Estimating the Causal Effect of Chloroquine Treatment on Mortality in Malaria Patients Using Marginal Structural Models

Acknowledgments

I am incredibly indebted to Professor Miguel Hernán and Dr. Jessica Young, both of whom supported me unconditionally as my thesis advisers. Everything that follows is a result of their dedicated, encouraging, and inspiring mentorship. From guiding me through the foundational elements of the field to helping place the final touches on this report, this journey truly would not have been possible without them.

Additionally, I would like to express my sincerest gratitude to Sarah Anoke for her generous support in reviewing my writing, Victoria Lin and Zilu Zhang for assisting with the technical implementation of my research, Professors Margo Levine and Sean Eddy for encouraging me to continue at important crossroads in this journey, and last but certainly not least, my friends and family for supporting me through it all.

Contents

Chapter 1

Introduction

1.1 Background

Randomized controlled trials (RCTs) have been heralded as the gold-standard in comparative effectiveness research [1]. In the fields of clinical medicine and public health, effectively designed RCTs allow for rigorous, unbiased, and externally valid statistical analyses on the effect of novel treatments or interventions. The earliest RCTs were often short, double-blinded, and tightly-controlled experiments on a small, targeted group of patients [2]. The nature of these RCTs ensured that factors such as confounding and selection bias could be effectively minimized among treatment and intervention groups [2]. In particular, pre-trial randomization ensured that baseline confounding was controlled for, and the short timelines mitigated the risk of patients adhering imperfectly to their assigned protocols or dropping out of the study [2].

Unfortunately, as novel treatments and interventions have become more complex, many of the crucial advantages of RCTs have been lost. Consider malaria, which despite being treatable, remains one of the most prevalent infectious diseases in the modern world [3]. Malaria is caused by a parasite that infects certain types of mosquitoes, which in turn feed on humans. When left untreated, malaria can be fatal or lead to serious

complications, such as hemolytic anemia, acute respiratory distress syndrome, acute kidney failure, and metabolic acidosis [3]. The treatment regimen for patients with malaria involves consistent and frequent dosing of oral chloroquine or artemisinin-based combination therapy (ACT) [4, 5, 6, 7]. Adherence to the assignment treatment regimen, including the frequency of dosing, duration of treatment, and number of pills, is a critical factor in determining patient outcomes [5, 6]. In particular, failure to adhere significantly increases the risk of acquiring anti-malarial resistance, treatment of which is more aggressive and occurs on a longer time-scale [3, 5, 6, 7].

In the case of diseases like malaria, the RCTs that are required to estimate unbiased and externally valid treatment effects often occur on a long time-scale and require patients representative of the population of interest. Unfortunately, such RCTs are commonly prone to post-randomization confounding and selection bias [2]. Incomplete adherence can create bias if unknown post-randomization confounders affect treatment decisions, and selection bias can arise if confounders or competing events lead patients to drop out of a study, miss certain visits, or otherwise become lost to follow up [2]. Absent patterns of perfect adherence and patient follow up, randomized trials become less rigorous and more akin to observational studies. Moreover, the *intention-to-treat* (ITT) comparisons often used in the analyses of such RCTs become problematic [1, 2].

ITT analyses estimate the differential effects of being assigned to a specific treatment group, irrespective of the treatment actually taken [1]. As a result, these analyses do not rely on assumptions regarding patient adherence or dosing patterns. In the analysis of more complex RCTs, such as one for patients with malaria, this oversimplification can be very misleading. Thus, ITT analyses may inaccurately capture the clinical effectiveness of a tested intervention [1]. In such cases, recent research has pointed to the potential of using different analytical methods and approaches to reduce bias and better estimate causal effectiveness. These include adjusting for confounders with methods like inverse probability weighting (IPW) to perform more robust per-protocol analyses [1, 2, 8, 9,

10, 11, 12].

Applying such approaches could account for RCTs in which participants adhere to their assigned treatment intermittently, adhere to treatment based on complex time-varying confounders, become lost to follow up, and beyond. In our research, our goal is to specifically examine the per-protocol effect of a binary time-varying treatment on the long-term risk of an outcome, such as all-cause mortality. Previous approaches have been unsuccessful in estimating this effect due to neglecting the presence of time-varying confounders (i.e. covariates that are themselves risk factors for mortality while also being predictive of a patient's future exposure to his/her assigned intervention) [1, 2, 9, 10, 11, 12]. The bias of these estimates is especially significant when the confounders themselves are influenced by a patient's past treatment history [12].

However, recent work by Robins, *et al.* and others have shown how the application of IPW estimation to appropriately specified marginal structural models (MSMs) do not demonstrate such bias and can in fact accurately estimate and recover causal effects [8, 9, 11, 12, 13]. Such studies constitute the foundation upon which our research is built.

1.2 Objectives

The objectives of our research are as follows: (1) Simulate complex longitudinal data with time-varying confounding from a a malaria RCT with a known causal effect; the patient data should exhibit both imperfect treatment adherence and losses to follow up, (2) Analyze the simulated data with IPW estimation to show that when all time-varying confounders and a patient's entire visit process are known, IPW can successfully recover the known treatment effect, and (3) Showcase situations in which even IPW estimation fails to recover the true causal effect. Ultimately, our hope is that our findings will help inform the design and implementation of future clinical studies to enable accurate and robust causal conclusions.

Chapter 2

Statistical Methodology

2.1 Notation and Key Definitions

The simulated data that we eventually use for our analysis comprise measurements over 60 time points for 2,000 patients who are randomly assigned to one of two groups. Specifically, the data for each patient consist of discretized, longitudinal measurements of the following random variables at time t: A_t, a binary time-varying indicator of exposure to an intervention; $\boldsymbol{L_t}$, a vector of binary and/or continuous time-varying covariates that affect exposure to that intervention; C_t, a binary time-varying indicator of study dropout; Y_t, a binary time-varying indicator of an outcome event; and U, a baseline, time-fixed common cause of the time-varying covariates and outcome event. Ultimately, our objective is to fit a correctly-specified model to our data and estimate its coefficients to uncover the causal effect of A_t on Y_t.

Definition: Let $\bar{A}_t = \{A_0, \ldots, A_t\}$ denote a sequence of actions (i.e. an exposure strategy) from time 0 to time t, and let $P[Y_t(\bar{a}_t)]$ specify the distribution of Y_t when $\bar{A}_t$ is fixed at a predetermined $\bar{a}_t$. A_t has a *causal effect* on Y_t if the distribution of $P[Y_t(\bar{a}_t)]$ depends on $\bar{a}_t$ [10].

In survival analyses, measurements that quantify these differences, such as hazard functions, are used to estimate causal effects [14].

We now turn to the related topic of *counterfactual outcomes*. Let $\bar{a}_t$ denote an arbitrary exposure strategy. Then, we define $Y_{t,\bar{a}_t}$ as the patient's counterfactual outcome history if he/she had followed, possibly contrary to reality, the specified exposure strategy [9, 15]. It is worth noting the connection between the counterfactual outcome $Y_{t,\bar{a}_t}$ as an instance of the distribution $P[Y_t(\bar{a}_t)]$. In the context of hazards, which we discuss in the following section, we analogously define $T_{\bar{a}_t}$ as the counterfactual time at which a patient experiences the outcome event, in the universe where he/she adopted the specified exposure strategy [9].

2.2 Marginal Structural Models

Marginal structural models (MSMs) are models that attempt to parameterize the counterfactual outcome distribution $P[Y_t(\bar{a}_t)]$ [10], as follows:

$$P[Y_t(\bar{a}_t)] = f(\bar{a}_t, \boldsymbol{p_t}) \tag{1}$$

Here, $\boldsymbol{p_t}$ is the vector of additional parameters required to appropriately specify the MSM, given the underlying structure of data. The term marginal indicates that the model is marginal over all time-varying and fixed covariates, and the term structural indicates the interventional structure of the distribution of Y_t, given the fixed sequence of treatment actions $\bar{A}_t = \bar{a}_t$ [9, 10, 11].

