, \mathbf{X}_{\mathcal{N}_k}) = \mathcal{C}(\mathbf{x}_k, \mathbf{x}_{\mathcal{N}_k}, |\mathcal{N}_k|).$$

One possible form of network weights are those formed using "absolute value" or "overlap" weights where the network weights for individual k in a network are:

$$w_k = |a_k - E[A_k|X_k = x_k]| * \Pi_l |a_l - E[A_l|X_l = x_l]|,$$

where $l \in \mathcal{N}_k$ [32]. In a network where all egos have the same number of neighbours, the joint distribution of the treatment of the ego and the number of their treated neighbours is balanced with respect to their covariates, and the weights are comparable.

Under this setting, several key assumptions must be made; in particular, that each individual is treated independently of others according to the same randomization rule [32]. This assumption simplifies estimating the group propensity score, since it can be expressed as the product of a given individual's estimated propensity score and each of their neighbours' [32]. Additionally, there is the assumption that each individual's propensity of treatment depends only on their own covariate values and not those of their neighbours, which may depend on the characteristics of the network at hand [32].

This modified dWOLS approach with network propensity weights has been implemented in simulation studies for isolated couples, i.e. when there is one ego and their corresponding neighbour (alter), as well as five types of fixed networks: a ring of points on a circle; a square lattice; an Erdos-Renyi network realization; a Barabasi-Albert network realization; and the Longleaf Pines spatial network [32]. These types of networks provide structure on defining the neighbours for a given node or ego, as a ring or square lattice are relatively simple structures with constant degree while a spatial network defines the neighbours of a given node based on the spatial distance between each other [32]. Thus, the dWOLS method offers much promise in determining optimal DTRs in networks with interference.

There are two key advantages of network propensity weights compared with the standard dWOLS approach: first, network weights offer more flexibility in cases where the treatment-free model contains interference terms, such as neighbours' covariates [32]. In this case, by decomposing the ego's and alters' treatment and treatment-free models, then blip parameter estimators are consistent provided either part of the ego's treatment or treatment-free model is correctly specified, and if the ego's treatment and treatment-free models are correctly specified but the alters' treatment-free model is not, then the spillover effects estimators can also be consistently estimated from dWOLS with network weights [32]. Secondly, incorporating neighbours' treatment model into the analysis can help increase the efficiency of blip parameter estimates when the alter's treatment-free model is mis-specified [32].

While this approach offers ease of implementation for partial interference when only within-group interference is taking place, several limitations may arise when between-group interference takes place. Jiang et al. [32] note that their approach assumes that individuals' propensities depend on their own covariate values, and this in turn could depend to a limited degree on the characteristics of the network of interest. In hierarchical datasets, interest may lie on accounting for characteristics of both the individual and group as a whole that can affect the resulting treatment status [15]. Additionally, Jiang et al. [32] proposed that network propensity weights can be extended to situations with different ego-degree mixture cases by dividing the network weight by the average value of the network weight over all nodes of the same degree, which in turn may reduce the variability of weights overall and boost estimation efficiency. By considering scenarios in which between-group interference only occurs between a select number of subgroups within a hierarchy, propensity scores with group-level covariates and random effects may offer a solution to both limitations by accounting for group-level features and correlation between individuals' treatments with ease of implementation. In Chapter 3, we will extend this approach by conducting simulation studies in which a mixed effects model propensity score is used to construct the dWOLS weights for hierarchical-structured observational

datasets under within- and between-group interference, and compare the performance to standard and network propensity weights.

2.7 Propensity Score Modelling for Multilevel Data

While the propensity score has frequently been explored in analyses with unstructured data, there has been work in extending its application to multilevel or hierarchical data through the incorporation of random effects. Arpino and Mealli [2] showed that estimation of the propensity score for hierarchical data with multilevel models incorporating random effects generally achieve good balancing in cluster level unobserved covariates, as does a single level model with dummy variables for each cluster. Furthermore, random-effects models are a desirable method for precision medicine research in that a random coefficient can be considered as a parameter that remains a characteristic constant for a particular patient while varying across other patients [30]. It is therefore intuitive to consider a mixed effects model approach for propensity score-based weights in the context of DTR methodology for network interference.

Previous work has also developed propensity-score weighted marginal, clustered, and doubly robust estimators in the context of hierarchical data [38]. These may be of particular interest in the context of network interference because accounting for the hierarchical nature of clustered data through propensity score-weighted estimators gives insight to different group-level characteristics that may affect treatment propensity and, in turn, influence the extent of interference occurring within one group versus another. For instance, both unit-level covariates of individual j in group i , denoted U_{ij} , and cluster-level covariates V_i can comprise $\mathbf{X}_{ij} = (U_{ij}, V_i)$ in the model:

$$\text{logit}(e_{ij}) = \delta_i + \mathbf{X}_{ij}\alpha,$$

where δ_i has a prior distribution such as $\delta_i \sim N(\delta_0, \sigma_\delta^2)$ [38]. Indeed, these authors have noted that accounting for clustering in propensity score analyses can account for violations of no unmeasured cluster-level confounders that result in cluster-level effects and violate the SUTVA [38]. Consideration must be given when determining whether to include a cluster-specific variable as a fixed effect or a random effect; in general, V_i should be included as a fixed effect if the number of clusters is relatively small, and as a random effect otherwise [85].

Tchetgen Tchetgen and Vanderweele [76] have also considered joint propensity scores using mixed effects models which account for both cluster-specific and individual-specific

variation. They note that an appeal of this approach is that it accounts not only for possible associations between treatments of two individuals, A_{ij} and $A_{ij'}$, in a group conditional on covariates $\mathbf{X}_i$, but also between treatments of two individuals from different groups [76]. Perez-Heydrich et al. [55] used inverse-probability weighted (IPW) estimators to estimate treatment effects in the presence of interference using data from an individually-randomized, placebo-controlled trial of cholera vaccination in Matlab, Bangladesh. When estimating the group-level propensity scores, they used a mixed effects logit model as proposed by Tchetgen Tchetgen and Vanderweele [76] to estimate the probability of trial participation, combined with the known probability of vaccine assignment (2/3) conditional on participation in the trial. A binary indicator for participation of individual j in group i , B_{ij} , was modelled as

$$h_{ij}(b_i) = P(B_{ij} = 1 | \mathbf{X}_{ij}, b_i) = \text{logit}^{-1}(\mathbf{X}_{ij}\phi_a + b_i),$$

where $\mathbf{X}_{ij}$ is a $1 \times p$ covariates of individual j in group i and b_i is the random effect for group i , where $b_i \sim N(0, \phi_b)$ [55]. The joint probability of treatment (i.e. vaccination) for group i was then modelled as

$$Pr(\mathbf{A}_i | \mathbf{X}_i; \phi) = \int \prod_{j=1}^{n_i} \{(2/3)h_{ij}(b_i)\}^{A_{ij}} \{1 - (2/3)h_{ij}(b_i)\}^{1-A_{ij}} f_b(b_i; \phi_b) db_i,$$

where $\mathbf{A}_i$ denotes the vector of treatment indicators for all individuals in group i and $\mathbf{X}_i = (\mathbf{X}_{i1}, \dots, \mathbf{X}_{in_i})^T$ is an $n_i \times p$ matrix of the covariates from all n_i individuals in group i . As such, the literature to date has proposed mixed effects models for the propensity score in analyses for observational data that involve within-group interference and, by extension, some degree of between-group interference as well.

2.7.1 Network Weights for Multilevel Data

To extend the network weights proposed by Jiang et al. [32] to settings where interference occurs within and between groups, we consider J distinct groups to which an individual can belong. For a given ego k in group j , $j = 1, \dots, J$, we let $\mathcal{N}_k$ consist of all alters relative to ego k such that $\mathcal{N}_k$ can be partitioned into all alters who also belonging to group j , denoted $\mathcal{N}_{\text{within},k}$, and all alters belonging to groups $j', j' \neq j$, denoted $\mathcal{N}_{\text{between},k}$; that is, $\mathcal{N}_{\text{within},k} \cup \mathcal{N}_{\text{between},k}$. In this way, $\mathcal{N}_{\text{within},k}$ consists of all within-group alters for ego k while $\mathcal{N}_{\text{between},k}$ consists of all between-group alters for ego k . We let $\mathbf{v}$ denote cluster-level variables shared by all individuals across the J different groups and let b_j denote a random intercept for group j . We can now define the standard propensity for ego k as

$$\pi_k = \mathbb{P}(A_k = 1 | \mathbf{x}_k, \mathbf{v}, b_j) = \text{logit}^{-1}(\mathbf{x}_{jk}\phi_a + \mathbf{v}\phi_b + b_j),$$

the standard propensity for alter $l \in \mathcal{N}_{\text{within},k}$ as

$$\pi_l = \mathbb{P}(A_l = 1 | \mathbf{x}_l, \mathbf{v}, b_j) = \text{logit}^{-1}(\mathbf{x}_{jl}\phi_a + \mathbf{v}\phi_b + b_j),$$

and the standard propensity for alter $l \in \mathcal{N}_{\text{between},k}$ as

$$\pi_l = \mathbb{P}(A_l = 1 | \mathbf{x}_l, \mathbf{v}, b_{j'}) = \text{logit}^{-1}(\mathbf{x}_{jl}\phi_a + \mathbf{v}\phi_b + b_{j'}).$$

As defined in 2.6, let $s_{k,a}$ represent the alters of individual k with treatment status 1, which can be partitioned into $s_{j,k,a} \cup s_{j',k,a}$ with $s_{j,k,a}$ representing within-group alters with treatment status 1 in group j and $s_{j',k,a}$ representing between-group alters with treatment status 1 in group j' . Then by assuming independent treatment assignment conditional on all cluster-level variables and random effects, the joint propensity score for ego k and their alters in $\mathcal{N}_k$ can be expressed as

$$\begin{aligned} \pi_{k,a_k,s_{k,a}}(\mathbf{x}_k, \mathcal{N}_k, \mathbf{x}_{\mathcal{N}_k}, \mathbf{v}, b) &= \mathbb{P}(a_k, s_{k,a} | \mathbf{x}_k, \mathcal{N}_k, \mathbf{x}_{\mathcal{N}_k}, \mathbf{v}, b) \\ &= \pi_k(\mathbf{x}_k)^{a_k} (1 - \pi_k(\mathbf{x}_k))^{1-a_k} \times \prod_{l \in s_{j,k,a}} \pi(\mathbf{x}_l) \prod_{l \in N_{\text{within},k}, \notin s_{j,k,a}} (1 - \pi(\mathbf{x}_l)) \times \\ &\quad \prod_{l \in s_{j',k,a}} \pi(\mathbf{x}_l) \prod_{l \in N_{\text{between},k}, \notin s_{j',k,a}} (1 - \pi(\mathbf{x}_l)). \end{aligned}$$

Using this network propensity function, one can construct network weights as outlined in 2.6, which will be demonstrated in simulation studies throughout Chapters 4 and 5.

Chapter 3

Optimizing Post-ED Visit Dispositions Among Opioid Overdose Patients Using Age-Tailored Guidelines with Dynamic Weighted Ordinary Least Squares

3.1 Introduction

To date, most applications of DTR methods have been done in contexts that use information from problem-specific clinical protocols, such as diagnosis, patient-specific symptoms, or other clinically relevant information. A systematic review of methods used to estimate DTRs using observational healthcare data found that primary focus areas included critical care, cardiovascular disease, and oncology [39]. While many DTR applications have focused on determining sequences of treatments to maximize a patient-specific physiological or behavioural outcome (such as the impact of the timing of introducing solid food into an infant’s diet on their BMI [80] or the use of e-cigarettes on cigarette dependence [32]), there is a gap in the literature on using DTRs for health service system optimization, where the treatments and outcome of interest can be a reflection of hospital resource allocation. In this chapter, we provide a motivating application in the context of managing OUD to show how the DTR framework can be used to identify patterns of post-ED dispositions in observational datasets and recommend more efficient approaches with an outcome of interest that reflects both patient care and healthcare burden, such as the number of physician

visits needed following ED disposition. This will also provide context for applications of the methodology introduced in Chapter 4, which extends these ideas to a more complex setting where the SUTVA is likely to be violated among pWOUD within social and spatial networks.

3.1.1 Overview of the Opioid Overdose Crisis in Canada

The opioid overdose crisis is an ongoing national public health crisis in Canada, with nearly 20,000 opioid-related poisoning ED visits reported in 2024 between January and September (72 per day on average) [12]. This crisis is not restricted to a particular region, but findings from 2016 showed that these rates were highest in western provinces such as BC and Alberta, as well as in Yukon and the Northwest Territories [40]. There are many contributing factors to the increase in opioid-related hospitalizations and overdose deaths, such as the over-prescribing of opioid medications; social-structural determinants of health such as homelessness and poverty; and the ongoing criminalization of drugs that potentially encourages people to consume drugs in less safe ways [72]. The COVID-19 pandemic of 2020 introduced new challenges to the opioid overdose crisis, as physical distancing measures posed difficulties for pWOUD to access health clinics and pharmacies for medications and to utilize supervised consumption services (SCS) [43]. Furthermore, health and economic factors such as mental health challenges associated with social isolation and income instability have been thought to place pWOUD at higher risk in the context of the COVID-19 pandemic [43]. The opioid overdose crisis in Canada is a complex issue that motivates the need to integrate evidence-based medication treatments with psychosocial interventions that support basic and personalized needs for pWOUD, which will improve their longevity and quality of life [22].

The opioid overdose crisis in Canada is detrimental not only to pWOUD, but also to the Canadian healthcare system. Understanding the characteristics of frequent emergency department users is critical to developing effective solutions to improve patient care and decrease financial burden of health care [47]. As national rates of mortality and hospitalizations continue to rise, the broader impact of opioids on ED service utilization needs to be studied to inform the urgency with which alternative interventions should be implemented [73]. This is necessary to improve patient care for people with OUD while ensuring that ED services can remain available to other groups of patients [73]. A population-based retrospective cohort study using administrative ED records from the National Ambulatory Care Reporting System (NACRS) database study showed a substantial increase in opioid-related ED visits and paramedic transports in Ontario between 2009 and 2019, with the largest increase of ED visits found in young adults (19 – 35 years) who also comprise the

largest age demographic that experience fatal opioid overdoses and poisonings [73]. Given that the proportion of ED patients admitted to hospital or a significant increase in returned ED visit within 30 days was shown to have remained stable over the 11 year period, these findings suggest that EDs are capable of managing and treating the vast majority of opioid-related patient conditions. However, the authors emphasize the need for interventions that improve patient-centered care while reducing use of scarce ED services [73]. Another study using data from between 2012 and 2017 found that the highest poisoning-related ED visit rates in BC were among individuals between the ages of 25 and 44 years for males and 15 to 19 years for females [54]. Studies such as these emphasize the importance of identifying subpopulations at high risk for poisoning-related ED presentations [54] and the need to implement timely strategies aimed at safely reducing opioid-related ED visits with alternative care models [73].

3.1.2 Motivating Example: A Personalized Treatment Approach to Opioid Overdose Harm Reduction and Prevention

The complex and chronic nature of OUD presents many opportunities for analysis within the framework of precision medicine, where an optimal sequence of services can be tailored based on individual characteristics such as age or living arrangement. While the need for person-centered approaches for prevention and treatment of pWOUD has been recognized, there are limited studies on how to incorporate patient preferences in these treatments, and patterns of engagement in prevention services and treatment are often complex [11]. One study found that personally-tailored opioid overdose prevention education and naloxone distribution (PTOEND) programs were effective in decreasing perceived treatment barriers, self-reported opioid-use, and overdose-risk behaviours [84]. In addition to including information on general risk factors, overdose recognition, and naloxone utilization, these programs included education about personal overdose-risk factors and medication for opioid use disorder (MOUD) [84]. A recent study by Treitler et al. [77] examined the association of an ED peer recovery support service with post-discharge addiction treatment initiation, repeat overdose, and acute care utilization, and their findings showed that implementation of ED peer recovery support service was associated with an increase in MOUD initiation and a decrease in repeat medically treated overdoses. However, treatment effects were heterogeneous and likely dependent on factors such as availability of other hospital- and community-based OUD services [77]. These studies highlight examples of patient-specific factors and service utilization that can be incorporated into personalized treatment strategies for pWOUD, as well as the ever-increasing need for integrative prevention and treatment approaches to combat the opioid overdose crisis.

To assess characteristics of frequent ED users in BC, Canada, Moe et al. [47] performed a retrospective study from the BC Provincial Overdose Cohort using health administrative databases from 2012/13 to 2015/2016. The median age across frequent ED users (defined as the top 10% ED users in each fiscal year) in 2015/16 was 54 (IQR: 35–73), while the median age across non-frequent users (defined as the bottom 90% ED users in each fiscal year) in 2015/16 was 50 (IQR: 33–66) [47]. The authors defined visit acuity (i.e. severity) scores using the Canadian Triage and Acuity Scale (CTAS) [31], and their study found that 81.3% of frequent users had a CTAS score of 3, 4, or 5, indicating a somewhat stable (i.e. non-emergent or life-threatening) presentation, and this percentage was similar among non-frequent users (82.6%) [47].

Another retrospective cohort study by Moe et al. [46] found that, across all patients in BC who had an opioid-related overdose between 2015 and 2016, the median age was 35.0 (IQR: 27.0–46.0), with 51.1% being in the age group of 14 – 35 years. Among patients with non-overdose related visits, the median age was 41.0 (IQR: 28.0–56.0), with 38.6% being in the age group of 14 – 35 years [46]. Their findings also showed that 65.6% of all patients in BC who had an opioid-related overdose between 2015 and 2016 had a CTAS score ranging from 3 to 5, compared to 82.7% among patients with non-overdose related visits [46].

These findings demonstrate the importance of implementing strategies that can reduce opioid-related ED visits among patients who do not present life-threatening conditions [73], and they motivate the following research question:

“Among patients who are in relatively stable conditions as reflected by their CTAS scores, how can the DTR framework be used to predict the optimal disposition (being discharged to home versus admitted or transferred to another facility) for an individual with OUD based on their age to minimize the number of physician visits needed following ED visits?”

Personalizing treatments for OUD in the form of service system utilization remains an open and challenging problem due to the complex nature of patient interactions with networks of providers and the variety of different services that may be available across different healthcare facilities. In the following section, we will illustrate an application of the dWOLS DTR as previously outlined in 2.3 to an analysis of synthetic data simulated from the demographic summaries provided in the retrospective cohort study by Moe et al. [46], and we emphasize that in this context, a patient’s “treatment” will refer to their disposition (being discharged to home or admitted to another facility) following an ED visit.

3.2 Application: Optimizing Disposition Rules for Opioid Overdose Patients with an Age-Tailored DTR

We demonstrate an application of dWOLS for predicting whether an individual would benefit from being discharged to home versus admitted or transferred to another facility following an overdose-related visit to the ED, using simulated data summarized in the retrospective cohort study by Moe et al. [46]. Their study analyzed administrative data for ED visits from linked NACRS data in the BC Provincial Overdose Cohort, which is a collection of linked administrative data for persons who experienced an opioid-related overdose between January 1, 2015 and December 31, 2017 using specific case definitions in multiple data sources, such as BC Emergency Health Services and NACRS [46]. Using these data, the authors’ study cohort consisted of persons 14 to 74 years of age who made at least one nonfatal overdose-related ED visit between January 1, 2015, and December 31, 2016, and they defined the index overdose visit as the first nonfatal, non-overdose-related ED visit [46]. Patient characteristics associated with index overdose visits to EDs in BC are summarized from 2015-16, including CTAS scores, age (median and IQR), sex, disposition after index visit, psychiatric comorbidities, material deprivation quintile, Elixhauser index (with and without mental health), and health care encounters such as physician visits made in the 365 days before the patient’s index visit (median and IQR).

The objective of this analysis is to identify whether the number of physician visits following an ED visit can be minimized by determining whether or not a patient would benefit from being discharged to home versus admitted or transferred to another facility based on their age at the time of the recorded ED visit. Due to several limitations with the data, and for the purpose of illustrating dWOLS, we only consider “age” as the tailoring variate from which an estimated treatment rule is constructed.

In this setting, we define the j -th “stage” of analysis as the time from the patient’s j -th recorded ED visit up to (but not including) their $j + 1$ visit, with stage $j = 1$ representing the time between a patient’s first visit (that is, their index visit as defined in [46]) up to their second ED visit. Since the characteristic summaries captured by Moe et al. [46] are for patients’ index visits from January 1, 2015 to December 31, 2016, the study does not explicitly provide demographic information for patients with more than one ED visit at stages $j = 2$ or onwards. To address these limitations, we propose several assumptions to assist with simulating necessary variables needed for a two-stage analysis:

- Patient covariates for sex, age, and presence of a psychiatric comorbidity remain unchanged from the onset of stage 1 to the onset of stage 2; the assumption of age remaining unchanged is reasonable if we assume that all patients had at least one

ED Visit (Stage)	Age at ED Visit	Disposition (1=Discharged, 0=Admitted)	Stage j Outcome
1	Age ₁	1	Number of physician visits between ED visits 1 and 2
2	Age ₂	1	Number of physician visits between ED visits 2 and 3
3	Age ₃	0	Number of physician visits between ED visits 3 and 4
...	...	...	...

Figure 3.1: Illustration of data for a two-stage analysis using simulated data from a patient with three recorded ED visits.

subsequent ED visit that occurred within a year from their index visit, although in practice this could be modified so that a patient’s age increased by at least 1 year by stage 2. Summaries of an alternative analysis under this assumption are available in Appendix B.

- The simulated CTAS score at stage 2 depends on frequencies of simulated CTAS scores observed at stage 1.
- The outcome at stage j is defined as the number of physician visits made between stages j and $j + 1$. For a two-stage analysis, we would determine whether a patient would benefit from being discharged to home versus admitted or transferred to another facility at their first or second recorded overdose, based on their age at each time, to minimize the total number of physician visits elapsed between ED visits 1 – 2 and ED visits 2 – 3.

This is elucidated in Figure 3.1.

3.2.1 Data Pre-Processing and Methods

Prior to the data cleaning and pre-processing stage, the cohort consisted of 3593 patients who had at least one nonfatal overdose-related ED visit between Jan. 1, 2015 and Dec. 31, 2016 [46]. Among these patients, 3133 were discharged to home or residence, 150 left against medical advice, and 310 were admitted or transferred to another facility following their index ED visit. The median age for all patients was 35 years (IQR: 27 – 46), and the




Figure 3.2: Overview of selection criteria for 100 simulated overdose datasets.

median number of physician visits made was 15 (IQR: 5–39). Among all patients who were discharged, 2.3% had a CTAS score of 1, 28.3% had a CTAS score of 2, 62.8% had a CTAS score of 3, 5.6% had a CTAS score of 4, 2.3% had a CTAS score of 4, 0.6% had a CTAS score of 5, and 0.4% had a missing or unknown CTAS score. 30% of discharged patients were female, and 46.5% of discharged patients were recorded as having a psychiatric comorbidity classified as anxiety, depression, schizophrenia, or a personality disorder.

Of the 310 patients who were admitted to hospital or transferred, 16.1% had a CTAS score of 1, 55.8% had a CTAS score of 2, 24.5% had a CTAS score of 3, 2.9% had a CTAS score of 4, and 0.65% had a missing or unknown CTAS score. Among these patients, 40.6% of these patients were female, and 74.5% were recorded as having a psychiatric comorbidity classified as anxiety, depression, schizophrenia, or a personality disorder.

From each simulated cohort, we chose to exclude all patients who left against medical advice since our “treatment” of interest is to determine whether a patient would benefit from either being discharged with approval or admitted or transferred to another facility for further care. Patients with CTAS scores of 4 or 5 were merged into one category due to lower counts in these groups relative to other CTAS scores in the dataset. Since the demographics from [46] are for index visits to the ED but no information for follow-up visits are available, we make several assumptions for simulating stage 2 variables. First, we simulate CTAS scores for stage 2 such that the probability of a patient with a CTAS score of 1 or 2 has an 80% probability of being assigned a CTAS score of 2, while the probability of a patient with a CTAS score of 3 or 4 has an 80% probability of having a CTAS score of 3 at stage 2. We let $\text{Age} \sim N(35, 14)$, and for patients who are discharged, the number of physician visits is represented as $Y \sim N(14, 4)$ while for patients who are admitted to hospital, $Y \sim N(15, 4)$. We assume patient covariates for sex, age, and presence of a psychiatric comorbidity remain the same from stage 1 to stage 2. We then predicted the probability of being discharged to home or residence at stage 2 ($A_2 = 1$) versus admitted or transferred ($A_2 = 0$) using a logistic model regressing disposition on CTAS score at stage 2, sex, age, and presence of a psychiatric comorbidity. Once these stage 2 covariates were simulated, we excluded patients with a CTAS score of 1 or 2 from analysis given that we wish to determine how hospital resources can be strategically allocated among patients with an overdose-related ED visit who exhibit “urgent” or “less urgent” conditions [73], rather than among those who require immediate resuscitation or are in emergency (CTAS scores of 1, and 2, respectively) [31]. Using these inclusion-exclusion criteria as outlined in Figure 3.2 and our assumptions for simulating stage 2 variables, our average sample size across 100 simulated datasets was 2246 (IQR: 2228 – 2268).

Across the 100 simulated datasets, the average median age at stages 1 and 2 was 36 years (IQR: (28 – 45)). Additional demographic summaries based on the simulated

CTAS and Disposition Stratified by Stage and Sex					
Stage	Sex	Triage level (CTAS)		Disposition	
		3	4	Discharged	Admitted
1	Male	90.9 %	9.1 %	96.7 %	3.3 %
	Female	91.1 %	8.9 %	94.9 %	5.1 %
2	Male	80.2 %	19.8 %	59.0 %	41.0 %
	Female	79.5 %	20.5 %	59.4 %	40.6 %

Table 3.1: Average proportions of CTAS scores and dispositions across males and females at stages 1 and 2 of analysis from 100 simulated datasets.

datasets are presented in Table 3.1. For each dataset, we estimate the probability of a patient being discharged to home or residence ($A = 1$), $\mathbb{E}[A_j|\mathbf{X}_j]$ at each stage j using a logistic regression model, where the probability of being discharged is dependent on sex, age, and presence of a psychiatric comorbidity. Our treatment-free model models the average number of physician visits between ED visits as a linear function of sex, age, CTAS score, and presence of a psychiatric comorbidity, while our blip model only contains age at ED visit as a tailoring variable, i.e. $\gamma(x_{ji}, a_{ji}) = a_{ji}(\psi_{j0} + \psi_{j1}x_{ji})$ where x_{ji} is patient i 's age at ED visit j . Implementation of dWOLS using these models was done using the R package `DTRreg` [82].

3.3 Results

The average blip parameter estimates and corresponding 95% confidence intervals from 100 simulated datasets are summarized in Table 3.2. To use these estimates to form an optimal treatment rule for post-ED disposition, one would identify a patient's age at stage $j = 1, 2$ and identify the sign of the equation $\hat{\psi}_{0j} + \hat{\psi}_{1j} \times \text{age at index visit}$. At stage 1, the optimal treatment rule would be "Discharge a patient to home or residence if $0.930 + 0.00134 \times \text{Age} > 0$," and since age is positive, this would translate to discharging all patients to home or residence regardless of their age. At stage 2, however, the optimal treatment rule would be "Discharge a patient to home or residence if $-0.886 + 0.0206 \times \text{Age} > 0$," i.e. a patient would be discharged to home or residence if their age is greater than 43 years

Average Cohort Findings			
Stage (j)	$\hat{\psi}_{0j}$	$\hat{\psi}_{1j}$	Mean j -th Stage Regret
1	0.930 (-1.831 ,3.693)	0.00134 (-0.0703, 0.073)	0.0556
2	-0.886 (-7.833, 6.060)	0.0206 (-0.158, 0.195)	0.561

Table 3.2: Mean blip parameter estimates and regrets from 100 simulations.

and admitted to hospital or transferred otherwise.

Using the above estimates, we calculate the proportion of all individuals who undertook the estimated optimal disposition at stages 1 and 2 based on the DTR tailored to age across the 100 simulated datasets. On average, in stage 1, 91.9% of all patients were “correctly” disposed or admitted as recommended by the estimated DTR, while at stage 2, this proportion decreased to 47.7%. The average median age of all patients who did receive optimal treatment in the first stage was 36 (IQR: 28, 45), and at stage two, the median was 40 (IQR: 34, 48). Among all patients who did not receive optimal treatment, the median age of those who did not receive optimal treatment in the first stage was 37 (IQR: 29, 46), and at the second stage, the median was 35 (IQR: 29, 42). Among patients who received the optimal recommended dispositions at stages 1 and 2, the average median number of physician visits was 13.94, while among those who did not receive the optimal recommended dispositions at stage 1 or stage 2, the the average median number of physician visits was 14.19

3.4 Discussion

DTRs are a key methodology in precision medicine research as a way to tailor treatments based on individual-specific information and improve patient outcomes. This is the first application of DTRs to health service system optimization, with the objective being to minimize the number of physician visits between successive opioid-related ED visits to reduce healthcare burden and improve allocation of scarce hospital resources. We note that

our findings are primarily for proof of concept, as our data face multiple limitations. First, due to needing to simulate data based on summary demographics for index visits to the ED, we are limited by the number of clinically-relevant variables that may be used to determine an optimal treatment rule. Although our analysis considered age as the tailoring variate due to these limitations, we note that it would be of interest to consider other factors such as patient symptoms, living arrangement, history of previously-received treatments for OUD, or self-reported satisfaction with use of prior services when constructing an optimal treatment rule. Utilizing comprehensive datasets containing these variables and clinical expertise on substance use treatment would be necessary to differentiate between associational and causal relationships, and this in turn would be considered when constructing treatment-free, treatment, and blip models for a stronger analysis. The simulation of stage 2 variables was based on several simplifying assumptions, among these being that an individual with a given CTAS score in the first stage was likely to maintain the same CTAS score in the second stage (approximately 80%). We also note that in practice, subsequent “stages” of treatment past the index ED visit may occur at different points in time for different patients; in our analysis, we defined a “stage” as the time between the j -th and $j + 1$ visits and assumed that the second ED visit occurred within a year of all patients’ first ED visit so that their ages remain unchanged. However, actual dates of ED visits likely differ among patients, and indeed, future work may involve a time-to-event analysis of similar data within a DTR framework using the `DWSurv` package in R [68], where the occurrence of an overdoses is an “event” of interest and distinguished from a given “stage” of treatment.

Under this set-up, our findings showed that the proportion of individuals whose observed disposition aligned with the predicted optimal disposition from the estimated DTR decreased in the second stage relative to the first stage. In other words, younger patients tended to be discharged to home or residence when they would have benefited from being admitted to hospital or transferred in the second ED visit. Although these findings were obtained from an analysis of simulated data under our previously-outlined assumptions, they nonetheless demonstrate how DTR methods such as `dWOLS` can be used to provide insights into characteristics of patients who may benefit from one post-ED disposition over another, in turn suggesting how treatment strategies can be effectively allocated to prioritize scarce hospital resources. For instance, younger adults may less likely be connected with community-based treatment, harm reduction programs, and other support that could help manage OUD and reduce the number of physician visits needed post-ED visit. While much work remains to be explored in health service system utilization among pWOD, our findings address an important gap in the DTR literature on personalizing treatment strategies for the purpose of healthcare burden management.

Chapter 4

dWOLS with Mixed Effects Propensity Score Models in Hierarchical Networks under Between- and Within-Group Interference

4.1 Research Objective

Our first research objective is to develop a modification of dynamic weighted ordinary least squares regression (dWOLS) for nested hierarchical data that accounts for both within- and between-group interference, extending on previous DTR methodology that accounts for general forms of interference in social networks [32]. To account for group-level variability in treatment assignment, we consider a network propensity score based on a logistic mixed-effects treatment model that incorporates fixed effects and random intercepts at two levels of grouping in a hierarchical network, and we show that our resulting network weights take the same form as in [32]. To illustrate the importance of accounting for both forms of interference that may take place in a hierarchical network, we compare our implementation of dWOLS with other modified dWOLS approaches that only consider within-group interference or that consider both forms of interference but use different weights. Thus, we expand on previous work by considering hierarchical datasets in which group-level variation can affect the propensity of treatment and assessing the performance of dWOLS when the interference structure is mis-specified or when different




Figure 4.1: DAG of simulation studies in *Study 1a*, where $\mathbf{X}_N$ denotes the individual-level covariates of the ego’s alters, $\mathbf{A}_N$ denotes the treatments of the ego’s alters, V denotes a group-level variable corresponding to the hierarchy that the ego and alter belong to, and Y_{ego} represents the ego’s outcome.

weighting approaches are used. We compare the performance of these different dWOLS implementations in terms of optimal treatment rates and blip and spillover parameter estimates in both one-stage and two-stage settings, with an important assumption that a given ego’s alters are randomly assigned but fixed across stages. Lastly, we conclude Chapter 4 with a real-world application to service class utilization among opioid overdose patients, where interference may take place within and between groups of individuals in the same geographical proximity.

4.2 Simulation Studies: One- and Two-Stage Settings

4.2.1 *Study 1a*

We first consider a one-stage scenario where hierarchical networks of data are generated as follows: given I different hierarchies (Level 3 of grouping), there are J distinct subgroups nested in each hierarchy (Level 2 of grouping) with n_{ij} individuals (Level 1) in each subgroup for a total network size of $N = I \times J \times n_{ij}$ individuals. For simplicity, we assume that the J number of subgroups is fixed across each hierarchy and n_{ij} is the same for each subgroup. We further assume that while there is both within- and between-group




Figure 4.2: Hierarchical network structure, with interference occurring within and between Level 2 groups as depicted using arrows within and between individuals at Level 1.

interference across the J different subgroups in a given hierarchy, no interference occurs between the hierarchies (Level 3 groups). That is, we restrict between-group interference to only occur at Level 2 in each hierarchy. Letting n_i denote the total number of units in hierarchy i , $i = 1, \dots, I$, we assume n_i is the same across hierarchies and n_{ij} is the same across subgroups within each hierarchy, i.e. $n_i = J \times n_{ij}$.

For our one-stage setting, we simulate I different hierarchies, each consisting of J different subgroups with n individuals belonging to each subgroup. In each hierarchy i , $i = 1, \dots, I$, we randomly sample n_{egos} individuals from the $J \times n_{ij}$ total individuals in the hierarchy, with the remaining $(J \times n_{ij}) - n_{\text{egos}}$ individuals randomly assigned as alters. This results in a total of $I \times n_{\text{egos}}$ individuals for analysis. When assigning alters, we ensure that each ego is guaranteed to have at least one alter as a neighbour; however, no restriction is placed on the subgroup to which an alter must belong. In this way, we allow for interference to occur both within subgroups (if, say, all of an ego's neighbours also belong to the same subgroup at level J) as well as between subgroups in the same hierarchy.