In our research, we focus on a specific type of MSM, a marginal structural Cox proportional hazards model [9, 8, 14]. In survival analysis, which broadly encompasses the family of statistical approaches used to study the occurrence or time to occurrence of an event of interest (often, all-cause mortality), the Cox proportional hazards model is particularly useful in simultaneously analyzing the relationships between several

predictors and a specified outcome [14]. In clinical studies, where patient survival is affected by numerous risk factors (or predictors) beyond the assigned intervention, the Cox proportional hazards model is a powerful tool for appropriately modeling such data and quantifying the effect size of each covariate.

Specifically, the Cox proportional hazards model parameterizes a hazard $h(t)$ as a function of time t [9, 14]. *Hazard functions* denote the probability that a patient observed at time t displays the event of interest (i.e. all-cause mortality) at time t [14]. For a set of n covariates $\{x_0, ...x_n\}$, a set of n coefficients $\{\theta_0, ...\theta_n\}$, and a baseline hazard h_0, a hazard function $h(t)$ can be expressed as:

$$h(t) = h_0(t) * e^{(\theta_0 x_1 + ... + \theta_n x_n)} \tag{2}$$

For a sequence of treatment actions $\bar{A}_t$ and a vector of time-independent baseline covariates $\boldsymbol{V}$, $h(t)$ can also be expressed as the conditional mortality rate [9]:

$$h(t) = h(t|\bar{A}_t, \boldsymbol{V}) = h_0(t) * e^{\theta_0 A_t + \boldsymbol{\theta_1 V}} \tag{3}$$

By taking the logarithm of the hazard function, we can estimate the $\{\theta_0, ...\theta_n\}$ coefficients with regression analysis. The transformed coefficients $e^{\theta_0}, ..., e^{\theta_n}$ are known as *hazard ratios* [14]. Hazard ratios greater than one are positively associated with the event of interest, hazard ratios less than one are negatively associated with the event of interest, and hazard ratios equal to one are described as having a null effect on the event of interest. Alternatively, a hazard ratio greater than one increases the hazard, and in the context of survival analysis, increases the probability of the outcome event(i.e. all-cause mortality). An important consequence of the Cox proportional hazards model is that due to cancellation of the baseline hazard term $h_0(t)$, when taking the ratio of two hazards at time t but with differing covariate values, the hazard for the event of interest in one group of patients (e.g. a treatment group) is proportional in a fixed,

time-independent manner to the hazard in another group (e.g. a control group) [14].

Let us return now to our counterfactual outcome $T_{\bar{a}_t}$, which represents the time at which a patient experiences the outcome event had he/she followed, possibly contrary to reality, a specified exposure strategy $\bar{a}_t$. We can now define a Cox proportional hazards model that specifies this counterfactual outcome for any strategy $\bar{a}_t$, as follows:

$$h_{T_{\bar{a}_t}}(t) = h(t|\bar{A}_t, \bar{\boldsymbol{L}}_{\boldsymbol{t}}, \boldsymbol{V}) = h_0(t) * e^{\theta_0 A_t + \boldsymbol{\theta_1 L_t} + \boldsymbol{\theta_2 V}} \tag{4}$$

In our model for $h(t)$ in **Equation 3**, the probability of the event of interest is conditional on only $\bar{A}_t$ and a vector of time-independent baseline covariates $\boldsymbol{V}$. In many clinical studies, however, given the presence of time-varying confounders, such a model would be misspecified [9, 10, 11, 12]. Instead, the model for $h_{T_{\bar{a}_t}}(t)$ in **Equation 4**, which also conditions on a vector of time-dependent covariates $\boldsymbol{L_t}$, would be more appropriate. However, this model as it stands is difficult to fit. Thus, we now turn to the issue of time-dependent confounding and how to account for it in MSMs using inverse probability weighting (also known as inverse probability of treatment weighting) [8, 9, 10].

2.3 Confounding and Inverse Probability Weighting

Recall that for a binary, time-varying intervention A_t having a sequence of treatment decisions $\{A_0, \ldots, A_t\}$, our ultimate goal is to estimate the causal effect of A_t on a binary, time-varying outcome event Y_t. This calculation becomes complicated in the presence of time-varying covariates $\boldsymbol{L_t}$, where for integer-valued $i \in \{0, \ldots, t\}$, A_i affects $\boldsymbol{L_i}$ and $\boldsymbol{L_i}$ and A_i affect future $\boldsymbol{L_{i+1}}$ and A_{i+1} [9, 10, 12]. For the sake of illustration, assume there exists an unmeasured time-fixed baseline common cause of the time-varying covariates and event of interest, which we have previously defined as U. While we can measure A_t, $\boldsymbol{L_t}$, and Y_t at every time point t, U is always unmeasured. Yet, it has a definite effect

Figure 1: Single time point DAG to illustrate the consequences of conditioning on $\boldsymbol{L}$

on both $\boldsymbol{L_t}$ and Y_t. A_t, however, is conditionally independent of U given past values of $\boldsymbol{L_{t-}}$ and A_{t-} [10]. At a single time point t, these relationships can be better represented graphically in a structure called a directed acyclical graph (DAG) (see **Figure 1**).

A DAG is a mathematical object containing a set of vertices and a set of directed edges that connect a subset of those vertices [16]. When used to represent causal structures - in which the lack of a directed edge from a vertex F to a vertex G indicates that F is not a direct cause of G (relative to the full set of variables) - DAGs are causal if all common causes of any pair of variables specified are included [16, 11].

Although it may seem intuitive to recover the causal effect of A_t on Y_t by conditioning on $\boldsymbol{L_t}$, the existence of U leads to a key bias that must not be neglected. In the simplified graph in **Figure 1**, conditioning on $\boldsymbol{L}$ creates an association between A and Y through U. Thus, in a situation where we have a known null effect of A on Y, conditioning on $\boldsymbol{L}$ may lead to the estimation of a biased nonzero effect between A and Y through the creation of a backdoor path between these two variables [10, 15]. Furthermore, if the time-varying covariate $\boldsymbol{L}$ has its own effect on the outcome of interest, estimating the effect of A on Y while conditioning on $\boldsymbol{L}$ will likely lead to an overestimation of the effect [10].

One method to resolve this bias and account for other challenges in RCT analysis (i.e. loss of patient follow up) is inverse probability weighting (IPW) [8, 9, 11, 10]. In IPW regression analyses, Y_t is regressed on A_t with each time point's data weighted by the inverse of the probability of intervention ($P[A_t = 1|\bar{A}_{t-1}, \bar{\boldsymbol{L}}_{\boldsymbol{t}}]$), given the history of treatment decisions ($\bar{A}_{t-1}$) and covariate levels ($\bar{\boldsymbol{L}}_{\boldsymbol{t}}$) up to that time t. In most cases, this probability cannot be empirically observed and must instead be estimated with an *exposure allocation model* [8].

To help build intuition, let us start with the single time point case. In this case, where each patient $j \in \{1, \ldots, N\}$ in an N-patient trial is observed at just one time point ($t = 0$), non-stabilized weights can thus be calculated as follows [8]:

$$nsw_j = \frac{1}{P[A_{j,0} = a_{j,0} | \boldsymbol{L_{j,0}} = \boldsymbol{l_{j,0}}]} \tag{5}$$

Stabilized weights can also be calculated, which have narrower distributions than non-stabilized weights and thus remove the more extreme weight values that can amplify the variance in the resulting estimate. Additionally, such weights enhance the statistical efficiency of IPW regression analysis [9]. In stabilized weights, the numerator measures the unconditional probability of the observed level of A_0, and both the numerator and denominator are conditioned, when applicable, on a set of baseline time-fixed covariates $\boldsymbol{V}$ that are associated with exposure of A_0 [8]:

$$sw_j = \frac{P[A_{j,0} = a_{j,0} | \boldsymbol{V} = \boldsymbol{v}]}{P[A_{j,0} = a_{j,0} | \boldsymbol{L_{j,0}} = \boldsymbol{l_{j,0}}, \boldsymbol{V} = \boldsymbol{v}]} \tag{6}$$

In the analysis of RCT data, exposure allocation model estimates for the numerator and denominator can be calculated by regressing A_t on the constant 1 and regressing A_t on $\boldsymbol{L_t}$, respectively [8]. If applicable, the vector of time-independent baseline covariates $\boldsymbol{V}$ should also be regressed on in the numerator and denominator.