We include an i -th level cluster-level covariate $V_i \sim N(1, 1)$ and an individual-level covariate $X_{ijk} \sim N(2, 1)$ in the treatment model for individual k in subgroup j nested within hierarchy i . Treatment status is modelled as $A_{ijk} \sim \text{Bernoulli}(\text{expit}(\pi_{ijk}))$ where $\text{logit}(\pi_{ijk}) = -1 + X_{ijk} + 0.5V_i + 0.25 \sin(X_{ijk}) + b_{0i} + b_{0ij}$, with $b_{0i} \sim N(0, \tau_3)$, $b_{0ij} \sim N(0, \tau_2)$ representing random intercepts at the I and J levels of grouping, respectively, and $b_{0i} \perp b_{0ij}$. By assuming independence of b_{0i}, b_{0ij} , it follows that $b = b_{0i} + b_{0ij} \sim N(0, \tau_3 + \tau_2) = N(0, \tau)$. Our outcomes are modelled as $Y_{ijk} \sim N(\mu_{ijk}, 1)$ where

$$\mu_{ijk} = X_{ijk} + X_{ijk}^2 + V_i + N_{\text{total},ijk} + N_{\text{total},ijk} V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 N_{\text{total},ijk} + N_{\text{total},ijk} V_i),$$

with $N_{\text{total},ijk}$ denoting the total number of treated alters relative to an ego, from both the same and different groups at the J -th level. The optimal treatment rule at stage 1 is:

$$A^{opt} = 1 \text{ if } \psi_0 + \psi_1 X_{ijk} + \psi_2 N_{\text{total},ijk} + \psi_3 N_{\text{total},ijk} V_i > 0, 0 \text{ otherwise.}$$

Using a similar structure to the group propensity model proposed by [55], we construct IPTW-style weights for the egos based on the joint probability of treatment with their assigned alters. For ego k in group j nested in group i :

$$w_{ijk} = h_{ijk}(b) \}^{A_{ijk}} \{1 - h_{ijk}(b)\}^{1-A_{ijk}} \times \prod_{l=1}^{N_{\text{total},ijk}} \{h_{ijl}(b)\}^{A_{ijl}} \{1 - h_{ijl}(b)\}^{1-A_{ijl}},$$

where $h_{ijk}(b) = \Pr(A_{ijk} = 1 | X_{ijk}, V_i, b; \hat{\alpha})$ is the estimated probability of treatment for individual j in group i conditional on covariates X_{ijk}, V_i and random effects $b = b_{0i} + b_{0ij}$.

We consider four methods for this analysis. In Method 1 (ME1), we implement dWOLS using network propensity weights but incorrectly assume that interference only occurs within groups so that $N_{\text{total},ijk}$ refers to the total number of treated alters relative to ego k within the same group j in hierarchy i . That is,

$$w_{ijk} = |a_{ijk} - E(A_{ijk} | X_{ijk} = x_{ijk}, V_i = v_i)| \times \prod_l^{N_{\text{within},ijk}} |a_{ljk} - E(A_{ljk} | X_{ljk} = x_{ljk}, V_i = v_i)|,$$

where $\prod_l^{N_{\text{within},ijk}}$ is taken over all of the ego's alters that are in the same group at Level 2, and

$$E[Y|x, v, a; \beta, \psi] = \beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 v + \xi_0 N_{\text{within}} + \xi_1 N_{\text{within}} v + a(\psi_0 + \psi_1 x + \psi_2 N_{\text{within}} + \psi_3 N_{\text{within}} v).$$

In Method 2 (ME2), we correctly specify the interference structure so that $N_{\text{total},ijk}$ includes treated alters relative to ego k within the same and/or possibly other groups j in hierarchy i . However, we use “standard” overlap weights that account only for the ego's treatment. That is,

$$w_{ijk} = |a_{ijk} - E(A_{ijk} | X_{ijk} = x_{ijk}, V_i = v_i)|,$$

and

$$E[Y|x, v, a; \beta, \psi] = \beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 v + \xi_0 N_{\text{total}} + \xi_1 N_{\text{total}} v + a(\psi_0 + \psi_1 x + \psi_2 N_{\text{total}} + \psi_3 N_{\text{total}} v).$$

In Method 3 (ME3), we correctly specify the interference structure so that $N_{\text{total},ijk}$ includes treated alters relative to ego k regardless if they are in the same or other groups j in hierarchy i . We use network propensity weights that account for the ego and all of their alters' treatments. That is,

$$w_{ijk} = |a_{ijk} - E(A_{ijk}|X_{ijk} = x_{ijk}, V_i = v_i)| \times \Pi_l^{N_{\text{total},ijk}} |a_{ljk} - E(A_{ljk}|X_{ljk} = x_{ljk}, V_i = v_i)|,$$

where $\Pi_l^{N_{\text{total},ijk}}$ is taken over all of the ego's alters at Level 3, and

$$E[Y|x, v, a; \beta, \psi] = \beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 v + \xi_0 N_{\text{total}} + \xi_1 N_{\text{total}} v + a(\psi_0 + \psi_1 x + \psi_2 N_{\text{total}} + \psi_3 N_{\text{total}} v).$$

Lastly, Method 4 (ME4) is the same as ME3, but rescales the network weights by dividing each ego's weight by the mean of all other ego's weights who have the same number of alters. This was proposed as an extension by Jiang et al. [32] to account for varying ego-alter degree mixtures and improve efficiency.

To assess the double-robustness property of the different dWOLS methods under misspecification of either the treatment or treatment-free models, we will then test these methods under the following four scenarios:

- Scenario 1: The treatment-free model is correctly specified, but the treatment model is mis-specified via omission of the non-linear term in X_{ijk} .
- Scenario 2: The treatment model is correctly specified but the treatment-free model is mis-specified via omission of the non-linear term in X_{ijk} .
- Scenario 3: Both the treatment- and treatment-free models are mis-specified via omission of the non-linear term in X_{ijk} .
- Scenario 4: Both the treatment- and treatment-free models are correctly specified.

Mixed effects model parameters for treatment in ME2, ME3, and ME4 were estimated using the "glmer" function in the `lme4` package in R. We set $\psi = (1, 1, -1, -1.5)$ as our blip parameters and simulate 1000 datasets with $I = 10$, $n = 50$, $n_{\text{egos}} = 50$, $\tau_3 = 1$, and $\tau_2 = 1$, varying the number of subgroups J in each hierarchy from 2 to 4. In this way, the total number of egos n_{egos} is fixed at $I \times n_{\text{egos}} = 500$, but the total number of alters available for each ego increases as J increases. For performance metrics, we define the mean optimal treatment rate as the proportion of individuals in the simulated datasets who would be optimally treated if they had received the treatment as determined by the corresponding method's decision rule, as outlined in [81]. We also include empirical standard error (EmpSE) estimates, defined by Morris et al. [49] as

J	Mean Optimal Treatment Rate			Empirical Standard Error (EmpSE)		
	2	3	4	2	3	4
Observed	0.581	0.458	0.394	0.0652	0.0627	0.0631
ME1	0.812	0.706	0.712	0.0507	0.0487	0.0532
ME2	0.970	0.978	0.982	0.0211	0.0123	0.0109
ME3	0.960	0.960	0.958	0.0296	0.0248	0.0273
ME4	0.960	0.965	0.967	0.0296	0.0210	0.0212

Table 4.1: Mean optimal treatment rates for Scenario 1 in *Study 1a*, where the treatment model is mis-specified via omission of the non-linear term in X_{ijk} .

$$\sqrt{\frac{1}{n_{\text{sim}}-1} \sum_{r=1}^{n_{\text{sim}}} (\hat{\theta}_r - \bar{\theta})^2},$$

where $\hat{\theta}_r$ is the estimated mean optimal treatment rate in simulated dataset r and $\bar{\theta}$ is the mean of $\hat{\theta}_r$ from all 1000 simulated datasets.

Our findings for the mean optimal treatment rates across the 1000 simulated datasets in the one-stage setting for Scenarios 1 and 2 are presented in Tables 4.1 and 4.2, and the results for Scenarios 3 and 4 can be found in Appendix C. Boxplots of blip and spillover parameter estimates are shown in Figures 4.3 to 4.6. Mis-specification of the interference term by only accounting for within-group alters results in lower optimal treatment rates as the number of subgroups J , and therefore the total number of potential alters from different subgroups, increases from 2 to 4. ME2, ME3, and ME4 correctly specify the interference exposure term in the blip, and therefore all three methods perform well in terms of the mean optimal treatment rates. Consistent with the work of Jiang et al. [32], we find that rescaling the network weights by the total number of alters assigned to an ego results in lower EmpSE for the mean optimal treatment rates. Interestingly, we note that ME2 tends to produce blip parameter estimates with lower variance than ME3 or ME4 as J increases from 2 to 4. We speculate that this is because network weights have been shown to offer some efficiency gains when the treatment-free model can be decomposed into separate components corresponding to the ego and alter [32], while the treatment-free model in *Study 1a* does not include individual-specific covariates relating to the alter.

4.2.2 *Study 1b*

In *Study 1a*, we considered an outcome model for the ego such that the treatment-free component did not contain individual-level covariates X relating to the alter. Our results




Figure 4.3: Boxplots of ψ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 1 in *Study 1a*, where the treatment model is mis-specified via omission of the non-linear term in X_{ijk} .




Figure 4.4: Boxplots of ψ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 2 in *Study 1a*, where the treatment-free model is mis-specified via omission of the non-linear term in X_{ijk} .

J	Mean Optimal Treatment Rate			Empirical Standard Error (EmpSE)		
	2	3	4	2	3	4
Observed	0.581	0.458	0.394	0.0652	0.0627	0.0631
ME1	0.797	0.700	0.709	0.0557	0.0556	0.0543
ME2	0.949	0.960	0.966	0.0360	0.0236	0.0200
ME3	0.932	0.930	0.931	0.0502	0.0423	0.0443
ME4	0.932	0.938	0.943	0.0502	0.0370	0.0340

Table 4.2: Mean optimal treatment rates for Scenario 2 in *Study 1a*, where the treatment-free model is mis-specified via omission of the non-linear term in X_{ijk} .




Figure 4.5: Boxplots of ξ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 1 in *Study 1a*, where the treatment model is mis-specified via omission of the non-linear term in X_{ijk} .




Figure 4.6: Boxplots of ξ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 2 in *Study 1a*, where the treatment-free model is mis-specified via omission of the non-linear term in X_{ijk} .

showed that, if the interference structure is correctly specified to account for both within- and between-group interference, then blip parameter estimates using ME2 with standard overlap weights had lower variance than those obtained using ME3 or ME4, which used network propensity weights. We speculate that this is because simulation studies by Jiang et al. [32] found that network propensity weights resulted in lower variances of $\hat{\psi}$ when the treatment-free model can be decomposed into the ego’s covariates (TF-ego) and the alter’s covariates (TF-alter), and at least one of TF-ego, TF-alter, or the treatment model is correctly specified. This in turn produced an efficiency gain with higher optimal treatment rates, specifically when TF-alter is mis-specified. Network weights were also shown to offer consistent estimators of spillover effect terms when the treatment model was correctly specified, unlike standard overlap weights.

In *Study 1b*, we now consider a more complex form of the outcome model in a setting where each ego in a hierarchy with $J = 2$ subgroups is assigned exactly one alter from the same subgroup and one alter from the other subgroup so that each ego is paired with two alters. The treatment-free model can be decomposed into the ego’s covariates (TF-ego), a within-group alter’s covariates (TF-alter-W), and a between-group alter’s covariates (TF-alter-B). The outcome model is then modelled as $Y_{ijk} \sim N(\mu_{ijk}, 1)$, where

$$\mu_{ijk} = \sin(X_{ego}) + X_{alter,within}^2 + \exp(X_{alter,between}) + V_i + \xi_0 N_{total,ijk} + \xi_1 N_{total,ijk} V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 N_{total,ijk} + N_{total,ijk} V_i).$$

We consider two scenarios for this analysis and compare ten methods in each. In Scenario (a), we correctly specify the same treatment model as in *Study 1a*, i.e. $A_{ijk} \sim \text{Bernoulli}(\text{expit}(\pi_{ijk}))$ where $\text{logit}(\pi_{ijk}) = -1 + X_{ijk} + 0.5V_i + 0.25 \sin(X_{ijk}) + b_{0i} + b_{0ij}$, with $b_{0i} \sim N(0, \tau_3)$, $b_{0ij} \sim N(0, \tau_2)$ representing random intercepts at the I and J levels of grouping, respectively, and $b_{0i} \perp b_{0ij}$. In Scenario (b), we incorrectly specify the treatment model via omission of the non-linear term in X . Table 4.3 outlines the different methods for comparison under each of these scenarios, where some or all components of the treatment-free model are mis-specified and the choice of weights (i.e. network weights or standard overlap weights) also varies across methods.

As in *Study 1a*, we set $\psi = (1, 1, -1, -1.5)$ as our blip parameters and simulate 1000 datasets with $I = 10$, $n_{ij} = 75$, $n_{egos} = 50$, $\tau_3 = 1$, and $\tau_2 = 1$, keeping the number of subgroups fixed at $J = 2$. In this way, the total number of egos n_{egos} is fixed at $I \times n_{egos} = 500$, and each ego has exactly two alters, one from their same J subgroup and one from the other subgroup.

The mean optimal treatment rates and EmpSE across all six methods for *Study 1b* are presented in Table 4.4. Consistent with the findings from simulations by Jiang et al. [32],




Figure 4.7: DAG of simulation studies in *Study 1c*, where $\mathbf{X}_{1,\mathcal{N}}$ and $\mathbf{X}_{2,\mathcal{N}}$ denote the individual-level covariates of the ego's alters at stages 1 and 2, respectively, while $\mathbf{A}_{1,\mathcal{N}}$ and $\mathbf{A}_{2,\mathcal{N}}$ denote the treatments of the ego's alters at stages 1 and 2, respectively.

network propensity weights offer some efficiency gains as evident by a higher mean optimal treatment rate compared to standard overlap weights in ME7 and ME8 across Scenarios (a) and (b), when TF-alter-B is mis-specified. In Appendix C, we also observe efficiency gains from network propensity weights when both TF-ego and TF-alter-B, or TF-alter-W and TF-alter-B, are mis-specified at the same time across Scenarios (a) and (b).

4.2.3 *Study 1c*

We now extend our one-stage simulations to a two-stage setting where the number of alters relative to an ego does not change from the first stage, but both the ego's and their alters' covariate information changes between stages. In these simulations, we again simulate I different hierarchies, each consisting of J different subgroups with n_{ij} individuals

Scenario (a): Treatment Model Correct					
Method	Network Weights	Standard Weights	TF-ego	TF-alter-W	TF-alter-B
ME1	✓	✗	✓	✓	✓
ME	✗	✓	✓	✓	✓
ME3	✓	✗	✗	✓	✓
ME4	✗	✓	✗	✓	✓
ME5	✓	✗	✓	✗	✓
ME6	✗	✓	✓	✗	✓
ME7	✓	✗	✓	✓	✗
ME8	✗	✓	✓	✓	✗
ME9	✓	✗	✗	✗	✗
ME10	✗	✓	✗	✗	✗
Scenario (b): Treatment Model Incorrect					
ME1	✓	✗	✓	✓	✓
ME2	✗	✓	✓	✓	✓
ME3	✓	✗	✗	✓	✓
ME4	✗	✓	✗	✓	✓
ME5	✓	✗	✓	✗	✓
ME6	✗	✓	✓	✗	✓
ME7	✓	✗	✓	✓	✗
ME8	✗	✓	✓	✓	✗
ME9	✓	✗	✗	✗	✗
ME10	✗	✓	✗	✗	✗

Table 4.3: Overview of methods for comparison under each scenario in *Study 1b*, with a ✓ indicating which weights have been implemented in a given method and which component of the outcome model has been correctly specified, while an ✗ indicates which weights have not been implemented and which component of the outcome model has been incorrectly specified.

Scenario (a)		
Method	Mean Optimal Treatment Rate	Empirical Standard Error (EmpSE)
Observed	0.421	0.0764
ME1	0.968	0.0209
ME2	0.979	0.0137
ME3	0.961	0.0265
ME4	0.974	0.0165
ME5	0.950	0.0358
ME6	0.965	0.0233
ME7	0.822	0.137
ME8	0.813	0.123
ME9	0.796	0.120
ME10	0.796	0.120
Scenario (b)		
Method	Mean Optimal Treatment Rate	Empirical Standard Error (EmpSE)
Observed	0.421	0.0764
ME1	0.967	0.0210
ME2	0.979	0.0137
ME3	0.961	0.0267
ME4	0.974	0.0164
ME5	0.951	0.0349
ME6	0.965	0.0232
ME7	0.828	0.134
ME8	0.812	0.122
ME9	0.795	0.121
ME10	0.795	0.121

Table 4.4: Optimal treatment rates under Scenarios (a) and (b) in *Study 1b*.

belonging to each subgroup, and we randomly sample n_{egos} individuals from the $J \times n_{ij}$ total individuals in the hierarchy, with the remaining $(J \times n_{ij}) - n_{\text{egos}}$ individuals randomly assigned as alters. In the first stage, we simulate cluster-level covariate $V_{1,i} \sim N(1, 1)$ and individual-level covariate $X_{1,ijk} \sim N(2, 1)$; in the second stage, the cluster- and individual-level covariates are, respectively, $V_{2,i} \sim N(1 + 0.5V_{1,i}, 1)$ and $X_{2,ijk} \sim N(X_{1,ijk}, 1)$.

As in the one-stage simulations, treatment status at stage 1 is modelled as

$$A_{1,ijk} \sim \text{Bernoulli}(\text{expit}(\pi_{1,ijk})),$$

where $\text{logit}(\pi_{1,ijk}) = -1 + X_{1,ijk} + 0.5V_{1,i} + 0.25 \sin(X_{1,ijk}) + b_{0i} + b_{0ij}$, with $b_{0i} \sim N(0, \tau_3)$, $b_{0ij} \sim N(0, \tau_2)$, and treatment status at stage 2 is modelled as

$$A_{2,ijk} \sim \text{Bernoulli}(\text{expit}(\pi_{2,ijk})),$$

where $\text{logit}(\pi_{2,ijk}) = -1 + X_{2,ijk} + 0.5V_{2,i} + 0.25 \sin(X_{2,ijk}) + b_{0i} + b_{0ij}$.