The longitudinal case with time-varying values of A_t and L_t is more complex. To capture the entire history $\bar{A}_t$ and $\bar{\boldsymbol{L}}_{\boldsymbol{t}}$ for each patient $j \in \{1, \ldots, N\}$ from time $i \in \{0, \ldots, t\}$, the weights at each time point are cumulatively multiplied for patient j, as follows [8, 9, 10]:

$$sw_{j,i} = \prod_{k=0}^{i} \frac{P[A_{j,k} = a_{jk} | \bar{A}_{j,k-1} = \bar{a}_{j,k-1}, \boldsymbol{V_j} = \boldsymbol{v_j}]}{P[A_{j,k} = a_{j,k} | \bar{A}_{j,k-1} = \bar{a}_{j,k-1}, \bar{\boldsymbol{L}}_{\boldsymbol{j,k}} = \bar{\boldsymbol{L}}_{\boldsymbol{j,k}}, \boldsymbol{V_j} = \boldsymbol{v_j}]} \tag{7}$$

With these weights calculated, a regression analysis can be performed (e.g. on a weighted

marginal structural Cox proportional hazards model) to estimate the causal effect of A_t on Y_t [8, 9, 10]. Note than in **Equations 6** and **7**, the discrete probability calculations assume a binary intervention (i.e. $A_t = 1$ specifies treatment and $A_t = 0$ specifies no treatment). Additionally, in certain survival analyses where the value of A_t becomes fixed after a certain switch (e.g. once patients start therapy, they remain on therapy), the stabilized weights for the corresponding patient simplify to 1 at every time point after the switch [8, 9]. Similar to the single time point case, estimates for the numerator and denominator of each patient-time weight can be calculated using exposure allocation models that regress the time-varying $A_{j,t}$ on the treatment decision history $\bar{A}_{j,t-1}$, time-varying covariate history $\bar{\boldsymbol{L}}_{\boldsymbol{j,t}}$ (only in the numerator), time-fixed covariate history $\bar{\boldsymbol{V}}_{\boldsymbol{j}}$ (if applicable), and the measurement time t [8]. The number of weights calculated depends on the number of regularly-spaced, discrete time points chosen to fit an MSM representing the underlying data.

When all time-varying covariates are known (i.e. when L_t is fully measured), IPW estimation re-weights the measured data to create a "pseudopopulation" in which $\boldsymbol{L_t}$ has no confounding effect on the relationship between A_t and Y_t, thus removing confounder bias from the calculation of the causal effect of A_t on Y_t [8, 9, 10, 17]. Referring back to the Cox proportional hazard function defined in **Equation 4**, the IPW-estimated effect of A_t on Y_t - denote this as $\hat{\theta_0}$ - is an unbiased estimator for θ_0. Moreover, this IPW estimator $\hat{\theta_0}$ will ultimately converge to a parameter θ_0, where e^{θ_0} is the hazard ratio expressing the causal effect of A_t on Y_t [9].

Thus, in the case of longitudinal data with time-varying confounding and no time-fixed $\boldsymbol{V}$, the weighted Cox proportional hazards function can be written as:

$$h_{T_{\bar{a}_t}}(t) = h(t|\bar{A}(t), \bar{\boldsymbol{L}}(\boldsymbol{t})) = h_0(t) * e^{\theta_0 A(t)} \tag{8}$$

Through inverse probability weighting, we have obtained an MSM that can be fit more easily ($\boldsymbol{L_t}$ has been effectively eliminated as a confounder).

Up until now, we have shown how IPW addresses the issue of time-varying confounding through the creation and analysis of an appropriate pseudopopulation. We will now discuss how IPW can help resolve the issue of informative loss to follow up. We use the term informative in the situation that a patient's covariate values contain information on whether or not they are likely to be lost to follow up ($\boldsymbol{L_t}$ is associated with C_t) [8, 9]. For example, if men are more likely to be lost to follow up than women, gender informs whether or not a given patient will be lost to follow up; loss in this case is informative and the corresponding data should be weighted accordingly.

In other words, if a patient is lost to follow up due to some covariate level, that loss is informative and should be weighted accordingly. Let $C_{j,t}$ be an indicator of patient j being lost to follow up at time t. Then, we can calculate stabilized inverse probability of dropout weights, analogous to the weights in **Equation 7** [8, 9]:

$$sw_{j,i}^{\dagger} = \prod_{k=0}^{i} \frac{P[C_{j,k} = c_{j,k} | \bar{C}_{j,k-1} = \bar{c}_{j,k-1}, \boldsymbol{V_j} = \boldsymbol{v_j}]}{P[C_{j,k} = c_{j,k} | \bar{C}_{j,k-1} = \bar{c}_{j,k-1}, \bar{\boldsymbol{L}}_{\boldsymbol{j,k}} = \bar{\boldsymbol{L}}_{\boldsymbol{j,k}}, \boldsymbol{V_j} = \boldsymbol{v_j}]} \tag{9}$$

With these new weights calculated, a final regression analysis can be performed (i.e. on a marginal structural Cox proportional hazards model weighted by $sw_{j,i} * sw_{j,i}^{\dagger}$) to estimate the unconfounded causal effect of A_t on Y_t [8, 9].

2.4 Malaria Treatment as Motivation

We now turn back to the case of malaria, a major infectious disease that leads to nearly 450,000 deaths annually, the vast majority of which are in children less than 5 years of age [3, 5]. As described in **Chapter 1**, analysis of malaria RCT data is likely to be complicated by the presence of time-varying confounding, imperfect treatment adherence, and losses to follow up. We will now make these complications concrete.

Consider malaria patients who have been infected by the parasite species *Plasmodium falciparum*. Unlike other *Plasmodium* species, *P. falciparum* can rapidly lead to a severe

disease phenotype [3]. *P. falciparum* parasites are also readily capable of acquiring drug resistance, emphasizing the importance of rapid treatment [3, 5, 7]. Standard treatment for patients with *P. falciparum* infection usually includes two weeks of oral chloroquine [3, 6]. In high-burden settings, it is important for patients to be monitored closely to ensure treatment adherence and mitigate the risk of relapse [5].

For the sake of our analysis, we start by defining Z as a time-fixed binary indicator of randomized assignment: $Z = 1$ if a patient is randomized into the "always treat arm" and $Z = 0$ if a patient is randomized into the "never treat arm". Accordingly, we then define A_t as a binary, time-varying treatment indicator for chloroquine therapy at time t: $A_t = 1$ if a patient is on chloroquine therapy and $A_t = 0$ if a patient is not. In our survival analysis of the effect of chloroquine treatment on the mortality of malaria patients, our event of interest Y_t is an indicator of all-cause mortality, while C_t is an indicator of a patient being lost to follow up in our study.

We now consider the following two time-varying confounders of A_t on Y_t: $L_{1,t}$, a continuously-measured covariate representing *P. falciparum* parasitemia levels (as a %) in a patient's blood [4], and $L_{2,t}$, a binary covariate representing the presence ($L_{2,t} = 1$) or absence ($L_{2,t} = 0$) of chloroquine resistance (CR). To illustrate how $L_{1,t}$ and $L_{2,t}$ are time-varying confounders, note that: 1) both $L_{1,t}$ and $L_{2,t}$ are risk factors for Y_t: high levels of parasitemia and the presence of CR are risk factors for death [3, 4, 5, 7]; 2) both $L_{1,t}$ and $L_{2,t}$ predict subsequent exposure to A_t: a patient with high levels of parasitemia but no CR would necessitate continued exposure to chloroquine therapy [5], while a patient who acquires CR might be exposed to an even more aggressive, yet still well tolerated, dose of therapy to overcome his/her drug-resistance [6]; and 3) exposure to A_t predicts both $L_{1,t}$ and $L_{2,t}$: while exposure to chloroquine would hopefully lower levels of parasitemia, heavy use of treatment can lead to the onset of CR [7].