As in *Study 1b*, our second-stage simulations incorporate the alters' covariates in the treatment-free component of the outcome model; specifically, $Y_{ijk} \sim N(\mu_{ijk}, 1)$, where

$$\begin{aligned} \mu_{ijk} = & \log |V_{1,i}| + \log |V_{2,i}| + \exp(X_{1,ijk}) + \exp(X_{2,ijk}) + |\tau(\mathbf{X}_{1,N_{\text{total}}})| + |\tau(\mathbf{X}_{2,N_{\text{total}}})| + \\ & N_{\text{total},ijk} + N_{\text{total},ijk}V_{1,i} + N_{\text{total},ijk} + N_{\text{total},ijk}V_{2,i} + A_{1,ijk}(\psi_0 + \psi_1X_{1,ijk} + \psi_2N_{\text{total},ijk} + \\ & N_{\text{total},ijk}V_{1,i}) + A_{2,ijk}(\psi_0 + \psi_1X_{2,ijk} + \psi_2N_{\text{total},ijk} + N_{\text{total},ijk}V_{2,i}), \end{aligned}$$

and where $\tau(\mathbf{X}_{1,N_{\text{total}}}) = \sum_{l \in \mathcal{N}_{\text{total}}} X_{1,ijl} / N_{\text{total}}$ and $\tau(\mathbf{X}_{2,N_{\text{total}}}) = \sum_{l \in \mathcal{N}_{\text{total}}} X_{2,ijl} / N_{\text{total}}$ denote the average of individual-level covariate measurements from the alters of ego k at stages 1 and 2, respectively. As in the one-stage simulations, the optimal treatment rule at stage 1 is:

$$A^{1,opt} = 1 \text{ if } \psi_0 + \psi_1X_{1,ijk} + \psi_2N_{\text{total},ijk} + \psi_3N_{\text{total},ijk}V_{1,i} > 0, 0 \text{ otherwise,}$$

and the optimal treatment rule at stage 2 is:

$$A^{2,opt} = 1 \text{ if } \psi_0 + \psi_1X_{2,ijk} + \psi_2N_{\text{total},ijk} + \psi_3N_{\text{total},ijk}V_{2,i} > 0, 0 \text{ otherwise.}$$

We set $\psi = (1, 1, -1, -1.5)$ as our blip parameters and simulate 1000 datasets with $I = 10$, $n_{ij} = 50$, $n_{\text{egos}} = 50$, $\tau_3 = 1$, and $\tau_2 = 1$, keeping the number of subgroups fixed at $J = 2$ so that each alter is assigned exactly one alter either from the same subgroup J or the other subgroup.

In Figures 4.8 and 4.9, we show that our results for the blip and spillover parameter estimates are consistent with theoretical results of consistency in the two-stage setting using ME2 and ME3 when the treatment model is mis-specified via omission of the non-linear term in X_{ijk} . Similarly, in Figures 4.10 and 4.11, we show consistency in the blip and spillover parameter estimates in the two-stage setting using ME3 when the treatment-free model is mis-specified by using linear terms in $X_{1,ijk}$, $X_{2,ijk}$, $V_{1,i}$, and $V_{2,i}$ in the treatment-free model.

4.3 Application to Service Class Utilization in Opioid Overdose Patients

4.3.1 Data Pre-Processing and Methods


We now demonstrate an application of our methods for optimizing service system utilization among opioid overdose patients using data simulated from demographic summaries of 3593 patients in the retrospective cohort study by Moe et al. [46]. Since 150 individuals left against medical advice and 13 individuals who were either discharged or admitted to hospital had unknown CTAS scores, we simulate stage 1 and stage 2 covariates for sex, CTAS scores, age, presence of a psychiatric comorbidity, and total number of physician visits for a total of 3430 individuals using the same approach as outlined in 3.2. We then subset the data to consider patients with a CTAS score of 3, 4, or 5. To generate a hierarchical structuring of the data, and to induce correlation in the dispositions among patients, we simulate random intercepts for 5 different provincial regions and 10 different facilities, with two facilities nested in each provincial region. randomly partition the patients into 5 different hierarchies each consisting of 2 different subgroups. For stages $j = 1, 2$, treatment is modelled as

$$A_{ijk} \sim \text{Bernoulli}(\pi_{ijk}),$$


where

$$\text{logit}(\pi_{ijk}) = 3 + 0.25 \times CTAS_j - Psych + 0.05 \times Sex + 0.016 \times Age + b_{0i} + b_{0ij},$$


with $b_{0i} \sim N(0, 1)$, $b_{0ij} \sim N(0, 1)$ representing random intercepts at the regional and facility levels of grouping, respectively, and $b_{0i} \perp b_{0ij}$. The coefficients for CTAS, presence of a psychiatric comorbidity, sex, and age were chosen based on an estimated logistic model for stage 1 disposition, but the random intercepts induce correlation between the dispositions of patients nested within the same facilities in a given region.




(a)



(b)



(c)



(d)

Figure 4.8: Boxplots of ψ parameter estimates for $J = 2$ in a two-stage setting in *Study 1c*, where the treatment model is mis-specified via omission of the non-linear term. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.




Figure 4.9: Boxplots of ξ parameter estimates for $J = 2$ in a two-stage setting in *Study 1c*, where the treatment model is mis-specified via omission of the non-linear term. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.







Figure 4.10: Boxplots of ψ parameter estimates for $J = 2$ in a two-stage setting in *Study 1c*, where the treatment-free model is mis-specified by using linear terms in $X_{1,ijk}$, $X_{2,ijk}$, $V_{1,i}$, and $V_{2,i}$. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.




(a)



(b)



(c)



(d)

Figure 4.11: Boxplots of ξ parameter estimates for $J = 2$ in a two-stage setting in *Study 1c*, where the treatment-free model is mis-specified by using linear terms in $X_{1,ijk}$, $X_{2,ijk}$, $V_{1,i}$, and $V_{2,i}$. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.

4.3.2 Ego and Alter Assignments Within and Between Facilities

Since our simulated data do not explicitly identify relationships or measures of contact between individuals, we perform a random ego-alter assignment assuming that individuals within the same provincial region can interact with other individuals who are either within a facility of the same name or possibly a different one within the same region. For simplicity, we assume that patients have not moved between provincial regions following their ED visits, and when assigning egos and alters, we randomly sample egos from each region so that each ego is guaranteed to have at least one alter who is either considered as “belonging” to the same facility or possibly a different facility but still within the same region. For instance, if there were 457 patients belonging to a given region, we would randomly designate 228 of these individuals as egos and all other remaining patients within that region as potential alters such that each ego has at exactly one alter. We again emphasize that this method for selecting egos and alters is strictly for proof of concept, and we do not consider the dates during which the egos and their randomly assigned alters had overdoses recorded; rather, we simply focus on treating the period of time after a recorded overdose as its own “stage” of treatment.

We perform 100 simulations of randomized ego-alter assignments to assess our blip and spillover parameter estimates under the different ego-alter assignments for each simulation. The average sample size was 2246, and the average number of egos was 560.

As in Chapter 3, we represent the j -th stage “treatment” received by a patient as a binary realization of whether that individual was discharged to home or residence following the j -th ED visit, i.e. $A_j = 1$ if the patient was discharged and $A_j = 0$ admitted to hospital or transferred. Our treatment-free model uses CTAS score, presence of a psychiatric comorbidity, age, sex, and and treatment of the ego’s assigned alter, i.e. $f(\mathbf{x}_j^\beta; \beta_j, \xi_j) = \beta_{0j} + \beta_{1j}CTAS_j + \beta_{2j}Psych_j + \beta_{3j}Age_j + \beta_{4j}Sex_j + \xi_{0j}A_j^{alter}$. Our tailoring variates are “age” and the treatment received by the ego’s alter, i.e. $\gamma(a_j, \mathbf{x}_j^\psi; \psi_j) = a_j(\psi_{0j} + \psi_{1j}Age_j + \psi_{2j}A_j^{alter})$.

4.3.3 Results

In stage 1, we found that across the 100 simulated datasets with different ego-alter assignments, if we accounted for both within- and between-group interference, 50.60% of patients had received the service class assignment as determined by the estimated DTR on average, and at stage 2, this average “optimal treatment rate” had stayed approximately the same at 49.92%. The optimal treatment rule at stage 1 can be interpreted as “At stage 1, regardless of whether a patient’s alter has been discharged or admitted to hospital,

Within- and Between-Group Interference						
ψ	Blip Estimates			ξ	Spillover Estimates	
	Tailoring Covariate	stage 1	stage 2		stage 1	stage 2
ψ_0	Disch	0.0560	-0.243	ξ_0	0.0654	0.0236
ψ_1	Disch*Alter	-0.0238	0.00785			
ψ_2	Disch*Age	0.000712	0.00638			

Within-Group Interference						
ψ	Blip Estimates			ξ	Spillover Estimates	
	Tailoring Covariate	stage 1	stage 2		stage 1	stage 2
ψ_0	Disch	0.00794	-0.226	ξ_0	0.105	0.0188
ψ_1	Disch*Alter	-0.0297	-0.00649			
ψ_2	Disch*Age	0.00172	0.00668			

Table 4.5: Mean blip and spillover parameter estimates from 100 different ego-alter assignments in the simulated opioid overdose dataset.

discharge the patient regardless of their age.” The optimal treatment rule at stage 2 can be interpreted as “At stage 2, if a patient’s alter has been discharged, discharge the patient if they are older than 37 years of age. If the patient’s alter has been admitted to hospital or transferred, then discharge the patient if they are older than 38 years of age.”

We compare these results to a scenario that only considers within-group interference. In this setting, the mean optimal treatment rates are similar, with 51.09% of patients having received the estimated optimal disposition at stage 1 and 50.68% of patients having received the estimated optimal disposition at stage 2. The optimal treatment rule at stage 1 can be interpreted as “At stage 1, regardless of whether a patient’s alter has been discharged or admitted to hospital, discharge the patient regardless of their age.” The optimal treatment rule at stage 2 can be interpreted as “At stage 2, if a patient’s alter has been discharged, discharge the patient if they are older than 35 years of age. If the patient’s alter has been admitted to hospital or transferred, then discharge the patient if they are older than 34 years of age.”

We emphasize that these estimated treatment rules are strictly for proof of concept due to limitations of the simulated data and therefore should not be implemented in practice without clinical expertise. However, they illustrate how accounting for the treatment of an alter who is either within the same or different facility can be incorporated into an application of dWOLS under hierarchical network interference.

4.4 Discussion

We have explored different variations of the dWOLS method in scenarios where both within- and between-group interference is taking place. Considering hierarchical structured data provides an intuitive way of depicting scenarios where both forms of interference may be occurring, and in our simulations, we showed that correctly specifying accounting for both forms of interference results is important for obtaining consistent blip parameter estimators and, subsequently, high optimal treatment rates. Limiting the number of groups through which between-group interference can occur is important for a computational perspective, particularly when using network propensity weights since the weights will become smaller as the number of neighbours for an individual increases [32]. Indeed, as the number of groups within a hierarchy increases, it becomes more important to rescale the network weights to ensure lower variability in the blip estimates. Nonetheless, our findings in the one-stage simulations show that ME2, ME3, and ME4, which all account for between-group interference with different weighting methods, result in higher mean optimal treatment rates than the dWOLS method ME1 that only considers within-group

interference.

Our simulations offer many directions in future work applying dWOLS to settings with between- and within-group interference. It may be of interest to consider multi-stage settings in which only within-group interference occurs in earlier stages of treatment, but later extends to between-group interference occurring at later stages as well. While our propensity weights are based on mixed effects logistic models when the treatment model is correctly specified, other weighting methods that account for the clustered nature of the dataset may be explored in future simulations. Additionally, our treatment models considered random intercepts to demonstrate some amount of group-specific variability at the baseline of treatment, whereas random slopes would indicate that there is group-specific variability in how a patient’s covariates affect their probability of treatment. Future analyses may consider the level of complexity needed in a treatment model with random and fixed effects when constructing standard and network weights, as this could affect the high variability and extreme outliers observed in some of our blip and spillover parameter estimates. There may also be more complex scenarios that could be modelled with a DAG for hierarchical network interference, where an alter’s outcome affects the outcome of an ego, or when the objective is to simultaneously optimize the outcome of the ego and the outcomes of their within-group alters (e.g. optimizing the outcomes of an ego and their alters who are attending the same facility for treatment) while accounting for some level of between-group interference.

Although we were limited by available patient-specific demographics from the simulated opioid overdose dataset, our findings nonetheless illustrate how DTRs with interference can be applied to post-ED visit disposition optimization. This remains a promising area of research for DTRs, as to the best of our knowledge, DTRs have yet to be explored in the context of health service system utilization. Considering the dispositions taken by close contacts when determining whether a given individual may benefit from being discharged to home or admitted to hospital following an ED visit may have important implications for resource allocation when hospital resources are scarce, and it also allows researchers to gain potential insights as to how an individual’s engagement with post-ED treatment is influenced by their social or spatial network.

Chapter 5

Application of dWOLS with Network Propensity Weights with Varying Extents of Interference

5.1 Introduction

In Chapter 4, we showed that failing to account for within- and between-group interference in a hierarchical network resulted in biased blip parameter estimates and, subsequently, lower optimal treatment rates, particularly as the number of subgroups (and therefore the potential for between-group neighbours in addition to within-group neighbours) increases. Our simulations and data analysis relied on the key assumption that every ego was randomly assigned to at least one alter, either from the same or a different group in the same hierarchy. In doing so, we provided contexts where the SUTVA was always violated, as interference was guaranteed to be present in each hierarchical network.

Determining whether interference is likely to be present in a dataset is often context-specific, and thus it is often up to the discretion of the researcher to assume whether the SUTVA is likely to be violated in a given setting. In the PATH analysis by Jiang et al. [32], interference was assumed to occur in each household - that is, their analysis assumed that each randomly-designated ego was assigned to an alter. However, an interesting question arises when interference may be occurring in some, but not all, households. For instance, if only 10% of households in the PATH dataset had individuals whose outcomes were affected by household members' treatments, and all other individuals in the remaining households did not have their outcomes affected by those whom they lived with, we may say that

“interference is occurring among 10% of individuals selected for analysis.” In cases such as this, where interference is only thought to occur between some (but not all) individuals, then the optimal treatment rules estimated from a “standard” dWOLS analysis that does not account for interference may still correctly estimate the optimal treatment rule most of the time. However, if interference were to take place in the majority of households, say 90%, then failing to account for interference in a dWOLS implementation may result in higher bias for the blip parameter estimates and, subsequently, the optimal treatment rates.

Furthermore, determining whether interference is present in a dataset may also be important if the optimal treatment rule depends on the presence of an alter whose treatment status is interacting with one of the ego’s covariates. Consider the following scenario where we simulate a patient-level variable X , a cluster-level variable V for a hierarchy with $J = 3$ groups, the probability of treatment assignment conditional on X and V with random intercepts $b_{0i} \sim N(0, 1)$ and $b_{0j} \sim N(0, 1)$, the total number of potential alters represented by $N_{\text{total},ijk}$, and an optimal treatment rule for an ego that depends on their alter’s treatment assignment:

- $X \sim N(2, 1)$
- $V \sim N(1, 1)$
- $P(A = 1|X = x, V = v) = \text{expit}(-1 + X + 0.5V + 0.25\sin(X)) + b_{0i} + b_{0j}$
- Optimal treatment rule: $-1 + X + N_{\text{total},ijk} - 2N_{\text{total},ijk}V > 0$

Under this set-up, we first consider the case where an ego has no alters and thus $N_{\text{total},ijk} = 0$. In this case, the optimal treatment rule would be “Treat with $A = 1$ if $X > 1$; otherwise, assign $A = 0$.” Given $P(X > 1) = 0.84$, patients would be treated optimally most of the time provided they have no alters and the overall probability of treatment assignment is close to 80%. However, this becomes more complicated once interference is taking place in the dataset. If an ego is assigned one or more alters, then the optimal treatment rule becomes “Treat with $A = 1$ if $X > 1 - N_{\text{total},ijk} + 2N_{\text{total},ijk}V$; otherwise, assign $A = 0$.” If the treatment status of all alters in $N_{\text{total},ijk}$ is 0, then the optimal treatment rule would be the same as the case where an ego has no alters.

However, the probability of an alter’s treatment status being 0 will depend on the overall probability of treatment assignment in the network, assuming each individual is treated independently of one another. The optimal treatment rule will also vary across individuals based on the number of egos with an assigned alter present in the dataset, depending

on who does or does not have an alter’s treatment assignment that must be taken into consideration in the optimal treatment rule. For instance, if the overall probability of treatment assignment (i.e. $P(A = 1|X = x, V = v)$) is high across the network, then even if some egos have alters with non-zero treatment assignments, the optimal treatment rule “Treat if $X > 1$, 0 otherwise” with $P(A = 1|X = x, V = v) = 0.84$ will likely be met for most individuals provided that the proportion of individuals designated as egos with assigned alters is relatively low (i.e. most individuals’ outcomes do not depend on the treatment assignment of alters).

Going back to the previous set-up, suppose 25% of the individuals of analysis in a dataset are assigned an alter and $P(A = 1|X = x, V = v) = 0.84$. Then the optimal treatment rule, for 75% of patients who do not have an alter, is “Treat with $A = 1$ if $X > 1$; otherwise, treat with $A = 0$.” Given that $P(X > 1) = 0.84$ and $P(A = 1|X = x, V = v) = 0.81$, these patients will likely receive the optimal treatment rule. For the remaining 25% of individuals assigned an alter, we must consider three cases: either all alters are treated, some alters are treated, or no alters are treated. Those whose alters are not treated have the same optimal treatment rule as those without an alter, while those with, say, one treated alter have the optimal treatment rule “Treat if $X > 2V$, 0 otherwise.” Then, given $X > 2V$ approximately 50% of the time, we would expect individuals with an assigned alter to be receiving the optimal treatment at a rate lower than those without an alter. However, by varying the overall probability of treatment assignment within a hierarchy, we can indirectly vary the rates of optimal treatment received among egos who have at least one alter.

In the sections that follow, we conduct simulation studies to illustrate how the extent and nature of interference occurring among individuals in a dataset can be analyzed by varying the overall proportion of individuals within a hierarchy who are assigned an alter either from the same or a different subgroup in the hierarchy. By varying the probability of treatment assignment and the proportion of individuals whose outcomes are affected by an alter’s treatment status, we can compare different implementations of dWOLS that fail to account for interference entirely, partially account for interference, or correctly specify the form of interference present. In *Study 5a*, we vary the proportion of individuals in a dataset with assigned alters who may be either from the same or a different subgroup in the hierarchy. In the second set of simulations, *Study 5b*, we ensure that every ego has one assigned alter, but we vary the proportion of egos who are assigned an alter from a different subgroup rather than the same subgroup of the ego. In this way, we vary the extent of between-group interference occurring relative to within-group interference.

5.2 Simulation Studies: Varying the Overall Proportion of Egos Assigned with Alters

5.2.1 *Study 2a*

Using the same hierarchical network setting as in Chapter 4, we consider I different hierarchies (Level 3 of grouping), with J distinct subgroups nested in each hierarchy (Level 2 of grouping) and n individuals (Level 1) in each J -level subgroup for a total network size of $N = I \times J \times n$ individuals. For simplicity, we assume that the J number of subgroups is fixed across each hierarchy and n is the same for each subgroup.

In Chapter 4, a critical design of our simulations was that each of the n_{ego} designated egos for analysis was guaranteed to be assigned to at least one alter, whether it was within the same or different subgroup at Level 2. In this setting, interference was present throughout the dataset, and mis-specification of the exposure mapping to account for a restricted (within-group only) form of interference resulted in lower optimal treatment rates as the number of subgroups J per hierarchy increased. Nonetheless, our methods still accounted for at least some form of interference occurring within the exposure mapping in the blip function. In the simulations that follow, we will determine the extent of interference that must be present in a dataset to warrant use of dWOLS methodology that accounts for interference.

As in Chapter 4, we consider both an i -th level cluster-level covariate $V_i \sim N(1, 1)$ and an individual-level covariate $X_{ijk} \sim N(2, 1)$ in the treatment model for individual k in subgroup j nested within hierarchy i . Treatment status is modeled as $A_{ijk} \sim \text{Bernoulli}(\pi_{ijk})$ where

$$\text{logit}(\pi_{ijk}) = -1 + \alpha_1 X_{ijk} + 0.5V_i + 0.25 \sin(X_{ijk}) + b_{0i} + b_{0ij},$$

with $b_{0i} \sim N(0, \tau_3)$, $b_{0ij} \sim N(0, \tau_2)$ representing random intercepts at the I and J levels of grouping, respectively, and $b_{0i} \perp b_{0ij}$. By assuming independence of b_{0i}, b_{0ij} , it follows that $b = b_{0i} + b_{0ij} \sim N(0, \tau_3 + \tau_2) = N(0, \tau)$. The outcomes are modelled as $Y_{ijk} \sim N(\mu_{ijk}, 1)$, where

$$\mu_{ijk} = X_{\text{ego},ijk} + X_{\text{ego},ijk}^2 + V_i + \xi_0 N_{\text{total},ijk} + \xi_1 N_{\text{total},ijk} V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 N_{\text{total},ijk} + \psi_3 N_{\text{total},ijk} V_i),$$

and we note that $100p\%$ of n_{egos} , $p = 0.20, 0.40, 0.60, 0.80$, have at least one alter (not necessarily treated). By setting α_1 to 0.80, or -0.55 , we can set average probability of

receiving treatment to 75% or 25%, and given that all individuals are treated independently of each other, this in turn can affect whether $N_{\text{total},ijk} > 0$ for egos who are assigned at least one alter.

We set $\psi = (\psi_0, \psi_1, \psi_2, \psi_3) = (-1, 1, 1, -2)$. For individuals without any assigned alters, the optimal treatment rule is

$$A^{opt} = 1 \text{ if } X_{ijk} > 1, 0 \text{ otherwise,}$$

and for individuals with assigned alters, the optimal treatment rule is

$$A^{opt} = 1 \text{ if } X_{ijk} > 1 - N_{\text{total},ijk} + 2N_{\text{total},ijk}V_i, 0 \text{ otherwise.}$$

By fixing the number of subgroups in each hierarchy to $J = 3$ and ensuring that all randomly selected egos whose outcomes are affected by alters' treatment are assigned exactly $J - 1 = 3 - 1 = 2$ alters, for each ego, $N_{\text{total}} = 0, 1$, or 2 , depending on which of the ego's alters received the treatment assignment of $A = 1$. This results in three different optimal treatment rules for egos who are assigned alters:

1. For egos where $N_{\text{total}} = 0$: $A^{opt} = 1$ if $X_{ijk} > 1$, 0 otherwise.
2. For egos where $N_{\text{total}} = 1$: $A^{opt} = 1$ if $X_{ijk} > 2V_i$, 0 otherwise.
3. For egos where $N_{\text{total}} = 2$: $A^{opt} = 1$ if $X_{ijk} > 4V_i - 1$, 0 otherwise.

We consider four methods for this analysis. In Method 1 (ME1), we implement a standard dWOLS approach that does not account for interference in the form of an exposure mapping in the treatment-free or blip components of the outcome model. That is,

$$E[Y|x, v, a; \beta, \psi] = \beta_0 + \beta_1x + \beta_2x^2 + \beta_3v + a(\psi_0 + \psi_1x).$$

In Method 2 (ME2), we consider an exposure mapping that determines whether an ego has at least one treated alter, whether it is from the same or a different subgroup in the hierarchy. Thus, rather than considering the total number of treated alters, the interference structure is represented as an indicator function for the presence of at least one treated alter. That is,

$$E[Y|x, v, a; \beta, \psi] = \beta_0 + \beta_1x + \beta_2x^2 + \beta_3v + \xi_0I(N_{\text{total}} > 1) + \xi_1I(N_{\text{total}} > 1)v + a[\psi_0 + \psi_1x + \psi_2I(N_{\text{total}} > 1) + \psi_3I(N_{\text{total}} > 1)v].$$

In Method 3 (ME3), we only account for within-group interference whereby $N_{\text{within},ijk}$ includes treated alters relative to ego k only if they are in the same group j in hierarchy i . That is,

$$E[Y|x, v, a; \beta, \psi] = \beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 v + \xi_0 N_{\text{within}} + \xi_1 N_{\text{within}} v + a(\psi_0 + \psi_1 x + \psi_2 N_{\text{within}} + \psi_3 N_{\text{within}} v).$$

Lastly, in Method 4 (ME4), we correctly account for both within- and between-group interference whereby $N_{\text{total},ijk}$ includes treated alters relative to ego k regardless if they are in the same or different group j in hierarchy i . That is,

$$E[Y|x, v, a; \beta, \psi] = \beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 v + \xi_0 N_{\text{total}} + \xi_1 N_{\text{total}} v + a(\psi_0 + \psi_1 x + \psi_2 N_{\text{total}} + \psi_3 N_{\text{total}} v).$$

To assess the double-robustness property of the different dWOLS methods under mis-specification of either the treatment or treatment-free models, we will then compare the optimal treatment rates, blip parameter estimates, and spillover estimates from these methods under the following two scenarios:

- Scenario 1: The treatment-free model is correctly specified, but the treatment model is mis-specified via omission of the non-linear term in X_{ijk} .
- Scenario 2: The treatment model is correctly specified but the treatment-free model is mis-specified via omission of the non-linear term in X_{ijk} .

Mixed effects model parameters for treatment were estimated using the “glmer” function in the “lme4” package in R. We set $\psi = (-1, 1, 1, -2)$ as our blip parameters and simulate 1000 datasets with $I = 10$, $n = 50$, $n_{\text{egos}} = 50$, $\tau_3 = 1$, and $\tau_2 = 1$, fixing the number of subgroups J in each hierarchy to $J = 3$. By setting $J = 3$, we allow for $N_{\text{total}} = 2$ for each alter and we can assess how well ME2 performs in settings of high (75%) or low (25%) probabilities of treatment assignment, as well as comparing its performance against ME3 and ME4 which directly model within- and between-group interference that may be taking place. Across all four methods, we use the same standard overlap weights, i.e.

$$w_{ijk} = |a_{ijk} - E(A_{ijk}|X_{ijk} = x_{ijk}, V_i = v_i)|,$$

rather than network propensity weights, as our outcome model does not contain alter covariates in X in the treatment-free component and we wish to primarily investigate mis-specification of the exposure mapping on blip and spillover estimates and the overall optimal treatment rates.

5.2.2 Results of One-Stage Settings: Varying the Overall Proportion of Egos Assigned with Alters

The mean optimal treatment rate and EmpSE for each method under different conditions of overall treatment assignment and varying values of p are shown in Tables 5.1-5.4, with blip parameter estimates shown in 5.1-5.4. When the treatment-free model is mis-specified and the overall probability of treatment assignment is relatively low (25%), we find that the mean optimal treatment rates from ME1 and ME3 are similar when the overall proportion of egos assigned to alters is relatively low (20%), while the mean optimal treatment rates from ME2 and ME4 remain similar as the overall proportion of egos assigned to alters increases from 0.20 to 0.80. Even as the proportion of egos assigned with alters, p , increases from 0.20 to 0.80, the optimal treatment rate from ME1 reaches a minimum of 64.1% when $p = 0.80$ compared to the much lower observed optimal treatment rate of 33.5%. Similarly, ME3 performs reasonably well as p increases even by failing to account for treatments of between-group alters, which can likely be attributed to most within- and between-group alters having a low probability of treatment assignment and, therefore, the same treatment status of $A = 0$. Our findings are similar when the overall probability of treatment assignment is relatively low (25%) and the treatment model is mis-specified.

When the treatment-free model is mis-specified and the overall probability of treatment assignment is relatively high (75%), we again find that the mean optimal treatment rates from ME1 and ME3 are similar when the overall proportion of egos assigned to alters is relatively low (20%), while the mean optimal treatment rates from ME2 and ME4 remain similar as the overall proportion of egos assigned to alters increases from 0.20 to 0.80. However, as the proportion of egos assigned with alters, p , increases from 0.20 to 0.80, the optimal treatment rate from ME1 becomes much lower so that only about 60% of individuals are being optimally treated when interference is not accounted for. The optimal treatment rate from ME3 drops below 80% when $p = 0.60$, whereas under the low treatment assignment probability, it dropped below 80% when $p = 0.80$.

5.2.3 Study 2b

In *Study 2b*, we extend the methods in *Study 2a* to a two-stage setting where the ego's outcome is $Y_{ijk} \sim N(\mu_{ijk}, 1)$, where

$$\begin{aligned} \mu_{ijk} = & \log |V_{1,i}| + \log |V_{2,i}| + X_{1,ijk} + X_{2,ijk} + N_{\text{total},ijk} + N_{\text{total},ijk} V_{1,i} + N_{\text{total},ijk} + \\ & N_{\text{total},ijk} V_{2,i} + A_{1,ijk}(\psi_0 + \psi_1 X_{1,ijk} + \psi_2 N_{\text{total},ijk} + N_{\text{total},ijk} V_{1,i}) + A_{2,ijk}(\psi_0 + \psi_1 X_{2,ijk} + \\ & \psi_2 N_{\text{total},ijk} + N_{\text{total},ijk} V_{2,i}), \end{aligned}$$

p	Mean Optimal Treatment Rate				Empirical Standard Error (EmpSE)			
	0.20	0.40	0.60	0.80	0.20	0.40	0.60	0.80
Observed	0.313	0.320	0.327	0.335	0.0380	0.0371	0.0378	0.0407
ME1	0.904	0.815	0.736	0.652	0.0398	0.0724	0.103	0.127
ME2	0.974	0.967	0.960	0.954	0.0156	0.0182	0.0206	0.0246
ME3	0.935	0.881	0.831	0.767	0.0277	0.0482	0.0647	0.0900
ME4	0.977	0.974	0.971	0.968	0.0152	0.0175	0.0187	0.0217


Table 5.1: Mean optimal treatment rates, Treatment Model Mis-specified, $P(A = 1|X = x, V = v) = 25\%$

p	Mean Optimal Treatment Rate				Empirical Standard Error (EmpSE)			
	0.20	0.40	0.60	0.80	0.20	0.40	0.60	0.80
Observed	0.648	0.597	0.550	0.492	0.0325	0.0338	0.0436	0.0608
ME1	0.843	0.722	0.649	0.602	0.0511	0.0744	0.0785	0.0706
ME2	0.970	0.961	0.954	0.946	0.0155	0.0176	0.0188	0.0207
ME3	0.905	0.829	0.777	0.739	0.0347	0.0537	0.0651	0.0552
ME4	0.978	0.976	0.975	0.974	0.0148	0.0151	0.0151	0.0157


Table 5.2: Mean optimal treatment rates, Treatment Model Mis-specified, $P(A = 1|X = x, V = v) = 75\%$

p	Mean Optimal Treatment Rate				Empirical Standard Error (EmpSE)			
	0.20	0.40	0.60	0.80	0.20	0.40	0.60	0.80
Observed	0.313	0.320	0.327	0.335	0.0380	0.0371	0.0378	0.0407
ME1	0.897	0.811	0.730	0.641	0.0431	0.0759	0.113	0.139
ME2	0.955	0.945	0.937	0.931	0.0296	0.0339	0.0356	0.0354
ME3	0.923	0.872	0.822	0.754	0.0336	0.0517	0.0708	0.101
ME4	0.958	0.950	0.945	0.940	0.0290	0.0330	0.0357	0.0343


Table 5.3: Mean optimal treatment rates, Treatment-Free Model Mis-specified, $P(A = 1|X = x, V = v) = 25\%$




(a)



(b)




(c)




(d)


Figure 5.1: Blip parameter estimates when the treatment model is mis-specified under a low probability of treatment assignment, $P(A = 1|X = x, V = v) = 25\%$.




(a)



(b)



(c)



(d)

Figure 5.2: Blip parameter estimates when the treatment model is mis-specified under a high probability of treatment assignment, $P(A = 1|X = x, V = v) = 75\%$.




Figure 5.3: Blip parameter estimates when the treatment free model is mis-specified under a low probability of treatment assignment, $P(A = 1|X = x, V = v) = 25\%$.




Figure 5.4: Blip parameter estimates when the treatment free model is mis-specified under a high probability of treatment assignment, $P(A = 1|X = x, V = v) = 75\%$.

p	Mean Optimal Treatment Rate				Empirical Standard Error (EmpSE)			
	0.20	0.40	0.60	0.80	0.20	0.40	0.60	0.80
Observed	0.648	0.597	0.550	0.492	0.0325	0.0338	0.0436	0.0608
ME1	0.841	0.721	0.648	0.600	0.0523	0.0779	0.0809	0.0730
ME2	0.955	0.945	0.939	0.932	0.0274	0.0280	0.0274	0.0279
ME3	0.899	0.825	0.772	0.735	0.0376	0.0568	0.0694	0.0586
ME4	0.961	0.956	0.956	0.954	0.0273	0.0280	0.0270	0.0264

Table 5.4: Mean optimal treatment rates, Treatment-Free Model Mis-specified, $P(A = 1|X = x, V = v) = 75\%$

with $V_{1,i} \sim N(1, 1)$ and $X_{1,ijk} \sim N(2, 1)$ representing simulated cluster- and individual-level covariates respectively in the first stage, and $V_{2,i} \sim N(1 + 0.5V_{1,i}, 1)$ and $X_{2,ijk} \sim N(X_{1,ijk}, 1)$ in the second stage.

As in *Study 2a*, we fix the number of subgroups in each hierarchy to $J = 3$ to ensure that all $100p\%$ of randomly selected egos whose outcomes are affected by alters' treatment are assigned exactly 2 alters, and we consider $p = 0.20$ so that few egos are assigned alters or $p = 0.80$ whereby the majority of egos are assigned alters. We set α_1 to 0.80 so that the average probability of receiving treatment, conditional on a hierarchy's random effects, is 75%. Figures 5.5-5.8 show the blip and spillover estimates when the treatment-free model is mis-specified through linearity in V .

5.3 Simulation Studies: Varying the Ratio of Egos Assigned Between-Group Alters to Egos Assigned Within-Group Alters

In Chapter 3, we saw that network propensity weights helped improve the optimal treatment rates and improved efficiency in blip/spillover parameter estimates if the alter's treatment-free model was mis-specified, whether the within- or between-group component. This was consistent with what we saw in Cong's paper. In the simulations that follow, we will consider scenarios where the ego is assigned exactly one alter who may be either from the same or different subgroup in the hierarchy. The treatment-free component of the model can be decomposed in the ego's and alter's, and we will vary the proportion of ego-alter assignments where the alter is from a different group. We will then compare network




Figure 5.5: Blip parameter estimates under ME2 in *Study 2b* when 20% of egos are assigned an alter in stages one and two ((a) and (b), respectively) versus when 80% of egos are assigned an alter ((c) and (d), respectively).




Figure 5.6: Blip parameter estimates under ME4 in *Study 2b* when 20% of egos are assigned an alter in stages one and two ((a) and (b), respectively) versus when 80% of egos are assigned an alter ((c) and (d), respectively).




Figure 5.7: Spillover parameter estimates under ME2 in *Study 2b* when 20% of egos are assigned an alter in stages one and two ((a) and (b), respectively) versus when 80% of egos are assigned an alter ((c) and (d), respectively).




Figure 5.8: Spillover parameter estimates under ME4 in *Study 2b* when 20% of egos are assigned an alter in stages one and two ((a) and (b), respectively) versus when 80% of egos are assigned an alter ((c) and (d), respectively).

propensity weights that only account for within-group neighbours versus that account for within- and between-group neighbours.