With these variables in mind, we now simulate longitudinal data for malaria patients enrolled in an RCT with a known causal effect and known time-varying confounding.

Chapter 3

Data Generation

3.1 Underlying DAG Structure

By generating simulations of RCT data in R that are reflective of real-world, complex randomized studies, our ultimate goal is to fit an IP-weighted marginal structural Cox proportional hazards model (as in **Equation 8**) to recover a known causal effect. In order to generate the full simulated data, we start by defining an underlying DAG.

Again, in our research, the variables involved in the simulated data include: (1) a continuous, time-fixed baseline common cause (e.g. underlying health status) U of the time-varying covariates and event of interest, (2) a vector $\boldsymbol{L_t} = [L_{1,t}, L_{2,t}]$ of time-varying covariates (e.g. *P. falciparum* parasitemia and CR), (3) a binary indicator A_t of treatment adherence (e.g. chloroquine exposure), (4) a binary indicator C_t of study dropout, and (5) a binary indicator Y_t of an event of interest (e.g. all-cause mortality).

In order to generate data in a way that would allow us to recover a known causal effect, we consider the null case, in which neither time-varying chloroquine treatment A_t nor the time-varying covariates $L_{1,t}$ or $L_{2,t}$ have any effect on mortality (Y_t) [1]. We can imagine this being the case for a subset of patients infected by a non-*P. falciparum* malaria parasite, for whom chloroquine treatment has no therapeutic effect. In the

null case, we start by randomizing each patient in our simulated trial into one of two treatment arms: n_1 patients are randomized into arm 1 ($Z = 1$) and n_2 patients are randomized in to arm 2 ($Z = 0$). Next, we generate the time-fixed baseline unmeasured $U \sim Unif(0, 1)$ for each patient, where U affects $L_{1,t}$, $L_{2,t}$, and Y_t. We then allow $L_{1,t}$ and $L_{2,t}$ to be affected by past exposure patterns of A_{t^-} and levels of $\boldsymbol{L_{t^-}}$, and allow A_t to be affected by past levels of $\boldsymbol{L_t}$ and A_{t^-}. As C_t is an indicator of informative censoring, we also allow C_t to be affected by levels of $\boldsymbol{L_t}$.

It is important to state explicitly that A_t has no dependence on U, for reasons described in **Chapter 2.2**. Further, data generation under this null scenario creates the foundational structure we will need to control for $\boldsymbol{L_t}$ to adjust for confounding caused by a baseline U. We do not end up directly including U in our analysis; rather, we use it as a basis for generating unmeasured confounding in our data. These relationships are summarized in the DAG in **Figure 2**. Note that while data are generated in R for each arm using the same methodology, certain parameters (see **Appendix**) are modified to differentiate the two groups' treatment exposure patterns.

Ultimately the simulated data are a function of the parameters $\vec{\alpha}$, $\vec{\beta_1}$, $\vec{\beta_2}$, $\vec{\eta}$, $\vec{\theta}$, and σ^2, as demonstrated in the **Appendix**.

3.2 Longitudinal Data Generation

For the first part of our analysis in **Chapter 4**, we consider data generated in R from a longitudinal RCT with two arms: arm 1 ($Z = 1$) is defined as the "always treat" group (perfect adherence would consist of A always equaling 1) and arm 2 ($Z = 0$) is defined as the "never treat" group (perfect adherence would consist of A always equaling 0). At baseline, 1000 subjects are randomly assigned to each arm. We then define parameters that specify associations among U, L_t, A_t, Y_t, and C_t, such that by the end of the simulated trial ($t = 60$ time points), the following practical conditions are met: 1) the

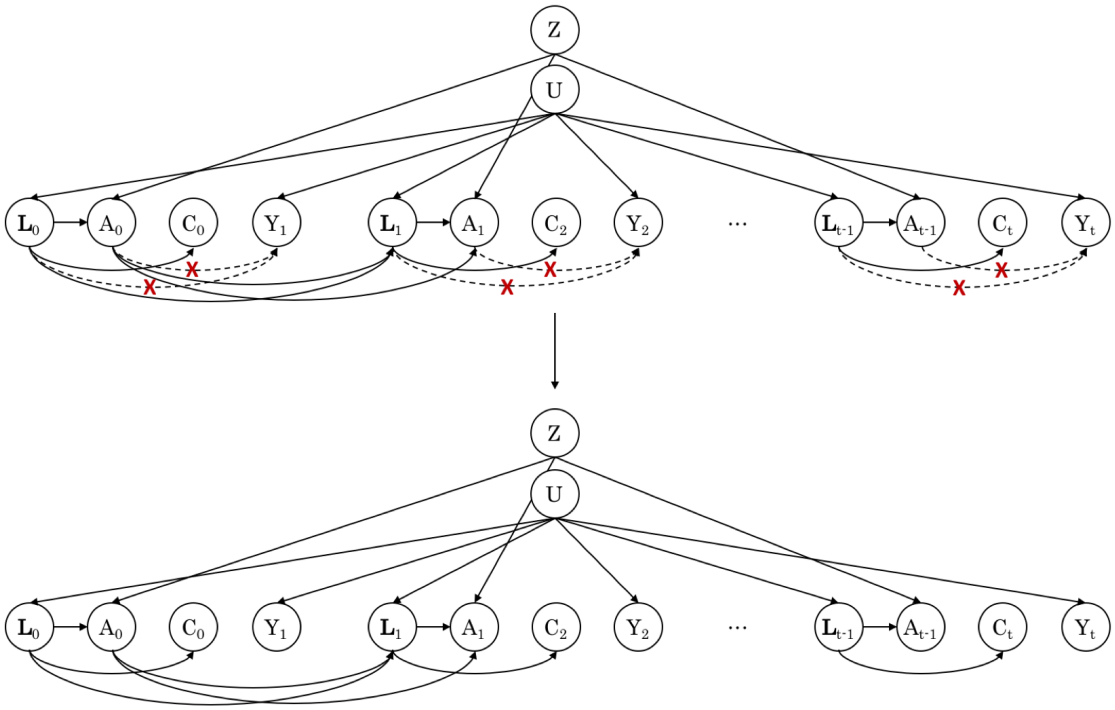

Figure 2: Underlying DAG for data generation under the null

overall proportion of patients who die ($Y_t = 1$ by $t = 60$) is around 10% or less; 2) the overall proportion of patients who dropout of the study by ($C_t = 1$ by $t = 60$) is around 20% or less; 3) the overall proportion of patients who do not perfectly adhere to their assigned exposure strategy by $t = 60$ is around 20-50%. We generate data with the assumption that once a patient is non-adherent, they stay non-adherent at all subsequent time points [8, 9]. For example, if a patient in arm 1 decides to stop taking chloroquine therapy at time $t = 30$ (i.e. $A_{30} = 0$), that patient will remain non-adherent at all subsequent time points (i.e. $A_t = 0$ for $t \geq 30$).

To generate data under the null (no effect of A_t or $\boldsymbol{L_t}$ on Y_t), the parameters specifying the association between A_t and Y_t (θ_2), $L_{1,t}$ and Y_t (θ_3), and $L_{2,t}$ and Y_t (θ_3) are all set to equal 0 (see **Appendix**). Finally, the normally distributed covariate $L_{1,t}$ is generated to reflect clinically meaningful levels of parasitemia in malaria patients [3].