Simulation studies from Jiang et al. [32] showed an efficiency gain from using network weights instead of individual overlap weights when both the ego's and alter's treatment models were correctly (or incorrectly) specified at the same time, but one of the ego's or alter's treatment-free model components were misspecified. In Chapter 3, we showed through simulation that if the treatment-free model of an ego's outcome model can be decomposed into three components corresponding to the ego, their within-group neighbours, and their between-group neighbours respectively, then network weights provided some efficiency gains and slightly higher optimal treatment rates compared to standard individual overlap weights if one or both of the within- or between-group alter's treatment-free components is mis-specified. In these simulations, we assigned each ego one alter from the same subgroup and one alter from another subgroup so that each ego had one within-group alter and one between-group alter.

We now consider the case where every ego is assigned exactly one alter who may either be from the same (within) group or a different (between) group in the same hierarchy, and the treatment-free model in the ego's outcome model now consists of a non-linear term in the alter's covariate. Such a scenario may arise when the researcher has information recorded for the ego's alter as well as the ego themselves, and it is up to their discretion to determine how weights are constructed for each ego. As in previous simulations, we consider both an i -th level cluster-level covariate $V_i \sim N(1, 1)$ and an individual-level covariate $X_{ijk} \sim N(2, 1)$ in the treatment model for individual k in subgroup j nested within hierarchy i . Treatment status is modelled as $A_{ijk} \sim \text{Bernoulli}(\text{expit}(\pi_{ijk}))$ where

$$\text{logit}(\pi_{ijk}) = -1 + \alpha_1 X_{ijk} + 0.5V_i + 0.25 \sin(X_{ijk}) + b_{0i} + b_{0ij},$$

with $b_{0i} \sim N(0, \tau_3)$, $b_{0ij} \sim N(0, \tau_2)$ representing random intercepts at the I and J levels of grouping, respectively, and $b_{0i} \perp b_{0ij}$. By assuming independence of b_{0i} , b_{0ij} , it follows that $b = b_{0i} + b_{0ij} \sim N(0, \tau_3 + \tau_2) = N(0, \tau)$.

The outcomes are modelled as $Y_{ijk} \sim N(\mu_{ijk}, 1)$, where

$$\mu_{ijk} = X_{\text{ego},ijk} + V_i + \exp(X_{\text{alter},ijk}) + A_{\text{alter},ijk} + A_{\text{alter},ijk}V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 A_{\text{alter},ijk} + \psi_3 A_{\text{alter},ijk}V_i)$$

and we note that $100p\%$ of n_{egos} , $p = 0.20, 0.40, 0.60, 0.80$, have a between-group alter while $100(1 - p)\%$ of n_{egos} have a within-group alter. The optimal treatment rule is

$$A^{opt} = 1 \text{ if } \psi_0 + \psi_1 X_{ijk} + \psi_2 A_{\text{alter},ijk} + \psi_3 A_{\text{alter},ijk} V_i > 0, 0 \text{ otherwise}$$

and we set $\psi = (\psi_0, \psi_1, \psi_2, \psi_3) = (-1, 1, 1, -2)$.

We consider four methods for this analysis. In Method 1 (ME1), the ego's treatment-free model is misspecified through non-linearity in their covariate, and we implement a dWOLS approach with network weights only if the ego has been assigned a within-group alter and standard weights otherwise; that is,

$$E[Y|X, X_{\text{alter}}, A, A_{\text{alter}}] = \sin(X_{\text{ego},ijk}) + V_i + \exp(X_{\text{alter},ijk} + A_{\text{alter},ijk} + A_{\text{alter},ijk} V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 A_{\text{alter},ijk} + \psi_3 A_{\text{alter},ijk} V_i)).$$

and network propensity weights are only constructed for egos assigned a within-group alter; that is,

$$w_{ijk} = |a_{ijk} - E(A_{ijk}|X_{ijk} = x_{ijk}, V_i = v_i)| \times |a_{ljk} - E(A_{ljk}|X_{ljk} = x_{ljk}, V_i = v_i)|,$$

if the ego has a within-group alter, and otherwise,

$$w_{ijk} = |a_{ijk} - E(A_{ijk}|X_{ijk} = x_{ijk}, V_i = v_i)|.$$

In Method 2 (ME2), the alter's treatment-free model is misspecified in their covariate, and we implement a dWOLS approach with network weights only if the ego has been assigned a within-group alter and standard weights otherwise; that is,

$$E[Y|X, X_{\text{alter}}, A, A_{\text{alter}}] = X_{\text{ego},ijk} + V_i + \sin(X_{\text{alter},ijk}) + A_{\text{alter},ijk} + A_{\text{alter},ijk} V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 A_{\text{alter},ijk} + \psi_3 A_{\text{alter},ijk} V_i).$$

and network propensity weights are only constructed for egos assigned a within-group alter; that is,

$$w_{ijk} = |a_{ijk} - E(A_{ijk}|X_{ijk} = x_{ijk}, V_i = v_i)| \times |a_{ljk} - E(A_{ljk}|X_{ljk} = x_{ljk}, V_i = v_i)|,$$

if the ego has a within-group alter, and otherwise,

$$w_{ijk} = |a_{ijk} - E(A_{ijk}|X_{ijk} = x_{ijk}, V_i = v_i)|.$$

In Method 3 (ME3), the ego's treatment-free model is misspecified through non-linearity in their covariate, and we implement a dWOLS approach with network weights that accounts for within- or between-group alters; that is,

$$E[Y|X, X_{\text{alter}}, A, A_{\text{alter}}] = \sin(X_{\text{ego},ijk}) + V_i + \exp(X_{\text{alter},ijk} + A_{\text{alter},ijk} + A_{\text{alter},ijk}V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 A_{\text{alter},ijk} + \psi_3 A_{\text{alter},ijk}V_i)).$$

and network propensity weights are constructed for all egos; that is,

$$w_{ijk} = |a_{ijk} - E(A_{ijk}|X_{ijk} = x_{ijk}, V_i = v_i)| \times |a_{ljk} - E(A_{ljk}|X_{ljk} = x_{ljk}, V_i = v_i)|.$$

In Method 4 (ME4), the alter's treatment-free model is misspecified in their covariate, and we implement a dWOLS approach with network weights that accounts for within- or between-group alters; that is,

$$E[Y|X, X_{\text{alter}}, A, A_{\text{alter}}] = X_{\text{ego},ijk} + V_i + \sin(X_{\text{alter},ijk} + A_{\text{alter},ijk} + A_{\text{alter},ijk}V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 A_{\text{alter},ijk} + \psi_3 A_{\text{alter},ijk}V_i)).$$

and network propensity weights are constructed for all egos; that is,

$$w_{ijk} = |a_{ijk} - E(A_{ijk}|X_{ijk} = x_{ijk}, V_i = v_i)| \times |a_{ljk} - E(A_{ljk}|X_{ljk} = x_{ljk}, V_i = v_i)|.$$

We will then test these methods under the following two scenarios:

- Scenario 1: The treatment model is mis-specified via omission of the non-linear term in X_{ijk} .
- Scenario 2: The treatment model is correctly specified.

Mixed effects model parameters for treatment were estimated using the “glmer” function in the “lme4” package in R. We set $\psi = (-1, 1, 1, -2)$ as our blip parameters and simulate 1000 datasets with $I = 10$, $n = 50$, $n_{\text{egos}} = 50$, $\tau_3 = 1$, and $\tau_2 = 1$, fixing the number of subgroups J in each hierarchy to $J = 2$. As in *Study 2b*, we set α_1 to 0.80 so that the average probability of receiving treatment, conditional on a hierarchy's random effects, is 75%.

5.3.1 Results of One-Stage Settings: Varying the Ratio of Egos Assigned Between-Group Alters to Egos Assigned Within-Group Alters

Our findings are summarized in Figures 5.9-5.11. As the proportion of between-group alters increases relative to the proportion of within-group alters assigned among egos, ME1 and ME3 perform well in terms of the mean optimal treatment rate and remain relatively stable. This suggests that when the ego’s component of the treatment-free model is mis-specified, failing to account for between-group alters in construction of the network propensity weights in ME1 still performs nearly as well as ME3. However, when the alter’s component of the treatment free model is mis-specified in ME2 and ME4, the mean optimal treatment rates tend to be lower with high variability, as further evidenced by the high variability in blip parameter estimates for ME2 and ME4. Our findings suggest that mis-specification of the network weights to only account for within-group alters (ME2) performs slightly worse in terms of mean optimal treatment rates than if between-group alters are also accounted for in the network weights (ME4) as the proportion of between-group alters to within-group alters increases.

5.4 Discussion

In this chapter, we explored a variety of simulations comparing different implementations of dWOLS when the overall proportion of individuals for analysis with assigned alters varies. Previous work on DTRs with interference has been done in settings where all individuals for whom the optimal DTR is being estimated has been assigned at least one alter, but it is possible that the outcomes of only some, but not all, individuals for analysis are influenced by treatments of alters. Determining whether interference is taking place in an observational dataset is not straightforward, as violations of the SUTVA will often be specific to the nature of the dataset at hand. However, we show that even when interference is present in an observational dataset unbeknownst to the researcher, mis-specification of the exposure mapping in a dWOLS implementation may still yield relatively high mean optimal treatment rates when the (1) the overall probability of treatment assignment in a hierarchy is relatively low or (2) the proportion of individuals assigned at least one alter is roughly less than 50%. Furthermore, even if the analyst is not sure whether an ego’s outcome is affected by alters from the same group or a different group, an exposure mapping that specifies the presence of at least one treated alter irrespective of their group can still yield high mean optimal treatment rates.







Figure 5.9: Comparison of mean optimal treatment rates when the treatment model is mis-specified (a) versus correctly specified (b).




(a)



(b)




(c)




(d)


Figure 5.10: Blip parameter estimates when the treatment model is mis-specified.




(a)



(b)



(c)



(d)

Figure 5.11: Blip parameter estimates when the treatment model is correctly specified.

Future work may investigate scenarios when the “extent” of interference present (i.e. the proportion of individuals for analysis with assigned alters) changes between stages. For instance, at the first stage of treatment, only 20% of individuals for analysis may be assigned an alter, but by stage two, this increases to 40%. Alternatively, the number of alters whose treatment affects an ego’s outcome may increase over stages. Investigating more complex situations such as these have important applications for infectious disease modelling and pandemic preparedness, such as when determining vaccination strategies for individuals in household and community settings at different stages of a pandemic or as an individual’s social network changes over time.

Chapter 6

Conclusion and Future Work

Interference in social and spatial networks poses many challenges for DTR estimation, and development of modified DTR methods that account for this phenomenon is a relatively new area of research. While much of the causal inference literature to date has assumed partial interference in settings that violate the SUTVA, there is a gap in exploring scenarios where both within- and between-group interference may be present in an observational dataset. Throughout this thesis, we contributed to these areas of research by providing a framework under which interference may occur within- and between-groups in hierarchical datasets, which was motivated by an analysis of post-ED treatment strategies for management of OUD, and by building on previous work by Jiang et al. [32] to demonstrate how dWOLS can be modified to account for both forms of interference in hierarchical datasets through a variety of simulation studies.

In Chapter 3, we introduced the opioid epidemic in Canada as a motivating context for application of DTR methodology to reduce burden on the healthcare system while improving patients' engagement with treatment. This topic was inspired by collaboration with a provincial health authority, and while the data from this project are not currently available due to a delayed ethics clearance, we used demographic summaries provided by Moe et al. [46] from the Provincial Overdose Cohort in BC to simulate realistic data for this area of study. Using data simulated under several key assumptions, we demonstrated how dWOLS can estimate sequences of post-ED dispositions with the goal of minimizing the overall number of physician visits needed following ED visits. Despite several important limitations with the data, such as limited patient-specific variables and lack of information on patients' previous treatment history, we nonetheless provide an extension of dWOLS to health system service utilization where the outcome of interest reflects a potential measure of healthcare system burden. This chapter also introduces the real-world context to which

the methodology outlined in Chapters 4 and 5 can be applied under several assumptions on the structure of interference that may be present.

In Chapter 4, we extended previous work by Jiang et al. [32] to show how dWOLS can be modified to account for interference in hierarchical networks where both within- and between-group interference may be taking place. We first demonstrated how failure to account for between-group interference in hierarchical networks through a mis-specified exposure mapping results in lower mean rates of predicted optimal treatment as the number of subgroups nested within a hierarchy - and therefore the potential for between-group interference to occur - increases. By introducing outcome models where the treatment-free component can be decomposed into three parts each consisting of covariates for the ego, a within-group alter, and a between-group alter, we showed how network propensity weights can offer some gains in efficiency compared to standard overlap weights depending on which components are simultaneously mis-specified or correctly specified. We concluded this chapter with an extension of the work in Chapter 3 in which we assumed that interference may be present in the simulated opioid overdose data and showed how estimated treatment rules differ depending on whether only within-group interference or both within- and between-group interference can occur within and between different facilities.

Lastly, Chapter 5 explored a variety of simulation studies demonstrating the extent to which interference may be present in a dataset to warrant use of modified dWOLS methodology depending on the overall probability of treatment within a hierarchy and the number of individuals whose outcome is affected by the treatments of their neighbours. We also considered scenarios where all egos are assigned at least one alter, but the proportion of within-group alters varies from the proportion of between-group alters. These simulation studies demonstrate the extent to which within- and between-group interference may be taking place in a dataset for mis-specification of the exposure mapping or network propensity weights to result in lower mean optimal treatment rates or greater variability in blip and spillover parameter estimates. An important extension of this work lies assessing the extent to which interference may be occurring in a dataset through specification of a model that predicts whether the outcome of an individual of interest may be affected by the treatments received by an alter within their same group or from a different group. For instance, if the probability of an individual within a facility interacting with another individual from a separate facility is relatively low, then it may be unnecessary to account for both within- and between-group interference when specifying the exposure mapping, and a “simpler” dWOLS analysis is likely to still perform well in terms of the mean optimal treatment rate.

Our findings motivate much future work in the areas of DTR estimation for post-ED service utilization and DTR estimation in hierarchical networks. Extensions to our simula-

tions in Chapter 4 can consider more complex treatment models that incorporate random slopes in addition to random intercepts to introduce group-specific variation in effects of a predictor variable across groups, as well as scenarios where the ego and within-group alters share the same treatment model but a between-group alter’s treatment model is different. Although our exposure mapping was defined as the total number of treated alters relative to an ego and ego-alter assignments were made at random within a hierarchy, it would be interesting to model ego-alter assignments based on the risk levels for the individuals, strengths of social network ties, or other information that would suggest the ego-alter relationship is partially known. Our work in Chapter 5 can be extended to incorporate information where the probability of each individual in a hierarchical network being assigned a within- or between-group alter can be estimated rather than pre-determined. For instance, two individuals who receive treatment in the same facility may be more likely to interact with each other if they share similar treatment history, the same living arrangement through shared housing, or if they both attend the same support group for post-ED visit behavioural counseling. DTR estimation under interference networks remains a challenging area of study, but our work contributes to a greater understanding of its applicability across a variety of settings.

References

- [1] Peter M. Aronow and Cyrus Samii. Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics*, 11(4):1912 – 1947, 2017.
- [2] Bruno Arpino and Fabrizia Mealli. The specification of the propensity score in multilevel observational studies. *Computational Statistics & Data Analysis*, 55(4):1770–1780, 2011.
- [3] Euan A Ashley. Towards precision medicine. *Nature Reviews Genetics*, 17(9):507–522, 2016.
- [4] Susan Athey, Dean Eckles, and Guido W Imbens. Exact p-values for network interference. *Journal of the American Statistical Association*, 113(521):230–240, 2018.
- [5] Laura B Balzer, Wenjing Zheng, Mark J van der Laan, and Maya L Petersen. A new approach to hierarchical data analysis: targeted maximum likelihood estimation for the causal effect of a cluster-level exposure. *Statistical Methods in Medical Research*, 28(6):1761–1780, 2019.
- [6] Jonas Bambi, Yudi Santoso, Hanieh Sadri, Ken Moselle, Abraham Rudnick, Stan Robertson, Ernie Chang, Alex Kuo, Joseph Howie, Gracia Yunruo Dong, et al. A methodological approach to extracting patterns of service utilization from a cross-continuum high dimensional healthcare dataset to support care delivery optimization for patients with complex problems. *BioMedInformatics*, 4(2):946–965, 2024.
- [7] Christina Bamia, Ian R. White, and Michael G. Kenward. Some consequences of assuming simple patterns for the treatment effect over time in a linear mixed model. *Statistics in Medicine*, 32(15):2585–2594, 2013.
- [8] Brian G. Barkley, Michael G. Hudgens, John D. Clemens, Mohammad Ali, and Michael E. Emch. Causal inference from observational studies with clustered interfer-

- ence, with application to a cholera vaccine study. *The Annals of Applied Statistics*, 14(3):1432 – 1448, 2020.
- [9] Weiwei Beckerleg and Joel Hudgins. Substance use-related emergency department visits and resource utilization. *Western Journal of Emergency Medicine*, 23(2):166, 2022.
 - [10] Richard Bellman. Functional equations in the theory of dynamic programming–vii. a partial differential equation for the fredholm resolvent. *Proceedings of the American Mathematical Society*, 8(3):435–440, 1957.
 - [11] Carlos Blanco, Tisha RA Wiley, Jacqueline J Lloyd, Marsha F Lopez, and Nora D Volkow. America’s opioid crisis: the need for an integrated public health approach. *Translational Psychiatry*, 10(1):167, 2020.
 - [12] Public Health Agency of Canada. Opioid- and stimulant-related harms in canada: Key findings, Mar 2025.
 - [13] Bibhas Chakraborty. *Statistical methods for dynamic treatment regimes : reinforcement learning, causal inference, and personalized medicine*. Statistics for Biology and Health. Springer, New York, NY, 2013.
 - [14] Laura C Chambers, Benjamin D Hallowell, Elizabeth A Samuels, Mackenzie Daly, Janette Baird, and Francesca L Beaudoin. An evaluation of the association between specific post-overdose care services in emergency departments and subsequent treatment engagement. *JACEP Open*, 4(1):e12877, 2023.
 - [15] Avital Cnaan, Nan M Laird, and Peter Slasor. Using the general linear mixed model to analyse unbalanced repeated measures and longitudinal data. *Statistics in Medicine*, 16(20):2349–2380, 1997.
 - [16] Francis S Collins and Harold Varmus. A new initiative on precision medicine. *New England Journal of Medicine*, 372(9):793–795, 2015.
 - [17] D. R. Cox. *Planning of experiments*. Wiley publication in applied statistics. Wiley, New York, 1958.
 - [18] Taylor Fleming, Jade Boyd, Koharu Loulou Chayama, Kelly R. Knight, and Ryan McNeil. Using alone at home: What’s missing in housing-based responses to the overdose crisis? *Harm Reduction Journal*, 21(1):24–7, 2024.

- [19] Laura Forastiere, Edoardo M Airoidi, and Fabrizia Mealli. Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, 116(534):901–918, 2021.
- [20] Laura Forastiere, Fabrizia Mealli, Albert Wu, and Edoardo Airoidi. Estimating causal effects under interference using bayesian generalized propensity scores. *arXiv preprint arXiv:1807.11038*, 2018.
- [21] Andrzej Gałeczki and Tomasz Burzykowski. Linear mixed-effects model. In *Linear Mixed-Effects Models Using R*, Springer Texts in Statistics, pages 245–273. Springer New York, New York, NY, 2013.
- [22] Tony P George, Lauren Welsh, Susan L Franchuk, and Franco J Vaccarino. Why integrating medications and psychosocial interventions is important to successfully address the opioid crisis in canada. *The Canadian Journal of Psychiatry*, 67(3):176–178, 2022.
- [23] A. Giffin, B. J. Reich, S. Yang, and A. G. Rappold. Generalized propensity score approach to causal inference with spatial interference. *Biometrics*, 79(3):2220–2231, 2023.
- [24] M Elizabeth Halloran and Claudio J Struchiner. Study designs for dependent happenings. *Epidemiology*, 2(5):331–338, 1991.
- [25] M. Elizabeth Halloran and Claudio J. Struchiner. Causal inference in infectious diseases. *Epidemiology (Cambridge, Mass.)*, 6(2):142–151, 1995.
- [26] Guanglei Hong and Stephen W Raudenbush. Evaluating kindergarten retention policy: A case study of causal inference for multilevel observational data. *Journal of the American Statistical Association*, 101(475):901–910, 2006.
- [27] Wallace J Hopp, Jun Li, and Guihua Wang. Big data and the precision medicine revolution. *Production and Operations Management*, 27(9):1647–1664, 2018.
- [28] Michael G Hudgens and M Elizabeth Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842, 2008.
- [29] Guido W Imbens and Donald B Rubin. Causal inference for statistics, social, and biomedical sciences: An introduction. In *Cambridge and New York: Cambridge University Press, 2015, pp. xix, 625*, pages xix–xix. 2015.

- [30] Francisco J Diaz, Hung-Wen Yeh, and Jose de Leon. Role of statistical random-effects linear models in personalized medicine. *Current Pharmacogenomics and Personalized Medicine (Formerly Current Pharmacogenomics)*, 10(1):22–32, 2012.
- [31] Michael J Murray. The canadian triage and acuity scale: A canadian perspective on emergency department triage. *Emergency Medicine*, 15(1):6–10, 2003.
- [32] Cong Jiang, Michael P Wallace, and Mary E Thompson. Dynamic treatment regimes with interference. *Canadian Journal of Statistics*, 51(2):469–502, 2023.
- [33] Janusz Kaczorowski, Jaunathan Bilodeau, Aaron M Orkin, Kathryn Dong, Raoul Daoust, and Andrew Kestler. Emergency department–initiated interventions for patients with opioid use disorder: a systematic review. *Academic Emergency Medicine*, 27(11):1173–1182, 2020.
- [34] Vishesh Karwa and Edoardo M Airoidi. A systematic investigation of classical causal inference strategies under mis-specification due to network interference. *arXiv preprint arXiv:1810.08259*, 2018.
- [35] Sina Khanmohammadi, Naiier Adibeig, and Samaneh Shanehbandy. An improved overlapping k-means clustering method for medical applications. *Expert Systems with Applications*, 67:12–18, 2017.
- [36] Michael S Kramer, Beverley Chalmers, Ellen D Hodnett, Zinaida Sevkovskaya, Irina Dzikovich, Stanley Shapiro, Jean-Paul Collet, Irina Vanilovich, Irina Mezen, Thierry Ducruet, et al. Promotion of breastfeeding intervention trial (probit): a randomized trial in the republic of belarus. *JAMA*, 285(4):413–420, 2001.
- [37] Navin Kumar, William Oles, Benjamin A Howell, Kamila Janmohamed, Selena T Lee, Melissa C Funaro, Patrick G O’Connor, and Marcus Alexander. The role of social network support in treatment outcomes for medication for opioid use disorder: A systematic review. *Journal of Substance Abuse Treatment*, 127:108367, 2021.
- [38] Fan Li, Alan M. Zaslavsky, and Mary Beth Landrum. Propensity score weighting with multilevel data. *Statistics in Medicine*, 32(19):3373–3387, 2013.
- [39] David Liang, Animesh Kumar Paul, Daniala L Weir, Vera HM Deneer, Russell Greiner, Arno Siebes, and Helga Gardarsdottir. Methods in dynamic treatment regimens using observational healthcare data: A systematic review. *Computer Methods and Programs in Biomedicine*, 263:108658, 2025.

- [40] Belzak Lisa and Halverson Jessica. Evidence synthesis-the opioid crisis in canada: a national perspective. *Health Promotion and Chronic Disease Prevention in Canada: Research, Policy and Practice*, 38(6):224, 2018.
- [41] Lan Liu, Michael G Hudgens, and Sylvia Becker-Dreps. On inverse probability-weighted estimators in the presence of interference. *Biometrika*, 103(4):829–842, 2016.
- [42] Nicholas T Longford and John A Nelder. Statistics versus statistical science in the regulatory process. *Statistics in Medicine*, 18(17-18):2311–2320, 1999.
- [43] Laura MacKinnon, M Eugenia Socías, and Geoff Bardwell. Covid-19 and overdose prevention: Challenges and opportunities for clinical practice in housing settings. *Journal of Substance Abuse Treatment*, 119:108153, 2020.
- [44] Charles F Manski. Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1):S1–S23, 2013.
- [45] Katrina Milaney, Jenna Passi, Lisa Zaretsky, Tong Liu, Claire M O’Gorman, Leslie Hill, and Daniel Dutton. Drug use, homelessness and health: responding to the opioid overdose crisis with housing and harm reduction services. *Harm Reduction Journal*, 18:1–10, 2021.
- [46] Jessica Moe, Mei Chong, Bin Zhao, Frank X Scheuermeyer, Roy Purssell, and Amanda Slaunwhite. Death after emergency department visits for opioid overdose in british columbia: a retrospective cohort analysis. *Canadian Medical Association Open Access Journal*, 9(1):E242–E251, 2021.
- [47] Jessica Moe, Fiona O’Sullivan, Margaret J McGregor, Michael J Schull, Kathryn Dong, Brian R Holroyd, Eric Grafstein, Corinne M Hohl, Johanna Trimble, and Kimberlyn M McGrail. Characteristics of frequent emergency department users in british columbia, canada: a retrospective analysis. *Canadian Medical Association Open Access Journal*, 9(1):E134–E141, 2021.
- [48] Erica EM Moodie, Robert W Platt, and Michael S Kramer. Estimating response-maximized decision rules with applications to breastfeeding. *Journal of the American Statistical Association*, 104(485):155–165, 2009.
- [49] Tim P Morris, Ian R White, and Michael J Crowther. Using simulation studies to evaluate statistical methods. *Statistics in Medicine*, 38(11):2074–2102, 2019.
- [50] Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 65(2):331–355, 2003.

- [51] Office of the Chief Coroner for Ontario/Ontario Forensic Pathology Service, Public Health Ontario, and Ontario Drug Policy Research Network. Opioid toxicity and access to treatment among adolescents and young adults in ontario. Technical report, Office of the Chief Coroner for Ontario/Ontario Forensic Pathology Service, Toronto, Ontario, 2023. Available from Public Health Ontario.
- [52] Shraddha Pai and Gary D Bader. Patient similarity networks for precision medicine. *Journal of Molecular Biology*, 430(18):2924–2938, 2018.
- [53] Georgia Papadogeorgou, Fabrizia Mealli, and Corwin M Zigler. Causal inference with interfering units for cluster and population level treatment allocation programs. *Biometrics*, 75(3):778–787, 2019.
- [54] Samantha Pawer, Fahra Rajabali, Alex Zheng, Ian Pike, Roy Pursell, Atousa Zargaran, and Shelina Babul. Socioeconomic factors and substances involved in poisoning-related emergency department visits in british columbia, canada. *Health Promotion and Chronic Disease Prevention in Canada: Research, Policy and Practice*, 41(7-8):211, 2021.
- [55] Carolina Perez-Heydrich, Michael G Hudgens, M Elizabeth Halloran, John D Clemens, Mohammad Ali, and Michael E Emch. Assessing effects of cholera vaccination in the presence of interference. *Biometrics*, 70(3):731–741, 2014.
- [56] Sophia Rabe-Hesketh, Anders Skrondal, and Andrew Pickles. Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects. *Journal of Econometrics*, 128(2):301–323, 2005.
- [57] Yoshihiko Raita, Carlos A Camargo Jr, Liming Liang, and Kohei Hasegawa. Big data, data science, and causal inference: A primer for clinicians. *Frontiers in Medicine*, 8:678047, 2021.
- [58] Nigel Rice and Andrew Jones. Multilevel models and health economics. *Health Economics*, 6(6):561–575, 1997.
- [59] James M. Robins, Miguel Ángel Hernán, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology (Cambridge, Mass.)*, 11(5):550–560, 2000.
- [60] James M. Robins, Danyu Lin, and Patrick J Heagerty. Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium in Biostatistics*, volume 179 of *Lecture Notes in Statistics*, pages 189–326. Springer New York, United States, 2004.

- [61] Paul R Rosenbaum. From association to causation in observational studies: The role of tests of strongly ignorable treatment assignment. *Journal of the American Statistical Association*, 79(385):41–48, 1984.
- [62] Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- [63] Susanne Rosthøj, Catherine Fullwood, Robin Henderson, and Syd Stewart. Estimation of optimal dynamic anticoagulation regimes from observational data: a regret-based approach. *Statistics in Medicine*, 25(24):4197–4215, 2006.
- [64] Donald B Rubin. Randomization analysis of experimental data: The fisher randomization test comment. *Journal of the American statistical association*, 75(371):591–593, 1980.
- [65] Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- [66] Nitika Sanger, Balpreet Panesar, Tea Rosic, Brittany Dennis, Alessia D’Elia, Alannah Hillmer, Caroul Chawar, Leen Naji, Jacqueline Hudson, M Constantine Samaan, et al. The future of precision medicine in opioid use disorder: inclusion of patient-important outcomes in clinical trials. *Brazilian Journal of Psychiatry*, 43:138–146, 2020.
- [67] Sharon Schwartz, Nicolle M Gatto, and Ulka B Campbell. Extending the sufficient component cause model to describe the stable unit treatment value assumption (sutva). *Epidemiologic Perspectives & Innovations*, 9(1):1–11, 2012.
- [68] Gabrielle Simoneau, Erica EM Moodie, Michael P Wallace, and Robert W Platt. Optimal dynamic treatment regimes with survival endpoints: introducing dwsurv in the r package dtrreg. *Journal of Statistical Computation and Simulation*, 90(16):2991–3008, 2020.
- [69] Michael E Sobel. What do randomized studies of housing mobility demonstrate? causal inference in the face of interference. *Journal of the American Statistical Association*, 101(476):1398–1407, 2006.
- [70] Thorkild IA Sørensen. Which patients may be harmed by good treatments? *The Lancet*, 348(9024):351–352, 1996.
- [71] Robert Stiratelli, Nan Laird, and James H. Ware. Random-effects models for serial observations with binary response. *Biometrics*, 40(4):961–971, 1984.

- [72] Carol Strike and Tara Marie Watson. Losing the uphill battle? emergent harm reduction interventions and barriers during the opioid overdose crisis in canada. *International Journal of Drug Policy*, 71:178–182, 2019.
- [73] Ryan P Strum, Shawn Mondoux, Fabrice I Mowbray, Paul Miller, Andrew Worster, Richard Ferron, and Andrew P Costa. Quantifying the escalating impact of paramedic transported emergency department visits for opioid-related conditions in ontario, canada: A population-based cohort study. *Plos one*, 18(9):e0291194, 2023.
- [74] Elizabeth A. Stuart. Matching methods for causal inference: A review and a look forward. *Statistical Science*, 25(1):1–21, 2010.
- [75] Lin Su, Wenbin Lu, and Rui Song. Modelling and estimation for optimal treatment decision with interference. *Stat*, 8(1):e219, 2019.
- [76] Eric J Tchetgen Tchetgen and Tyler J VanderWeele. On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21(1):55–75, 2012.
- [77] Peter Treitler, Stephen Crystal, Joel Cantor, Sujoy Chakravarty, Anna Kline, Cory Morton, Kristen Gilmore Powell, Suzanne Borys, and Nina A Cooperman. Emergency department peer support program and patient outcomes after opioid overdose. *JAMA Network Open*, 7(3):e243614–e243614, 2024.
- [78] Angela L Vaughan, Trent L Lalonde, and Michael A Jenkins-Guarnieri. Assessing student achievement in large-scale educational programs using hierarchical propensity scores. *Research in Higher Education*, 55:564–580, 2014.
- [79] Geert Verbeke. *Linear mixed models for longitudinal data*. Springer series in statistics. Springer, New York, 1st ed. 2000. edition, 2009.
- [80] Michael P Wallace and Erica EM Moodie. Doubly-robust dynamic treatment regimen estimation via weighted least squares. *Biometrics*, 71(3):636–644, 2015.
- [81] Michael P Wallace, Erica EM Moodie, and David A Stephens. Reward ignorant modeling of dynamic treatment regimes. *Biometrical Journal*, 60(5):991–1002, 2018.
- [82] MP Wallace, Erica EM Moodie, David A Stephens, Gabrielle Simoneau, Shannon T Holloway, Juliana Schulz, and Maintainer Shannon T Holloway. Package ‘dtrreg’. *R package version*, 1, 2016.

- [83] Jack Wilkinson, Kellyn F Arnold, Eleanor J Murray, Maarten van Smeden, Kareem Carr, Rachel Sippy, Marc de Kamps, Andrew Beam, Stefan Konigorski, Christoph Lippert, et al. Time to reality check the promises of machine learning-powered precision medicine. *The Lancet Digital Health*, 2(12):e677–e680, 2020.
- [84] Theresa Winhusen, Christine Wilder, Michael S Lyons, Jeff Theobald, Frankie Kropp, and Daniel Lewis. Evaluation of a personally-tailored opioid overdose prevention education and naloxone distribution intervention to promote harm reduction and treatment readiness in individuals actively using illicit opioids. *Drug and Alcohol Dependence*, 216:108265, 2020.
- [85] Shu Yang. Propensity score weighting for causal inference with clustered data. *Journal of Causal Inference*, 6(2):20170027, 2018.
- [86] Corwin M Zigler, Francesca Dominici, and Yun Wang. Estimating causal effects of air quality regulations using principal stratification for spatially correlated multivariate intermediate outcomes. *Biostatistics*, 13(2):289–302, 2012.
- [87] Corwin M Zigler and Georgia Papadogeorgou. Bipartite causal inference with interference. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 36(1):109, 2021.

APPENDICES

Appendix A

Appendix of Chapter 2

A.1 Linear Mixed Models

The classical linear mixed model with a single level of grouping can be formulated as follows:

$$\mathbf{Y}_i = \mathbf{X}_i\beta + \mathbf{Z}_i\mathbf{b}_i + \epsilon_i,$$

where $\mathbf{Y}_i$ represents an n_i -dimensional response vector for individual i , X_i and Z_i represent $(n_i \times p)$ and $(n_i \times q)$ dimensional matrices of known covariates, β is a p -dimensional vector for the fixed effects, $\mathbf{b}_i$ is a q -dimensional vector for the random effects, and ϵ_i is an n_i -dimensional vector of residual components [79], with $\mathbf{b}_i \perp \epsilon_i$ and

$$\mathbf{b}_i \sim N(0, \mathcal{D}), \epsilon_i \sim N(0, \Sigma_i).$$

Here, $\mathcal{D}$ is a general $(q \times q)$ matrix, while Σ_i is an $(n_i \times n_i)$ matrix that depends on i only through the dimension of n_i [79]. Under these conditions, $\mathbf{Y}_i|\mathbf{b}_i \sim N(X_i\beta + Z_i\mathbf{b}_i, \Sigma_i)$, and the marginal density of $\mathbf{Y}_i$ is given by

$$f(\mathbf{y}_i) = \int f(\mathbf{y}_i|\mathbf{b}_i)f(\mathbf{b}_i)d\mathbf{b}_i,$$

i.e. the density function of an n_i -dimensional vector with mean $X_i\beta$ and covariance matrix $V_i = Z_i\mathcal{D}Z_i' + \Sigma_i$ [79].

In a linear mixed model, $\mathbf{b}_i$ can contain random effects corresponding to random intercepts, random slopes, or both. If one assumes a random intercept-only model, Z_i becomes a vector of ones and $\mathbf{b}_i$ contains only one element, e.g. b_i , with variance σ_1^2 [7]. In this case, our model is structured as follows:

$$\mathbf{Y}_i = \mathbf{X}_i\beta + \mathbf{1}\mathbf{b}_i + \epsilon_i,$$

where $b_i \sim N(0, \sigma_1^2)$ and $\epsilon_i \sim N(0, \Sigma_i)$ [7].