Thus, for each of the 2,000 patients enrolled in our simulated study, values of Z, U, $L_{1,t}$, $L_{2,t}$, A_t, C_t, and Y_t are generated, as shown in **Table 1**. Based on 50 iterations

Z	Patient	Time	U	L_1 (%)	L_2	A	C	Y
1	1	0	0.028	2.5	0	1	0	0
1	1	1	0.028	8.6	1	1	0	0
1	1	2	0.028	2.8	0	1	0	0
...	...	...	...	...	...	...	...	...
0	2000	57	0.125	6.8	1	1	0	0
0	2000	58	0.125	9.0	1	1	0	0
0	2000	59	0.125	5.1	1	1	0	0

Table 1: Sample longitudinal data generated; last patient is treatment non-adherent ($A_t = 1$)

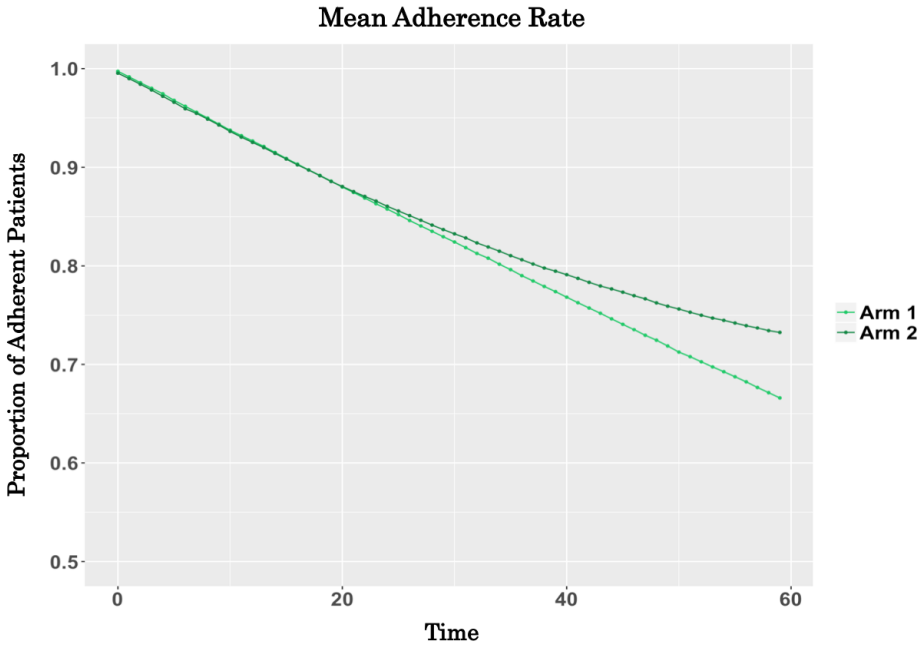

Figure 3: Mean adherence rates by treatment arm in sample longitudinal data

of simulating full 2,000-patient data sets, we observed that in arm 1, by time $t = 60$, the mean treatment adherence rate was 66.6%, mean mortality rate was 8.6%, and mean study dropout rate was 13.2%. As for arm 2, by time $t = 60$, the mean treatment adherence rate was 73.2%, mean mortality rate was 8.5%, and mean study dropout rate was 14.7% (see **Figures 3-5**). Each of these quantities satisfy the conditions above.

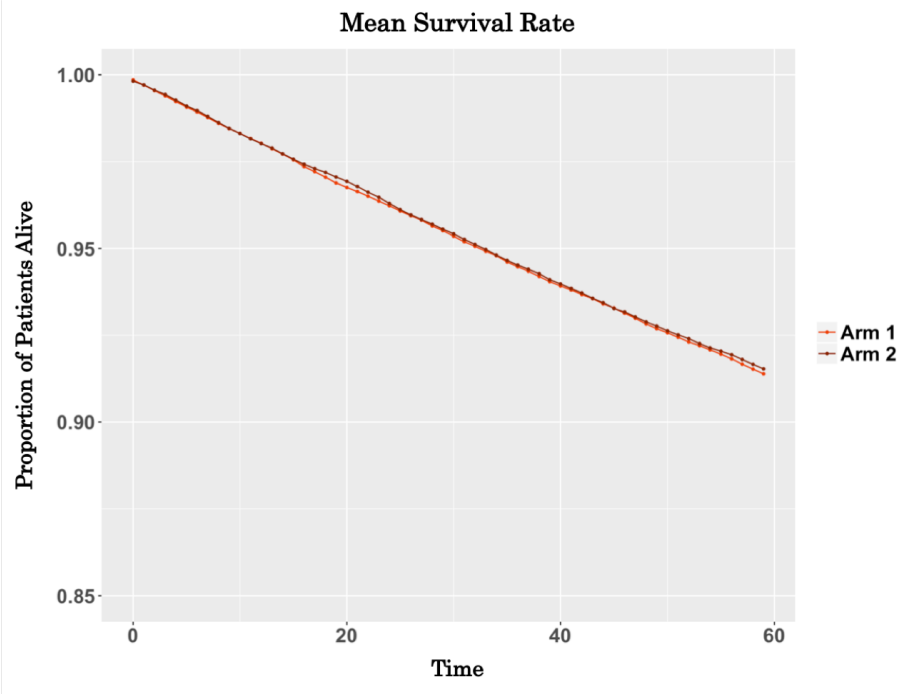

Figure 4: Mean survival rates by treatment arm in sample longitudinal data

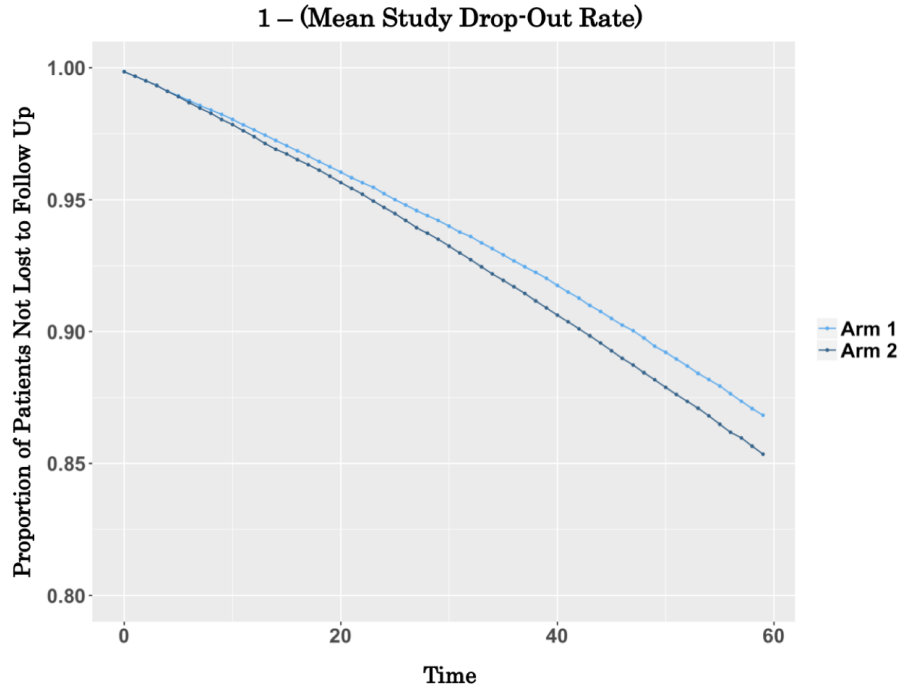

Figure 5: Mean study non-dropout rates by treatment arm in sample longitudinal data

Chapter 4

Results

Having generated data under the null for a longitudinal RCT, we now shift our attention back to fitting our marginal structural Cox proportional hazards model and using IPW estimation to recover the null causal effect.

4.1 Longitudinal Study

Using the `ipw` package in `R` [8], we first calculate subject-specific time-varying weights for each row of uncensored patient-time data. Then, we fit our IP-weighted Cox proportional hazards model on the 2,000-patient longitudinal ($t = 60$ time points per patient) RCT data that we simulated in **Chapter 3**.

In our IPW analysis in the longitudinal case, two forms of censoring were performed after each data set was generated. First, data for non-adherent patients (i.e. patients in Arm 1 neglecting to take chloroquine therapy or patients in Arm 0 taking therapy) were censored prior to IPW estimation. This was done to restrict our causal analysis only to patients who adhered to their assigned treatment strategy (i.e. emphper-protocol analysis) [1]; stabilized weights were then calculated for each patient as specified in **Equation 7**. Second, data for patients who were lost to follow up were weighted by the inverse probability of dropping out, as specified in **Equation 9**. As such non-adherence

is nonrandom and instead dependent on $\boldsymbol{L_t}$, it is important to re-weight the data accordingly [8, 9]. Note that in our analysis, we did not have any time-fixed covariates $\boldsymbol{V}$ on which to condition A_t or C_t.

To estimate the causal effect of A on Y, we performed two analyses on each of the 50 fully simulated data sets: 1) an intention-to-treat (ITT) analysis regressing Y_t on a patient's assigned treatment arm ($Z = 1$ for always treat; or $Z = 0$ for never treat); and 2) a weighted survival analysis regressing Y_t on informatively censored and weighted values of A_t. Stabilized weights in this case were calculated as the product of **Equations 7** and **9**.

Our analyses reveal that our Cox proportional hazards model with IPW estimation is effective at recovering the known null (causal) effect in the underlying data: $\theta_2 = 0.010$, corresponding to a hazard ratio of 1.043 (**Table 2**). In other words, returning to our story of chloroquine treatment for patients with malaria, patients in the "always treat" group (arm 1) have approximately the same risk of death as patients in the "never treat" group (arm 2); indeed, chloroquine has a null effect on mortality. As the stabilized weights in the per-protocol IPW estimate offer a best-fit guess to account for imperfect adherence and losses to follow up, there is an unavoidable degree of uncertainty in the resulting causal effect. The ensuing ITT estimate, however, does not include this uncertainty, as it ignores treatment adherence when regressing Y_t on Z.

Thus, although our IPW result is an important validation of our statistical methodology, it is not particularly beneficial in the null case. After all, a key advantage of traditional ITT analyses is that, irrespective of patients' adherence patterns, they are able to accurately estimate the effect of A_t on Y_t in the null case [1]. As expected, our analyses confirm the success of the ITT regression in estimating the null causal effect: $\theta_2 = 0.009$ (**Table 2**).

Note: We define 95% confidence intervals unconventionally. In particular, the confidence intervals that we report in **Table 2** constitute the 0.025 and 0.975 quantiles

Data Analyzed	Mean IPW Estimate	95% CI
ITT estimator	0.009	(-0.253, 0.312)
IPW estimator	0.010	(-0.277, 0.313)

Table 2: Recovering the null causal effect in our simulated longitudinal case

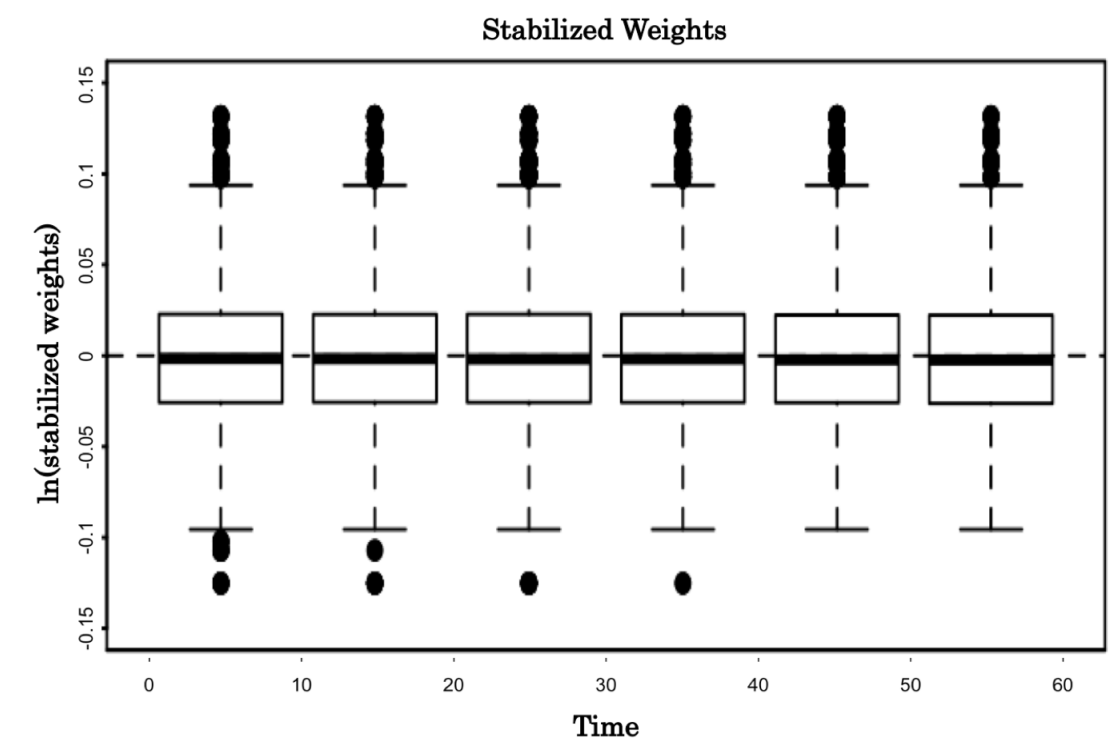

Figure 6: Spread of stabilized weights in IPW estimation to fit longitudinal RCT data

of the distribution of IPW and ITT estimates across all simulation iterations.

Finally, it is illustrative to graphically examine the weights used in the IPW analysis. **Figure 6** displays the spread of stabilized weights at each time point. While the majority of observations appear to be weighted equally (assigned a weight $sw_{j,i}$ of 1), there are certainly observations weighted higher and lower (e.g. data for patients who were informatively censored or non-adherent to their assigned treatment) which contribute to the adjusting of time-varying covariates in the per-protocol IPW estimate.

Our results in this section successfully validate our statistical methodology and data generating mechanism in **Chapters 2** and **3**. The result in the next section, however, constitutes the key finding of our research.

4.2 Incomplete Patient History

Our goal now is to showcase settings in which even the application of IPW estimation with MSMs can lead to biased conclusions. We start by generating a new 2,000 patient data set, with 1,000 patients assigned to arm 1 ($Z = 1$: once treatment initiated, always treat) and 1,000 patients assigned to arm 2 ($Z = 0$: never treat). The data simulated in this example allowed for nonzero values of the parameters θ_2, θ_3, and θ_4 (see **Appendix**), specifying the associations between A_t and Y_t, $L_{1,t}$ and Y_t, and $L_{2,t}$ and Y_t, respectively. Thus, the underlying data being fitted in this section has nonzero causal effect.

In the first phase of analysis, we fit the full simulated data set (once again using the `ipw` package in `R` [8]) with a weighted Cox proportional hazards model. In this analysis, patients' entire history (from time $i = \{0, \ldots, t\}$) of treatment decisions, time-varying covariate levels, and outcome events are known. This is the same approach that we used to recover the true causal effect in the previous section. The resulting mean hazard ratio summarizing the causal effect of A_t on Y_t, after repeating the analysis with 50 simulated data sets, was 1.383 (**Table 3**). In other words, in our story of chloroquine treatment for patients with malaria, patients enrolled in this study who were assigned to be exposed to chloroquine treatment (arm 1) have approximately 1.4 times the risk of death as patients who were assigned to not be exposed to chloroquine treatment (arm 2).

Unfortunately, no RCT can be conducted where every patient's complete history is known. Let us assume that chloroquine treatment in a trial is administered over a 5-day regimen. The above analysis would require a setting in which each patient's full treatment, *P. falciparum* parasitemia, and chloroquine resistance history are recorded every two hours for the 5-day treatment period ($t = 60$ time points). In the real-world, this is highly impractical. What is more likely is that a patient would come in for single check-up each day, where his/her levels of treatment adherence, *P. falciparum*

Data Analyzed	Mean Hazard Ratio Estimate	95% CI
Complete data (IPW estimator)	1.383	(0.941, 1.753)
Subsetted data (IPW estimator)	0.936	(0.716, 1.232)

Table 3: Incomplete patient history leading to bias in estimate of hazard ratio

parasitemia, and chloroquine resistance would be measured and recorded. We thus imitate this approach by shrinking the fully simulated data in each iteration from above, extracting only the data at times $t = 0, 12, 24, 36, 48$ and accounting for patients' mortality and study dropout indicators. We then re-fit a new weighted Cox proportional hazards model. While we hope the results are similar to those in the first phase, the resulting mean hazard ratio summarizing the causal effect of chloroquine treatment (A_t) on mortality (Y_t), after repeating the analysis with the 50 subsetted data sets, is 0.936 (**Table 3**).