When considering both random intercepts and random slopes in a linear mixed effects model, the matrix Z_i takes on the following form:

$$Z_i = \begin{pmatrix} 1 & z_{i,12} \\ 1 & z_{i,22} \\ \dots & \\ 1 & z_{i,n_i2} \end{pmatrix}$$

and $\mathbf{b}_i = (b_{i1}, b_{i2})^T$ [7]. Our model is then structured as follows:

$$\mathbf{Y}_i = \mathbf{X}_i\beta + \mathbf{Z}_i\mathbf{b}_i + \epsilon_i,$$

where $\mathbf{b}_i \sim N(0, D)$ and $\epsilon_i \sim N(0, \Sigma_i)$, with $D = \begin{pmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{pmatrix}$ [7].

It should be noted that, rather than considering repeated measurements from individuals, linear mixed models can also be applied to hierarchical survey data, hierarchical multivariate outcomes, time series, and so forth [79]. The linear mixed model as outlined above can be extended to multilevel hierarchical data with N first-level groups indexed $i = 1, \dots, N$, each with n_i subgroups containing n_{ij} observations, as modelled below:

$$\mathbf{y}_{ij} = \mathbf{X}_{ij}\beta + \mathbf{Z}_{1,ij}\mathbf{b}_i + \mathbf{Z}_{2,ij}\mathbf{b}_{ij} + \epsilon_{ij},$$

with

$$\mathbf{b}_i \sim N_{q_1}(0, \mathcal{D}_1), \mathbf{b}_{ij} \sim N_{q_2}(0, \mathcal{D}_2), \epsilon_i \sim N_{n_i}(0, \mathcal{R}_{ij})$$

where the random vectors $\mathbf{b}_i$, $\mathbf{b}_{ij}$, and ϵ_i are independent of each other [21]. Additionally, mixed effects models can also be extended to cases where the outcome of interest is binary rather than continuous; letting $p_{ij} = Pr(Y_{ij} = 1)$ and λ_i representing an n_i -dimensional vector of $\text{logit}(p_{ij})$, we have

$$\lambda_i = \mathbf{X}_i\beta + \mathbf{Z}_i\mathbf{b}_i + \epsilon_i,$$

where all other terms are defined as previously [71].

Appendix B

Appendix of Chapter 3

B.1 Application of dWOLS to Simulated Opioid Overdose Data with Changes to Age in Stage 2

In Section 3.2, we simulated data for pWOUD using demographics from Moe et al. [46] under a key assumption that a patient’s age remained the same between Stages 1 and 2. Here, we provide an alternate analysis where some patients’ age may increase by 1 year since their last ED visit.

As in Section 3.2, from each simulated cohort, we exclude all patients who left against medical advice, and patients with CTAS scores of 4 or 5 are merged into one category. We simulate CTAS scores for stage 2 such that the probability of a patient with a CTAS score of 1 or 2 has an 80% probability of being assigned a CTAS score of 2, while the probability of a patient with a CTAS score of 3 or 4 has an 80% probability of having a CTAS score of 3 at stage 2. $\text{Age} \sim N(35, 14)$, and for patients who are discharged, the number of physician visits is represented as $Y \sim N(14, 4)$ while for patients who are admitted to hospital, $Y \sim N(15, 4)$. We assume patient covariates for sex and presence of a psychiatric comorbidity remain the same from stage 1 to stage 2, while approximately 70% of patients have a one-year increase in age since their index visit. We then predicted the probability of being discharged to home or residence at stage 2 ($A_2 = 1$) versus admitted or transferred ($A_2 = 0$) using a logistic model regressing disposition on CTAS score at stage 2, sex, age, and presence of a psychiatric comorbidity. Once these stage 2 covariates were simulated, we excluded patients with a CTAS score of 1 or 2 from analysis. Using these inclusion-exclusion criteria and assumptions, our average sample size across 100 simulated datasets was 2246 (IQR: 2228 – 2268).

CTAS and Disposition Stratified by Stage and Sex					
Stage	Sex	Triage level (CTAS)		Disposition	
		3	4	Discharged	Admitted
1	Male	90.9 %	9.1 %	96.7 %	3.3 %
	Female	91.1 %	8.9 %	94.9 %	5.1 %
2	Male	80.2 %	19.8 %	59.0 %	41.0 %
	Female	79.6 %	20.4 %	59.4 %	40.6 %

Table B.1: Average proportions of CTAS scores and dispositions across males and females at stages 1 and 2 of analysis from 100 simulated datasets.

Across the 100 simulated datasets, the average median age at stage 1 was 36.09 years (IQR: (27.57 – 45.06)), while the average median age at stage 2 was 36.99 years (IQR: (28.47 – 45.96)). Additional demographic summaries based on the simulated datasets are presented in Table B.1. For each dataset, we estimate the probability of a patient being discharged to home or residence ($A = 1$), $\mathbb{E}[A_j|\mathbf{X}_j]$ at each stage j using a logistic regression model, where the probability of being discharged is dependent on sex, age, and presence of a psychiatric comorbidity. Our treatment-free model models the average number of physician visits between ED visits as a linear function of sex, age, CTAS score, and presence of a psychiatric comorbidity, while our blip model only contains age at ED visit as a tailoring variable, i.e. $\gamma(x_{ji}, a_{ji}) = a_{ji}(\psi_{j0} + \psi_{j1}x_{ji})$ where x_{ji} is patient i 's age at ED visit j . Implementation of dWOLS using these models was done using the R package DTRreg [82].

B.2 Results

The average blip parameter estimates and corresponding 95% confidence intervals from 100 simulated datasets are summarized in Table B.2. To use these estimates to form an optimal DTR treatment rule for service class usage, one would identify a patient's age at stage $j = 1, 2$ and identify the sign of the equation $\hat{\psi}_{0j} + \hat{\psi}_{1j} \times \text{age}$ at index visit. At stage 1, the optimal treatment rule would be “Discharge a patient to home or residence if $3.946 + 0.000335 \times \text{Age} > 0$,” and since age is positive, this would translate to discharging

Average Cohort Findings			
Stage (j)	$\hat{\psi}_{0j}$	$\hat{\psi}_{1j}$	Mean j -th Stage Regret
1	3.946 (1.1494 ,6.742)	0.000335 (-0.0720, 0.0726)	0.152
2	-0.807 (-7.951, 6.337)	0.0178 (-0.166, 0.198)	0.563

Table B.2: Mean blip parameter estimates and regrets from 100 simulations.

all patients to home or residence regardless of their age. At stage 2, however, the optimal treatment rule would be “Discharge a patient to home or residence if $-0.807 + 0.0178 \times \text{Age} > 0$,” i.e. a patient would be discharged to home or residence if their age is greater than 45.34 years and admitted to hospital or transferred otherwise.

Using the above estimates, we calculate the proportion of all individuals who undertook the estimated optimal disposition at stages 1 and 2 based on the DTR tailored to age across the 100 simulated datasets. On average, in stage 1, 96.15% of all patients were “correctly” disposed or admitted as recommended by the estimated DTR, while at stage 2, this proportion decreased to 47.21%. The average median age of all patients who received optimal treatment in the first stage was 36.09 (IQR: 27.59, 45.05), and at stage two, the median was 40.42 (IQR: 33.73, 47.61). Among all patients who did not receive optimal treatment, the median age of those who did not receive optimal treatment in the first stage was 37 (IQR: 29, 46), and at the second stage, the median was 35 (IQR: 29, 42). Among patients who received the optimal recommended dispositions at stages 1 and 2, the average median number of physician visits was 13.94, while among those who did not receive the optimal recommended dispositions at stage 1 or stage 2, the the average median number of physician visits was 14.70

Appendix C

Appendix of Chapter 4

C.1 Additional Results from *Study 1a*

In *Study 1a* in 4.2, we considered four different implementations of dWOLS with treatment model

$$\text{logit}(\pi_{ijk}) = -1 + X_{ijk} + 0.5V_i + 0.25 \sin(X_{ijk}) + b_{0i} + b_{0ij}, \quad b_{0i} \sim N(0, \tau_3), b_{0ij} \sim N(0, \tau_2)$$

and outcomes modelled as $Y_{ijk} \sim N(\mu_{ijk}, 1)$ where

$$\mu_{ijk} = X_{ijk} + X_{ijk}^2 + V_i + N_{\text{total},ijk} + N_{\text{total},ijk}V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 N_{\text{total},ijk} + N_{\text{total},ijk}V_i),$$

with $N_{\text{total},ijk}$ denoting the total number of treated alters relative to an ego, from both the same and different groups at the J -th level. In Tables C.1 and C.2, and in Figures C.1-C.4, we present results for Scenarios 3 and 4, where both the treatment and treatment-free models are incorrectly specified (Scenario 3) or correctly specified (Scenario 4).

J	Mean Optimal Treatment Rate			Empirical Standard Error (EmpSE)		
	2	3	4	2	3	4
Observed	0.581	0.458	0.394	0.0652	0.0627	0.0631
ME1	0.793	0.697	0.709	0.0549	0.0572	0.0551
ME2	0.944	0.957	0.965	0.0373	0.0254	0.0209
ME3	0.929	0.929	0.930	0.0507	0.0428	0.0444
ME4	0.929	0.937	0.942	0.0507	0.0380	0.0347

Table C.1: Mean optimal treatment rates for Scenario 3 in *Study 1a*, where both the treatment and treatment-free models are incorrectly specified via omission of the non-linear term in X_{ijk} .

J	Mean Optimal Treatment Rate			Empirical Standard Error (EmpSE)		
	2	3	4	2	3	4
Observed	0.581	0.458	0.394	0.0652	0.0627	0.0631
ME1	0.812	0.705	0.712	0.0508	0.0489	0.0532
ME2	0.970	0.978	0.982	0.0211	0.0122	0.0109
ME3	0.960	0.960	0.958	0.0298	0.0247	0.0273
ME4	0.960	0.965	0.968	0.0298	0.0208	0.0208

Table C.2: Mean optimal treatment rates for Scenario 4 in *Study 1a*, where both the treatment and treatment-free models are correctly specified.




Figure C.1: Boxplots of ψ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 3, where both the treatment and treatment-free models are misspecified via omission of the non-linear term in X_{ijk} .




Figure C.2: Boxplots of ψ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 4, where both the treatment and treatment-free models are correctly specified.




Figure C.3: Boxplots of ξ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 3, where both the treatment and treatment-free models are misspecified via omission of the non-linear term in X_{ijk} .





Figure C.4: Boxplots of ξ parameter estimates as J varies from 2 to 4 across ME2, ME3, and ME4 under Scenario 4, where both the treatment and treatment-free models are correctly specified.

C.2 Additional Results from *Study 1b*


In *Study 1b*, we considered a setting where each ego in a hierarchy with $J = 2$ subgroups is assigned exactly one alter from the same subgroup and one alter from the other subgroup, and the treatment-free component of the ego's outcome model is more complex, containing covariates from each alter. The treatment-free model can be decomposed into the ego's covariates (TF-ego), a within-group alter's covariates (TF-alter-W), and a between-group alter's covariates (TF-alter-B). The outcome model is then modelled as $Y_{ijk} \sim N(\mu_{ijk}, 1)$, where

$$\mu_{ijk} = \sin(X_{ego}) + X_{alter,within}^2 + \exp(X_{alter,between}) + V_i + \xi_0 N_{total,ijk} + \xi_1 N_{total,ijk} V_i + A_{ijk}(\psi_0 + \psi_1 X_{ijk} + \psi_2 N_{total,ijk} + N_{total,ijk} V_i).$$


In Scenario (a), we correctly specify the same treatment model as in *Study 1a*, i.e. $A_{ijk} \sim \text{Bernoulli}(\pi_{ijk})$ where $\text{logit}(\pi_{ijk}) = -1 + X_{ijk} + 0.5V_i + 0.25 \sin(X_{ijk}) + b_{0i} + b_{0ij}$, with $b_{0i} \sim N(0, \tau_3)$, $b_{0ij} \sim N(0, \tau_2)$ representing random intercepts at the I and J levels of grouping, respectively, and $b_{0i} \perp b_{0ij}$. In Scenario (b), the treatment model is misspecified via omission of the non-linear term in X . We provide boxplots of the blip and spillover parameter estimates in Figures C.9-C.12, and in Tables C.3 and C.2, we provide extensions of Tables 4.3 and 4.4 to consider cases where only one component of the treatment-free model is correctly specified.




(a)



(b)




(c)




(d)


Figure C.5: Boxplots of ψ parameter estimates under Scenario (a) in *Study 1b*.




(a)



(b)



(c)



(d)

Figure C.6: Boxplots of ψ parameter estimates under Scenario (b) in *Study 1b*.




Figure C.7: Boxplots of ξ parameter estimates under Scenario (a) in *Study 1b*.

C.3 Additional Results from *Study 1c*

In *Study 1c* in Section 4.2, we provided blip and spillover parameter estimates in a two-stage setting where treatment status at stage 1 is modelled as

$$A_{1,ijk} \sim \text{Bernoulli}(\pi_{1,ijk}),$$

where $\text{logit}(\pi_{1,ijk}) = -1 + X_{1,ijk} + 0.5V_{1,i} + 0.25 \sin(X_{1,ijk}) + b_{0i} + b_{0ij}$, with $b_{0i} \sim N(0, \tau_3)$, $b_{0ij} \sim N(0, \tau_2)$, and treatment status at stage 2 is modelled as

$$A_{2,ijk} \sim \text{Bernoulli}(\pi_{2,ijk}),$$

where $\text{logit}(\pi_{2,ijk}) = -1 + X_{2,ijk} + 0.5V_{2,i} + 0.25 \sin(X_{2,ijk}) + b_{0i} + b_{0ij}$.

The alters' covariates are incorporated in the treatment-free component of the outcome model; specifically, $Y_{ijk} \sim N(\mu_{ijk}, 1)$, where

$$\begin{aligned} \mu_{ijk} = & \log |V_{1,i}| + \log |V_{2,i}| + \exp(X_{1,ijk}) + \exp(X_{2,ijk}) + |\tau(\mathbf{X}_{1,N_{\text{total}}})| + |\tau(\mathbf{X}_{2,N_{\text{total}}})| + \\ & N_{\text{total},ijk} + N_{\text{total},ijk}V_{1,i} + N_{\text{total},ijk} + N_{\text{total},ijk}V_{2,i} + A_{1,ijk}(\psi_0 + \psi_1X_{1,ijk} + \psi_2N_{\text{total},ijk} + \\ & N_{\text{total},ijk}V_{1,i}) + A_{2,ijk}(\psi_0 + \psi_1X_{2,ijk} + \psi_2N_{\text{total},ijk} + N_{\text{total},ijk}V_{2,i}), \end{aligned}$$






Figure C.8: Boxplots of ξ parameter estimates under Scenario (b) in *Study 1b*.


and where $\tau(\mathbf{X}_{1,N_{\text{total}}}) = \sum_{l \in \mathcal{N}_{\text{total}}} X_{1,ijl}/N_{\text{total}}$ and $\tau(\mathbf{X}_{2,N_{\text{total}}}) = \sum_{l \in \mathcal{N}_{\text{total}}} X_{2,ijl}/N_{\text{total}}$ denote the average of individual-level covariate measurements from the alters of ego k at stages 1 and 2, respectively. In Figures C.13 and C.14, we show the blip and spillover parameter estimates from a scenario where the treatment model is correctly specified, but the treatment-free model is mis-specified in terms of the cluster-level variable only by using linear terms in V_i .




(a)



(b)




(c)




(d)


Figure C.9: Boxplots of ψ parameter estimates under Scenario (a) in Table C.3.




(a)



(b)



(c)



(d)

Figure C.10: Boxplots of ψ parameter estimates under Scenario (b) in Table C.3.

Scenario (a): Treatment Model Correct					
Method	Network Weights	Standard Weights	TF-ego	TF-alter-W	TF-alter-B
ME1	✓	✗	✓	✗	✗
ME2	✗	✓	✓	✗	✗
ME3	✓	✗	✗	✓	✗
ME4	✗	✓	✗	✓	✗
ME5	✓	✗	✗	✓	✗
ME6	✗	✓	✗	✓	✗
Scenario (b): Treatment Model Incorrect					
ME1	✓	✗	✓	✗	✗
ME2	✗	✓	✓	✗	✗
ME3	✓	✗	✗	✓	✗
ME4	✗	✓	✗	✓	✗
ME5	✓	✗	✗	✓	✗
ME6	✗	✓	✗	✓	✗

Table C.3: Extension of *Study 1b* where only one of the ego's, alter-W's, or alter-B's treatment-free model is correctly specified.

Scenario (a)		
Method	Mean Optimal Treatment Rate	Empirical Standard Error (EmpSE)
Observed	0.421	0.0764
ME1	0.816	0.139
ME2	0.810	0.122
ME3	0.943	0.0399
ME4	0.959	0.0264
ME5	0.794	0.142
ME6	0.792	0.122
Scenario (b)		
Method	Mean Optimal Treatment Rate	Empirical Standard Error (EmpSE)
Observed	0.421	0.0764
ME1	0.821	0.137
ME2	0.810	0.122
ME3	0.942	0.0393
ME4	0.959	0.0262
ME5	0.798	0.141
ME6	0.791	0.122

Table C.4: Optimal treatment rates under Scenarios (a) and (b) in *Study 1b* when only one component of the treatment-free model (the ego's or one of the alter's) is correctly specified




Figure C.11: Boxplots of ξ parameter estimates under Scenario (a) in Table C.3.




Figure C.12: Boxplots of ξ parameter estimates under Scenario (b) in Table C.3.




Figure C.13: Boxplots of ψ parameter estimates for $J = 2$ in a two-stage setting under Scenario 2 in *Study 1c* where the treatment-free model is mis-specified by using linear terms in V_i only. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.




Figure C.14: Boxplots of ξ parameter estimates for $J = 2$ in a two-stage setting under Scenario 2 in *Study 1c* where the treatment-free model is mis-specified by using linear terms in V_i only. Figures (a) and (b) are results obtained using ME2, while Figures (c) and (d) are results obtained using network absolute value weights in ME3.