Note: Again, the 95% confidence intervals that we report in **Table 3** are unconventional. In particular, they constitute the 0.025 and 0.975 quantiles of the distribution of hazard ratio estimates for the complete and subsetted RCT data across all simulation iterations.

The results are markedly different. If we assume the researchers in the simulated study only had access to the data in the second phase of our analysis, they would likely draw the conclusion that chloroquine treatment slightly reduces the hazard of mortality in malaria patients; conservatively, they would conclude a null effect on survival. However, based on our analysis of the patient's full history, we know the true effect of chloroquine on patients in this study is an increased mortality hazard. Such a bias could have grave implications in clinical settings.

Chapter 5

Conclusion

Overall, our research demonstrates the power of statistical methods like IP-weighting with marginal structural models to ascertain the true causal effect of a time-varying intervention A_t on a time-varying outcome Y_t in the presence of time-dependent confounding, imperfect treatment adherence, and loss to follow up. Such an approach is a significant improvement over traditional ITT analyses, which only consider a patient's assigned treatment arm [1, 2].

However, in cases where experimenters only have access to limited patient data, it is still not possible to obtain an exact analytic solution to a marginal structural model using the IPW methods described in this paper. Rather, as we showed in **Chapter 4**, per-protocol IPW causal effect estimates when only limited patient data are known can offer a misleading conclusion compared to the true causal effect.

The path ahead suggests two alternatives: 1) roll out technologies that enable the continuous monitoring and evaluation of treatment decisions and a set of pre-determined time-varying covariates in all patients over the course of a study or 2) identify settings and conditions under which the bias of causal effects - when only limited patient data are known - is minimized. Although the former approach may seem intractable, mobile health technologies are offering encouraging results, even in high-burden, low-resource

settings [18]. As cell phones - albeit not smart phones - are nearly ubiquitous in such contexts, patient data can be more regularly fielded and recorded virtually.

Nonetheless, our future work aims to tackle the latter approach. Ultimately, the hope is that such information will aid in the design and development of more robust RCTs. In the case of infectious diseases like malaria, designing such trials could be pivotal in streamlining the validation and accelerating the roll-out of promising interventions in overburdened, under-resourced settings.

Appendix

Data Generating Mechanism

In our underlying data generating mechanism, recall that the simulated data are a function of the parameters $\vec{\alpha}$, $\vec{\beta_1}$, $\vec{\beta_2}$, $\vec{\eta}$, $\vec{\theta}$, and σ^2. Here, we outline the explicit mechanism by which the data are generated. Notice how logistic probabilities (inverse logit function probabilities) frequently appear, as logistic regression is used to predict binary outcome variables [19]. We start by randomizing patients into one of two treatment arms and generating data at time $t = 0$:

1. For the $n_1 + n_2$ patients enrolled in the study, n_1 patients are randomized into arm 1 ($Z = 1$) and n_2 patients are randomized in to arm 2 ($Z = 0$); Z is a time-fixed binary indicator of randomized assignment. Depending on the treatment arm assigned, the $\vec{\alpha}$ parameters are modified accordingly.

2. U is generated as a uniformly distributed random variable: $U \sim Unif(0, 1)$

3. $L_{1,0}$ is generated as a normally distributed random variable: $L_{1,0} \sim N(\mu(u), \sigma^2)$ with mean $\mu(u)$ and variance σ^2. At $t = 0$, $\mu(u) = \mathrm{E}[L_{1,0}|U = u, Y_0 = C_0 = 0] = \beta_{1,0} + \beta_{1,1}u$.

4. $L_{2,0}$ is generated as a Bernoulli random variable: $L_{2,0} \sim Bern(p(l_{1,0}, u))$, where $p(l_{1,0}, u) = \Pr[L_{2,0} = 1|L_{1,0} = l_{1,0}, U = u, C_0 = Y_0 = 0] = \frac{e^{\beta_{2,0}+\beta_{2,1}u+\beta_{2,2}l_{1,0}}}{1+e^{\beta_{2,0}+\beta_{2,1}u+\beta_{2,2}l_{1,0}}}$

5. A_0 is generated as a Bernoulli random variable $A_0 \sim Bern(p^{z=1}(l_{1,0}, l_{2,0}))$, where $p^{z=1}(l_{1,0}, l_{2,0}) = \Pr[A_0 = 1|L_{1,0} = l_{1,0}, L_{2,0} = l_{2,0}, Z = 1, C_0 = Y_0 = 0] =$

$\frac{e^{\alpha_0^{z=1}+\alpha_1^{z=1}l_{1,0}+\alpha_2^{z=1}l_{2,0}}}{1+e^{\alpha_0^{z=1}+\alpha_1^{z=1}l_{1,0}+\alpha_2^{z=1}l_{2,0}}}$

6. C_1 is generated as a Bernoulli random variable $C_1 \sim Bern(p(a_0, l_{1,0}, l_{2,0}))$, where $p(a_0, l_{1,0}, l_{2,0}) = \Pr[C_1 = 1|A_0 = a_0, L_{1,0} = l_{1,0}, L_{2,0} = l_{2,0}, C_0 = Y_0 = 0] = \frac{e^{\eta_0+\eta_1 a_0+\eta_2 l_{1,0}+\eta_3 l_{2,0}}}{1+e^{\eta_0+\eta_1 a_0+\eta_2 l_{1,0}+\eta_3 l_{2,0}}}$

7. Y_1 is generated as a Bernoulli random variable $Y_1 \sim Bern(p_y(a_0, l_{1,0}, l_{2,0}, u))$, where $p_y(a_0, l_{1,0}, l_{2,0}, u) = \Pr[Y_1 = 1|A_0 = a_0, L_{1,0} = l_{1,0}, L_{2,0} = l_{2,0}, U = u, C_1 = Y_0 = 0] = \frac{e^{\theta_0+\theta_1 u+\theta_2 a_0+\theta_3 l_{1,0}+\theta_4 l_{2,0}}}{1+e^{\theta_0+\theta_1 u+\theta_2 a_0+\theta_3 l_{1,0}+\theta_4 l_{2,0}}}$

    - When generating data under the null, the parameters specifying the association between A and Y (θ_2), L_1 and Y (θ_3), and L_2 and Y (θ_4) should all be set to equal 0

8. If $C_1 = 1$ or $Y_1 = 1$ or data is only being generated at one time point per patient, stop generating dating for this patient; otherwise, continue generating data for $t = 1, ..., T$ as follows:

9. $L_{1,t}$ is generated as a normally distributed random variable: $L_{1,t} \sim N(\mu_t(\bar{a}_{t-1}, \bar{l}_{1,t-1}, \bar{l}_{2,t-1}, u), \sigma^2)$ with mean $\mu_t(\bar{a}_{t-1}, \bar{l}_{1,t-1}, \bar{l}_{2,t-1}, u)$ and variance σ^2; $\mu_t(\bar{a}_{t-1}, \bar{l}_{1,t-1}, \bar{l}_{2,t-1}, u) = \mathrm{E}[L_{1,t}|\bar{A}_{t-1} = \bar{a}_{t-1}, \bar{L}_{1,t-1} = \bar{l}_{1,t-1}, \bar{L}_{2,t-1} = \bar{l}_{2,t-1}, U = u, Y_t = C_t = 0] = \beta_{1,0}+\beta_{1,1}u+\beta_{1,2}a_{t-1}+\beta_{1,3}cumavg(\bar{a}_{t-2})+\beta_{1,4}cumavg(\bar{l}_{1,t-1})+\beta_{1,5}l_{2,t-1}+\beta_{1,7}t$.

10. $L_{2,t}$ is generated as a Bernoulli random variable: $L_{2,t} \sim Bern(p_t(l_{1,t}, \bar{a}_{t-1}, \bar{l}_{1,t-1}, \bar{l}_{2,t-1}, u))$, where $p_t(l_{1,t}, \bar{a}_{t-1}, \bar{l}_{1,t-1}, \bar{l}_{2,t-1}, u) = \Pr[L_{2,t} = 1|L_{1,t} = l_{1,t}, \bar{A}_{t-1} = \bar{a}_{t-1}, \bar{L}_{1,t-1} = \bar{l}_{1,t-1}, \bar{L}_{2,t-1} = \bar{l}_{2,t-1}, U = u, Y_t = C_t = 0] =$

$$\frac{e^{\beta_{2,0}+\beta_{2,1}u+\beta_{2,2}cumavg(\bar{l}_{1,t})+\beta_{2,3}l_{2,t-1}+\beta_{2,5}a_{t-1}+\beta_{2,6}cumavg(\bar{a}_{t-2})+\beta_{2,7}t}}{1+e^{\beta_{2,0}+\beta_{2,1}u+\beta_{2,2}cumavg(\bar{l}_{1,t})+\beta_{2,3}l_{2,t-1}+\beta_{2,5}a_{t-1}+\beta_{2,6}cumavg(\bar{a}_{t-2})+\beta_{2,7}t}}$$

11. A_t is generated as a Bernoulli random variable $A_t \sim Bern(p_t^{z=1}(\bar{l}_{1,t}, \bar{l}_{2,t}, \bar{a}_{t-1}))$,

where $p_t^{z=1}(\bar{l}_{1,t}, \bar{l}_{2,t}, \bar{a}_{t-1}) = \Pr[A_t = 1 | \bar{L}_{1,t} = l_{1,t}, \bar{L}_{2,t} = l_{2,t}, \bar{A}_{t-1} = \bar{a}_{t-1}, Z = 1, C_t = Y_t = 0] = \frac{e^{\alpha_0^{z=1} + \alpha_1^{z=1} cumavg(\bar{l}_{1,t}) + \alpha_2^{z=1} l_{2,t} + \alpha_3^{z=1} t}}{1 + e^{\alpha_0^{z=1} + \alpha_1^{z=1} cumavg(\bar{l}_{1,t}) + \alpha_2^{z=1} l_{2,t} + \alpha_3^{z=1} t}}$

- However, once a patient is non-adherent to their assigned intervention, they remain non-adherent at all subsequent time points. Thus, if a patient in arm 1 goes off chloroquine treatment at time t, $A(t)$ and all subsequent values of A will equal 0. Similarly, if a patient in arm 2 starts chloroquine treatment at time t, $A(t)$ and all subsequent values of A will equal 1.

12. C_{t+1} is generated as a Bernoulli random variable $C_{t+1} \sim Bern(p_t(\bar{a}_t, \bar{l}_{1,t}, \bar{l}_{2,t}))$, where $p_t(\bar{a}_t, \bar{l}_{1,t}, \bar{l}_{2,t}) = \Pr[C_{t+1} = 1 | \bar{A}_t = \bar{a}_t, \bar{L}_{1,t} = \bar{l}_{1,t}, \bar{L}_{2,t} = \bar{l}_{2,t}, C_t = Y_t = 0] = \frac{e^{\eta_0 + \eta_1 cumavg(\bar{a}_t) + \eta_2 cumavg(\bar{l}_{1,t}) + \eta_3 l_{2,t} + \eta_5 t}}{1 + e^{\eta_0 + \eta_1 cumavg(\bar{a}_t) + \eta_2 cumavg(\bar{l}_{1,t}) + \eta_3 l_{2,t} + \eta_5 t}}$

13. Y_{t+1} is generated as a Bernoulli random variable $Y_{t+1} \sim Bern(p_{yt}(\bar{a}_t, \bar{l}_{1,t}, \bar{l}_{2,t}, u))$, where $p_{yt}(\bar{a}_t, \bar{l}_{1,t}, \bar{l}_{2,t}, u) = \Pr[Y_{t+1} = 1 | \bar{A}_t = \bar{a}_t, \bar{L}_{1,t} = \bar{l}_{1,t}, \bar{L}_{2,t} = \bar{l}_{2,t}, U = u, C_{t+1} = Y_t = 0] = \frac{e^{\theta_0 + \theta_1 u + \theta_2 cumavg(\bar{a}_t) + \theta_3 cumavg(\bar{l}_{1,t}) + \theta_4 l_{2,t} + \theta_6 t}}{1 + e^{\theta_0 + \theta_1 u + \theta_2 cumavg(\bar{a}_t) + \theta_3 cumavg(\bar{l}_{1,t}) + \theta_4 l_{2,t} + \theta_6 t}}$

    - When generating data under the null, the parameters specifying the association between A and Y (θ_2), L_1 and Y (θ_3), and L_2 and Y (θ_4) should all be set to equal 0

14. If $C_{t+1} = 1$ or $Y_{t+1} = 1$ or data is only being generated for t time points per patient, stop generating dating for this patient; otherwise, continue generating data according to steps 8-12.

References

[1] **Hernán MA, Hernández-Díaz S**. "Beyond the intention-to-treat in comparative effectiveness research". In: *Clinical Trials* 9 (2012), pp. 48–55.

[2] **Hernán MA, Hernández-Díaz S, Robins JM**. "Randomized trials analyzed as observational studies". In: *Annals of Internal Medicine* 159.8 (2013), pp. 560–563.

[3] **CDC**. *Treatment of malaria: Guidelines For clinicians (United States)*. June 2015. URL: `https://www.cdc.gov/malaria/`.

[4] **Moore KA, Simpson JA, *et al.*** "Safety of artemisinins in first trimester of prospectively followed pregnancies: an observational study". In: *The Lancet Infectious Diseases* 16.5 (2016), pp. 576–583.

[5] **White NJ, Pongtavornpinyo W, *et al.*** "Hyperparasitaemia and low dosing are an important source of anti-malarial drug resistance". In: *Malaria Journal* 8.253 (2009).

[6] **Ursing J, Rombo L, *et al.*** "High-dose chloroquine for treatment of chloroquine-resistant *Plasmodium falciparum* malaria". In: *The Journal of Infectious Diseases* 213.8 (2016), pp. 1315–1321.

[7] **Wellems TE, Plowe CV**. "Chloroquine-resistant malaria". In: *The Journal of Infectious Diseases* 184.6 (2001), pp. 770–776.

[8] **Van der Wal WM, Geskus RB**. "ipw: An R package for inverse probability weighting". In: *Journal of Statistical Software* 43.13 (2011), pp. 1–23.

[9] **Hernán MA, Brumback B, Robins JM**. "Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men". In: *Epidemiology* 11.5 (2000), pp. 561–570.

[10] **Havercroft WG, Didelez V**. "Simulating from marginal structural models with time-dependent confounding". In: *Statistics in Medicine* 31 (2012), pp. 4190–4206.

[11] **Robins JM, Hernán MA, Brumback B**. "Marginal structural models and causal inference in epidemiology". In: *Epidemiology* 11.5 (2000), pp. 550–560.

[12] **Young JG, Hernán MA, *et al.*** "Simulation from structural survival models under complex time-varying data structures". pp. 1-6.

[13] **Goldfeld K**. *When there's a fork in the road, take it. Or, taking a look at marginal structural models*. Dec. 2017. URL: `https://www.rdatagen.net/post/when-a-covariate-is-a-confounder-and-a-mediator/`.

[14] **STHDA**. *Cox proportional-hazards model*. URL: http://www.sthda.com/english/wiki/cox-proportional-hazards-model.

[15] **Hernán MA, Robins JM**. *Causal Inference*. forthcoming. Boca Raton: Chapman & Hall/CRC, 2018.

[16] **Scheines R**. "An introduction to causal inference". pp. 1-16.

[17] **Goldfeld K**. *When you use inverse probability weighting for estimation, what are the weights actually doing?* Dec. 2017. URL: https://www.rdatagen.net/post/inverse-probability-weighting-when-the-outcome-is-binary/.

[18] **Lee S, Cho Y, Kim S**. "Mapping mHealth (mobile health) and mobile penetrations in sub-Saharan Africa for strategic regional collaboration in mHealth scale-up: An application of exploratory spatial data analysis". In: *Globalization and Health* 13.63 (2017), pp. 1–11.

[19] **Prabhakaran S**. *Logistic regression*. 2016. URL: http://r-statistics.co/Logistic-Regression-With-R.html.